VARIATIONAL ENTROPIC OPTIMAL TRANSPORT

Roman Dyachenko[1] **Nikita Gushchin**[2,3] **Kirill Sokolov**[4] **Petr Mokrov**[2]
Evgeny Burnaev [2,3] **Alexander Korotin** [2,3]

[1]Higher School of Economics, Moscow, Russia
[2]Applied AI Institute, Moscow, Russia
[3]AXXX, Russia
[4]Lomonosov Moscow State University, Moscow, Russia

rrdiachenko@edu.hse.ru, i.nikita.gushchin@gmail.com

ABSTRACT

Entropic optimal transport (EOT) in continuous spaces with quadratic cost is a classical tool for solving the domain translation problem. In practice, recent approaches optimize a weak dual EOT objective depending on a single potential, but doing so is computationally not efficient due to the intractable log-partition term. Existing methods typically resolve this obstacle in one of two ways: by significantly restricting the transport family to obtain closed-form normalization (via Gaussian-mixture parameterizations), or by using general neural parameterizations that require simulation-based training procedures. We propose Variational Entropic Optimal Transport (VarEOT), based on an exact variational reformulation of the log-partition $\log \mathbb{E}[\exp(\cdot)]$ as a tractable minimization over an auxiliary positive normalizer. This yields a differentiable learning objective optimized with stochastic gradients and avoids the necessity of MCMC simulations during the training. We provide theoretical guarantees, including finite-sample generalization bounds and approximation results under universal function approximation. Experiments on synthetic data and unpaired image-to-image translation demonstrate competitive or improved translation quality, while comparisons within the solvers using the same weak dual EOT objective show the benefit of the proposed optimization principle.

1 INTRODUCTION

Entropic Optimal Transport (EOT) with quadratic cost is a well-established mathematical framework with strong theoretical properties which found wide application in generative modeling and especially for unpaired domain translation. Despite this, the practical adoption of entropic transport methods has been limited by the lack of efficient and flexible algorithms. Existing approaches typically suffer from at least one of the following drawbacks: they are not simulation-free and rely on costly sampling during training (Mokrov et al., 2024); they require adversarial optimization (Gushchin et al., 2024b); they involve training a sequence of models rather than a single objective (Shi et al., 2023); they impose restrictive parametric forms on the transport plan (Korotin et al., 2024); or they turn out to be too sensitive to the entropic regularization strength (Daniels et al., 2021).

In our paper, we make a decisive step towards solving the drawbacks of existing EOT methods, and propose a novel simulation-free solver based on an innovative variational reformulation of the weak dual EOT objective. We present the following **main contributions**:

1. **Variational dual objective** (§3.1). We derive an equivalent reformulation of the weak dual objective of entropic OT with quadratic cost in which the intractable log-partition term is replaced by an exact variational minimization over an auxiliary positive normalizer.

2. **Simulation-free training solver** (§3.2). Building on this reformulation, we propose a *simulation-free* variational solver that jointly learns the dual potential and the auxiliary normalizer, enabling fully simulation-free training with neural parameterizations (no MCMC).

3. **Learning guarantees** (§3.3). We provide finite-sample learning guarantees for recovery of the entropic OT plan, decomposing error into estimation and approximation terms, and show vanishing approximation error under universal function approximation.
4. **Evaluation** (§5). We evaluate our solver on synthetic data and unpaired image-to-image translation. We highlight gains by comparing against solvers optimizing the same weak dual objective.

2 BACKGROUND

This section provides the necessary background on entropic optimal transport and its optimization using the weak dual reformulation. In §2.1, we recall the entropic optimal transport problem with quadratic cost, introduce its weak dual formulation, and describe the structure of the optimal transport plan. In §2.2, we clarify our learning setup. In §2.3, we review existing weak dual solvers.

2.1 ENTROPIC OPTIMAL TRANSPORT WITH THE QUADRATIC COST

Let $p_0, p_1 \in \mathcal{P}_{\mathrm{ac}}(\mathbb{R}^D)$ be absolutely continuous Borel probability measures on $\mathbb{R}^D$. Let $\Pi(p_0, p_1)$ denote the set of couplings (transport plans) on $\mathbb{R}^D \times \mathbb{R}^D$ with marginals p_0 and p_1. We write $\pi(x_0, x_1)$ for a plan density; H is the differential entropy. The Entropic Optimal Transport (EOT) problem is given by:

$$\mathrm{EOT}_\varepsilon(p_0, p_1) \stackrel{\text{def}}{=} \min_{\pi \in \Pi(p_0,p_1)} \left\{ \mathop{\mathbb{E}}_{\pi(x_0,x_1)} \left[\frac{\|x_0 - x_1\|^2}{2} \right] - \varepsilon \int_{\mathbb{R}^D} H\big(\pi(\cdot \mid x_0)\big) p_0(x_0)\, dx_0 \right\}. \tag{1}$$

The entropic regularization term in (1) is due to (Mokrov et al., 2024). Equivalent forms exist (Cuturi, 2013; Léonard, 2014) that differ only by additive constants not affecting the solution.

The EOT problem admits a unique minimizer π^*, referred to as the EOT plan. While the primal formulation (1) is conceptually appealing, it is computationally inconvenient, since enforcing the marginal constraints $\pi \in \Pi(p_0, p_1)$ requires optimizing over a complex set of probability measures.

Weak dual form of EOT. Objective (1) admits the following weak dual representation (Mokrov et al., 2024, Theorem 1):

$$\sup_f \underbrace{\left\{ \mathop{\mathbb{E}}_{p_1(x_1)} f(x_1) - \varepsilon \mathop{\mathbb{E}}_{p_0(x_0)} \log Z(f, x_0) \right\}}_{\stackrel{\text{def}}{=} \mathcal{L}(f)}, \tag{2}$$

the $\sup$ is taken over integrable functions $f : \mathbb{R}^D \to \mathbb{R}$ and

$$Z(f, x_0) \stackrel{\text{def}}{=} \int_{\mathbb{R}^D} \exp\left(\frac{f(x_1) - \frac{1}{2}\|x_0 - x_1\|^2}{\varepsilon} \right) dx_1 \tag{3}$$

is the partition function.

Optimal transport plan. Let f^* be an optimizer of (2), the corresponding optimal transport plan π^* (Mokrov et al., 2024, Theorem 1) can be recovered from it. By the disintegration with respect to the source marginal p_0 we have $\pi^*(x_0, x_1) = \pi^*(x_1 \mid x_0)\, p_0(x_0)$, where:

$$\pi^*(x_1 \,|\, x_0) = \frac{1}{Z(f^*, x_0)} \exp\left(\frac{f^*(x_1) - \frac{1}{2}\|x_0 - x_1\|^2}{\varepsilon} \right). \tag{4}$$

2.2 COMPUTATIONAL EOT SETUP

In practice, the source and target distributions, p_0, p_1, as well as the EOT objective (1), could be expressed and treated in different ways. To avoid possible misunderstanding, below we formalize our **practical learning setup**:

We assume that source and target distributions p_0 and p_1 are accessible only by a limited number of i.i.d. empirical samples (datasets) $\{x_0^1, x_0^2, \dots x_0^N\} \sim p_0$; $\{x_1^1, x_1^2, \dots x_1^M\} \sim p_1$. Our aim is to approximate the optimal conditional plan $\pi^*(\cdot|x_0)$ (eq. (4)) between entire *distributions* p_0 and p_1. The recovered solution should provide the *out-of-sample* estimation, i.e., allow generating samples from $\pi^*(\cdot|x_0^{\text{new}})$, where x_0^{new} is a new sample from p_0 which is not necessarily present in the train dataset.

This setup falls within **continuous** OT, in contrast to discrete OT (Cuturi, 2013; Peyré et al., 2019), which computes correspondence directly between provided source and target samples and does not naturally accommodate out-of-sample estimation.

2.3 Existing Weak Dual Formulation Solvers

The practical optimization of problem (2) remains challenging due to the partition function $Z(f, x_0)$, which is intractable to compute exactly. Below we review two representative approaches.

