QDSB: QUANTIZED DIFFUSION SCHRÖDINGER BRIDGES

Anonymous authors
Paper under double-blind review

ABSTRACT

Learning generative models in settings where the source and target distributions are only specified through unpaired samples is gaining in importance. Here, one frequently-used model are Schrödinger bridges (SB), which represent the most likely evolution between both endpoint distributions. To accelerate training, simulation-free SBs avoid the path simulation of the original SB models. However, learning simulation-free SBs requires paired data; a coupling of the source and target samples is obtained as the solution of the entropic optimal transport (OT) problem. As obtaining the optimal global coupling is infeasible in many practical cases, the entropic OT problem is iteratively solved on minibatches instead. Still, the repeated cost remains substantial and the locality can distort the global transport geometry. We propose *quantized diffusion Schrödinger bridges* (QDSB), which compute the endpoint coupling on anchor-quantized endpoint distributions and lift the resulting plan back to original data points through cell-wise sampling. We show that the regularized optimal coupling is stable w.r.t. anchor quantization, with an error controlled by the quality of the anchor approximation. In real-world experiments, QDSB matches the sample quality of existing baselines, requiring substantially less time. Code and data are available at `github.com/anon1248/qdsb`.

1 INTRODUCTION

Generative modeling is increasingly used to learn the evolution of one data distribution to another, and has found applications in image (Rombach et al., 2022; Zhang et al., 2025), audio (Kong et al., 2021; Lemercier et al., 2024), and video synthesis (Blattmann et al., 2023; Wang et al., 2025). Well-known approaches, such as denoising diffusion probabilistic models (Ho et al., 2020), score-based stochastic differential equations (Song et al., 2021), flow matching (Lipman et al., 2023), rectified flows (Liu et al., 2023), and consistency models (Song et al., 2023), train neural networks that represent a map evolving a simple *reference* distribution to an observed data distribution. In this setting, the reference distribution is typically available in closed form and the learned map allows generating new data.

Parallel to this setting, a core problem in generative modeling is learning a mapping from one observed *data* distribution to another data distribution, naturally arising in unpaired image translation (Zhu et al., 2017; Liu et al., 2017) and in modeling cell-state dynamics (Tong et al., 2020). Here, one observes samples from two endpoint distributions that are either paired or unpaired; we focus on the latter. More precisely, we study the setting where one observes independent samples from a source and target distribution (both unknown), and the goal is to find the most likely evolution between both distributions, given a reference process. Schrödinger bridges (De Bortoli et al., 2021), which connect diffusion-based models (Song et al., 2021) with optimal transport (OT; Villani 2009; 2021), provide a principled framework to tackle this challenging problem.

However, estimating the Schrödinger bridge requires paired samples, that is, knowledge of which source sample should be mapped to which target sample. In the unpaired case, this correspondence must be constructed from the data observed; the estimated pairing—or coupling—has key implications on the quality of the learned map. While the original approach to learning the Schrödinger bridge involves simulating the data of intermediate evolution steps, which is computationally costly, recent simulation-free methods make Schrödinger bridge training substantially more practical. In particular, Tong et al. (2024) showed that, by using the known decomposition of the Schrödinger bridge (with

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Brownian motion reference) into a weighted combination of Brownian bridges (Föllmer, 1988), a so-called simulation-free estimation is possible; their SF2M approach replaces costly simulation by score and flow matching against these Brownian bridges.

Yet, this simplification still requires paired samples. To obtain a paired sample given an unpaired one, one solves the entropic OT problem (Cuturi, 2013) between the full source and target samples. In practice, this is infeasible due to the size of the samples; the pairing is commonly estimated on minibatches only. However, the minibatch heuristic may fail to capture the global transport geometry and can become overly diffuse under regularization (Tong et al., 2024, Appendix A.1). Further, the minibatch heuristic addresses the runtime problem only partially as the entropic OT problem is solved in every minibatch, rather than once globally. As a result, the pairing step remains a computational and statistical bottleneck in simulation-free bridge training.

We address this bottleneck by proposing an anchor-based quantization for learning simulation-free Schrödinger bridges. Our approach first quantizes each endpoint distribution onto a finite set of well-chosen anchors, then computes the entropic OT between the induced discrete measures, and finally samples original endpoints from matched anchor cells for bridge training. Importantly, the simulation-free training objective itself remains unchanged, since the model is still supervised by Brownian bridges between raw data points rather than between anchor locations. This yields a cheaper pairing procedure while allowing explicit control of the quality of the resulting OT approximation.

Our **contributions** are as follows.

- *Quantized couplings for simulation-free bridges.* We introduce an anchor-based coupling method for simulation-free Schrödinger bridge training from unpaired samples. The method computes entropic OT on quantized endpoint distributions and lifts the resulting plan back to original endpoint pairs through cell-wise sampling.
- *Stability analysis.* We bound the influence of the perturbation of the solution of the entropic OT problem induced by the anchor quantization. The bound in Theorem 3.3 and Lemma 3.4 provide a geometric principle that guides anchor selection in practice.
- *Empirical evaluation.* We study the proposed coupling on toy and real-world tasks. The experiments indicate that our method preserves the quality of simulation-free bridge training while improving the time-performance trade-off of the coupling stage. All code and data is openly available at `github.com/anon1248/qdsb`, including code to reproduce all tables and figures.

Outline. Section 2 discusses related work, while Section 3 reviews simulation-free Schrödinger bridges and presents our quantized coupling method. Section 4 reports experiments, and Section 5 concludes. All proofs are deferred to the appendices.

2 Related work

Instead of mapping one observed distribution to a simple reference distribution, our work lies in the data-to-data setting, where the goal is to connect two observed endpoint distributions. While non-SB-based methods exist in this setting (Zhu et al., 2017; Liu et al., 2017; Tong et al., 2020), we focus our discussion on SB-based approaches. Recent SB solvers include diffusion SB matching (Shi et al., 2023), generalized SB matching (Liu et al., 2024), and light and optimal SB matching (Gushchin et al., 2024). A closely related extension is LightSBB-M (Alouadi et al., 2026), which considers the Schrödinger-Bass bridge. Among simulation-free approaches, the direct starting point of our work is SF2M (Tong et al., 2024). As detailed above, its remaining bottleneck is the computational time of solving an entropic OT problem in each training minibatch, which we address in this work.

Classical OT (Monge, 1781; Kantorovich, 1942) provides the coupling foundation of both entropic OT and SBs, and entropic regularization makes OT practical through Sinkhorn iterations (Cuturi, 2013). However, large-scale OT remains expensive, which motivates approximations based on minibatches (Fatras et al., 2021), low-rank couplings (Scetbon et al., 2021), or anchor points (Lin et al., 2021). Recent work improves minibatch OT itself, for example, through unbalanced minibatch OT (Fatras et al., 2021), mini-batch partial OT (Nguyen et al., 2022b), or hierarchical couplings between minibatches (Nguyen et al., 2022a). Our method is orthogonal to these as they improve the minibatch OT surrogate, while leaving the surrounding training algorithm unchanged and thus

still requiring the repeated solving of the entropic OT problem. An alternative to reduce transport complexity is using a finite set of representatives (Lin et al., 2021), but then the inherent limitations of simulation-free training remain unchanged. Our approach introduces an anchor-quantized coupling for the distinct problem of constructing endpoint couplings within simulation-free SB training rather than solving OT in general, and therefore has the benefit of not requiring the repeated solving of OT problems in each minibatch.

3 Simulation-free Schrödinger bridges with quantized couplings

This section introduces QDSB, our quantized coupling procedure for simulation-free SB training from unpaired samples. To set the stage, we recall the simulation-free objective of Tong et al. (2024).

3.1 Preliminaries

Notation. Let $d \in \mathbb{N}$, equip $\mathbb{R}^d$ with a fixed norm $\|\cdot\|$, and let $c(x,y) = \|x-y\|$ denote the associated distance function. Given $a \in [1,\infty)$, let $\mathcal{P}_a(\mathbb{R}^d)$ denote the set of Borel probability measures μ on $\mathbb{R}^d$ with finite a-th moment, that is, $\int c(x,y)^a \,\mathrm{d}\mu(x) < \infty$ for $y \in \mathbb{R}^d$. For $a = \infty$, $\mathcal{P}_\infty(\mathbb{R}^d)$ is the set of measures with bounded support. For $\mu, \nu \in \mathcal{P}_a(\mathbb{R}^d)$, let $\Pi(\mu,\nu)$ denote the set of couplings with marginals μ and ν, and $W_a(\mu,\nu)$ denote their associated a-Wasserstein distance: for $a \in [1,\infty)$, $W_a(\mu,\nu)^a = \inf_{\pi\in\Pi(\mu,\nu)} \int c(x,y)^a \,\mathrm{d}\pi(x,y)$; for $a=\infty$, $W_\infty(\mu,\nu) = \inf_{\pi\in\Pi(\mu,\nu)} \operatorname{ess\,sup}_{(x,y)\sim\pi} c(x,y)$. For a Borel measurable map $T : \mathbb{R}^d \to \mathbb{R}^d$ and probability measure μ, we write $T_\#\mu$ for the pushforward of μ under T.

Schrödinger bridge (SB). Fix $d \in \mathbb{N}$, $a \in [1,\infty]$, and let $q_0, q_1 \in \mathcal{P}_a(\mathbb{R}^d)$. The goal of the SB problem is to find the most likely evolution between q_0 and q_1, that is, it is the solution $P^* = \arg\min_P \mathrm{KL}(P\|Q)$ subject to $p_0 = q_0$ and $p_1 = q_1$, where P is a probability measure on continuous $[0,1] \to \mathbb{R}^d$ functions that is absolutely continuous w.r.t. a given reference measure Q, and law p with p_t denoting the time-t marginal of P with $t \in [0,1]$.

