
PAIRFLOW: CLOSED-FORM SOURCE-TARGET COUPLING FOR FEW-STEP GENERATION IN DISCRETE FLOW MODELS

006 **Anonymous authors**

007 Paper under double-blind review

ABSTRACT

013 We introduce PAIRFLOW, a lightweight preprocessing step for training Discrete
014 Flow Models (DFMs) to achieve few-step sampling without requiring a pretrained
015 teacher. DFMs have recently emerged as a new class of generative models for
016 discrete data, offering strong performance. However, they suffer from slow sam-
017 pling due to their iterative nature. Existing acceleration methods largely depend on
018 finetuning, which introduces substantial additional training overhead. PAIRFLOW
019 addresses this issue with a lightweight preprocessing step. Inspired by ReFlow
020 and its extension to DFMs, we train DFMs from coupled samples of source and
021 target distributions, without requiring any pretrained teacher. At the core of our
022 approach is a closed-form inversion for DFMs, which allows efficient construction
023 of paired source-target samples. Despite its extremely low cost, taking only up to
024 1.7% of the compute needed for full model training, PAIRFLOW matches or even
025 surpasses the performance of two-stage training involving finetuning. Furthermore,
026 models trained with our framework provide stronger base models for subsequent
027 distillation, yielding further acceleration after finetuning. Experiments on molecular
028 data as well as binary and RGB images demonstrate the broad applicability and
029 effectiveness of our approach.

1 INTRODUCTION

031 Discrete Flow Models (DFMs) (Campbell et al., 2024; Gat et al., 2024) have recently emerged as a
032 promising class of generative models, extending the idea of Flow Models (FMs) for continuous data
033 to the discrete domain. By adapting flow-based principles to categorical structures, DFMs provide
034 a principled and efficient way to capture complex discrete distributions through iterative sampling.
035 They have shown success across a variety of applications, particularly in scientific domains such as
036 molecule generation (Ramakrishnan et al., 2014; Irwin et al., 2012), where DFMs offer a natural
037 framework for modeling chemical structures and generating novel candidates.

038 Analogous to FMs in the continuous domain, a key challenge of DFMs is the long computation
039 time for generation due to their iterative sampling nature. Recent work (Deschenaux & Gulcehre,
040 2025; Hayakawa et al., 2025; Sahoo et al., 2025; Yoo et al., 2025) have sought to accelerate the
041 generative process through distillation-based finetuning, which builds on ideas originally developed
042 for continuous flow matching. Notably, ReFlow (Liu & Gong, 2023a) is a well-known technique for
043 FMs that pairs samples from the source (prior) distribution and the target distribution by running
044 the generative process of a pretrained model and using the resulting pairs for finetuning. Recently,
045 this idea has also been extended to DFMs (Yoo et al., 2025) to reduce conditional total correlation
046 through finetuning, thereby enabling few-step generation.

047 Despite these promising results in acceleration, distillation-based methods incur substantial finetuning
048 overhead, amounting to about 10–20% of the time required to train the base model from scratch. In
049 other words, the gain in generation speed comes at the expense of considerable additional training cost.
050 To our knowledge, no prior work has addressed this training-time cost when pursuing inference-time
051 acceleration. This raises a natural question: can we achieve speedups comparable to distillation-based
052 approaches while requiring only a lightweight preprocessing phase that requires orders of magnitude
053 less compute, on the order of tens of GPU minutes?

054 We propose PAIRFLOW, a training framework for DFM_s that enables few-step sampling by con-
055 structing paired source–target samples using closed-form velocities. **While inspired by ReDi-style**
056 **coupling-driven training, our approach eliminates the need for a pretrained teacher by using**
057 **closed-form formulations and achieves acceleration without finetuning.** The algorithm for com-
058 puting source–target pairs is fully parallelizable and requires at most 1.7% of compute needed for
059 full model training. Despite relying only on a lightweight preprocessing step, PAIRFLOW attains
060 performance comparable to, or even superior to state-of-the-art distillation-based techniques, which
061 can require up to 143 times more computation. Furthermore, models trained with our technique
062 provide stronger bases for subsequent distillation, delivering additional performance gains while
063 incurring only minimal preprocessing cost.

064 At the core of our framework is the simulation of probability paths connecting source (prior) and target
065 (data) distributions in discrete spaces, made possible by closed-form expressions of velocities. While
066 closed-form forward velocities have been studied for flow models in continuous domains (Karras et al.,
067 2022; Bertrand et al., 2025), they have, to the best of our knowledge, neither been explored for DFM_s
068 nor applied to identifying suitable source–target pairs in the context of distillation-based acceleration,
069 as in ReDi (Yoo et al., 2025). In this work, we investigate this idea for the first time. For DFM_s,
070 with a particular focus on uniform-state models (Sahoo et al., 2025; Schiff et al., 2025) equipped
071 with a self-correcting mechanism, we show that the closed-form forward velocity is determined
072 by the Hamming distance, which measures the number of differing tokens between two sequences.
073 Using this velocity, samples from the source (latent) distribution can be mapped to given target
074 samples. However, because multiple source samples may map to the same target, covering all targets
075 through coupling would require an impractically large number of source samples. To overcome
076 this, we derive the corresponding backward velocity in closed form and leverage it to simulate
077 backward probability paths that efficiently map data points to source points, making pair discovery
078 computationally efficient.

079 In our experiments, we show that the proposed framework enables few-step sampling across diverse
080 discrete domains, including molecular data (Ramakrishnan et al., 2014; Irwin et al., 2012) and 2D
081 images, exemplified by MNIST-Binary (LeCun et al., 2002) and CIFAR-10 (Krizhevsky et al., 2009).
082 On the QM9 (Ramakrishnan et al., 2014) and ZINC-250k (Irwin et al., 2012) datasets, **PAIRFLOW**
083 **not only improves the base model but also performs comparably to, or even better than, distilled**
084 **models that require up to 143 \times more compute during finetuning, compared to our lightweight**
085 **preprocessing algorithm.** Similar improvements are observed on MNIST-Binary, where models
086 paired with PAIRFLOW achieve performance comparable to those using DCD (Sahoo et al., 2025)
087 and ReDi (Yoo et al., 2025), while being up to 35 \times faster. Furthermore, after subsequent distillation,
088 base models trained with pairs generated by our method consistently achieve higher performance,
089 underscoring the importance of well-constructed source–target pairings.

090 2 RELATED WORK

091 2.1 DISCRETE FLOW MODELS

092 The concept of Flow Matching (Lipman et al., 2023) has recently been extended to the discrete
093 flow-based models (Gat et al., 2024; Campbell et al., 2024; Sahoo et al., 2024; Schiff et al., 2025),
094 demonstrating its flexibility across high-dimensional and structured data (Bai et al., 2025; Chang
095 et al., 2022; Weber et al., 2024; Arriola et al., 2025; Nie et al.; Yu et al., 2023; Lee et al., 2025;
096 Campbell et al., 2024; Wang et al., 2025). Among these, uniform-state models (Sahoo et al., 2025;
097 Schiff et al., 2025) have been studied for their self-correcting properties, which enable recovery
098 from errors introduced during parallel decoding. However, their performance degrades markedly in
099 few-step settings, posing a key limitation for efficient generation under tight compute budgets.

100 2.2 ACCELERATING DISCRETE FLOW MODELS

101 Several approaches have been proposed to accelerate sampling with DFM_s. Park et al. (2024) directly
102 optimize sampling timesteps for improved parallelism while mitigating decoding errors. Hayakawa
103 et al. (2025) highlight the importance of capturing dimensional correlations for faster sampling and
104 introduce mixture models to this end, at the cost of additional loss terms that complicate optimization.
105 Another line of work adapts techniques from continuous domains, as in Sahoo et al. (2025), that
106 propose a discrete analogue of Consistency Distillation (CD) (Song et al., 2023) by leveraging the
107 duality between uniform-state and continuous Gaussian models. Most relevant to our approach is

108 ReDi (Yoo et al., 2025), which draws inspiration from the concept of straight flows in ReFlow (Liu &
 109 Gong, 2023a) and iteratively optimizes pairs of data and noise samples.
 110

111 3 PRELIMINARIES

113 In this section, we provide a brief overview of flow matching for generative modeling of discrete
 114 data (Sec. 3.1), followed by a rectification technique (Yoo et al., 2025) that enables faster generation
 115 (Sec. 3.2) by reducing total correlation errors.
 116

117 3.1 DISCRETE FLOW MATCHING

119 The goal of Discrete Flow Matching (DFM) (Campbell et al., 2024; Gat et al., 2024) is to learn a
 120 probability path $p_t(\cdot)$ that connects a known, easy-to-sample source distribution $p(x)$ to an unknown
 121 target distribution $q(x)$, both defined over a discrete state space. Once $p_t(\cdot)$ is known, samples from q
 122 can be generated by drawing $x_0 \sim p$ and transporting it along the path.
 123

124 Specifically, consider a sequence $x = (x^1, x^2, \dots, x^N)$ of N tokens, where each token takes values
 125 in a vocabulary $\mathcal{V} = \{1, 2, \dots, K\}$ of size K . A sequence x then resides in the product space \mathcal{V}^N .
 126 We denote by $\Delta^K = \{p \in \mathbb{R}^K \mid \sum_i p_i = 1, p_i \geq 0\}$ the probability simplex of dimension $K - 1$, on
 127 which distributions over \mathcal{V} are defined.
 128

129 Given a target probability path $p_t(x) : \mathcal{V}^N \times [0, 1] \rightarrow [0, 1]$ with an associated velocity field
 130 $v_t(x) : \mathcal{V}^N \times [0, 1] \rightarrow \mathbb{R}^{N \times K}$, we introduce a network $p_{1|t}^\theta(x) : \mathcal{V}^N \times [0, 1] \rightarrow (\Delta^K)^N$ to
 131 approximate $v_t(x)$. Its parameters θ are optimized via the DFM objective (Gat et al., 2024):
 132

$$133 \mathcal{L}_{\text{DFM}}(\theta) = - \sum_{i \in \{1, \dots, N\}} \mathbb{E}_{t \sim \mathcal{U}[0, 1], x_0 \sim p, x_1 \sim q, z \sim p_t(x|x_0, x_1)} \log p_{1|t}^\theta(x_1^i | z), \quad (1)$$

134 where $p_{1|t}^\theta(x_1^i | z)$ denotes the learned probability denoiser, which predicts the categorical distribution
 135 of the clean token x_1^i given an intermediate sequence z . Here, the conditional probability path
 136 $p_t(z|x_0, x_1)$ generates samples z by interpolating between a data point $x_1 \sim q$ and a source sample
 137 $x_0 \sim p$. Assuming independence across tokens in sequence x , the conditional density factorizes as
 138

$$139 p_t(z|x_0, x_1) = \prod_{i=1}^N p_t(z^i|x_0, x_1). \quad (2)$$

140 As token-wise conditional paths $p_t(z^i|x_0, x_1)$, Gat et al. (2024) employ the mixture path of form:
 141

$$142 p_t(z^i|x_0, x_1) = (1 - \kappa_t) \delta_{x_0}(z^i) + \kappa_t \delta_{x_1}(z^i), \quad (3)$$

143 where the *scheduler* $\kappa_t = \kappa(t)$ is a monotonically increasing function over $t \in [0, 1]$ satisfying
 144 $\kappa_0 = 0$ and $\kappa_1 = 1$. For notational convenience, given $x, y \in \mathcal{V}^N$, we define the Dirac delta $\delta_y(x)$ as
 145

$$146 \delta_y(x) = \prod_{i=1}^N \delta_{y^i}(x^i), \text{ where } \delta_{y^i}(x^i) = \begin{cases} 1 & x^i = y^i \\ 0 & x^i \neq y^i \end{cases}. \quad (4)$$

147 We also use the shorthand $\delta_y(x^i) = \delta_{y^i}(x^i)$. After optimizing θ , the learned model parameterizes an
 148 approximation of the marginal velocity field:
 149

$$150 v_t^\theta(x^i, z) = \frac{\dot{\kappa}_t}{1 - \kappa_t} \left[p_{1|t}^\theta(x^i | z) - \delta_z(x^i) \right], \quad (5)$$

151 where $\dot{\kappa}_t = \frac{\partial \kappa_t}{\partial t}$. This learned velocity field $v_t^\theta(x^i, z)$ then transports samples over the interval $[0, 1]$
 152 to simulate trajectories along $p_t(\cdot)$ and thereby generate samples. Each update step is defined as:
 153

$$154 x_{t+h}^i \sim \text{Cat} \left(x_{t+h}^i; \delta_{x_t^i}(\cdot) + h \cdot v_t^\theta(x_{t+h}^i, x_t) \right), \quad (6)$$

155 where $h > 0$ is the step size.
 156

162 3.2 STRAIGHTENING PROBABILITY PATHS FOR ACCELERATED SAMPLING
 163

164 The concept of straight probability paths was originally introduced in the continuous domain to
 165 enable accelerated sampling. Prior work (Liu & Gong, 2023b) identified curved probability paths as
 166 a key challenge in few-step sampling: when velocity fields are evaluated only at coarse time steps,
 167 numerical integration deviates from the true trajectories. Liu & Gong (2023b) addressed this issue
 168 through “rectification,” in which a student flow model is trained on source–target pairs generated by a
 169 teacher model, effectively yielding significantly straighter probability paths.

170 In the discrete setting, this challenge of *path curvature* translates to capturing the *statistical correlations*
 171 between tokens. Since DFM_s approximate exponentially large joint transitions through fully
 172 factorized per-token updates, a mismatch inevitably arises between the true joint transition and its
 173 product-form approximation. This discrepancy becomes especially detrimental during few-step generation,
 174 where highly correlated tokens must be updated simultaneously. To address this, prior works
 175 have primarily relied on distillation-based approaches (Hayakawa et al., 2025; Sahoo et al., 2025;
 176 Deschenaux & Gulcehre, 2025), aiming to better capture these correlations by explicitly transferring
 177 multi-step dependencies from a stronger teacher model.

178 Yoo et al. (2025) formalized this factorization mismatch via conditional Total Correlation (TC),
 179 defined as:

180
$$\text{TC}_\pi(x_s|x_t) = \mathbb{E}_{x_t} \left[D_{\text{KL}} \left(p_{s|t}(x_s|x_t) \parallel \prod_{i=1}^N p_{s|t}(x_s^i|x_t) \right) \right], \quad (7)$$

181 which serves as a metric for the factorization error. Crucially, Yoo et al. (2025) interpret this factorization
 182 error as the discrete analog of path curvature: minimizing TC is equivalent to “straightening” the
 183 trajectory by decoupling token transitions. Analogous to ReFlow (Liu & Gong, 2023b), which rectifies
 184 continuous paths, they demonstrate that reducing Eqn. 7 requires iteratively refining the source–target
 185 coupling $\pi(x_0, x_1)$. To achieve this, they employ an iterative distillation process, alternating between
 186 generating improved pairs using the current model and optimizing \mathcal{L}_{DFM} . This procedure effectively
 187 finds a “statistically straight” coupling that enables efficient few-step generation.

188 4 PAIRFLOW
 189

190 For DFM_s, ReDi (Yoo et al., 2025) improves sample quality in few-step generation by rectifying
 191 source–target pairs. However, it relies on samples from a pretrained model followed by costly
 192 retraining or finetuning. We take this one step further and pose the following question: What if these
 193 pairs could be generated directly from the data, without relying on a pretrained model or sampling
 194 from it?

195 To address this question, we propose a principled approach for discovering well-aligned source–target
 196 pairs without relying on pretrained models, enabling models trained on such pairs to achieve strong
 197 performance with few-step sampling. Our method, termed PAIRFLOW, leverages closed-form velocity
 198 fields that can be computed directly from the data samples, requiring only prior knowledge of the
 199 source distribution. We assume this distribution to be uniform, a choice extensively studied in recent
 200 work (Sahoo et al., 2025; Schiff et al., 2025), as models trained under this prior naturally acquire
 201 self-correcting properties.

202 In Sec. 4.1, we introduce the closed-form forward velocity for discrete flow matching (Gat et al.,
 203 2024). In Sec. 4.2, we extend this to the closed-form backward velocity and propose an algorithm for
 204 discovering well-aligned source–target pairs during the preprocessing phase.

205 4.1 FINDING PAIRS VIA CLOSED-FORM FORWARD VELOCITY FIELDS
 206

207 As discussed in Sec. 3.1, DFM_s (Campbell et al., 2024; Gat et al., 2024) aim to learn a marginal
 208 velocity field $v_t(\cdot)$ that induces a probability path $p_t(\cdot)$, transporting the source distribution $p_0(\cdot)$ to
 209 the target distribution $q(\cdot) = p_1(\cdot)$, which is unknown in practice. Instead, we only have access to
 210 a finite dataset of M samples $\{d_m\}_{m=1}^M$. This empirical distribution $\tilde{q}(x)$ can be represented as a
 211 mixture of Dirac deltas centered at the observed samples:

212
$$q(x) \approx \tilde{q}(x) = \frac{1}{M} \sum_{m=1}^M \delta_{d_m}(x). \quad (8)$$

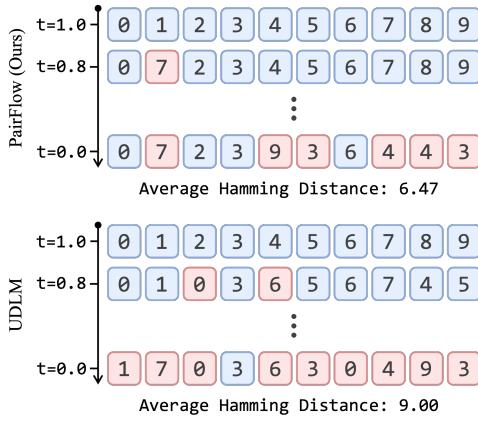


Figure 1: Illustrations of data inversion in PAIRFLOW and the standard corruption process in UDLM. PAIRFLOW achieves a lower average Hamming distance (6.47 vs. 9.0), promoting straighter paths during training.

Algorithm 1: PAIRFLOW

```

1 Input: Dataset  $\{d_m\}_{m=1}^M$ , Steps  $T$ 
2 Output: Pairs  $\pi = \{(x_{0,m}, x_{1,m})\}_{m=1}^M$ 
3 Initialize  $\pi \leftarrow \emptyset$ 
4 for  $m \leftarrow 1$  to  $M$  do
5    $x_{1,m} \leftarrow d_m$ 
6    $x \leftarrow x_{1,m}$ 
7   for  $t \leftarrow 1$  to  $T$  do
8     Compute  $p_{0|t}(\cdot|x)$  and  $\check{v}_t(\cdot, x)$ 
9     Sample
10     $z \sim \text{Cat}(z; \delta_x(\cdot) - h \cdot \check{v}_t(\cdot, x))$ 
11     $x \leftarrow z$ 
12   $x_{0,m} \leftarrow x$ 
13   $\pi \leftarrow \pi \cup \{(x_{0,m}, x_{1,m})\}$ 
14 return  $\pi$ 

```

For continuous domains, Karras et al. (2022); Bertrand et al. (2025) have shown that the velocity field transporting p_0 to q can be derived in closed form when both distributions admit tractable density expressions. To the best of our knowledge, this idea has not been explored in discrete domains; in the following, we derive the closed-form velocity field for discrete domains for the first time.