EgNOT solver. The authors of EgNOT (Mokrov et al., 2024) solve (2) by parametrizing f_θ by a neural network and deriving the gradient of the weak dual objective $\mathcal{L}(f_\theta)$:

$$\nabla_\theta \mathcal{L}(f_\theta) = \mathop{\mathbb{E}}_{p_1(x_1)} [\nabla_\theta f_\theta(x_1)] - \mathop{\mathbb{E}}_{p_0(x_0)} \left[\mathop{\mathbb{E}}_{\pi_\theta(x_1|x_0)} [\nabla_\theta f_\theta(x_1)] \right], \tag{5}$$

where $\pi_\theta(x_1 \mid x_0) = \frac{1}{Z(f_\theta, x_0)} \exp\left(\frac{f_\theta(x_1) - \frac{1}{2}\|x_0 - x_1\|^2}{\varepsilon} \right)$. While flexible, **this approach is *not simulation-free***: each evaluation of the loss or its gradient requires sampling from the model distribution $\pi_\theta(x_1 \mid x_0)$ via MCMC methods (Girolami & Calderhead, 2011; Hoffman et al., 2014; Samsonov et al., 2022), which can be computationally expensive and sensitive to hyperparameters.

LightSB solver. An alternative strategy is proposed in LightSB (Korotin et al., 2024). The authors introduce adjusted potential v_θ and parametrization:

$$\pi_\theta(x_1|x_0) = \frac{\exp\left(\langle x_0, x_1 \rangle / \varepsilon\right) v_\theta(x_1)}{c_\theta(x_0)}, \tag{6}$$

where $c_\theta(x_0) \stackrel{\text{def}}{=} \int_{\mathbb{R}^D} \exp\left(\langle x_0, x_1 \rangle / \varepsilon\right) v_\theta(x_1) dx_1$ is the normalization. To circumvent the intractability of $c_\theta(x_0)$, the authors parameterize $v_\theta(x_1)$ as a Gaussian mixture $v_\theta(x_1) = \sum_{k=1}^{K'} \alpha_k \mathcal{N}(x_1|r_k, \varepsilon S_k)$, yielding a closed-form normalization. While this leads to a fully tractable and simulation-free objective, it **restricts the expressiveness of the method**.

Summary. Existing weak dual solvers trade off expressiveness and tractability: EgNOT supports flexible potentials but requires MCMC during training, while LightSB is simulation-free but restricts the conditional plan family. Below, we present VarEOT, which takes the **best of both worlds** by enabling simulation-free training *without* restricting $\pi(x_1 \,|\, x_0)$ to a narrow parametric family.

3 Variational Entropic Optimal Transport

In this section, we introduce our variational approach to entropic optimal transport. In §3.1, we derive a new variational dual formulation. In §3.2, we present a practical learning algorithm. In §3.3, we provide learning guarantees. All proofs are in Appendix A.

3.1 New Variational Dual Formulation of EOT

Our goal is to propose a weak dual solver that, unlike EgNOT and LightSB, does not require simulation during training and allows for expressive parameterization. A key challenge is differentiating through the partition function $\log Z(f, x_0)$: as a logarithm of an expectation, it cannot be unbiasedly estimated from finite samples in a straightforward way. To overcome this, we adopt a variational approximation for the logarithm.

Proposition 3.1 (Variational bound for the partition function)**.** *The logarithm of partition function* $\log Z(f, x_0)$ *admits the variational upper bound:*

$$\log Z(f, x_0) \leq -1 + \xi(x_0) + \frac{D}{2} \log(2\pi\varepsilon) + \mathop{\mathbb{E}}_{z \sim \mathcal{N}(0, I)} \left[\exp\Big(\frac{f(x_0 + \sqrt{\varepsilon} z)}{\varepsilon} - \xi(x_0) \Big) \right], \tag{7}$$

where $\xi : \mathbb{R}^D \to \mathbb{R}$ is an arbitrary integrable function. The upper bound is tight when

$$\xi^*(x_0) = \log Z(f, x_0) - \frac{D}{2}\log(2\pi\varepsilon). \tag{8}$$

Thanks to the obtained estimate, we can obtain a tractable simulation-free loss:

Theorem 3.2 (Variational dual form of EOT)**.** *Let*

$$\mathcal{L}(f,\xi) \stackrel{def}{=} \varepsilon\Big(1 - \frac{D}{2}\log(2\pi\varepsilon)\Big) + \mathop{\mathbb{E}}_{x_1\sim p_1}[f(x_1)] - \varepsilon \mathop{\mathbb{E}}_{x_0\sim p_0}[\xi(x_0)] - \varepsilon \mathop{\mathbb{E}}_{x_0\sim p_0}\Big[\mathop{\mathbb{E}}_{z\sim\mathcal{N}(0,I)} \exp\Big(\frac{f(x_0+\sqrt{\varepsilon}z)}{\varepsilon} - \xi(x_0)\Big)\Big]. \tag{9}$$

Then the entropic optimal transport weak dual formulation admits the following variational form:

$$\mathrm{EOT}_\varepsilon(p_0, p_1) = \sup_f \mathcal{L}(f) = \sup_{f,\xi} \mathcal{L}(f,\xi). \tag{10}$$

The optimal solution (f^*, ξ^*), where ξ^* is set by (8)*, recovers*

$$\pi^*(x_1|x_0) = \frac{(2\pi\varepsilon)^{-\frac{D}{2}}}{\exp(\xi^*(x_0))}\exp\left(\frac{f^*(x_1) - \frac{1}{2}\|x_0 - x_1\|^2}{\varepsilon}\right).$$

This novel variational dual form overcomes the original problem of estimating the log partition function $\log Z(f, x_0)$ in the weak dual form (2).

For convenience we define:

$$\pi^{f,\xi}(x_0,x_1) \stackrel{\text{def}}{=} \frac{(2\pi\varepsilon)^{-\frac{D}{2}}p_0(x_0)}{\exp(\xi(x_0))}\exp\left[\frac{f(x_1) - \frac{1}{2}\|x_0-x_1\|^2}{\varepsilon}\right];$$
$$\pi^f(x_0,x_1) \stackrel{\text{def}}{=} \pi^{f,\xi^*}(x_0,x_1).$$

The following theorem establishes that the gap between the optimal and current VarEOT objective value directly corresponds to the KL discrepancy between the recovered measure and the EOT plan.

Theorem 3.3. *For any measurable functions f and ξ,*

$$\varepsilon\mathrm{KL}\big(\pi^* \,\big\|\, \pi^f\big) \le \varepsilon\,\mathrm{KL}\big(\pi^* \,\big\|\, \pi^{f,\xi}\big) = \mathcal{L}^* - \mathcal{L}(f,\xi), \tag{11}$$

where $\mathrm{KL}(\cdot\|\cdot)$ denotes the KL divergence between non-negative measures (see Definition A.1 in Appendix A), and $\mathcal{L}^*$ is the optimal value of weak dual EOT objective (9)*.*

Theorem 3.3 certifies that optimizing objective (9) directly enables us to approximate the ground truth EOT plan. Additionally, eq. (11) suggests that at inference it is better to use π^f, not $\pi^{f,\xi}$, which we exploit in our practical implementation.

3.2 Practical Algorithm

Training. We parametrize both the potential f and the auxiliary variational function ξ by neural networks f_θ and ξ_ψ. Given mini-batches $\{x_i^0\}_{i=1}^{N_0} \sim p_0$ and $\{x_j^1\}_{j=1}^{N_1} \sim p_1$, the expectations in eq. (9) are approximated via Monte Carlo with i.i.d. Gaussian noise $z_{i,k} \sim \mathcal{N}(0, I)$, yielding the empirical loss (omitting the additive constant):

$$\widehat{\mathcal{L}}(f_\theta, \xi_\psi) \stackrel{\text{def}}{=} \frac{1}{N_1}\sum_{j=1}^{N_1} f_\theta(x_j^1) - \varepsilon\frac{1}{N_0}\sum_{i=1}^{N_0}\xi_\psi(x_i^0) - \frac{\varepsilon}{N_0 K}\sum_{i,k}\exp\left(\frac{f_\theta(x_i^0 + \sqrt{\varepsilon}\,z_{i,k})}{\varepsilon} - \xi_\psi(x_i^0)\right). \tag{12}$$

Importantly, **this approximation does not require sampling from the model distribution itself**, in contrast to energy-based approaches like EgNOT. The training is summarized in Algorithm 1.

