

000 001 002 003 004 005 RETHINKING THE FLOW-BASED GRADUAL DOMAIN 006 ADAPTION: A SEMI-DUAL TRANSPORT PERSPECTIVE 007 008 009

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

ABSTRACT

028 Gradual domain adaptation (GDA) aims to mitigate domain shift by progressively
029 adapting models from the source domain to the target domain via intermediate
030 domains. However, real intermediate domains are often unavailable or ineffect-
031 tive, necessitating the synthesis of intermediate samples. Flow-based models have
032 recently been used for this purpose by interpolating between source and target dis-
033 tributions; however, their training typically relies on sample-based log-likelihood
034 estimation, which can discard useful information and thus degrade GDA perfor-
035 mance. The key to addressing this limitation is constructing the intermediate do-
036 mains via samples directly. To this end, we propose an Entropy-regularized Semi-
037 dual Unbalanced Optimal Transport (E-SUOT) framework to construct interme-
038 diate domains. Specifically, we reformulate flow-based GDA as a Lagrangian
039 dual problem and derive an equivalent semi-dual objective that circumvents the
040 need for likelihood estimation. However, the dual problem leads to an unstable
041 min–max training procedure. To alleviate this issue, we further introduce entropy
042 regularization to convert it into a more stable alternative optimization procedure.
043 Based on this, we propose a novel GDA training framework and provide theoret-
044 ical analysis in terms of stability and generalization. Finally, extensive experiments
045 are conducted to demonstrate the efficacy of the E-SUOT framework.
046

1 INTRODUCTION

047 Unsupervised Domain Adaptation (UDA) (Pan & Yang, 2010; Tzeng et al., 2017; Long et al., 2015;
048 Courty et al., 2014; 2017a), which transfers knowledge from a well-trained source domain to a re-
049 lated yet unlabeled target domain, is of great importance across fundamental application areas. For
050 example, in recommender systems (Liu et al., 2023; Zheng et al., 2024), a cold-start user has no
051 interaction history with new items, so domain adaptation helps transfer user and item knowledge
052 from an existing system to improve recommendations. Similar scenarios occur in machine trans-
053 lation, where a model trained on high-resource language pairs like English-French can be adapted
054 to translate between English and low-resource languages with limited parallel data (Gazdieve et al.,
055 2023). These scenarios highlight the importance of conducting UDA to bridge domain gaps and
056 ensure reliable performance in real-world applications.

057 Despite these methodological advances, directly performing UDA can be brittle when the
058 source–target shift is substantial or class overlap is weak. In such cases, one-shot alignment of-
059 ten degrades discriminability and amplifies pseudo-label errors during self-training. This challenge
060 motivates a transition from the traditional UDA setting to the Gradual Domain Adaptation (GDA)
061 setting (He et al., 2024), where adaptation proceeds through a sequence of intermediate distributions
062 that progressively bridge the domain gap. A key aspect of generating intermediate domains in GDA
063 is to interpolate between the source and target domains. Various methods have been proposed to
064 construct such intermediate domains, among which flow-based approaches (Kobyzev et al., 2020;
065 Papamakarios et al., 2021) have attracted increasing attention, primarily due to their property of
066 preserving probability density along the transformation path, thereby enabling consistent and stable
067 probability densities without distortion or loss of information. To drive the samples from the source
068 domain towards those of the target domain, it is necessary to design an appropriate driving force,
069 typically derived from a discrepancy metric. Among these metrics, f -divergence (Sason & Verdú,
070 2016) is most widely used due to its computational efficiency, empirical effectiveness, and principled
071 formulation within the framework of geometry for probability distributions (Amari, 2016).

Despite the success of flow-based approaches in GDA (Sagawa & Hino, 2025; Zhuang et al., 2024; Zeng et al., 2025), we argue that directly applying standard flow-based models leads to suboptimal performance. Specifically, existing flow-based frameworks utilizing f -divergence often require the explicit estimation of target domain probability density functions (PDFs) from available target samples (In our setting, for simplicity, we treat both log PDF and its gradient, also known as **score function**, as forms of density estimation, since they characterize the underlying data distribution.) (Vincent, 2011; Santambrogio, 2017; Ambrosio et al., 2005), whereas the subsequent GDA process relies on these estimated (normalized / unnormalized) PDFs to drive the source-to-target transfer. For example, Zhuang et al. (2024) estimate the unnormalized target domain PDF in the **score function form and generate intermediate domains via Langevin dynamics**. Consequently, the quality of the intermediate domain heavily depends on the accuracy of the estimated target PDF; if this estimation is inaccurate, the performance of the downstream task is likely to suffer significantly.

To address these limitations, we propose a novel flow-based GDA framework E-SUOT, which leverages the semi-dual formulation of gradient flows. Rather than explicitly estimating PDFs, we recast flow evolution as an optimization problem that combines an f -divergence term with a Wasserstein distance regularization term, enabling sample transport toward the target domain without reliance on PDF estimation. However, as the semi-dual reformulation inherently leads to an adversarial training paradigm that can compromise stability and performance, we introduce entropy regularization to the objective to guarantee the stability of the training process. Based on this, we summarize the algorithm for E-SUOT-based intermediate domain generation, prove the convergence of our E-SUOT framework, and empirically demonstrate its effectiveness on representative GDA tasks. Extensive experiments validate that E-SUOT achieves superior performance compared with existing methods.

Contributions. The main contributions of this paper are summarized as follows:

- We develop a semi-dual formulation for intermediate domain generation in flow-based GDA, which eliminates the need for explicit estimation of the target-domain PDF—whether normalized or unnormalized—or its score-based representation.
- We introduce an entropy regularization term to address the unstable issue inherent in the semi-dual formulation, resulting in the novel and stable E-SUOT framework.
- We conducted various experiments to demonstrate the superiority of the proposed E-SUOT approach compared to prevalent approaches.

2 PRELIMINARIES

2.1 SETTINGS AND NOTATIONS

In GDA, we consider a labeled source domain, $T - 1$ unlabeled intermediate domains, and an unlabeled target domain. Let the input space be \mathcal{X} and the label space be \mathcal{Y} . We denote inputs as $x \in \mathcal{X}$ and labels as $y \in \mathcal{Y}$. We index the domains by $t \in \{0, 1, \dots, T\}$, where $t = 0$ denotes the source domain and $t = T$ denotes the target domain. Each domain induces a marginal distribution p_t over \mathcal{X} . Let \mathcal{H} be a hypothesis class of classifiers $h : \mathcal{X} \rightarrow \mathcal{Y}$. We assume that each domain admits a labeling function $q_t \in \mathcal{H}$. Given a loss function $\mathcal{L} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_{\geq 0}$, the generalization error of h on domain t is defined as $\varepsilon_{p_t}(h) = \mathbb{E}_{p_t(x)}[\mathcal{L}(h(x), q_t(x))]$. A source classifier $q_0 \in \mathcal{H}$ can be learned via supervised learning on the source domain with minimal error $\varepsilon_{p_0}(q_0)$. The objective of GDA is to evolve q_0 through the intermediate domains to a classifier h_T so as to minimize the target error $\varepsilon_{p_T}(h_T)$.

2.2 FLOWS FOR INTERMEDIATE DOMAIN GENERATION

A flow describes the time-dependent evolution of particles induced by a smooth invertible (diffeomorphic) map. Based on this, the intermediate domains can be seen as a discretization of a continuous flow linking source and target distributions. This motivates flow-based models, which evolve a distribution over a fixed time horizon while preserving normalization, and are thus well-suited for GDA. From the flow perspective, intermediate domains are generated by the following ordinary differential equation:

$$\frac{dx_t}{dt} = v_t(x_t) = -\nabla \frac{\delta \mathbb{D}[p(x_t), p_T(x)]}{\delta p(x_t)}, \quad x_{t=0} = x_0, \quad (1)$$

108 where $p(x_t)$ is the (empirical) PDF induced by $\{x_{t,i}\}_{i=1}^N$, and we desire the law $p(x_T)$ to approximate
 109 the target $p_T(x)$. Here $v_t : \mathcal{X} \rightarrow \mathcal{X}$ is the velocity field. The core design problem is to choose
 110 v_t so that $p(x_t) \xrightarrow{t \rightarrow T} p_T(x)$. A principled approach is to define v_t as the steepest descent direction
 111 of some discrepancy functional $\mathbb{D}[p(x_t), p_T(x)]$ between $p(x_t)$ and $p_T(x)$ as demonstrated in the
 112 second equal sign in Eq. (1). Notably, $\delta/\delta p$ denotes the first variation, and the second equality sign
 113 is called “gradient flow”.

114 Among various choices, f -divergences are favored in GDA for their task-aligned objectives, stable
 115 probability-preserving dynamics, and efficient computation when compared to alternatives such as
 116 Sinkhorn divergence and maximum mean discrepancy (Glaser et al., 2021). For an f -divergence,
 117

$$118 \quad \mathbb{D}_f[p(x_t), p_T(x)] = \int f\left(\frac{p(x_t)}{p_T(x)}\right) p_T(x) dx, \quad (2)$$

120 with $f : (0, \infty) \rightarrow \mathbb{R}$ convex and $f(x) = 0$ if and only if $x = 1$. A canonical example is the
 121 Kullback–Leibler (KL) divergence with $f(u) = u \log u$. In this case,
 122

$$123 \quad v_t(x_t) = \nabla \log p_T(x) - \nabla \log p(x_t), \quad (3)$$

124 and, in the weak partial differential equation sense (Evans, 2022; Liu, 2017), the induced dynamics
 125 yield the classical *Langevin dynamic* (Welling & Teh, 2011; Santambrogio, 2017).

126 Intuitively, applying the forward Euler scheme with step size η to the gradient flow in Eq. (1) under
 127 an f -divergence yields a discrete-time generation for the intermediate domain, which is equivalent
 128 to solving a 2-Wasserstein-distance–regularized optimization problem as (see Section C.1):
 129

$$130 \quad x_{t+\eta} = x_t - \eta \nabla \frac{\delta \mathbb{D}_f[p(x_t), p_T(x)]}{\delta p(x_t)} \Rightarrow p(x_{t+\eta}) = \arg \min_{\rho(x) \in \mathcal{P}_2(\mathbb{R}^D)} \frac{1}{2\eta} \mathcal{W}_2^2(\rho(x), p(x_t)) + \mathbb{D}_f[\rho(x), p_T(x)], \quad (4)$$

133 where $\mathcal{P}_2(\mathbb{R}^D)$ denotes the Wasserstein space (Villani et al., 2009), which is the set of the distributions
 134 with finite second moment. Here \mathcal{W}_2 is the 2-Wasserstein distance, whose definition is given
 135 as follows:

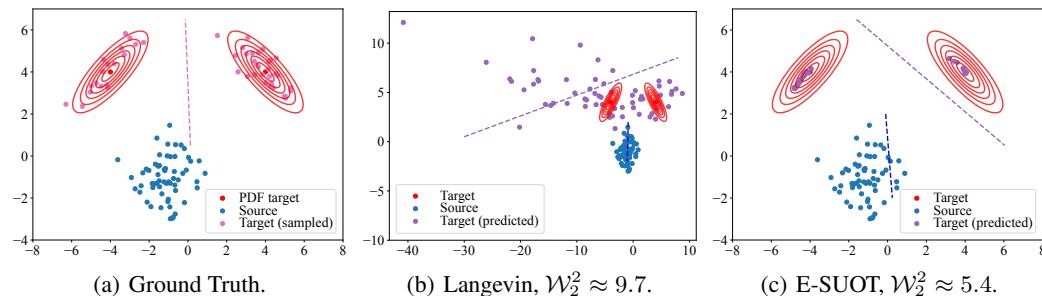
$$136 \quad \mathcal{W}_2^2(\rho, \xi) = \inf_{\pi \in \Pi(\rho, \xi)} \iint \|x - y\|_2^2 \pi(x, y) dx dy, \quad (5)$$

138 and $\Pi(\rho, \xi)$ is the set of joint distribution on $\mathbb{R}^D \times \mathbb{R}^D$ with marginal distributions ρ and ξ .

140 3 METHODOLOGY

142 3.1 MOTIVATION ANALYSIS

144 Flow-based approaches, exemplified by gradient-flow methods, interpolate between the source and
 145 target distributions by gradually minimizing a discrepancy measure, typically an f -divergence, be-
 146 tween the two domains. The success of these methods in GDA tasks critically depends on accurately
 147 estimating the target distribution’s probability density function (PDF). Given a reliable estimate, one
 148 can construct a velocity field that progressively pushes source samples toward the target distribution.