We base our framework on the assumption of a uniform prior distribution over the discrete state space \mathcal{V}^N , defined as $p_0(x) = \mathcal{U}^N$, where $\mathcal{U} = \text{Cat}(\cdot; \frac{1}{K})$ denotes the uniform distribution over the dictionary \mathcal{V} . For the empirical target distribution $\tilde{q}(x)$ introduced in Eqn. 8, we show in App. A.1 that the closed-form denoiser $p_{1|t}(x^i|z)$ and its associated velocity field $\hat{v}_t(x^i, z)$ are given by:

$$p_{1|t}(x^i|z) = \frac{\sum_{m=1}^M \delta_{d_m^i}(x^i) \gamma^{-h(d_m, z)}}{\sum_{m=1}^M \gamma^{-h(d_m, z)}} \implies \hat{v}_t(x^i, z) = \frac{\kappa_t}{1 - \kappa_t} [p_{1|t}(x^i|z) - \delta_z(x^i)] \quad (9)$$

where $\gamma = (1 + (K-1)\kappa_t)/(1 - \kappa_t)$, K denotes the vocabulary size, and $h(s, z) = N - \sum_{i=1}^N \delta_{s^i}(z^i)$ is the Hamming distance between sequences s and z , i.e., the number of positions at which they differ. The token-wise denoiser $p_{1|t}(x^i|z)$ above is a weighted mixture of Dirac deltas, where sequences closer to z under the Hamming distance contribute more. Intuitively, the forward velocity field $\hat{v}_t(x^i, z)$ pulls each token toward those from dataset sequences most similar to z . The most direct way to construct source-target pairs using $\hat{v}_t(x^i, z)$ is to sample $x_0 \sim p_0(x)$ and evolve it along the velocity field until it reaches a dataset point x_1 . In practice, however, the generated data points fail to fully cover $\tilde{q}(x)$, requiring an impractically large number of source samples to achieve sufficient coverage. Our empirical results, reported in App. C.1, support this claim and motivate the exploration of a more efficient alternative, which we present in the following section.

4.2 FINDING PAIRS VIA CLOSED-FORM BACKWARD VELOCITY FIELDS

We address this issue by *backtracing* trajectories along $p_t(\cdot)$, starting from $\tilde{q}(x)$ and progressing toward the source distribution $p_0(\cdot)$. Unlike the forward construction in Sec. 4.1, this guarantees that all data points in $\tilde{q}(x)$ are included in the resulting pairs by design. As illustrated at the top of Fig. 1, PAIRFLOW inverts data samples toward the source distribution, assumed to be uniform. Unlike the standard corruption process used by UDLM (Schiff et al., 2025) shown at the bottom of Fig. 1, the source samples obtained by PAIRFLOW remain closer to the original data in terms of Hamming distance. This helps the model learn to recover data with fewer token transitions during training, effectively approximating the straight probability paths explored in ReFlow (Liu & Gong, 2023a) and ReDi (Yoo et al., 2025).

The remaining challenge is to derive the closed-form backward velocity that governs this process. This can be obtained by following a construction analogous to Sec. 4.1. Specifically, we first derive

270 Table 1: Dataset and training statistics. N denotes the number of tokens per sample, K the dictionary
 271 size, $|X_1|$ the dataset size, and T_* the runtime of each method (in minutes, measured in wall-clock
 272 time with RTX A6000). The numbers in parentheses are the proportion of time relative to T_{Base} .

Dataset	N	K	$ X_1 $	T_{Base}	T_{DCD}	T_{ReDi}	T_{PAIRFLOW}
MNIST-Binary	768	2	60,000	80 (100.0%)	40 (50.0%)	49 (61.0%)	1.4 (1.7%)
CIFAR-10	3,072	256	100,000	6720 (100.0%)	360 (5.3%)	468 (6.9%)	20 (0.3%)
QM9	32	40	127,190	450 (100.0%)	115 (24.8%)	100 (22.2%)	0.8 (0.2%)
ZINC-250k	72	74	224,568	1,110 (100.0%)	211 (19.0%)	194 (17.4%)	13 (1.2%)

279 the closed-form noise predictor $p_{0|t}(x^i|z)$:

$$282 \quad p_{0|t}(x^i|z) = \delta_z(x^i) - \frac{\kappa_t(K\delta_{x^i}(z^i) - 1)}{1 + (K - 1)\kappa_t} \frac{\sum_{m=1}^M \delta_{d_m^i}(z^i) \gamma^{-h(d_m, z)}}{\sum_{m=1}^M \gamma^{-h(d_m, z)}}, \quad (10)$$

285 with a detailed derivation provided in App. A.2. Substituting this into the definition of the backward
 286 velocity field from Gat et al. (2024)

$$288 \quad \check{v}_t(x^i, z) = \frac{\dot{\kappa}_t}{\kappa_t} [\delta_z(x^i) - p_{0|t}(x^i|z)], \quad (11)$$

290 we obtain the desired closed-form expression

$$292 \quad \check{v}_t(x^i, z) = \frac{\dot{\kappa}_t(K\delta_{x^i}(z^i) - 1)}{1 + (K - 1)\kappa_t} \frac{\sum_{m=1}^M \delta_{d_m^i}(z^i) \gamma^{-h(d_m, z)}}{\sum_{m=1}^M \gamma^{-h(d_m, z)}}. \quad (12)$$

295 The second term in Eqn. 10 computes the conditional likelihood of the i -th token taking value $x^i \in \mathcal{V}$
 296 given the current sequence z , marginalized over all dataset $\{d_m\}_{m=1}^M$. The contribution of each data
 297 sample d_m is determined by its proximity to z under the Hamming distance $h(d_m, z)$, assigning
 298 higher weight to tokens with greater local consensus. Consequently, updating with $\check{v}_t(x^i, z)$ pushes
 299 the sample away from the data distribution and toward the source distribution $p_0(x)$. Using $\check{v}_t(x^i, z)$,
 300 we construct pairs $\{(x_{0,m}, x_{1,m})\}_{m=1}^M$ by initializing from data points $\{d_m\}_{m=1}^M$ (equivalently,
 301 $\{x_{1,m}\}_{m=1}^M$) and iteratively applying the backward update rule in Eqn. 6 for a fixed number of
 302 iterations T . The overall procedure is summarized in Alg. 1.

304 5 EXPERIMENTAL RESULTS

306 We validate the effectiveness of the proposed method and the source–target pairs it discovers across
 307 several discrete generative modeling benchmarks involving molecular data and images. We first
 308 summarize the experimental setup in Sec. 5.1. In Sec. 5.2 and Sec. 5.3, we compare our method
 309 against baselines in molecular and image generation, respectively. In Sec. 5.4, we further demonstrate
 310 that models trained with pairs discovered by our method not only achieve improved performance
 311 directly, but also benefit subsequent distillation phases, as the resulting base model provides a stronger
 312 initialization for existing distillation techniques.

313 5.1 EXPERIMENT SETUP

315 **Baselines.** Across multiple benchmarks, we compare our approach against state-of-the-art discrete
 316 flow models, including MDLM (Sahoo et al., 2024) and UDLM (Schiff et al., 2025). Since our method
 317 is based on a uniform source distribution, we adopt UDLM (Schiff et al., 2025), the leading uniform-
 318 state model, as our base and denote UDLM trained with pairs generated by Alg. 1 as PAIRFLOW
 319 throughout the remainder of this section. In addition, we compare against these models augmented
 320 with distillation-based techniques that require additional finetuning, Discrete Consistency Distillation
 321 (DCD) (Sahoo et al., 2025) and ReDi (Yoo et al., 2025), denoted throughout this section by the suffixes
 322 “+ DCD” and “+ ReDi”. The detailed training setups of these models, such as hyperparameters, are
 323 provided in App. B. Additionally, we report the performance of the same base model trained on pairs
 324 formed by each data point and a source sample randomly drawn from the uniform distribution with
 325 our detailed experimental results in App. D.

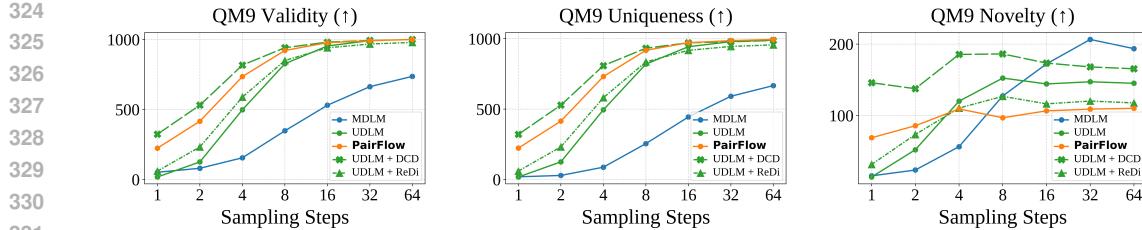


Figure 2: Step-wise performance analysis on the QM9 dataset (Ramakrishnan et al., 2014). Each plot reports the number of valid (left), unique (middle), and novel (right) SMILES strings (Weininger, 1988) out of 1,024 generated samples. Best viewed when zoomed in.

Benchmarks. We evaluate our method across a diverse set of discrete generative modeling benchmarks, covering both molecule and image generation tasks. For molecule generation, we experiment with the QM9 (Ramakrishnan et al., 2014) and ZINC-250k (Irwin et al., 2012) datasets. For image generation, we use the MNIST dataset (LeCun et al., 2002) with binarized pixel values (denoted MNIST-Binary) and the CIFAR-10 dataset (Krizhevsky et al., 2009), where pixel intensities are scaled to 8-bit integers, and horizontal flip augmentation is applied. Dataset statistics, including sample size, vocabulary size, and overall dataset size, are summarized in Tab. 1.

Evaluation Setup. For molecular generation, we follow Schiff et al. (2025) and evaluate the validity, uniqueness, and novelty of generated molecules. Specifically, we sample 1,024 SMILES strings (Weininger, 1988), convert them into molecular graphs, and compute these metrics. All results are averaged over 10 trials, with further details provided in App. D. For image generation, we report Fréchet Inception Distance (FID) (Heusel et al., 2017) and Inception Score (IS) (Salimans et al., 2016). FID is computed with 1,000 images for MNIST-Binary, and both FID and IS are computed with 5,000 generated images for CIFAR-10. The training dataset is used as the reference for FID computation. Across all experiments, we vary the number of sampling steps to evaluate performance in both low- and high-NFE settings. In particular, we generate samples using 1 – 64 steps for molecular benchmarks (QM9 and ZINC-250k) and MNIST-Binary benchmark, and 8 – 1024 steps for CIFAR-10 (Krizhevsky et al., 2009), as models yielded excessively high FIDs under extremely low-step settings.

5.2 MOLECULE GENERATION

We begin by benchmarking unconditional molecule generation, where models are tasked with generating SMILES strings (Weininger, 1988) that represent molecules. As illustrated in Fig. 2 and Fig. 3, which summarize validity (left), uniqueness (middle), and novelty (right), PAIRFLOW consistently improves upon its base model UDLM (Schiff et al., 2025), yielding substantial gains in few-step settings. It facilitates 1-step generation on QM9 (Ramakrishnan et al., 2014), a challenging setting that requires capturing all token-wise dependencies simultaneously. In this case, validity increases from 17.5 to 223.4, corresponding to a 12.8 \times improvement. Similar trends are observed in the 2-step and 4-step settings, with validity improving by 231% and 47.6%, respectively. As shown in Fig. 2 (left), this improvement is particularly significant: the 2-step and 4-step validities of PAIRFLOW are comparable to the 4-step and 8-step validities achieved by UDLM (Schiff et al., 2025). Comparable improvements are also seen on the ZINC-250k (Irwin et al., 2012) dataset.

Remarkably, **PAIRFLOW introduces minimal overhead—less than 2% of the training cost as shown in Tab. 1—and requires no pretrained models, yet achieves performance comparable to, and in some cases surpassing, models distilled from the same base using DCD (Sahoo et al., 2025) and ReDi (Yoo et al., 2025)**, both of which rely on pretrained models and finetuning. On both QM9 (Ramakrishnan et al., 2014) and ZINC-250k (Irwin et al., 2012), PAIRFLOW consistently outperforms UDLM + ReDi across all few-step settings, achieving substantially higher 2-step validities on QM9 (232.4 vs. 416.0) and ZINC-250k (75.9 vs. 146.3). At the same time, PAIRFLOW matches the performance of UDLM + DCD, with comparable 2-step validities on QM9 (416.0 vs. 530.8). This is particularly notable given that the additional preprocessing cost of PAIRFLOW amounts to only 0.69% on QM9 (Ramakrishnan et al., 2014) and 6.16% on ZINC-250k (Irwin et al., 2012), relative to the full cost of DCD (Sahoo et al., 2025). Detailed numerical results with standard deviations are reported in App. D.

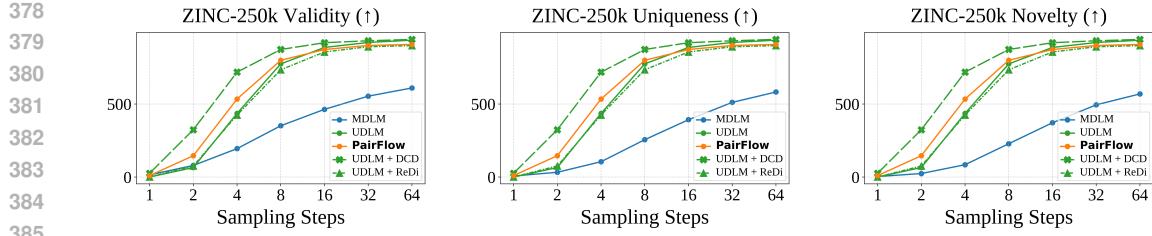


Figure 3: Step-wise performance analysis on the ZINC-250k dataset (Irwin et al., 2012). Each plot reports the number of valid (left), unique (middle), and novel (right) SMILES strings (Weininger, 1988) out of 1,024 generated samples. Best viewed when zoomed in.

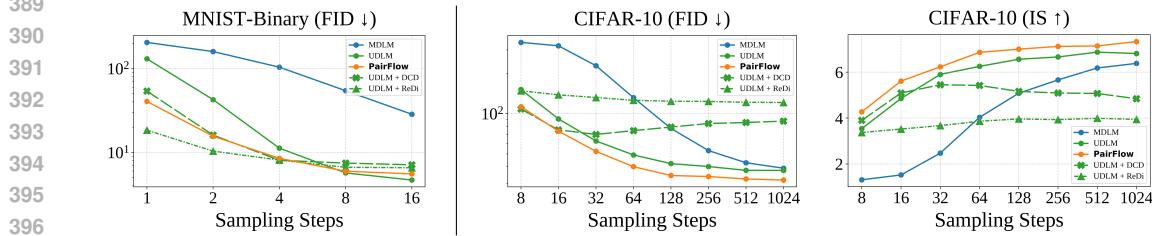


Figure 4: Step-wise performance analysis on discretized image datasets. From left to right: FID on MNIST-Binary (LeCun et al., 2002), FID on CIFAR-10 (Krizhevsky et al., 2009), and Inception Scores (IS) on CIFAR-10. Best viewed when zoomed in.

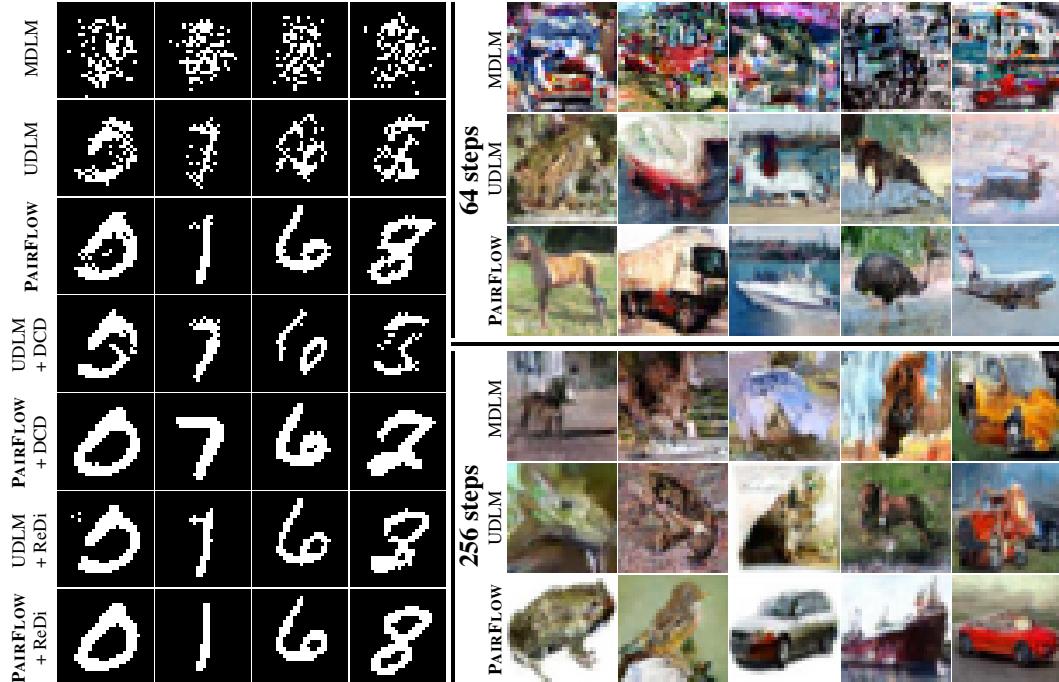


Figure 5: Qualitative results of 1-step generation on MNIST-Binary (28×28 ; left) and 64-step (top right) and 256-step (bottom right) generation on CIFAR-10 (32×32).

5.3 IMAGE GENERATION

We further extend our experiments to image domains where each pixel has discretized intensities. As in Sec. 5.2, we evaluate model performance across multiple sampling steps and summarize the results in Fig. 4. Qualitative samples for MNIST-Binary and CIFAR-10 are shown in Fig. 5. Both qualitative and quantitative results show that PAIRFLOW improves the performance of UDLM (Schiff et al., 2025) and, in few-step settings, achieves performance comparable to DCD (Sahoo et al., 2025).