Simulation-free training (SF2M). Tong et al. (2024) consider the Brownian reference measure Q given by the law of the scaled multivariate Brownian motion σB, where B is a standard Brownian motion and $\sigma > 0$. In that case, P^* admits a representation as a mixture of Brownian bridges

$$P^*((x_t)_{t\in[0,1]}) = \int Q((x_t)_{t\in[0,1]} \mid x_0, x_1) \,\mathrm{d}\pi^{2\sigma^2}_{q_0,q_1}(x_0,x_1), \tag{1}$$

where $Q(\cdot \mid x_0, x_1)$ is the Brownian bridge connecting x_0 and x_1. It is weighted by the entropic OT coupling $\pi^{2\sigma^2}_{q_0,q_1}$ between q_0 and q_1, that is, $\pi^{\tau}_{q_0,q_1}$ is the minimizer of

$$\mathcal{S}_\tau(q_0,q_1) = \inf_{\pi\in\Pi(q_0,q_1)} \left[\int c(x_0,x_1) \,\mathrm{d}\pi(x_0,x_1) + \tau \,\mathrm{KL}(\pi\|q_0\otimes q_1) \right], \tag{2}$$

where $\tau = 2\sigma^2$ in (1). In this case, it is known (Tong et al., 2024) that to approximate P^* it suffices to learn the so-called drift $u^\circ_t(x \mid x_0,x_1) = \frac{1-2t}{t(1-t)}(x - (tx_1 + (1-t)x_0)) + (x_1 - x_0)$ and score $\nabla_x \log p_t(x \mid x_0,x_1) = \frac{tx_1+(1-t)x_0-x}{\sigma^2 t(1-t)}$ of the underlying stochastic differential equation. Indeed, Tong et al. (2024) then fit neural networks approximating each, that is, they minimize

$$\mathcal{L}_\theta = \mathbb{E}_{t\sim\mathcal{U}(0,1), z\sim q(z), x\sim p_t(x|z)}\big[\|v_\theta(t,x) - u^\circ_t(x \mid z)\|^2 + \lambda(t)^2 \|s_\theta(t,x) - \nabla_x \log p_t(x \mid z)\|^2\big], \tag{3}$$

with $z = (x_0,x_1)$, $v_\theta(t,x)$ a neural network approximating the drift term, $s_\theta(t,x)$ a neural network approximating the score term, θ the network parameters, and $\lambda(t) > 0$ a weighting function.

Since only unpaired i.i.d. samples from q_0 and q_1 are observed, the endpoint coupling $q(x_0,x_1)$ is unknown and must be constructed. Tong et al. (2024) achieve this by solving a separate OT problem on each training minibatch and use the resulting matches as samples $z \sim q(x_0,x_1)$. However, minibatch OT is only a surrogate for the desired global coupling: subsampled estimators need not preserve the exact marginals, and when combined with entropic regularization they can produce overly dense, nearly uniform transport plans that obscure the data geometry (Tong et al., 2024, Appendix A.1). The quality of the learned bridge critically depends on how $q(x_0,x_1)$ is constructed. This observation motivates our improved coupling procedure proposed in the following.

3.2 QDSB: QUANTIZED DIFFUSION SCHRÖDINGER BRIDGES

We replace the minibatch OT heuristic in Tong et al. (2024) with a coupling computed on a quantized representation of the endpoint distributions. For each endpoint distribution q_i, $i \in \{0, 1\}$, we select a finite anchor set $A_i \subseteq \mathbb{R}^d$ and define a map $T_i : \mathbb{R}^d \to A_i$ that assigns each point to its closest anchor. This induces the discrete pushforward measure $\tilde{q}_i := (T_i)_{\#} q_i$. We then solve the entropic OT problem once between $\tilde{q}_0$ and $\tilde{q}_1$ and use the resulting anchor-level OT plan to sample endpoint pairs for the loss in (3).

While we explain this procedure more thoroughly in Section 3.2.2, Figure 1 briefly illustrates this construction. The large blue and green points denote the source and target anchors, and the gray links denote the OT plan between them. During training, endpoint pairs are sampled from matched anchor cells. The anchors and transport links are weighted to match the empirical marginals of the original endpoint distributions. The orange curves indicate example trajectories.

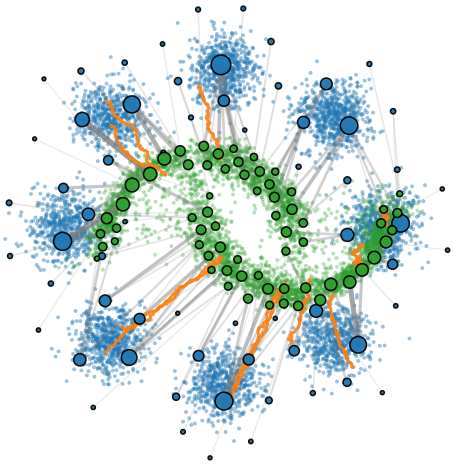

Figure 1: Illustration of QDSB on 2D toy datasets, see Section 4.

The networks are still trained on original endpoint samples rather than on anchor coordinates. At the empirical level, let $C_i(\alpha)$ denote the set of training samples assigned to anchor $\alpha \in A_i$, and let $\hat{q}_i(\cdot \mid \alpha)$ be the uniform distribution on that cell. Each anchor $\alpha \in A_i$ therefore carries mass $|C_i(\alpha)|/n$, and the anchor-level OT problem is solved between weighted anchor (and thus discrete) measures rather than between uniform measures on A_0 and A_1. The anchor-level transport plan $\pi^{\tau}_{\tilde{q}_0,\tilde{q}_1}$ induces the coupling $q_{\text{anc}}(x_0, x_1) = \sum_{\alpha\in A_0} \sum_{\beta\in A_1} \pi^{\tau}_{\tilde{q}_0,\tilde{q}_1}(\alpha, \beta)\, \hat{q}_0(x_0 \mid \alpha)\, \hat{q}_1(x_1 \mid \beta)$. By construction, q_{anc} has the empirical marginals of the two endpoint datasets; the approximation therefore affects only how source and target points are paired, not which endpoint samples are seen.

This construction separates the pairing step from the sample generation step used during training. The anchor-level transport plan determines which source cells are matched to which target cells, while sampling inside these cells provides original endpoint pairs for the Brownian-bridge supervision. As a result, the expensive OT computation is moved outside the inner optimization loop and solved on a much smaller discrete problem. The next two subsections justify this.

3.2.1 THEORY: ANCHOR QUANTIZATION STABILITY

The central approximation underlying QDSB is that the entropic optimal transport problem is solved after first quantizing the endpoint distributions onto finite anchor sets. The key theoretical question is therefore whether the resulting anchor-level transport plan remains sufficiently close to the transport plan that would be obtained from the unquantized marginals. In the following, we answer this question in the affirmative with only mild assumptions on the transport cost and the regularized transport problem. Specifically, Assumption 3.1 controls the sensitivity of the transport cost to perturbations of the marginals (admissibility). The admissibility requirement is mild and includes standard choices such as the squared Euclidean cost. Assumption 3.2 is the transport inequality of Eckstein & Nutz (2022) that relates the stability of the objective values to the stability of optimal couplings. Similarly, this requirement is also mild and satisfied in the common case of $a = 2$ and sub-Gaussian marginals (see discussion under (Eckstein & Nutz, 2022, Lemma 3.10)).

Assumption 3.1 (Admissible cost)**.** Let $a \in [1, \infty]$. The cost function $c : \mathbb{R}^d \times \mathbb{R}^d \to [0, \infty)$ is Borel measurable and satisfies the admissibility condition of Eckstein & Nutz (2022, Definition 3.3) with constant $L \geq 0$, that is, for any probability measures $r_0, r_0', r_1, r_1' \in \mathcal{P}_a(\mathbb{R}^d)$, and any couplings $\pi \in \Pi(r_0, r_1)$, $\pi' \in \Pi(r_0', r_1')$, it holds that

$$\left| \int c \, \mathrm{d}(\pi - \pi') \right| \leq L W_a'(\pi, \pi') \text{ with } W_a'(\pi, \pi') := \begin{cases} (W_a(r_0, r_0')^a + W_a(r_1, r_1')^a)^{\frac{1}{a}} & \text{if } a < \infty, \\ \max\{W_a(r_0, r_0'), W_a(r_1, r_1')\} & \text{if } a = \infty. \end{cases}$$

Assumption 3.2 (Transport inequality)**.** Let $a \in [1, \infty]$ and $b \in [1, \infty)$ with $b \leq a$. The endpoint distributions $q_0, q_1 \in \mathcal{P}_a(\mathbb{R}^d)$ satisfy the transport inequality Eq. (I_q) of Eckstein & Nutz (2022) with constant $C_b \geq 0$, that is, $W_b'(\pi, \pi') \leq C_b \, \mathrm{KL}(\pi, \pi')^{1/(2b)}$ for all $\pi, \pi' \in \Pi(q_0, q_1)$.

Assumptions 3.1 and 3.2 provide the ingredients needed to control the effect of anchor quantization on the transport problem in (2). The first assumption bounds the perturbation of the objective value induced by changes in the marginals, while the second converts this value control into a bound on the corresponding optimal couplings. Under these assumptions, Theorem 3.3 shows that the resulting approximation error scales directly with the quantization error of the anchor projections.

Theorem 3.3 (Anchor stability)**.** *Let $i \in \{0, 1\}$, $a \in [1, \infty]$, $q_i \in \mathcal{P}_a(\mathbb{R}^d)$, and $\tilde{q}_i := (T_i)_{\#} q_i$, where $T_i : \mathbb{R}^d \to A_i$ are measurable projections onto finite anchor sets $A_i \subseteq \mathbb{R}^d$. Define the quantization error as $\varepsilon_i := (\int \|x_i - T_i(x_i)\|^a \, \mathrm{d}q_i(x_i))^{1/a}$ if $a < \infty$, and $\varepsilon_i := \operatorname{ess\,sup}_{x_i \sim q_i} \|x_i - T_i(x_i)\|$ if $a = \infty$. Set $\Delta_a := (\varepsilon_0^a + \varepsilon_1^a)^{1/a}$ if $a < \infty$, and $\Delta_a := \max\{\varepsilon_0, \varepsilon_1\}$ if $a = \infty$. Recall the entropic OT objective $\mathcal{S}_\tau$ from (2). Let $\pi^\tau_{q_0,q_1}$ and $\pi^\tau_{\tilde{q}_0,\tilde{q}_1}$ denote the corresponding optimal couplings w.r.t. $\mathcal{S}_\tau(q_0, q_1)$ and $\mathcal{S}_\tau(\tilde{q}_0, \tilde{q}_1)$, respectively. Further, assume that $c : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ is Borel measurable and admissible with constant $L \geq 0$ (Assumption 3.1), and that the transport inequality in Assumption 3.2 holds for (q_0, q_1) with $C_b \geq 0$. Then, the following bounds hold:*

(i) ***Endpoint perturbation bound:*** $W_a(q_0, \tilde{q}_0) \leq \varepsilon_0$ *and* $W_a(q_1, \tilde{q}_1) \leq \varepsilon_1$.

