

000 001 002 003 004 005 JOINT DISTILLATION FOR FAST LIKELIHOOD EVALU- 006 ATION AND SAMPLING IN FLOW-BASED MODELS 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030

ABSTRACT

031 Log-likelihood evaluation enables important capabilities in generative models, in-
032 cluding model comparison, certain fine-tuning objectives, and many downstream
033 applications. Yet paradoxically, some of today’s best generative models – diffu-
034 sion and flow-based models – still require hundreds to thousands of neural func-
035 tion evaluations (NFEs) to compute a single likelihood. While recent distillation
036 methods have successfully accelerated sampling to just a few steps, they achieve
037 this at the cost of likelihood tractability: existing approaches either abandon like-
038 lihood computation entirely or still require expensive integration over full trajec-
039 tories. We present fast flow joint distillation (F2D2), a framework that simulta-
040 neously reduces the number of NFEs required for both sampling and likelihood
041 evaluation by two orders of magnitude. Our key insight is that in continuous nor-
042 malizing flows, the coupled ODEs for sampling and likelihood are computed from
043 a shared underlying velocity field, allowing us to jointly distill both the sampling
044 trajectory and cumulative divergence using a single model. F2D2 is modular,
045 compatible with existing flow-based few-step sampling models, and requires only
046 an additional divergence prediction head. Experiments demonstrate F2D2’s capa-
047 bility of achieving accurate log-likelihood with few-step evaluations while main-
048 taining high sample quality, solving a long-standing computational bottleneck in
049 flow-based generative models. As an application of our approach, we propose a
050 lightweight self-guidance method that enables a 2-step MeanFlow model to out-
051 perform a 1024 step teacher model with only a single additional backward NFE.
052
053

1 INTRODUCTION

034 Log-likelihood evaluation and likelihood-based inference have long been fundamental to statistical
035 modeling and machine learning, serving as the backbone for parameter estimation (Fisher, 1922),
036 model selection (Akaike, 1974), and hypothesis testing (Neyman & Pearson, 1933). In the era of
037 generative AI, the ability to efficiently evaluate log-likelihood (log-density) has become even more
038 critical, as it directly enables key post-training techniques including reinforcement learning and pre-
039 ference optimization, where likelihoods are important for methods like PPO, DPO and GRPO (Schul-
040 man et al., 2017; Ouyang et al., 2022; Rafailov et al., 2023; Shao et al., 2024). Beyond these ap-
041 plications, optimizing log-likelihood also encourages generative models to capture all modes of the
042 data distribution, avoiding mode collapse that plagues adversarial approaches (Razavi et al., 2019).

043 While likelihood evaluation is useful for modern generative modeling, the most successful gener-
044 ative models for images and video (Rombach et al., 2022; Black Forest Labs, 2025; OpenAI,
045 2024; Polyak et al., 2024; Google DeepMind, 2025), namely diffusion and flow matching mod-
046 els, suffer from a critical weakness: computing likelihood requires prohibitively expensive iterative
047 neural function evaluations (NFEs). In particular, discrete-time diffusion models like DDPM (Ho
048 et al., 2020; Nichol & Dhariwal, 2021) require summing up variational bounds across all timesteps,
049 which needs hundreds to thousands of forward passes to compute a single likelihood. Similarly,
050 continuous-time formulations like score SDE (Song et al., 2020) and flow matching (Lipman et al.,
051 2022; Albergo & Vanden-Eijnden, 2022; Liu et al., 2022; Albergo et al., 2023) also must inte-
052 grate the divergence along the learned (probability) flow ODE, which typically requires numerical
053 integration with 100-1000 NFEs for accurate likelihood evaluation. While advanced solvers can
significantly reduce NFEs (Karras et al., 2022), they fundamentally cannot escape the integration
requirement and produce vastly inaccurate results when restricted to very few steps (≤ 10 NFEs).

054 This computational burden makes many likelihood-based finetuning objectives, model comparison,
 055 and downstream applications prohibitively expensive for modern diffusion/flow matching models.
 056

057 Interestingly, diffusion and flow matching models faced the same NFE bottleneck for *sampling* when
 058 they were first introduced, where they initially required 1000+ steps to generate a single image.
 059 Research addressing this issue has been remarkably successful, with methods that learn to skip
 060 multiple steps, either through distillation or self-consistency objectives, emerging as particularly
 061 powerful solutions (Salimans & Ho, 2022; Song et al., 2023; Kim et al., 2023; Frans et al., 2024;
 062 Geng et al., 2025; Boffi et al., 2025b;a). However, despite achieving few-step sampling, most of
 063 these methods completely lose the ability to compute likelihoods, and the few methods that preserve
 064 likelihood computation (e.g. Kim et al. (2023)) still require integrating over the entire trajectory with
 065 hundreds of NFEs, making fast likelihood evaluation impossible. Thus, while practitioners have
 066 solutions to fast sampling for diffusion and flow matching models, fast log-likelihood evaluation
 067 remains an unsolved problem. Here, we show that it is possible to achieve *both* at the same time.

068 **Contributions.** Our key insight is that, in flow matching and continuous normalizing flows (CNFs)
 069 in general (Chen et al., 2018), computing exact likelihoods requires solving coupled ODEs: the
 070 sampling trajectory $\frac{d}{dt}x_t = v_\theta(x_t, t)$ and the log-density evolution $\frac{d}{dt}\log p_t(x_t) = -\text{div}(v_\theta(x_t, t))$,
 071 both depending on the same learned velocity v_θ . Since the divergence term can be viewed as another
 072 output derived from the same velocity model, we can learn to distill both the flow trajectory and its
 073 corresponding divergence computation simultaneously within a single model. By jointly optimizing
 074 for both accurate few-step sampling and log-likelihood evaluation, we can potentially achieve a
 075 model that succeeds at both tasks.

076 Based on these insights, we propose fast flow joint distillation (F2D2), a simple and modular frame-
 077 work for jointly learning fast sampling and fast log-likelihood evaluation in flow-based models. Our
 078 key idea is to leverage the flow map framework (Boffi et al., 2025a) to train a single model to pre-
 079 dict both the sampling trajectory and cumulative divergence in parallel using a joint self-distillation
 080 objective, ensuring both outputs learn to skip the numerous steps in training. This makes F2D2
 081 plug-and-play with any CNF-based few-step sampling method like shortcut models and MeanFlow,
 082 and requires only a new divergence prediction head alongside the existing velocity prediction. To
 083 our knowledge, F2D2 is the first method to enable accurate few-step log-likelihood evaluation in
 084 diffusion/CNF-based generative models, solving a long-standing limitation of these frameworks.

085 We demonstrate that our method produces both calibrated likelihoods and high quality samples with
 086 few-step NFEs on image datasets CIFAR-10 (Krizhevsky et al.) and ImageNet 64 × 64 (Deng et al.,
 087 2009). We show that our F2D2 are compatible with and can be directly apply to pre-trained shortcut
 088 models, MeanFlow and a new distillation method we propose in this paper.

089 As an application of our method, we introduce maximum likelihood self-guidance, a lightweight
 090 test-time intervention which uses rapid likelihood evaluation to optimize over generated samples,
 091 requiring only an additional forward and backward pass through the model. Remarkably, we show
 092 that F2D2 with maximum likelihood self-guidance instantiated with 2-step MeanFlow achieves
 093 lower FID than a 1024-step flow matching model of the same size on CIFAR-10. This proof of
 094 concept demonstrates the expanded algorithmic sandbox enabled by rapid likelihood evaluation.

095 2 BACKGROUND

096 Let p_{data} denote the data distribution with samples $x \in \mathbb{R}^d$. We consider a time variable $t \in [0, 1]$
 097 where $t = 0$ corresponds to a simple noise distribution $p_0 = \mathcal{N}(0, I)$ and $t = 1$ corresponds to the
 098 data distribution $p_1 = p_{\text{data}}$. We denote the marginal distribution at time t as $p_t(x)$.

099 **Flow Matching.** Flow matching (Lipman et al., 2022; Albergo & Vanden-Eijnden, 2022; Liu et al.,
 100 2022; Albergo et al., 2023) is a scalable training method for generative modeling that learns a time-
 101 dependent velocity field $v_\theta : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ to transport samples from a simple noise distribution
 102 p_0 to the data distribution p_1 . Along a straight-line path $x_t = (1-t)x_0 + tx_1$ that linearly interpolates
 103 between a noise sample $x_0 \sim p_0$ and a data sample $x_1 \sim p_1$, it models the evolution dynamic with
 104 the ordinary differential equation ODE:
 105

$$106 \frac{d}{dt}\hat{x}_t = v_\theta(\hat{x}_t, t), \quad x_0 \sim p_0 \quad (2.1)$$

108 where \hat{x}_t arises from integrating the learned flow model. Since the velocity along this path is simply
 109 $x_1 - x_0$, flow matching minimizes the regression objective:
 110

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{t \sim [0,1], x_0 \sim p_0, x_1 \sim p_1} [\|v_\theta(x_t, t) - (x_1 - x_0)\|^2] \quad (2.2)$$

112 This objective encompasses diffusion models as a special case with different interpolation
 113 schemes (Gao et al., 2024). For sampling, it solves the ODE from $t = 0$ to $t = 1$ with numeri-
 114 cal solvers like Euler or dopri5, and typically requires 100-1000 NFEs for high quality samples.

115 **Continuous Normalizing Flows and Likelihood Computation.** A flow-matching model is a spe-
 116 cial case of a continuous normalizing flow (CNF) (Chen et al., 2018), which transports data from
 117 an initial distribution $x_0 \sim p_0$ to an estimated distribution by integrating an ODE. In the case of
 118 flow-based models, this is precisely Eq. (2.1). An advantage of the CNF formalism is the ability to
 119 explicitly compute likelihoods via the coupled system of ODEs:

$$\frac{d}{dt} \begin{bmatrix} \hat{x}_t \\ \log p_{t;\theta}(\hat{x}_t) \end{bmatrix} = \begin{bmatrix} v_\theta(\hat{x}_t, t) \\ -\text{div}(v_\theta(\hat{x}_t, t)) \end{bmatrix} \quad (2.3)$$

123 Above, $\text{div}(v_\theta(\hat{x}_t, t)) = \text{Tr}(\nabla_{\hat{x}_t} v_\theta(\hat{x}_t, t))$ denotes the divergence of the velocity field v_θ , and $p_{t;\theta}$
 124 represents the likelihood of \hat{x}_t under Eq. (2.1), which we note depends on model parameter via v_θ .
 125 Integrating backwards from $t = 1$ (data) to $t = 0$ (noise) with initial conditions $[x_1, 0]^\top$, we obtain:

$$\log p_1(x_1) = \log p_0(\hat{x}_0) + \int_1^0 \text{div}(v_\theta(\hat{x}_t, t)) dt = \log p_0(\hat{x}_0) - \int_0^1 \text{div}(v_\theta(\hat{x}_t, t)) dt, \quad (2.4)$$

126 where \hat{x}_0 and the intermediate \hat{x}_t 's are obtained by integrating the flow backward from x_1 .
 127

128 Likelihood evaluation is typically expensive, requiring both careful, finely-discretized integration of
 129 an ODE across time steps, and a computation of the divergence term whose exact computation (or
 130 the variance of its randomized estimator (Grathwohl et al., 2018)) scales at least linearly in ambient
 131 dimension d . Thus, likelihood evaluation is far more computationally burdensome than sampling.

132 **Few-Step Flow-based Models.** To address the computational expense of multiple ODE integrations
 133 in *sampling*, recent few-step flow-based models (Kim et al., 2023; Frans et al., 2024; Geng et al.,
 134 2025; Boffi et al., 2025a;b) learn to directly predict the outcome of integrating the ODE in Eq. (2.1)
 135 using only a small number of function evaluations (NFEs). These methods can be viewed as sharing
 136 a common strategy of learning to predict the *flow map* of the underlying ODE.

137 **Definition 2.1** (Flow Map). Given an ODE $dx_t = v(x_t, t)dt$, the flow map $\Phi : \mathbb{R}^d \times [0, 1]^2 \rightarrow \mathbb{R}^d$
 138 is the solution operator that maps any state at time t to its corresponding state at time s :

$$\Phi(x_t, t, s) = x_t + \int_t^s v(x_\tau, \tau) d\tau = x_s \quad (2.5)$$

139 After learning the flow map with network parameter θ , one can directly perform few-step sampling
 140 by discretizing the time interval $[0, 1]$ into K steps with timesteps $0 = t_0 < t_1 < \dots < t_K = 1$, and
 141 iteratively applying the learned flow map: $\hat{x}_{t_{i+1}} = \Phi_\theta(\hat{x}_{t_i}, t_i, t_{i+1})$ for $i = 0, \dots, K-1$, starting
 142 from $x_0 \sim p_0$. This reduces sampling from hundreds of ODE solver steps to just K NFEs (typically
 143 $K < 10$), as each application of Φ directly predicts the integrated result over the interval $[t_i, t_{i+1}]$
 144 without explicit numerical integration.

3 METHOD

145 We propose to jointly accelerate both sampling and *likelihood evaluation* by learning a flow-map
 146 on the joint ODE system described in in Eq. (2.3). Again, $p_0 = \mathcal{N}(0, I)$ is the source distribution,
 147 $p_1 = p_{\text{data}}$. p_t represents the marginal distribution of the interpolant $x_t = tx_1 + (1-t)x_0$, v denotes
 148 the ground truth velocity and $p_{t;\theta}$ is the distribution of \hat{x}_t under the learned flow model Eq. (2.1).
 149 Our aim is to design model which supports two key capabilities:

1. **Fast sampling:** Draw a \hat{x}_1 from a trained flow model, using a few number of NFEs,
 $K_{\text{samp}} < 10$.
2. **Fast likelihood evaluation:** Evaluate the log likelihood of either model samples \hat{x}_1 or data
 samples x_1 using a few number of NFEs, $K_{\text{ll}} < 10$.

162 3.1 FAST FLOW JOINT DISTILLATION (F2D2): PARAMETRIZING A JOINT FLOW MAP
163

164 Our key insight is that we can apply few-step flow-based models to the ODE in Eq. (2.3) which
165 jointly parametrizes sampling and likelihood evaluation. Following Boffi et al. (2025a); Frans et al.
166 (2024); Geng et al. (2025), we adopt a linear parametrization of the flow map:

$$167 \quad \Phi_\theta(\hat{x}_t, t, s) = \hat{x}_t + (s - t)u_\theta(\hat{x}_t, t, s) \quad (3.1)$$

169 where $u_\theta : \mathbb{R}^d \times [0, 1]^2 \rightarrow \mathbb{R}^d$ predicts the average velocity that directly transports states from time
170 t to time s and ideally $u_\theta(x_t, t, s) \approx \frac{1}{s-t} \int_t^s v(x_\tau, \tau) d\tau$. With this parametrization, we recover an
171 estimate of the instantaneous velocity as $u_\theta(x_t, t, t)$ in the $s \rightarrow t$ limit, and obtain simple conditions
172 for valid flow maps.

173 **Proposition 3.1** (Flow Map Conditions (Boffi et al., 2025a)). *An operator $\Phi(x, t, s) = x + (s - t)u(x, t, s)$ is a valid flow map if and only if for all $(t, s) \in [0, 1]^2$ and for all $x \in \mathbb{R}^d$, $u(x, t, t) = v(x, t)$ and any of the following conditions holds:*

- 177 (a) Φ solves the **Lagrangian equation** $\partial_s \Phi(x, t, s) = u(\Phi(x, t, s), s, s)$.
- 178 (b) Φ solves the **Eulerian equation** $\partial_t \Phi(x, t, s) + \nabla_x \Phi(x, t, s)u(x, t, t) = 0$.
- 179 (c) Φ satisfies the **semigroup property** $\Phi(\Phi(x, t, r), r, s) = \Phi(x, t, s)$ for $t < r < s$.

181 Let $z_t = \log p_t(x_t) \in \mathbb{R}$ and $\hat{z}_t = \log p_{t;\theta}(\hat{x}_t) \in \mathbb{R}$ denote the log likelihood, we can then separately
182 parametrize the flow maps for the two subsystems in Eq. (2.3) as

$$184 \quad \Phi_{X;\theta_X}(\hat{x}_t, t, s) = \hat{x}_t + (s - t)u_{\theta_X}(\hat{x}_t, t, s), \quad (3.2)$$

$$185 \quad \Phi_{Z;\theta_Z}(\hat{x}_t, \hat{z}_t, t, s) = \hat{z}_t + (s - t)D_{\theta_Z}(\hat{x}_t, t, s)$$

187 Here $u_{\theta_X}(\hat{x}_t, t, s)$ still estimates the average velocity, and $D_{\theta_Z}(x_t, t, s)$ approximates the average
188 divergence $D_{\theta_Z}(x_t, t, s) \approx -\frac{1}{s-t} \int_t^s \text{div}(v(x_\tau, \tau)) d\tau$ along the true trajectory between t and s .

189 Notice that average divergence depends *only* on x_t , not z_t . The fact that x_t is sufficient in our
190 parametrization follows from the joint ODE Eq. (2.3), where the evolution of the likelihood z_t is
191 determined by the divergence of the first flow evaluated at x_t .

192 Therefore, denoting the joint state at time t as $y_t = (x_t, z_t)^\top$, we can then parametrize the joint flow
193 map using shared parameter θ as

$$195 \quad \Phi_{Y;\theta}(\hat{y}_t, t, s) = \begin{bmatrix} \Phi_X(\hat{x}_t, t, s) \\ \Phi_Z(\hat{x}_t, \hat{z}_t, t, s) \end{bmatrix} = \hat{y}_t + (s - t)f_\theta(\hat{x}_t, t, s), \quad (3.3)$$

$$196 \quad f_\theta(\hat{x}_t, t, s) = \begin{bmatrix} u_\theta(\hat{x}_t, t, s) \\ D_\theta(\hat{x}_t, t, s) \end{bmatrix}$$

200 Above, the networks for u_θ and D_θ share the same backbone with separate prediction heads for their
201 respective components. The exact architecture is described in Appendix C.

202 **Theoretical Justification.** To justify the parameterization Eq. (3.3), we recall a property denoted
203 the tangent condition by Boffi et al. (2025a), leveraged by Kim et al. (2023); Geng et al. (2025) to
204 recover the instantaneous velocity and divergence in the $s \rightarrow t$ limit:

205 **Lemma 3.2** (Tangent Condition). *The flow map $\Phi_Y(y, t, s)$ for the joint system Eq. (2.3) satisfies
lim _{$t \rightarrow s$} $\partial_s \Phi_Y(y, t, s) = f(x, s, s) = (v(x, t), -\text{div}(v(x_t, t)))^\top$.*

208 We can then characterize valid joint flow maps under our parametrization as the following:

210 **Proposition 3.3** (Characterization of the Joint Flow Map). *Let $\Phi_Y(y, t, s) = y + (s - t)f(x, t, s)$
211 satisfy $f(x, s, s) = (v(x, t), -\text{div}(v(x_t, t)))^\top$ denotes the dynamics for the joint sampling and
212 likelihood system Eq. (2.3). Then, $\Phi_Y(y, t, s)$ is the flow map for the joint system if and only if
213 $\forall (y, t, s) \in \mathbb{R}^{d+1} \times [0, 1]^2$, any of the following conditions are satisfied:*

- 214 (a) Φ_Y solves the **Lagrangian equation** $\partial_s \Phi_Y(y, t, s) = f(\Phi_Y(y, t, s), s, s)$.
- 215 (b) Φ solves the **Eulerian equation** $\partial_t \Phi_Y(y, t, s) + \nabla_y \Phi_Y(y, t, s)f(y, t, t) = 0$.