159 Figure 1: Illustrative Example: Comparison between Langevin Dynamics and E-SUOT, where the
 160 corresponding decision boundary is shown as a dashed line.

161 However, directly estimating the PDF from target domain data is generally ill-posed (Vincent et al.,
 162 2010; Song et al., 2020). When the estimate is inaccurate, the induced velocity field can push

samples into low-probability regions of the target distribution, causing a substantial shift between the generated and true target domains and degrading downstream task performance. To illustrate this issue, we compare ground-truth target samples with those obtained via Langevin dynamics and E-SUOT in Figs. 1(a) to 1(c). The PDF for the target domain is estimated using denoised score matching (Vincent, 2011). In addition, we also report the Wasserstein distance between the predicted and ground-truth samples (relative to Fig. 1(a)) in the captions of Figs. 1(b) and 1(c), which constituted the lower generalization bound for GDA tasks. From Figs. 1(b) and 1(c), it is evident that when the estimated log-likelihood function is inaccurate, the samples generated for the target distribution deviate substantially from the ground truth and yield a large Wasserstein distance, which may ultimately limit performance on GDA tasks. *Although neither method perfectly recovers the ground truth distribution, E-SUOT shows a clearer alignment with the major modes and the corresponding decision boundary of the target domain, resulting in a substantially lower Wasserstein distance compared to Langevin dynamics.* In summary, the key questions addressed in this paper can be summarized as follows: How can we generate intermediate domains without compromising the accuracy of the target domain? How can robust intermediate domain generation be achieved within this framework? Does this approach improve the performance for GDA task?

3.2 DUAL-FORM TRANSPORTATION FOR INTERMEDIATE DOMAIN GENERATION

As shown in Eq. (4), simulating the gradient flow to generate intermediate domains is precisely equivalent to solving a Wasserstein-distance-regularized optimization problem. This insight opens up a practical alternative: *instead of explicitly estimating the target domain’s probability density, one can guide source samples by directly tackling this optimization formulation.* Thus, we have the following proposition regarding the solution property of the problem defined in Eq. (4):

Proposition 1. *Consider the following primal problem:*

$$\mathcal{L}^{\text{Primal}} = \arg \min_{\rho(x) \in \mathcal{P}_2(\mathbb{R}^D)} \frac{1}{2\eta} \mathcal{W}_2^2(\rho(x), p(x_t)) + \mathbb{D}_f[\rho(x), p_T(x)]. \quad (6)$$

This problem is equivalent to the following semi-dual formulation:

$$\mathcal{L}^{\text{SemiDual}} = \sup_w \mathbb{E}_{p(x_t)} \left[\inf_{\mathbf{T}} \left(\frac{1}{2\eta} \|\mathbf{T}(x_t) - x_t\|_2^2 - w(\mathbf{T}(x_t)) \right) \right] - \mathbb{E}_{p_T(x)} [f^*(-w(x))], \quad (7)$$

where $w : \mathbb{R}^D \rightarrow \mathbb{R}$ is a measurable continuous function, $\mathbf{T} : \mathbb{R}^D \rightarrow \mathbb{R}^D$ is the transport map, and f^ denotes the convex conjugate of f , defined as $f^*(z) := \sup_{y \geq 0} (zy - f(y))$.*

Importantly, the structure of the semi-dual problem ensures that both $p_t(x)$ and $p_T(x)$ are involved only through expectation operators, rather than through explicit density evaluations. This enables the use of Monte Carlo methods to approximate all necessary integrals, thereby eliminating the need for access to the density function—particularly for the target domain—when constructing intermediate distributions. Practically, following prior works (Korotin et al., 2023; Choi et al., 2023; 2024), we can parameterize both the dual potential w and the transport map \mathbf{T} by neural networks, denoted as w_ϕ and \mathbf{T}_θ respectively. The models are trained in an alternating adversarial scheme to learn the sequence of maps $\{\mathbf{T}_{\theta,t}\}_{t=0}^{T-1}$, which can be applied to generate intermediate domains progressively.

3.3 ROBUST TRAINING PROCEDURE FOR SEMI-DUAL FORM TRANSPORTATION

While Section 3.2 provides a semi-dual form of the gradient flow problem that avoids explicit PDF estimation in target domain, naively training $\mathcal{L}^{\text{SemiDual}}$ in Eq. (7) is intrinsically unstable because of its composite ‘sup-inf’ structure. This instability is not merely algorithmic: the objective itself may be non-identifiable. We formalize this phenomenon by proving that the dual problem can have non-unique optima, as the following theorem shows:

Proposition 2. *The semi-dual formulation in Eq. (7) admits non-unique optimal solutions.*

To address this issue, we incorporate an entropy regularization term into the primal objective Eq. (6), which leads to the following proposition:

216 **Proposition 3.** Let $\kappa(x_t, x) := p(x_t) p_T(x)$ denote the reference joint PDF. The entropy-regularized
 217 primal problem is

$$\begin{aligned} 219 \quad \mathcal{L}^{E\text{-Primal}} &= \arg \min_{\rho \in \mathcal{P}_2(\mathbb{R}^D)} \frac{1}{2\eta} \mathcal{W}_2^2(\rho(x), p(x_t)) + \mathbb{D}_f[\rho(x), p_T(x)] \\ 220 \quad &+ \epsilon \iint \pi(x_t, x) [\log \frac{\pi(x_t, x)}{\kappa(x_t, x)} - 1] dx_t dx, \end{aligned} \quad (8)$$

223 and is equivalent to the semi-dual optimization problem

$$225 \quad \mathcal{L}^{E\text{-SemiDual}} = \sup_w -\epsilon \mathbb{E}_{p(x_t)} [\log \mathbb{E}_{p_T(x)} (\exp(\frac{w(x) - \frac{1}{2\eta} \|x - x_t\|_2^2}{\epsilon}))] - \mathbb{E}_{p_T(x)} [f^*(-w(x))], \quad (9)$$

227 where $w : \mathbb{R}^D \rightarrow \mathbb{R}$ and f^* are as defined in Proposition 1.

229 On this basis, we provide a theoretical guarantee of uniqueness for the semi-dual objective in Eq. (9):

230 **Proposition 4.** The semi-dual formulation in Eq. (9) admits a unique optimal solution.

231 Notably, as seen in Eq. (9), the semi-dual objective depends solely on the potential w . Consequently,
 232 we can optimize a single model, which lowers the computational burden. We therefore parameterize
 233 w by a neural network w_ϕ and carry out the optimization.

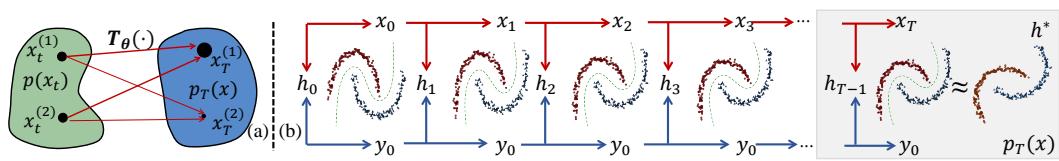
235 Finally, conditioned on the resulting w_ϕ , we subsequently optimize the transport map $\mathbf{T}_\theta(x)$ via the
 236 following objective based on Eq. (7):

$$237 \quad \arg \min_{\theta} \frac{1}{2\eta} \|x_t - \mathbf{T}_\theta(x_t)\|_2^2 - w_\phi(\mathbf{T}_\theta(x_t)). \quad (10)$$

239 Notably, we denote our approach as ‘‘E-SUOT’’, as the derivation of \mathbf{T}_θ is grounded in the Entropy-
 240 regularized Semi-dual Unbalanced Optimal Transport framework.

242 3.4 OVERALL WORKFLOW FOR E-SUOT

244 Although Sections 3.2 and 3.3 have presented the E-SUOT framework for intermediate domain
 245 generation, they do not provide a unified view of the overall workflow for generating intermediate
 246 domains. To address this, we summarize the complete procedure in Algorithm 1 (Due to page
 247 limit, the complete algorithm and other detailed information are summarized in Appendix E) and
 248 the corresponding illustration is given in Fig. 2. As shown in the algorithm, the construction of w_ϕ
 249 and \mathbf{T}_θ are performed as separate steps, corresponding to Fig. 2(a), and are illustrated in Lines 3–6
 250 and Lines 7–10, respectively. By iteratively executing the procedure described in Lines 3–10, we
 251 obtain a sequence of transport maps, $\mathcal{T} = \{\mathbf{T}_{\theta,t}\}_{t=0}^{T-1}$, which progressively transport samples from
 252 the source domain to the target domain, as we demonstrate in Fig. 2(b).



260 **Figure 2:** The illustration of the proposed E-SUOT: (a) the unbalanced OT formulation used to
 261 solve the transport map $\mathbf{T}_\theta(\cdot)$ at time t , where thicker arrows and larger points indicate higher mass
 262 flows, and (b) the evolution process from the source to the target domain (This figure is conceptually
 263 inspired by Zhuang et al. (2024)).

265 Once the transport map sequence $\mathcal{T} = \{\mathbf{T}_{\theta,t}\}_{t=0}^{T-1}$ has been obtained, we proceed to train the classi-
 266 fier h in a stage-wise manner along the transport path. Specifically, at each intermediate step t , we
 267 first map samples x_t from the current domain to the next intermediate domain x_{t+1} using the corre-
 268 sponding transport map $\mathbf{T}_{\theta,t}$. We then update or train the model h_t using the mapped data x_{t+1} as
 269 input. By iteratively applying this procedure for $t = 0, \dots, T-1$, the model is progressively adapted
 along the sequence of intermediate domains, ultimately bridging the source and target domains.

270 **Algorithm 1** Overall Workflow for Constraining E-SUOT-based Intermediate Domain Generation

271 **Input:** Source domain samples: $\{(x_0^{(i)}, y_0^{(i)})\}_{i=1}^N$, target domain samples: $\{(x_T^{(i)}, y_T^{(i)})\}_{i=1}^N$, entropy

272 regularization strength: ϵ , step size: η , number of intermediate domain $T - 1$, neural network batch

273 size \mathcal{B} , and neural network training epochs: \mathcal{E} .

274 **Output:** The set of transportation map: $\mathcal{T} = \{\mathbf{T}_{\theta,t}\}_{t=0}^{T-1}$.

275 1: $\mathcal{T} \leftarrow \emptyset$.

276 2: **for** $t = 0$ **to** $T - 1$ **do**

277 3: **for** $e = 1$ **to** \mathcal{E} **do**

278 4: Sample a batch $\{x_t^{(i)}\}_{i=1}^{\mathcal{B}} \sim \{(x_t^{(i)}, y_t^{(i)})\}_{i=1}^N$ and $\{x_T^{(i)}\}_{i=1}^{\mathcal{B}} \sim \{(x_T^{(i)}, y_T^{(i)})\}_{i=1}^N$.

279 5: Update $w_{\phi,t}$ by: $\phi \leftarrow \arg \min_{\phi} \frac{\epsilon}{\mathcal{B}} \sum_{j=1}^{\mathcal{B}} \log \frac{1}{\mathcal{B}} \sum_{i=1}^{\mathcal{B}} [\exp(\frac{w_{\phi,t}(x_T^{(j)}) - \frac{1}{2\eta} \|x_t^{(j)} - x_T^{(i)}\|_2^2}{\epsilon})] +$

280 $\frac{1}{\mathcal{B}} \sum_{j=1}^{\mathcal{B}} f^*(-w_{\phi,t}(x_T^{(j)}))$.

281 6: **end for**

282 7: **for** $e = 1$ **to** \mathcal{E} **do**

283 8: Sample a batch $\{x_t^{(i)}\}_{i=1}^{\mathcal{B}} \sim \{(x_t^{(i)}, y_t^{(i)})\}_{i=1}^N$.

284 9: Update $\mathbf{T}_{\theta,t}$ by: $\theta \leftarrow \arg \min_{\theta} \frac{1}{\mathcal{B}} \sum_{i=1}^{\mathcal{B}} \frac{1}{2\eta} \|x_t^{(i)} - \mathbf{T}_{\theta,t}(x_t^{(i)})\|_2^2 - w_{\phi,t}(\mathbf{T}_{\theta,t}(x_t^{(i)}))$.

285 10: **end for**

286 11: $x_{t+1}^{(i)} \leftarrow \mathbf{T}_{\theta,t}(x_t^{(i)}), \forall i \in \{1, \dots, N\}$.