Algorithm 1: Training procedure for Variational Entropic Optimal Transport (VarEOT)

Input : samples from distributions p_0 and p_1;
entropy regularization parameter $\varepsilon > 0$;
batch sizes N_0, N_1; number of noise samples K;
potential network $\hat{f}_\theta : \mathbb{R}^D \to \mathbb{R}$;
auxiliary network $\hat{\xi}_\psi : \mathbb{R}^D \to \mathbb{R}$.

Output : trained potential $\hat{f}_{\theta^*}$.

for *each training iteration* **do**

Sample mini-batches $\{x_i^0\}_{i=1}^{N_0} \sim p_0$,;
$\{x_j^1\}_{j=1}^{N_1} \sim p_1$;
Sample i.i.d. noise variables $z_{i,k} \sim \mathcal{N}(0, I)$ for $i = 1, \dots, N_0$, $k = 1, \dots, K$;
Compute the empirical loss $\hat{\mathcal{L}}$ according to eq. (12);
Update ψ, θ by using the gradients $\nabla_\psi \hat{\mathcal{L}}, \nabla_\theta \hat{\mathcal{L}}$;

Algorithm 2: Langevin sampling from VarEOT

Input : $x_0 \sim p_0$; trained f_θ; $\varepsilon > 0$; steps S; step size $\eta > 0$.

Output : sample $x_1 \sim \pi(\cdot \mid x_0)$.

Initialize $x_1^{(0)} = x_0$;

for $s = 1$ **to** S **do**

Sample $z^{(s)} \sim \mathcal{N}(0, I)$;
$h = \nabla_{x_1} f_\theta(x_1^{(s-1)}) - \frac{1}{\varepsilon}(x_1^{(s-1)} - x_0)$;
$x_1^{(s)} \leftarrow x_1^{(s-1)} + \eta h + \sqrt{2\eta}\, z^{(s)}$;

return $x_1^{(S)}$

Inference. Following eq. (11), after training we only use f_θ for inference. The conditional transport plan is implicitly defined via eq. (4). To generate samples, we employ Langevin dynamics targeting $\pi(x_1 \mid x_0) \propto \exp\big(f_\theta(x_1) - \frac{1}{2\varepsilon}\|x_1 - x_0\|^2\big)$. The sampling procedure is detailed in Algorithm 2. In all experiments, 10^1–10^3 Langevin steps with proper step size suffice for high-quality samples.

Method comparison. Table 1 compares VarEOT to EgNOT and LightSB. VarEOT is *simulation-free* (unlike EgNOT) and does not impose a *restricted parameterization* (unlike LightSB).

Table 1: Comparison of dual solvers for entropic OT.

Method	Simulation-free training	Not restricted parameterization
EgNOT	×	✓
LightSB	✓	×
VarEOT **(ours)**	✓	✓

3.3 Finite Sample Learning Guarantees

We quantify the discrepancy between the transport plan recovered by VarEOT and the ground truth EOT solution. Our method works with finite samples and potentials restricted to parametric classes, introducing several sources of error: *finite-sample error*, *function class restriction*, and *optimization error*. Our theoretical analysis focuses on the first two.

In the results below, fix an activation $\sigma : \mathbb{R} \to \mathbb{R}$ that is continuous and non-polynomial (e.g., ReLU, GELU). For each $n \geq 1$, let

$$\mathcal{F}_n := \mathsf{NN}_n^\sigma(D, 1), \quad \Xi_n := \mathsf{NN}_n^\sigma(D, 1), \tag{13}$$

denote neural-network classes of increasing capacity. We often write $\mathcal{F}$ and Ξ, suppressing index n.

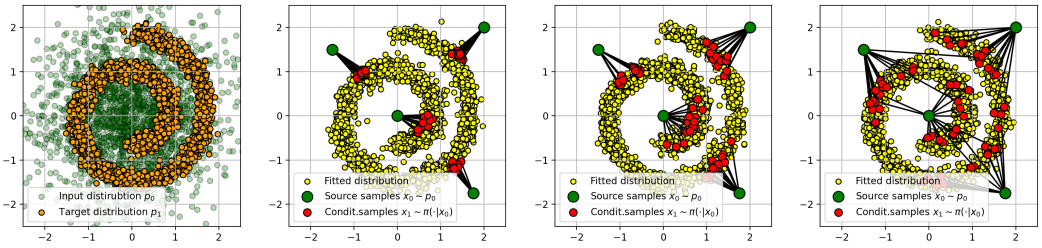

(a) $x_0 \sim p_0$, $x_1 \sim p_1$. (b) $\varepsilon = 0.01$. (c) $\varepsilon = 0.1$. (d) $\varepsilon = 1.0$.

Figure 1: Optimal plan learned with VarEOT **(ours)** in *Gaussian→Swiss roll* example.

Proposition 3.4. *The following bound holds:*

$$\varepsilon\mathbb{E}\Big[\mathrm{KL}\big(\pi^* \,\|\, \pi^{\hat{f}}\big)\Big] \leq \underbrace{2\mathbb{E} \sup_{f\in\mathcal{F},\, \xi\in\Xi} \Big|\mathcal{L}(f,\xi) - \widehat{\mathcal{L}}(f,\xi)\Big|}_{\textit{Estimation error}} + \underbrace{\mathcal{L}^* - \sup_{f\in\mathcal{F},\xi\in\Xi} \mathcal{L}(f,\xi)}_{\textit{Approximation error}}, \tag{14}$$

where $\hat{f} = \arg\max_{f\in\mathcal{F}} \widehat{\mathcal{L}}\big(f, Z(f,\cdot) - \frac{D}{2}\log(2\pi\varepsilon)\big)$.

Theorem 3.5 (Bound on estimation error)**.** *Assume that* $\mathcal{F} \subset \mathrm{Lip}_{L_f,M_f}(X_1)$ *and* $\Xi \subset \mathrm{Lip}_{L_\xi,M_\xi}(X_0)$. *Then*

$$\mathbb{E} \sup_{f\in\mathcal{F},\, \xi\in\Xi} \Big|\mathcal{L}(f,\xi) - \widehat{\mathcal{L}}(f,\xi)\Big| \leq O\Big(N^{-\frac{1}{1+D}}\Big) + O\Big(M^{-\frac{1}{1+D}}\Big) \tag{15}$$

where hidden constants depend on $D, L_f, L_\xi, M_f, M_\xi$ *and diameters of* X_0, X_1.

Remark 3.6. Assuming Lipschitz regularity is well motivated: neural networks with Lipschitz activations are Lipschitz when weight norms are controlled (e.g., via spectral normalization).

Theorem 3.7 (Vanishing approximation error)**.** *Let the sequences* $\{\mathcal{F}_n\}_{n\geq 1}$ *and* $\{\Xi_n\}_{n\geq 1}$ *be as in* (13). *Then for any* $\delta > 0$ *there exists* $n \in \mathbb{N}$ *such that* $\mathcal{L}^* - \sup_{f\in\mathcal{F}_n,\, \xi\in\Xi_n} \mathcal{L}(f,\xi) < \delta$.

Summary. With enough data (large N, M) the estimation gap becomes small, and with sufficiently expressive neural networks the approximation error becomes small.

4 Related Work

We review existing continuous EOT approaches and briefly compare them with our solver.

Weak dual EOT (a.k.a. *semi*-dual EOT) methods (Mokrov et al., 2024; Korotin et al., 2024) optimize eq. (2) and are the most similar to VarEOT. The detailed discussion is given in §2.3. Our method overcomes their limitations while retaining their advantages.