216 **Algorithm 1** Shortcut-F2D2 Training

217 1: **for** each training step **do**

218 2: $x_1 \sim p_{\text{data}}, x_0 \sim p_0, (t, s) \sim \mathcal{U}([0, 1]^2)$ with $t < s$

219 3: $r \leftarrow (t + s)/2$

220 4: $x_t \leftarrow (1 - t)x_0 + tx_1$

221 5: $x_r \leftarrow x_t + (r - t)u_\theta(x_t, t, r)$

222 6: $\mathcal{L}_{\text{VM-SC}}(\theta) \leftarrow \|u_\theta(x_t, t, t) - (x_1 - x_0)\|^2$

223 7: $\mathcal{L}_{\text{u-SC}}(\theta) \leftarrow \|u_\theta(x_t, t, s) - \frac{1}{2}\text{sg}(u_\theta(x_t, t, r) + u_\theta(x_r, r, s))\|^2$

224 8: $D_t \leftarrow \text{sg}(\text{div}(u_\theta(x_t, t, t)))$

225 9: $\mathcal{L}_{\text{div-SC}}(\theta) \leftarrow \|D_\theta(x_t, t, t) + D_t\|^2$

226 10: $\mathcal{L}_{\text{D-SC}}(\theta) \leftarrow \|D_\theta(x_t, t, s) - \frac{1}{2}\text{sg}(D_\theta(x_t, t, r) + D_\theta(x_r, r, s))\|^2$

227 11: $\mathcal{L}_{\text{SC-F2D2}}(\theta) \leftarrow \mathcal{L}_{\text{VM-SC}}(\theta) + \mathcal{L}_{\text{u-SC}}(\theta) + \mathcal{L}_{\text{div-SC}}(\theta) + \mathcal{L}_{\text{D-SC}}(\theta)$

228 12: Update θ w.r.t. $\mathcal{L}_{\text{SC-F2D2}}(\theta)$

229 13: **end for**

230 14: **return** θ

231
232 (c) Φ satisfies the **semigroup property** $\Phi_Y(y, t, s) = \Phi_Y(\Phi_Y(y, t, r), r, s)$ for $t < r < s$.

233
234 We provide the full analysis of the characterization in Appendix A. We refer to the family of
235 algorithms which learns a joint map of this characterization as **fast flow joint distillation (F2D2)**.
236 Notably, we can derive four separate training objectives – one pair for the sampling subsystem and
237 the other pair for the likelihood subsystem – and jointly optimizing them yields a valid flow map
238 for Eq. (2.3). The general F2D2 training objective is:

239
$$\mathcal{L}_{\text{F2D2}}(\theta) := \mathcal{L}_{\text{VM}}(\theta) + \mathcal{L}_{\text{u}}(\theta) + \mathcal{L}_{\text{div}}(\theta) + \mathcal{L}_{\text{D}}(\theta) \quad (3.4)$$

240 where the first two terms optimize for the sampling flow map Φ_X : \mathcal{L}_{VM} enforces the instantaneous
241 velocity matching (i.e. the tangent condition, which is often enforced by the flow matching loss in
242 practice), while \mathcal{L}_{u} enforces one of the flow map conditions from Proposition 3.1 for Φ_X . Similarly,
243 the last two terms optimize for the likelihood flow map Φ_Z : \mathcal{L}_{div} matches the instantaneous divergence,
244 and \mathcal{L}_{D} ensures that Φ_Z satisfies the conditions needed for the joint flow map Φ_Y to be valid
245 according to the conditions in Proposition 3.3.

246
247 3.2 INSTANTIATING F2D2 WITH SHORTCUT AND MEANFLOW

248 Though our method is, in principle, compatible with any flow map-based method, we instantiate
249 our formulation for Shortcut Models (Frans et al., 2024), based on the semigroup property, and for
250 MeanFlow (Geng et al., 2025), based on the Eulerian equation.

251
252 3.2.1 JOINT SHORTCUT: SHORTCUT-F2D2

253 Shortcut models (Frans et al., 2024) enforce the semigroup property (Proposition 3.1 (c)) by applying
254 it to the midpoint between timesteps t and s . This amounts to the shortcut self-consistency loss:

255
$$\mathcal{L}_{\text{u-SC}}(\theta) = \mathbb{E}_{t < s, x_t} \left[\|u_\theta(x_t, t, s) - \frac{1}{2}\text{sg}(u_\theta(x_t, t, r) + u_\theta(\Phi_{X;\theta}(x_t, t, r), r, s))\|^2 \right] \quad (3.5)$$

256 where $\text{sg}(\cdot)$ denotes stop-gradient. Combined with the tangent condition (at $t = s$),

257
$$v_\theta(x, t) = u_\theta(x, t, t), \quad (3.6)$$

258 and training with the flow matching loss as \mathcal{L}_{VM} ,

259
$$\mathcal{L}_{\text{VM-SC}}(\theta) := \mathbb{E}_{t \sim [0, 1], x_0 \sim p_0, x_1 \sim p_1} [\|u_\theta(x_t, t, t) - (x_1 - x_0)\|^2] \quad (3.7)$$

260 this self-consistency loss $\mathcal{L}_{\text{u-SC}}$ serves as \mathcal{L}_{u} and yields a valid flow map by enforcing both the
261 tangent condition and the semigroup property (see Corollary B.1 in Appendix B for details).

262 We convert this method to a joint self-distillation method by introducing the two additional losses:

263
$$\mathcal{L}_{\text{div-SC}}(\theta) := \mathbb{E} [\|D_\theta(x_t, t, t) + \text{div}(u_\theta^-(x_t, t, t))\|^2] \quad (3.8)$$

264
$$\mathcal{L}_{\text{D-SC}}(\theta) := \mathbb{E}_{t < s} \left[\|D_\theta(x_t, t, s) - \frac{1}{2}\text{sg}(D_\theta(x_t, t, r) + D_\theta(\Phi_{X;\theta}(x_t, t, r), r, s))\|^2 \right] \quad (3.9)$$

270 The first loss proceeds by analogy to Eq. (3.6), where the correct D_θ is precisely the instantaneous
 271 velocity field in the Z -component, and this is precisely $\text{div}(u_\theta)$ by Eq. (2.3). The second loss
 272 simply enforces the semigroup property on the Z -component. Take then together, we arrive at
 273 the shortcut variant of F2D2, Shortcut-F2D2, obtained by minimizing the loss

$$\mathcal{L}_{\text{SC-F2D2}}(\theta) := \mathcal{L}_{\text{VM-SC}}(\theta) + \mathcal{L}_{\text{u-SC}}(\theta) + \mathcal{L}_{\text{div-SC}}(\theta) + \mathcal{L}_{\text{D-SC}}(\theta). \quad (3.10)$$

276 We demonstrate the derivation of Shortcut-F2D2 in Appendix A and provide the pseudocode for
 277 Shortcut-F2D2 in Algorithm 1.

279 3.2.2 JOINT MEANFLOW: MEANFLOW-F2D2

280 Alternatively, MeanFlow (Geng et al., 2025) learns the time-averaged velocity $u_\theta(x, t, s)$ by enforcing
 281 the so-called MeanFlow identity $u_\theta(x_t, s, t) = v(x_t, t) + (s-t) \frac{d}{dt} u_\theta(x_t, s, t)$, which we show in
 282 Corollary B.2 in Appendix B solves the Eulerian equation (Proposition 3.1 (b)). To our knowledge,
 283 this is the first proof of this fact.

284 For efficiency, $\frac{d}{dt} u(x, s, t)$ can be computed as $\partial_t u(x, t, s) + \nabla_x u(x, t, s) v(x, t)$ obtained via a
 285 Jacobian–vector product (JVP). The training objective is then to enforce this identity by regressing
 286 the model’s prediction to the target implied thereby:

$$\mathcal{L}_{\text{MF}}(\theta) = \mathbb{E}_{t < s, x_0, x_1} \left[\|u_\theta(x_t, t, s) - \text{sg}((s-t)(\partial_t u(x_t, t, s) + \nabla_x u(x_t, t, s)v(x_t, t)) + v(x_t, t))\|^2 \right], \quad (3.11)$$

290 where $v(x_t, t) = x_1 - x_0$ is the ground truth instantaneous velocity of the interpolant. Importantly,
 291 Eq. (3.11) encapsulates both \mathcal{L}_{VM} and \mathcal{L}_{u} since \mathcal{L}_{MF} recovers the flow matching loss when $s = t$.

292 To extend this method to the joint system, we introduce the additional loss

$$\begin{aligned} \mathcal{L}_{\text{div-MF}}(\theta) = \mathbb{E} & \left[\|D_\theta(x_t, t, s) - \text{sg}((s-t)(\partial_t D_\theta(x_t, t, s) \right. \\ & \left. + \nabla_x D_\theta(x_t, t, s)v(x_t, t)) - \text{div}(u_\theta(x_t, t, t)))\|^2 \right], \end{aligned} \quad (3.12)$$

297 By analogy to \mathcal{L}_{MF} , $\mathcal{L}_{\text{div-MF}}$ obviates the need for explicit divergence matching term \mathcal{L}_{div} .
 298 MeanFlow-F2D2 then amounts to training the objective

$$\mathcal{L}_{\text{MF-F2D2}}(\theta) = \mathcal{L}_{\text{MF}}(\theta) + \mathcal{L}_{\text{div-MF}}(\theta). \quad (3.13)$$

301 We demonstrate the derivation of MeanFlow-F2D2 in Appendix A.

303 3.3 PRACTICAL DESIGN CHOICES

305 While instantaneous velocity supervision for sampling is straightforward to obtain from data, obtaining
 306 reliable and tractable supervision for the instantaneous divergence presents significant challenges.
 307 We address these through several key practical considerations.

308 **Parameter Sharing.** Since both X and Z components of our joint flow map derive from the
 309 same underlying velocity field, learning to predict both components simultaneously is fundamentally
 310 learning two transformations of the same underlying dynamics. As a result, we efficiently
 311 parametrize the joint flow map using a shared backbone network with two separate prediction heads.
 312 This parameter sharing architecture ensures both outputs are derived from a consistent representation
 313 of the flow dynamics and reduces the number of parameters compared to training separate models.

314 **Hutchinson Trace Estimator.** Computing divergence terms $\text{div}(v(x_t, t))$ requires $O(d)$ backward
 315 passes for exact computation, making training prohibitively expensive for high-dimensional data.
 316 Following standard practice (Grathwohl et al., 2018; Lipman et al., 2022; Song et al., 2020), we
 317 employ the Hutchinson trace estimator $\text{div}(v) \approx \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)}[\epsilon^\top \nabla_x v \cdot \epsilon]$ which provides unbiased
 318 estimates with only $O(1)$ computational cost per training step.

319 **Staged Training with Warm Start.** Since divergence supervision depends on having accurate
 320 velocity predictions, we adopt a staged training approach. In practice, we pre-train the sampling
 321 velocity component u_θ alone using existing flow map distillation techniques, which provides a
 322 good initialization for joint training later. Optionally, we can also pre-train a teacher flow matching
 323 model v_ϕ that serves as a reliable source of divergence supervision, replacing the potentially noisy
 $\text{div}(u_\theta(x_t, t, t))$ with the more accurate $\text{div}(v_\phi(x_t, t))$ during joint distillation.

324 **Algorithm 2** Maximum Likelihood Self-Guidance Sampling with F2D2

```

325
326 1:  $x_0 \sim p_0$ 
327 2:  $D \leftarrow D_\theta(x_0, 0, 1)$ 
328 3:  $\mathcal{L}_{\text{NLL}} \leftarrow -\log p_0(x_0) - D$ 
329 4:  $x_0 \leftarrow \text{Adam}(x_0, D)$   $\triangleright$  One step Adam update w.r.t.  $x_0$  optimizing  $\mathcal{L}_{\text{NLL}}$ 
330 5:  $t_0, \dots, t_{K_{\text{samp}}} \leftarrow \text{linspace}(0, 1, K_{\text{samp}} + 1)$ 
331 6: for  $i = 0, \dots, K_{\text{samp}} - 1$  do
332 7:    $u \leftarrow u_\theta(x_i, t_i, t_{i+1})$ 
333 8:    $x_{i+1} \leftarrow x_i + (t_{i+1} - t_i)u$ 
334 9: end for
335 10: return  $x_1$ 
336

```

337 **Shortcut-Distill-F2D2.** To further improve training stability and performance, we propose
 338 Shortcut-Distill, a shortcut model variant that combines the semigroup flow map with a learned
 339 teacher instantaneous velocity. Our three-stage pipeline consists of: (1) **Teacher pre-training:**
 340 Train v_ϕ using standard flow matching; (2) **Shortcut-Distill:** Warm start θ with the teacher parameters
 341 and replace $\mathcal{L}_{\text{div-sc}}$ with teacher supervision: $\mathbb{E}_t [\|u_\theta(x_t, t, t) - v_\phi(x_t, t)\|^2]$; (3) **Joint dis-
 342 tillation:** Warm start θ from sampling distillation, add divergence head and train both components
 343 jointly. This approach maintains the semigroup condition while leveraging a pre-trained velocity
 344 field to ensure the joint flow map are well-aligned.

345
 346 3.4 APPLICATION: MAXIMUM LIKELIHOOD SELF-GUIDANCE WITH F2D2
 347

348 Now that we have access to log-likelihood computation with few NFEs, we can explore various
 349 applications. One particularly interesting one is using the one-step divergence prediction (combined
 350 with the source distribution’s log-likelihood) as a pseudo-likelihood objective for inference-time
 351 optimization. Specifically, we can optimize the initial noise x_0 to improve sample quality before
 352 running the sampling procedure. This approach resembles reward-based initial noise optimization
 353 for one-step generation models (Eyring et al., 2024), except we do not require external reward
 354 models. Instead, we obtain the guidance signal from the model’s own likelihood prediction head –
 355 effectively performing self-guidance at inference time to improve sample quality. This maximum
 356 likelihood self-guidance sampling algorithm is described in Algorithm 2.
 357

358 4 RELATED WORKS
 359

360 **Likelihood computation in diffusion and flow models.** While diffusion and flow-based mod-
 361 els excel at sample generation, their likelihood evaluation remains computationally expensive.
 362 Discrete-time diffusion models compute likelihoods through variational bounds requiring hundreds
 363 of NFEs (Ho et al., 2020; Nichol & Dhariwal, 2021). Continuous formulations enable exact likeli-
 364 hood via the probability flow ODE (Song et al., 2020) but require numerical integration with 100-
 365 1000 NFEs. Prior research has explored various techniques to improve likelihood estimation (Grath-
 366 wohl et al., 2018; Song et al., 2021), but they still require many NFEs for accurate evaluation. In
 367 parallel, normalizing flows (Rezende & Mohamed, 2015) offer exact and tractable log-likelihoods
 368 by using specialized network architecture designs and change-of-variables formula. Many recent
 369 efforts (Zhai et al., 2024; Ho et al., 2019; Chen et al., 2019) aim at scaling up the normalizing flow
 370 principles through better dequantization and and advance architectures.

371 **Accelerating sampling in flow-based models.** Reducing sampling costs has been a major focus
 372 in diffusion and flow matching research. Advanced ODE solvers (Karras et al., 2022; Lu et al.,
 373 2022) leverage the semi-linear structure of the probability flow ODE to reduce discretization error.
 374 Distillation methods (Salimans & Ho, 2022; Sauer et al., 2024) and self-distillation models (Song
 375 & Dhariwal, 2023; Zhou et al., 2025) provide alternative solutions by training student models to
 376 match teacher trajectories or training to perform self-bootstrapping with fewer steps. In particular,
 377 consistency models (Song et al., 2023) and consistency trajectory models (Kim et al., 2023) learn
 378 direct mappings from any point along the trajectory to data.

378
 379
 380
 381
 382
 383
 Table 1: NLL and FID results on CIFAR-10 dataset with different numbers of Euler steps. The flow
 matching model here, which achieves BPD 3.12 as the NLL with 1024 steps and FID 2.60 with 200
 steps, is also the teacher model we use in our Shortcut-Distill. For NLL, the closer to the teacher
 result (3.12 BPD) the better, and for FID, the lower the better. We denote the best results in **bold**, the
 second best with underlines, the overall best results in and invalid predictions in gray color.

Method	8 Steps		4 Steps		2 Steps		1 Step	
	NLL	FID	NLL	FID	NLL	FID	NLL	FID
Flow Matching	-9.93	20.63	-24.01	64.27	-52.85	146.24	-111.19	313.54
Shortcut Model	-12.07	7.10	-28.03	9.63	-60.01	16.04	-124.15	27.28
Shortcut-Distill (Ours)	-11.42	5.01	-26.82	5.41	-57.72	7.13	-119.42	12.75
MeanFlow	-12.33	<u>4.34</u>	-22.73	<u>5.14</u>	-47.31	<u>2.84</u>	-97.71	2.80
Shortcut-F2D2 (Ours)	<u>3.07</u>	8.78	3.26	10.21	2.73	15.58	0.20	27.35
Shortcut-Distill-F2D2 (Ours)	3.12	5.68	<u>2.87</u>	5.96	<u>2.38</u>	7.35	<u>1.62</u>	13.76
MeanFlow-F2D2 (Ours)	1.91	3.78	1.31	4.37	1.29	2.59	3.33	<u>3.02</u>

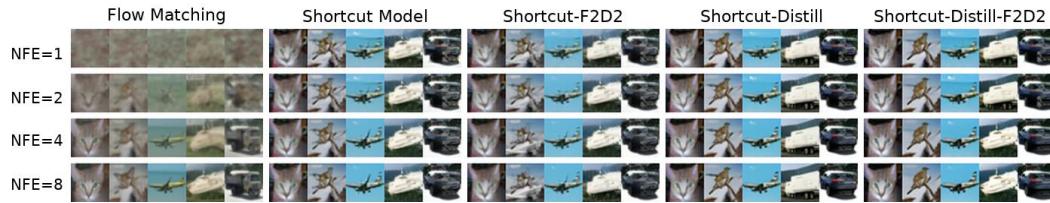


Figure 1: CIFAR-10 generated samples from different models with different numbers of steps.