287 12: $\mathcal{T} \leftarrow \mathcal{T} \cup \{\mathbf{T}_{\theta,t}\}$

288 13: **end for**

291 3.5 THEORETICAL ANALYSIS

293 Notably, our derivation sidesteps the explicit estimation of the PDF of the target domain by leveraging
 294 the semi-dual formulation. This naturally leads to two important questions: (1) Can the proposed
 295 E-SUOT framework transport the source domain sufficiently close to the target domain? (2) How
 296 does the model perform on the target domain after transport?

297 To address the first question, we present the following theorem, which quantitatively characterizes
 298 the discrepancy between $\rho(x)$ and $p_T(x)$:

300 **Theorem 5.** *The optimal solution $\rho^*(x)$ to problem defined in Eq. (8) satisfies the following bound:*

$$301 \mathbb{D}_f[\rho^*(x), p_T(x)] \leq \mathcal{W}_2(p(x_t), p_T(x)). \quad (11)$$

302 From Theorem 5, we observe that as t increases, the transported PDF $\rho(x)$ progressively becomes
 303 similar to $p_T(x)$. Based on this result, we present the following theorem, which provides a theoretical
 304 guarantee for the model’s performance on the target domain:

306 **Theorem 6.** *Under mild assumptions, the E-SUOT-based GDA ensures that the target domain generalization error is upper-bounded by the following inequality:*

$$308 \varepsilon_{p_T}(h_T) \leq \varepsilon_{p_0}(h_0) + \varepsilon_{p_0}(h_T^*) + \iota \zeta \mathcal{C} + \mathcal{S}_{\text{stat}}, \quad (12)$$

309 where ι is the Lipschitz constant of the loss function, ζ is the Lipschitz constant bound for hypotheses
 310 in \mathcal{H} , \mathcal{C} aggregates the cumulative domain transportation and label continuity costs along the
 311 adaptation path, and $\mathcal{S}_{\text{stat}}$ is the statistical error term.

313 4 EXPERIMENTAL RESULTS

314 4.1 EXPERIMENTAL SETUP

317 **Datasets:** We conduct case studies on four datasets. Specifically, for GDA task, the datasets are
 318 “Portraits” (Kumar et al., 2020), “MNIST 45°” and “MNIST 60°” (LeCun, 1998; Deng, 2012).
 319 For UDA task, we conduct experiment on the “Office-Home” dataset (Venkateswara et al., 2017).
 320 Detailed information about these datasets is given in Appendix E.1.

321 **Implementation:** Following prior work (Zhuang et al., 2024; Sagawa & Hino, 2025), we employ
 322 semi-supervised UMAP to produce low-dimensional embeddings while preserving class discrim-
 323 inability. Unless stated otherwise, we use the KL divergence in the implementation of the E-SUOT.
 324 Additional details are available in Appendix E.2.

324 4.2 BASELINE COMPARISON RESULTS
325

326 We first compare our proposed approach with several existing GDA-based methods, including Self-
327 training, GST (4 intermediate domains) (Kumar et al., 2020), GOAT (He et al., 2024), CNF (Sagawa
328 & Hino, 2025), and GGF (Zhuang et al., 2024). The detailed information of the experiments are
329 provided in Section E. All baseline models for the GDA task are evaluated with five groups of
330 random seeds, each repeated three times.

331 Table 1: Baseline comparison on the GDA task with standard deviation and p -value.
332

333 Method	334 Portraits			335 MNIST 45°			336 MNIST 60°		
	337 Accuracy (%)	338 Δ	339 p -value	340 Accuracy (%)	341 Δ	342 p -value	343 Accuracy (%)	344 Δ	345 p -value
Source	71.2	-	-	58.4	-	-	36.8	-	-
Self Train	<u>77.4</u> [*] \pm 5.02E-2	\uparrow 8.7%	1.25E-47	<u>58.7</u> [*] \pm 2.24E-2	\uparrow 0.5%	4.99E-50	<u>39.9</u> [*] \pm 2.00E-2	\uparrow 8.5%	1.19E-48
GST (4)	<u>76.1</u> [*] \pm 6.00E-2	\uparrow 6.9%	1.98E-21	<u>59.2</u> [*] \pm 2.45E-2	\uparrow 1.3%	1.21E-22	<u>39.9</u> [*] \pm 1.00E-2	\uparrow 8.5%	2.14E-23
GOAT	<u>74.9</u> [*] \pm 6.21E-1	\uparrow 5.3%	2.27E-17	<u>65.0</u> [*] \pm 1.05E-1	\uparrow 11.3%	4.07E-20	<u>37.2</u> [*] \pm 8.43E-2	\uparrow 1.1%	1.08E-19
CNF	80.0* \pm 1.85E0	\uparrow 12.4%	1.79E-15	<u>57.6</u> [*] \pm 1.08E0	\downarrow 1.4%	4.60E-16	<u>41.8</u> [*] \pm 1.92E0	\uparrow 13.5%	2.90E-14
GGF	<u>83.4</u> [*] \pm 8.79E-1	\uparrow 17.2%	9.18E-17	<u>57.7</u> [*] \pm 6.55E-1	\downarrow 1.2%	6.15E-17	<u>40.8</u> [*] \pm 8.35E-1	\uparrow 11.0%	1.04E-15
E-SUOT	86.4 [*] \pm 8.72E-2	\uparrow 21.5%	8.88E-21	72.1 [*] \pm 4.62E-1	\uparrow 23.4%	1.55E-17	51.0 [*] \pm 5.81E-1	\uparrow 38.6%	2.48E-16

346 *Kindly Note:* “*” marks the variants that E-SUOT outperforms significantly at p -value < 0.05 over paired
347 sample t -test. Δ denotes the performance change relative to the initial classifier. The accuracy is reported as
348 mean (%) \pm 1.0 \times standard deviation error. In addition, **Bolded** results are the best results. Underlined results
349 are the second best results.

350 As shown in Table 1, our proposed E-SUOT framework consistently outperforms the current state-
351 of-the-art GDA approaches on all evaluated datasets. These results demonstrate the effectiveness
352 and superiority of the E-SUOT framework. In addition, we observe that flow-based methods, such
353 as CNF and GGF, generally achieve top-2 performance on most datasets, highlighting the potential
354 of incorporating flow-based methods in GDA tasks. However, we also note that flow-based methods,
355 occasionally underperform. This observation suggests that flow-based GDA, which requires explicit
356 PDF estimation on target domain, may have inherent limitations, as discussed in Section 3.1.

357 On this basis, we further evaluate the performance of E-SUOT under the UDA task by comparing
358 it to the DANN (Ganin & Lempitsky, 2015), MSTN (Xie et al., 2018), GVB-GD (Cui et al.,
359 2020), RSDA (Gu et al., 2020; 2022), LAMBDA (Le et al., 2021), SENTRY (Prabhu et al., 2021),
360 FixBi (Na et al., 2021), CST (Liu et al., 2021a), CoVi (Na et al., 2022), and GGF (Zhuang et al.,
361 2024) on the Office-Home dataset (Venkateswara et al., 2017). Detailed information on the ex-
362 periments is provided in the Section E.2. The results, summarized in Table 2, show that E-SUOT
363 outperforms most existing methods on the majority of UDA tasks and achieves the highest overall
364 average performance. Despite not achieving top performance on every individual task, E-SUOT
365 attains the highest average result across the board, confirming its overall superiority in UDA task.

366 Table 2: Accuracy (%) comparison on the Office-Home dataset under the UDA setting.

367 Method	368 Ar \rightarrow Cl	369 Ar \rightarrow Pr	370 Ar \rightarrow Rw	371 Cl \rightarrow Ar	372 Cl \rightarrow Pr	373 Cl \rightarrow Rw	374 Pr \rightarrow Ar	375 Pr \rightarrow Cl	376 Pr \rightarrow Rw	377 Rw \rightarrow Ar	378 Rw \rightarrow Cl	379 Rw \rightarrow Pr	380 Avg.
DANN	45.6	59.3	70.1	47.0	58.5	60.9	46.1	43.7	68.5	63.2	51.8	76.8	57.6
MSTN	49.8	70.3	76.3	60.4	68.5	69.6	61.4	48.9	75.7	70.9	55.0	81.1	65.7
GVB-GD	57.0	74.7	79.8	64.6	74.1	74.6	65.2	55.1	81.0	74.6	59.7	84.3	70.4
RSDA	53.2	77.7	81.3	66.4	74.0	76.5	<u>67.9</u>	53.0	82.0	75.8	57.8	85.4	70.9
LAMDA	57.2	78.4	<u>82.6</u>	66.1	80.2	81.2	65.6	55.1	82.8	71.6	59.2	83.9	72.0
SENTRY	61.8	77.4	80.1	66.3	71.6	74.7	66.8	63.0	80.9	74.0	66.3	84.1	72.3
FixBi	58.1	77.3	80.4	67.7	79.5	<u>78.1</u>	65.8	57.9	81.7	<u>76.4</u>	62.9	86.7	72.7
CST	59.0	79.6	83.4	68.4	77.1	76.7	68.9	56.4	83.0	75.3	62.2	85.1	72.9
CoVi	58.5	78.1	80.0	<u>68.1</u>	<u>80.0</u>	77.0	66.4	60.2	82.1	76.6	<u>63.6</u>	<u>86.5</u>	73.1
GGF	59.4	75.6	81.7	67.6	77.6	78.0	67.4	61.0	82.7	75.9	62.4	85.4	72.9
E-SUOT	<u>61.6</u>	<u>79.3</u>	81.8	67.6	77.7	78.1	67.4	<u>61.2</u>	<u>82.9</u>	76.3	62.5	85.2	73.5
Win Counts	9	9	8	6	7	8	7	9	9	8	7	6	10

381 *Kindly Note:* **Bolded** and underlined results are the first and second best results, respectively.
382

378
379

4.3 INVESTIGATION OF THE UOT FORMULATION

380
381
382
383
384
385
386
387
388
389
390
391
392

While our main contribution lies in introducing the semi-dual UOT formulation to analyze and improve the flow-based GDA approach, we further investigate “why the UOT-based method performs better than the vanilla OT formulation”. To this end, we conduct experiments under label-shift and missing-class scenarios using the Portrait dataset (binary classification task). Specifically, we resample the target domain to vary the class prior $p(y = 1)$; when $p(y = 1) = 0.0$ or $p(y = 1) = 1.0$, a missing-class situation is realized. The comparison results between vanilla OT and UOT are reported in Table 3. Here, “E-SOT” denotes entropy-regularized semi-dual optimal transport. From the table, we observe that in the missing-class cases ($p(y = 1) = 0.0$ or $p(y = 1) = 1.0$), the vanilla OT formulation not only fails to improve but even degrades the classifier’s performance in the target domain. Moreover, under moderate label shift ($p(y = 1)$ between 0.3 and 0.9), the performance of vanilla OT fluctuates strongly and lacks stability, whereas UOT consistently improves performance. These observations demonstrate that incorporating the unbalanced OT formulation provides a more robust and effective approach for handling domain adaptation under label distribution mismatch.

393 Table 3: Comparison of vanilla and unbalanced OT formulations for GDA task on Portraits dataset.

Method	$p(y = 1) = 0.0$	$p(y = 1) = 0.1$	$p(y = 1) = 0.2$	$p(y = 1) = 0.3$	$p(y = 1) = 0.4$	
	Accuracy (%) Δ					
Initial	35.4	-	41	-	47.1	-
E-SOT	$55.4 \pm 3.15E-1$	$\uparrow 56.41\%$	$61.1 \pm 2.01E-1$	$\uparrow 48.94\%$	$67.1 \pm 2.45E-1$	$\uparrow 42.55\%$
E-SUOT	$64.5 \pm 9.09E-2$	$\uparrow 82.32\%$	$78.2 \pm 6.55E-2$	$\uparrow 90.50\%$	$74.5 \pm 1.02E-1$	$\uparrow 58.28\%$
Method	$p(y = 1) = 0.6$	$p(y = 1) = 0.7$	$p(y = 1) = 0.8$	$p(y = 1) = 0.9$	$p(y = 1) = 1.0$	
	Accuracy (%) Δ					
Initial	73.8	-	80	-	85.9	-
E-SOT	$75.2 \pm 3.08E-2$	$\uparrow 1.95\%$	$74.1 \pm 4.80E-2$	$\downarrow 7.33\%$	$80.0 \pm 2.21E-2$	$\downarrow 6.89\%$
E-SUOT	$79.1 \pm 2.56E-2$	$\uparrow 7.24\%$	$84.0 \pm 3.46E-2$	$\uparrow 5.05\%$	$87.8 \pm 1.12E-2$	$\uparrow 2.18\%$

402 *Kindly Note:* Δ denotes performance change percentage compared to E-SUOT with entropy regularization and KL divergence. The
403 accuracy is reported as mean (%) $\pm 1 \times$ standard deviation error. For source domain, $p(y = 1) = 0.63$. The acronym E-SOT stands
404 for entropy-regularized semi-dual optimal transport.