(ii) ***Entropic value stability:*** $|\mathcal{S}_\tau(q_0, q_1) - \mathcal{S}_\tau(\tilde{q}_0, \tilde{q}_1)| \leq L\Delta_a$.

(iii) ***Entropic coupling stability:*** $W_b'(\pi^\tau_{q_0,q_1}, \pi^\tau_{\tilde{q}_0,\tilde{q}_1}) \leq 2^{1/b-1/a}\Delta_a + C_b(2L\Delta_a/\tau)^{1/(2b)}$.

Theorem 3.3 has three practical implications. First, quantizing each endpoint distribution only moves mass by at most the quantization error ε_i. Second, the entropic OT value changes at most linearly in the combined error scale Δ_a. Third, the Wasserstein distance between the optimal regularized coupling itself and its quantized version is bounded. As a result, the anchor-level transport plan remains close to the original transport plan whenever the anchor cells are sufficiently small.

Lemma 3.4 turns the abstract quantity Δ_a into a geometric design principle for choosing anchors. If every point lies within radius r_i of its assigned anchor, then all stability terms are controlled by these radii. Consequently, anchor selection should emphasize coverage of the endpoint distributions rather than focusing exclusively on high-density regions. This is exactly the metric k-center problem (Har-Peled, 2011): Given a finite set of points $V \subseteq \mathbb{R}^d$ and number of anchors k, its goal is to find a set $C \subseteq V$ with $|C| = k$ that minimizes the coverage radius $r^* = \max_{u \in V} \min_{v \in C} c(u, v)$. In other words, for all points $x \in V$ and their associated anchors $y_x \in C$, it holds that $c(x, y_x) \leq r^*$. The greedy farthest-first traversal algorithm approximately solves the k-center problem and has an approximation factor of 2, that is, $r \leq 2r^*$ (elaborated in Theorem B.3). We note that a 2-approximation is the best possible approximation with polynomial runtime.

Lemma 3.4 (Radius-based bound)**.** *Assume the setting of Theorem 3.3. Suppose that, for each $i \in \{0, 1\}$, the anchor projection T_i satisfies $\|x_i - T_i(x_i)\| \leq r_i$ for q_i-almost every x_i. Then, (i) $\varepsilon_i \leq r_i$, and (ii) $\Delta_a \leq (r_0^a + r_1^a)^{1/a}$ if $a < \infty$, and $\Delta_a \leq \max\{r_0, r_1\}$ if $a = \infty$.*

If the projections T_i satisfy the radius condition in Lemma 3.4, the bounds in Theorem 3.3 can be read directly in terms of geometric coverage of the endpoint distributions.

Remark 3.5. Lemma 3.4 (i) and (ii) can be substituted directly into Theorem 3.3 (i)–(iii). Hence, once the anchor sets A_0 and A_1 induce T_1 and T_2 achieving the desired coverage radii, stability of the quantized couplings is solely controlled by the granularity of the anchor discretization.

At the same time, the usual limitations of high-dimensional metric quantization arise: if distances become less informative, obtaining small coverage radii may require many anchors. Nevertheless, our empirical results, which include an image translation experiment in a 512-dimensional latent space, indicate that the approximation remains effective in practice. One reason is that the anchor-level transport plan only determines which source and target cells interact, while the Brownian bridge supervision is still computed from original endpoints sampled within these cells. Hence, when the cells are sufficiently small, re-pairing points within a cell does not substantially change the training.

Algorithm 1 QDSB: Quantized diffusion Schrödinger bridges.

Input: Samples $\{x_0^{(i)}\}_{i=1}^n$ of source distribution q_0, samples $\{x_1^{(i)}\}_{i=1}^n$ of target distribution q_1, anchor set cardinality $k \in \mathbb{N}$, anchor refresh interval $t_{\text{refresh}} \in \mathbb{N}$;
Output: Trained score network s_θ and drift network v_θ;

1: ▷ *Anchor construction*
2: $A_0, A_1 \leftarrow$ Select k points each from $\{x_0^{(i)}\}_{i=1}^n$ and $\{x_1^{(i)}\}_{i=1}^n$ using farthest-first traversal;
3: Compute empirical OT plan $\pi^\tau_{\tilde{q}_0,\tilde{q}_1}$ between the anchor measures induced by A_0 and A_1;
4: ▷ *Training*
5: Initialize score network s_θ and drift network v_θ;
6: **while** training **do**
7: Refresh anchor sets A_0, A_1, and anchor coupling $\pi^\tau_{\tilde{q}_0,\tilde{q}_1}$ every t_{refresh} epochs (see lines 1–3);
8: Sample $t \in [0,1]$ and an anchor pair $(\tilde{\alpha}, \tilde{\beta}) \sim \pi^\tau_{\tilde{q}_0,\tilde{q}_1}$;
9: Sample $\tilde{x}_0 \sim \hat{q}_0(\cdot \mid \tilde{\alpha})$, $\tilde{x}_1 \sim \hat{q}_1(\cdot \mid \tilde{\beta})$, set $z \leftarrow (\tilde{x}_0, \tilde{x}_1)$, and sample $x \sim p_t(x \mid z)$;
10: Compute (3), that is, $\mathcal{L}_\theta \leftarrow \|v_\theta(t,x) - u_t^\circ(x \mid z)\|^2 + \lambda(t)^2 \|s_\theta(t,x) - \nabla \log p_t(x \mid z)\|^2$;
11: Update θ using backpropagation on $\mathcal{L}_\theta$;

3.2.2 Algorithm

Algorithm 1 summarizes the proposed training procedure. It consists of an outer coupling stage, which constructs and periodically refreshes the anchor-level transport plan, and an inner training stage, which reuses this plan to generate endpoint pairs for the simulation-free loss.

The **first stage** constructs the anchor representation of the two endpoint datasets. We select anchors independently on source and target data via farthest-first traversal, which is motivated in Lemma 3.4 and detailed in Algorithm 2 in the appendix. Given the anchor sets, nearest-anchor assignment induces source cells $C_0(\alpha)$ for $\alpha \in A_0$ and target cells $C_1(\beta)$ for $\beta \in A_1$, together with the corresponding empirical measures supported on the anchors and used for solving the (entropic) OT problem.

The **second stage** amortizes the transport plan computation across many gradient updates. Instead of solving an OT problem for each training minibatch, QDSB computes a discrete entropic OT plan on the anchor supports only and refreshes it every t_{refresh} epochs. This substantially reduces the computational overhead while still allowing the anchor geometry and coupling to adapt over the course of training. Within the training loop, the anchor-level transport plan is lifted back to pairs of original data points. Sampling $(\tilde{\alpha}, \tilde{\beta})$ first chooses a matched pair of source and target cells, and sampling uniformly from $C_0(\tilde{\alpha})$ and $C_1(\tilde{\beta})$ then restores the empirical variability that would be lost if we trained on anchors alone. This step is crucial in practice, since the score and drift supervision are computed from Brownian bridges between actual endpoints rather than between anchors. Given $z = (\tilde{x}_0, \tilde{x}_1)$, an intermediate state is sampled from the associated Brownian bridge, and the networks are trained to regress toward the closed-form drift and score. Increasing the number of anchors refines the geometry captured by the transport plan and recovers the original objective in the limit.

3.3 Positioning w.r.t. minibatch OT

Tong et al. (2024) use minibatch OT to approximate the endpoint coupling, making it natural to contrast this approach with QDSB. Minibatch OT replaces q_0 and q_1 by uniform empirical measures on random subsamples and solves OT only on these supports. Fatras et al. (2021) analyze the resulting averaged surrogate and Monte Carlo estimator. The connection to QDSB becomes explicit when minibatches are interpreted as random anchor sets A_0 and A_1. Given these sampled supports, Theorem 3.3 and Lemma 3.4 show that the induced coupling error is controlled by the corresponding quantization scale Δ_a, and equivalently by the covering radii of A_0 and A_1. The key differences are that QDSB selects anchors to explicitly minimize this radius and constructs weighted pushforward measures $(T_i)_\# q_i$ rather than uniform minibatch weights. The resulting stability guarantees make QDSB more suitable than random minibatch OT approximations for simulation-free SB training.

Table 1: MMD results on three toy settings (Table 1a) and three real-world datasets (Table 1b). We report mean and standard deviation over five seeds, and highlight the best value in each column.

(a) MMD results on 2D toy experiments. Lower values are better.

Method	After 10 s			After 60 s			500 epochs + total time			
	8G→M	G→M	G→8G	8G→M	G→M	G→8G	8G→M	G→M	G→8G	Time
QDSB (ours)	**0.024** ± 0.007	**0.015** ± 0.010	**0.015** ± 0.015	**0.019** ± 0.007	**0.004** ± 0.005	**0.010** ± 0.016	0.023 ± 0.010	0.010 ± 0.009	**0.010** ± 0.016	105 s
DSB (2021)	0.181 ± 0.022	0.240 ± 0.041	0.208 ± 0.010	0.238 ± 0.015	0.165 ± 0.018	0.261 ± 0.013	0.202 ± 0.014	0.168 ± 0.023	0.259 ± 0.033	**103 s**
DSBM (2023)	0.081 ± 0.005	0.049 ± 0.005	0.103 ± 0.025	0.182 ± 0.046	0.083 ± 0.025	0.061 ± 0.021	0.139 ± 0.029	0.128 ± 0.027	0.103 ± 0.039	193 s
SF2M (2024)	0.040 ± 0.007	0.028 ± 0.004	0.041 ± 0.009	0.024 ± 0.008	0.020 ± 0.005	**0.010** ± 0.017	**0.022** ± 0.016	**0.006** ± 0.008	0.011 ± 0.016	205 s
SF2M + mPOT	0.109 ± 0.008	0.202 ± 0.008	0.108 ± 0.004	0.092 ± 0.008	0.188 ± 0.009	0.063 ± 0.006	0.090 ± 0.012	0.181 ± 0.007	0.057 ± 0.006	258 s
LightSB-M (2024)	0.669 ± 0.000	0.633 ± 0.002	0.673 ± 0.001	0.671 ± 0.001	0.620 ± 0.002	0.669 ± 0.001	0.644 ± 0.009	0.225 ± 0.006	0.603 ± 0.004	550 s

(b) MMD results on real-world cellular biology experiments. Lower values are better.