404
 405 **Flow map-based methods.** Flow maps provide a general framework to model the solution op-
 406 erator of ODEs, enabling direct prediction of integrated trajectories (Kim et al., 2023; Boffi et al.,
 407 2025a;b). Recent works exploit this structure for few-step sampling. For example, as we have shown
 408 above, shortcut models (Frans et al., 2024) imposes semigroup property to learn the flow maps while
 409 MeanFlow (Geng et al., 2025) and Align Your Flow (Sabour et al., 2025) enforce Eulerian condi-
 410 tions. While these methods successfully reduce sampling to less than 10 NFEs, they either abandon
 411 likelihood computation entirely or still require full trajectory integration for likelihood evaluation.

5 EXPERIMENTS

5.1 SETUPS

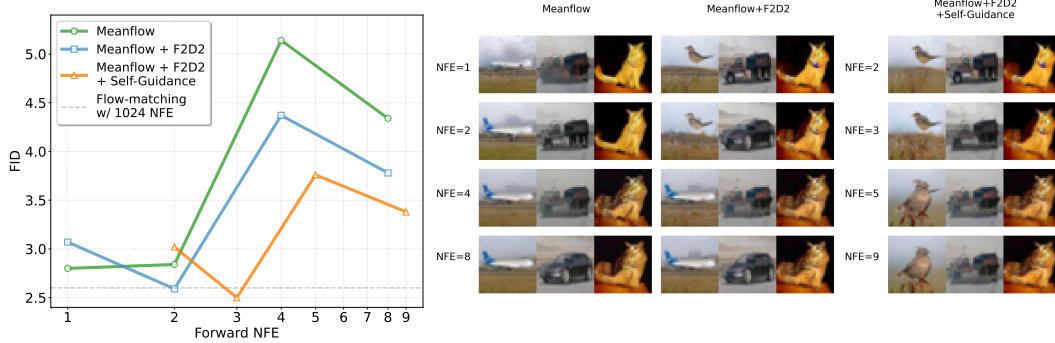
416 We empirically verify the effectiveness of our method on image datasets CIFAR-10 (Krizhevsky
 417 et al.) and ImageNet 64 × 64 (Deng et al., 2009). We evaluate the sample quality using Fréchet In-
 418 ception Distance (FID) (Heusel et al., 2017) on 50K generated images. The negative log-likelihood
 419 (NLL) is measured in bits per dimension (BPD) on the entire test set of CIFAR-10 and a randomly
 420 sampled 10K subset of the ImageNet test set. We compare our method against flow matching (Lip-
 421 man et al., 2022), shortcut models (Frans et al., 2024) and MeanFlow (Geng et al., 2025) as base-
 422 lines, and augment the later two for joint distillation. All models are unconditionally trained on both
 423 datasets. We use 1, 2, 4, 8 Euler steps for sampling and likelihood evaluation in both experiments.
 424 Implementation details about our method and the baselines can be found in Appendix C.

5.2 RESULTS

425 **CIFAR-10** Table 1 and Figure 1 show the quantitative and qualitative comparison on CIFAR-10 re-
 426 spectively. As we can observe, flow matching yields poor FID and invalid NLL estimates in few-step
 427 setting. Shortcut model and MeanFlow achieve significantly better FID and are able to compute NLL
 428 for their ability to recover instantaneous velocity, but their NLL values remain invalid. Incorporat-
 429 ing our proposed F2D2 brings the NLL estimations to a calibrated range close to the teacher’s BPD
 430 across different settings. In particular, both Shortcut-F2D2 and Shortcut-Distill-F2D2 substantially
 431

432 Table 2: Negative log-likelihood (NLL) measured in BPD and FID results on ImageNet 64×64
 433 dataset with different numbers of Euler steps. The flow matching model here, which achieves BPD
 434 3.34 as the NLL with 1024 steps and FID 13.09 with 200 steps, is also the teacher model we use in
 435 our Shortcut-Distill. For NLL, the closer to the teacher result (3.34 BPD) the better, and for FID,
 436 the lower the better. We denote the best results in **bold** and invalid predictions in gray color.
 437

438	Method	8 Steps		4 Steps		2 Steps		1 Step	
		NLL	FID	NLL	FID	NLL	FID	NLL	FID
440	Flow Matching	-6.41	31.60	-15.87	68.55	-35.23	170.00	-74.54	363.39
441	Shortcut-Distill (Ours)	-9.03	19.47	-22.30	21.73	-49.01	28.12	-102.07	42.72
442	Shortcut-Distill-F2D2 (Ours)	3.51	21.91	3.94	24.05	3.97	29.83	1.54	44.02



456 (a) FID comparison across different num- (b) Example samples from various MeanFlow-based models us-
 457 bers of forward NFEs. ing different numbers of forward NFEs.
 458

459 Figure 2: Results of MeanFlow-based methods on CIFAR-10.
 460
 461

462 improves NLL compared to plain their original counterparts while maintaining competitive FID,
 463 indicating that F2D2 can provide reasonable likelihood estimates without sacrificing much sam-
 464 ple quality. Finally, MeanFlow-F2D2 shows FID improvements relative to the original MeanFlow
 465 while simultaneously producing calibrated NLL, demonstrating that F2D2’s potential in providing
 466 complementary training signals that are beneficial to both components.

467 **ImageNet 64×64** Shown in Table 2, flow matching quickly degenerates under few-step sampling,
 468 with invalid NLL and extremely poor FID. Shortcut-Distill improves the few-step FID but still pro-
 469 duces invalid NLL. By contrast, Shortcut-Distill-F2D2 achieves both competitive FID and mean-
 470 ingful likelihoods close to the teacher’s BPD of 3.34 across all step counts. These results further
 471 confirm F2D2’s ability for simultaneous fast sampling and fast likelihood evaluation.

472 5.3 MAXIMUM LIKELIHOOD SELF-GUIDANCE WITH MEANFLOW-F2D2

473 Figure 2 shows the FID and qualitative comparison among different methods built upon MeanFlow
 474 using the number of forward NFE on CIFAR-10. As we can observe, our F2D2 improves the
 475 model’s inference time scaling ability. With additional self-guidance, the model not only surpasses
 476 the baseline MeanFlow performance but also outperforms a 1024-step flow matching model of the
 477 same size, demonstrating the effectiveness its own likelihood predictions as valid signals to guide
 478 the sampling process toward higher-quality generations.

481 5.4 2D CHECKERBOARD

482 In this section, we present a set of comparison of log-likelihood evaluation results on 2D checker-
 483 board, a synthetic dataset with analytically tractable ground truth log-likelihood. As we can observe
 484 in Figure 3, without F2D2, both flow matching and shorcut model catastrophically fail at few-step
 485 log-likelihood estimation. On the other hand, our Shortcut-F2D2 is able to accurately recover the

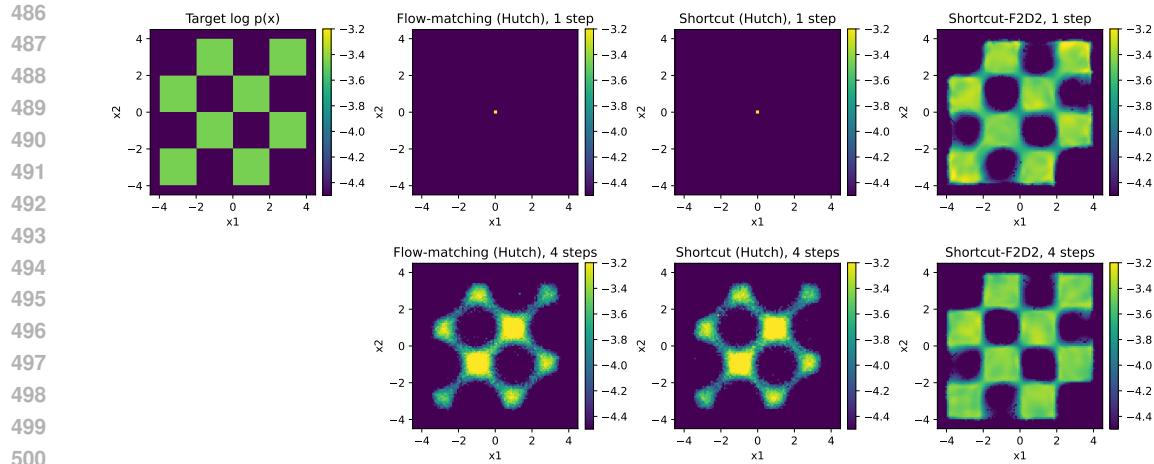


Figure 3: Log-likelihood comparison on 2D checkerboard dataset among different models.

target density distribution even with only 1 NFE, preserving both spatial structure and density values. This directly validates that, while traditional methods fail at few-step likelihood evaluation, F2D2 enables accurate likelihood via joint flow map distillation.

6 CONCLUSION

We present fast flow joint distillation (F2D2), a simple and modular framework that enables both fast sampling and fast likelihood evaluation in flow-based generative models. By jointly distilling the sampling trajectory and divergence computation into a unified flow map, our method simultaneously achieves accurate likelihood evaluation and high sample quality with just a few NFEs. Our experiments on CIFAR-10 and ImageNet 64×64 demonstrate that F2D2 maintains accurate likelihood estimates while preserving sample quality when applied to existing few-step methods including shortcut models and MeanFlow. The efficiency gains from F2D2 enable new algorithmic possibilities, as illustrated by our maximum likelihood self-guidance method, which enables a 2-step MeanFlow model to outperform a 1024-step flow matching model of the same size on CIFAR-10. As flow-based models continue to scale, we believe that efficient likelihood evaluation alongside fast sampling will become increasingly important for enabling new training objectives, model analysis techniques, and downstream applications that require both capabilities.

REPRODUCIBILITY STATEMENT

We provide proofs to our theoretical results in Appendix A and B. We also provide the implementation details to reproduce our algorithm and experimental results in Section 5 and Appendix C.

REFERENCES

Hirotugu Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974.

Michael S Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.

Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions, 2023. URL <https://arxiv.org/abs/2303.08797>.

Black Forest Labs. Flux.1 kontext: Flow matching for in-context image generation and editing in latent space. *arXiv preprint arXiv:2506.15742*, 2025. URL <https://arxiv.org/abs/2506.15742>.

540 Nicholas M. Boffi, Michael S. Albergo, and Eric Vanden-Eijnden. How to build a consistency
 541 model: Learning flow maps via self-distillation, 2025a. URL <https://arxiv.org/abs/2505.18825>.

542

543 Nicholas M. Boffi, Michael S. Albergo, and Eric Vanden-Eijnden. Flow map matching with stochastic
 544 interpolants: A mathematical framework for consistency models, 2025b. URL <https://arxiv.org/abs/2406.07507>.

545

546 Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary
 547 differential equations. *Advances in neural information processing systems*, 31, 2018.

548

549 Ricky TQ Chen, Jens Behrmann, David K Duvenaud, and Jörn-Henrik Jacobsen. Residual flows for
 550 invertible generative modeling. *Advances in neural information processing systems*, 32, 2019.

551

552 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hier-
 553 archical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*,
 554 pp. 248–255, 2009. doi: 10.1109/CVPR.2009.5206848.

555

556 Luca Eyring, Shyamgopal Karthik, Karsten Roth, Alexey Dosovitskiy, and Zeynep Akata. Reno:
 557 Enhancing one-step text-to-image models through reward-based noise optimization. *Advances in
 558 Neural Information Processing Systems*, 37:125487–125519, 2024.

559

560 Ronald A. Fisher. On the mathematical foundations of theoretical statistics. *Philosophical Transac-
 561 tions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical
 562 Character*, 222:309–368, 1922.

563

564 Kevin Frans, Danijar Hafner, Sergey Levine, and Pieter Abbeel. One step diffusion via shortcut
 565 models. *arXiv preprint arXiv:2410.12557*, 2024.

566

567 Ruiqi Gao, Emiel Hoogeboom, Jonathan Heek, Valentin De Bortoli, Kevin P. Murphy, and Tim
 568 Salimans. Diffusion meets flow matching: Two sides of the same coin. 2024. URL <https://diffusionflow.github.io/>.

569

570 Zhengyang Geng, Mingyang Deng, Xingjian Bai, J Zico Kolter, and Kaiming He. Mean flows for
 571 one-step generative modeling. *arXiv preprint arXiv:2505.13447*, 2025.

572

573 Google DeepMind. Veo 3: High-fidelity text- and image-to-video generation with audio.
 574 Model card, Google DeepMind, may 2025. URL <https://storage.googleapis.com/deepmind-media/Model-Cards/Veo-3-Model-Card.pdf>. Accessed 2025-09-21.

575

576 Will Grathwohl, Ricky TQ Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. Ffjord:
 577 Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint
 578 arXiv:1810.01367*, 2018.

579

580 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter.
 581 Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in
 582 neural information processing systems*, 30, 2017.

583

584 Jonathan Ho, Xi Chen, Aravind Srinivas, Yan Duan, and Pieter Abbeel. Flow++: Improving flow-
 585 based generative models with variational dequantization and architecture design. In *International
 586 conference on machine learning*, pp. 2722–2730. PMLR, 2019.

587

588 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
 589 neural information processing systems*, 33:6840–6851, 2020.

590

591 Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-
 592 based generative models. *Advances in neural information processing systems*, 35:26565–26577,
 593 2022.

594

595 Dongjun Kim, Chieh-Hsin Lai, Wei-Hsiang Liao, Naoki Murata, Yuhta Takida, Toshimitsu Uesaka,
 596 Yutong He, Yuki Mitsufuji, and Stefano Ermon. Consistency trajectory models: Learning proba-
 597 bility flow ode trajectory of diffusion. *arXiv preprint arXiv:2310.02279*, 2023.

594 Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced re-
 595 search). URL <http://www.cs.toronto.edu/~kriz/cifar.html>.

596

597 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
 598 for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.

599 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and
 600 transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.

601

602 Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast
 603 ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in neural*
 604 *information processing systems*, 35:5775–5787, 2022.

605 Jerzy Neyman and Egon S. Pearson. On the problem of the most efficient tests of statistical hypothe-
 606 ses. *Philosophical Transactions of the Royal Society A*, 231:289–337, 1933.

607

608 Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models.
 609 In *International conference on machine learning*, pp. 8162–8171. PMLR, 2021.

610

611 OpenAI. Sora: Video generation models as world simulators. Technical Report, OpenAI, 2024. URL <https://openai.com/index/video-generation-models-as-world-simulators/>. Accessed: 2024-02-15.

612

613 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
 614 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to fol-
 615 low instructions with human feedback. *Advances in neural information processing systems*, 35:
 616 27730–27744, 2022.

617

618 Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv
 619 Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, David Yan, Dhruv Choudhary, Dingkang
 620 Wang, Geet Sethi, Guan Pang, Haoyu Ma, Ishan Misra, Ji Hou, Jialiang Wang, Kiran Ja-
 621 gadeesh, Kunpeng Li, Luxin Zhang, Mannat Singh, Mary Williamson, Matt Le, Matthew Yu,
 622 Mitesh Kumar Singh, Peizhao Zhang, Peter Vajda, Quentin Duval, Rohit Girdhar, Roshan Sum-
 623 baly, Sai Saketh Rambhatla, Sam Tsai, Samaneh Azadi, Samyak Datta, Sanyuan Chen, Sean
 624 Bell, Sharadh Ramaswamy, Shelly Sheynin, Siddharth Bhattacharya, Simran Motwani, Tao Xu,
 625 Tianhe Li, Tingbo Hou, Wei-Ning Hsu, Xi Yin, Xiaoliang Dai, Yaniv Taigman, Yaqiao Luo, Yen-
 626 Cheng Liu, Yi-Chiao Wu, Yue Zhao, Yuval Kirstain, Zecheng He, Zijian He, Albert Pumarola,
 627 Ali Thabet, Artsiom Sanakoyeu, Arun Mallya, Baishan Guo, Boris Araya, Breena Kerr, Car-
 628 leigh Wood, Ce Liu, Cen Peng, Dimitry Vengertsev, Edgar Schonfeld, Elliot Blanchard, Felix
 629 Juefei-Xu, Fraylie Nord, Jeff Liang, John Hoffman, Jonas Kohler, Kaolin Fire, Karthik Sivaku-
 630 mar, Lawrence Chen, Licheng Yu, Luya Gao, Markos Georgopoulos, Rashel Moritz, Sara K.
 631 Sampson, Shikai Li, Simone Parmeggiani, Steve Fine, Tara Fowler, Vladan Petrovic, and Yuming
 632 Du. Movie gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720*, 2024.
 633 URL <https://arxiv.org/abs/2410.13720>. Submitted October 2024.

634

635 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
 636 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances*
 637 *in neural information processing systems*, 36:53728–53741, 2023.

638

639 Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with
 640 vq-vae-2. *Advances in neural information processing systems*, 32, 2019.

641

642 Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *Interna-
 643 tional conference on machine learning*, pp. 1530–1538. PMLR, 2015.

644

645 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 646 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-
 647 ence on computer vision and pattern recognition*, pp. 10684–10695, 2022.

648

649 Amirmojtaba Sabour, Sanja Fidler, and Karsten Kreis. Align your flow: Scaling continuous-time
 650 flow map distillation. *arXiv preprint arXiv:2506.14603*, 2025.

651

652 Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv*
 653 *preprint arXiv:2202.00512*, 2022.

648 Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion dis-
649 tillation. In *European Conference on Computer Vision*, pp. 87–103. Springer, 2024.
650

651 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
652 optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

653 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
654 Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathemati-
655 cal reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
656

657 Yang Song and Prafulla Dhariwal. Improved techniques for training consistency models. *arXiv
658 preprint arXiv:2310.14189*, 2023.

659 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
660 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint
661 arXiv:2011.13456*, 2020.

662 Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of
663 score-based diffusion models. *Advances in neural information processing systems*, 34:1415–
664 1428, 2021.

665 Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. 2023.

666 Shuangfei Zhai, Ruixiang Zhang, Preetum Nakkiran, David Berthelot, Jiatao Gu, Huangjie Zheng,
667 Tianrong Chen, Miguel Angel Bautista, Navdeep Jaitly, and Josh Susskind. Normalizing flows
668 are capable generative models. *arXiv preprint arXiv:2412.06329*, 2024.

669 Linqi Zhou, Stefano Ermon, and Jiaming Song. Inductive moment matching. *arXiv preprint
670 arXiv:2503.07565*, 2025.

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 **A CHARACTERIZATION OF THE JOINT FLOW MAP**
 703

704 In this section, we provide an analytical characterization of the joint flow map for the combined
 705 sampling and likelihood dynamics Eq. (2.3). For the simplicity of notations, we define the joint
 706 system

$$\begin{aligned} \frac{d}{dt}x_t &= v(x_t, t), & x(0) &= x_0 \sim \mathcal{N}(0, I) \\ \frac{d}{dt}z_t &= -\text{div}(v(x_t, t)), & z_0 &= \log p_0(x_0) \end{aligned} \quad (\text{A.1})$$

712 where we will use the shorthand z_t for $\log p_t(x_t)$. Moreover, we define the right-hand side
 713 of Eq. (A.1) as $g(y, t)^\top = (v(x, t), -\text{div}(v(x_t, t)))^\top$ where $y = (x, z)^\top$.