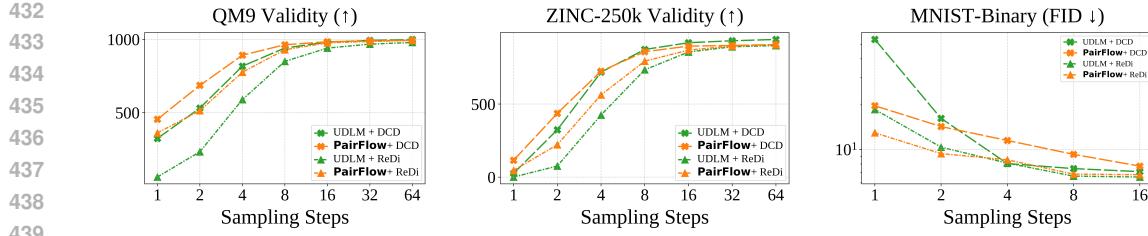


Figure 6: Step-wise performance analysis of distilled models on molecular and image datasets. From left to right: number of valid molecules on QM9, number of valid molecules on ZINC-250k, and FID on MNIST-Binary.

On MNIST-Binary (Fig. 4, left), PAIRFLOW achieves an FID of 40.59 in the 1-step setting, equivalently a 68.9% improvement over UDLM (Schiff et al., 2025). Consistent gains are observed across other few-step settings as well: at 2 steps, FID is reduced by 63.3% (15.61 vs. 42.54), and at 4 steps by 24.4% (8.51 vs. 11.25). On CIFAR-10 (Krizhevsky et al., 2009), where FID (Heusel et al., 2017) and IS (Salimans et al., 2016) are reported in Fig. 4 (middle and right), PAIRFLOW likewise outperforms the base UDLM, validating the effectiveness of the discovered source–target pairs.

As in molecular generation (Sec. 5.2), **PAIRFLOW performs comparable to distillation-based acceleration methods.** On MNIST-Binary, it achieves a lower FID than UDLM+DCD in the 1-step setting (40.59 vs. 53.84) and comparable performance at 2 steps (15.61 vs. 16.09). Likewise, PAIRFLOW performs competitively with UDLM+ReDi at 2 steps (15.61 vs. 10.36), while requiring substantially less compute than both. As summarized in Tab. 1, DCD (Sahoo et al., 2025) requires 40 minutes (T_{DCD}) and ReDi (Yoo et al., 2025) takes 49 minutes, whereas the preprocessing phase of PAIRFLOW completes in just 1.4 minutes ($T_{PAIRFLOW}$), yielding $28.6\times$ and $35\times$ speedups, respectively. On the CIFAR-10 benchmark, both DCD (Sahoo et al., 2025) and ReDi (Yoo et al., 2025) degrade model performance, as indicated by the higher FID in Fig. 4 (middle) and lower IS in Fig. 4 (right). The results in Tab. 13 suggest that, overall, acceleration methods do not work well on CIFAR-10. We hypothesize that this issue arises from the low performance of the teacher model, which negatively affects the student model when applying acceleration methods. Detailed results are reported in App. D.

5.4 DISTILLING MODELS TRAINED WITH ALIGNED PAIRS

While PAIRFLOW alone achieves performance comparable to, or even exceeding, distillation-based techniques (Yoo et al., 2025; Sahoo et al., 2025), as shown in Sec. 5.2 and Sec. 5.3, we further emphasize that it also serves as a strong initialization for subsequent distillation, yielding even greater performance gains when combined with existing methods. Crucially, this incurs negligible additional cost relative to the overall time required for distillation.

We validate this by distilling PAIRFLOW, trained on QM9 (Ramakrishnan et al., 2014), ZINC-250k (Irwin et al., 2012), and MNIST-Binary, using DCD (Sahoo et al., 2025) and ReDi (Yoo et al., 2025), and comparing their performance against distilled models whose teachers were the base UDLM (Schiff et al., 2025). As shown in Fig. 6, **student models distilled from PAIRFLOW, denoted PAIRFLOW+DCD and PAIRFLOW+ReDi, push the frontier of performance previously achieved by distillation-based techniques.** For example, on the QM9 dataset (Ramakrishnan et al., 2014), PAIRFLOW+DCD substantially improves validity over UDLM+DCD (453.8 vs. 323 for 1 step, 685.8 vs. 530.8 for 2 steps). A similar trend is observed for PAIRFLOW+ReDi on ZINC-250k (Irwin et al., 2012), yielding higher scores in both 1-step (46.3 vs. 0.7) and 2-step (221.5 vs. 75.9) generation. Importantly, as summarized in Tab. 1, these gains are achieved at only minimal additional preprocessing cost: 3.15% of the average runtime of distillation on MNIST-Binary, 0.77% on QM9, and 6.42% on ZINC-250k.

6 CONCLUSION

We have presented PAIRFLOW, a novel approach to accelerating the generative process of Discrete Flow Models (DFMs) through a lightweight preprocessing step performed prior to training. Our preprocessing, which couples source and target samples, requires only up to 1.7% of the base model training cost, making it at least 20 \times more efficient than finetuning, while still achieving comparable or even superior performance. The key enabler is the closed-form inversion, which eliminates the need for a pretrained teacher model.

486 ETHICS STATEMENT

487
488 We comply with the ICLR Code of Ethics. Our study relies exclusively on publicly available models
489 and datasets and involves no human subjects, user data, or personally identifiable information. We
490 acknowledge possible misuse of generative AI and urge responsible use of our method.

491 REPRODUCIBILITY STATEMENT

492
493 We plan to release the code upon publication. The pseudocode of our algorithm is provided in Alg. 1.
494 For reproducibility, we provide detailed proofs of the closed-form forward and backward velocities
495 in App. A. The experimental setup and hyperparameters used in our experiments are also described
496 in App. B.

497 REFERENCES

498
499 Marianne Arriola, Aaron Gokaslan, Justin T Chiu, Zhihan Yang, Zhixuan Qi, Jiaqi Han, Sub-
500 ham Sekhar Sahoo, and Volodymyr Kuleshov. Block diffusion: Interpolating between autoregres-
501 sive and diffusion language models. In *ICLR*, 2025.

502
503 Jinbin Bai, Tian Ye, Wei Chow, Enxin Song, Xiangtai Li, Zhen Dong, Lei Zhu, and Shuicheng Yan.
504 Meisisonic: Revitalizing masked generative transformers for efficient high-resolution text-to-image
505 synthesis. In *ICLR*, 2025.

506
507 Quentin Bertrand, Anne Gagneux, Mathurin Massias, and Rémi Emonet. On the closed-form of flow
508 matching: Generalization does not arise from target stochasticity. *arXiv*, 2025.

509
510 Andrew Campbell, Jason Yim, Regina Barzilay, Tom Rainforth, and Tommi Jaakkola. Generative
511 flows on discrete state-spaces: Enabling multimodal flows with applications to protein co-design.
512 In *ICML*, 2024.

513
514 Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative
515 image transformer. In *CVPR*, 2022.

516
517 Ciprian Chelba, Tomas Mikolov, Mike Schuster, Qi Ge, Thorsten Brants, Philipp Koehn, and Tony
518 Robinson. Exploring the limits of language modeling. *arXiv preprint arXiv:1312.3005*, 2013.

519
520 Justin Deschenaux and Caglar Gulcehre. Beyond Autoregression: Fast LLMs via Self-Distillation
521 Through Time. In *ICLR*, 2025.

522
523 Itai Gat, Tal Remez, Neta Shaul, Felix Kreuk, Ricky TQ Chen, Gabriel Synnaeve, Yossi Adi, and
524 Yaron Lipman. Discrete flow matching. In *NeurIPS*, 2024.

525
526 Satoshi Hayakawa, Yuhta Takida, Masaaki Imaizumi, Hiromi Wakaki, and Yuki Mitsufuji. Distillation
527 of discrete diffusion through dimensional correlations. In *ICML*, 2025.

528
529 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs
530 trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017.

531
532 John J Irwin, Teague Sterling, Michael M Mysinger, Erin S Bolstad, and Ryan G Coleman. ZINC: a
533 free tool to discover chemistry for biology. *Journal of chemical information and modeling*, 52(7):
534 1757–1768, 2012.

535
536 Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative
537 adversarial networks. In *CVPR*, 2019.

538
539 Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-
540 based generative models. In *NeurIPS*, 2022.

541
542 Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *ICLR*, 2014.

543
544 Alex Krizhevsky et al. Learning multiple layers of features from tiny images. 2009.

545
546 Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to
547 document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 2002.

540 Seul Lee, Karsten Kreis, Srimukh Prasad Veccham, Meng Liu, Danny Reidenbach, Yuxing Peng,
541 Saeed Gopal Paliwal, Weili Nie, and Arash Vahdat. Genmol: A drug discovery generalist with
542 discrete diffusion. In *ICML*, 2025.

543 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
544 for generative modeling. In *ICLR*, 2023.

545 Xingchao Liu and Liu Qiang Gong, Chengyue. Flow straight and fast: Learning to generate and
546 transfer data with rectified flow. In *ICLR*, 2023a.

547 Xingchao Liu and Liu Qiang Gong, Chengyue. Flow straight and fast: Learning to generate and
548 transfer data with rectified flow. In *ICLR*, 2023b.

549 Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, JUN ZHOU, Yankai Lin,
550 Ji-Rong Wen, and Chongxuan Li. Large language diffusion models. In *ICLR 2025 Workshop on*
551 *Deep Generative Model in Machine Learning: Theory, Principle and Efficacy*.

552 Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov,
553 Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning
554 robust visual features without supervision. *TMLR*, 2024.

555 Yong-Hyun Park, Chieh-Hsin Lai, Satoshi Hayakawa, Yuhta Takida, and Yuki Mitsufuji. Jump your
556 steps: Optimizing sampling schedule of discrete diffusion models. In *ICLR*, 2024.

557 Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. Quantum
558 chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1):1–7, 2014.

559 Subham Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin Chiu,
560 Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language
561 models. In *NeurIPS*, 2024.

562 Subham Sekhar Sahoo, Justin Deschenaux, Aaron Gokaslan, Guanghan Wang, Justin T Chiu, and
563 Volodymyr Kuleshov. The diffusion duality. In *ICML*, 2025.

564 Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen.
565 Improved techniques for training GANs. In *NeurIPS*, 2016.

566 Yair Schiff, Subham Sekhar Sahoo, Hao Phung, Guanghan Wang, Sam Boshar, Hugo Dalla-torre,
567 Bernardo P de Almeida, Alexander M Rush, Thomas PIERROT, and Volodymyr Kuleshov. Simple
568 guidance mechanisms for discrete diffusion models. In *ICLR*, 2025.

569 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
570 Poole. Score-based generative modeling through stochastic differential equations. In *International
571 Conference on Learning Representations*, 2021.

572 Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *ICML*, 2023.

573 Chenyu Wang, Masatoshi Uehara, Yichun He, Amy Wang, Avantika Lal, Tommi Jaakkola, Sergey
574 Levine, Aviv Regev, Tommaso Biancalani, et al. Fine-tuning discrete diffusion models via reward
575 optimization with applications to dna and protein design. In *ICLR*, 2025.

576 Mark Weber, Lijun Yu, Qihang Yu, Xueqing Deng, Xiaohui Shen, Daniel Cremers, and Liang-Chieh
577 Chen. Maskbit: Embedding-free image generation via bit tokens. 2024.

578 David Weininger. Smiles, a chemical language and information system. 1. introduction to methodol-
579 ogy and encoding rules. *Journal of Chemical Information and Computer Sciences*, 28(1):31–36,
580 1988. doi: 10.1021/ci00057a005.

581 Jaehoon Yoo, Wonjung Kim, and Seunghoon Hong. Redi: Rectified discrete flow. In *NeurIPS*, 2025.

582 Lijun Yu, Yong Cheng, Kihyuk Sohn, José Lezama, Han Zhang, Huiwen Chang, Alexander G
583 Hauptmann, Ming-Hsuan Yang, Yuan Hao, Irfan Essa, et al. Magvit: Masked generative video
584 transformer. In *CVPR*, 2023.

594 A PROOF FOR CLOSED-FORM INVERSION

596 In this section, we present the detailed derivations of the closed-form forward velocity Eqn. 9 and
 597 backward velocity Eqn. 12 introduced in Sec. 4. Sec. A.1 provides the proof of the closed-form
 598 forward velocity, while Sec. A.2 presents the proof of the closed-form backward velocity. Both
 599 derivations are based on the assumption of a uniform source distribution.

601 A.1 PROOF OF CLOSED-FORM FORWARD VELOCITY IN DISCRETE FLOW MODELS

603 Let $x, z \in \mathcal{V}^N$ be sequences of tokens $x^i, z^i \in \mathcal{V}$ for $i \in 1, \dots, N$, where each token takes values
 604 from the discrete vocabulary $\mathcal{V} = \{1, \dots, K\}$. We begin with the expression of the forward velocity
 605 given in Eqn. 9:

$$606 \hat{v}_t(x^i, z) = \frac{\dot{\kappa}_t}{1 - \kappa_t} [p_{1|t}(x^i|z) - \delta_z(x^i)]. \quad (13)$$

609 We first derive the closed-form expression for the probability denoiser $p_{1|t}(x^i|z)$:

610 Using Bayes' rule,

$$612 p_{1|t}(x^i|z) = \sum_{x_0, x_1} \delta_{x_1^i}(x^i) p_t(x_0, x_1|z) \quad (14)$$

$$615 = \sum_{x_0, x_1} \delta_{x_1^i}(x^i) \frac{p_t(x_0, x_1, z)}{p_t(z)} \quad (15)$$

$$617 = \frac{\sum_{x_0, x_1} \delta_{x_1^i}(x^i) p_t(x_0, x_1, z)}{\sum_{x_0, x_1} p_t(x_0, x_1, z)}. \quad (16)$$

620 We factor the joint as

$$621 p_t(x_0, x_1, z) \propto p_0(x_0) p_1(x_1) p_t(z|x_0, x_1) \quad (17)$$

$$623 p_t(z|x_0, x_1) = \prod_{j=1}^N [\kappa_t \delta_{x_1^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)] \quad (18)$$

626 and use the empirical target

$$628 p_1(x_1) = \frac{1}{M} \sum_{m=1}^M \delta_{d_m}(x_1) \quad (\text{Eqn. 8}). \quad (19)$$

631 Since $p_0(x_0)$ is constant, it cancels between the numerator and denominator of Eqn. 16, yielding

$$633 p_{1|t}(x^i|z) = \frac{\sum_{x_0, x_1} \delta_{x_1^i}(x^i) p_1(x_1) p_t(z|x_0, x_1)}{\sum_{x_0, x_1} p_1(x_1) p_t(z|x_0, x_1)} \\ 635 = \frac{\sum_{m=1}^M \sum_{x_0} \delta_{d_m^i}(x^i) \prod_{j=1}^N [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)]}{\sum_{m=1}^M \sum_{x_0} \prod_{j=1}^N [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)]}. \quad (20)$$

639 According to this expression, the probability denoiser $p_{1|t}(x^i|z)$ can be interpreted as a weighted
 640 sum over all data points x_1 , given by the term

$$642 \sum_{x_0} \prod_j [\kappa_t \delta_{x_1^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)]. \quad (21)$$

644 Let $h(s, z)$ denote the Hamming distance between two sequences s and z , defined as

$$646 h(s, z) = N - \sum_j \delta_{s^j}(z^j), \quad (22)$$

648 and let $h_+(s, z)$ represent the similarity between the sequences, defined as
649

650
$$h_+(s, z) = \sum_j \delta_{s^j}(z^j) = N - h(s, z). \quad (23)$$

651
652

653 The weight is computed only when d_m^i coincides with the target token x^i (i.e., $\delta_{d_m^i}(x^i) = 1$). In this
654 case, the term can be expressed as:
655

656
$$\sum_{x_0} \prod_j \left[\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j) \right] \quad (24)$$

657

658
$$= \sum_{k=0}^{h_+(d_m, z)} \binom{h_+(d_m, z)}{k} (1 - \kappa_t)^{N - h_+(d_m, z)} ((K - 1) \kappa_t)^{h_+(d_m, z) - k}. \quad (25)$$

659
660

661 To understand this transition, we first note that d_m and z are fixed in this scope, while x_0 is independent
662 across each dimension and follows a uniform distribution. This implies that we only need to consider
663 x_0 . For an arbitrary dimension j , the cases can be divided into four possibilities, and the corresponding
664 values of $\left[\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j) \right]$ are as follows:
665

666 **Case 1.** $d_m^j = z^j, x_0^j = z^j, \left[\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j) \right] = 1.$
667

668 **Case 2.** $d_m^j = z^j, x_0^j \neq z^j, \left[\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j) \right] = \kappa_t.$
669

670 **Case 3.** $d_m^j \neq z^j, x_0^j = z^j, \left[\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j) \right] = 1 - \kappa_t.$
671

672 **Case 4.** $d_m^j \neq z^j, x_0^j \neq z^j, \left[\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j) \right] = 0.$
673

674 Note that **Case 4** makes the term inside the product $\left[\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j) \right]$ equal to zero.
675

676 Thus, we only need to consider x_0 for which no dimension falls into **Case 4**. Among the $|K|^N$
677 possible choices of x_0 , only $|K|^{h_+(d_m, z)}$ satisfy $x_0^j = z^j$ for all dimensions where $d_m^j \neq z^j$. We then
678 classify the remaining cases according to the Hamming distance between x_0 and d_m . Note that the
679 maximum value of $h_+(x_0, d_m)$ is $h_+(d_m, z)$. Let k denote an integer in the range 0 to $h_+(d_m, z)$.
680 Then, the number of x_0 satisfying $h_+(x_0, d_m) = k$ is $\binom{h_+(d_m, z)}{k} (K - 1)^{h_+(d_m, z) - k}$, and in this case
681 the product term becomes
682

683
$$\prod_j \left[\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j) \right] = (\kappa_t)^{h_+(d_m, z) - k} (1 - \kappa_t)^{N - h_+(d_m, z)}. \quad (26)$$

684
685

686 We can then arrive at the equation above by summing over all possible k . Resuming the proof, the
687 term can be further simplified as follows:
688

689
$$\sum_{x_0} \prod_j \left[\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j) \right] \quad (27)$$

690

691
$$= \sum_{k=0}^{h_+(d_m, z)} \binom{h_+(d_m, z)}{k} (1 - \kappa_t)^{N - h_+(d_m, z)} ((K - 1) \kappa_t)^{h_+(d_m, z) - k} \quad (28)$$