Method	After 10 s			After 60 s			1000 epochs + total time			
	EB	Cite	Multi	EB	Cite	Multi	EB	Cite	Multi	Time
QDSB (ours)	**0.198** ± 0.023	**0.183** ± 0.004	**0.224** ± 0.054	**0.193** ± 0.026	**0.179** ± 0.003	0.231 ± 0.060	0.193 ± 0.026	0.183 ± 0.005	0.232 ± 0.061	**74 s**
DSB (2021)	0.264 ± 0.093	0.195 ± 0.034	0.295 ± 0.084	0.254 ± 0.053	0.219 ± 0.039	0.313 ± 0.146	0.221 ± 0.029	0.245 ± 0.044	0.318 ± 0.140	103 s
DSBM (2023)	0.233 ± 0.031	0.198 ± 0.056	0.247 ± 0.033	0.221 ± 0.055	0.189 ± 0.029	0.230 ± 0.018	0.219 ± 0.054	0.189 ± 0.025	0.231 ± 0.024	92 s
SF2M (2024)	0.214 ± 0.021	0.185 ± 0.016	0.229 ± 0.051	0.197 ± 0.019	**0.179** ± 0.008	0.223 ± 0.052	0.190 ± 0.028	0.180 ± 0.007	0.232 ± 0.062	318 s
SF2M + mPOT	0.216 ± 0.032	0.186 ± 0.014	**0.224** ± 0.031	0.215 ± 0.031	0.186 ± 0.007	**0.222** ± 0.036	0.206 ± 0.039	0.185 ± 0.003	**0.226** ± 0.045	468 s
LightSB-M (2024)	0.451 ± 0.072	0.371 ± 0.113	0.376 ± 0.045	0.426 ± 0.069	0.344 ± 0.119	0.359 ± 0.048	**0.189** ± 0.027	**0.176** ± 0.017	0.244 ± 0.070	2742 s

4 EXPERIMENTS

In this section, after summarizing the SB baselines, we benchmark our method in synthetic and real-world settings.

Experimental setup. We compare QDSB with five bridge baselines: DSB (De Bortoli et al., 2021), DSBM (Shi et al., 2023), SF2M (Tong et al., 2024), SF2M combined with mini-batch partial OT (mPOT) (Nguyen et al., 2022b), and LightSB-M (Gushchin et al., 2024). DSB is the original diffusion SB method, DSBM replaces the simulation-based training step by a matching objective, SF2M is the simulation-free baseline on which our method builds, SF2M + mPOT strengthens the minibatch coupling step, and LightSB-M is a recent SB solver. For a fair comparison, all methods use the same backbone models and the same training hyperparameters. The time measurements are conducted on the same machine (AMD Ryzen 9 7900X 12-Core, NVIDIA GeForce RTX 4090) with only one experiment running at a time. The Maximum Mean Discrepancy (MMD) values are computed between the groundtruth target samples and the predicted target samples from simulated trajectories of the trained models. We use the Radial Basis Function Kernel (RBF-Kernel) with its parameter set using the median distance heuristic; the RBF parameter is fixed per dataset.

2D toy experiments. Following related work, we first evaluate all methods on three standard 2D toy settings between the 8Gaussians distribution (blue in Figure 1), the Moons distribution (green in Figure 1), and a standard two-dimensional Gaussian distribution. This yields the three pairwise tasks 8G→M, G→M, and G→8G shown in Table 1a and Figure 2a–c. Table 1a shows that QDSB achieves

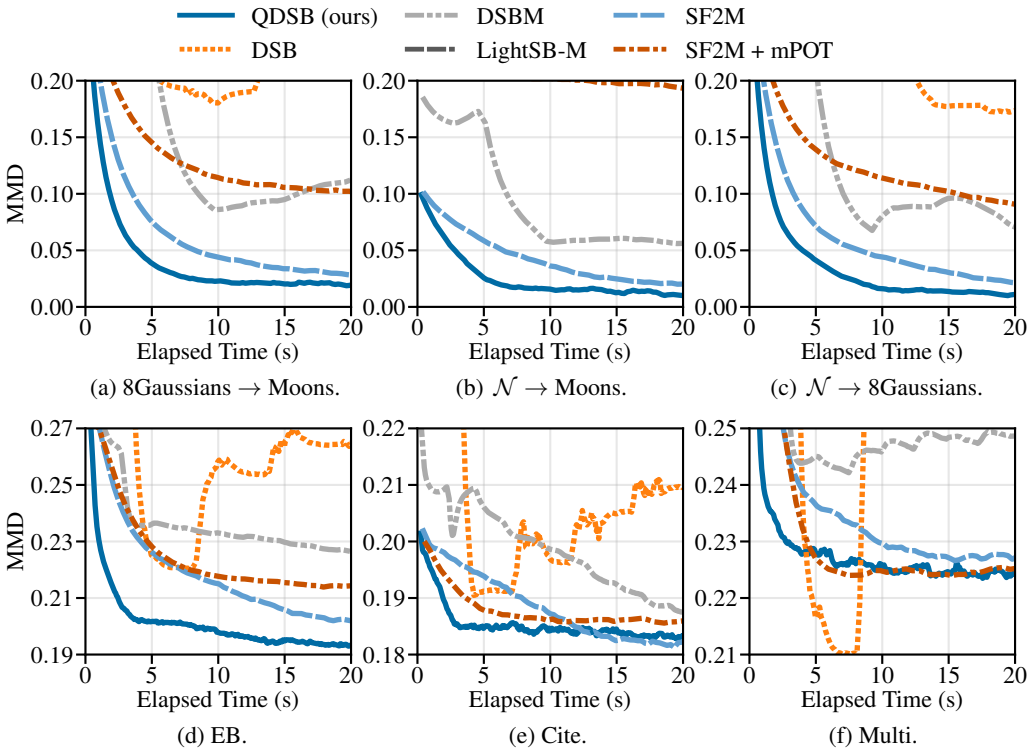

Figure 2: Time-quality trade-off curves measured by MMD over wall-clock time. Panels (a)–(c) show the 2D toy datasets, and panels (d)–(f) show the real-world cellular biology datasets. Curves closest to the bottom-left corner indicate the best time-quality trade-off.

the lowest MMD after 10 s and after 60 s on all three tasks, and it has comparable results after 500 epochs. Yet, it requires less total wall-clock time compared to most other approaches to complete training. Figure 2a–c confirms this over the full training horizon: the QDSB curves stay closest to the bottom-left corner in all cases, indicating the best time-quality trade-off among the methods. We note that DSB fluctuates strongly in some experiments because of alternating between fitting neural networks to the forward and backward process, respectively.

Real-world cellular biology experiments. We consider three real-world cellular biology datasets from Tong et al. (2024). EB is an embryoid body differentiation dataset with five population measurements over 30 days, while Cite and Multi are multimodal single-cell benchmarks (Luecken et al., 2025). Following Tong et al. (2024), we use four measurements on days 2, 3, 4, and 7. Table 1b shows that QDSB achieves the best MMD after 10 s on all three datasets, and it remains best on EB and tied best on Cite after 60 s. At the longest budget, some baselines obtain slightly lower final MMD values, but only at a substantially higher runtime. Figure 2d–f shows that QDSB reaches a strong performance faster and therefore provides the best overall time-quality trade-off on these tasks.

Sensitivity and ablation experiments. Figure 3 illustrates the effect of the number of anchors per marginal on the 8Gaussians-to-Moons task, and reports the resulting MMD (Figure 3a), the median coverage radius of the selected anchors (Figure 3b), and their relation (Figure 3c), averaged over five seeds. As the number of anchors increases, the MMD improves consistently (except for the slight bump at $k = 512$, which might be an artifact of the random initial anchor selection). This is consistent with our theory, since a finer anchor quantization yields a more accurate coupling approximation. Beyond about $k = 16$, however, the gains become smaller, which matches the fact that the coverage radius decreases more slowly afterwards. The case $k = 1$ also serves as an ablation experiment. This is because with only one anchor per marginal, all samples lie in a single anchor cell, so the induced coupling is equivalent to a random OT plan not using any geometric structure. Accordingly, $k \in \{1, 2, 4\}$ gives the worst quality as such few anchors do not convey enough geometry.

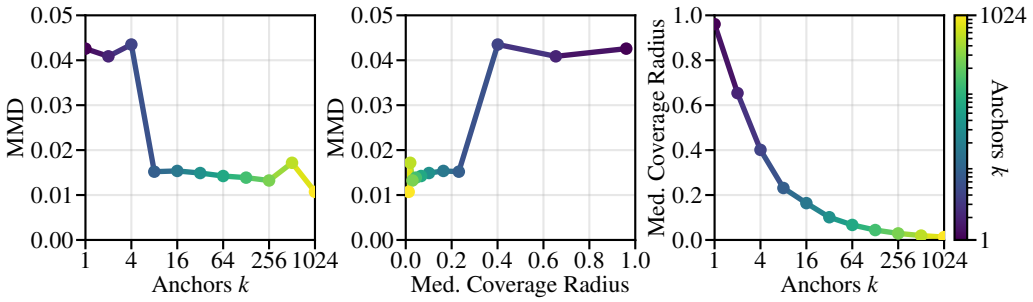

(a) Anchors vs. MMD. (b) Coverage vs. MMD. (c) Anchors vs. coverage.