714 We then define the flow maps for the two subsystems in Eq. (A.1) as

$$\begin{aligned} \Phi_X(x, t, s) &= x + (s - t)u(x, t, s), \\ \Phi_Z(x, z, t, s) &= z + (s - t)D(x, t, s) \end{aligned} \quad (\text{A.2})$$

718 In Eq. (A.2), we note that the hierarchical structure in Eq. (A.1) is explicit, and that the function D
 719 only depends on x and not on z . We then define the joint flow map as

$$\begin{aligned} \Phi_Y(y, t, s) &= \begin{bmatrix} \Phi_X(x, t, s) \\ \Phi_Z(x, z, t, s) \end{bmatrix} = y + (s - t)f(x, t, s), \\ f(x, t, s) &= \begin{bmatrix} u(x, t, s) \\ D(x, t, s) \end{bmatrix} \end{aligned} \quad (\text{A.3})$$

726 We first recall the simple tangent identity denoted the tangent condition by Boffi et al. (2025a), also
 727 leveraged by (Kim et al., 2023; Frans et al., 2024; Geng et al., 2025), which allows us to recover the
 728 instantaneous velocity and divergence in the $s \rightarrow t$ limit:

729 **Lemma A.1** (Tangent Condition). *The flow map $\Phi_Y(y, t, s)$ for the joint system Eq. (A.1) satisfies*
 730 $\lim_{t \rightarrow s} \partial_s \Phi_Y(y, t, s) = g(y, s) = f(x, s, s)$. *In particular, $u(x, s, s) = v(x, t)$ and $D(x, s, s) =$*
 731 $-\text{div}(v(x, s))$.

733 *Proof.* The proof follows by application of Lemma 2.1 from Boffi et al. (2025a). \square

735 Now, given Eq. (A.3), we may now state the following proposition, which is based on an identity
 736 similar to Lemma A.1 in reverse.

738 **Proposition A.2** (Characterization of the Joint Flow Map). *Let $\Phi_Y(y, t, s) = y + (s - t)f(x, t, s)$*
 739 *satisfy $f(x, s, s) = g(y, s)$ where $g(y, t)^\top = (v(x, t), -\text{div}(v(x_t, t)))^\top$ denotes the dynamics for*
 740 *the joint sampling and likelihood system Eq. (A.1). Then, $\Phi_Y(y, t, s)$ is the flow map for the joint*
 741 *system if and only if any of the following conditions are satisfied:*

742 1. (Lagrangian condition) $Z_{s,t}$ satisfies the Lagrangian equation

$$\partial_s \Phi_Y(y, t, s) = f(\Phi_Y(y, t, s), s, s) \quad \forall (y, t, s) \in \mathbb{R}^{d+1} \times [0, 1]^2. \quad (\text{A.4})$$

746 2. (Eulerian condition) $\Phi_Y(y, t, s)$ satisfies the Eulerian equation

$$\partial_t \Phi_Y(y, t, s) + \nabla_y \Phi_Y(y, t, s) f(y, t, t) = 0, \quad \forall (y, t, s) \in \mathbb{R}^{d+1} \times [0, 1]^2. \quad (\text{A.5})$$

750 3. (Semigroup condition) $Z_{s,t}$ satisfies the semigroup property

$$\Phi_Y(y, t, s) = \Phi_Y(\Phi_Y(y, t, r), r, s), \quad \forall (y, t, r, s) \in \mathbb{R}^{d+1} \times [0, 1]^3, t < r < s. \quad (\text{A.6})$$

752 *Proof.* The proof follows by application of Proposition 2.2 from Boffi et al. (2025a) applied to the
 753 joint system. \square

756 Proposition A.2 gives three characterizations of the joint flow map that may be used to devise few-
 757 step flow-based training algorithms. In each case, by definition of Eq. (A.3), the X block reduces
 758 to the flow map characterizations introduced in Boffi et al. (2025b;a) for the sampling system. The
 759 second block for the likelihood dynamics is new, which we focus on now to instantiate the resulting
 760 equations.

761 **Lagrangian likelihood equation.** By inspection, Eq. (A.4) leads to the equation
 762

$$763 D(x, t, s) = -\text{div}(v(\Phi_X(x, t, s), s)) - (s - t)\partial_s D(\Phi_X(x, t, s), t, s). \quad (\text{A.7})$$

764 Squaring the residual leads to the objective function
 765

$$766 \mathcal{L}(\hat{D}) = \mathbb{E} \left[\|\hat{D}(x_t, t, s) - (-\text{div}(v(\Phi_X(x_t, t, s), s)) - (s - t)\partial_s D(\Phi_X(x_t, t, s), t, s))\|^2 \right] \quad (\text{A.8})$$

768 At training time, because we don't have access to the true $-\text{div}(v)$ or the true sampling flow map
 769 $\Phi_X(x, t, s)$, we may replace them by their self-consistent estimates,

$$770 \mathcal{L}(\theta) = \mathbb{E} \left[\|D_\theta(x_t, t, s) - \text{sg}(-\text{div}(u_\theta(\Phi_{X;\theta}(x_t, t, s), s)) - (s - t)\partial_s D_\theta(\Phi_{X;\theta}(x_t, t, s), t, s))\|^2 \right] \quad (\text{A.9})$$

772 where we have also placed a stopgrad operator to avoid backpropagation through Jacobian-vector
 773 products and to control the flow of information from the teacher to the student. This gives the
 774 Lagrangian likelihood self-distillation algorithm.

775 **Eulerian likelihood equation.** To derive our Eulerian schemes, we first note that
 776

$$777 \nabla_y \Phi_Y(y, t, s) = \begin{bmatrix} \nabla_x \Phi_X(x, s, t) & 0 \\ \nabla_x \Phi_Z(x, z, t, s) & \nabla_z \Phi_Z(x, z, t, s) \end{bmatrix}. \quad (\text{A.10})$$

779 Hence to compute the second component of the Eulerian equation, we must collect some simple
 780 algebraic identities,

$$782 \begin{aligned} \partial_t \Phi_Z(x, z, t, s) &= -D(x, t, s) + (s - t)\partial_t D(x, t, s), \\ 783 \nabla_x \Phi_Z(x, z, t, s) &= (s - t)\nabla_x D(x, t, s), \\ 784 \nabla_z \Phi_Z(x, z, t, s) &= I. \end{aligned} \quad (\text{A.11})$$

785 Using the above, we find that the Eulerian relation for Y becomes
 786

$$787 -D(x, t, s) + (s - t)\partial_t D(x, t, s) + (s - t)\nabla_x D(x, t, s)v(x, t) - \text{div}(v(x, t)) = 0. \quad (\text{A.12})$$

788 We may enforce this equation by minimizing the square residual,

$$790 \mathcal{L}(\hat{D}) = \mathbb{E} \left[\|D(x_t, t, s) - ((s - t)\partial_t D(x_t, t, s) + (s - t)\nabla_x D(x_t, t, s)v(x_t, t) - \text{div}(v(x_t, t)))\|^2 \right]. \quad (\text{A.13})$$

793 At training time, we again place a $\text{sg}(\cdot)$ operator to avoid backpropagating through the derivatives,

$$794 \mathcal{L}(\theta) = \mathbb{E} \left[\|D_\theta(x_t, t, s) - \text{sg}((s - t)\partial_t D_\theta(x_t, t, s) + (s - t)\nabla_x D_\theta(x_t, t, s)v(x_t, t) - \text{div}(v(x_t, t)))\|^2 \right]. \quad (\text{A.14})$$

795 In the above, we do not have access to the ideal $v(x_t, t)$ nor $-\text{div}(v(x_t, t))$. However, we observe
 796 that because of the placement of $\text{sg}(\cdot)$, resulting gradient will be linear in $v(x_t, t)$, so that we may
 797 replace it by its Monte Carlo estimate. In practice, this reduces to conditional-OT flow matching.
 798 Second, we replace $-\text{div}(v(x_t, t))$ by the self-consistent estimate $-\text{div}(u_\theta(x_t, t, t))$, leading to
 799

$$800 \mathcal{L}(\theta) = \mathbb{E} \left[\|D_\theta(x_t, t, s) - \text{sg}((s - t)\partial_t D_\theta(x_t, t, s) + (s - t)\nabla_x D_\theta(x_t, t, s)(x_1 - x_0) - \text{div}(u_\theta(x_t, t, t)))\|^2 \right]. \quad (\text{A.15})$$

802 This gives the MeanFlow-F2D2 algorithm.

803 **Semigroup property.** Last, we consider the semigroup approach. The second block is given by,
 804

$$805 Y_{s,t}(z) = Y_{u,t}(Y_{s,u}(z)). \quad (\text{A.16})$$

$$806 \Phi_Y(y, t, s) = \Phi_Y(\Phi_Y(y, t, r), r, s) \quad (\text{A.17})$$

807 Writing this out using Eq. (A.3), we find that
 808

$$809 \begin{aligned} z + (s - t)D(t, s) &= \Phi_Z(x, z, t, r) + (s - r)D(\Phi_Z(x, z, t, r), r, s), \\ \iff z + (s - t)D(x, t, s) &= z + (r - t)D(x, t, r) + (s - r)D(\Phi_X(x, t, r), r, s), \end{aligned} \quad (\text{A.18})$$

810 Setting $r = \frac{1}{2}(t+s)$ recovers a continuous limit of shortcut models, as shown in Boffi et al. (2025a).
 811 In this case, Eq. (A.18) becomes

812

$$813 D(x, t, s) = \frac{1}{2} (D(x, t, r) + D(\Phi_X(x, t, r), r, s)), \quad (A.19)$$

814

815 Squaring the residual gives

816

$$817 \mathcal{L}(\hat{D}) = \mathbb{E} \left[\|\hat{D}(x_t, t, s) - \frac{1}{2} (D(x_t, t, r) + D(\Phi_X(x_t, t, r), r, s))\|^2 \right], \quad (A.20)$$

818

819 where as in the above Eulerian and Lagrangian approaches, we have replaced the ideal flow map X
 820 by the self-consistent estimate. Again, to control the flow of information we may place a stopgrad,

821

$$822 \mathcal{L}(\theta) = \mathbb{E} \left[\|D_\theta(x_t, t, s) - \frac{1}{2} \text{sg}(D_\theta(x_t, t, r) + D_\theta(\Phi_{X;\theta}(x_t, t, r), r, s))\|^2 \right], \quad (A.21)$$

823

824 which gives our Shortcut-F2D2 algorithm.

826 B ADDITIONAL THEORETICAL ANALYSIS

827 **Corollary B.1.** *Shortcut models enforce semigroup property with their self-consistency loss.*

828 *Proof.* The same proof as the semigroup property for Φ_Y holds for Φ_X . \square

829 **Corollary B.2.** *MeanFlow (Geng et al., 2025) directly enforces the Eulerian condition with Mean-
 830 Flow identity.*

831 *Proof.* Since $u(x_t, t, t) = v(x_t, t)$ is implied by the linear parametrization in Equation 3.1, we can
 832 have

833

$$\begin{aligned} \partial_t \Phi(x_t, t, s) + \nabla_{x_t} \Phi(x_t, t, s) u(x_t, t, t) &= 0 \\ \partial_t (x_t + (s-t)u(x_t, t, s)) + \nabla_{x_t} (x_t + (s-t)u(x_t, t, s)) v(x_t, t) &= 0 \\ -u(x_t, t, s) + (s-t)\partial_t u(x_t, t, s) + (I + (s-t)\nabla_{x_t} u(x_t, t, s)) v(x_t, t) &= 0 \\ (s-t)(\partial_t u(x_t, t, s) + \nabla_{x_t} u(x_t, t, s)v(x_t, t)) + v(x_t, t) &= u(x_t, t, s) \\ (s-t)\frac{d}{dt}u(x_t, t, s) + v(x_t, t) &= u(x_t, t, s) \end{aligned}$$

834

835 which is exactly MeanFlow identity. \square

	Unconditional CIFAR-10	ImageNet-64 × 64
Noise embedding	Positional	Positional
Channels	128	192
Channels multiple	1,2,2,2	1,2,3,4
Attention resolution	16	32,16,8
Residual blocks per resolution	4	3
Dropout	0.13	0.1
Batch size	512	1024
GPUs	8 L40S	8 L40S
Iterations	150k	125k
Learning Rate	1e-4 (constant)	1e-4 (constant)
Warmup Steps	30k (linear warmup)	60k (linear warmup)
Precision	float32	float32
Optimizer	RAdam	RAdam
EMA rate	0.9999	0.9999

Table 3: Training hyperparameters for flow-matching.

C IMPLEMENTATION DETAILS

Flow-matching. We use linear interpolation to generate training targets. Flow-matching models are trained on both unconditional CIFAR-10 and ImageNet-64 × 64 using the configurations in Table 3, with 200 standard Euler steps for sampling. This model serves as the teacher model for the second distillation stage.

Shortcut Model. We reimplement the shortcut model (Frans et al. (2024)) for CIFAR-10. Following their method, we use discrete timesteps and set 1/128 as the smallest unit of time for approximating the ODE. We consider 8 possible shortcut lengths ranging from $(1, 1/2, \dots, 1/128)$. Similarly, we divide the batch into 3/4 for flow-matching training and 1/4 for self-consistency training. However, instead of DiT, we use a U-Net backbone with configurations in Table 3. We train for 100k iterations on CIFAR-10 and sample with 1, 2, 4, and 8 standard Euler steps. Additionally, we follow the original paper and parametrize the model to take x_t , t and $s - t$ as inputs.

Shortcut-Distill. We use the velocity predicted by the flow-matching model as the target for the flow-matching loss, instead of $x_{\text{data}} - x_{\text{noise}}$. All other configurations remain the same as in the Shortcut Model.

MeanFlow. We directly use the pre-trained model weights and parametrization provided by the official PyTorch repository to conduct the CIFAR-10 MeanFlow experiments. The provided pre-trained MeanFlow model has the same size as all other models we implement for CIFAR-10.

Log-likelihood estimation with vanilla flow map models. Since Flow maps can recover instantaneous velocity through their tangent condition $u_\theta(x_t, t, t) \approx v(x_t, t)$, flow map models trained only for sampling can, in principle, evaluate likelihood by computing $\text{div}(u_\theta(x_t, t, t))$ and solving for Eq 2.3 and 2.4. As a result, we use Euler solver with this formulation to produce the NLL estimation of the baseline Shortcut, Shortcut-Distill and MeanFlow models.

F2D2 for Shortcut Model and Shortcut-Distill. For F2D2 training, we add a scalar head after the UNet decoder. The scalar head is implemented as an MLP: the decoder’s final feature map is first flattened, then passed through two fully connected layers with SiLU activations, and finally projected to a single scalar. This head maps the spatial feature representation into a scalar value, which we use to predict divergence.

For CIFAR-10, the hidden sizes are 128 and 64, while for ImageNet-64 × 64 they are 64 and 16. We linearly warm-start the model with teacher weights and train it using the flow-matching loss, with the teacher model taken from either the Shortcut Model or the Shortcut Model-Distill. We also adopt the same discrete timesteps as in the Shortcut Model. To balance the four losses, we scale down

918 Table 4: Sampling speed (wall-clock time and per-image latency) for different NFEs.
919

920 Dataset	921 Hardware	922 8 NFE	923 4 NFE	924 2 NFE	925 1 NFE
CIFAR-10	4×L40S	146 s (2.92 ms/img)	83 s (1.66 ms/img)	51 s (1.02 ms/img)	40 s (0.80 ms/img)
ImageNet-64×64	8×L40S	289 s (5.78 ms/img)	160 s (3.20 ms/img)	95 s (1.90 ms/img)	63 s (1.26 ms/img)

926 the divergence distillation targets by a factor of 20,000 for CIFAR-10 and 300,000 for ImageNet-
927 64×64. We train Shortcut-Distill-F2D2 for 10k iterations on CIFAR-10, 10k iterations on ImageNet-
928 64×64. We also train Shortcut-F2D2 for 10k iterations on CIFAR-10. We stop the training when
929 the calibrated value of bpd is reached. All other configurations remain the same. Although our
930 derivation matches $-\text{div}(v)$, our implementation parametrizes to predict $\text{div}(v)$, which yields an
931 equivalent formulation as long as the training loss and the sampling process as the appropriate signs.
932

933 **F2D2 for MeanFlow.** We use the same scalar head for CIFAR-10 with Shortcut-Distill-F2D2 and
934 train for an additional 50 epochs using the same configurations as in the original PyTorch imple-
935 mentation. We use the pre-trained flow matching model to provide the instantaneous divergence
936 supervision.

937 **Maximum Likelihood Self-Guidance with F2D2.** We first predict the negative likelihood using
938 randomly sampled noise with 1-step divergence prediction and take this negative likelihood as the
939 loss. We then update the noise with one step of Adam, using a learning rate of 1×10^{-3} for 1-step
940 sampling and 5×10^{-3} for 2-, 4-, and 8-step sampling. We then evaluate FID using 1, 2, 4, and 8
941 standard Euler steps for sampling.

942 **BPD.** We use the same method as Lipman et al. (2022) to compute the BPD. To compute BPD with
943 F2D2 models, we discretize the interval from 1 to 0 into equal segments, while the uniform step size
944 serves as the second time input.

945 **2D Checkerboard.** We follow the experimental setup of Boffi et al. (2025a). The model is a 4-layer
946 MLP with 512 hidden units per layer and GELU activations. Both the flow-matching baseline and
947 the Shortcut model are trained for 150k iterations with a batch size of 100k and an initial learning
948 rate of 10^{-3} , using a square-root decay schedule after 35k steps. For the Shortcut model, each
949 batch is divided in a 3 : 1 ratio between the flow-matching loss and the self-distillation loss. For
950 Shortcut-F2D2, we extend the output layer with an additional dimension and initialize training from
951 the pretrained Shortcut model, running an extra 27k iterations.

952 **Computation Cost and Runtime.** We conduct all our training on an 8-GPU L40S node. To train
953 a typical F2D2 CIFAR10 model, it takes around 1 day for the Flow-matching teacher, 1 day for the
954 vanilla Shortcut or Shortcut-Distill baselines, and 5 hours for the F2D2 finetuning on top of it. For
955 ImageNet64x64 models, it takes around 8 days for the flow matching teacher, 6 days for the vanilla
956 shortcut or shortcut-distill model, and 5 hours for the F2D2 finetuning. For inference run time, we
957 provide the wall-clock time to generate 50k samples for each dataset using different numbers of
958 NFEs in Table 4.

959 D ADDITIONAL ABLATIONS

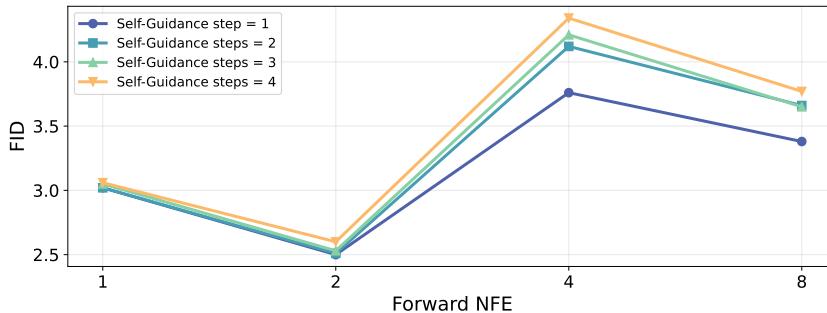
960 In this section, we present additional ablation study to further investigate the behavior of our F2D2
961 models.