Dual EOT methods (Genevay et al., 2016; Seguy et al., 2018; Daniels et al., 2021) optimize a *pair* of dual potentials (u, v) via the Sinkhorn algorithm. While bearing certain resemblance to VarEOT (e.g., simulation-free training), these approaches have notable pitfalls: (i) dual potentials do not fully recover $\pi^*(\cdot|x_0)$ without an auxiliary score-based model (Daniels et al., 2021) or heuristics such as barycentric projection (Genevay et al., 2016; Seguy et al., 2018); (ii) they are unstable under small ε (Daniels et al., 2021, §5.1). VarEOT provides a more user-friendly framework with less engineering.

Schrödinger bridge (SB) methods (De Bortoli et al., 2021; Vargas et al., 2021; Chen et al., 2022; Gushchin et al., 2023; Shi et al., 2023; Tong et al., 2024; Gushchin et al., 2024b; De Bortoli et al., 2024; Gushchin et al., 2024a) cast EOT dynamically and recover a solution as a stochastic differential equation. Compared to VarEOT, the **majority** are not simulation-free at training. Exceptions are (Tong et al., 2024) (based on mini-batch EOT approximation) and (Gushchin et al., 2024a) (restricted Gaussian mixture approximation).

5 Experimental Illustrations

We evaluate VarEOT on both synthetic and real-world data. Section 5.1 presents two-dimensional experiments and Section 5.2 provides evaluation on unpaired image-to-image translation. Technical details are in Appendix B.

(a) Qualitative comparison for *Male* $\rightarrow$ *Female* translation with $\varepsilon = 1.0$.

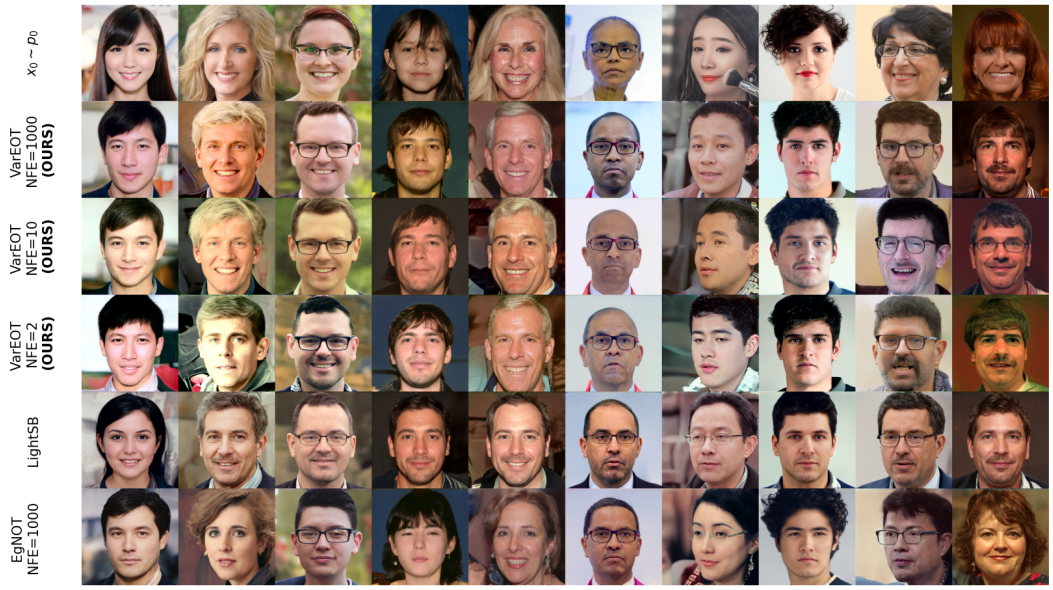

(b) Qualitative comparison for *Female* $\rightarrow$ *Male* translation with $\varepsilon = 1.0$.

Figure 2: Qualitative comparisons for image-to-image translation tasks. Top to bottom: input, VarEOT (ours), LightSB, EgNOT. Input images are ranked by encoder-decoder reconstruction quality (LPIPS), showing top-ranked examples.

5.1 Two-Dimensional Examples

We consider transporting a Gaussian source to a Swiss Roll target for $\varepsilon \in \{10^{-2}, 10^{-1}, 10^{0}\}$ (Figure 1). For small ε, the transport is nearly deterministic. As ε increases, the transport becomes progressively more stochastic with broader conditional distributions.

5.2 Unpaired Image-to-Image Translation

Unpaired image-to-image translation is a standard benchmark in EOT and Schrödinger Bridge literature (Zhu et al., 2017; Daniels et al., 2021; Chen et al., 2021). We follow the widely adopted ALAE protocol (Korotin et al., 2024; Theodoropoulos et al., 2024; Gushchin et al., 2024a)

ε	Task	FID ↓						LPIPS ↓					
		LightSB	EgNOT		VarEOT **(ours)**			LightSB	EgNOT		VarEOT **(ours)**		
			NFE=10	NFE=1000	NFE=2	NFE=10	NFE=1000		NFE=10	NFE=1000	NFE=2	NFE=10	NFE=1000
0.1	M→F	8.77	13.4	**7.63**	181.44	11.70	12.74	0.582	0.6	0.585	0.872	0.580	**0.570**
	F→M	10.87	14.9	**6.33**	248.76	17.00	13.98	0.597	0.603	0.598	0.833	0.586	**0.578**
	A→C	15.42	40.28	**10.52**	247.81	50.35	17.33	0.577	0.586	0.588	0.839	**0.545**	0.577
	C→A	13.46	18.47	**10.4**	253.59	21.48	20.13	0.594	0.602	0.608	0.865	0.591	**0.590**
0.5	M→F	19.09	35.4	14.07	**9.04**	9.50	9.77	0.611	0.718	0.594	0.619	0.600	**0.598**
	F→M	25.81	51.73	11.66	20.01	10.02	**9.94**	0.628	0.695	0.613	**0.609**	0.619	0.616
	A→Y	22.30	74.25	**18**	78.85	26.38	25.99	0.608	0.702	0.603	**0.593**	0.617	0.613
	C→A	22.70	40.38	16.95	17.83	**16.55**	16.74	**0.614**	0.7	0.629	0.632	0.616	**0.614**
1.0	M→F	22.63	39.41	17.58	10.38	**9.59**	9.71	0.637	0.749	0.634	0.633	0.613	**0.610**
	F→M	20.97	66.64	26.21	15.21	10.52	**10.28**	0.649	0.72	0.634	0.654	0.633	**0.630**
	A→C	24.43	60.57	28.45	53.39	**16.60**	16.83	0.634	0.724	0.623	0.700	0.610	**0.608**
	C→A	24.87	51.48	23.68	**15.73**	16.55	16.70	0.637	0.733	0.661	0.638	0.623	**0.619**
10.0	M→F	31.85	37.61	26.68	**21.84**	21.97	22.38	0.680	0.755	0.735	0.625	0.621	**0.620**
	F→M	34.47	63.73	36	**25.55**	28.84	29.24	0.693	0.758	0.734	0.637	0.632	**0.630**
	A→C	**31.85**	92.6	64.85	49.84	50.72	51.20	0.680	0.767	0.763	0.619	**0.615**	**0.615**
	C→A	35.35	51.74	35.53	**31.01**	32.84	34.07	0.683	0.748	0.741	0.639	0.636	**0.633**

Table 2: Quantitative comparison in the ALAE latent space. FID and LPIPS (lower is better) for four tasks across ε values. VarEOT results shown for different Langevin inference steps (NFE). Best per row in **bold**.

(Kornilov et al., 2024; Gazdieva et al., 2024; Han et al., 2025) based on the ALAE autoencoder (Pidhorskyi et al., 2020) trained on 1024×1024 FFHQ images (Karras et al., 2019). The first 60K images are used for training, split into (*male*, *female*) and (*child*, *adult*) subsets. We extract 512-dimensional latent representations and consider 4 setups: M→F, F→M, A→C, and C→A.

Training. We learn a latent entropic optimal transport plan $\pi_\theta(z_1 \mid z_0)$ using our variational formulation, trained entirely in latent space without paired data.

Inference. To translate a test image x_0^{new}, we (i) encode $z_0^{\text{new}} = \text{Enc}(x_0^{\text{new}})$, (ii) sample $z_1 \sim \pi_\theta(z_1 \mid z_0^{\text{new}})$ via Langevin dynamics, and (iii) decode $x_1 = \text{Dec}(z_1)$.