692

693
$$= (1 - \kappa_t)^{N - h_+(d_m, z)} \sum_{k=0}^{h_+(d_m, z)} \binom{h_+(d_m, z)}{k} ((K - 1) \kappa_t)^k \quad (29)$$

694

695
$$= (1 - \kappa_t)^N \left(\frac{(K - 1) \kappa_t + 1}{1 - \kappa_t} \right)^{h_+(d_m, z)} \quad (30)$$

696

697
$$= (1 - \kappa_t)^N \left(1 + \frac{\kappa_t}{1 - \kappa_t} K \right)^{h_+(d_m, z)}. \quad (31)$$

698

We define $\gamma := 1 + \frac{\kappa_t}{1-\kappa_t} K$, and by substituting this simplified expression into the noise predictor above, we finally obtain Eqn. 10.

$$p_{1|t}(x^i|z) = \frac{\sum_{m=1}^M \delta_{d_m^i}(x^i) \sum_{x_0} \prod_j [\kappa_t \delta_{d_m^j}(z^j) + (1-\kappa_t) \delta_{x_0^j}(z^j)]}{\sum_{m=1}^M \sum_{x_0} \prod_j [\kappa_t \delta_{d_m^j}(z^j) + (1-\kappa_t) \delta_{x_0^j}(z^j)]} \quad (32)$$

$$= \frac{\sum_{m=1}^M \delta_{d_m^i}(x^i) (1-\kappa_t)^N \left(1 + \frac{\kappa_t}{1-\kappa_t} K\right)^{h_+(d_m, z)}}{\sum_{m=1}^M (1-\kappa_t)^N \left(1 + \frac{\kappa_t}{1-\kappa_t} K\right)^{h_+(d_m, z)}} \quad (33)$$

$$= \frac{\sum_{m=1}^M \delta_{d_m^i}(x^i) \left(1 + \frac{\kappa_t}{1-\kappa_t} K\right)^{h_+(d_m, z)}}{\sum_{m=1}^M \left(1 + \frac{\kappa_t}{1-\kappa_t} K\right)^{h_+(d_m, z)}} \quad (34)$$

$$= \frac{\sum_{m=1}^M \delta_{d_m^i}(x^i) \gamma^{-h(d_m, z)}}{\sum_{m=1}^M \gamma^{-h(d_m, z)}}. \quad (35)$$

A.2 PROOF OF CLOSED-FORM BACKWARD VELOCITY IN DISCRETE FLOW MODELS

Similarly to the proof of the closed-form forward velocity in Sec. A.1, we start from the backward velocity in Eqn. 11:

$$\check{v}_t(x^i, z) = \frac{\dot{\kappa}_t}{\kappa_t} [\delta_{z^i}(x^i) - p_{0|t}(x^i|z)]. \quad (36)$$

We derive the closed-form noise predictor as follows:

$$p_{0|t}(x^i|z) = \sum_{x_0, x_1} \delta_{x_0^i}(x^i) p_t(x_0, x_1|z) \quad (37)$$

$$= \sum_{x_0, x_1} \delta_{x_0^i}(x^i) \frac{p_t(x_0, x_1, z)}{p_t(z)} \quad (38)$$

$$= \frac{\sum_{x_0, x_1} \delta_{x_0^i}(x^i) p_t(x_0, x_1, z)}{p_t(z)} \quad (39)$$

$$= \frac{\sum_{x_0, x_1} \delta_{x_0^i}(x^i) p_t(x_0, x_1, z)}{\sum_{x_0, x_1} p_t(x_0, x_1, z)}, \quad (40)$$

The last expression is further expanded to:

$$p_{0|t}(x^i|z) = \frac{\sum_{x_0, x_1} \delta_{x_0^i}(x^i) p_t(x_0, x_1, z)}{\sum_{x_0, x_1} p_t(x_0, x_1, z)} \quad (41)$$

$$= \frac{\sum_{x_0, x_1} \delta_{x_0^i}(x^i) p_1(x_1) p_t(z|x_0, x_1)}{\sum_{x_0, x_1} p_1(x_1) p_t(z|x_0, x_1)} \quad (42)$$

$$= \frac{\sum_{m=1}^M \sum_{x_0} \delta_{x_0^i}(x^i) \prod_{j=1}^N [\kappa_t \delta_{d_m^j}(z^j) + (1-\kappa_t) \delta_{x_0^j}(z^j)]}{\sum_{m=1}^M \sum_{x_0} \prod_{j=1}^N [\kappa_t \delta_{d_m^j}(z^j) + (1-\kappa_t) \delta_{x_0^j}(z^j)]}. \quad (43)$$

For the denominator, we use the same formula as in Eqn. 24:

$$\sum_{x_0} \prod_j [\kappa_t \delta_{d_m^j}(z^j) + (1-\kappa_t) \delta_{x_0^j}(z^j)] = (1-\kappa_t)^N \left(1 + \frac{\kappa_t}{1-\kappa_t} K\right)^{h_+(d_m, z)}. \quad (44)$$

To simplify the numerator, we decompose the expression by separating the i -th dimension:

$$\sum_{x_0} \delta_{x_0^i}(x^i) \prod_j [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)] \quad (45)$$

$$= \sum_{\substack{x_0 \text{ with} \\ x_0^i = x^i}} \prod_j [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)] \quad (46)$$

$$= \sum_{\substack{x_0 \text{ with} \\ x_0^i = x^i}} [\kappa_t \delta_{d_m^i}(z^i) + (1 - \kappa_t) \delta_{x_0^i}(z^i)] \prod_{j \neq i} [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)] \quad (47)$$

$$= \sum_{\substack{x_0 \text{ with} \\ x_0^i = x^i}} [\kappa_t \delta_{d_m^i}(z^i) + (1 - \kappa_t) \delta_{x_0^i}(z^i)] \prod_{j \neq i} [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)] \quad (48)$$

$$= \underbrace{[\kappa_t \delta_{d_m^i}(z^i) + (1 - \kappa_t) \delta_{x_0^i}(z^i)]}_{(\square)} \underbrace{\sum_{\substack{x_0 \text{ with} \\ x_0^i = x^i}} \prod_{j \neq i} [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)]}_{(\triangle)} \quad (49)$$

We can break it down into four cases for the i -th dimension term, along with their corresponding values for the (\square) part.

Case 1. $d_m^i = z^i, z^i = x^i, [\kappa_t \delta_{d_m^i}(z^i) + (1 - \kappa_t) \delta_{x_0^i}(z^i)] = 1$.

Case 2. $d_m^i = z^i, z^i \neq x^i, [\kappa_t \delta_{d_m^i}(z^i) + (1 - \kappa_t) \delta_{x_0^i}(z^i)] = \kappa_t$.

Case 3. $d_m^i \neq z^i, z^i = x^i, [\kappa_t \delta_{d_m^i}(z^i) + (1 - \kappa_t) \delta_{x_0^i}(z^i)] = 1 - \kappa_t$.

Case 4. $d_m^i \neq z^i, z^i \neq x^i, [\kappa_t \delta_{d_m^i}(z^i) + (1 - \kappa_t) \delta_{x_0^i}(z^i)] = 0$.

The main difference here is that, for the numerator, we only consider x_0 such that $x_0^i = x^i$. **Case 4** makes the entire product term zero, so it is excluded. For remaining dimension $j \neq i$, the Hamming distance between x_1 and z is fixed for this term, it can be treated as a constant. we can follow the same logical flow as [derivation of Eqn. 26 in Sec. A.1](#) to parse the term in each the three cases with additional constraint: x_0 with $x_0^i = x^i$. Note that the following equations correspond to the (\triangle) part, restricted to the dimensions $j \neq i$.

Case 1. Same as dimension $N - 1$ with $h_+(d_m, z) - 1$ since $d_m^i = z^i$:

$$(1 - \kappa_t)^{(N-1) - (h_+(d_m, z) - 1)} \sum_{k=0}^{h_+(d_m, z) - 1} \binom{h_+(d_m, z) - 1}{k} ((K - 1) \kappa_t)^{h_+(d_m, z) - 1 - k} \quad (50)$$

$$= (1 - \kappa_t)^{N - h_+(d_m, z)} (1 + (K - 1) \kappa_t)^{h_+(d_m, z) - 1}. \quad (51)$$

Case 2. Same as dimension $N - 1$ with $h_+(d_m, z) - 1$ (same as above):

$$(1 - \kappa_t)^{(N-1) - (h_+(d_m, z) - 1)} \sum_{k=0}^{h_+(d_m, z) - 1} \binom{h_+(d_m, z) - 1}{k} ((K - 1) \kappa_t)^{h_+(d_m, z) - 1 - k} \quad (52)$$

$$= (1 - \kappa_t)^{N - h_+(d_m, z)} (1 + (K - 1) \kappa_t)^{h_+(d_m, z) - 1}. \quad (53)$$

Case 3. Same as dimension $N - 1$ with $h_+(d_m, z)$ since $d_m^i \neq z^i$:

$$(1 - \kappa_t)^{(N-1) - h_+(d_m, z)} \sum_{k=0}^{h_+(d_m, z)} \binom{h_+(d_m, z)}{k} ((K - 1) \kappa_t)^{h_+(d_m, z) - k} \quad (54)$$

$$= (1 - \kappa_t)^{N - h_+(d_m, z) - 1} (1 + (K - 1) \kappa_t)^{h_+(d_m, z)}. \quad (55)$$

810 **Case 4.** Do not need to consider for dimension $j \neq i$ because $\left[\kappa_t \delta_{d_m^i}(z^i) + (1 - \kappa_t) \delta_{x_0^i}(z^i) \right] = 0$.
 811

812 Note that the denominator corresponds to the same term derived in Sec. A.1. We divide the analysis
 813 into two categories. First, we consider the case where $z^i = x^i$, which corresponds to **Cases 1** and
 814 **3**. The joint expression, incorporating both the $j = i$ (□) and $j \neq i$ (△) cases, can be written as
 815 $\delta_{d_m^i}(z^i) \cdot \text{Case 1} + (1 - \delta_{d_m^i}(z^i)) \cdot \text{Case 3}$. The sum of these two cases can be written as:
 816

$$\frac{\sum_{m=1}^M \left[\delta_{d_m^i}(z^i) \frac{1}{1-\kappa_t} + (1 - \delta_{d_m^i}(z^i)) \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right) \right] (1 - \kappa_t)^N \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z)-1}}{\sum_{m=1}^M (1 - \kappa_t)^N \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z)}} \quad (56)$$

$$= \frac{\sum_{m=1}^M \left[\delta_{d_m^i}(z^i) \frac{1}{1-\kappa_t} + (1 - \delta_{d_m^i}(z^i)) \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right) \right] \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z)-1}}{\sum_{m=1}^M \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z)}} \quad (57)$$

$$= \frac{\sum_{m=1}^M \left[\delta_{d_m^i}(z^i) \frac{1}{1-\kappa_t+K\kappa_t} + (1 - \delta_{d_m^i}(z^i)) \right] \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z)}}{\sum_{m=1}^M \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z)}}. \quad (58)$$

830 We define $\gamma := 1 + \frac{\kappa_t}{1-\kappa_t} K$,

$$= \frac{\sum_{m=1}^M \left[\delta_{d_m^i}(z^i) \frac{1}{1-\kappa_t+K\kappa_t} + (1 - \delta_{d_m^i}(z^i)) \right] \gamma^{h_+(d_m, z)}}{\sum_{m=1}^M \gamma^{h_+(d_m, z)}} \quad (59)$$

$$= \sum_{m=1}^M \left[\delta_{d_m^i}(z^i) \frac{1}{1 - \kappa_t + K\kappa_t} + (1 - \delta_{d_m^i}(z^i)) \right] \frac{\gamma^{h_+(d_m, z)}}{\sum_{m=1}^M \gamma^{h_+(d_m, z)}} \quad (60)$$

$$= 1 - \frac{\kappa_t(K-1)}{1 - \kappa_t + K\kappa_t} \sum_{m=1}^M \delta_{d_m^i}(z^i) \frac{\gamma^{h_+(d_m, z)}}{\sum_{m=1}^M \gamma^{h_+(d_m, z)}}. \quad (61)$$

842 Second, when $z^i \neq x^i$, the expression of **Case 2** can be simplified as follows:
 843

$$\frac{\sum_{m=1}^M \delta_{d_m^i}(z^i) \frac{\kappa_t}{1-\kappa_t} (1 - \kappa_t)^N \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z)-1}}{\sum_{m=1}^M (1 - \kappa_t)^N \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z)}} \quad (62)$$

$$= \frac{\sum_{m=1}^M \delta_{d_m^i}(z^i) \frac{\kappa_t}{1-\kappa_t} \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z)-1}}{\sum_{m=1}^M \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z)}} \quad (63)$$

$$= \frac{\sum_{m=1}^M \delta_{d_m^i}(z^i) \frac{\kappa_t}{1-\kappa_t+K\kappa_t} \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z)}}{\sum_{m=1}^M \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z)}} \quad (64)$$

$$= \frac{\sum_{m=1}^M \delta_{d_m^i}(z^i) \frac{\kappa_t}{1-\kappa_t+K\kappa_t} \gamma^{h_+(d_m, z)}}{\sum_{m=1}^M \gamma^{h_+(d_m, z)}} \quad (65)$$

$$= \sum_{m=1}^M \delta_{d_m^i}(z^i) \frac{\kappa_t}{1 - \kappa_t + K\kappa_t} \frac{\gamma^{h_+(d_m, z)}}{\sum_{m=1}^M \gamma^{h_+(d_m, z)}} \quad (66)$$

$$= \frac{\kappa_t}{1 - \kappa_t + K\kappa_t} \sum_{m=1}^M \delta_{d_m^i}(z^i) \frac{\gamma^{h_+(d_m, z)}}{\sum_{m=1}^M \gamma^{h_+(d_m, z)}}. \quad (67)$$

864 We can then express the closed-form backward velocity as follows:
 865

$$866 \quad \check{v}_t(x^i, z) = \frac{\dot{\kappa}_t}{\kappa_t} [\delta_{z^i}(x^i) - p_{0|t}(x^i|z)] \quad (68)$$

$$867 \quad = \begin{cases} \frac{\dot{\kappa}_t(K-1)}{1-\kappa_t+K\kappa_t} \sum_{m=1}^M \delta_{d_m^i}(z^i) \frac{\gamma^{h+(d_m, z)}}{\sum_{m=1}^M \gamma^{h+(d_m, z)}} & \text{if } x^i = z^i \\ \frac{-\dot{\kappa}_t}{1-\kappa_t+K\kappa_t} \sum_{m=1}^M \delta_{d_m^i}(z^i) \frac{\gamma^{h+(d_m, z)}}{\sum_{m=1}^M \gamma^{h+(d_m, z)}} & \text{otherwise} \end{cases} \quad (69)$$

$$871 \quad = \frac{\dot{\kappa}_t(\delta_{x^i}(z^i)K-1)}{1-\kappa_t+K\kappa_t} \sum_{m=1}^M \delta_{d_m^i}(z^i) \frac{\gamma^{-h(d_m, z)}}{\sum_{m=1}^M \gamma^{-h(d_m, z)}}. \quad (70)$$

875 Since $\kappa_t = 1$ for $t = 1$, then $\gamma \rightarrow \infty$. So this equation is formally not defined at $t = 1$. Nevertheless,
 876 as $\lim_{t \rightarrow 1}$, the weighted sum over power of γ is dominated by the maximum term, which converges
 877 to 1. Hence, the expression can be rigorously interpreted as $\lim_{t \rightarrow 1} \check{v}_t(x^i, z)$, and in practice, this
 878 limiting value is used for sampling at $t = 1$.

880 B EXPERIMENT DETAILS

882 In Tab. 2, we summarize the hyperparameters used in the experiments presented in Sec. 5, covering
 883 both training and finetuning configurations for each dataset. All reported samples were generated
 884 using the greedy-tail denoiser described in (Sahoo et al., 2025). We employed an implementation of
 885 the closed-form backward velocity that is optimized at the CUDA level.

886 For the CIFAR-10 dataset, we follow the same setting as baseline (Schiff et al., 2025). Tab. 3 reports
 887 the FID and IS of baseline and PAIRFLOW measured with 1,000 steps and 50K samples, which are
 888 consistent with the results reported in Tab. 6 of (Schiff et al., 2025) and PAIRFLOW outperforms it.

890 Table 2: Summary of the training settings used in Sec. 5. Specifically, “Sampling Steps” under PAIR-
 891 FLOW and ReDi (Yoo et al., 2025) indicate the number of steps taken to generate pairs, “Teacher
 892 Update Period” under DCD (Sahoo et al., 2025) denotes the number of fine-tuning iterations between
 893 updates, when the teacher model is replaced by the current student model. “# Pairs” under ReDi (Yoo
 894 et al., 2025) denotes the number of pairs for the fine-tuning.

	MNIST-Binary	QM9	CIFAR10	ZINC-250k
Training Iterations	10K	50K	300K	200K
Data Dimension	28×28	32	$32 \times 32 \times 3$	72
Batch Size	256	1024	512	256
Network Architecture	U-Net	Transformer	U-Net	Transformer
Parameter Count	25.8M	92.4M	35.8M	92.4M
	PAIRFLOW			
Sampling Steps	20	20	20	64
	DCD (Sahoo et al., 2025)			
Training Iterations	5K	10K	50K	30K
Teacher Update Period	1K	2K	10K	5K
	ReDi (Yoo et al., 2025)			
Training Iterations	5K	10K	50K	30K
# Pairs	10K	20K	10K	20K
Sampling Steps	256	256	1024	256

913 C ADDITIONAL EXPERIMENTS

915 C.1 COVERAGE OF TRAINING DATASET BY SAMPLING WITH FORWARD VELOCITY

917 As discussed in Sec. 4.1, constructing pairs using the closed-form forward velocity with a training
 918 dataset of size $|X_1|$ incurs significantly higher cost to achieve full coverage compared to using

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000

Table 3: FID (Heusel et al., 2017) and IS (Salimans et al., 2016) of UDLM (Schiff et al., 2025) and PAIRFLOW on the CIFAR-10 dataset (Krizhevsky et al., 2009).

	FID	IS
UDLM	33.65	6.96
PAIRFLOW	28.07	7.37

Table 4: Total Correlation measure with pairs sampled from UDLM (Schiff et al., 2025) and PAIRFLOW trained on QM9 (Ramakrishnan et al., 2014).

	Total Correlation
UDLM	31.87
PAIRFLOW	30.72

Table 5: Summary of training set sizes $|X_1|$ for each dataset, the number of unique samples \bar{k} obtained by simulating paths using the closed-form forward velocity in Eqn. 9, and the corresponding coverage values: empirical (COV) and theoretically predicted (COV_{Pred}).