962 With variance around 1.8×10^6 each time step, the Hutchinson estimator does produce values that are
963 significantly larger in magnitude than the sampling range. Therefore, as we have mentioned in the
964 Appendix C, we follow common deep learning practices and apply a 5×10^{-5} scaling factor to the
965 divergence target in order for the neural network to better learn the prediction. Here in Table 5, we
966 present an ablation study on the divergence scaling factor to validate this choice of hyperparameter.
967 As we can observe, appropriate scaling enables simultaneous calibrated likelihood prediction and
968 high sample quality, while scaling too strong can produce invalid NLL and scaling too weak can
969 degrade overall performance.

970 Figure 4 shows the self-guidance sampling algorithm with multi-step Adam optimization. As we
971 can observe, while adding optimization steps does not further improve the FID, the performance

972
973
974 Table 5: Ablation study on the divergence scaling factor.
975
976
977
978
979

Scaling	8 Steps		4 Steps		2 Steps		1 Step	
	NLL	FID	NLL	FID	NLL	FID	NLL	FID
5×10^{-4}	3.27	5.16	2.43	6.31	1.26	7.06	-2.92	13.16
5×10^{-5}	3.12	5.68	2.87	5.96	2.38	7.35	1.62	13.76
5×10^{-6}	5.38	15.57	5.44	25.20	5.50	10.96	5.60	23.18

980
981
982
983
984
985
986
987
988
989
990
991
992 Figure 4: FID results of multi-step self-guidance sampling on CIFAR10.
993
994

995 also does not collapse. We are interested in other variants of our self-guidance algorithm to better
996 improve the performance as exciting future work directions.

997 Finally, we conduct additional study on the effect that different loss components have on the sample
998 quality as well as likelihood estimation accuracy. In particular, we slightly modify Equation 3.10 to
999 add an additional weighting scalar λ to $\mathcal{L}_{\text{div-sc}} + \mathcal{L}_{\text{D-sc}}$ and compare the performance under various
1000 weighting.

$$\mathcal{L}_{\text{SC-F2D2}}(\theta) := \mathcal{L}_{\text{VM-sc}}(\theta) + \mathcal{L}_{\text{u-sc}}(\theta) + \lambda(\mathcal{L}_{\text{div-sc}}(\theta) + \mathcal{L}_{\text{D-sc}}(\theta)) \quad (\text{D.1})$$

1001 Table 6 shows the Shortcut-Distill-F2D2 results with $\lambda = 0.01, 0.1, 1, 10$, with $\lambda = 1$ equivalent
1002 to our original setting and the performance of Shortcut-Distill without F2D2 as reference. As we
1003 can observe, adding an additional loss weighting can potentially help F2D2 achieve a better sweet
1004 spot in balancing sample quality and likelihood estimation: While our original equal weighting loss
1005 ($\lambda = 1$) obtains the most calibrated likelihood, it slightly penalizes the FID scores. Similarly,
1006 significantly down-weighting ($\lambda = 0.01$) further improves FID but loses the ability to produce accurate
1007 likelihood. However, when choosing an appropriate λ (in particular when $\lambda = 0.1$), this additional
1008 weighting can enable the model to produce both FID scores that match or even slightly improve the
1009 original Shortcut-Distill performance, and relatively calibrated log-likelihood estimation.
1010

1011 The experiment in Table 6 opens the door to many intriguing future research directions that can po-
1012 tentially further improve the performance. These include more fine-grained loss weighting scheme
1013 and developing training schedules similar to the ones proposed in [Frans et al. \(2024\)](#) and [Geng et al. \(2025\)](#), where the flow matching objective and the self-consistency objective are optimized in
1014 separate portions of the training.
1015

1016 Table 6: Ablation study on different loss weighting on $\mathcal{L}_{\text{div-sc}} + \mathcal{L}_{\text{D-sc}}$ in Shortcut-Distill-F2D2.
1017

Scaling	8 Steps		4 Steps		2 Steps		1 Step	
	NLL	FID	NLL	FID	NLL	FID	NLL	FID
Shortcut-Distill	-11.42	5.01	-26.82	5.41	-57.72	7.13	-119.42	12.75
10	1.47	11.66	0.82	11.23	-0.07	10.11	-2.19	18.27
1	3.12	5.68	2.87	5.96	2.38	7.35	1.92	13.76
0.1	2.59	5.01	2.00	5.33	1.78	7.01	0.68	13.04
0.01	2.04	4.87	1.51	5.28	0.59	6.76	-1.68	12.70

1026 E ADDITIONAL RESULTS

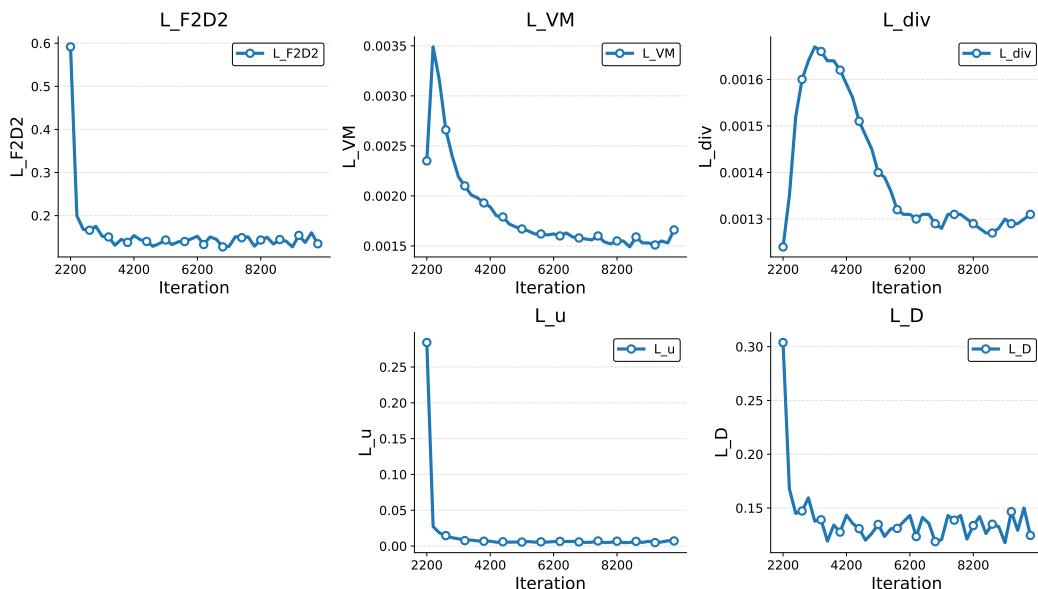
1028 In this section, we present additional experimental results. Specifically, in Table 7 we showcase the
 1029 NLL per sample error w.r.t. the teacher prediction on CIFAR10. As we can observe, the F2D2 variants
 1030 lower the error by $7\text{-}171\times$ in comparison to the baselines. This evaluation directly verifies that
 1031 our F2D2 can produce sample-level calibrated likelihood estimations that are suitable for practical
 1032 applications, not merely matching the summary statistics.

1033 Figure 5 demonstrate the training dynamics of our Shortcut-Distill-F2D2 model with the training
 1034 loss curve, where we can observe stable training with expected small fluctuations native to uniformly
 1035 random timestep selection at each iteration.

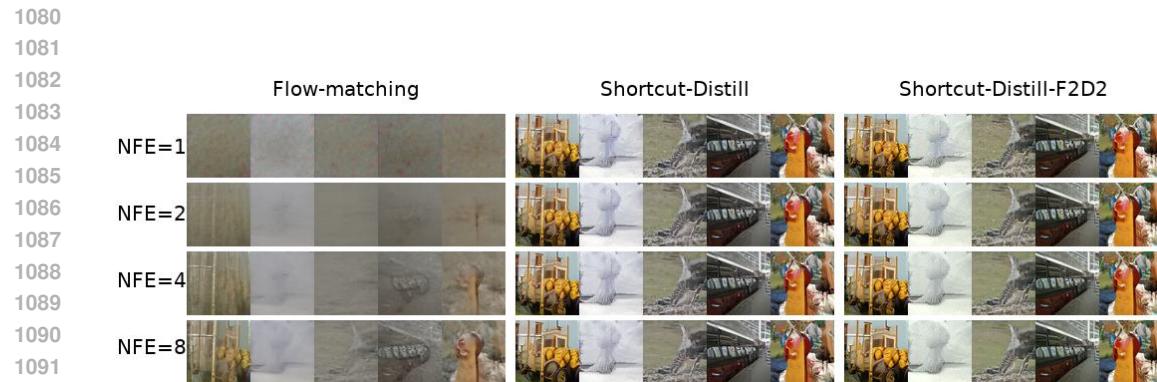
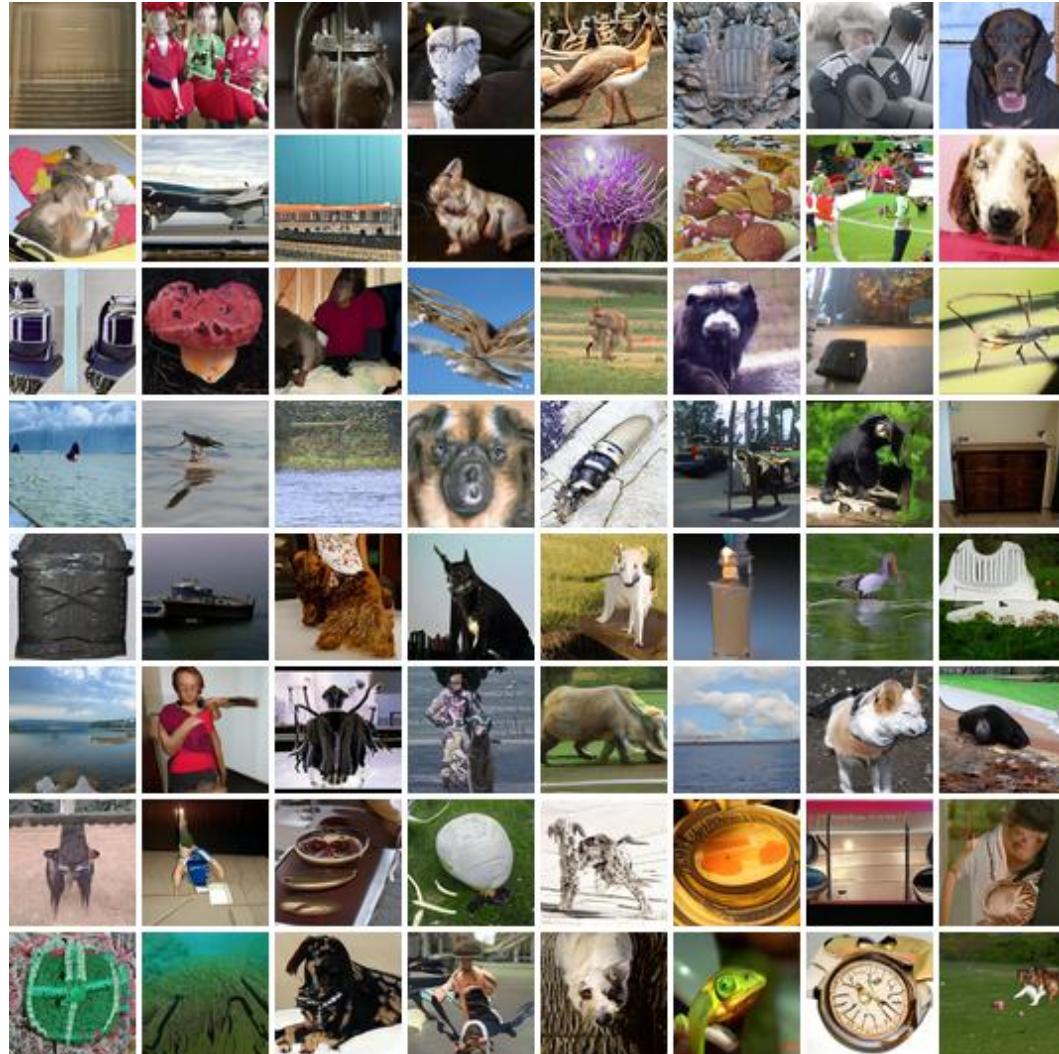
1036 We also provide additional qualitative results in Figure 6,7,8,9,10,11,12,13,14,15,16,17,18,19,20.
 1037 All images shown in this section are non-cherry picked results.

1039 Table 7: NLL Mean absolute per sample error w.r.t the teacher’s prediction in BPD on CIFAR10. It
 1040 shows that the F2D2 variants lower the error by $7\text{-}171\times$ in comparison to the baselines.

Method	8 steps	4 steps	2 steps	1 step
Shortcut	15.11	30.97	62.85	126.58
Shortcut-Distill	14.47	29.79	60.61	121.99
Meanflow	12.04	24.20	49.42	100.99
Shortcut-F2D2	<u>0.52</u>	<u>0.55</u>	<u>1.01</u>	3.08
Shortcut-Distill-F2D2	0.41	0.50	0.90	<u>1.69</u>
MeanFlow-F2D2	1.62	1.78	1.79	0.59



1071 Figure 5: The Shortcut-Distill-F2D2 loss curve on CIFAR10.

Figure 6: Imagenet 64×64 unconditional generation.Figure 7: 8-step unconditional ImageNet 64×64 generation with our Shortcut-Distill.

1132

1133

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143

1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177

1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

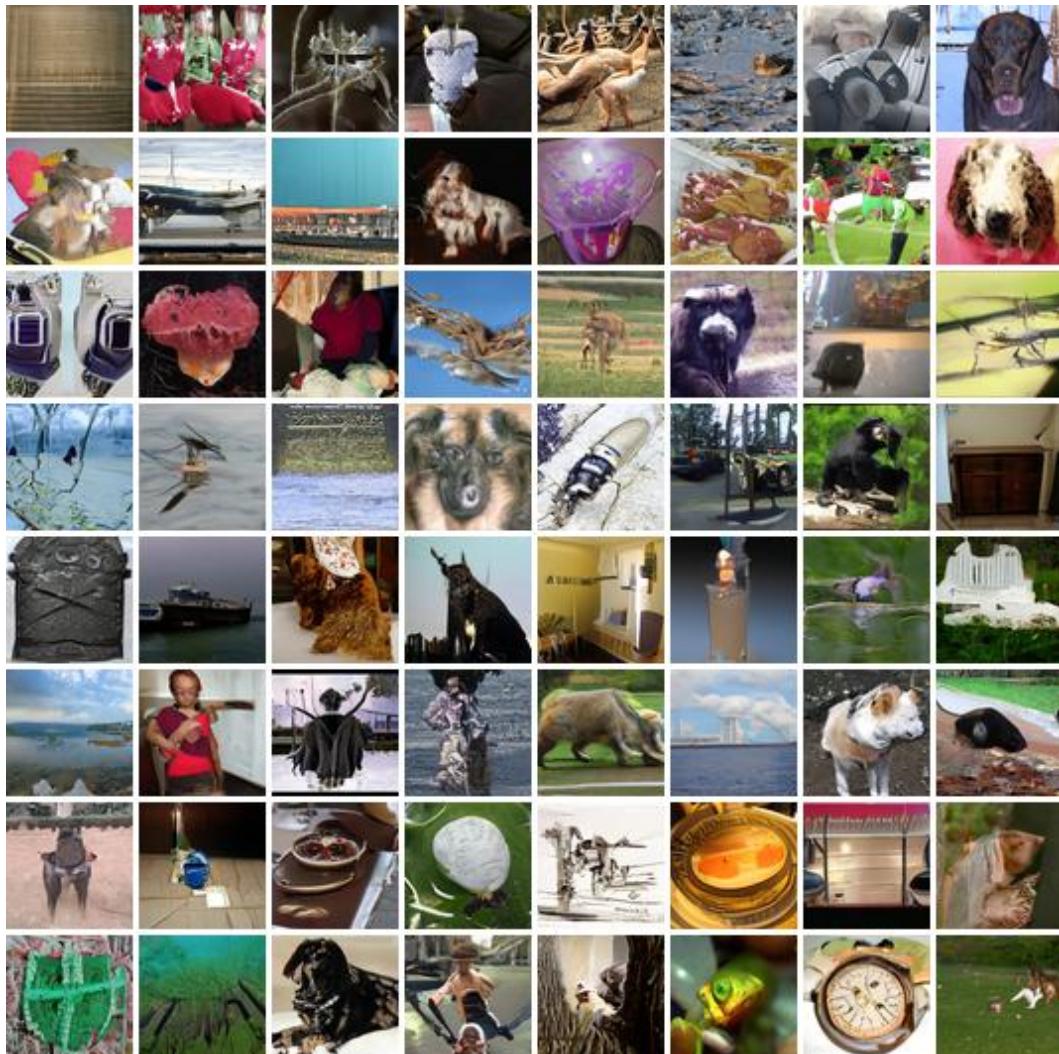


Figure 8: 2-step unconditional ImageNet 64×64 generation with our Shortcut-Distill.