Evaluation. We use FID (Heusel et al., 2017) and LPIPS (Zhang et al., 2018). FID is computed between translated and ALAE-reconstructed target images. LPIPS measures input-output similarity.

Results. Figure 2 shows qualitative results for M→F and F→M with $\varepsilon = 1.0$. Quantitative results are in Table 2, covering $\varepsilon \in \{0.1, 0.5, 1.0, 10.0\}$ across all four translation directions with different Langevin inference steps. VarEOT demonstrates competitive or improved performance across regimes. Additional qualitative results for other directions are in Appendix C. The effect of ε on diversity is in Figure 5, and FID sensitivity to Langevin parameters is in Figure 6.

6 Discussion

Potential impact. By introducing an exact variational reformulation of the weak dual EOT objective, VarEOT makes a step toward more efficient entropic transport algorithms that avoid key limitations of existing methods, such as simulation-based training, adversarial objectives, and restrictive parametric assumptions.

Limitations. While VarEOT enables simulation-free training, sampling at inference time via Langevin dynamics is still required, making generation quality dependent on step size and the number of steps. Additionally, the exponential terms in the objective may cause numerical instability if not handled carefully.

Reproducibility Statement

Technical details including architecture specifications, hyperparameters, and dataset preprocessing are provided in Appendix B. All proofs are provided in Appendix A.

REFERENCES

Tianrong Chen, Guan-Horng Liu, and Evangelos Theodorou. Likelihood training of schrödinger bridge using forward-backward SDEs theory. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=nioAdKCEdXB.

Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Stochastic control liaisons: Richard sinkhorn meets gaspard monge on a schrodinger bridge. *SIAM Review*, 63(2):249–313, 2021.

Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.

Max Daniels, Tyler Maunu, and Paul Hand. Score-based generative neural networks for large-scale optimal transport. *Advances in neural information processing systems*, 34:12955–12965, 2021.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in neural information processing systems*, 34:17695–17709, 2021.

Valentin De Bortoli, Iryna Korshunova, Andriy Mnih, and Arnaud Doucet. Schrodinger bridge flow for unpaired data translation. *Advances in Neural Information Processing Systems*, 37: 103384–103441, 2024.

Milena Gazdieva, Arip Asadulaev, Evgeny Burnaev, and Alexander Korotin. Light unbalanced optimal transport. *Advances in Neural Information Processing Systems*, 37:93907–93938, 2024.

Aude Genevay, Marco Cuturi, Gabriel Peyré, and Francis Bach. Stochastic optimization for large-scale optimal transport. *Advances in neural information processing systems*, 29, 2016.

Mark Girolami and Ben Calderhead. Riemann manifold langevin and hamiltonian monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2): 123–214, 2011.

Lee-Ad Gottlieb, Aryeh Kontorovich, and Robert Krauthgamer. Adaptive metric dimensionality reduction. *Theoretical Computer Science*, 620:105–118, 2016. doi: 10.1016/j.tcs.2015.10.040. Full version available as arXiv:1302.2752.

Nikita Gushchin, Alexander Kolesov, Alexander Korotin, Dmitry P Vetrov, and Evgeny Burnaev. Entropic neural optimal transport via diffusion processes. *Advances in Neural Information Processing Systems*, 36:75517–75544, 2023.

Nikita Gushchin, Sergei Kholkin, Evgeny Burnaev, and Alexander Korotin. Light and optimal schrödinger bridge matching. In *Forty-first International Conference on Machine Learning*, 2024a.

Nikita Gushchin, Daniil Selikhanovych, Sergei Kholkin, Evgeny Burnaev, and Aleksandr Korotin. Adversarial schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 37: 89612–89651, 2024b.

Dong-Sig Han, Jaein Kim, Hee Bin Yoo, and Byoung-Tak Zhang. Variational online mirror descent for robust learning in schr\" odinger bridge. *arXiv preprint arXiv:2504.02618*, 2025.

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

Matthew D Hoffman, Andrew Gelman, et al. The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *J. Mach. Learn. Res.*, 15(1):1593–1623, 2014.

Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4401–4410, 2019.

Alexander Kolesov, Petr Mokrov, Igor Udovichenko, Milena Gazdieva, Gudmund Pammer, Anastasis Kratsios, Evgeny Burnaev, and Alexander Korotin. Energy-guided continuous entropic barycenter estimation for general costs. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*, pp. 107513–107546. Neural Information Processing Systems Foundation, Inc., 2024.

Nikita Kornilov, Petr Mokrov, Alexander Gasnikov, and Aleksandr Korotin. Optimal flow matching: Learning straight trajectories in just one step. *Advances in Neural Information Processing Systems*, 37:104180–104204, 2024.

Alexander Korotin, Nikita Gushchin, and Evgeny Burnaev. Light schrödinger bridge. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=WhZoCLRWYJ.

Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*, volume 23 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3)*. Springer, 1991.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

Christian Léonard. A survey of the schrödinger problem and some of its connections with optimal transport, 2014. ISSN 1078-0947. URL https://www.aimsciences.org/article/id/d5bcf817-901d-4104-b7da-eade7847c53e.

Petr Mokrov, Alexander Korotin, Alexander Kolesov, Nikita Gushchin, and Evgeny Burnaev. Energy-guided entropic neural optimal transport. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=d6tUsZeVs7.

Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.

Stanislav Pidhorskyi, Donald A Adjeroh, and Gianfranco Doretto. Adversarial latent autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14104–14113, 2020.

Sergey Samsonov, Evgeny Lagutin, Marylou Gabrié, Alain Durmus, Alexey Naumov, and Eric Moulines. Local-global MCMC kernels: the best of both worlds. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL https://openreview.net/forum?id=zb-xfApk4ZK.

Vivien Seguy, Bharath Bhushan Damodaran, Remi Flamary, Nicolas Courty, Antoine Rolet, and Mathieu Blondel. Large scale optimal transport and mapping estimation. In *International Conference on Learning Representations*, 2018.

Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.

Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36:62183–62223, 2023.

Panagiotis Theodoropoulos, Nikolaos Komianos, Vincent Pacelli, Guan-Horng Liu, and Evangelos A Theodorou. Feedback schr\" odinger bridge matching. *arXiv preprint arXiv:2410.14055*, 2024.

Alexander Y 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*, pp. 1279–1287. PMLR, 2024.

Francisco Vargas, Pierre Thodoroff, Austen Lamacraft, and Neil Lawrence. Solving schrödinger bridges via maximum likelihood. *Entropy*, 23(9):1134, 2021.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 586–595, 2018.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 2223–2232, 2017.

A PROOFS

Notation. Throughout the appendix, we use the shorthand $Z_f := Z(f, \cdot)$ and $\zeta(x_0) := \xi(x_0) + \frac{D}{2}\log(2\pi\varepsilon)$.

Definition A.1. Let μ and ν be non-negative measures on $\mathbb{R}^d$, with μ absolutely continuous w.r.t. ν. The Kullback-Leibler divergence is defined by

$$\mathrm{KL}(\mu\|\nu) := \int \left[\frac{d\mu}{d\nu}\log\left(\frac{d\mu}{d\nu}\right) - \frac{d\mu}{d\nu} + 1\right] d\nu.$$

Definition A.2 (Rademacher complexity)**.** Let μ be a distribution on X, and let $\mathcal{H}$ be a class of functions $h : X \to \mathbb{R}$. The Rademacher complexity is defined by

$$\mathfrak{R}_N(\mathcal{H}, \mu) := \frac{1}{N}\,\mathbb{E}\left[\sup_{h\in\mathcal{H}} \sum_{i=1}^{N} \sigma_i h(x_i)\right],$$

where $\sigma_i \in \{-1, +1\}$ are independent Bernoulli random variables and x_i are i.i.d. samples from μ.

The following symmetrization bound is standard; see, e.g., Shalev-Shwartz & Ben-David (2014, Lemma 26.2).