405
406

4.4 ABLATION STUDIES

407
408
409
410
411
412
413
414

We perform ablation studies from two perspectives: the training strategy for T_θ and the choice of f -divergence. For the *training strategy*, we 1). examine the effect of removing the entropy regularization term—reducing the method to the adversarial training strategy in Eq. (7), and 2). evaluate a barycentric projection approach analogous to flow matching (Lipman et al., 2023), where the transport plan is first estimated and then used to project source samples toward the target, subsequently being refined during training. For the objective *functional*, we study different parameterizations of f^* , such as employing non-decreasing convex functions like 1) Softplus, and also compare the 2) χ^2 divergence and the 3) identity function. More detailed information on these experiments’ implementation is provided in Appendix E.3. The ablation study results are summarized in Table 4.

415

Table 4: Ablation study results on GDA setting with standard deviation and p -value.

Dataset	Portraits		MNIST 45°		MNIST 60°						
	Metric	Accuracy (%) Δ	p-value	Accuracy (%) Δ	p-value	Accuracy (%) Δ	p-value				
Training	Adversarial	KL	$74.8^* \pm 3.10E0$	$\downarrow 13.4\%$	$1.83E-03$	$52.0^* \pm 3.60E0$	$\downarrow 27.8\%$	$3.05E-04$	$34.9^* \pm 4.10E0$	$\downarrow 31.5\%$	$4.59E-03$
	Barycentric	KL	$83.9^* \pm 9.42E-1$	$\downarrow 3.0\%$	$4.05E-03$	$62.5^* \pm 1.06E0$	$\downarrow 13.3\%$	$1.77E-04$	$38.3^* \pm 4.46E0$	$\downarrow 24.8\%$	$6.57E-03$
Functional	Entropy	Softplus	$80.1^* \pm 3.53E0$	$\downarrow 7.3\%$	$2.24E-02$	$59.7^* \pm 1.14E0$	$\downarrow 17.2\%$	$7.08E-05$	$38.2^* \pm 1.10E0$	$\downarrow 25.1\%$	$6.27E-06$
	Entropy	χ^2	$79.8 \pm 6.07E0$	$\downarrow 7.7\%$	$9.57E-02$	$60.2^* \pm 1.43E0$	$\downarrow 16.5\%$	$8.03E-05$	$42.4^* \pm 3.47E0$	$\downarrow 16.9\%$	$7.19E-03$
	Entropy	Identity	$81.2^* \pm 1.71E0$	$\downarrow 6.1\%$	$3.13E-03$	$59.6^* \pm 1.25E0$	$\downarrow 17.4\%$	$8.62E-05$	$39.6^* \pm 2.34E0$	$\downarrow 22.3\%$	$2.59E-04$
	Entropy	KL	$86.4 \pm 8.72E-2$	-	-	$72.1 \pm 4.62E-1$	-	-	$51.0 \pm 5.81E-1$	-	-

422
423

Kindly Note: “*” marks the variants that E-SUOT outperforms significantly at p -value < 0.05 over paired sample t -test. Δ denotes
424 performance change percentage compared to E-SUOT with entropy regularization and KL divergence. The accuracy is reported as mean
425 (%) $\pm 1 \times$ standard deviation error.

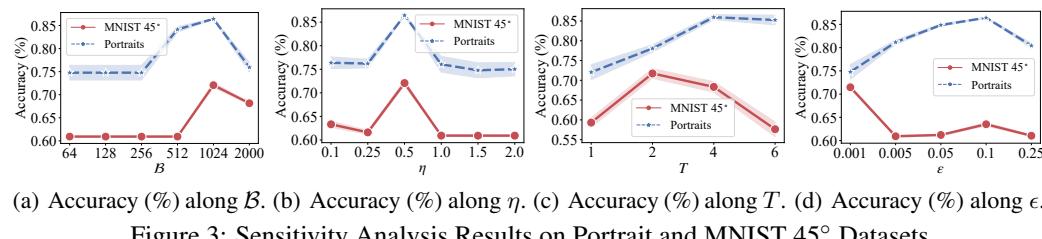
426
427
428
429
430
431

From Table 4, we find that adversarial training performs the worst, underscoring the importance of entropy regularization for model training in Section 3.3. While barycentric mapping is competitive, it struggles on complex datasets such as MNIST 45° and MNIST 60°, highlighting the need for the semi-dual formulation. Additionally, alternatives to KL divergence—especially Softplus—cause significant performance drops, emphasizing the importance of proper divergence selection. We also observe that replacing KL divergence with alternatives such as χ^2 divergence, the identity function, or particularly Softplus results in substantial performance degradation, further illustrating that choosing a suitable discrepancy to drive the evolution of source domain to target domain is critical for promising the performance of GDA.

432 4.5 SENSITIVITY ANALYSIS
433

434 From Figs. 3(a) to 3(d), we systematically investigate the sensitivity of our E-SUOT model with
435 respect to key hyperparameters, including batch size \mathcal{B} , discretization step size η , simulation steps
436 T , and entropy regularization strength ϵ on the Portraits and MNIST 45° datasets.

437 Specifically, as shown in Fig. 3(a), we observe that increasing the batch size \mathcal{B} initially improves
438 model performance; however, after a certain point, further increasing the batch size leads to a perfor-
439 mance decline. This pattern suggests that, in the simulation of WGF-based approaches (including
440 ours), careful selection of batch size is crucial: if \mathcal{B} is too small, stochastic sampling noise may
441 dominate and degrade the results; conversely, excessively large \mathcal{B} can cause the model to overfit
442 and diminish its performance. A similar trend is found when varying the discretization step size η ,
443 as illustrated in Fig. 3(b). A small step size may prevent the simulation trajectory from adequately
444 reaching the target distribution within a finite number of steps, limiting learning efficiency. On the
445 other hand, a step size that is too large introduces significant discretization error, which again results
446 in poor model performance. Furthermore, as demonstrated in Fig. 3(c), increasing the number of
447 simulation steps T also produces a non-monotonic effect: beyond a certain threshold, more steps
448 actually undermine performance. This is likely because aligning the feature/target distributions too
449 strictly does not necessarily correspond to optimal performance in the target domain, thus further
450 justifying our introduction of divergence-based regularization to relax strict alignment constraints
451 compared to traditional OT-based methods. Finally, as shown in Fig. 3(d), the entropy regulariza-
452 tion parameter ϵ also significantly influences results. We observe that varying ϵ can lead to diverse
453 performance outcomes, highlighting the importance of properly investigating and tuning the entropy
454 regularization strength in practical applications. In conclusion, our sensitivity study underscores the
455 importance of carefully selecting the batch size \mathcal{B} , step size η , and end time T for E-SUOT perfor-
456 mance, and further indicates that the entropy regularization strength ϵ is dataset-dependent and thus
457 warrants systematic validation on the target dataset to achieve optimal E-SUOT performance.



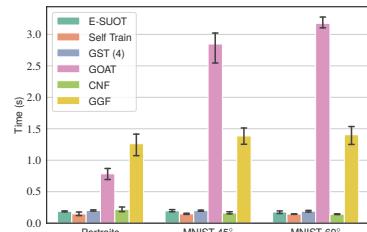
463 (a) Accuracy (%) along \mathcal{B} . (b) Accuracy (%) along η . (c) Accuracy (%) along T . (d) Accuracy (%) along ϵ .
464

465 Figure 3: Sensitivity Analysis Results on Portrait and MNIST 45° Datasets.
466

467 4.6 COMPUTATIONAL TIME COMPARISON
468

469 In this subsection, we further analyze the empirical time complexity of the proposed E-SUOT ap-
470 proach in comparison with alternative methods on the GDA task. The computational time results
471 are presented in Fig. 4.

472 As shown in Fig. 4, the GOAT approach is the most time-
473 consuming on larger datasets, while GGF takes more time
474 on smaller datasets; both consistently rank among the top
475 two in terms of computation cost. This can be attributed to
476 their inherent algorithmic structures: GOAT involves solv-
477 ing the exact optimal transport problem, which becomes
478 computationally prohibitive as the dataset size increases. In
479 contrast, GGF relies on the forward Euler method, which
480 requires a very small step size—and therefore a large num-
481 ber of iterations—to avoid significant simulation errors, re-
482 sulting in higher computational overhead even on smaller
483 datasets. Notably, the computational time of our proposed
484 E-SUOT remains stable as dataset size grows. This effi-
485 ciency stems from directly parameterizing the transport map
486 using a single forward pass through a neural network and the
487 JKO scheme, a variant of backward discretization approach, re-
488quiring only a few steps to achieve the desired performance.



489 Figure 4: Computational time (s).
490

486 5 RELATED WORKS
487488 5.1 GRADUAL DOMAIN ADAPTATION
489

490 GDA seeks to bridge the distributional gap between source and target domains by leveraging a se-
491 quence of intermediate domains, thereby enabling more fine-grained adaptation. Early works have
492 explored self-training strategies (Kumar et al., 2020), adversarial objectives (Wang et al., 2020), and
493 provided generalization bounds under gradual distribution shifts (Kumar et al., 2020; Dong et al.,
494 2022; Wang et al., 2022). However, these approaches often depend on the availability of discrete
495 intermediate domains (Chen & Chao, 2021). To address this, optimal transport approaches (Ab-
496 nar et al., 2021; He et al., 2024) have been leveraged to construct intermediate domains along the
497 Wasserstein geodesic, ensuring minimal distributional discrepancy in the adaptation process. More
498 recently, flow-based GDA has emerged, which explicitly models domain evolution and synthesizes
499 continuous intermediate distributions via parametric flows. For instance, Sagawa & Hino (2025)
500 uses continuous normalizing flows to parameterize domain trajectories as ODEs in the data space,
501 while Zhuang et al. (2024) incorporates label information into this evolution and employs gradient
502 flows to realize the steepest transformation from source to target domain. Nevertheless, flow-based
503 methods still require explicit estimation of the target domain’s PDF to guide the evolution, and
504 inaccuracies in this estimation can lead to performance drops in target domain. To address this lim-
505 itation, we reformulate the flow-based approach from a semi-dual formulation (see Proposition 1),
506 which unifies the flow-based and optimal transport methods. Building on this, we further propose a
507 convergence-guaranteed approach with the help of entropy regularization (Proposition 3) and ana-
508 lyze its generalization error (see Theorem 6). during the evolution of the flow and proposed gradient
509

510 5.2 SEMI-DUAL FORMULATION OF GRADIENT FLOWS
511

512 Gradient flow (Santambrogio, 2017), which seeks to optimize a specified functional in the space of
513 probability measures, has played a critical role in both sampling and optimization algorithm design.
514 For gradient flows induced by f -divergences (with the KL divergence being the notable example),
515 such as Langevin sampling (Welling & Teh, 2011), have been extensively explored to generate sam-
516 ples that progressively transition from the source domain toward the target domain. However, these
517 methods typically assume access to an exact (unnormalized) PDF for the target distribution (Liu &
518 Wang, 2016; Liu, 2017), which is often infeasible in practice when only samples are available. To
519 overcome this, several approaches have explored dual formulations of f -divergence (Nguyen et al.,
520 2007; 2010), which avoid explicit density estimation for the target domain and instead optimize pri-
521 mal formulation (Korotin et al., 2023; Rout et al., 2022; Fan et al., 2022; Gazdieveva et al., 2023; Choi
522 et al., 2023; 2024). These dual-formulation methods, however, generally require adversarial opti-
523 mization characterized by a composite “sup-inf” structure in order to properly approximate the dual
524 objective when implemented with neural networks (Nowozin et al., 2016; Arjovsky et al., 2017).
525 Our work differs from these approaches in two key aspects. First, we provide a theoretical analysis
526 from the perspective of the non-uniqueness of optimal solutions in Proposition 2, highlighting that
527 such adversarial formulations can suffer from this issue, which may hinder training stability. Build-
528 ing upon this insight, we introduce the entropy regularization that transforms the adversarial game
529 into an alternative paradigm in Proposition 3, and further prove that this regularization ensures the
530 stability via the uniqueness of the optima in Proposition 4 and convergence in Theorem 5.