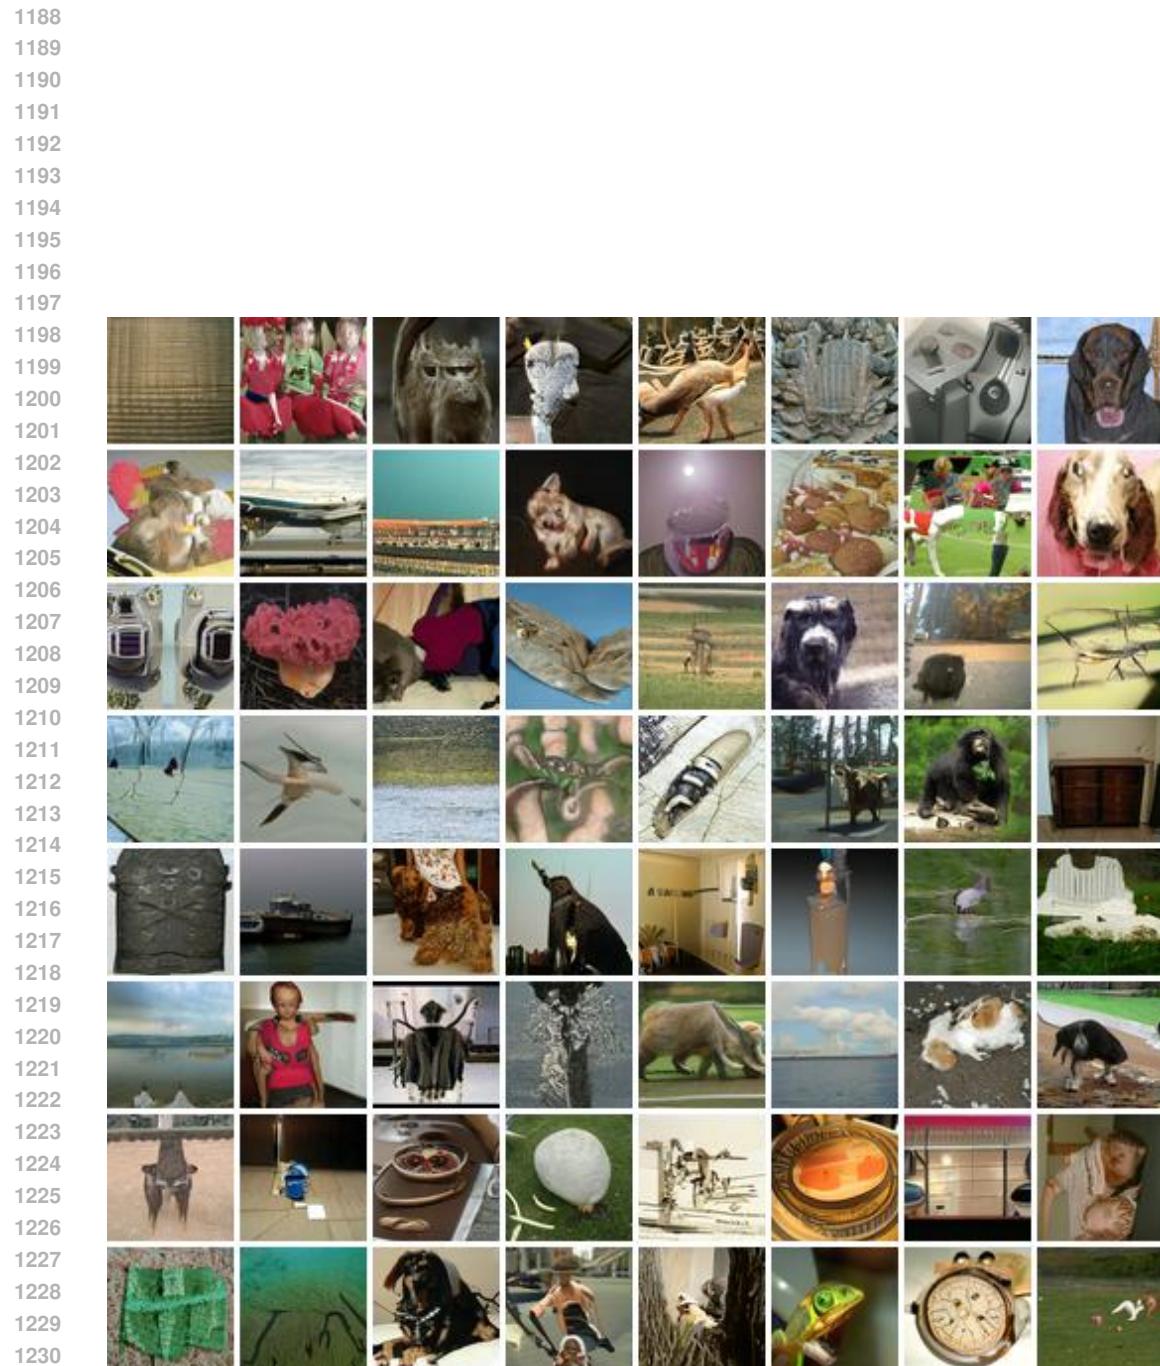
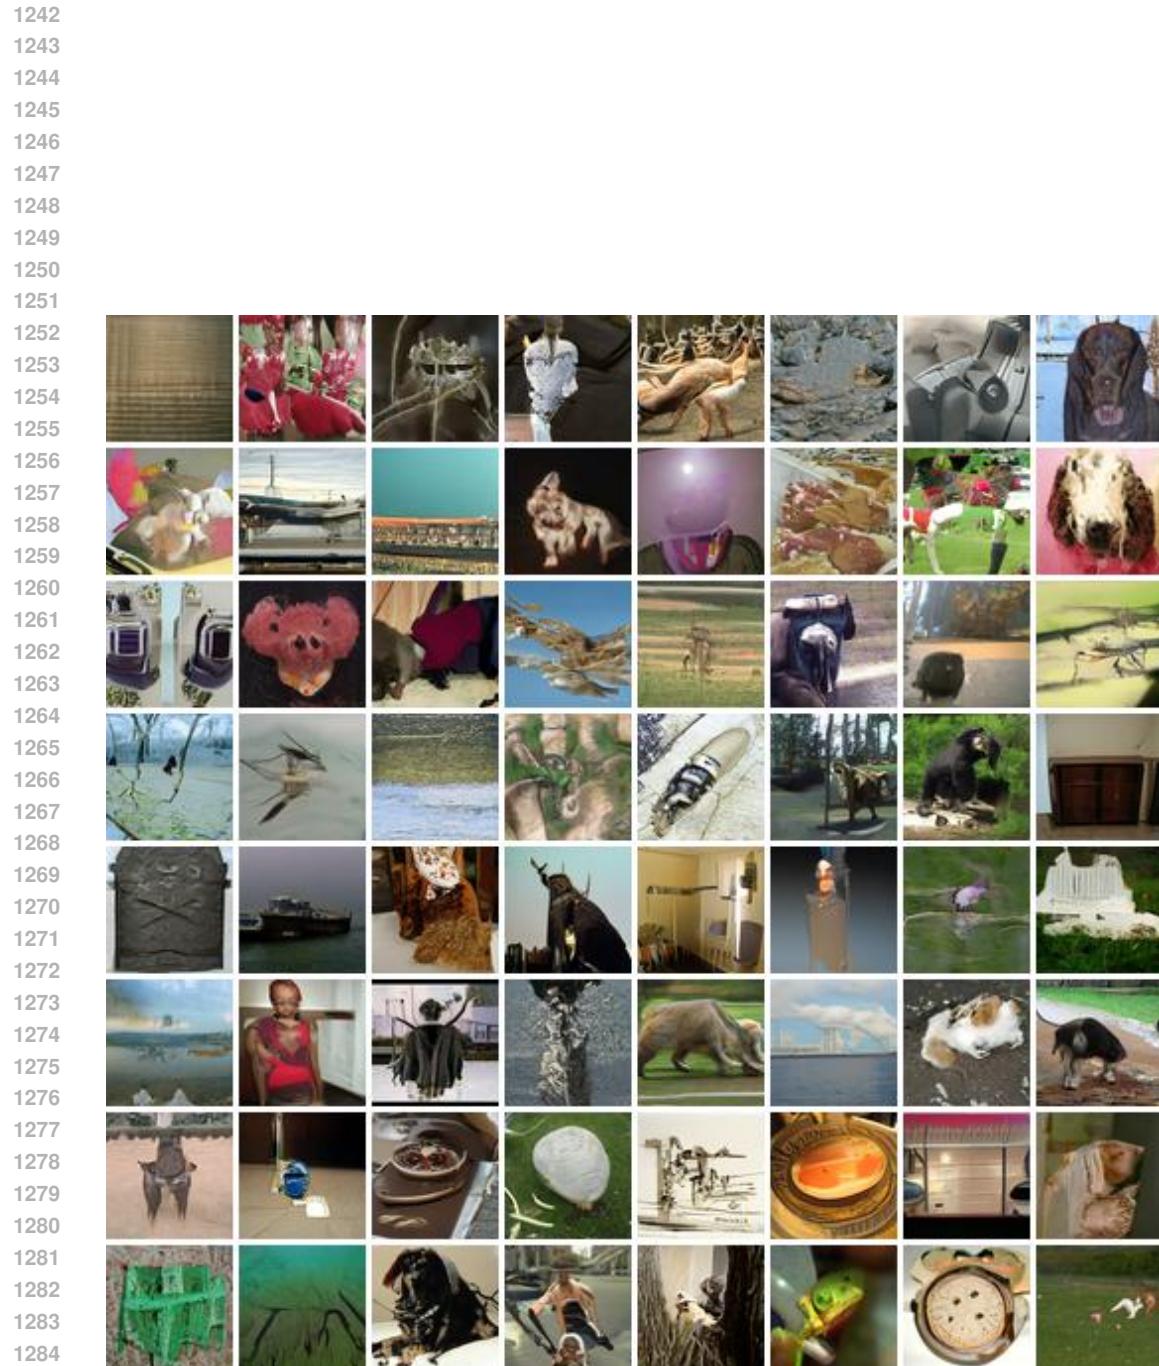


Figure 9: 8-step unconditional ImageNet 64 × 64 generation with our Shortcut-Distill-F2D2.

Figure 10: 2-step unconditional ImageNet 64×64 generation with our Shortcut-Distill-F2D2.

1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

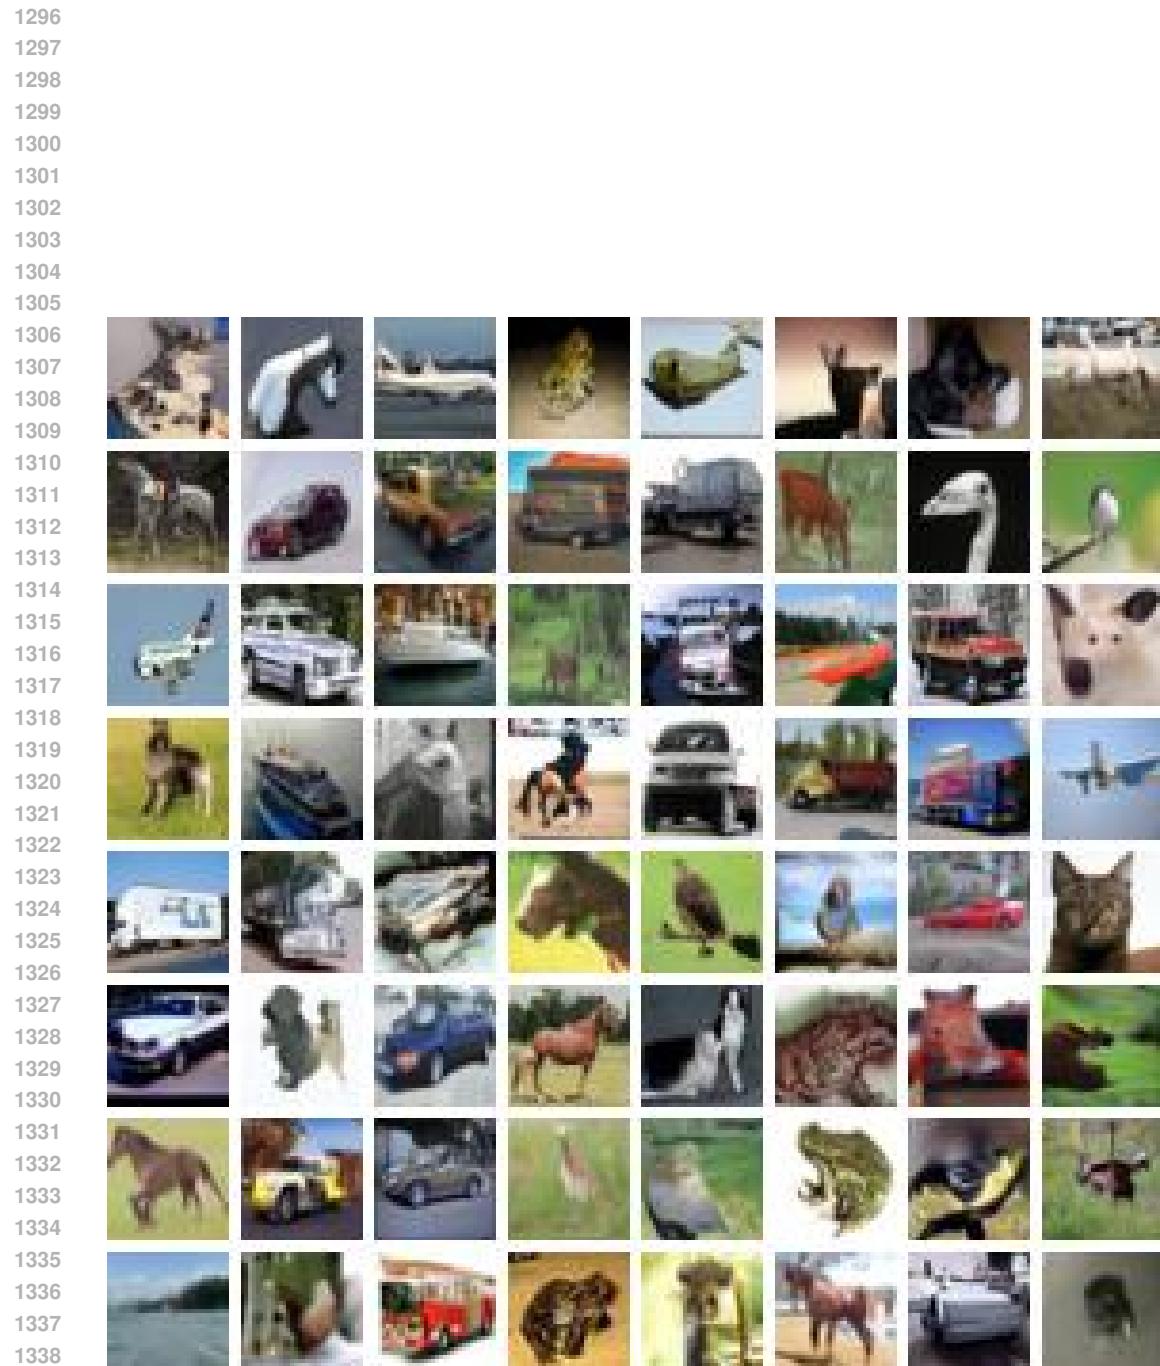


Figure 11: 8-step unconditional CIFAR-10 generation with our Shortcut-Distill.

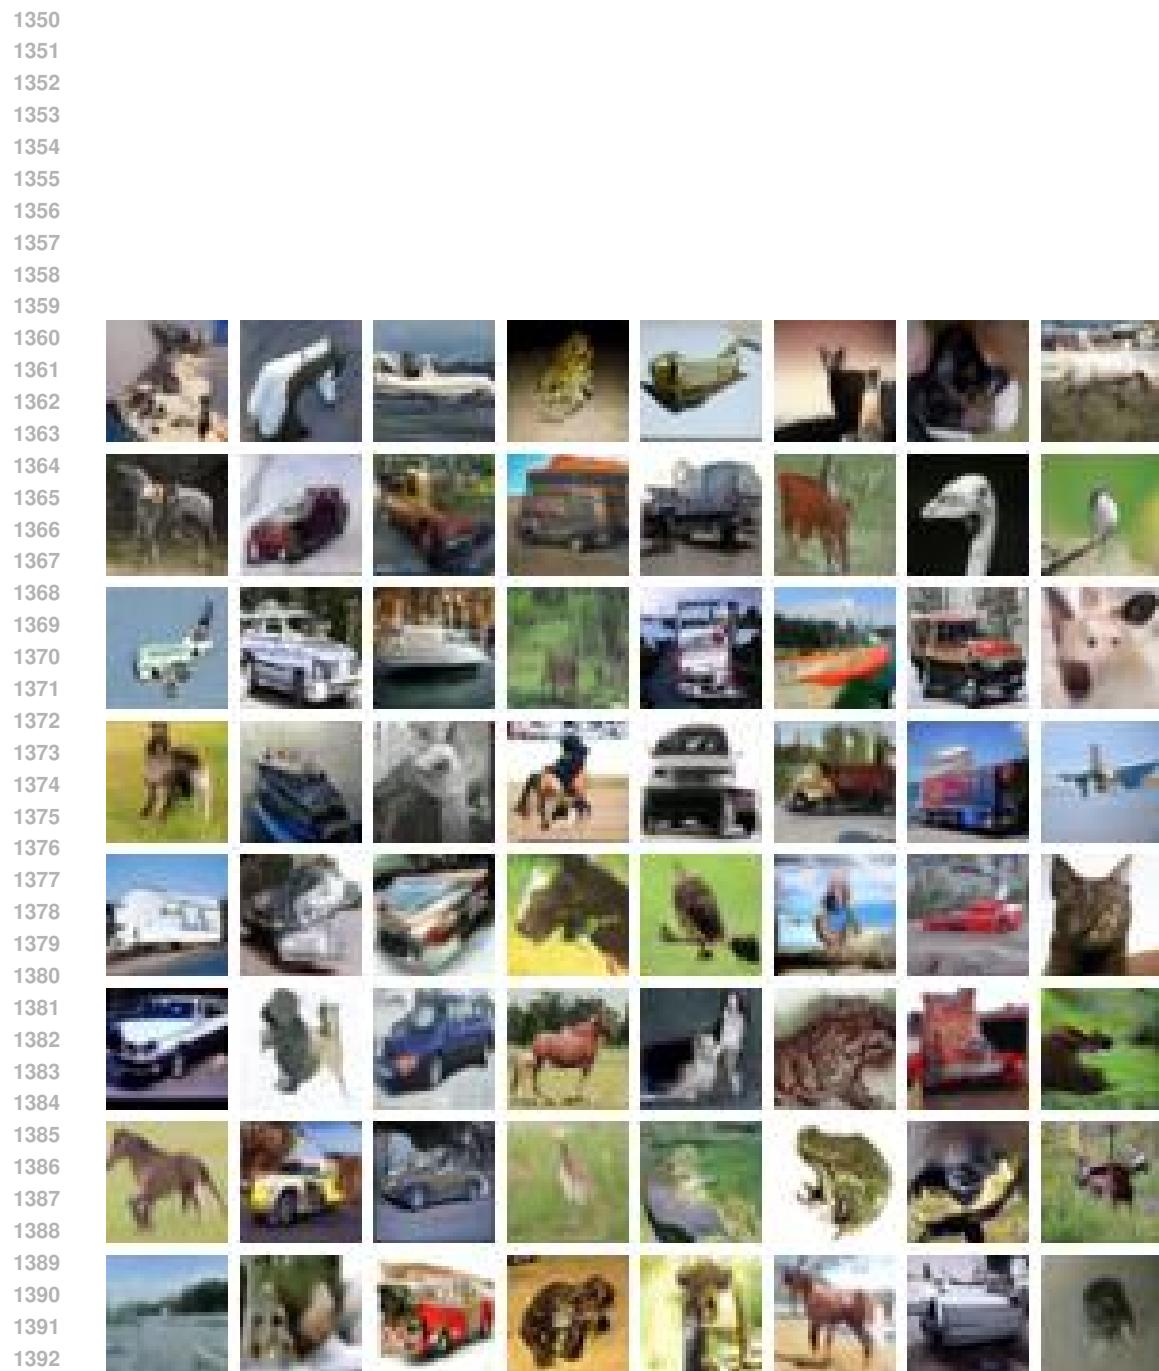


Figure 12: 2-step unconditional CIFAR-10 generation with our Shortcut-Distill.