Figure 3: Sensitivity and ablation experiments on the 8Gaussians to Moons dataset. Each point represents a separate QDSB training run over 1000 epochs, averaged over five different seeds.

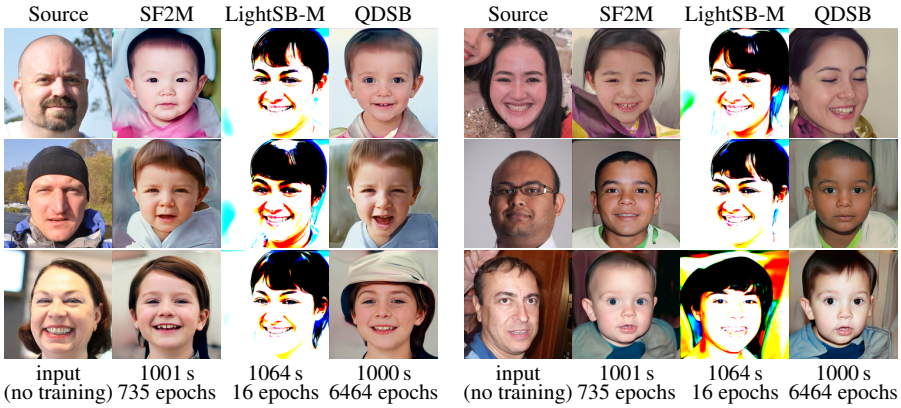

Figure 4: Qualitative comparison of adult-to-child image translation on FFHQ with a fixed time budget of 1000 s. The left column shows adult input images, and the remaining columns show the translated child images. Additional examples are provided in Appendix C.

Image translation with fixed time budget. Following the image-generation experiments of Gushchin et al. (2024); Alouadi et al. (2026), we consider image translation on FFHQ (Karras et al., 2019), which contains unpaired images of adults and children. The task is to learn a map between the adult and child distributions. We encode all images with the ALAE encoder (Pidhorskyi et al., 2020) and run all methods in the resulting 512-dimensional latent space. All methods are trained under the same time budget of 1000 s. Figure 4 shows the resulting samples. We were not able to obtain usable results from DSB and DSBM in this setting. Within the time budget, QDSB completes 6464 epochs, while LightSB-M completed only 16 epochs. This gap is reflected clearly in the image quality: LightSB-M remains far from a usable translation regime, whereas SF2M and especially QDSB produce coherent images. Overall, QDSB gives the strongest qualitative results.

5 Conclusion

We proposed a quantized coupling framework for simulation-free SB training from unpaired samples. QDSB replaces repeated minibatch OT solves by a single entropic transport problem on anchor-quantized endpoint distributions and then lifts the resulting anchor-level transport plan back to the original data via cell-wise sampling. We provided a theoretical analysis of this approximation, showing that both the entropic transport value and the optimal regularized coupling remain stable under quantization and that anchor coverage directly controls the induced error. Empirically, experiments on toy benchmarks and cellular biology datasets demonstrate that QDSB preserves the generation quality of existing simulation-free baselines while substantially improving wall-clock time needed for training. Overall, quantized couplings offer a simple, principled, and effective way to alleviate the coupling bottleneck in simulation-free SB training.

AI USE STATEMENT

We used AI models to polish the manuscript text, assist in creating experimental code, and prepare the plots. We verified all AI-generated code and text and take responsibility for the final content, claims, and artifacts of this work.

ETHICS STATEMENT

This paper presents basic research on generative modeling and, to our knowledge, raises no ethical concerns specific to the proposed method or experiments. Nevertheless, advances in generative modeling can have both beneficial and harmful downstream uses. Applications involving image or cellular data should consider data provenance, privacy, consent, and potential bias in their particular context. We support responsible use of the method in accordance with the ICLR Code of Ethics.

REPRODUCIBILITY STATEMENT

Section 3 describes the method and coupling algorithm, and Section 4 presents the evaluation setup and results. Appendix A contains complete proofs, while Appendices C and D provide further results and experimental details. Code and data for reproducing the reported tables and figures are available at `github.com/anon1248/qdsb`.

REFERENCES

Alexandre Alouadi, Pierre Henry-Labordère, Grégoire Loeper, Othmane Mazhar, Huyên Pham, and Nizar Touzi. LightSBB-M: Bridging Schrödinger and Bass for generative diffusion modeling. *CoRR*, abs/2601.19312, 2026.

Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 22563–22575, 2023.

Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 26, pp. 2292–2300, 2013.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion Schrödinger bridge with applications to score-based generative modeling. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, pp. 17695–17709, 2021.

Stephan Eckstein and Marcel Nutz. Quantitative stability of regularized optimal transport and convergence of Sinkhorn's algorithm. *SIAM Journal on Mathematical Analysis*, 54(6):5922–5948, 2022.

Kilian Fatras, Thibault Sejourne, Rémi Flamary, and Nicolas Courty. Unbalanced minibatch optimal transport; applications to domain adaptation. In *International Conference on Machine Learning (ICML)*, volume 139, pp. 3186–3197, 2021.

Hans Föllmer. Random fields and diffusion processes. In *École d'Été de Probabilités de Saint-Flour XV–XVII, 1985–87*, pp. 101–203. Springer, 1988.

Nikita Gushchin, Sergei Kholkin, Evgeny Burnaev, and Alexander Korotin. Light and optimal Schrödinger bridge matching. In *International Conference on Machine Learning (ICML)*, volume 235, pp. 17100–17122, 2024.

Sariel Har-Peled. *Geometric Approximation Algorithms*. American Mathematical Society, 2011.

Jonathan Ho, Ajay N. Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pp. 6840–6851, 2020.

Leonid Kantorovich. On the translocation of masses. In *Proceedings of the USSR Academy of Sciences*, volume 37, pp. 199–201, 1942.

Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4401–4410, 2019.

Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations (ICLR)*, 2021.

Jean-Marie Lemercier, Julius Richter, Simon Welker, Eloi Moliner, Vesa Välimäki, and Timo Gerkmann. Diffusion models for audio restoration: A review. *IEEE Signal Processing Magazine*, 41(6):72–84, 2024.

Chi-Heng Lin, Mehdi Azabou, and Eva Dyer. Making transport more robust and interpretable by moving data through a small number of anchor points. In *International Conference on Machine Learning (ICML)*, volume 139, pp. 6631–6641, 2021.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *International Conference on Learning Representations (ICLR)*, 2023.

Guan-Horng Liu, Yaron Lipman, Maximilian Nickel, Brian Karrer, Evangelos Theodorou, and Ricky T. Q. Chen. Generalized Schrödinger bridge matching. In *International Conference on Learning Representations (ICLR)*, 2024.

Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, pp. 700–708, 2017.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *International Conference on Learning Representations (ICLR)*, 2023.

Malte D. Luecken, Scott Gigante, Daniel B. Burkhardt, Robrecht Cannoodt, Daniel C. Strobl, Nikolay S. Markov, Luke Zappia, Giovanni Palla, Wesley Lewis, Daniel Dimitrov, Michael E. Vinyard, D. S. Magruder, Michaela F. Mueller, Alma Andersson, Emma Dann, Qian Qin, Dominik J. Otto, Michal Klein, Olga Borisovna Botvinnik, Louise Deconinck, Kai Waldrant, Sai Nirmayi Yasa, Artur Szałata, Andrew Benz, Zhijian Li, Jonathan M. Bloom, Angela Oliveira Pisco, Julio Saez-Rodriguez, Drausin Wulsin, Luca Pinello, Yvan Saeys, Fabian J. Theis, and Smita Krishnaswamy. Defining and benchmarking open problems in single-cell analysis. *Nature Biotechnology*, 43(7): 1035–1040, 2025.

Gaspard Monge. Mémoire sur la théorie des déblais et des remblais. *Mémoires de mathématique et de physique, presentés à l'Académie royale des sciences*, pp. 666–704, 1781.

Khai Nguyen, Dang Nguyen, Quoc Dinh Nguyen, Tung Pham, Hung Bui, Dinh Phung, and Trung Le. On transportation of mini-batches: A hierarchical approach. In *International Conference on Machine Learning (ICML)*, volume 162, pp. 16091–16119, 2022a.

Khai Nguyen, Dang Nguyen, The-Anh Vu-Le, Tung Pham, and Nhat Ho. Improving mini-batch optimal transport via partial transportation. In *International Conference on Machine Learning (ICML)*, volume 162, pp. 16656–16690, 2022b.

Stanislav Pidhorskyi, Donald A. Adjeroh, and Gianfranco Doretto. Adversarial latent autoencoders. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14104–14113, 2020.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10674–10685, 2022.

Meyer Scetbon, Marco Cuturi, and Gabriel Peyré. Low-rank Sinkhorn factorization. In *International Conference on Machine Learning (ICML)*, volume 139, pp. 9344–9354, 2021.

Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion Schrödinger bridge matching. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, 2023.

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

Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *International Conference on Machine Learning (ICML)*, volume 202, pp. 32211–32252, 2023.

Alexander Tong, Jessie Huang, Guy Wolf, David van Dijk, and Smita Krishnaswamy. TrajectoryNet: A dynamic optimal transport network for modeling cellular dynamics. In *International Conference on Machine Learning (ICML)*, volume 119, pp. 9526–9536, 2020.

Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free Schrödinger bridges via score and flow matching. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 238, pp. 1279–1287, 2024.

Cédric Villani. *Optimal Transport: Old and New*. Springer, 2009.

Cédric Villani. *Topics in Optimal Transportation*. American Mathematical Society, 2021.

Yaohui Wang, Xinyuan Chen, Xin Ma, Shangchen Zhou, Ziqi Huang, Yi Wang, Ceyuan Yang, Yinan He, Jiashuo Yu, Peiqing Yang, Yuwei Guo, Tianxing Wu, Chenyang Si, Yuming Jiang, Cunjian Chen, Chen Change Loy, Bo Dai, Dahua Lin, Yu Qiao, and Ziwei Liu. LaVie: High-quality video generation with cascaded latent diffusion models. *International Journal of Computer Vision (IJCV)*, 133(5):3059–3078, 2025.