531 6 CONCLUSIONS
532

533 In this paper, we addressed the challenge in flow-based GDA, namely the reliance on explicit esti-
534 mation of the target domain PDF inherited from traditional f -divergence formulations. To overcome
535 this, we reformulated the flow simulation as an optimization problem augmented with a Wasserstein
536 regularization term. Building on this, we derived a novel semi-dual formulation that avoids explicit
537 estimation of the target density. However, we observed that the resulting semi-dual structure intro-
538 duces instability due to its composite ‘sup-inf’ structure. To address this, we proposed an entropy
539 regularization term that eliminates the inner inf operator, thereby restoring stability and ensuring
540 uniqueness of the optimal solution. Based on these insights, we developed a new GDA framework
541 called “E-SUOT” and provided theoretical guarantees for its convergence and generalization. Fi-
542 nally, extensive experiments validate the effectiveness and practical advantages of our approach.

540 ETHICS STATEMENT
541

542 The authors have read and comply with the ICLR Code of Ethics. This research does not involve
543 human subjects or personally identifiable information, and uses only publicly available datasets under
544 their respective licenses. We do not foresee harmful or dual-use implications from the proposed
545 methods. There are no conflicts of interest or undisclosed sponsorship.

546
547 REPRODUCIBILITY STATEMENT
548

549 The anonymous downloadable source code is available at: https://anonymous.4open.science/r/E_SUOT_GDA-9240/. For theoretical results, the derivations proof of the claims
550 are included in Appendix C. Based on this, a detailed overall workflow for the proposed E-SUOT is
551 summarized in Appendix D. For datasets used in our experiments, we provide a complete description
552 of the dataset statistics and processing work flow in Appendix E.

553 REFERENCES
554

555 Samira Abnar, Rianne van den Berg, Golnaz Ghiasi, Mostafa Dehghani, Nal Kalchbrenner, and
556 Hanie Sedghi. Gradual domain adaptation in the wild: When intermediate distributions are absent.
557 *arXiv preprint arXiv:2106.06080*, pp. 1–15, 2021.

558 Ralph Abraham, Jerrold E Marsden, and Tudor Ratiu. *Manifolds, tensor analysis, and applications*,
559 volume 75. Springer Science & Business Media, 2012.

560 Shun-ichi Amari. *Information geometry and its applications*, volume 194. Springer, 2016.

561 Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the*
562 *space of probability measures*. Springer, 2005.

563 Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks.
564 In *Proc. Int. Conf. Mach. Learn.*, pp. 214–223. PMLR, 2017.

565 Francis Bach. *Learning theory from first principles*. MIT Press, Cambridge, USA, 2024.

566 Heinz H Bauschke and Patrick L Combettes. Fenchel–rockafellar duality. In *Convex Analysis and*
567 *Monotone Operator Theory in Hilbert Spaces*, pp. 247–262. Springer, 2017.

568 Lorenz T Biegler. *Nonlinear programming: concepts, algorithms, and applications to chemical*
569 *processes*. SIAM, 2010.

570 Clément Bonet, Christophe Vauthier, and Anna Korba. Flowing datasets with wasserstein over
571 wasserstein gradient flows. In *Proc. Int. Conf. Mach. Learn.*, pp. 1–57. PMLR, 2025.

572 Paul Brunzema, Mikkel Jordahn, John Willes, Sebastian Trimpe, Jasper Snoek, and James Harrison.
573 Bayesian optimization via continual variational last layer training. In *Proc. Int. Conf. Learn.*
574 *Represent.*, pp. 1–30, 2025.

575 John Charles Butcher. *Numerical methods for ordinary differential equations*. John Wiley & Sons,
576 New York, USA, 2016.

577 Kenneth F. Caluya and Abhishek Halder. Gradient flow algorithms for density propagation in
578 stochastic systems. *IEEE Trans. Autom. Control.*, 65(10):3991–4004, 2020.

579 Kenneth F. Caluya and Abhishek Halder. Wasserstein proximal algorithms for the Schrödinger
580 bridge problem: Density control with nonlinear drift. *IEEE Trans. Autom. Control.*, 67(3):1163–
581 1178, 2022. doi: 10.1109/TAC.2021.3060704.

582 Hong-You Chen and Wei-Lun Chao. Gradual domain adaptation without indexed intermediate do-
583 mains. In *Proc. Adv. Neural inf. Process. Syst.*, volume 34, pp. 8201–8214, 2021.

584 Lenaic Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Unbalanced op-
585 timal transport: Dynamic and kantorovich formulations. *J. Funct. Anal.*, 274(11):3090–3123,
586 2018.

594 Jaemoo Choi, Jaewoong Choi, and Myungjoo Kang. Generative modeling through the semi-dual
 595 formulation of unbalanced optimal transport. In *Proc. Adv. Neural Inf. Process. Syst.*, volume 36,
 596 pp. 42433–42455, 2023.

597 Jaemoo Choi, Jaewoong Choi, and Myungjoo Kang. Scalable wasserstein gradient flow for gen-
 598 erative modeling through unbalanced optimal transport. In *Proc. Int. Conf. Mach. Learn.*, pp.
 600 8629–8650. PMLR, 2024.

601 Jaemoo Choi, Jaewoong Choi, and Dohyun Kwon. Overcoming spurious solutions in semi-dual
 602 neural optimal transport: A smoothing approach for learning the optimal transport plan. In *Proc.*
 603 *Int. Conf. Mach. Learn.*, pp. 1–21, 2025.

604 Nicolas Courty, Rémi Flamary, and Devis Tuia. Domain adaptation with regularized optimal trans-
 605 port. In *Joint European conference on machine learning and knowledge discovery in databases*,
 606 pp. 274–289. Springer, 2014.

607 Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. Joint distribution
 608 optimal transportation for domain adaptation. In *Proc. Adv. Neural Inf. Process. Syst.*, volume 30,
 609 pp. 1–10, 2017a.

610 Nicolas Courty, Rémi Flamary, Devis Tuia, and Alain Rakotomamonjy. Optimal transport for
 611 domain adaptation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(9):1853–1865, 2017b. doi:
 612 10.1109/TPAMI.2016.2615921.

613 Shuhao Cui, Shuhui Wang, Junbao Zhuo, Chi Su, Qingming Huang, and Qi Tian. Gradually vanishing
 614 bridge for adversarial domain adaptation. In *Proc. Int. Conf. Mach. Learn.*, pp. 12455–12464,
 615 2020.

616 Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In C.J. Burges,
 617 L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger (eds.), *Proc. Adv. Neural Inf. Pro-
 618 cess. Syst.*, volume 26, pp. 1–9. Curran Associates, Inc., 2013.

619 Li Deng. The MNIST database of handwritten digit images for machine learning research [best of
 620 the web]. *IEEE Signal Process. Mag.*, 29(6):141–142, 2012. doi: 10.1109/MSP.2012.2211477.

621 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *Proc.*
 622 *Adv. Neural Inf. Process. Syst.*, volume 34, pp. 8780–8794, 2021.

623 Jing Dong, Shiji Zhou, Baoxiang Wang, and Han Zhao. Algorithms and theory for supervised
 624 gradual domain adaptation. *Trans. Mach. Learn. Res.*, pp. 1–14, 2022.

625 Lawrence C Evans. *Partial Differential Equations*, volume 19. American Mathematical Society,
 626 Amsterdam, Netherlands, 2022.

627 Jiaojiao Fan, Qinsheng Zhang, Amirhossein Taghvaei, and Yongxin Chen. Variational Wasserstein
 628 gradient flow. In *Proc. Int. Conf. Mach. Learn.*, pp. 6185–6215. PMLR, 2022.

629 Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In
 630 *Proc. Int. Conf. Mach. Learn.*, pp. 1–10. PMLR, 2015.

631 Milena Gazdieva, Alexander Korotin, Daniil Selikhnovych, and Evgeny Burnaev. Extremal domain
 632 translation with neural optimal transport. In *Proc. Adv. Neural Inf. Process. Syst.*, volume 36, pp.
 633 40381–40413, 2023.

634 Pierre Glaser, Michael Arbel, and Arthur Gretton. Kale flow: A relaxed kl gradient flow for proba-
 635 bilities with disjoint support. In *Proc. Adv. Neural Inf. Process. Syst.*, volume 34, pp. 8018–8031,
 636 2021.

637 Xiang Gu, Jian Sun, and Zongben Xu. Spherical space domain adaptation with robust pseudo-label
 638 loss. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, June 2020.

639 Xiang Gu, Jian Sun, and Zongben Xu. Unsupervised and semi-supervised robust spherical space
 640 domain adaptation. *IEEE Trans. Pattern Anal. Mach. Intell.*, pp. 1–1, 2022. doi: 10.1109/TPAMI.
 641 2022.3158637.

648 James Harrison, John Willes, and Jasper Snoek. Variational Bayesian last layers. In *Proc. Int. Conf.*
 649 *Learn. Represent.*, pp. 1–30, 2024.

650

651 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recogni-
 652 tion. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, pp. 770–778, 2016.

653

654 Yifei He, Haoxiang Wang, Bo Li, and Han Zhao. Gradual domain adaptation: Theory and algo-
 655 rithms. *J. Mach. Learn. Res.*, 25(361):1–40, 2024.

656

657 Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen,
 658 et al. LoRA: Low-rank adaptation of large language models. In *Proc. Int. Conf. Learn. Represent.*,
 659 pp. 1–13, 2022.

660

661 Rie Johnson and Tong Zhang. Composite functional gradient learning of generative adversarial
 662 models. In *Proc. Int. Conf. Mach. Learn.*, pp. 2371–2379. PMLR, 2018.

663

664 Rie Johnson and Tong Zhang. A framework of composite functional gradient methods for generative
 665 adversarial models. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(1):17–32, 2021. doi: 10.1109/
 666 TPAMI.2019.2924428.

667

668 Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the Fokker–
 669 Planck equation. *SIAM J. Math. Anal.*, 29(1):1–17, 1998.

670

671 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proc. Int. Conf.*
 672 *Learn. Represent.*, pp. 1–8, 2015.

673

674 Ivan Kobyzev, Simon JD Prince, and Marcus A Brubaker. Normalizing flows: An introduction and
 675 review of current methods. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(11):3964–3979, 2020.

676

677 Alexander Korotin, Lingxiao Li, Aude Genevay, Justin M Solomon, Alexander Filippov, and Evgeny
 678 Burnaev. Do neural optimal transport solvers work? a continuous wasserstein-2 benchmark. In
 679 *Proc. Adv. Neural Inf. Process. Syst.*, volume 34, pp. 14593–14605, 2021.

680

681 Alexander Korotin, Daniil Selikhanovich, and Evgeny Burnaev. Neural optimal transport. In *Proc.*
 682 *Int. Conf. Learn. Represent.*, pp. 1–34, 2023.

683

684 Ananya Kumar, Tengyu Ma, and Percy Liang. Understanding self-training for gradual domain
 685 adaptation. In *Proc. Int. Conf. Mach. Learn.*, pp. 5468–5479. PMLR, 2020.

686

687 Trung Le, Tuan Nguyen, Nhat Ho, Hung Bui, and Dinh Phung. Lamda: Label matching deep
 688 domain adaptation. In *Proc. Int. Conf. Mach. Learn.*, pp. 6043–6054. PMLR, 2021.

689

690 Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.

691

692 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow match-
 693 ing for generative modeling. In *Proc. Int. Conf. Learn. Represent.*, pp. 1–28, 2023.

694

695 Chang Liu, Jingwei Zhuo, Pengyu Cheng, Ruiyi Zhang, and Jun Zhu. Understanding and accelerat-
 696 ing particle-based variational inference. In *Proc. Int. Conf. Mach. Learn.*, pp. 4082–4092. PMLR,
 697 2019.

698

699 Hong Liu, Jianmin Wang, and Mingsheng Long. Cycle self-training for domain adaptation. *Proc.*
 700 *Adv. Neural Inf. Process. Syst.*, 34:22968–22981, 2021a.

701

702 Qiang Liu. Stein variational gradient descent as gradient flow. In *Proc. Adv. Neural inf. Process.*
 703 *Syst.*, volume 30, pp. 1–15, 2017.

704

705 Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference
 706 algorithm. In *Proc. Adv. Neural Inf. Process. Syst.*, volume 29, pp. 1–9, 2016.

707

708 Weiming Liu, Jiajie Su, Chaochao Chen, and Xiaolin Zheng. Leveraging distribution alignment via
 709 stein path for cross-domain cold-start recommendation. In *Proc. Adv. Neural Inf. Process. Syst.*,
 710 volume 34, pp. 19223–19234, 2021b.

702 Weiming Liu, Xiaolin Zheng, Jiajie Su, Longfei Zheng, Chaochao Chen, and Mengling Hu. Contrastive proxy kernel Stein path alignment for cross-domain cold-start recommendation. *IEEE Trans. Knowl. Data Eng.*, 35(11):11216–11230, 2023. doi: 10.1109/TKDE.2022.3233789.