Lemma A.3 (Representativeness estimation)**.** *Let $\mathcal{H}$ be a class of measurable functions $h : X \to \mathbb{R}$ that are integrable with respect to μ. Then for i.i.d. samples $x_1, \ldots, x_N \sim \mu$,*

$$\mathbb{E}\sup_{h\in\mathcal{H}}\left|\mathbb{E}_\mu[h] - \tfrac{1}{N}\sum_{i=1}^{N} h(x_i)\right| \le 2\,\mathfrak{R}_N(\mathcal{H}, \mu).$$

The next lemma is Talagrand's contraction lemma (Ledoux & Talagrand, 1991, Theorem 4.12).

Lemma A.4. *Let $\psi_i : \mathbb{R} \to \mathbb{R}$ be L-Lipschitz functions. Then for any collection of points x_i,*

$$\mathbb{E}_\sigma\Big[\sup_{h\in\mathcal{H}} \sum_{i=1}^{N} \sigma_i\,\psi_i(h(x_i))\Big] \le L\,\mathbb{E}_\sigma\Big[\sup_{h\in\mathcal{H}} \sum_{i=1}^{N} \sigma_i\,h(x_i)\Big].$$

Lemma A.5 (Two-dimensional lemma for products of classes)**.** *Let $\psi : \mathbb{R}^2 \to \mathbb{R}$ be such that* (17) *holds on the range of values $\{(u(x_0), v(x_0)) :\ u \in \mathcal{U},\ v \in \mathcal{V},\ x_0 \in X\}$. Then*

$$\mathfrak{R}_N\big(\{\psi(u(\cdot), v(\cdot)) :\ u \in \mathcal{U},\ v \in \mathcal{V}\},\ p\big) \le L_u\,\mathfrak{R}_N(\mathcal{U}, p) + L_v\,\mathfrak{R}_N(\mathcal{V}, p).$$

Proof. Fix a sample $x_1, \ldots, x_N$ and consider the empirical Rademacher complexity. Let $u_i = u(x_i)$ and $v_i = v(x_i)$. Consider the decomposition $\psi(u_i, v_i) = \big(\psi(u_i, v_i) - \psi(0, v_i)\big) + \psi(0, v_i)$. Then

$$\sup_{u\in\mathcal{U}, v\in\mathcal{V}} \sum_{i=1}^{N} \sigma_i\psi(u_i, v_i) \le \sup_{u,v} \sum_{i=1}^{N} \sigma_i\big(\psi(u_i, v_i) - \psi(0, v_i)\big) + \sup_{v} \sum_{i=1}^{N} \sigma_i\psi(0, v_i).$$

For fixed v, the functions $t \mapsto \psi(t, v_i) - \psi(0, v_i)$ are L_u-Lipschitz and vanish at $t = 0$. By Lemma A.4,

$$\mathbb{E}_\sigma \sup_{u} \sum_{i=1}^{N} \sigma_i\big(\psi(u_i, v_i) - \psi(0, v_i)\big) \le L_u\,\mathbb{E}_\sigma \sup_{u} \sum_{i=1}^{N} \sigma_i u_i.$$

The right-hand side does not depend on v, so we may add $\sup_v$ without changing it. Similarly, since $t \mapsto \psi(0, t)$ is L_v-Lipschitz, we again apply Lemma A.4 to obtain

$$\mathbb{E}_\sigma \sup_{v} \sum_{i=1}^{N} \sigma_i\psi(0, v_i) \le L_v\,\mathbb{E}_\sigma \sup_{v} \sum_{i=1}^{N} \sigma_i v_i.$$

Dividing by N and taking expectation over the sample yields the claim. □

Lemma A.6. *For any integrable f and ζ, the following identity holds:*

$$\mathcal{L}(f) - \mathcal{L}(f,\zeta) = \varepsilon\, \mathbb{E}_{x_0\sim p_0}\Big[\phi\Big(\frac{Z_f(x_0)}{\exp\zeta(x_0)}\Big)\Big],$$

where $\phi(t) := t - \log t - 1 \geq 0$. In particular, $\mathcal{L}(f) \geq \mathcal{L}(f,\zeta)$, and equality holds if and only if $\zeta(x_0) = Z_f(x_0)$ for p-a.e. x_0.

Proof.

$$\mathcal{L}(f) - \mathcal{L}(f,\zeta) = -\varepsilon\mathbb{E}_p[\log Z_f] + \varepsilon\mathbb{E}_p[\log\zeta] + \varepsilon\mathbb{E}_p[Z_f/\zeta] - \varepsilon = \varepsilon\mathbb{E}_p\Big[\tfrac{Z_f}{\zeta} - \log\tfrac{Z_f}{\zeta} - 1\Big].$$

□

A.1 Proof of Proposition 3.1

Proof. Make the change of variables $x_1 = x_0 + \sqrt{\varepsilon}\, z$, so that $dx_1 = \varepsilon^{D/2} dz$ and hence

$$Z(f,x_0) = (2\pi\varepsilon)^{D/2}\, \mathbb{E}_{z\sim\mathcal{N}(0,I)}\left[\exp\left(\frac{f(x_0+\sqrt{\varepsilon}z)}{\varepsilon}\right)\right] = \mathbb{E}[\exp(A(z))]\,,$$

where $A(z) := \frac{f(x_0+\sqrt{\varepsilon}z)}{\varepsilon} + \underbrace{\frac{D}{2}\log(2\pi\varepsilon)}_{\stackrel{\text{def}}{=} C}$. Then for any $\xi(x_0)\in\mathbb{R}$,

$$\log Z = \log\mathbb{E}[e^A] = (\xi + C) + \log\mathbb{E}[e^{A-(\xi+C)}] \leq (\xi + C) + \big(\mathbb{E}[e^{A-(\xi+C)}] - 1\big),$$

using $\log u \leq u - 1$ for all $u > 0$. Equality holds iff $\mathbb{E}[e^{A-(\xi+C)}] = 1$, i.e. when $\xi + C = \log\mathbb{E}[e^A] = \log Z(f,x_0)$. □

A.2 Proof of Theorem 3.2

Proof. Recall the weak dual form of entropic OT (for the quadratic cost)

$$\mathrm{EOT}_\varepsilon(p_0,p_1) = \sup_f\Big\{\mathbb{E}_{x_1\sim p_1}[f(x_1)] - \varepsilon\,\mathbb{E}_{x_0\sim p_0}\big[\log Z(f,x_0)\big]\Big\},$$

where $Z(f,x_0) = \int \exp\big((f(x_1) - \frac{1}{2}\|x_1 - x_0\|^2)/\varepsilon\big)\, dx_1$. By Proposition 3.1 (with $C = \frac{D}{2}\log(2\pi\varepsilon)$), for any $\xi(x_0)$ and all x_0,

$$\log Z(f,x_0) \leq C - 1 + \xi(x_0) + \mathbb{E}_{z\sim\mathcal{N}(0,I)}\Big[\exp\Big(\frac{f(x_0+\sqrt{\varepsilon}z)}{\varepsilon} - \xi(x_0)\Big)\Big].$$

Plugging this upper bound into the weak dual yields $\mathrm{EOT}_\varepsilon(p_0,p_1) \geq \sup_{f,\xi}\mathcal{L}(f,\xi)$. Conversely, since the inequality holds for every ξ, for each fixed f we have $\mathcal{L}(f,\xi) \leq \mathbb{E}_{p_1}[f] - \varepsilon\mathbb{E}_{p_0}[\log Z(f,x_0)]$, hence $\sup_{f,\xi}\mathcal{L}(f,\xi) \leq \mathrm{EOT}_\varepsilon(p_0,p_1)$. The bound is tight at $\xi(x_0) + C = \log Z(f,x_0)$ (Proposition 3.1), so equality holds. □

A.3 Proof of Theorem 3.3

Proof. Let f^* be a maximizer of the dual problem (2), and set $\exp\zeta^*(x_0) := Z_{f^*}(x_0)$. Therefore,

$$\varepsilon\int\log\left(\frac{d\pi^*}{d\pi^{f,\zeta}}\right)d\pi^* = \varepsilon\int_{X_0}(\zeta(x_0) - \zeta^*(x_0))\,dp(x_0) + \int_{X_1}\big(f^*(x_1) - f(x_1)\big)\,dp_1(x_1),$$

where we used that π^* has marginals p_0 and p_1.