	QM9	ZINC-250k	MNIST-Binary	CIFAR-10
$ X_1 $	127,190	224,568	60,000	100,000
\bar{k}	77,104	140,779	37,711	63,117
COV	60.62%	62.68%	62.85%	63.11%
COV _{Pred}			63.21%	

the backward velocity. Let k denote the number of source samples drawn from the source (prior) distribution, assumed to be uniform in our work. The probability that a given element in the training set is selected is $\left(1 - \frac{1}{|X_1|}\right)^k$. Accordingly, we denote by \bar{k} the number of unique samples among the k draws, whose expectation is: $\sum_{i=1}^{|X_1|} \left[1 - \left(1 - \frac{1}{|X_1|}\right)^k\right] = |X_1| \left[1 - \left(1 - \frac{1}{|X_1|}\right)^k\right]$. In addition, we define the coverage as the ratio between the number of unique elements obtained through this sampling procedure and the training set size: $\text{COV} = \bar{k}/|X_1|$.

To validate our claim in Sec. 4.1, we sample $k = |X_1|$ data points x_1 by transporting source samples x_0 , independently drawn from the uniform distribution, along the velocity field defined in Eqn. 9. Using these samples, we evaluate the coverage following the definition above. The dataset sizes, number of unique samples among the generated samples, and the empirical and theoretical coverages are summarized in Tab. 5. These findings indicate that, even when sampling the same number of points as the training set size, only about 63% of the training distribution can be recovered in practice. Achieving full coverage would therefore require a substantially larger number of samples, introducing significantly higher computational cost. Motivated by this finding, we instead propose tracing backward from data samples, using a closed-form velocity field that we derive for this purpose (Sec. 4.2).

C.2 TOTAL CORRELATION ANALYSIS OF CLOSED-FORM VELOCITY

As in Sec. 3.2, Yoo et al. (2025) demonstrated that iteratively refining the joint distribution of source-target pairs in discrete flow models improves few-step performance by reducing the total correlation of the model. In this section, we measure the total correlation following their methodology. Specifically, we perform sampling with neural networks, including UDLM (Schiff et al., 2025) and PAIRFLOW trained on QM9 (Ramakrishnan et al., 2014), starting from identical initial states x_0 but with varying random seeds. We randomly select 20,000 initial states x_0 , and for each x_0 , we generate 10 samples with a step size of 256. As shown in Tab. 4, PAIRFLOW achieves a lower total correlation, consistent with the improved performance observed in few-step sampling, as discussed above.

D DETAILED EXPERIMENTAL RESULTS

In this section, we provide the detailed experimental results corresponding to those summarized in Sec. 5. For the molecular datasets (QM9 (Ramakrishnan et al., 2014) and ZINC-250k (Irwin et al., 2012)), we generate 1,024 samples across varying timesteps and evaluate validity, uniqueness, and novelty. Reported values are averaged over 10 trials, with standard deviations also included. For the

972 image domain, we report FID on MNIST-Binary (LeCun et al., 2002), and both FID (Heusel et al.,
973 2017) and IS (Salimans et al., 2016) on CIFAR-10 (Krizhevsky et al., 2009). Detailed experimental
974 settings are provided in Sec. 5.1.

975 Results on QM9 are presented in Tables 6, 7, and 8, reporting validity, uniqueness, and novelty,
976 respectively. Corresponding results on ZINC-250k are shown in Tables 9, 10, and 11. Finally, results
977 for the image datasets are summarized in Tab. 12 (FID on MNIST-Binary) and Tab. 13 (FID and IS
978 on CIFAR-10). The FID measured on MNIST-Binary (LeCun et al., 2002), FID and IS measured on
979 CIFAR-10 (Krizhevsky et al., 2009), are summarized in Tab. 12 and Tab. 13, respectively.
980
981
982
983
984
985
986
987
988
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 Table 6: Validity scores (\uparrow) on QM9 (Ramakrishnan et al., 2014) for various methods across different
 1027 steps. Best values per column are highlighted in bold.

Method	1	2	4	8	16	32	64
MDLM	51.4 \pm 5.7	80.0 \pm 11.1	154.8 \pm 10.0	347.5 \pm 11.5	530.0 \pm 11.6	662.9 \pm 16.5	736.4 \pm 16.4
UDLM	17.5 \pm 3.2	125.5 \pm 11.8	497.6 \pm 8.3	826.6 \pm 10.3	953.5 \pm 6.1	991.9 \pm 4.2	1000.1 \pm 3.5
Random	47.1 \pm 5.7	194.6 \pm 8.2	554.2 \pm 13.3	858.3 \pm 17.5	962.0 \pm 6.7	989.9 \pm 7.6	998.6 \pm 5.3
PAIRFLOW	223.4 \pm 12.7	416.0 \pm 12.4	734.9 \pm 7.2	921.5 \pm 11.0	977.1 \pm 3.9	990.9 \pm 5.9	1000.2 \pm 4.5
UDLM + DCD	323.0 \pm 19.5	530.8 \pm 20.0	816.6 \pm 14.4	941.1 \pm 8.5	981.0 \pm 4.8	993.0 \pm 3.6	999.4 \pm 4.7
PAIRFLOW + DCD	453.8 \pm 16.4	685.8 \pm 16.7	891.6 \pm 11.9	963.1 \pm 7.8	983.7 \pm 8.5	989.3 \pm 3.5	993.2 \pm 5.7
UDLM + Redi	59.7 \pm 8.8	232.4 \pm 9.2	588.4 \pm 15.8	849.6 \pm 14.2	940.5 \pm 8.5	967.5 \pm 5.2	978.8 \pm 5.2
PAIRFLOW + Redi	361.0 \pm 115.2	512.6 \pm 44.2	775.7 \pm 10.0	929.1 \pm 11.6	976.2 \pm 4.5	985.6 \pm 6.7	993.1 \pm 7.1

1040 Table 7: Uniqueness scores (\uparrow) on QM9 (Ramakrishnan et al., 2014) for various methods across
 1041 different steps. Best values per column are highlighted in bold.

Method	1	2	4	8	16	32	64
MDLM	18.9 \pm 3.0	28.8 \pm 3.3	87.4 \pm 8.4	254.7 \pm 10.3	443.8 \pm 16.4	591.0 \pm 18.6	666.9 \pm 17.5
UDLM	17.5 \pm 3.2	125.4 \pm 11.7	495.0 \pm 8.2	819.5 \pm 11.4	943.0 \pm 5.7	979.1 \pm 5.0	990.0 \pm 4.7
Random	47.1 \pm 5.7	194.5 \pm 8.1	551.2 \pm 12.5	853.1 \pm 16.9	953.5 \pm 7.8	981.1 \pm 6.7	989.6 \pm 5.7
PAIRFLOW	223.0 \pm 12.3	414.7 \pm 12.0	731.4 \pm 6.9	917.4 \pm 11.8	971.6 \pm 4.3	986.2 \pm 5.3	994.8 \pm 5.2
UDLM + DCD	320.9 \pm 18.7	528.0 \pm 19.7	808.5 \pm 12.7	932.3 \pm 8.1	970.9 \pm 5.8	978.3 \pm 4.9	987.2 \pm 4.3
PAIRFLOW + DCD	451.8 \pm 15.7	681.6 \pm 16.5	886.5 \pm 12.6	957.8 \pm 7.6	978.7 \pm 9.1	983.4 \pm 3.9	989.0 \pm 5.3
UDLM + Redi	59.7 \pm 8.8	231.6 \pm 9.5	581.0 \pm 15.1	834.7 \pm 11.4	917.5 \pm 9.6	944.8 \pm 6.2	956.1 \pm 5.2
PAIRFLOW + Redi	359.5 \pm 113.1	507.7 \pm 43.0	765.3 \pm 8.8	913.5 \pm 10.2	959.7 \pm 5.8	968.8 \pm 8.4	973.0 \pm 9.6

1054 Table 8: Novelty scores (\uparrow) on QM9 (Ramakrishnan et al., 2014) for various methods across different
 1055 steps. Best values per column are highlighted in bold.

Method	1	2	4	8	16	32	64
MDLM	15.4 \pm 3.1	23.6 \pm 2.6	56.3 \pm 5.2	127.4 \pm 9.4	172.2 \pm 12.4	206.5 \pm 7.4	193.6 \pm 10.4
UDLM	13.8 \pm 2.9	52.0 \pm 8.2	120.0 \pm 3.8	152.4 \pm 9.1	144.2 \pm 12.4	147.2 \pm 9.7	145.1 \pm 9.0
Random	26.9 \pm 5.0	66.5 \pm 7.1	128.8 \pm 13.0	156.8 \pm 10.6	145.5 \pm 9.1	146.5 \pm 7.0	150.1 \pm 8.7
PAIRFLOW	68.8 \pm 7.8	85.6 \pm 10.0	109.2 \pm 7.8	96.8 \pm 9.9	106.5 \pm 12.5	108.9 \pm 9.4	110.0 \pm 9.9
UDLM + DCD	145.9 \pm 7.9	137.5 \pm 9.8	185.5 \pm 10.9	186.2 \pm 13.7	173.4 \pm 15.5	168.0 \pm 10.4	165.4 \pm 11.6
PAIRFLOW + DCD	110.3 \pm 6.3	136.1 \pm 12.0	146.2 \pm 12.9	139.2 \pm 9.6	131.8 \pm 11.7	140.1 \pm 9.6	133.2 \pm 7.9
UDLM + Redi	31.4 \pm 8.0	73.3 \pm 7.1	110.4 \pm 12.2	126.8 \pm 8.9	116.3 \pm 8.4	120.2 \pm 9.0	117.6 \pm 7.1
PAIRFLOW + Redi	84.2 \pm 11.3	92.0 \pm 6.9	101.1 \pm 9.6	98.6 \pm 8.4	98.8 \pm 13.9	95.8 \pm 8.3	98.9 \pm 7.5

1068 Table 9: Validity scores (\uparrow) on ZINC-250k (Irwin et al., 2012) for various methods across different
 1069 steps. Best values per column are highlighted in bold.

Method	1	2	4	8	16	32	64
MDLM	15.0 \pm 4.0	79.4 \pm 4.7	194.6 \pm 15.1	351.3 \pm 20.6	463.7 \pm 17.7	553.9 \pm 10.7	610.1 \pm 17.6
UDLM	0.3 \pm 0.5	65.2 \pm 8.2	435.7 \pm 14.4	775.1 \pm 19.5	887.3 \pm 12.7	921.5 \pm 8.5	937.3 \pm 3.9
Random	0.6 \pm 0.9	68.3 \pm 10.7	351.2 \pm 15.8	569.4 \pm 16.6	611.0 \pm 16.3	602.4 \pm 13.3	571.0 \pm 13.2
PAIRFLOW	9.9 \pm 2.3	146.3 \pm 10.4	533.9 \pm 13.9	799.4 \pm 9.2	873.2 \pm 14.1	901.0 \pm 14.2	907.8 \pm 7.7
UDLM + DCD	25.7 \pm 4.7	323.9 \pm 12.5	718.2 \pm 13.5	873.5 \pm 8.4	919.8 \pm 10.0	933.1 \pm 5.9	942.2 \pm 4.3
PAIRFLOW + DCD	114.9 \pm 14.3	436.3 \pm 16.5	725.1 \pm 11.5	858.2 \pm 10.0	896.5 \pm 8.2	900.9 \pm 9.5	907.1 \pm 13.7
UDLM + Redi	0.7 \pm 0.8	75.9 \pm 7.9	424.8 \pm 16.6	734.4 \pm 8.6	856.3 \pm 11.3	892.2 \pm 10.4	900.1 \pm 10.8
PAIRFLOW + Redi	46.3 \pm 6.3	221.5 \pm 11.0	562.8 \pm 12.7	793.6 \pm 8.4	869.3 \pm 14.2	897.4 \pm 9.1	907.0 \pm 10.5

1080 Table 10: Uniqueness scores (\uparrow) on ZINC-250k (Irwin et al., 2012) for various methods across
 1081 different steps. Best values per column are highlighted in bold.

Method	1	2	4	8	16	32	64
MDLM	7.4 \pm 2.7	32.7 \pm 2.4	105.1 \pm 8.3	256.0 \pm 18.3	392.6 \pm 16.6	511.3 \pm 14.0	582.4 \pm 17.3
UDLM	0.3 \pm 0.5	65.2 \pm 8.2	435.7 \pm 14.4	775.1 \pm 19.5	887.3 \pm 12.7	921.5 \pm 8.5	937.2 \pm 3.8
Random	0.6 \pm 0.9	68.3 \pm 10.7	351.2 \pm 15.8	569.4 \pm 16.6	611.0 \pm 16.3	602.4 \pm 13.3	571.0 \pm 13.2
PAIRFLOW	9.9 \pm 2.3	146.3 \pm 10.4	533.9 \pm 13.9	799.4 \pm 9.2	873.2 \pm 14.1	901.0 \pm 14.2	907.8 \pm 7.7
UDLM + DCD	25.7 \pm 4.7	323.9 \pm 12.5	718.2 \pm 13.5	873.5 \pm 8.4	919.8 \pm 10.0	933.1 \pm 5.9	942.2 \pm 4.3
PAIRFLOW + DCD	114.9 \pm 14.3	436.3 \pm 16.5	725.1 \pm 11.5	858.2 \pm 10.0	896.5 \pm 8.2	900.9 \pm 9.5	907.1 \pm 13.7
UDLM + Redi	0.7 \pm 0.8	75.9 \pm 7.9	424.8 \pm 16.6	734.3 \pm 8.7	856.3 \pm 11.3	892.2 \pm 10.4	900.0 \pm 10.9
PAIRFLOW + Redi	46.3 \pm 6.3	221.5 \pm 11.0	562.8 \pm 12.7	793.6 \pm 8.4	869.3 \pm 14.2	897.4 \pm 9.1	907.0 \pm 10.5

1094 Table 11: Novelty scores (\uparrow) on ZINC-250k (Irwin et al., 2012) for various methods across different
 1095 steps. Best values per column are highlighted in bold.

Method	1	2	4	8	16	32	64
MDLM	3.8 \pm 2.7	24.2 \pm 2.6	84.5 \pm 7.4	228.3 \pm 17.1	372.0 \pm 15.2	494.5 \pm 15.2	569.3 \pm 17.1
UDLM	0.3 \pm 0.5	65.2 \pm 8.2	435.7 \pm 14.4	775.1 \pm 19.5	887.3 \pm 12.7	921.3 \pm 8.8	936.9 \pm 4.1
Random	0.6 \pm 0.9	68.3 \pm 10.7	351.2 \pm 15.8	569.4 \pm 16.6	611.0 \pm 16.3	602.4 \pm 13.3	571.0 \pm 13.2
PAIRFLOW	9.9 \pm 2.3	146.3 \pm 10.4	533.9 \pm 13.9	799.4 \pm 9.2	873.2 \pm 14.1	901.0 \pm 14.2	907.8 \pm 7.7
UDLM + DCD	25.7 \pm 4.7	323.9 \pm 12.5	718.2 \pm 13.5	873.5 \pm 8.4	919.7 \pm 9.9	933.0 \pm 5.8	942.2 \pm 4.3
PAIRFLOW + DCD	114.9 \pm 14.3	436.3 \pm 16.5	725.1 \pm 11.5	858.2 \pm 10.0	896.4 \pm 8.3	900.9 \pm 9.5	907.1 \pm 13.7
UDLM + Redi	0.7 \pm 0.8	75.9 \pm 7.9	424.8 \pm 16.6	734.3 \pm 8.7	856.3 \pm 11.3	892.1 \pm 10.5	900.0 \pm 10.9
PAIRFLOW + Redi	46.3 \pm 6.3	221.5 \pm 11.0	562.8 \pm 12.7	793.5 \pm 8.3	869.3 \pm 14.2	897.4 \pm 9.1	907.0 \pm 10.5

1108 Table 12: FID (\downarrow) on MNIST-Binary (LeCun et al., 2002) for various methods across different steps.
 1109 Best values per column are bolded.

Method	1	2	4	8	16	32	64
MDLM	204.64	159.26	103.74	54.41	28.51	12.31	7.01
UDLM	130.57	42.54	11.25	5.70	4.69	4.77	5.01
Random	128.57	36.59	9.41	5.60	5.00	5.10	5.19
PAIRFLOW	40.58	15.61	8.50	5.97	5.55	5.24	5.17
UDLM + DCD	53.84	16.09	8.06	7.46	7.12	6.52	6.65
PAIRFLOW + DCD	19.51	14.20	11.47	9.28	7.75	7.82	8.42
UDLM + ReDi	18.44	10.35	8.11	6.65	6.56	6.55	6.49
PAIRFLOW + ReDi	12.90	9.38	8.54	6.85	6.79	6.96	6.94

E EXPERIMENT ON CONTINUOUS FLOW MATCHING

Alongside our main experiments in the discrete setting, we also demonstrate the potential of our method to extend to continuous domains, as illustrated by the toy experiment presented below. Here, we denote by PAIRFLOW a continuous flow model trained on source–target pairs constructed using the continuous variant of the algorithm described in Sec. 4.2.

E.1 CLOSED-FORM VELOCITY IN CONTINUOUS FLOW MATCHING

Setup. Let $X_0 \sim p_0$ (source), $X_1 \sim q$ (target) be independent random variables in \mathbb{R}^N and consider the linear probability path

$$X_t = (1 - t)X_0 + tX_1, \quad t \in [0, 1]. \quad (71)$$

1134 Table 13: FID (\downarrow) and IS (\uparrow) on CIFAR-10 (Krizhevsky et al., 2009) for various methods across
 1135 different timesteps. Best values per column are bolded.