703

704

705 Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *Proc. Int. Conf. Mach. Learn.*, pp. 1–9. PMLR, 2015.

706

707

708 Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, pp. 1–63, 2018.

709

710

711 Jaemin Na, Heechul Jung, Hyung Jin Chang, and Wonjun Hwang. FixBi: Bridging domain spaces for unsupervised domain adaptation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, pp. 1094–1103, 2021.

712

713

714 Jaemin Na, Dongyoon Han, Hyung Jin Chang, and Wonjun Hwang. Contrastive vicinal space for unsupervised domain adaptation. In *Proc. Eur. Conf. Comput. Vis.*, pp. 92–110. Springer, 2022.

715

716 Kirill Neklyudov, Jannes Nys, Luca Thiede, Juan Carrasquilla, Qiang Liu, Max Welling, and Alireza Makhzani. Wasserstein quantum Monte Carlo: a novel approach for solving the quantum many-body schrödinger equation. In *Proc. Adv. Neural Inf. Process. Syst.*, volume 36, pp. 63461–63482, 2023.

717

718

719

720 XuanLong Nguyen, Martin J Wainwright, and Michael Jordan. Estimating divergence functionals and the likelihood ratio by penalized convex risk minimization. In *Proc. Adv. Neural Inf. Process. Syst.*, volume 20, pp. 1–8, 2007.

721

722

723 XuanLong Nguyen, Martin J Wainwright, and Michael I Jordan. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Trans. Inf. Theory.*, 56(11):5847–5861, 2010.

724

725

726 Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f -GAN: Training generative neural samplers using variational divergence minimization. In *Proc. Adv. Neural Inf. Process. Syst.*, volume 29, pp. 1–9, 2016.

727

728

729

730 Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.*, 22(10):1345–1359, 2010. doi: 10.1109/TKDE.2009.191.

731

732

733 George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *J. Mach. Learn. Res.*, 22(57):1–64, 2021.

734

735

736 Min Sue Park, Cheolhyeong Kim, Hwijae Son, and Hyung Ju Hwang. The deep minimizing movement scheme. *J. Comput. Phys.*, 494:112518, 2023.

737

738

739 Michaël Perrot, Nicolas Courty, Rémi Flamary, and Amaury Habrard. Mapping estimation for discrete optimal transport. In *Proc. Adv. Neural Inf. Process. Syst.*, volume 29, pp. 1–9, 2016.

740

741

742 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.

743

744 Viraj Prabhu, Shivam Khare, Deeksha Kartik, and Judy Hoffman. SENTRY: Selective entropy optimization via committee consistency for unsupervised domain adaptation. In *Proc. IEEE Int. Conf. Comput. Vis.*, pp. 8558–8567, 2021.

745

746

747 Litu Rout, Alexander Korotin, and Evgeny Burnaev. Generative modeling with optimal transport maps. In *Proc. Int. Conf. Learn. Represent.*, pp. 1–22, 2022.

748

749

750 Shogo Sagawa and Hideitsu Hino. Gradual domain adaptation via normalizing flows. *Neural Comput.*, 37(3):522–568, 2025.

751

752 Filippo Santambrogio. {Euclidean, metric, and Wasserstein} gradient flows: an overview. *Bulletin of Mathematical Sciences*, 7:87–154, 2017.

753

754

755 Igal Sason and Sergio Verdú. f -divergence inequalities. *IEEE Trans. Inf. Theory.*, 62(11):5973–6006, 2016.

756 Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algo-*
 757 *rithms*. Cambridge university press, London, 2014.
 758

759 Jiaxin Shi, Chang Liu, and Lester Mackey. Sampling with mirrored Stein operators. *Proc. Int. Conf.*
 760 *Learn. Represent.*, pp. 1–26, 2022.

761 Stephen Simons. Minimax theorems and their proofs. In *Minimax and applications*, pp. 1–23.
 762 Springer, Berlin, Germany, 1995.
 763

764 Maurice Sion. On general minimax theorems. *Pacific J. Math.*, 8(4):171–176, 1958.

765 Marta Skreta, Tara Akhound-Sadegh, Viktor Ohanesian, Roberto Bondesan, Alan Aspuru-Guzik,
 766 Arnaud Doucet, Rob Brekelmans, Alexander Tong, and Kirill Neklyudov. Feynman-Kac correc-
 767 tors in diffusion: Annealing, guidance, and product of experts. In *Proc. Int. Conf. Mach. Learn.*,
 768 pp. 1–44. PMLR, 2025.

769

770 Yang Song, Sahaj Garg, Jiaxin Shi, and Stefano Ermon. Sliced score matching: A scalable approach
 771 to density and score estimation. In *Proc. Conf. Uncertainty in Artificial Intelligence*, pp. 574–584.
 772 PMLR, 2020.

773 Masashi Sugiyama and Motoaki Kawanabe. *Machine learning in non-stationary environments:*
 774 *Introduction to covariate shift adaptation*. MIT press, New York, 2012.

775

776 Masashi Sugiyama, Matthias Krauledat, and Klaus-Robert Müller. Covariate shift adaptation by
 777 importance weighted cross validation. *J. Mach. Learn. Res.*, 8(5):1–21, 2007.

778 George Brinton Thomas, Maurice D Weir, Joel Hass, Christopher Heil, and Antonio Behn. *Thomas'*
 779 *calculus: early transcendentals*, volume 15. Pearson Boston, London, U.K., 2014.

780

781 Hugo Touchette. Legendre-fenchel transforms in a nutshell. *URL <http://www.maths.qmul.ac.uk/ht/archive/lfth2.pdf>*, pp. 25, 2005.

782

783 Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain
 784 adaptation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, pp. 2962–2971, 2017. doi:
 785 10.1109/CVPR.2017.316.

786

787 Vladimir N Vapnik. An overview of statistical learning theory. *IEEE Trans. Neural. Netw.*, 10(5):
 788 988–999, 1999.

789

790 Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep
 791 hashing network for unsupervised domain adaptation. In *Proc. IEEE Conf. Comput. Vis. Pattern*
 792 *Recognit.*, pp. 5018–5027, 2017.

793

794 Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.

795

796 Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Comput.*,
 797 23(7):1661–1674, 2011.

798

799 Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol.
 800 Stacked denoising autoencoders: Learning useful representations in a deep network with a local
 801 denoising criterion. *J. Mach. Learn. Res.*, 11:3371–3408, 2010. ISSN 1532-4435.

802

803 Fangyikang Wang, Huminhao Zhu, Chao Zhang, Hanbin Zhao, and Hui Qian. GAD-PVI: A general
 804 accelerated dynamic-weight particle-based variational inference framework. In *Proc. AAAI Conf.*
 805 *Artif. Intell.*, volume 37, pp. 1–29, 2023.

806

807 Hao Wang, Hao He, and Dina Katabi. Continuously indexed domain adaptation. In *Proc. Int. Conf.*
 808 *Mach. Learn.*, pp. 9898–9907, 2020.

809

810 Haoxiang Wang, Bo Li, and Han Zhao. Understanding gradual domain adaptation: Improved anal-
 811 ysis, optimal path and beyond. In *Proc. Int. Conf. Mach. Learn.*, pp. 22784–22801. PMLR, 2022.

812

813 Max Welling and Yee W Teh. Bayesian learning via stochastic gradient Langevin dynamics. In
 814 *Proc. Int. Conf. Mach. Learn.*, pp. 681–688, 2011.

810 Shaoan Xie, Zibin Zheng, Liang Chen, and Chuan Chen. Learning semantic representations for
 811 unsupervised domain adaptation. In *Proc. Int. Conf. Mach. Learn.*, pp. 1–10. PMLR, 2018.

812

813 Haoyun Yin, Yixuan Qiu, and Xiao Wang. Wasserstein coresets via sinkhorn loss. *Transactions on*
 814 *Machine Learning Research*, pp. 1–40, 2025.

815 Zhichen Zeng, Ruizhong Qiu, Wenxuan Bao, Tianxin Wei, Xiao Lin, Yuchen Yan, Tarek F Abdelza-
 816 her, Jiawei Han, and Hanghang Tong. Pave your own path: Graph gradual domain adaptation on
 817 fused gromov-wasserstein geodesics. *arXiv preprint arXiv:2505.12709*, 2025.

818

819 Chao Zhang, Zhijian Li, Xin Du, and Hui Qian. DPVI: A dynamic-weight particle-based variational
 820 inference framework. In Lud De Raedt (ed.), *Proc. Int. Joint Conf. Artif. Intell.*, pp. 4900–4906.
 821 International Joint Conferences on Artificial Intelligence Organization, 7 2022. Main Track.

822 Xiaolin Zheng, Weiming Liu, Chaochao Chen, Jiajie Su, Xinting Liao, Mengling Hu, and Yanchao
 823 Tan. Mining user consistent and robust preference for unified cross domain recommendation.
 824 *IEEE Trans. Knowl. Data Eng.*, 36(12):8758–8772, 2024. doi: 10.1109/TKDE.2024.3446581.

825

826 Jia-Jie Zhu. Inclusive KL minimization: A Wasserstein-Fisher-Rao gradient flow perspective. In
 827 *Proc. Int. Conf. Learn. Represent. Frontiers in Probabilistic Inference: Learning meets Sampling*
 828 (*Workshop*), 2025.

829 Zhan Zhuang, Yu Zhang, and Ying Wei. Gradual domain adaptation via gradient flow. In *Proc. Int.*
 830 *Conf. Learn. Represent.*, pp. 1–26, 2024.

831

832

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

Contents for Appendix

864	A Nomenclature	18
865		
866		
867	B Mathematical Background on Optimal Transport	19
868		
869		
870	C Theoretical Derivation	20
871	C.1 Derivation of Eq. (4)	20
872	C.2 Derivation of Proposition 1	22
873	C.3 Derivation of Proposition 2	23
874	C.4 Derivation of Proposition 3	25
875	C.5 Derivation of Proposition 4	27
876	C.6 Derivation of Theorem 5	27
877	C.7 Derivation of Theorem 6	29
878	C.8 Discussions on the selection of step size η	32
879	D Detailed Algorithm of the E-SUOT Framework	35
880		
881	E Detailed Information for Experiments	35
882	E.1 Dataset Descriptions	35
883	E.2 Experimental Settings	36
884	E.3 Detailed Information for Ablation Studies	37
885	F Additional Experimental Results	39
886	F.1 Impact of Early Transport Steps on Adaptation Performance	39
887	F.2 Performance Improvement Varying Different UDA Feature Extractors	39
888	F.3 Investigation on the UMAP’s Dimension	40
889		
890	G Limitations & Future Directions and Broader Impact	41
891	G.1 Limitations & Future Directions	41
892	G.2 Broader Impact Statement	41
893		
894	H LLM Usage Statement	42
895		
896		
897		
898		
899		
900		
901		
902		
903		
904		
905		
906		
907		
908		
909		
910		
911		
912		
913		
914		
915		
916		
917		

918 **A NOMENCLATURE**
919
920
921922 **Table A.1: Technical terminology table.**

923 Symbol	924 Description	925 Symbol	926 Description
T	Intermediate domain number	\mathcal{A}	Upper bound on the L_2 norm of the first variation of the KL divergence
Δ	Performance difference	\mathcal{B}	Upper bound on the L_2 norm of the gradient of the first variation of the KL divergence
$\ \cdot\ _2$	L_2 norm	\mathcal{H}_0	Light-tail constant
$\arg \min$	Argument of the minimum	∇	Gradient operator
T	Transportation map	ω	Parameter of classifier
T^{-1}	Inverse transformation of the transportation map	ϕ	Parameter of potential function
δ	Variation operator	π	Transportation plan
\det	Determinant	$\rho^*(x)$	Optimal PDF
ϵ	Entropy regularization coefficient	\sup	Supremum
η	Discretization stepsize	θ	Parameter of transportation map
\exp	Exponential function	δ	Dirac delta mass
\inf	Infimum	ε	Generalization error
ι	Lipschitz constant of the loss function	\hat{h}	Logit layer of classifier
$\kappa(x, y)$	Reference joint PDF	\hat{y}	Predicted label
λ_1	Coefficient of unbalanced optimal transport	ζ	Lipschitz constant bound for hypotheses in \mathcal{H}
λ_2	Coefficient of unbalanced optimal transport	c	Cost matrix
$\mathbb{D}_f[\rho(x), p_T(x)]$	f -divergence of distribution $q(x)$ with-respect-to distribution $p(x)$	$f^*(x)$	Convex conjugate function of $f(x)$
$\mathbb{D}_{\text{KL}}[\rho(x), p_T(x)]$	Kullback-Leibler divergence of distribution $q(x)$ with-respect-to distribution $p(x)$	h	Classifier
$\mathbb{E}_q(x)[f(x)]$	Expectation of function $f(x)$ with-respect-to distribution $q(x)$	t	Time index
\mathbb{I}	Indicator function	u	Kantorovich potential function
B	Batch size	v_t	Velocity field
C	Cumulative domain transportation and label continuity costs along the adaptation path	w	Kantorovich potential function
\mathcal{H}	Hypothesis class for classifier	x	Input
\mathcal{L}	Loss function	y	Label
$\mathcal{P}_2(\mathbb{R}^D)$	D-dimensional Wasserstein space	GDA	Gradual domain adaptation
$\mathcal{S}_{\text{stat}}$	Statistical error term.	JKO	Jordan-Kinderlehrer-Otto
\mathcal{T}	Set of transportation map	KL divergence	Kullback-Leibler divergence
\mathcal{W}_p	p -Wasserstein distance	OT	Optimal transport
\mathcal{X}	Input space	PDF	Probability density function
\mathcal{Y}	Label space	UDA	Unsupervised domain adaptation
N	Sample size	UOT	Unbalanced optimal transport
d	Differential operator		
softmax	softmax function		