Next, we compute the total mass of $\pi^{f,\zeta}$:

$$\int_{X_0}\int_{X_1} p_0(x)\exp\left(\frac{f(x_1) - c(x_0,x_1)}{\varepsilon} - \zeta(x_0)\right)dx_1\,dx_0 = \int_{X_0}\frac{Z_f(x_0)}{\exp\zeta(x_0)}\,dp(x_0).$$

Since π^* is a probability measure, $\pi^*(X_0\times X_1) = 1$. Combining,

$$\varepsilon\,\mathrm{KL}(\pi^*\|\pi^{f,\zeta}) = \left(\int_{X_1} f^*\,dp_1 - \varepsilon\int_{X_0}\zeta^*\,dp\right) - \left(\int_{X_1} f\,dp_1 - \varepsilon\int_{X_0}\zeta\,dp_0 - \varepsilon\int_{X_0}\frac{Z_f}{\exp\zeta}\,dp_0 + \varepsilon\right).$$

The first bracket equals $\mathcal{L}(f^*) = \mathcal{L}^*$, while the second bracket is $\mathcal{L}(f,\zeta)$. □

A.4 Proof of Proposition 3.4

Proof. Recall $\hat{f} = \operatorname{argmax}_{f\in\mathcal{F}} \widehat{\mathcal{L}}(f, Z(f,\cdot))$, $\hat{\zeta} = \log Z_{\hat{f}}$, then

$$\begin{aligned}
\varepsilon\mathbb{E}\Big[\mathrm{KL}\Big(\pi^* \,\Big\|\, \pi^{\hat{f},\hat{\zeta}}\Big)\Big] &= \mathcal{L}^* - \mathcal{L}(\hat{f},\hat{\zeta}) \\
&= \mathcal{L}^* - \sup_{f\in\mathcal{F},\zeta\in\Xi} \mathcal{L}(f,\zeta) + \sup_{f\in\mathcal{F},\zeta\in\Xi} \mathcal{L}(f,\zeta) - \sup_{f\in\mathcal{F},\zeta\in\Xi} \widehat{\mathcal{L}}(f,\zeta) \\
&+ \underbrace{\sup_{f\in\mathcal{F},\zeta\in\Xi} \widehat{\mathcal{L}}(f,\zeta) - \widehat{\mathcal{L}}(\hat{f},\hat{\zeta})}_{\leq 0} + \widehat{\mathcal{L}}(\hat{f},\hat{\zeta}) - \mathcal{L}(\hat{f},\hat{\zeta}) \\
&\leq 2 \sup_{f\in\mathcal{F},\zeta\in\Xi} |\mathcal{L}(f,\zeta) - \widehat{\mathcal{L}}(f,\zeta)| + \mathcal{L}^* - \sup_{f\in\mathcal{F},\zeta\in\Xi} \mathcal{L}(f,\zeta).
\end{aligned}$$

□

A.5 Proof of Theorem 3.5

Fix classes $\mathcal{F} \subset \mathrm{Lip}_{L_f,M_f}(X_1)$ and $\Xi \subset \mathrm{Lip}_{L_\zeta,M_\zeta}(X_0)$ so:

$$\|f\|_\infty \leq M_f \;\; \forall f \in \mathcal{F}, \quad 0 < \delta_{\mathcal{E}} := \exp(-M_\zeta) \leq \exp\zeta(x) \leq M_{\mathcal{E}} := \exp(M_\zeta) \;\; \forall \zeta \in \Xi. \tag{16}$$

We introduce $\mathcal{E} := \{\exp(\zeta) : \zeta \in \Xi\}$. Our goal is to bound $\Delta := \sup_{f\in\mathcal{F},\,\zeta\in\Xi} \Big|\mathcal{L}(f,\zeta) - \widehat{\mathcal{L}}(f,\zeta)\Big|$.

By the triangle inequality

$$\Delta \leq \underbrace{\sup_{f\in\mathcal{F}} \left|\mathbb{E}_q[f] - \frac{1}{M}\sum_{j=1}^{M} f(y_j)\right|}_{\Delta_1} + \varepsilon \underbrace{\sup_{\zeta\in\Xi} \left|\mathbb{E}_p[\zeta] - \frac{1}{N}\sum_{i=1}^{N} \zeta(x_i)\right|}_{\Delta_2} + \varepsilon \underbrace{\sup_{f\in\mathcal{F},\,\zeta\in\Xi} \left|\mathbb{E}_p\left[\frac{Z_f}{\zeta}\right] - \frac{1}{N}\sum_{i=1}^{N}\left[\frac{Z_f}{\exp\zeta}\right](x_i)\right|}_{\Delta_3}.$$

By Lemma A.3, $\mathbb{E}[\Delta_1] \leq 2\,\mathfrak{R}_M(\mathcal{F},q)$ and $\mathbb{E}[\Delta_2] \leq 2\,\mathfrak{R}_N(\Xi,p)$.

Consider $\mathcal{Z} := \{Z(f,\cdot) : \; f \in \mathcal{F}\}$ and $\mathcal{H}_{\mathrm{frac}} := \{Z(f,\cdot)/\exp\zeta : \; f \in \mathcal{F},\; \zeta \in \Xi\}$. From the bounds we get $0 < Z(f,x_0) \leq (2\pi\varepsilon)^{D/2} \cdot e^{M_f/\varepsilon} =: M_Z$.

Define $\psi(u,v) = u/v$ on $u \in [0, M_Z]$, $v \in [\delta_{\mathcal{E}}, M_{\mathcal{E}}]$. Then $|\partial\psi/\partial u| = 1/v \leq 1/\delta_{\mathcal{E}} =: L_u$ and $|\partial\psi/\partial v| = u/v^2 \leq M_Z/\delta_{\mathcal{E}}^2 =: L_v$, so:

$$|\psi(u,v) - \psi(u',v')| \leq L_u|u-u'| + L_v|v-v'|. \tag{17}$$

Applying Lemma A.3 and Lemma A.5, we obtain

$$\mathbb{E}[\Delta_3] \leq 2\left(\frac{1}{\delta_{\mathcal{E}}}\mathfrak{R}_N(\mathcal{Z},p) + \frac{M_Z M_{\mathcal{E}}}{\delta_{\mathcal{E}}^2}\mathfrak{R}_N(\Xi,p)\right).$$

Collecting bounds:

$$\mathbb{E}[\Delta] \leq 2\,\mathfrak{R}_M(\mathcal{F},q) + \frac{2\varepsilon}{\delta_{\mathcal{E}}}\,\mathfrak{R}_N(\Xi,p) + 2\varepsilon\left(\frac{1}{\delta_{\mathcal{E}}}\mathfrak{R}_N(\mathcal{Z},p) + \frac{M_Z M_{\mathcal{E}}}{\delta_{\mathcal{E}}^2}\mathfrak{R}_N(\Xi,p)\right).$$

By Gottlieb et al. (2016, Theorem 4.3), $\mathfrak{R}(\mathcal{F},p_1) \leq O(M^{-1/(1+D)})$ and $\mathfrak{R}(\Xi,p_0) \leq O(N^{-1/(1+D)})$. By Kolesov et al. (2024, Theorem 4.5), $\mathfrak{R}(\mathcal{Z},p_0) \leq O(N^{-1/2}) \leq O(N^{-1/(1+D)})$.

A.6 Proof of Theorem 3.7

Proof. Let $X_0 := \mathrm{supp}(p_0)$ and $X_1 := \mathrm{supp}(p_1)$ be compact, and $c(x_0,x_1) = \frac{1}{2}\|x_0 - x_1\|^2$. Fix $\delta > 0$. We use that $\mathcal{L}$ is invariant under additive shifts: $\mathcal{L}(f+t) = \mathcal{L}(f)$ for any $t \in \mathbb{R}$.