Method	1	2	4	8	16	32	64	128	256	512	1024
FID (\downarrow)											
MDLM	407.31	359.97	340.43	340.98	321.79	228.34	131.63	77.44	52.89	42.88	39.04
UDLM	340.47	321.98	255.95	151.60	91.31	62.43	49.06	42.26	40.18	37.63	37.60
Random	328.83	314.64	220.16	127.66	82.76	58.42	45.18	39.50	36.78	35.16	35.07
PAIRFLOW	269.87	260.29	192.62	112.78	73.87	52.22	40.06	34.42	33.83	32.42	31.85
UDLM + DCD	318.33	282.38	204.15	108.30	75.56	69.87	74.62	79.35	84.34	85.79	87.87
PAIRFLOW + DCD	223.86	190.15	139.46	95.91	82.11	80.70	86.59	100.16	114.25	123.41	129.45
UDLM + ReDi	251.02	212.20	171.24	148.18	138.18	131.94	125.81	123.77	123.39	121.52	121.06
PAIRFLOW + ReDi	275.09	250.87	184.45	119.04	89.46	74.46	66.22	61.83	61.35	60.10	59.53
IS (\uparrow)											
MDLM	1.21	1.22	1.24	1.31	1.52	2.47	4.03	5.08	5.66	6.18	6.38
UDLM	1.32	1.48	2.11	3.54	4.85	5.90	6.25	6.56	6.66	6.87	6.81
Random	1.37	1.52	2.47	4.01	5.24	6.08	6.56	6.75	6.87	7.07	7.05
PAIRFLOW	1.72	1.80	2.70	4.27	5.61	6.23	6.86	7.00	7.12	7.14	7.33
UDLM + DCD	1.49	1.60	2.23	3.90	5.09	5.45	5.42	5.16	5.09	5.07	4.84
PAIRFLOW + DCD	2.21	2.38	3.23	4.38	4.99	5.13	4.97	4.48	4.26	3.82	3.81
UDLM + ReDi	1.85	2.35	2.96	3.37	3.52	3.67	3.86	3.96	3.93	3.99	3.94
PAIRFLOW + ReDi	1.80	1.95	2.91	4.18	5.02	5.65	5.88	6.03	6.06	6.09	6.30

1158 Table 14: FID of PAIRFLOW on MNIST (LeCun et al., 2002) with continuous values, measured using
 1159 FID over 50K samples across various timesteps. Best values are bolded.

Method	1	2	4	8	16	32	64
CondOT	398.43	91.17	27.34	10.80	5.81	3.99	3.16
PAIRFLOW	74.89	14.40	6.89	3.78	2.70	2.42	2.37
CondOT+RF	32.70	9.12	5.46	4.24	3.93	3.56	3.24
PAIRFLOW+RF	28.61	4.15	3.01	2.87	2.89	2.91	2.94

1170 For flow matching with the linear path Eqn. 71, the optimal velocity field equals the conditional drift:
 1171

$$v_t(x) = \mathbb{E}[X_1 - X_0 | X_t = x]. \quad (72)$$

1175 We derive a closed form of Eqn. 72 that is directly computable from p_0 and q .

1178 **Derivation.** By Bayes' rule with a Dirac constraint for the linear relation Eqn. 71,

$$p(x_0, x_1 | X_t = x) \propto p_0(x_0) q(x_1) \delta(x - (1-t)x_0 - tx_1). \quad (73)$$

1184 Integrating out x_0 using $\delta(Ay - b) = |\det A|^{-1} \delta(y - A^{-1}b)$ with $A = (1-t)I$ gives

$$p(x_1 | X_t = x) \propto q(x_1) (1-t)^{-D} p_0\left(\frac{x - tx_1}{1-t}\right). \quad (74)$$

1188 Hence,

$$1190 \quad v_t(x) = \frac{\iint (x_1 - x_0) p(x_0, x_1 | X_t = x) dx_0 dx_1}{\iint p(x_0, x_1 | X_t = x) dx_0 dx_1} \quad (75)$$

$$1194 \quad = \frac{\iint (x_1 - x_0) p_0(x_0) q(x_1) \delta(x - (1-t)x_0 - tx_1) dx_0 dx_1}{\iint p_0(x_0) q(x_1) \delta(x - (1-t)x_0 - tx_1) dx_0 dx_1} \quad (76)$$

$$1198 \quad = \frac{\int q(x_1) p_0\left(\frac{x-tx_1}{1-t}\right) \left(x_1 - \frac{x-tx_1}{1-t}\right) dx_1}{\int q(x_1) p_0\left(\frac{x-tx_1}{1-t}\right) dx_1}. \quad (77)$$

1203 Observing $x_1 - \frac{x-tx_1}{1-t} = \frac{x_1-x}{1-t}$, we obtain the compact form

$$1205 \quad v_t(x) = \frac{1}{1-t} \frac{\int q(x_1) p_0\left(\frac{x-tx_1}{1-t}\right) (x_1 - x) dx_1}{\int q(x_1) p_0\left(\frac{x-tx_1}{1-t}\right) dx_1}. \quad (78)$$

1210 When we have a dataset with samples $\{d_m\}_{m=1}^M$, the target distribution q is approximated by the
1211 empirical measure $q(x_1) \approx \frac{1}{M} \sum_{m=1}^M \delta_{d_m}(x_1)$, then Eqn. 78 reduces as follow:

$$1213 \quad v_t(x) = \frac{1}{1-t} \frac{\sum_{m=1}^M p_0\left(\frac{x-td_m}{1-t}\right) (d_m - x)}{\sum_{m=1}^M p_0\left(\frac{x-td_m}{1-t}\right)}. \quad (79)$$

1217 When p_0 is standard Gaussian, $p_0(y) = (2\pi)^{-D/2} \exp\left(-\frac{1}{2}\|y\|_2^2\right)$, the normalizing constants cancel
1218 in Eqn. 79, yielding the closed form velocity:

$$1220 \quad v_t(x) = \frac{1}{1-t} \frac{\sum_{m=1}^M \exp\left(-\frac{1}{2}\left\|\frac{x-td_m}{1-t}\right\|_2^2\right) (d_m - x)}{\sum_{m=1}^M \exp\left(-\frac{1}{2}\left\|\frac{x-td_m}{1-t}\right\|_2^2\right)}. \quad (80)$$

1225 This formulation has already been introduced in previous works (Karras et al., 2022; Bertrand et al.,
1226 2025); however, to the best of our knowledge, no prior work has extended this idea to designing
1227 couplings for accelerating flow models using the re-flow technique (Liu & Gong, 2023a). In the
1228 continuous domain, the backward velocity can be obtained directly by flipping the sign of the forward
1229 velocity. In contrast, in the discrete domain, the corresponding expression does not converge as
1230 $\lim_{t \rightarrow 1}$, and thus the backward velocity cannot be employed for sampling starting from data points.
1231 Therefore, in this section, we perform experiments using the forward velocity.

1232 E.2 CONTINUOUS FLOW MATCHING ON MNIST

1235 We train rectified flow models on MNIST (LeCun et al., 2002) using two pairing strategies: (i)
1236 independent pairing (baseline) and (ii) closed-form pairing as described in Sec. E.1. We adopt
1237 CondOT (Lipman et al., 2023) as our base flow model, which is originally trained with a independent
1238 pairing. We denote the variant of CondOT trained on pairs generated by the closed-form forward
1239 velocity as PAIRFLOW. To enable a few-step sampling, we subsequently apply rectification distillation
1240 (ReFlow (Liu & Gong, 2023a)) to each pretrained model, denoted by the suffix “+RF”.

1241 We use an NCSN++-style U-Net backbone (Song et al., 2021) with a base width of 64 and 3
1242 downsampling stages (doubling channels at each stage), optimized using Adam (Kingma & Ba, 2014)

1242 with a learning rate of 2×10^{-4} . The pretraining takes 500 epochs. The distillation stage requires
1243 200 epochs with a learning rate of 2×10^{-5} .
1244

1245 Tab. 14 summarizes performance at various sampling steps. Without distillation, closed-form pairing
1246 (PAIRFLOW) yields significantly better FID in the few-step settings and maintains the performance
1247 in the many-step settings, relative to the baseline. With distillation (ReFlow (Liu & Gong, 2023a),
1248 our method still shows better performance: PAIRFLOW+RF achieves a lower FID in every sampling
1249 budget than ReFlow applied to the baseline. These results show that closed-form pairing benefits both
1250 undistilled and distilled flow models, with especially large gains when the sampling steps are small.
1251

1252 E.3 CONTINUOUS RECTIFIED FLOWS ON DIMENSION-VARYING SYNTHETIC DATA

1253 To assess scalability, we construct an N -fold product of the standard two-moons distribution, yielding
1254 a dataset in \mathbb{R}^{2N} . We consider dimensions $d \in \{2, 4, 8, 16, 32, 64, 128, 256\}$ (i.e., $d = 2N$) and train
1255 rectified flow models with and without closed-form pairing under a common training setup. The
1256 architecture is a simple transformer-based encoder with depth 8, where the hidden size increases with
1257 dimension as 32, 64, 128, 192, 256, 384, 512, 768, respectively.
1258

1259 For the synthetic experiments we report the Chamfer distance (log scale) between 50,000 training
1260 datapoints and 5,000 generated samples. Since the dataset is an N -fold product of 2D two-moons,
1261 Chamfer distance is computed using only the first two coordinates to keep the metric scale consistent
1262 across d and measure fidelity to the base 2D geometry.
1263

1264 Fig. 7 shows the quantitative results. At low dimensions, closed-form pairing yields substantial
1265 improvements over the independently paired baseline. However, as the data dimension increases, we
1266 observe that the magnitude of the improvement decreases. This trend suggests a practical limitation
1267 of closed-form pairing for high-dimensional continuous data.
1268

1269 F ADDITIONAL QUALITATIVE RESULTS

1270 In addition to Fig. 4 in the main paper, we further visualize the 1-step and 2-step generation results
1271 for MNIST-Binary (LeCun et al., 2002) in Fig. 9 and Fig. 10. As discussed in Sec. 5.3, PAIRFLOW
1272 outperforms the base models (Schiff et al., 2025; Sahoo et al., 2024) and achieves comparable
1273 quality to the base models combined with acceleration methods (Sahoo et al., 2025; Yoo et al., 2025).
1274 Additional visualizations for CIFAR-10 (Krizhevsky et al., 2009) with 64- and 256-step generations
1275 are shown in Fig. 8, demonstrating that our method outperforms the other base models.
1276

1278 G THE USE OF LARGE LANGUAGE MODELS

1279 We used a large language model for grammar checking and minor language polishing during the
1280 writing process.
1281

1284 H MORE CLEAR PROOF FOR CLOSED-FORM BACKWARD VELOCITY

1285 In this section, we present a clearer and more refined version of the proof for the closed-form
1286 backward velocity, improving upon the original exposition in Sec. A.2.
1287

1289 H.1 PROOF OF CLOSED-FORM BACKWARD VELOCITY IN DISCRETE FLOW MODELS

1290 Similarly to the proof of the closed-form forward velocity in Sec. A.1, we start from the backward
1291 velocity in Eqn. 11:
1292

$$1294 \check{v}_t(x^i, z) = \frac{\dot{\kappa}_t}{\kappa_t} [\delta_{z^i}(x^i) - p_{0|t}(x^i|z)]. \quad (81)$$

1296 We derive the closed-form noise predictor as follows:
1297

$$1298 \quad p_{0|t}(x^i|z) = \sum_{x_0, x_1} \delta_{x_0^i}(x^i) p_t(x_0, x_1|z) \quad (82)$$

$$1299 \quad = \sum_{x_0, x_1} \delta_{x_0^i}(x^i) \frac{p_t(x_0, x_1, z)}{p_t(z)} \quad (83)$$

$$1300 \quad = \frac{\sum_{x_0, x_1} \delta_{x_0^i}(x^i) p_t(x_0, x_1, z)}{p_t(z)} \quad (84)$$

$$1301 \quad = \frac{\sum_{x_0, x_1} \delta_{x_0^i}(x^i) p_t(x_0, x_1, z)}{\sum_{x_0, x_1} p_t(x_0, x_1, z)}, \quad (85)$$

1302 The last expression is further expanded to:
1303

$$1304 \quad p_{0|t}(x^i|z) = \frac{\sum_{x_0, x_1} \delta_{x_0^i}(x^i) p_t(x_0, x_1, z)}{\sum_{x_0, x_1} p_t(x_0, x_1, z)} \quad (86)$$

$$1305 \quad = \frac{\sum_{x_0, x_1} \delta_{x_0^i}(x^i) p_1(x_1) p_t(z|x_0, x_1)}{\sum_{x_0, x_1} p_1(x_1) p_t(z|x_0, x_1)} \quad (87)$$

$$1306 \quad = \frac{\sum_{m=1}^M \sum_{x_0} \delta_{x_0^i}(x^i) \prod_{j=1}^N [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)]}{\sum_{m=1}^M \sum_{x_0} \prod_{j=1}^N [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)]}. \quad (88)$$

1307
1308 For the denominator, we use the same formula as in Eqn. 31:
1309

$$1310 \quad \sum_{m=1}^M \sum_{x_0} \prod_{j=1}^N [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)] = \sum_{m=1}^M (1 - \kappa_t)^N \left(1 + \frac{\kappa_t}{1 - \kappa_t} K \right)^{h_+(d_m, z)}. \quad (89)$$

1311 Next, for the numerator, we can rewrite it as:
1312

$$1313 \quad \sum_{m=1}^M \sum_{x_0} \delta_{x_0^i}(x^i) \prod_{j=1}^N [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)] \quad (90)$$

$$1314 \quad = \sum_{m=1}^M \sum_{\substack{x_0 \text{ with} \\ x_0^i = x^i}} \prod_j [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)]. \quad (91)$$

1315 The i -th index should be considered separately as x_0^i is set to be equal to x^i . Separating the $j = i$ term from the product yields
1316

$$1317 \quad \sum_{m=1}^M \sum_{\substack{x_0 \text{ with} \\ x_0^i = x^i}} [\kappa_t \delta_{d_m^i}(z^i) + (1 - \kappa_t) \delta_{x_0^i}(z^i)] \prod_{j \neq i} [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)] \quad (92)$$

$$1318 \quad = \sum_{m=1}^M \sum_{\substack{x_0 \text{ with} \\ x_0^i = x^i}} [\kappa_t \delta_{d_m^i}(z^i) + (1 - \kappa_t) \delta_{x_0^i}(z^i)] \prod_{j \neq i} [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)] \quad (93)$$

$$1319 \quad = \sum_{m=1}^M [\kappa_t \delta_{d_m^i}(z^i) + (1 - \kappa_t) \delta_{x_0^i}(z^i)] \sum_{\substack{x_0 \text{ with} \\ j \neq i}} \prod_{\substack{x_0^i = x^i}} [\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j)]. \quad (94)$$

1320 Since the i -th coordinate of x_0 is fixed to x^i , the summation over x_0 with $x_0^i = x^i$ no longer depends on this index. Consequently, when we consider the summation only over the remaining coordinates $j \neq i$, the resulting expression takes exactly the same form as the computation presented in Sec. A.1 (Eqn. 27-Eqn. 31). The only differences are that (i) the effective dimensionality of the product is

reduced from N to $N - 1$, and (ii) the matching count term must exclude the i -th coordinate, yielding $h_+(d_m, z)$ to $h_+(d_m, z) - \delta_{d_m^i}(z^i)$. Reflecting these adjustments, we obtain

$$\sum_{\substack{x_0 \text{ with } j \neq i \\ x_0^i = x^i}} \prod \left[\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j) \right] = (1 - \kappa_t)^{N-1} \left(1 + \frac{\kappa_t}{1 - \kappa_t} K \right)^{h_+(d_m, z) - \delta_{d_m^i}(z^i)}, \quad (95)$$

and

$$\begin{aligned}
& \sum_{m=1}^M \left[\kappa_t \delta_{d_m^i}(z^i) + (1 - \kappa_t) \delta_{x^i}(z^i) \right] \sum_{\substack{x_0 \text{ with } j \neq i \\ x_0^i = x^i}} \left[\kappa_t \delta_{d_m^j}(z^j) + (1 - \kappa_t) \delta_{x_0^j}(z^j) \right]
\end{aligned} \tag{96}$$

$$= \sum_{m=1}^M [\kappa_t \delta_{d_m^i}(z^i) + (1 - \kappa_t) \delta_{x^i}(z^i)] (1 - \kappa_t)^{N-1} \left(1 + \frac{\kappa_t}{1 - \kappa_t} K\right)^{h_+(d_m, z) - \delta_{d_m^i}(z^i)} \quad (97)$$

$$= \sum_{m=1}^M \left[\frac{\kappa_t}{1-\kappa_t} \delta_{d_m^i}(z^i) + \delta_{x^i}(z^i) \right] (1-\kappa_t)^N \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z) - \delta_{d_m^i}(z^i)}. \quad (98)$$

Using this nominator, we can denote the closed-form noise predictor

$$p_{0|t}(x^i|z) = \frac{\sum_{m=1}^M \left[\frac{\kappa_t}{1-\kappa_t} \delta_{d_m^i}(z^i) + \delta_{x^i}(z^i) \right] (1-\kappa_t)^N \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z) - \delta_{d_m^i}(z^i)}}{\sum_{m=1}^M (1-\kappa_t)^N \left(1 + \frac{\kappa_t}{1-\kappa_t} K \right)^{h_+(d_m, z)}} \quad (99)$$

$$= \sum_{m=1}^M \left[\frac{\kappa_t}{1-\kappa_t} \delta_{d_m^i}(z^i) + \delta_{x^i}(z^i) \right] \frac{\left(1 + \frac{\kappa_t}{1-\kappa_t} K\right)^{h_+(d_m, z) - \delta_{d_m^i}(z^i)}}{\sum_{m'=1}^M \left(1 + \frac{\kappa_t}{1-\kappa_t} K\right)^{h_+(d_{m'}, z)}} \quad (100)$$

$$= \sum_{m=1}^M \left[\frac{\kappa_t}{1-\kappa_t} \delta_{d_m^i}(z^i) + \delta_{x^i}(z^i) \right] \left[1 - \frac{K\kappa_t \delta_{d_m^i}(z^i)}{1 + (K-1)\kappa_t} \right] \frac{\gamma^{h+(d_m, z)}}{\sum_{m'=1}^M \gamma^{h+(d_{m'}, z)}} \quad (101)$$

$$= \delta_{x^i}(z^i) - \frac{\kappa_t (K \delta_{x^i}(z^i) - 1)}{1 + (K-1)\kappa_t} \sum_{m=1}^M \delta_{d_m^i}(z^i) \frac{\gamma^{h+(d_m, z)}}{\sum_{m'=1}^M \gamma^{h+(d_{m'}, z)}}, \quad (102)$$

where $\gamma := 1 + \frac{\kappa_t}{1-\kappa_t} K$.

We can then express the closed-form backward velocity as follows:

$$\check{v}_t(x^i, z) = \frac{\dot{\kappa}_t}{\kappa_t} [\delta_{z^i}(x^i) - p_{0|t}(x^i|z)] \quad (103)$$

$$= \frac{\dot{\kappa}_t (K \delta_{x^i}(z^i) - 1)}{1 + (K-1)\kappa_t} \sum_{m=1}^M \delta_{d_m^i}(z^i) \frac{\gamma^{h_+(d_m, z)}}{\sum_{m'=1}^M \gamma^{h_+(d_{m'}, z)}} \quad (104)$$

$$= \frac{\dot{\kappa}_t(K\delta_{x^i}(z^i) - 1)}{1 + (K-1)\kappa_t} \sum_{m=1}^M \delta_{d_m^i}(z^i) \frac{\gamma^{-h(d_m, z)}}{\sum_{m'=1}^M \gamma^{-h(d_{m'}, z)}}. \quad (105)$$

Since $\kappa_t = 1$ for $t = 1$, then $\gamma \rightarrow \infty$. So this equation is formally not defined at $t = 1$. Nevertheless, as $\lim_{t \rightarrow 1}$, the weighted sum over power of γ is dominated by the maximum term, which converges to 1. Hence, the expression can be rigorously interpreted as $\lim_{t \rightarrow 1} \check{v}_t(x^i, z)$, and in practice, this limiting value is used for sampling at $t = 1$.