1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

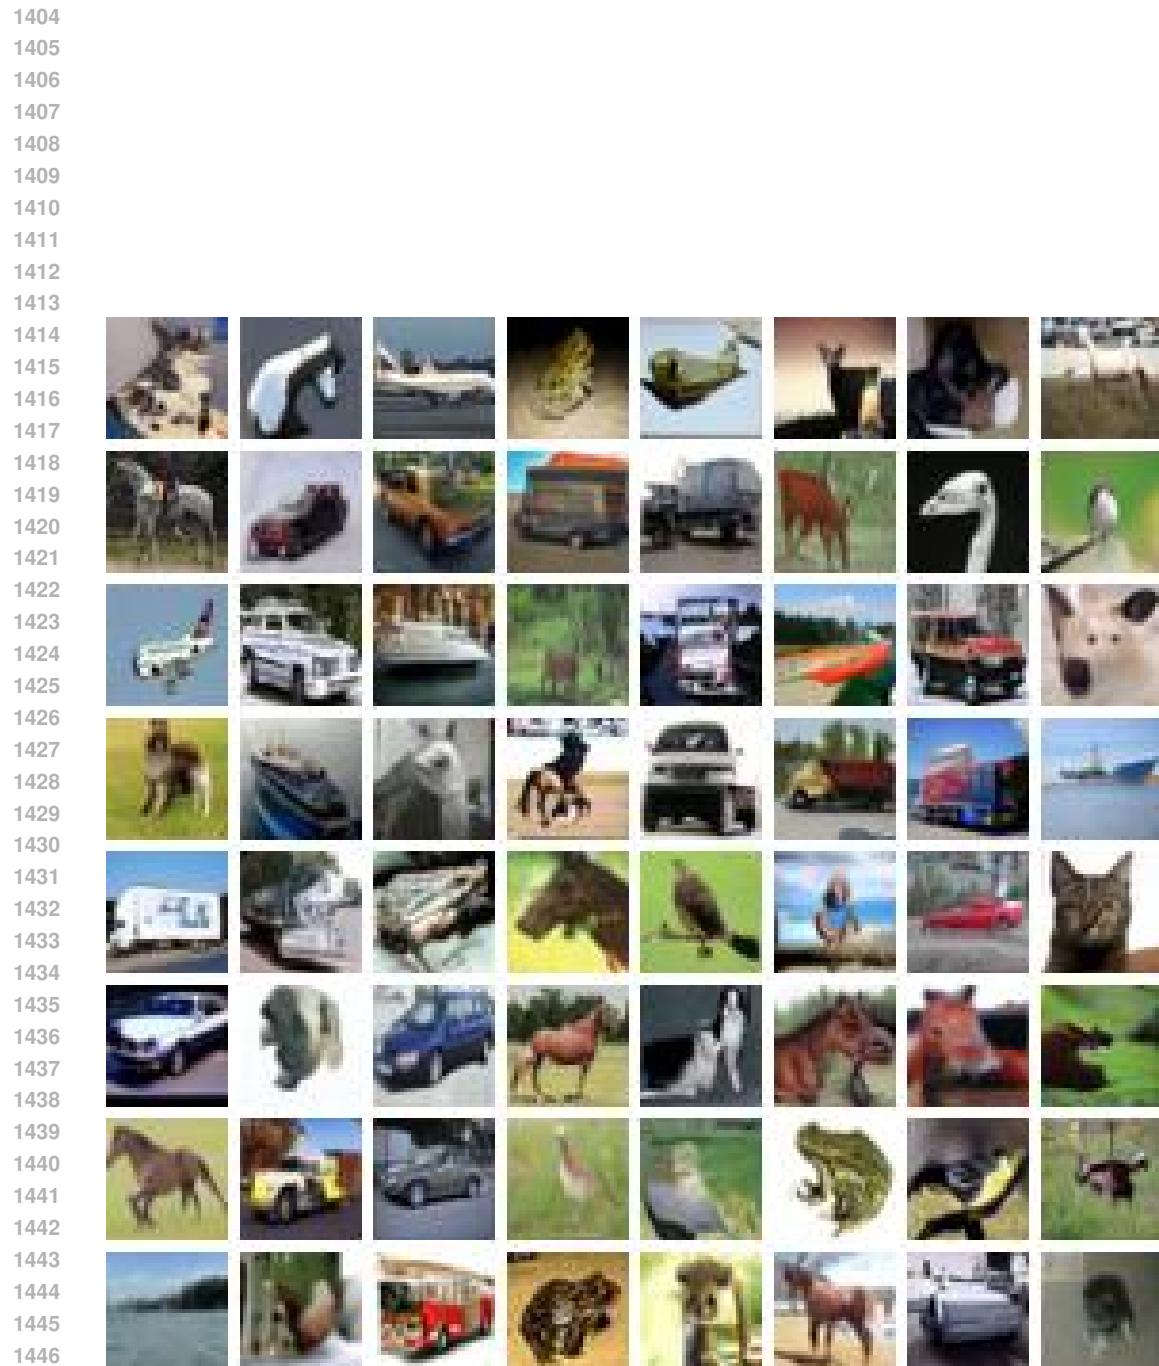


Figure 13: 8-step unconditional CIFAR-10 generation with our Shortcut-Distill-F2D2.

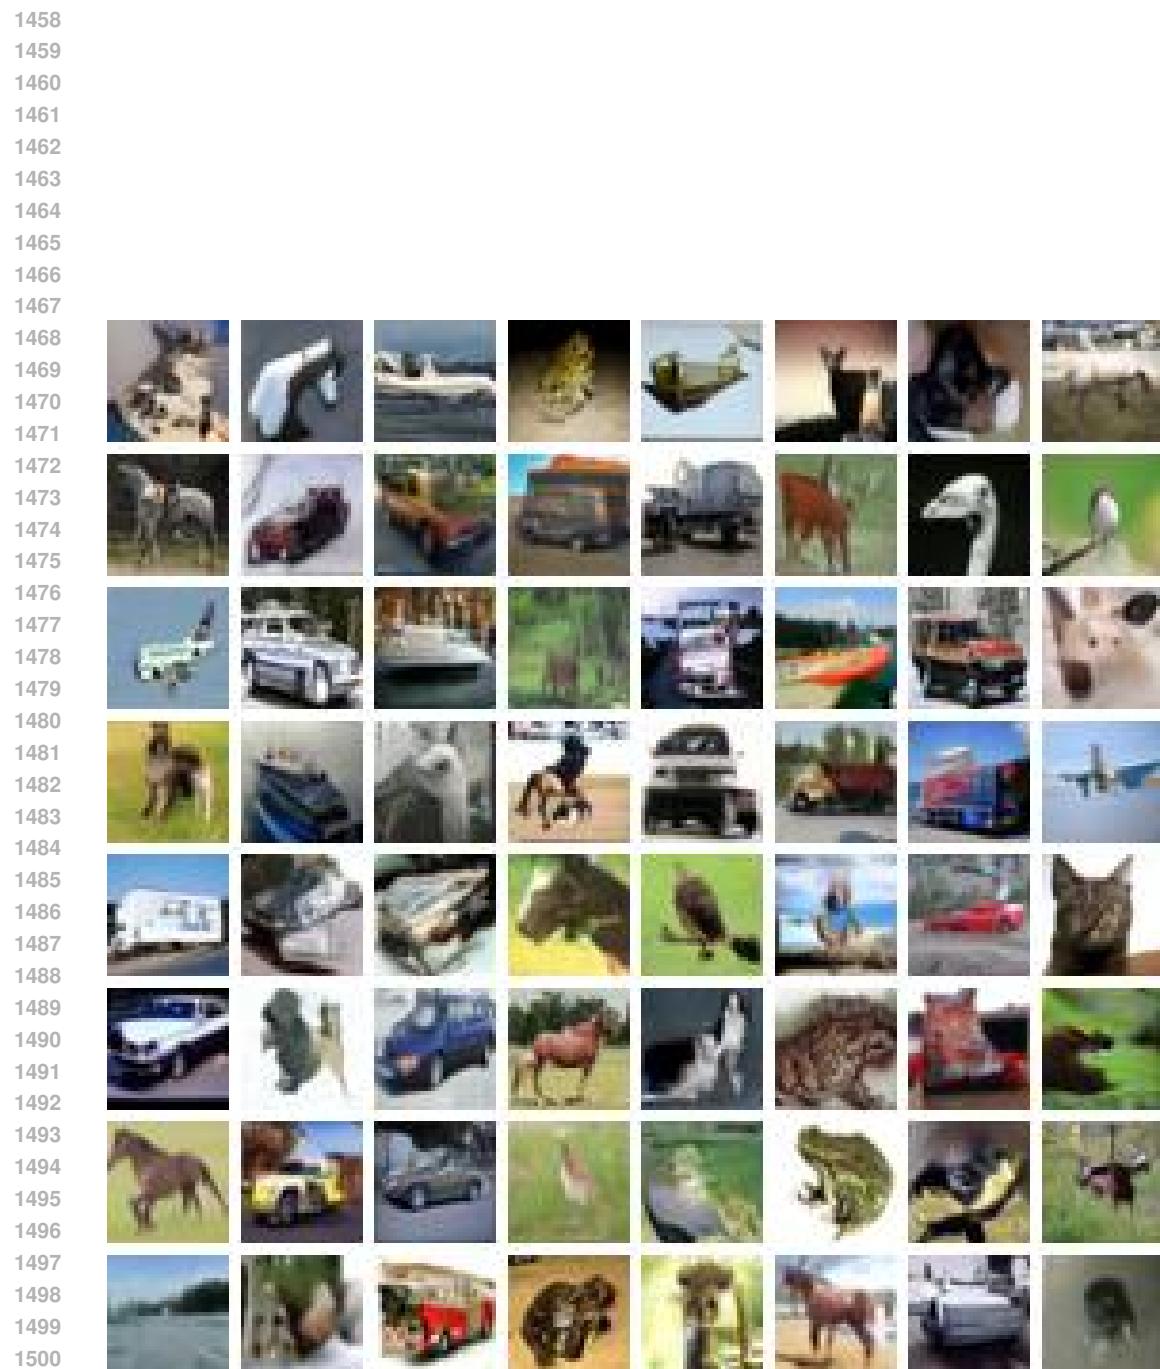


Figure 14: 2-step unconditional CIFAR-10 generation with our Shortcut-Distill-F2D2.

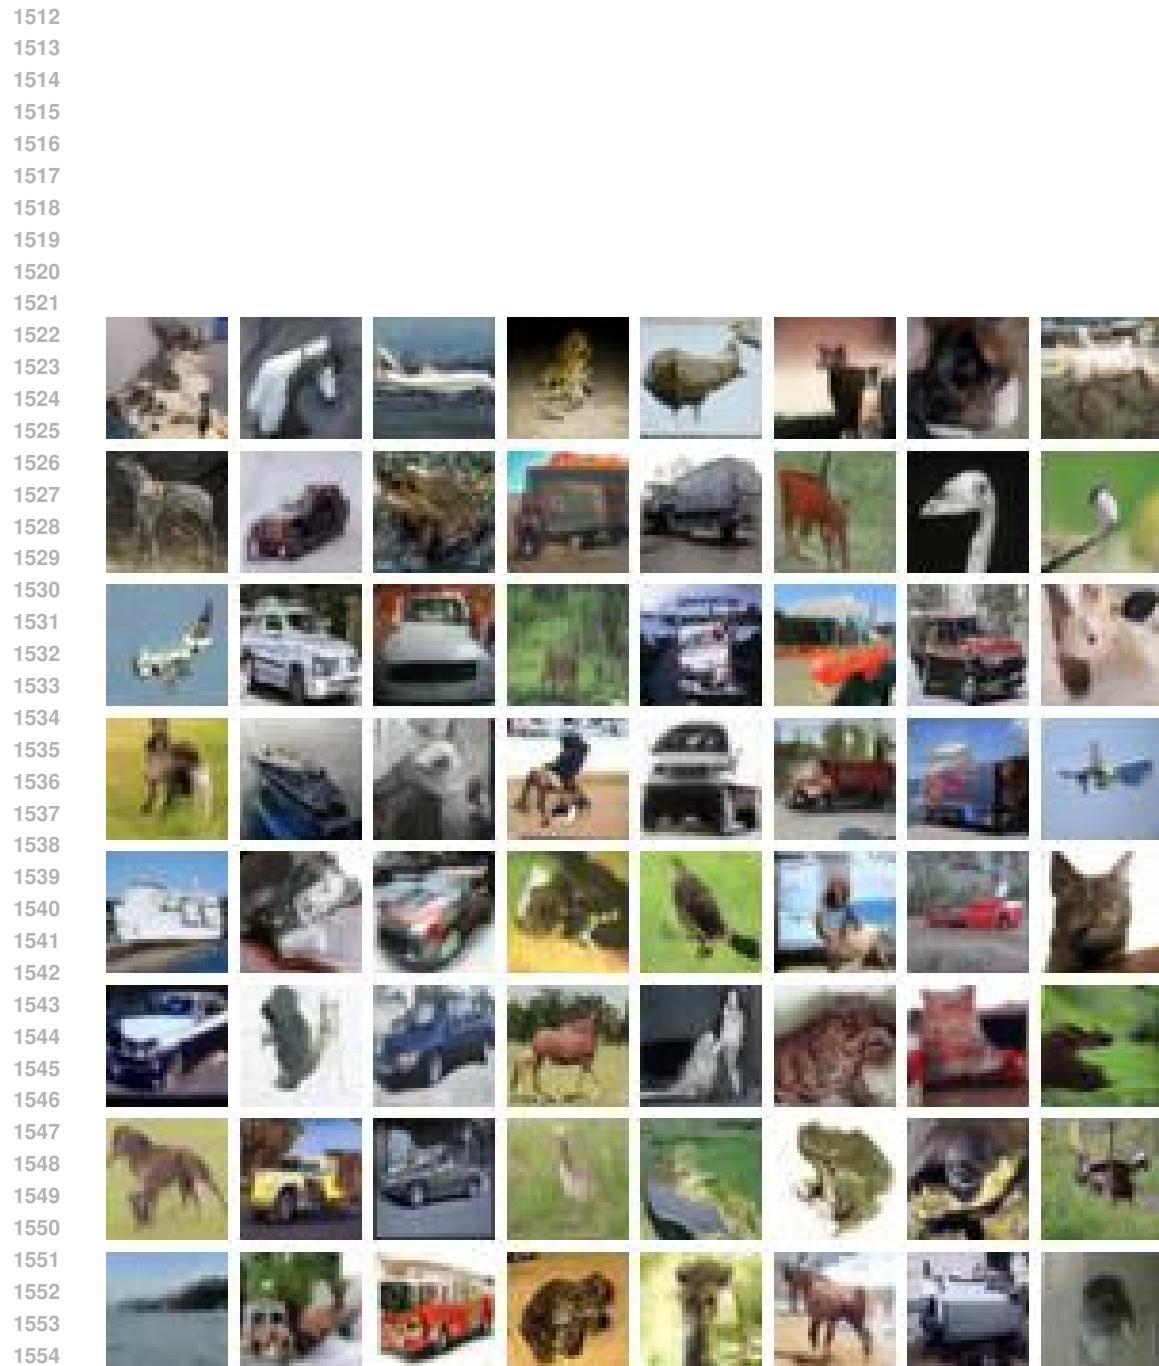


Figure 15: 8-step unconditional CIFAR-10 generation with our Shortcut-F2D2.

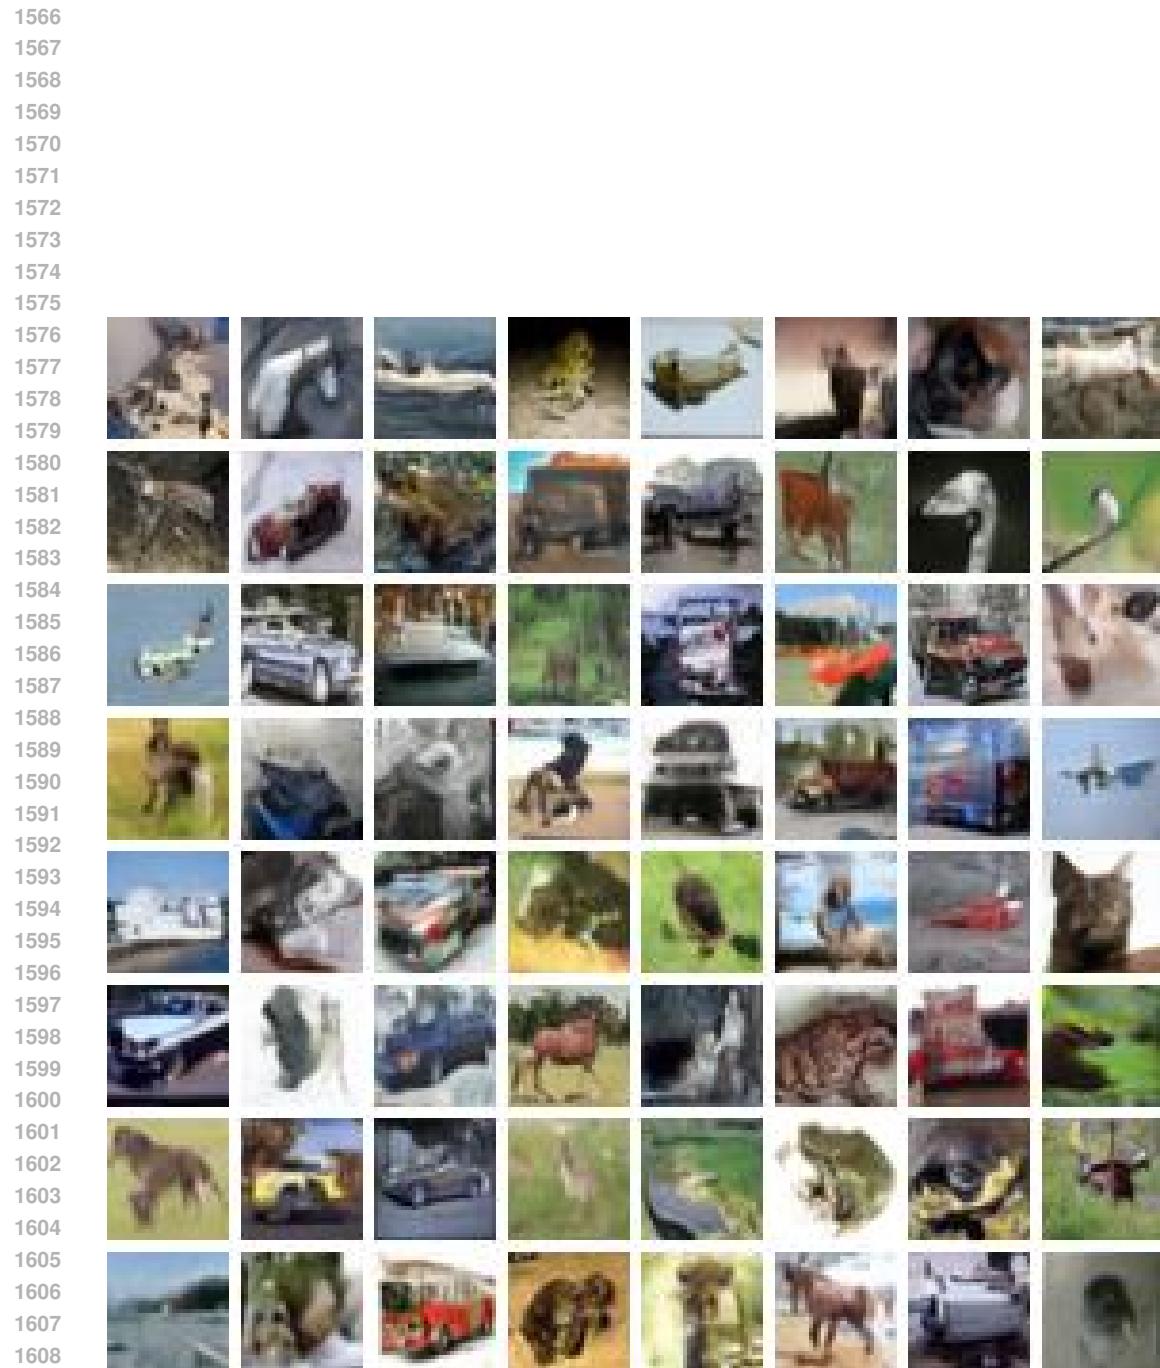


Figure 16: 2-step unconditional CIFAR-10 generation with our Shortcut-F2D2.