Jinjin Zhang, Qiuyu Huang, Junjie Liu, Xiefan Guo, and Di Huang. Diffusion-4k: Ultra-high-resolution image synthesis with latent diffusion models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 23464–23473, 2025.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *International Conference on Computer Vision (ICCV)*, pp. 2223–2232, 2017.

A PROOFS

This section collects all proofs of our results stated in the main text.

A.1 PROOF OF THEOREM 3.3

Part (i)**.** The proof of part (i) first establishes that $x \mapsto (x, T_i(x))$ describes a coupling. Then, we directly upper-bound the Wasserstein distance.

Let $i \in \{0, 1\}$ and $\Phi_i : \mathbb{R}^d \to \mathbb{R}^d \times \mathbb{R}^d$ with $\Phi_i(x) := (x, T_i(x))$. Since T_i is measurable, Φ_i is measurable. Define the measure $\gamma_i := (\Phi_i)_{\#} q_i$ on $\mathbb{R}^d \times \mathbb{R}^d$ equipped with its Borel sigma-algebra, and let $\zeta_x(x, y) := x$ and $\zeta_y(x, y) := y$ denote the coordinate maps. Then,

$$(\zeta_x)_{\#}\gamma_i = (\zeta_x)_{\#}(\Phi_i)_{\#} q_i = (\zeta_x \circ \Phi_i)_{\#} q_i = (\mathrm{Id}_{\mathbb{R}^d})_{\#} q_i = q_i, \text{ and}$$
$$(\zeta_y)_{\#}\gamma_i = (\zeta_y)_{\#}(\Phi_i)_{\#} q_i = (\zeta_y \circ \Phi_i)_{\#} q_i = (T_i)_{\#} q_i = \tilde{q}_i.$$

Therefore, γ_i is a coupling between q_i and $\tilde{q}_i$, that is, $\gamma_i \in \Pi(q_i, \tilde{q}_i)$. Then, we discern two cases.

- For $a < \infty$, we have

$$W_a(q_i, \tilde{q}_i)^a \overset{(*)}{=} \inf_{\pi \in \Pi(q_i, \tilde{q}_i)} \int \|x - y\|^a \, \mathrm{d}\pi(x, y) \overset{(**)}{\leq} \int \|x - y\|^a \, \mathrm{d}\gamma_i(x, y)$$
$$\overset{(***)}{=} \int \|x - T_i(x)\|^a \, \mathrm{d}q_i(x) = \varepsilon_i^a,$$

  where $(*)$ uses the definition of the Wasserstein metric, $(**)$ upper-bounds the infimum with the coupling γ_i, and $(***)$ uses the definition of Φ_i and the change-of-variables formula. Since all terms are non-negative, taking the a-th root gives $W_a(q_i, \tilde{q}_i) \leq \varepsilon_i$.

- Similarly, for $a = \infty$, we have

$$W_\infty(q_i, \tilde{q}_i) = \inf_{\pi \in \Pi(q_i, \tilde{q}_i)} \operatorname*{ess\,sup}_{(x,y)\sim\pi} \|x - y\| \leq \operatorname*{ess\,sup}_{(x,y)\sim\gamma_i} \|x - y\| = \operatorname*{ess\,sup}_{x \sim q_i} \|x - T_i(x)\| = \varepsilon_i.$$

Part (ii)**.** The proof proceeds by reducing $\mathcal{S}_\tau$ to the normalization used by Eckstein & Nutz (2022) and then applying their value-stability theorem (recalled in Theorem B.1).

Define the rescaled cost function $\bar{c} := \frac{1}{\tau} c$ and the rescaled objective function

$$\bar{\mathcal{S}}(\alpha, \beta) := \inf_{\pi \in \Pi(\alpha, \beta)} \int \bar{c} \, \mathrm{d}\pi(x, y) + \mathrm{KL}(\pi \| \alpha \otimes \beta),$$

for any $\alpha, \beta \in \mathcal{P}_a(\mathbb{R}^d)$. Since $\tau > 0$, we have $\mathcal{S}_\tau(\alpha, \beta) = \tau \bar{\mathcal{S}}(\alpha, \beta)$ because the positive constant τ commutes with taking the infimum.

We next show that the assumed admissibility of c with constant L implies that of $\bar{c}$ with constant L/τ. Indeed, let $r_0, r_0', r_1, r_1' \in \mathcal{P}_a(\mathbb{R}^d)$, $\pi \in \Pi(r_0, r_1)$, and $\pi' \in \Pi(r_0', r_1')$. Then,

$$\left| \int \bar{c} \, \mathrm{d}(\pi - \pi') \right| \overset{(*)}{=} \left| \int \frac{1}{\tau} c \, \mathrm{d}(\pi - \pi') \right| \overset{(**)}{=} \frac{1}{\tau} \left| \int c \, \mathrm{d}(\pi - \pi') \right| \overset{(***)}{\leq} \frac{L}{\tau} W_a'(\pi, \pi'), \tag{4}$$

where $(*)$ holds by the definition of $\bar{c}$, $(**)$ by the homogeneity of the integral, and $(***)$ uses Assumption 3.1.

Define the marginal perturbation size as

$$\delta_a := \begin{cases} (W_a(q_0, \tilde{q}_0)^a + W_a(q_1, \tilde{q}_1)^a)^{1/a} & \text{for } a < \infty, \\ \max\{W_\infty(q_0, \tilde{q}_0), W_\infty(q_1, \tilde{q}_1)\} & \text{for } a = \infty. \end{cases}$$

By part (i), $W_a(q_0, \tilde{q}_0) \leq \varepsilon_0$ and $W_a(q_1, \tilde{q}_1) \leq \varepsilon_1$. Therefore, if $a < \infty$,

$$\delta_a = (W_a(q_0, \tilde{q}_0)^a + W_a(q_1, \tilde{q}_1)^a)^{1/a} \overset{(*)}{\leq} (\varepsilon_0^a + \varepsilon_1^a)^{1/a} = \Delta_a,$$

where $(*)$ holds as $x \mapsto x^a$ and $x \mapsto x^{1/a}$ are non-decreasing on $[0, \infty)$. If $a = \infty$,

$$\delta_\infty = \max\{W_\infty(q_0, \tilde{q}_0), W_\infty(q_1, \tilde{q}_1)\} \overset{(**)}{\leq} \max\{\varepsilon_0, \varepsilon_1\} = \Delta_\infty,$$

where $(**)$ holds by the properties of the maximum.

We now apply Theorem B.1 with marginals (q_0, q_1) and $(\tilde{q}_0, \tilde{q}_1)$ and cost $\bar{c}$. Since $\bar{c}$ satisfies the admissibility condition with constant L/τ, the theorem yields

$$\left|\bar{\mathcal{S}}(q_0, q_1) - \bar{\mathcal{S}}(\tilde{q}_0, \tilde{q}_1)\right| \leq \frac{L}{\tau}\delta_a.$$

Putting all steps together, we obtain

$$|\mathcal{S}_\tau(q_0, q_1) - \mathcal{S}_\tau(\tilde{q}_0, \tilde{q}_1)| = \tau\left|\bar{\mathcal{S}}(q_0, q_1) - \bar{\mathcal{S}}(\tilde{q}_0, \tilde{q}_1)\right| \leq \tau\frac{L}{\tau}\delta_a \leq L\Delta_a.$$

Part (iii)**.** Recall $\bar{c}$, $\bar{\mathcal{S}}$, and δ_a from the proof of part (ii). Since scaling the objective by a positive constant does not change its minimizers, $\pi^\tau_{q_0,q_1}$ and $\pi^\tau_{\tilde{q}_0,\tilde{q}_1}$—assumed optimal for $\mathcal{S}_\tau$—are the optimal couplings of $\bar{\mathcal{S}}$.

We now apply Theorem B.2 with marginals (q_0, q_1) and $(\tilde{q}_0, \tilde{q}_1)$, cost $\bar{c}$, admissibility constant $\bar{L} := L/\tau$ (established in (4)), and the transport inequality constant C_b from Assumption 3.2. This yields

$$\begin{aligned} W_b'(\pi^\tau_{q_0,q_1}, \pi^\tau_{\tilde{q}_0,\tilde{q}_1}) &\overset{(*)}{\leq} 2^{1/b-1/a}\delta_a + C_b(2\bar{L}\delta_a)^{1/(2b)} \overset{(**)}{=} 2^{1/b-1/a}\delta_a + C_b(2L\delta_a/\tau)^{1/(2b)} \\ &\overset{(***)}{\leq} 2^{1/b-1/a}\Delta_a + C_b(2L\Delta_a/\tau)^{1/(2b)}, \end{aligned}$$

where $(*)$ applies Theorem B.2, $(**)$ substitutes $\bar{L}$, and $(***)$ holds by the intermediate results in the proof of part (ii) and as the map

$$x \mapsto 2^{1/b-1/a}x + C_b(2Lx/\tau)^{1/(2b)}$$

is non-decreasing on $[0, \infty)$. This completes the proof of part (iii).

A.2 Proof of Lemma 3.4

Part (i)**.** By assumption, for each $i \in \{0, 1\}$, $\|x_i - T_i(x_i)\| \leq r_i$ for q_i-almost every x_i. Then, we discern two cases.

- For $a < \infty$, we have
$$\varepsilon_i^a \overset{(*)}{=} \int \|x_i - T_i(x_i)\|^a \,\mathrm{d}q_i(x_i) \overset{(**)}{\leq} \int r_i^a \,\mathrm{d}q_i(x_i) \overset{(***)}{=} r_i^a,$$
where $(*)$ substitutes the definition of ε_i, $(**)$ uses the assumption and that $x \mapsto x^a$ is non-decreasing on $[0, \infty)$, and $(***)$ simplifies the constant in the integral.
- For $a = \infty$, we have
$$\varepsilon_i \overset{(*)}{=} \operatorname*{ess\,sup}_{x_i \sim q_i} \|x_i - T_i(x_i)\| \overset{(**)}{\leq} r_i,$$
where $(*)$ substitutes the definition of ε_i and $(**)$ directly uses the assumption.