I RE-FLOW ITERATION RESULTS

This section details the results of applying the iterative re-flow procedure (Liu & Gong, 2023b; Yoo et al., 2025) on the OM9 dataset (Ramakrishnan et al., 2014). We generated 1,024 samples

1404 across various timesteps and re-flow iterations, assessing validity, uniqueness, and novelty. The
1405 configuration for subsequent re-flow iterations follows the same protocol as our main experiment.
1406 As shown in Tab. 15, Tab. 16, and Tab. 17, all metrics are averaged over 10 independent runs with
1407 standard deviations provided. We observed that iterative re-flow consistently enhances few-step
1408 generation capabilities. Notably, PAIRFLOW demonstrates superior performance over the baseline; it
1409 outperforms UDLM not only at equivalent iteration levels but also surpasses UDLM with multiple
1410 re-flow iterations even when PAIRFLOW uses one or no additional iterations.

1411 **Table 15:** validity scores (\uparrow) on QM9 (Ramakrishnan et al., 2014) for UDLM and PAIRFLOW across
1412 varying rectification steps and NFEs (1 to 64). The best and second-best values per column are
1413 highlighted in **bold** and underlined, respectively.

Method	1	2	4	8	16	32	64
UDLM	17.5 ± 3.2	125.5 ± 11.8	497.6 ± 8.3	826.6 ± 10.3	953.5 ± 6.1	<u>991.9 ± 4.2</u>	1000.1 ± 3.5
+ Re-Flow	59.7 ± 8.8	232.4 ± 9.2	588.4 ± 15.8	849.6 ± 14.2	940.5 ± 8.5	967.5 ± 5.2	978.8 ± 5.2
+ Re-Flow	160.9 ± 11.8	368.0 ± 11.8	673.3 ± 11.2	878.5 ± 11.6	945.4 ± 8.1	967.9 ± 7.4	975.8 ± 8.0
+ Re-Flow	280.2 ± 6.7	470.7 ± 18.6	742.2 ± 18.7	897.9 ± 9.4	945.3 ± 5.2	965.0 ± 5.9	972.6 ± 7.1
PAIRFLOW	223.4 ± 12.7	416.0 ± 12.4	734.9 ± 7.2	921.5 ± 11.0	<u>977.1 ± 3.9</u>	990.9 ± 5.9	<u>1000.2 ± 4.5</u>
+ Re-Flow	361.0 ± 115.2	512.6 ± 44.2	775.7 ± 10.0	929.1 ± 11.6	976.2 ± 4.5	985.6 ± 6.7	993.1 ± 7.1
+ Re-Flow	443.4 ± 13.4	598.6 ± 18.4	823.0 ± 16.0	935.2 ± 7.7	969.0 ± 5.7	984.1 ± 6.0	989.0 ± 3.8
+ Re-Flow	<u>529.2 ± 10.6</u>	<u>688.0 ± 11.3</u>	<u>863.6 ± 9.5</u>	<u>945.9 ± 4.5</u>	975.7 ± 6.1	982.6 ± 7.1	990.2 ± 7.2

1424
1425 **Table 16:** Uniqueness scores (\uparrow) on QM9 (Ramakrishnan et al., 2014) for UDLM and PAIRFLOW
1426 across varying rectification steps and NFEs (1 to 64). The best and second-best values per column are
1427 highlighted in **bold** and underlined, respectively.

Method	1	2	4	8	16	32	64
UDLM	17.5 ± 3.2	125.4 ± 11.7	495.0 ± 8.2	819.5 ± 11.4	943.0 ± 5.7	<u>979.1 ± 5.0</u>	990.0 ± 4.7
+ Re-Flow	59.7 ± 8.8	231.6 ± 9.5	581.0 ± 15.1	834.7 ± 11.4	917.5 ± 9.6	944.8 ± 6.2	956.1 ± 5.2
+ Re-Flow	159.4 ± 10.9	363.3 ± 12.2	657.2 ± 9.3	846.3 ± 11.8	910.8 ± 7.3	930.5 ± 9.3	938.8 ± 11.7
+ Re-Flow	275.6 ± 7.2	456.3 ± 17.7	712.1 ± 15.7	856.8 ± 10.3	899.0 ± 4.9	907.0 ± 10.9	918.9 ± 5.7
PAIRFLOW	223.0 ± 12.3	414.7 ± 12.0	731.4 ± 6.9	<u>917.4 ± 11.8</u>	<u>971.6 ± 4.3</u>	<u>986.2 ± 5.3</u>	<u>994.8 ± 5.2</u>
+ Re-Flow	359.5 ± 113.1	507.7 ± 43.0	765.3 ± 8.8	913.5 ± 10.2	959.7 ± 5.8	968.8 ± 8.4	973.0 ± 9.6
+ Re-Flow	437.9 ± 13.8	586.7 ± 17.9	801.0 ± 15.4	906.5 ± 10.1	939.1 ± 6.2	955.9 ± 7.2	960.2 ± 6.5
+ Re-Flow	<u>516.9 ± 10.8</u>	<u>662.8 ± 13.4</u>	<u>828.7 ± 8.4</u>	903.2 ± 4.7	932.5 ± 9.5	942.8 ± 4.8	949.6 ± 8.4

1438
1439 **Table 17:** Novelty scores (\uparrow) on QM9 (Ramakrishnan et al., 2014) for UDLM and PAIRFLOW across
1440 varying rectification steps and NFEs (1 to 64). The best and second-best values per column are
1441 highlighted in **bold** and underlined, respectively.

Method	1	2	4	8	16	32	64
UDLM	13.8 ± 2.9	52.0 ± 8.2	120.0 ± 3.8	<u>152.4 ± 9.1</u>	<u>144.2 ± 12.4</u>	<u>147.2 ± 9.7</u>	<u>145.1 ± 9.0</u>
+ Re-Flow	31.4 ± 8.0	73.3 ± 7.1	110.4 ± 12.2	<u>126.8 ± 8.9</u>	116.3 ± 8.4	120.2 ± 9.0	117.6 ± 7.1
+ Re-Flow	61.1 ± 6.5	103.0 ± 8.6	<u>129.0 ± 9.7</u>	124.6 ± 12.5	<u>128.3 ± 5.6</u>	<u>128.4 ± 10.6</u>	122.8 ± 9.4
+ Re-Flow	91.2 ± 9.2	<u>116.4 ± 8.1</u>	<u>127.2 ± 8.3</u>	124.2 ± 9.3	121.9 ± 6.0	117.0 ± 4.9	121.8 ± 8.0
PAIRFLOW	68.8 ± 7.8	85.6 ± 10.0	109.2 ± 7.8	96.8 ± 9.9	106.5 ± 12.5	108.9 ± 9.4	110.0 ± 9.9
+ Re-Flow	84.2 ± 11.3	92.0 ± 6.9	101.1 ± 9.6	98.6 ± 8.4	98.8 ± 13.9	95.8 ± 8.3	98.9 ± 7.5
+ Re-Flow	100.8 ± 7.7	109.5 ± 8.7	101.0 ± 6.8	101.1 ± 7.8	94.5 ± 6.6	95.6 ± 9.6	99.0 ± 8.8
+ Re-Flow	<u>114.6 ± 9.4</u>	108.0 ± 5.9	106.5 ± 8.9	96.2 ± 9.8	95.6 ± 9.0	94.8 ± 7.1	95.7 ± 8.2

J SUBSET PAIRING RESULTS

1453
1454 In this section, we present comprehensive experimental results from applying our subset-partition
1455 pairing technique to the ZINC-250k molecular dataset (Irwin et al., 2012). Following the same
1456

1458 protocol as our main experiments, we generated 1,024 samples across varying timesteps and par-
 1459 tition counts to evaluate validity, uniqueness, and novelty. The only deviation from the standard
 1460 method is the pairing strategy; here, we calculate the closed-form backward velocity exclusively
 1461 within each subset to reduce the computational cost of Eqn. 10. All reported metrics, summarized
 1462 in Tab. 18, Tab. 19, and Tab. 20, are averaged over 10 independent runs with corresponding standard
 1463 deviations. Additionally, we report the pairing-time cost for each subset-partition configuration. The
 1464 results demonstrate that our subset-pairing algorithm effectively reduces the computational time for
 1465 pairing, while maintaining performance comparable to the full-set baseline.

1466
 1467 Table 18: Validity scores (\uparrow) on the Zinc-250k dataset (Irwin et al., 2012) evaluated across different
 1468 subset partitions and NFEs (1 to 64). The best and second-best values per column are highlighted in
 1469 **bold** and underlined, respectively. T_{PAIRFLOW} denotes the runtime (in minutes) of each configuration,
 1470 measured in wall-clock time using an RTX A6000 GPU.

Method	T_{PAIRFLOW}	1	2	4	8	16	32	64
UDLM	0	0.3 ± 0.5	65.2 ± 8.2	435.7 ± 14.4	775.1 ± 19.5	887.3 ± 12.7	921.5 ± 8.5	937.3 ± 3.9
Random	0	0.6 ± 0.9	68.3 ± 10.7	351.2 ± 15.8	569.4 ± 16.6	611.0 ± 16.3	602.4 ± 13.3	571.0 ± 13.2
PAIRFLOW (Full)	13m	9.9 ± 2.3	146.3 ± 10.4	533.9 ± 13.9	799.4 ± 9.2	873.2 ± 14.1	901.0 ± 14.2	907.8 ± 7.7
PAIRFLOW (2-Sub)	6m	10.6 ± 2.8	<u>145.7 ± 13.5</u>	<u>530.6 ± 20.3</u>	802.5 ± 6.9	882.6 ± 7.1	902.4 ± 13.4	911.5 ± 9.2
PAIRFLOW (4-Sub)	2.9m	<u>12.1 ± 3.3</u>	142.5 ± 5.2	<u>509.7 ± 11.0</u>	780.9 ± 14.4	858.9 ± 7.9	886.7 ± 8.7	899.0 ± 10.5
PAIRFLOW (8-Sub)	1.5m	12.3 ± 2.5	141.1 ± 6.7	510.9 ± 17.2	766.8 ± 12.6	857.5 ± 8.4	886.9 ± 9.6	896.5 ± 6.3

1471
 1472 Table 19: Uniqueness scores (\uparrow) on Zinc-250k (Irwin et al., 2012) evaluated across different subset
 1473 partitions and NFEs (1 to 64). The best and second-best values per column are highlighted in **bold**
 1474 and underlined, respectively. T_{PAIRFLOW} denotes the runtime (in minutes) of each configuration.

Method	T_{PAIRFLOW}	1	2	4	8	16	32	64
UDLM	0	0.3 ± 0.5	65.2 ± 8.2	435.7 ± 14.4	775.1 ± 19.5	887.3 ± 12.7	921.5 ± 8.5	937.2 ± 3.8
Random	0	0.6 ± 0.9	68.3 ± 10.7	351.2 ± 15.8	569.4 ± 16.6	611.0 ± 16.3	602.4 ± 13.3	571.0 ± 13.2
PAIRFLOW (Full)	13m	9.9 ± 2.3	146.3 ± 10.4	533.9 ± 13.9	799.4 ± 9.2	873.2 ± 14.1	901.0 ± 14.2	907.8 ± 7.7
PAIRFLOW (2-Sub)	6m	10.6 ± 2.8	<u>145.7 ± 13.5</u>	<u>530.6 ± 20.3</u>	802.5 ± 6.9	882.6 ± 7.1	902.4 ± 13.4	911.5 ± 9.2
PAIRFLOW (4-Sub)	2.9m	<u>12.1 ± 3.3</u>	142.5 ± 5.2	<u>509.7 ± 11.0</u>	780.9 ± 14.4	858.9 ± 7.9	886.7 ± 8.7	899.0 ± 10.5
PAIRFLOW (8-Sub)	1.5m	12.3 ± 2.5	141.1 ± 6.7	510.9 ± 17.2	766.8 ± 12.6	857.5 ± 8.4	886.9 ± 9.6	896.5 ± 6.3

1479
 1480 Table 20: Novelty scores (\uparrow) on Zinc-250k (Irwin et al., 2012) evaluated across different subset
 1481 partitions and NFEs (1 to 64). The best and second-best values per column are highlighted in **bold**
 1482 and underlined, respectively. T_{PAIRFLOW} denotes the runtime (in minutes) of each configuration.

Method	T_{PAIRFLOW}	1	2	4	8	16	32	64
UDLM	0	0.3 ± 0.5	65.2 ± 8.2	435.7 ± 14.4	775.1 ± 19.5	887.3 ± 12.7	921.3 ± 8.8	936.9 ± 4.1
Random	0	0.6 ± 0.9	68.3 ± 10.7	351.2 ± 15.8	569.4 ± 16.6	611.0 ± 16.3	602.4 ± 13.3	571.0 ± 13.2
PAIRFLOW (Full)	13m	9.9 ± 2.3	146.3 ± 10.4	533.9 ± 13.9	799.4 ± 9.2	873.2 ± 14.1	901.0 ± 14.2	907.8 ± 7.7
PAIRFLOW (2-Sub)	6m	10.6 ± 2.8	<u>145.7 ± 13.5</u>	<u>530.6 ± 20.3</u>	802.5 ± 6.9	882.6 ± 7.1	902.4 ± 13.4	911.5 ± 9.2
PAIRFLOW (4-Sub)	2.9m	<u>12.1 ± 3.3</u>	142.5 ± 5.2	<u>509.7 ± 11.0</u>	780.9 ± 14.4	858.8 ± 8.0	886.7 ± 8.7	899.0 ± 10.5
PAIRFLOW (8-Sub)	1.5m	12.3 ± 2.5	141.1 ± 6.7	510.9 ± 17.2	766.8 ± 12.6	857.5 ± 8.4	886.9 ± 9.6	896.5 ± 6.3

K APPLICATION FOR MORE COMPLEX SYSTEMS

1508 In this section, we evaluate our method on a higher-dimensional dataset. Specifically, we use the
 1509 FFHQ (Karras et al., 2019) dataset, downsampled to 64×64 . Following the same protocol as in
 1510 our main experiments, we generate 5,000 samples across varying timesteps and report the FID

1512 computed against the training set. The results of this experiment are provided in Tab. 21. All training
 1513 hyperparameters are kept identical to those used in the CIFAR-10 experiments described in Section 5.
 1514

1515 **Table 21:** Comparison of FID scores (\downarrow) on FFHQ (Karras et al., 2019) downsampled to 64×64
 1516 resolution across extended NFE steps (1 to 1024). Best values per column are highlighted in bold.
 1517

Method	1	2	4	8	16	32	64	128	256	512	1024
UDLM	403.04	399.26	363.97	273.31	153.71	97.87	74.85	63.93	59.28	55.99	55.30
PAIRFLOW	394.14	368.36	329.13	243.88	140.05	90.85	69.67	59.86	56.52	54.19	53.18

1521
 1522 We adopt the LM1B (Chelba et al., 2013) dataset to evaluate our method under a substantially larger
 1523 vocabulary size and training corpus. The text corpus is segmented into sequences of varying lengths
 1524 ($N = 16, 32, 64, 128$), while keeping the total number of training samples fixed ($|X_1| \approx 3.5M$). To
 1525 assess generation quality, we compute generative perplexity using GPT-2 Large and entropy on 1,024
 1526 generated samples for each NFE setting. The results are summarized in Tab. 22 and Tab. 23. For
 1527 training, we follow the network hyperparameter configuration of (Schiff et al., 2025), modifying only
 1528 the number of training iterations for each sequence dimensionality.
 1529

1530 **Table 22:** Generative Perplexity (\downarrow) on LM1B (Chelba et al., 2013) measured with GPT2-large across
 1531 varying lengths (N) and their corresponding training iterations (Iter.) over NFE steps 4 to 1024. Best
 1532 values are highlighted in bold.
 1533

N	Iter.	Method	4	8	16	32	64	128	256	512	1024
16	200k	UDLM	299.18	225.92	207.17	195.82	200.77	197.04	199.12	195.37	198.22
		PAIRFLOW	242.22	208.04	200.99	190.36	191.74	199.45	188.84	196.91	198.12
32	200k	UDLM	263.93	192.78	167.85	167.49	155.68	150.52	152.40	151.74	154.02
		PAIRFLOW	218.48	172.27	156.35	150.53	143.83	145.77	142.54	141.04	147.57
64	400k	UDLM	214.07	150.59	130.49	120.19	117.90	116.23	112.24	113.77	115.11
		PAIRFLOW	174.78	138.94	123.06	115.71	114.73	112.92	111.29	107.06	110.83
128	600k	UDLM	169.61	123.48	105.13	98.94	97.89	94.92	93.75	94.12	93.59
		PAIRFLOW	167.90	121.09	102.16	96.61	93.93	91.51	90.21	89.09	89.07

1544 **Table 23:** Comparison of Entropy (\uparrow) on LM1B (Chelba et al., 2013) across varying lengths (N) and
 1545 training iterations (Iter.) over NFE steps 4 to 1024. Best values are highlighted in bold.
 1546

N	Iter.	Method	4	8	16	32	64	128	256	512	1024
16	200k	UDLM	2.46	2.49	2.50	2.50	2.50	2.51	2.50	2.51	2.50
		PAIRFLOW	2.48	2.49	2.51	2.52	2.52	2.53	2.52	2.53	2.52
32	200k	UDLM	3.05	3.09	3.12	3.13	3.13	3.13	3.13	3.13	3.13
		PAIRFLOW	3.06	3.12	3.13	3.14	3.15	3.15	3.15	3.15	3.16
64	400k	UDLM	3.57	3.63	3.67	3.68	3.69	3.70	3.70	3.69	3.70
		PAIRFLOW	3.57	3.65	3.69	3.70	3.71	3.71	3.71	3.72	3.71
128	600k	UDLM	3.98	4.09	4.14	4.16	4.17	4.17	4.17	4.18	4.18
		PAIRFLOW	4.00	4.11	4.16	4.18	4.19	4.20	4.20	4.19	4.20