Assume $\mathcal{L}^* < \infty$ and let f^* be an optimizer with $\sup_{x_1\in X_1} f^*(x_1) = 0$. Since X_1 is compact, pick a Lipschitz function $g : X_1 \to \mathbb{R}$ with $\|g - f^*\|_{\infty,X_1} \leq \eta$ and $\sup_{X_1} g = 0$, where $\eta > 0$ will be chosen. Set $B := \|g\|_{\infty,X_1}$ and choose $M > B$.

Define the L-Lipschitz extension $\widetilde{g}_L(x_1) := \sup_{z \in X_1}(g(z) - L\|x_1 - z\|)$ and the bounded Lipschitz potential $f_{L,M}(x_1) := \max\{-M, \widetilde{g}_L(x_1)\}$. Writing $t := (M + B)/L$ and $X_1^t := \{x_1 : \mathrm{dist}(x_1, X_1) \le t\}$, we have $f_{L,M} = -M$ outside X_1^t.

For $x_0 \in X_0$, the uniform approximation implies $e^{-\eta/\varepsilon} Z_{f^*}(x_0) \le \int_{X_1} e^{(g - c(x_0,\cdot))/\varepsilon}\, dx_1 \le e^{\eta/\varepsilon} Z_{f^*}(x_0)$. Since $\lambda(X_1) > 0$, we get a uniform lower bound $\underline{Z} > 0$, and

$$\mathcal{L}(f^*) - \mathcal{L}(f_{L,M}) \le 2\eta + \varepsilon \, \frac{\lambda(X_1^t \setminus X_1) + e^{-M/\varepsilon}(\pi\varepsilon)^{D/2}}{e^{-\eta/\varepsilon}\underline{Z}}.$$

Since $\lambda(X_1^t \setminus X_1) \to 0$ as $t \to 0$, choose η, M, and L so that $\mathcal{L}(f_{L,M}) \ge \mathcal{L}^* - \delta/2$.

For n large enough that $L_n \ge L$ and $M_n \ge M$, we have $f_{L,M} \in \mathcal{F}_n$. The one-potential objective can be written in two-potential form with $\zeta_f = -\varepsilon \log Z_f$, and for bounded f this ζ_f is bounded and Lipschitz. For n large enough, $\zeta_f \in \Xi_n$ and $\sup_{f \in \mathcal{F}_n,\, \zeta \in \Xi_n} \mathcal{L}(f, \zeta) \ge \sup_{f \in \mathcal{F}_n} \mathcal{L}(f) \ge \mathcal{L}^* - \delta/2$, yielding the result. □

B Details of the Experiments

B.1 VarEOT: Optimization and Architecture

Optimization. Training uses AdamW (Loshchilov & Hutter, 2017) with lr $= 10^{-4}$, $\beta_1 = 0.7$, $\beta_2 = 0.8$, weight decay 10^{-4}. EMA with momentum 0.999 is used for evaluation. All models are trained for 10^4 gradient steps.

Network architecture. The transport potential is parameterized by an MLP with four fully connected layers, SiLU activations, hidden width 256, input dimension $d_{\text{in}} = 512$ (for latent-space experiments), and output dimension 1.

Training setup. Batch size 256. Monte Carlo samples $K = 256$.

Langevin inference. For SwissRoll: 1000 steps, step size 10^{-3}. For ALAE: step size 0.5 (NFE=2), 0.1 (NFE=10), 10^{-3} (NFE=1000). Step sizes selected based on FID (Figure 6).

Update schedule. Simultaneous optimization of f_θ and ξ_ψ at every step.

B.2 Baseline Methods

LightSB. $K = 10$ Gaussian components, lr $= 10^{-3}$, batch size 128, 10^4 steps.

EgNOT. Same network architecture as VarEOT for fair comparison. Adam with lr $= 5 \cdot 10^{-5}$, $(\beta_1, \beta_2) = (0, 0.999)$, 10^4 steps, batch size 128.

B.3 Image Data and Preprocessing

We use the official ALAE implementation (`https://github.com/podgorskiy/ALAE`), FFHQ annotations from `https://github.com/DCGM/ffhq-features-dataset`, LightSB from `https://github.com/ngushchin/LightSB`, and EgNOT from `https://github.com/PetrMokrov/Energy-guided-Entropic-OT`.

C Additional Experimental Results

C.1 Other Unpaired Image-to-Image Translation Setups

We provide qualitative results for A→C, and C→A setups with $\varepsilon = 1.0$ in Figures 3–4, confirming consistent behavior across translation directions.

Figure 3: Qualitative comparison for *Adult* $\rightarrow$ *Child* translation with $\varepsilon = 1.0$. Top to bottom: input, VarEOT (ours), LightSB, EgNOT.

Figure 4: Qualitative comparison for *Child* $\rightarrow$ *Adult* translation with $\varepsilon = 1.0$. Top to bottom: input, VarEOT (ours), LightSB, EgNOT.

C.2 Dependence on ε

Figure 5 shows how VarEOT's solution depends on ε in the M$\rightarrow$F experiment. Diversity increases with ε.

C.3 Effect of Langevin Inference Parameters

Figure 6 shows FID as a function of Langevin step size and number of steps for different ε in the M$\rightarrow$F task.

(a) VarEOT M→F, $\varepsilon = 0.1$. Almost no diversity.

(b) VarEOT M→F, $\varepsilon = 0.5$. Reasonable diversity.

(c) VarEOT M→F, $\varepsilon = 1.0$. Moderate diversity.

(d) VarEOT M→F, $\varepsilon = 10.0$. High diversity.

Figure 5: VarEOT with NFE=10 in M→F for different ε.

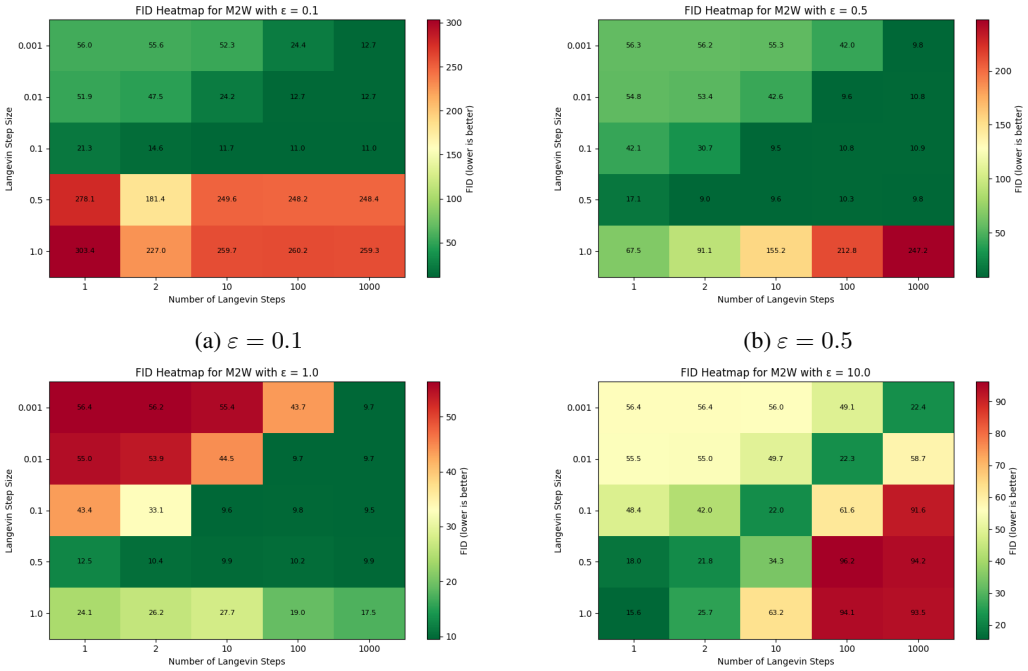

(a) $\varepsilon = 0.1$ (b) $\varepsilon = 0.5$

(c) $\varepsilon = 1.0$ (d) $\varepsilon = 10.0$

Figure 6: FID heatmaps for M→F in the ALAE latent space. Each heatmap shows FID vs. Langevin step size and number of inference steps for a fixed ε.