Part (ii)**.** Similar to the proof of part (i), we discern two cases.

- For $a < \infty$, we have
$$\Delta_a^a \overset{(*)}{=} \varepsilon_0^a + \varepsilon_1^a \overset{(**)}{\leq} r_0^a + r_1^a,$$
where $(*)$ substitutes the definition of Δ_a and $(**)$ uses part (i) and that $x \mapsto x^a$ is non-decreasing on $[0, \infty)$.
- For $a = \infty$, we have
$$\Delta_\infty \overset{(*)}{=} \max\{\varepsilon_0, \varepsilon_1\} \overset{(**)}{\leq} \max\{r_0, r_1\},$$
where $(*)$ substitutes the definition of Δ_a and $(**)$ follows by the properties of the maximum.

B External theorems

This section collects theorems that we use in the proofs of our statements.

Theorem B.1 (Eckstein & Nutz 2022, Theorem 3.7 (ii), $N = 2$)**.** *Let $a \in [1, \infty]$, $r_0, r_0', r_1, r_1' \in \mathcal{P}_a(\mathbb{R}^d)$, the cost function c admissible with constant $L \geq 0$ (Assumption 3.1), and*

$$\bar{\mathcal{S}}(\alpha, \beta) := \inf_{\pi \in \Pi(\alpha,\beta)} \left[\int c \, \mathrm{d}\pi(x, y) + \mathrm{KL}(\pi \| \alpha \otimes \beta) \right]. \tag{5}$$

Then,

$$|\bar{\mathcal{S}}(r_0, r_1) - \bar{\mathcal{S}}(r_0', r_1')| \leq L\delta_a, \text{ where } \delta_a := \begin{cases} (W_a(r_0, r_0')^a + W_a(r_1, r_1')^a)^{1/a} & \text{for } a < \infty, \\ \max\{W_\infty(r_0, r_0'), W_\infty(r_1, r_1')\} & \text{for } a = \infty. \end{cases} \tag{6}$$

Theorem B.2 (Eckstein & Nutz 2022, Theorem 3.11, $N = 2$)**.** *Let $a \in [1, \infty]$, $b \in [1, \infty)$ with $b \leq a$, $r_0, r_0', r_1, r_1' \in \mathcal{P}_a(\mathbb{R}^d)$. Suppose that (r_0, r_1) satisfy Assumption 3.2 with constant C_b, and that the cost function c satisfies Assumption 3.1. Recall $\bar{\mathcal{S}}$ from* (5) *and δ_a from* (6)*. Then, the optimizers $\pi_{r_0,r_1}, \pi_{r_0',r_1'}$ of $\bar{\mathcal{S}}(r_0, r_1)$ and $\bar{\mathcal{S}}(r_0', r_1')$, respectively, satisfy*

$$W_b'(\pi_{r_0,r_1}, \pi_{r_0',r_1'}) \leq 2^{1/b-1/a}\delta_a + C_b(2L\delta_a)^{1/(2b)}.$$

The next result provides guarantees on the farthest-first traversal approach to the k-center problem, recalled in Algorithm 2 for self-completeness.

Algorithm 2 Greedy k-center: Farthest-first traversal.

Input: Finite point set $\{x^{(i)}\}_{i=1}^n \subseteq \mathbb{R}^d$, anchor set cardinality $k \in \{1, \ldots, n\}$;
Output: Anchor set $A \subseteq \{x^{(i)}\}_{i=1}^n$ with $|A| = k$;

1: Select an arbitrary initial point $x_{\text{init}} \in \{x^{(i)}\}_{i=1}^n$ and set $A \leftarrow \{x_{\text{init}}\}$;
2: **while** $|A| < k$ **do**
3: $\quad x_{\text{next}} \leftarrow \arg\max_{x \in \{x^{(i)}\}_{i=1}^n} \min_{\alpha \in A} \|x - \alpha\|$;
4: $\quad A \leftarrow A \cup \{x_{\text{next}}\}$;
5: **return** A;

Theorem B.3 (Har-Peled 2011, Theorem 4.3)**.** *Given a finite set of points $V \subseteq \mathbb{R}^d$ and number of anchors k, the goal of the k-center clustering problem is to find a set $C \subseteq V$ with $|C| = k$ that minimizes the associated coverage radius*

$$r_C = \max_{u \in V} \min_{v \in C} c(u, v).$$

The optimal coverage radius is denoted by r^*. The greedy k-center algorithm in Algorithm 2 is a 2-approximation to the optimal k-center clustering problem, that is, $r_A \leq 2r^*$, where r_A uses the anchor set A output by the algorithm.

C Additional experimental results

Figures 5 and 6 provide additional qualitative examples for the fixed-budget FFHQ experiment from Figure 4. They use the same adult-to-child translation task, the same 512-dimensional ALAE latent space, and the same wall-clock budget of 1000 s. As in the main figure, we only report SF2M, LightSB-M, and QDSB because we were not able to obtain usable results for DSB and DSBM. The qualitative trend is unchanged: QDSB produces the most consistent child-face translations, SF2M often gives reasonable outputs, and LightSB-M remains unstable given this budget.

D Additional experimental details

This section augments Section 4 by summarizing the main training details and hyperparameter choices of all reported methods.

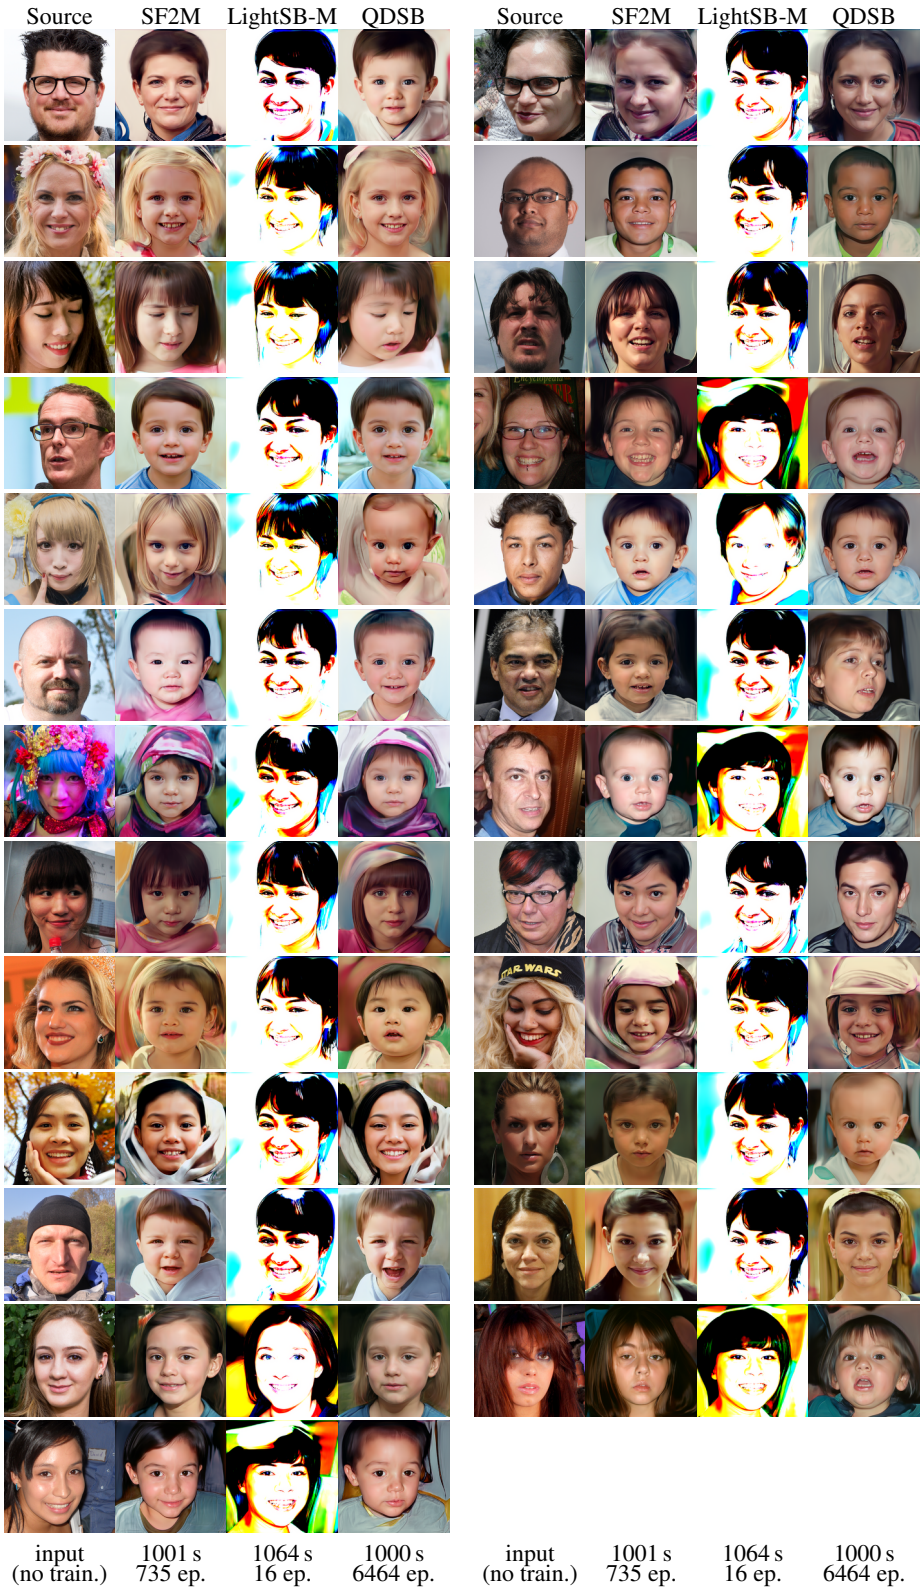

Figure 5: Additional qualitative results for the fixed-budget FFHQ experiment from Figure 4. Each row shows an adult input together with the corresponding outputs of SF2M, LightSB-M, and QDSB under the same training-time budget.