961
962
963
964
965
966
967
968
969
970
971

972 B MATHEMATICAL BACKGROUND ON OPTIMAL TRANSPORT 973

974 We begin by reviewing the relevant background of optimal transport, based on references (Villani
975 et al., 2009; Peyré et al., 2019). Assume continuous variables with densities: source $\rho(x)$ supported
976 on \mathcal{X} , target $\xi(y)$ supported on \mathcal{Y} , and a cost $c(x, y) \geq 0$. We search for a joint probability density
977 function which is called transport plan $\pi(x, y) \geq 0$ such that:

$$979 \int \pi(x, y) dy = \rho(x), \quad (B.1a)$$

$$980 \int \pi(x, y) dx = \xi(y), \quad (B.1b)$$

983 and minimize expected cost:

$$985 \inf_{\pi \geq 0} \iint c(x, y) \pi(x, y) dy dx, \quad (B.2)$$

987 where $c(x, y)$ is the cost function, for example, squared Euclidean norm: $c(x, y) = \|x - y\|_2^2$.
988 Notably, when $c(x, y)$ is chosen as the squared Euclidean distance, the resulting optimal transport
989 cost corresponds to the squared Wasserstein-2 distance between the two PDFs.

990 Introducing potentials $u(x)$ and $w(y)$ as Lagrange multipliers for the marginal constraints, we get:

$$992 \sup_{u, w} \left[\int u(x) \rho(x) dx + \int w(y) \xi(y) dy \right] \quad \text{s.t.} \quad u(x) + w(y) \leq c(x, y) \quad \forall x, y. \quad (B.3)$$

995 Intuitively, u and w are “prices”; the constraint ensures the total price never exceeds the cost func-
996 tion. In addition, u and w are also called “(Kantorovich) potential” in optimal transport.

997 Based on this, we can eliminate one potential via the c -transform as follows:

$$999 w^c(x) := \inf_y c(x, y) - w(y). \quad (B.4)$$

1001 Based on this, we get the semi-dual formulation of optimal transport problem (Korotin et al., 2021;
1002 2023; Choi et al., 2023; 2024; 2025) which maximizes over one potential:

$$1004 \sup_w \int w^c(x) \rho(x) dx + \int w(y) \xi(y) dy. \quad (B.5)$$

1006 Notably, when total mass may differ or we allow creation/destruction of mass, we can relax marginal
1007 constraints using the f -divergence-based penalty terms (Chizat et al., 2018; Zhang et al., 2022).
1008 Specifically, we still want to optimize $\pi(x, y) \geq 0$, but we will penalize deviations of the induced
1009 marginals $\tilde{\rho}(x) := \int \pi(x, y) dy$ and $\tilde{\xi}(y) := \int \pi(x, y) dx$ from $\rho(x)$ and $\xi(y)$:

$$1011 \min_{\pi \geq 0} \iint c(x, y) \pi(x, y) dy dx + \lambda_1 \mathbb{D}_f(\tilde{\rho}(x), \rho(x)) + \lambda_2 \mathbb{D}_f(\tilde{\xi}(y), \xi(y)), \quad (B.6)$$

1014 where $\mathbb{D}_f(\tilde{\rho}(x), \rho(x)) = \int \rho(x) f\left(\frac{\tilde{\rho}(x)}{\rho(x)}\right) dx$ and $\lambda_{1,2} > 0$.

1016 In addition, using the convex conjugate f^* , the dual problem becomes

$$1018 \max_{u, w} - \int \rho(x) f_1^*(-u(x)) dx - \int \xi(y) f_2^*(-w(y)) dy \quad \text{s.t.} \quad u(x) + w(y) \leq c(x, y) \quad \forall x, y, \quad (B.7)$$

1020 where f_1, f_2 are the chosen divergences on each side.

1021 Similarly, we can eliminate one potential via the c -transform as follows:

$$1023 \max_w - \int \rho(x) f_1^*(-w^c(x)) dx - \int \xi(y) f_2^*(-w(y)) dy, \quad w^c(x) = \inf_y \{c(x, y) - w(y)\}. \quad (B.8)$$

1025 Based on this, we obtain the semi-dual formulation of the unbalanced optimal transport problem.

1026 **C THEORETICAL DERIVATION**
 1027

1028 **C.1 DERIVATION OF EQ. (4)**
 1029

1030 In this subsection, we want to derive the following equivalent relationship in the main content to
 1031 uphold the rigor of our manuscript:

1032
$$x_{t+\eta} = x_t - \eta \nabla \frac{\delta \mathbb{D}_f[p(x_t), p_T]}{\delta p(x_t)} \Rightarrow p(x_{t+\eta}) = \arg \min_{\rho(x) \in \mathcal{P}_2(\mathbb{R}^D)} \frac{1}{2\eta} \mathcal{W}_2^2(\rho(x), p(x_t)) + \mathbb{D}_f[\rho(x), p_T(x)]. \quad (C.1)$$

 1033
 1034
 1035

1036 Notably, the optimization problem given by the right-hand-side of the abovementioned equation is
 1037 also called Jordan-Kinderlehrer-Otto canonical form (Jordan et al., 1998; Caluya & Halder, 2020;
 1038 2022) or minimum movement scheme (Park et al., 2023). Before conducting the derivation, it is
 1039 necessary to introduce the definition of Wasserstein distance. The squared 2-Wasserstein distance
 1040 \mathcal{W}_2^2 can be defined by finding a transport map $\mathbf{T} : \mathbb{R}^D \rightarrow \mathbb{R}^D$ that minimizes the average cost of
 1041 transporting mass from $\rho(x)$ to $\xi(x)$ as follows:

1042
$$\mathcal{W}_2^2(\rho, \xi) = \inf_{\mathbf{T} : \mathbf{T}_\# \rho(x) = \xi(x)} \int \|x - \mathbf{T}(x)\|_2^2 \rho(x) dx, \quad (C.2)$$

 1043

1044 where $\mathbf{T}_\#$ indicates the pushforward measure, and the expression for $\mathbf{T}(x)$ is defined as follows:

1045
$$\mathbf{T}(x) = x + \eta v_t(x). \quad (C.3)$$

 1046

1047 Meanwhile, during the transportation, the differential equation that delineates PDF of the evolution
 1048 process driven by Eq. (1) is called *continuity equation*, defined as follows:

1049
$$\frac{\partial \rho(x_t)}{\partial t} = -\nabla \cdot [v_t(x_t) \rho(x_t)]. \quad (C.4)$$

 1050
 1051

1052 Building on Eqs. (C.3) and (C.4), and discretizing the continuity equation in the time domain using
 1053 the forward Euler scheme (Butcher, 2016; Evans, 2022), we obtain:

1054
$$\rho(x) = \rho(x_t) - \eta \nabla \cdot (\rho(x_t) v_t(x_t)) + \mathcal{O}(\eta^2). \quad (C.5)$$

 1055

1056 Taking the functional derivative of $\mathbb{D}_f[\rho(x), p_T(x)]$ with respect to $\rho(x)$, we get:

1057
$$\begin{aligned} \frac{d}{d\eta} \mathbb{D}_f[\rho(x), p_T(x)] &= \frac{d}{d\eta} \int p_T(x) f\left(\frac{\rho(x)}{p_T(x)}\right) dx \\ &= \int p_T(x) \frac{d}{d\eta} f\left(\frac{\rho(x)}{p_T(x)}\right) dx \\ &= \int p_T(x) f'\left(\frac{\rho(x)}{p_T(x)}\right) \frac{1}{p_T(x)} \frac{\partial \rho(x)}{\partial \eta} dx \\ &\stackrel{(i)}{=} \int \frac{\delta \mathbb{D}_f}{\delta \rho(x)} \frac{\partial \rho(x)}{\partial \eta} dx \end{aligned} \quad (C.6)$$

 1058
 1059
 1060
 1061
 1062
 1063
 1064
 1065
 1066
 1067
 1068

1069 Here, step (i) is based on comparing the first variation:

1070
$$\begin{aligned} \delta \mathbb{D}_f[\rho; \sigma] &= \frac{d}{d\varepsilon} \bigg|_{\varepsilon=0} \int p_T(x) f\left(\frac{\rho(x) + \varepsilon \sigma(x)}{p_T(x)}\right) dx \\ &= \int p_T(x) f'\left(\frac{\rho(x)}{p_T(x)}\right) \frac{1}{p_T(x)} \sigma(x) dx \quad (\text{chain rule, } p_T \text{ fixed}) \\ &= \int f'\left(\frac{\rho(x)}{p_T(x)}\right) \sigma(x) dx, \end{aligned}$$

 1071
 1072
 1073
 1074
 1075
 1076
 1077

1078 with the definition of functional derivative:

1079
$$\delta \mathbb{D}_f[\rho; \sigma] = \int \frac{\delta \mathbb{D}_f}{\delta \rho(x)} \sigma(x) dx,$$

1080 where $\sigma(x)$ denotes an arbitrary perturbation function. Inserting Eq. (C.5) into Eq. (C.6), we get
1081

$$\begin{aligned}
1082 \frac{d}{d\eta} \mathbb{D}_f[\rho(x), p_T(x)] &= \int \frac{\delta \mathbb{D}_f}{\delta \rho(x)} [-\nabla \cdot (\rho(x)v_t(x))] dx \\
1083 &= \int \frac{\delta \mathbb{D}_f}{\delta \rho(x)} [-v_t^\top(x) \nabla \rho(x) - \rho(x) \nabla \cdot v_t(x)] dx \\
1084 &\stackrel{(ii)}{=} \int \left(-\nabla \cdot \left[\frac{\delta \mathbb{D}_f}{\delta \rho(x)} \rho(x)v_t(x) \right] + \rho(x) v_t^\top(x) \nabla \frac{\delta \mathbb{D}_f}{\delta \rho(x)} \right) dx \\
1085 &\stackrel{(iii)}{=} \int \rho(x) v_t^\top(x) \nabla \frac{\delta \mathbb{D}_f[\rho(x), p_T(x)]}{\delta \rho(x)} dx.
\end{aligned} \tag{C.7}$$

1091 Step (ii) is based on the chain rule:
1092

$$\begin{aligned}
1093 \nabla \cdot \left[\frac{\delta \mathbb{D}_f}{\delta \rho(x)} \rho(x)v_t(x) \right] &= \frac{\delta \mathbb{D}_f}{\delta \rho(x)} \rho(x) [\nabla \cdot v_t(x)] \\
1094 &\quad + \frac{\delta \mathbb{D}_f}{\delta \rho(x)} v_t^\top(x) \nabla \rho(x) \\
1095 &\quad + [\nabla \frac{\delta \mathbb{D}_f}{\delta \rho(x)}]^\top [\rho(x)v_t(x)].
\end{aligned} \tag{C.8}$$

1096 Step (iii) uses a mild regularity assumption (Abraham et al., 2012; Liu et al., 2019; Shi et al., 2022)
1097 on $\frac{\delta \mathbb{D}_f}{\delta \rho(x)} \rho(x)v_t(x)$, for example rapid decay as $x \rightarrow \infty$, so that
1098

$$\int -\nabla \cdot \left[\frac{\delta \mathbb{D}_f}{\delta \rho(x)} \rho(x)v_t(x) \right] dx = 0. \tag{C.9}$$