1610
1611
1612
1613
1614
1615
1616
1617
1618
1619

1620
1621
1622
1623
1624
1625
1626
1627
1628
1629

1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663

1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673

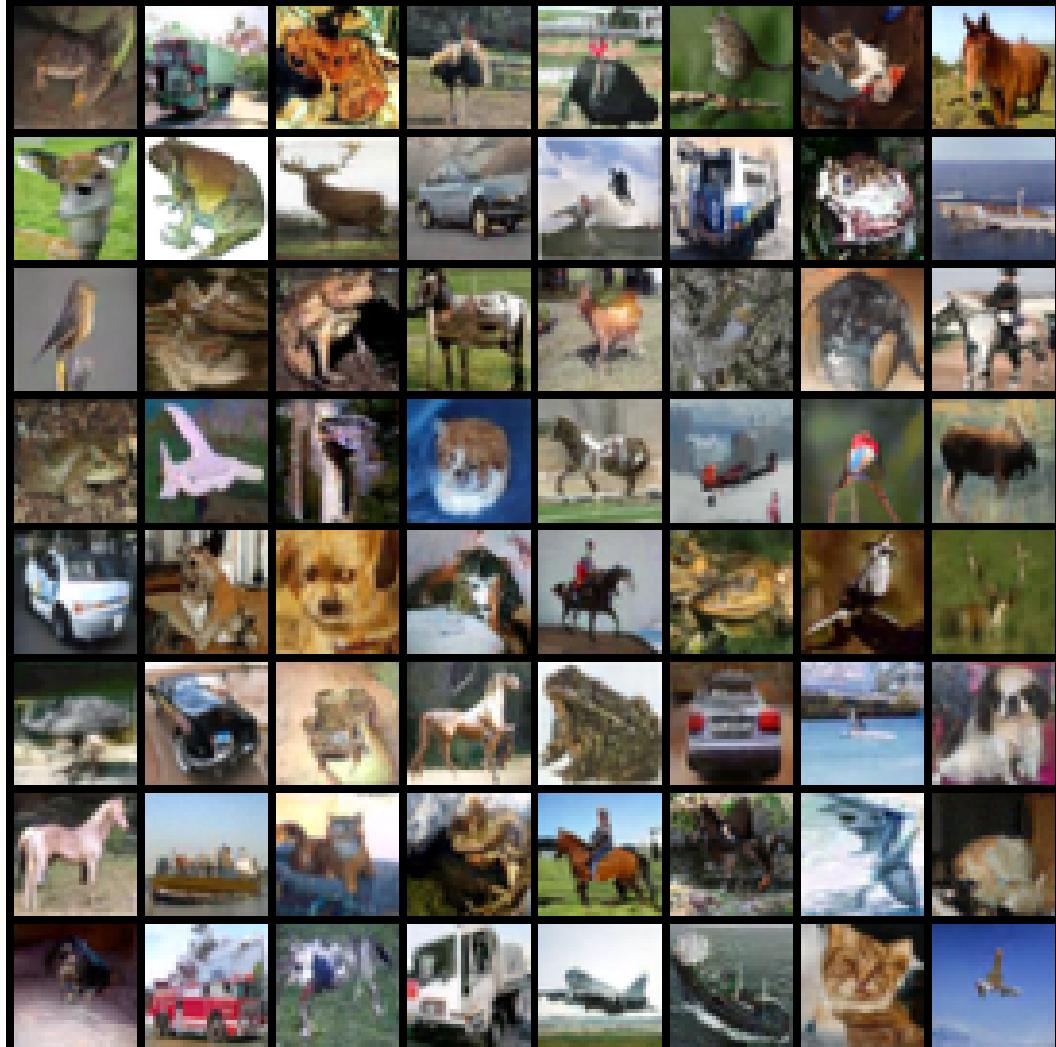


Figure 17: 2-step unconditional CIFAR-10 generation with MeanFlow-F2D2 (Ours).

1674
1675
1676
1677
1678
1679
1680
1681
1682
1683

1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717

1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

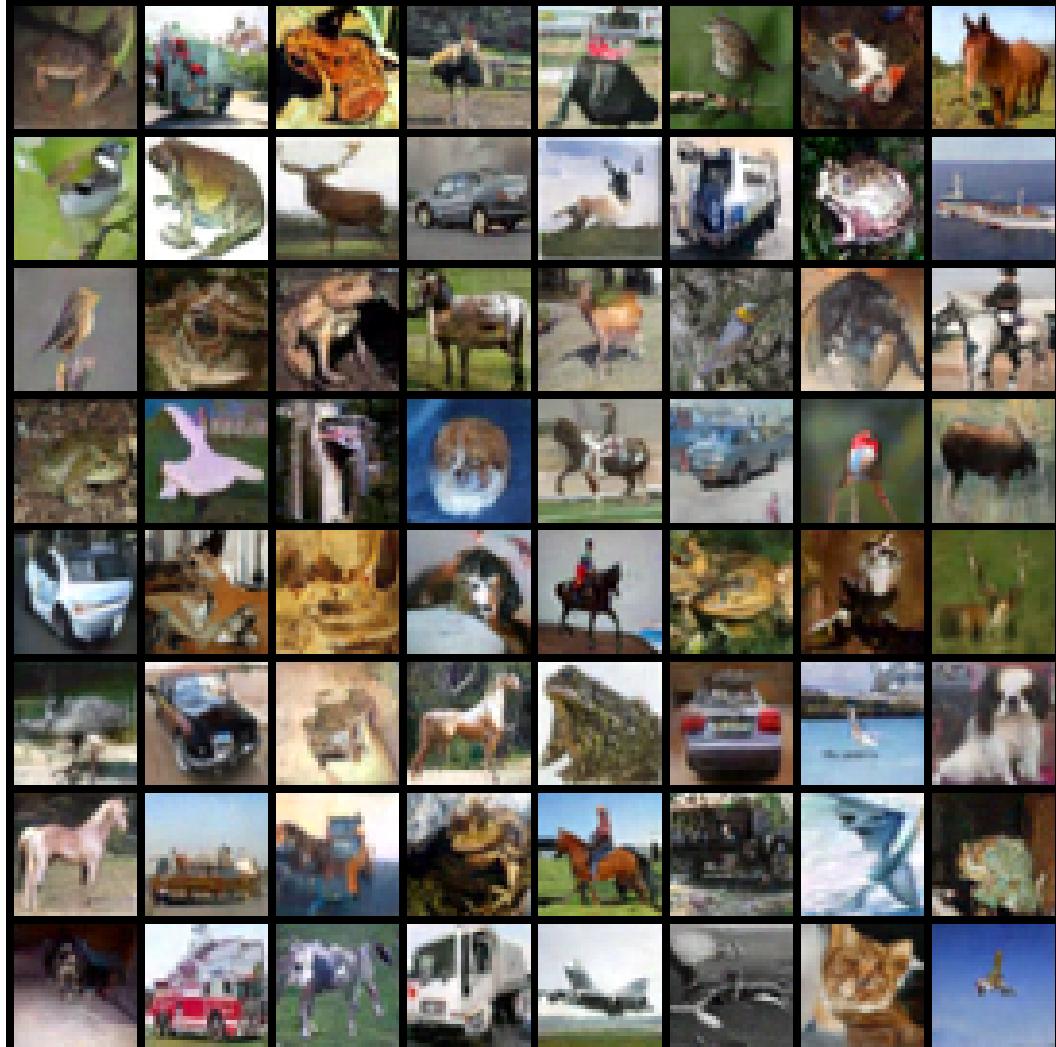


Figure 18: 1-step unconditional CIFAR-10 generation with MeanFlow-F2D2 (Ours).

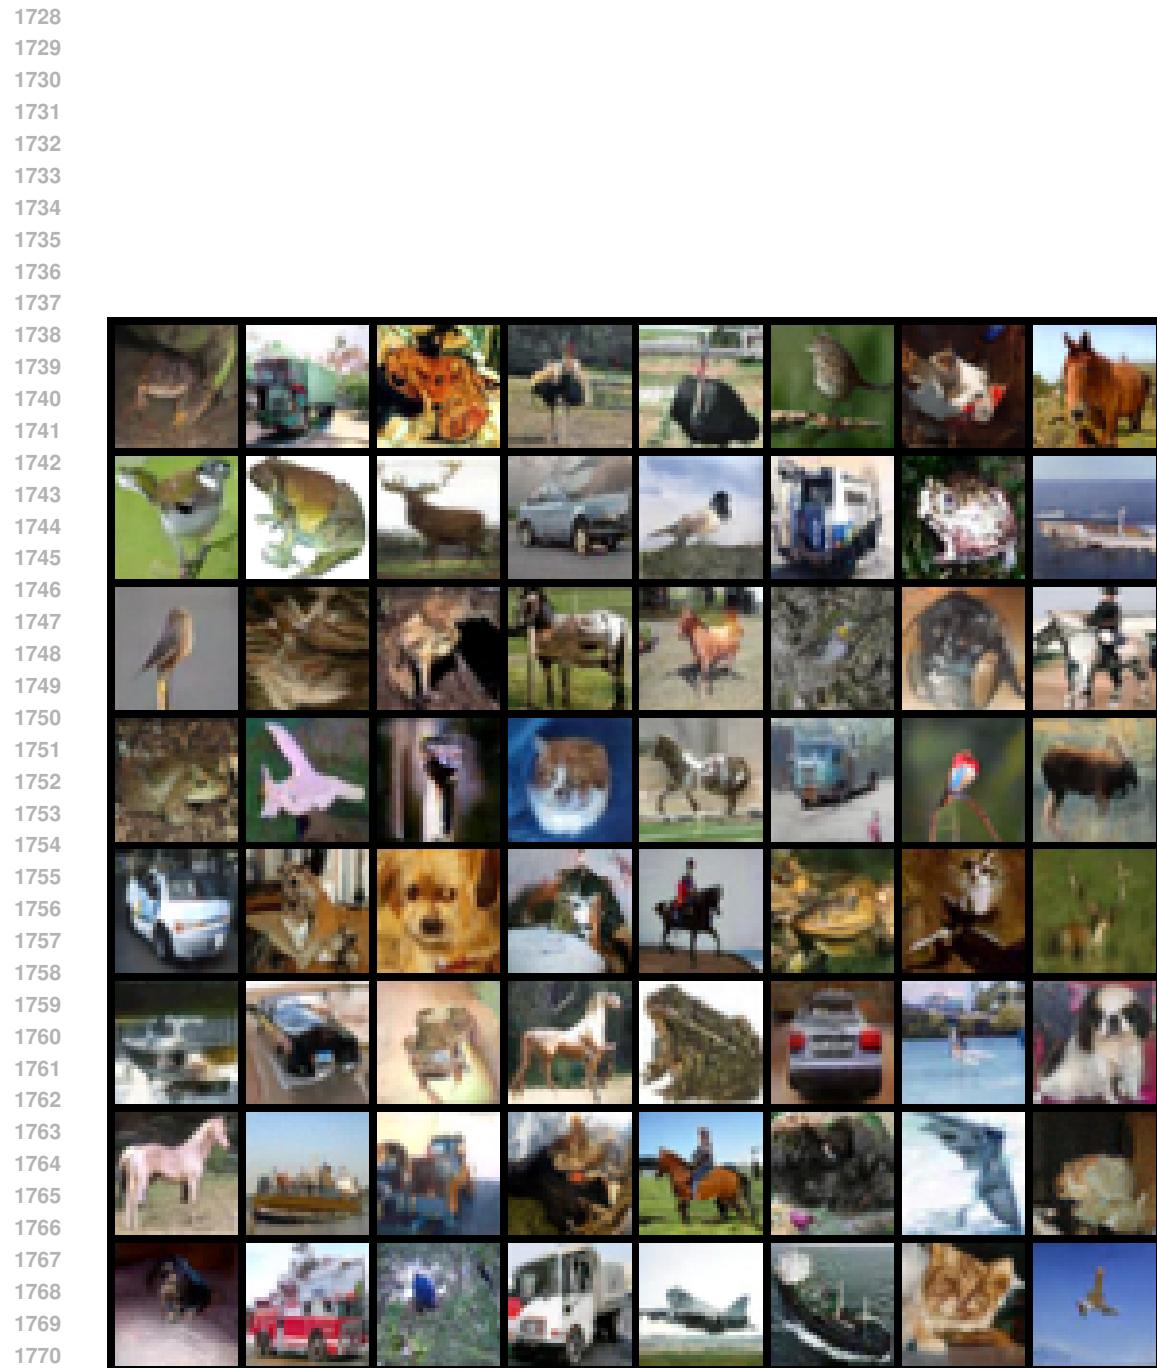


Figure 19: 2-step unconditional CIFAR-10 generation with our MeanFlow-F2D2-Self-Guidance.

1782
1783
1784
1785
1786
1787
1788
1789
1790
1791

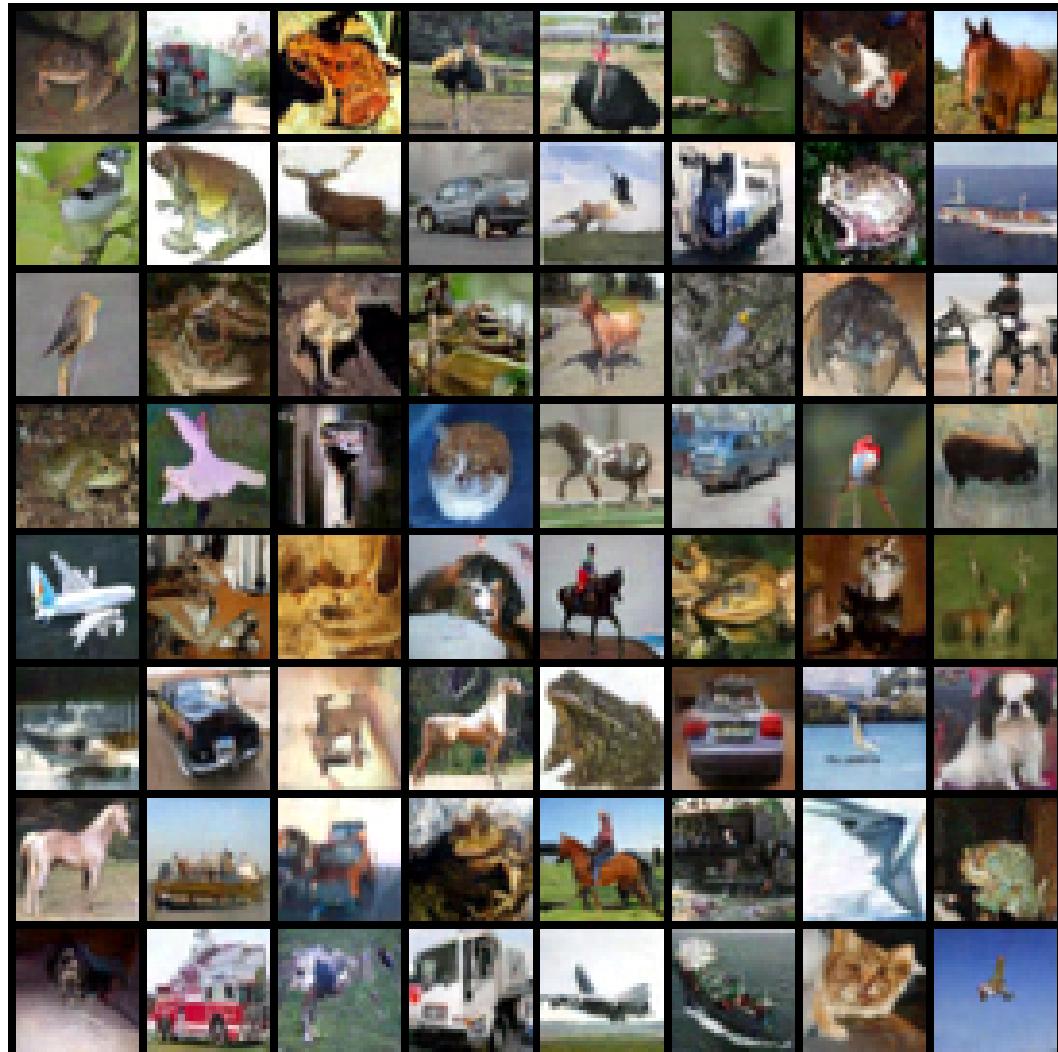


Figure 20: 1-step unconditional CIFAR-10 generation with our MeanFlow-F2D2-Self-Guidance.

1827
1828
1829
1830
1831
1832
1833
1834
1835

1836 F LIMITATIONS AND FUTURE WORKS

1837
1838 In this section, we discuss the limitations and future works of our method. First, F2D2 train-
1839 ing requires careful early stopping: once the calibrated BPD value is reached, further training can
1840 potentially lead to overfitting or degraded likelihood estimation. Future work could address this
1841 through improved network architecture and auxiliary regularization. Second, on the ImageNet-
1842 64 × 64 dataset, due to computational resources constraints, we only train with reduced model size
1843 and insufficient training iterations rather than an ideal large-scale configuration. We expect training
1844 with longer duration and larger models can further improve the performance. We also restrict our-
1845 selves to unconditional generation, which is substantially more challenging than conditional setups.
1846 Finally, the performance of our method is sensitive to divergence target scaling and its practical
1847 effectiveness in other architectures or modalities remains to be fully validated in future works.
1848

1849 G THE USE OF LARGE LANGUAGE MODELS (LLMs)

1850
1851 We use Large Language Models (LLMs) to refine and polish the manuscript. LLMs also support
1852 our code debugging, but they are not involved in developing the algorithms or conducting the exper-
1853 iments. The authors take full responsibility for the content of manuscript.
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