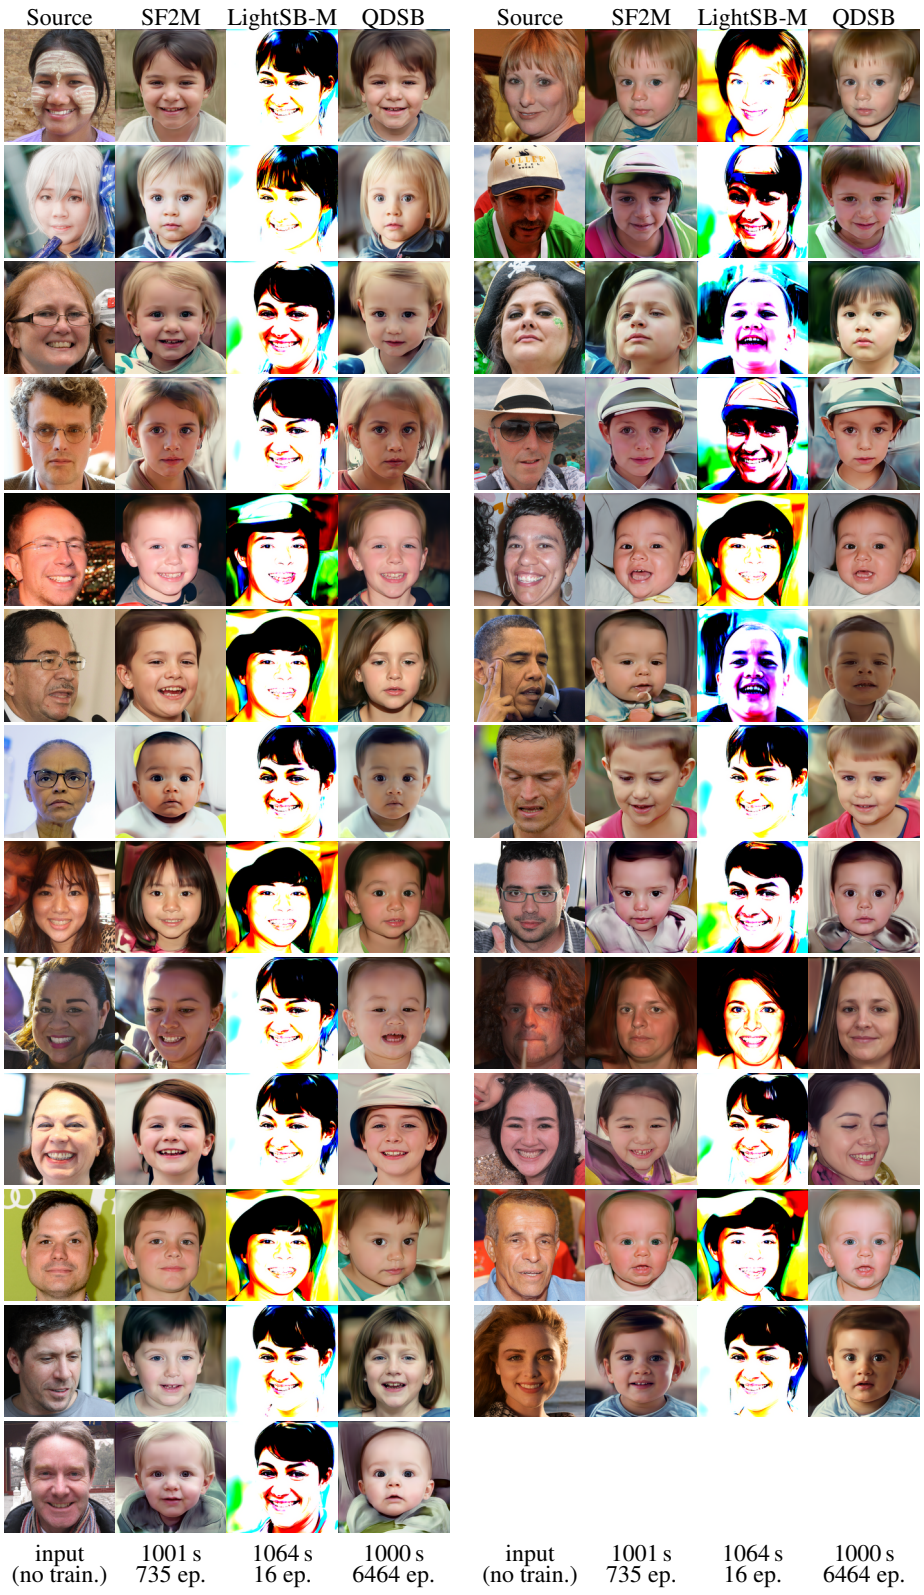

Figure 6: More qualitative results for the fixed-budget FFHQ experiment from Figure 4. The same pattern as in Figure 5 can be observed across additional examples.

Common setup. All code and data are available at `github.com/anon1248/qdsb`. We ran the experiments on a machine with an AMD Ryzen 9 7900X CPU and an NVIDIA GeForce RTX 4090 GPU. The toy and single-cell benchmarks were run on CPU, whereas the FFHQ latent-space experiment was run on GPU. For the toy, single-cell, and sensitivity results we report averages over five seeds $\{0, 1, 2, 3, 4\}$, while the FFHQ image experiment is qualitative and uses a single seed $\{0\}$. Whenever time-quality curves are recorded, elapsed time is measured with evaluation overhead removed, so Figures 2 and 3 reflect training time rather than checkpoint time. In the toy suite, the curves are averaged over seeds, and in the single-cell benchmarks they are averaged over seeds and leave-one-out tasks. For the quantitative experiments, quality is measured by MMD with an RBF kernel whose bandwidth is chosen by the median heuristic on 4096 reference points. For the toy and single-cell benchmarks, final MMD is computed on the full evaluation populations. All experiments were run sequentially to avoid interference. Running all experiments takes about one week on the hardware described above.

Model architecture and optimization. Across all methods, we use the same time-dependent four-layer MLP backbone. The layers are $d - 64 - 64 - d$, where d is the input dimension and the hidden layers use a width of 64. We use the AdamW optimizer with learning rate 10^{-4} and weight decay 10^{-2}, batch size 256, and rollout batch size 2048. LightSB-M uses its standard Gaussian-mixture parameterization. In the FFHQ latent experiment, all methods use the AdamW optimizer with learning rate 10^{-3} and weight decay 10^{-2}, batch size 512, and rollout batch size 2048.

2D toy experiments. The three tasks are 8Gaussians→Moons, Gaussian→Moons, and Gaussian→8Gaussians. Each endpoint distribution provides 16384 training samples and 4096 evaluation samples. In the main toy benchmark, all methods are trained for 500 epochs and MMD is recorded every epoch. We use an exact numerical OT solver with squared Euclidean cost, $\sigma = 0.25$, and 100 Euler-Maruyama steps per unit time. Our approach uses 256 anchors per marginal, farthest-first anchor selection, anchor refresh every 100 epochs, and random sampling within matched anchor cells. SF2M+mPOT differs from SF2M only in the endpoint coupling and uses a partial transport fraction of 0.8. LightSB-M uses 128 Gaussian components and a minimum covariance scale of 10^{-2}. For DSB and DSBM, we choose the inner iteration count automatically so that the total number of parameter updates matches the SF2M update budget. In the reported toy runs, DSB uses 20 IPF rounds, 20 Langevin steps, and constant step size $\gamma = 0.01$. DSBM uses 20 outer iterations, $\sigma = 0.25$, 100 rollout steps per unit time, and loss weighting. The sensitivity experiment in Figure 3 keeps the QDSB setup fixed, varies the number of anchors from 1 to 1024, and trains for 1000 epochs for each of the anchor counts.

Single-cell experiments. For EB, we use the five observed timepoints with nominal times $0, 1, 2, 3, 4$ and leave out the three intermediate timepoints. For Cite and Multi, we use donor 13176 at days $2, 3, 4, 7$ and leave out days 3 and 4. All methods operate on a 5-dimensional representation obtained by fitting an incremental PCA with 100 components on the raw data, whitening the resulting representation, and then truncating to five dimensions, following the setup of Tong et al. (2020). All single-cell methods are trained for 1000 epochs and MMD is recorded every epoch. We use an exact numerical OT solver with squared Euclidean cost, discrete model times, and 100 Euler-Maruyama steps per unit time. The bridge noise values are 0.25 for EB and Cite, and 1.0 for Multi. Our approach uses 256 anchors per marginal, farthest-first anchor selection, anchor refresh every 100 epochs, and random endpoint sampling within cells. SF2M+mPOT again differs from SF2M only in the coupling sampler and uses a partial transport fraction of 0.8. LightSB-M uses 512 Gaussian components and a minimum covariance scale of 10^{-2}. For DSB and DSBM, we again choose the inner iteration count automatically so that the total number of parameter updates matches the SF2M update budget. In the reported single-cell runs, DSB uses 20 IPF rounds, 20 Langevin steps, and a constant step size of $\gamma = 0.01$. DSBM uses 20 outer iterations, 20 rollout steps per unit time, and loss weighting.

Image translation experiment. We use the FFHQ data split into adult age groups 20–29, 30–39, and 40–49 and child age groups 0–2, 3–6, and 7–9, with evaluation fraction 0.1 and split seed 0. The images are encoded using a pretrained ALAE checkpoint (`model_157.pth`), which gives a 512-dimensional latent representation. All reported image runs use unpaired adult and child latents, a fixed wall-clock budget of 1000 s, and 50 decoded examples per method. All methods use the same 512-dimensional MLP backbone described above. Our approach uses 512 anchors per marginal, no anchor refresh, and random endpoint sampling within cells. LightSB-M uses 10 Gaussian components and a minimum covariance scale of 10^{-2}. We also ran DSB and DSBM in this setting, but they did not produce usable image translations and are therefore omitted from the qualitative figures.

For completeness, the image-space DSB configuration uses 50 inner iterations, 20 IPF rounds, 20 Langevin steps, and a constant step size of $\gamma = 0.01$. The image-space DSBM configuration uses hidden dimension 512, 50 inner iterations, 20 outer iterations, 100 rollout steps per unit time, and loss weighting.