L ANALYSIS FOR THE OVERRFITTING IN IMAGE DOMAINS

1563 In this section, we evaluate our method using FID computed on the test sets of two image domains:
 1564 CIFAR-10 (Krizhevsky et al., 2009) and MNIST-Binary (LeCun et al., 2002). For CIFAR-10, we
 1565 additionally report FID scores measured with DINOv2 (Oquab et al., 2024). The overall results are
 summarized in Tab. 24 and Tab. 25. Across all evaluation metrics, the performance trend is consistent

1566 with our main findings—PAIRFLOW delivers improved generation quality over the baseline, with
 1567 especially strong gains in the few-step generation regime.
 1568
 1569
 1570

1571 **Table 24:** FID and FID-Dino scores (\downarrow) on test dataset for CIFAR-10 (Krizhevsky et al., 2009)
 1572 comparison across extended NFE steps (1 to 1024). Best values are bolded.
 1573

NFE	1	2	4	8	16	32	64	128	256	512	1024
Method	FID (\downarrow)										
UDLM	306.45	296.77	266.64	178.11	114.04	80.40	62.70	53.83	50.96	47.61	47.16
PAIRFLOW	235.65	247.18	209.94	137.16	94.24	67.53	51.48	43.79	42.44	40.85	39.92
Method	FID-DINOv2 (\downarrow)										
UDLM	2448.46	2410.65	1975.58	1344.44	959.39	755.80	646.47	598.53	598.54	553.17	560.75
PAIRFLOW	2059.26	1988.21	1626.08	1127.30	828.09	623.89	530.70	486.59	484.27	470.33	470.05

1581
 1582
 1583
 1584
 1585
 1586 **Table 25:** FID scores (\downarrow) on the MNIST-Binary (LeCun et al., 2002) test set across various NFE steps
 1587 (1 to 64). Best values are bolded.
 1588

Method	1	2	4	8	16	32	64
UDLM	129.05	42.17	11.42	6.18	5.13	5.37	5.50
PAIRFLOW	42.87	17.37	9.62	6.36	5.80	5.51	5.24
UDLM+DCD	57.85	19.23	9.82	8.41	7.87	7.12	7.38
PAIRFLOW+DCD	19.56	13.06	10.90	8.55	7.40	7.22	7.85
UDLM+ReDi	19.08	10.79	8.77	7.01	6.89	6.57	6.61
PAIRFLOW+ReDi	13.73	9.59	8.98	7.24	7.22	6.98	7.12

1597
 1598
 1599
 1600 We further assess potential memorization by measuring nearest-neighbor distances with respect to the
 1601 training set. For MNIST-Binary (LeCun et al., 2002), we compute pixel-wise ℓ_2 distances, whereas
 1602 for CIFAR-10 (Krizhevsky et al., 2009), we evaluate both ℓ_2 distance and cosine similarity between
 1603 features extracted using DINOv2 (Oquab et al., 2024). As summarized in Tab. 26 and Tab. 27, across
 1604 all evaluation settings, the nearest-neighbor distances of PAIRFLOW are comparable to or slightly
 1605 larger than those of the baseline. These results support the conclusion that our method does not suffer
 1606 from severe overfitting or excessive memorization of the training data.
 1607
 1608

1609 **Table 26:** Comparison of ℓ_2 and Dino (Oquab et al., 2024) Cosine nearest neighbor distance on the
 1610 CIFAR-10 (Krizhevsky et al., 2009) training set across extended NFE steps (1 to 1024). Best values
 1611 are bolded.
 1612

NFE	1	2	4	8	16	32	64	128	256	512	1024
Metric	ℓ_2 (\uparrow)										
UDLM	7.97	9.13	10.06	10.06	9.40	8.94	8.63	8.41	8.29	8.29	8.22
PAIRFLOW	8.03	8.76	9.42	9.56	9.18	8.75	8.63	8.55	8.51	8.52	8.52
Metric	Cosine(DINOv2) (\downarrow)										
UDLM	0.242	0.227	0.231	0.241	0.237	0.238	0.235	0.237	0.237	0.235	0.237
PAIRFLOW	0.245	0.228	0.235	0.239	0.236	0.232	0.232	0.230	0.232	0.235	0.233

1620 **Table 27: Comparison of ℓ_2 nearest neighbor distance on the MNIST-Binary (LeCun et al., 2002)**
1621 training set across extended NFE steps (1 to 64). Best values are bolded.

1622

Method	1	2	4	8	16	32	64
UDLM	8.18	7.06	6.55	6.42	6.27	6.29	6.25
PAIRFLOW	7.36	6.95	6.63	6.42	6.26	6.34	6.30
UDLM+DCD	7.24	6.78	6.57	6.50	6.31	6.40	6.39
PAIRFLOW+DCD	7.64	7.34	7.11	6.91	6.74	6.76	6.79
UDLM+ReDi	7.14	6.81	6.49	6.24	6.10	6.14	6.08
PAIRFLOW+ReDi	6.84	6.57	6.35	6.09	5.95	5.97	5.96

1631

1632

M LIMITATIONS AND FUTURE WORK

1633

1635 We hope this work initiates broader discussion on reducing training compute while still enabling
1636 fast generation in generative models. Such efficiency can have a significant impact, from reducing
1637 energy consumption in training large-scale generative models to contributing to the democratization
1638 of foundation model development.

1639 A natural follow-up question to our work is whether the same idea can be applied to continuous
1640 Flow Matching (FM). We have evaluated this extension on continuous FM models, with results
1641 provided in App. E. Our experiments with synthetic data show that the method is effective for
1642 relatively low-dimensional data, while its advantage a bit diminishes for higher-dimensional data.
1643 We will further investigate the effect of our method on continuous data, where we hypothesize that a
1644 substantially larger number of source–target pairs will be required. Nonetheless, we emphasize that
1645 even in this initial exploration of accelerating flow models through well-aligned pairing, PAIRFLOW
1646 is particularly well-suited for low-dimensional discrete data, which includes many forms of scientific
1647 data such as molecular and protein structures.

1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673

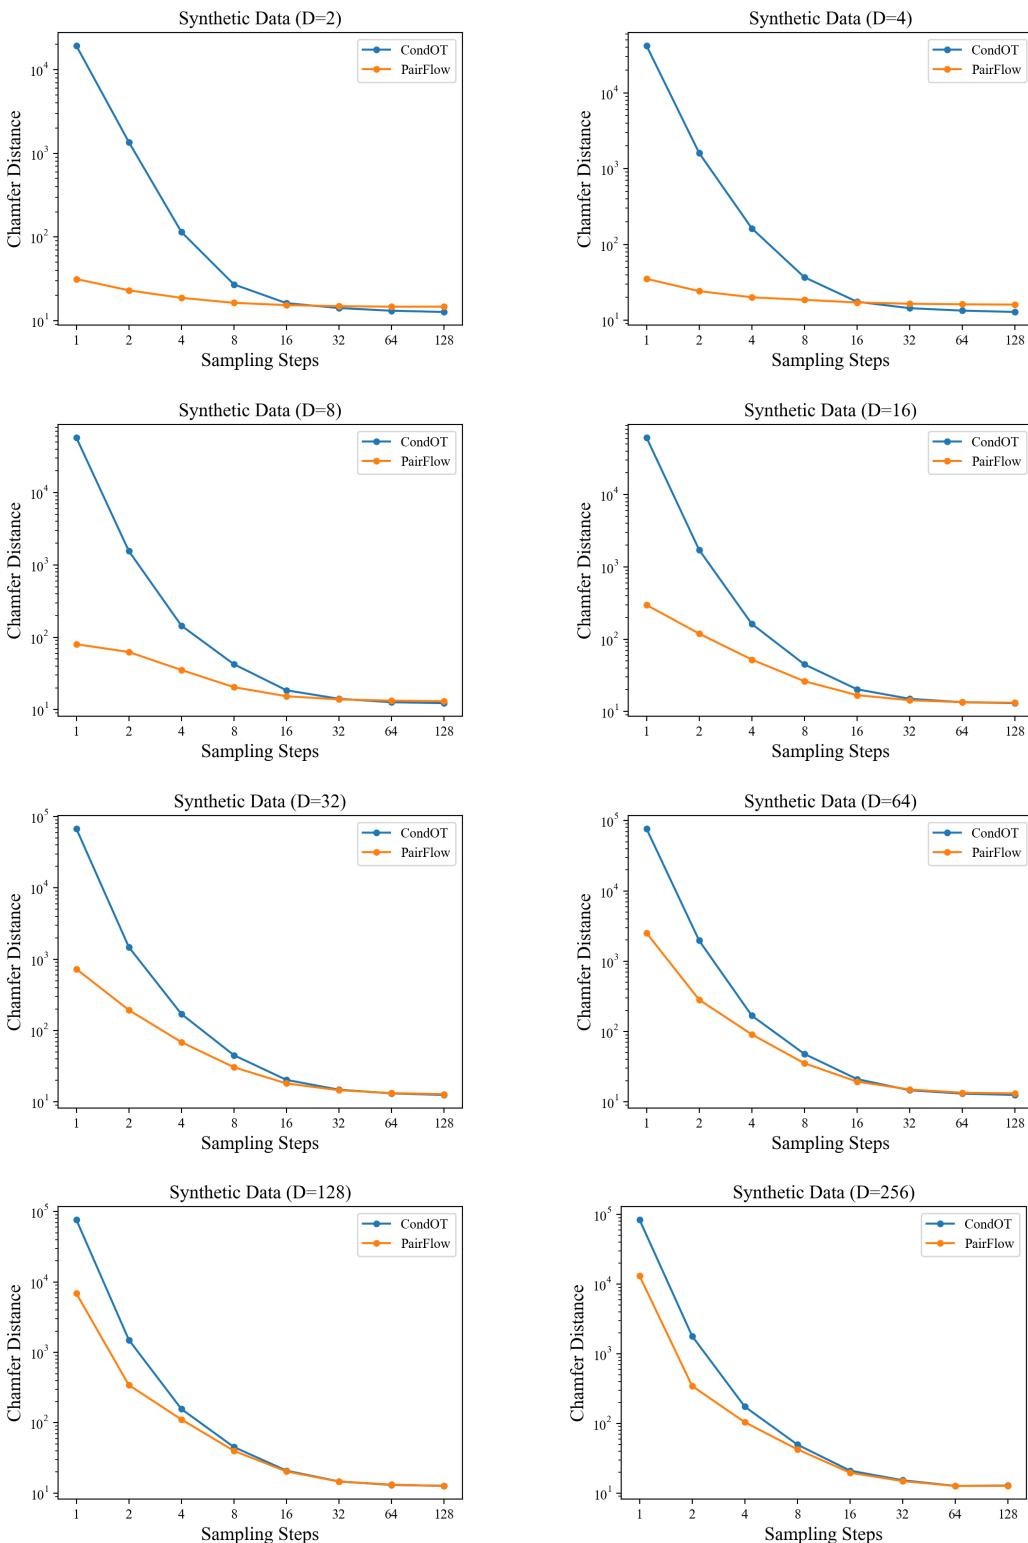


Figure 7: Step-wise performance analysis on synthetic N-fold two-moons. Each plot shows Chamfer distance (log scale) vs. sampling steps for different dimensions d . Closed-form pairing (PAIRFLOW) consistently outperforms standard CondOT—especially at few sampling steps—while the margin shrinks as d increases, indicating diminishing gains in high dimensions.

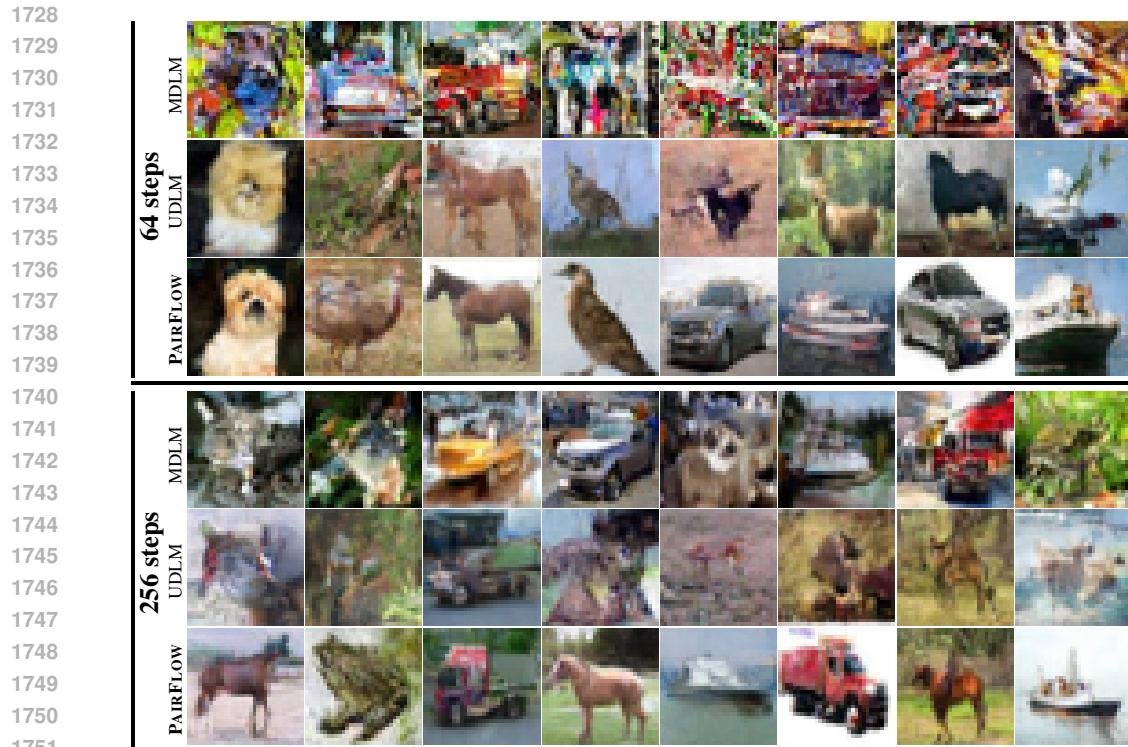
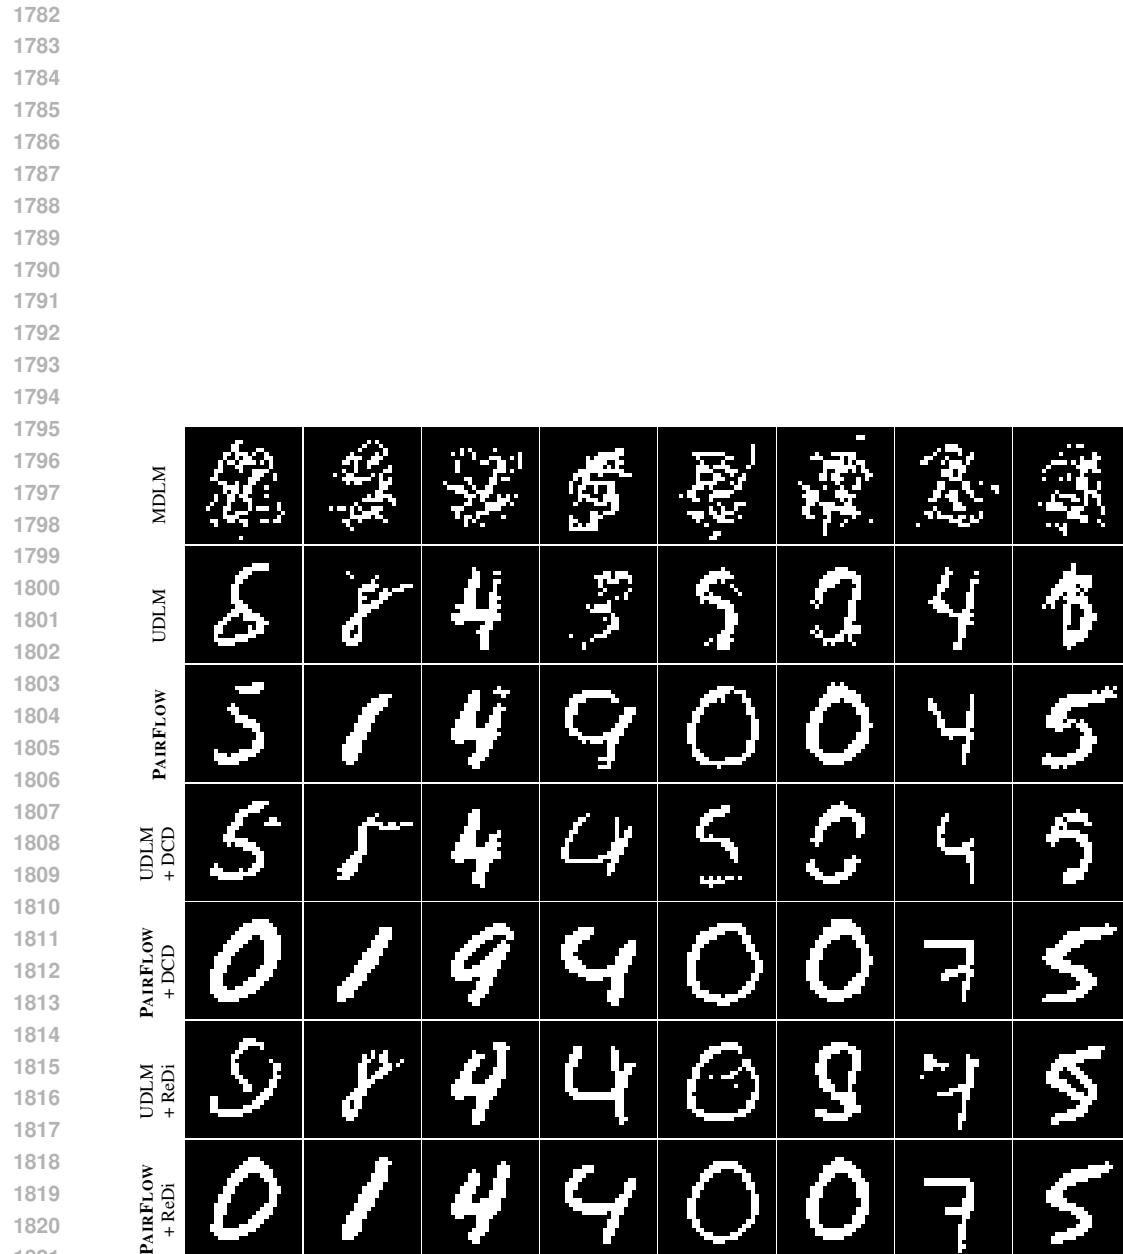


Figure 8: Additional qualitative results of 64-step and 256-step generation on CIFAR-10 (32 × 32).

1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2



1822 Figure 10: Additional qualitative results of 2-step generation on MNIST-Binary (28 × 28).
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835