1100 Consequently, $\mathbb{D}_f[\rho(x), p_T(x)]$ can be expanded as follows when $\eta \rightarrow 0$:
1101

$$\mathbb{D}_f[\rho(x), p_T(x)] = \mathbb{D}_f[p(x_t), p_T(x)] + \eta \int p(x_t) v_t^\top(x_t) \nabla \frac{\delta \mathbb{D}_f[p(x_t), p_T(x)]}{\delta p(x_t)} dx. \tag{C.10}$$

1102 For the squared 2-Wasserstein distance, we get:
1103

$$\mathcal{W}_2^2(\rho(x), p(x_t)) = \int p(x_t) \|x - \mathbf{T}^*(x_t)\|_2^2 dx = \eta^2 \int p(x_t) \|v_t^*(x_t)\|_2^2 dx \leq \eta^2 \int p(x_t) \|v_t(x_t)\|_2^2 dx, \tag{C.11}$$

1104 where $\mathbf{T}^*(x)$ and $v_t^*(x)$ are the optimal transportation map and optimal velocity field. Since $v_t(x)$
1105 is not the optimal velocity field, we obtain the last inequality. Based on Eqs. (C.10) and (C.11), we
1106 finally reach the following result:
1107

$$\begin{aligned}
1108 \mathbb{D}_f[\rho(x), p_T(x)] &+ \frac{1}{2\eta} \mathcal{W}_2^2(\rho(x), p(x_t)) - \mathbb{D}_f[p(x_t), p_T(x)] \\
1109 &\leq \mathbb{D}_f[\rho(x), p_T(x)] + \frac{\eta}{2} \mathbb{E}_{p(x_t)} [\|v_t(x_t)\|_2^2] + \eta \int \cdot [p(x_t) v_t^\top(x_t) \nabla \frac{\delta \mathbb{D}_f[p(x_t), p_T(x)]}{\delta p(x_t)}] dx - \mathbb{D}_f[\rho(x), p_T(x)] \\
1110 &\leq \underbrace{\frac{\eta}{2} \mathbb{E}_{p(x_t)} [\|\nabla \frac{\delta \mathbb{D}_f[p(x_t), p_T(x)]}{\delta p(x_t)}\|_2^2]}_{\geq 0} + \frac{\eta}{2} \mathbb{E}_{p(x_t)} [\|v_t(x_t)\|_2^2] + \eta \int p(x_t) v_t^\top(x_t) \nabla \frac{\delta \mathbb{D}_f[p(x_t), p_T(x)]}{\delta p(x_t)} dx \\
1111 &= \frac{\eta}{2} \mathbb{E}_{p(x_t)} \{ \|v_t(x_t) + \nabla \frac{\delta \mathbb{D}_f[p(x_t), p_T(x)]}{\delta p(x_t)}\|_2^2 \}.
\end{aligned} \tag{C.12}$$

1112 Consequently, the optimal velocity field that reduces the upper bound of the optimization problem
1113 defined by the right-hand-side of Eq. (4) can be given as follows:
1114

$$v_t^*(x_t) = -\nabla \frac{\delta \mathbb{D}_f[p(x_t), p_T(x)]}{\delta p(x_t)}, \tag{C.13}$$

1115 which implies that the left-hand side of Eq. (C.1) is a sufficient condition for the optimality of its
1116 right-hand side.
1117

1134 C.2 DERIVATION OF PROPOSITION 1
11351136 **Proposition (1).** Consider the following primal problem:

1137
1138
$$\mathcal{L}^{Primal} = \arg \min_{\rho(x) \in \mathcal{P}_2(\mathbb{R}^D)} \frac{1}{2\eta} \mathcal{W}_2^2(\rho(x), p(x_t)) + \mathbb{D}_f[\rho(x), p_T(x)]. \quad (\text{C.14})$$

1139

1140 This problem is equivalent to the following semi-dual formulation:

1141
1142
$$\mathcal{L}^{SemiDual} = \sup_w \mathbb{E}_{p(x_t)} \left[\inf_{\mathbf{T}} (\|\mathbf{T}(x_t) - x_t\|_2^2 - w(\mathbf{T}(x_t))) \right] - \mathbb{E}_{p_T(x)} [f^*(-w(x))], \quad (\text{C.15})$$

1143

1144 where $w : \mathbb{R}^D \rightarrow \mathbb{R}$ is a measurable continuous function, $\mathbf{T} : \mathbb{R}^D \rightarrow \mathbb{R}^D$ is the transport map, and
1145 f^* denotes the convex conjugate of f , defined as $f^*(z) := \sup_{y \geq 0} (zy - f(y))$.1146 *Proof.* Eq. (C.14) can be reformulated as follows:

1147
1148
$$\inf_{\pi \in \mathbb{R}_+^{D \times D}} \frac{1}{2\eta} \iint \|x_t - x\|_2^2 \pi(x_t, x) dx_t dx + \int f\left(\frac{\rho(x)}{p_T(x)}\right) p_T(x) dx, \quad (\text{C.16a})$$

1149

1150
1151 s.t. $p(x_t) = \int \pi(x_t, x) dx, \quad \rho(x) = \int \pi(x_t, x) dx_t. \quad (\text{C.16b})$
1152

1153 Based on this, we introduce the Lagrangian multiplier [Biegler \(2010\)](#) $u(x_t)$ and $w(x)$ to handle the
1154 equality constraints given by Eq. (C.16b) as follows:

1155
1156
$$\begin{aligned} \mathcal{L} &= \frac{1}{2\eta} \iint \|x_t - x\|_2^2 \pi(x_t, x) dx_t dx + \int f\left(\frac{\rho(x)}{p_T(x)}\right) p_T(x) dx \\ &\quad + \int u(x_t) [p(x_t) - \int \pi(x_t, x) dx] dx_t + \int w(x) [\rho(x) - \int \pi(x_t, x) dx_t] dx \\ &= \iint \left[\frac{1}{2\eta} \|x_t - x\|_2^2 - u(x_t) - w(x) \right] \pi(x_t, x) dx_t dx \\ &\quad + \int u(x_t) p(x_t) dx_t + \int w(x) \rho(x) + f\left(\frac{\rho(x)}{p_T(x)}\right) p_T(x) dx. \end{aligned} \quad (\text{C.17})$$

1157
1158
1159
1160
1161
1162
1163
1164

1165 On this basis, the dual function can be given as follows due to the linear independent structure of
1166 problem defined by Eq. (C.17):

1167
1168
$$\begin{aligned} g(u, w) &= \inf_{\pi(x_t, x)} \iint \left[\frac{1}{2\eta} \|x_t - x\|_2^2 - u(x_t) - w(x) \right] \pi(x_t, x) dx_t dx \\ &\quad + \int u(x_t) p(x_t) dx_t + \inf_{\rho(x)} \int \left[w(x) \frac{\rho(x)}{p_T(x)} + f\left(\frac{\rho(x)}{p_T(x)}\right) \right] p_T(x) dx \\ &= \inf_{\pi(x_t, x)} \iint \left[\frac{1}{2\eta} \|x_t - x\|_2^2 - u(x_t) - w(x) \right] \pi(x_t, x) dx_t dx \\ &\quad + \int u(x_t) p(x_t) dx_t - \int p_T(x) f^*(-w(x)) dx. \end{aligned} \quad (\text{C.18})$$

1169
1170
1171
1172
1173
1174
1175
1176
1177
1178

1179 where the last line uses the Legendre–Fenchel conjugate ([Touchette, 2005](#); [Caluya & Halder, 2020](#)).
1180 Writing $y(x) = \rho(x)/p_T(x)$ and using separability, we have

1181
1182
$$\begin{aligned} \inf_{\rho(x)} \int \left[w(x) \frac{\rho(x)}{p_T(x)} + f\left(\frac{\rho(x)}{p_T(x)}\right) \right] p_T(x) dx &= \int \inf_{y \geq 0} (w(x) y + f(y)) p_T(x) dx \\ &= - \int \sup_{y \geq 0} ((-w(x)) y - f(y)) p_T(x) dx \\ &= - \int p_T(x) f^*(-w(x)) dx. \end{aligned} \quad (\text{C.19})$$

1183
1184
1185
1186
1187

Suppose that $\frac{1}{2\eta}\|x_t - x\|_2^2 - u(x_t) - w(x) < 0$ for some pair (x_t, x) . In this case, concentrating all the mass of $\pi(x_t, x)$ at this point drives the Lagrangian in Eq. (C.18) to $-\infty$. To avoid such degenerate solutions, it is necessary to impose the condition $\frac{1}{2\eta}\|x_t - x\|_2^2 - u(x_t) - w(x) \geq 0$ almost everywhere. Consequently, the dual problem can be written as

$$\sup_{u(x_t) + w(x) \leq \frac{1}{2\eta}\|x_t - x\|_2^2} \left\{ \int u(x_t)p(x_t) dx_t - \int p_T(x)f^*(-w(x)) dx \right\}. \quad (\text{C.20})$$

Equivalently, introducing the convex indicator function ℓ , this becomes

$$\sup_{u,w} \left\{ \int u(x_t)p(x_t) dx_t - \int p_T(x)f^*(-w(x)) dx - \ell\left(u(x_t) + w(x) \leq \frac{1}{2\eta}\|x_t - x\|_2^2\right) \right\}. \quad (\text{C.21})$$

Since f^* is convex, non-decreasing, and differentiable, and because $\|x_t - x\|_2^2 \geq 0$, the choice $u(x_t) \equiv -1$ and $w(x) \equiv -1$ ensures all terms in Eq. (C.21) are finite. By Fenchel–Rockafellar’s theorem (Bauschke & Combettes, 2017), strong duality therefore holds. Moreover, by complementary slackness the optimal plan π^* assigns zero mass to pairs where $\frac{1}{2\eta}\|x_t - x\|_2^2 - u^*(x_t) - w^*(x) > 0$, implying that $\frac{1}{2\eta}\|x_t - x\|_2^2 = u^*(x_t) + w^*(x)$ π^* -almost everywhere. Hence,

$$u^*(x_t) = \inf_x \left(\frac{1}{2\eta}\|x_t - x\|_2^2 - w^*(x) \right).$$

Substituting this into the dual yields the semi-dual formulation

$$\sup_{w(x)} \left\{ \int \inf_x \left[\frac{1}{2\eta}\|x_t - x\|_2^2 - w(x) \right] p(x_t) dx_t - \int p_T(x)f^*(-w(x)) dx \right\}. \quad (\text{C.22})$$

Defining the transport map via the c -transform as

$$\begin{aligned} \mathbf{T}^*(x_t) &\in \arg \min_x \left(\frac{1}{2\eta}\|x_t - x\|_2^2 - w(x) \right) \\ \iff \inf_x \left(\frac{1}{2\eta}\|x_t - x\|_2^2 - w(x) \right) &= \frac{1}{2\eta}\|x_t - \mathbf{T}^*(x_t)\|_2^2 - w(\mathbf{T}^*(x_t)), \end{aligned} \quad (\text{C.23})$$

and substituting Eq. (C.23) into Eq. (C.22), we obtain the final semi-dual objective

$$\mathcal{L}^{\text{SemiDual}} = \sup_w \mathbb{E}_{p(x_t)} \left[\left\| \frac{1}{2\eta} \mathbf{T}^*(x_t) - x_t \right\|_2^2 - w(\mathbf{T}^*(x_t)) \right] - \mathbb{E}_{p_T(x)} [f^*(-w(x))], \quad (\text{C.24})$$

It should be pointed out that there is no closed-form expression of the optimal $\mathbf{T}^*(x_t)$ for each $w(x)$ (Korotin et al., 2023; Choi et al., 2023). Hence, the optimization $\mathbf{T}(x_t)$ for each $w(x)$ is required, and we reach the final semi-dual objective as follows based on Eq. (C.24):

$$\mathcal{L}^{\text{SemiDual}} = \sup_w \mathbb{E}_{p(x_t)} \left[\inf_{\mathbf{T}} \left(\frac{1}{2\eta} \|\mathbf{T}(x_t) - x_t\|_2^2 - w(\mathbf{T}(x_t)) \right) \right] - \mathbb{E}_{p_T(x)} [f^*(-w(x))].$$

□

C.3 DERIVATION OF PROPOSITION 2

Proposition (2). *The semi-dual formulation in Eq. (7) admits non-unique optimal solutions.*

Proof. Consider the discrete optimal transport setting with a single source point (x_t in Eq. (7)) and two symmetric target points (x in Eq. (7)). Augment the dual objective with an f -divergence term acting only on the target potential w , but not on the source potential u . Then the dual optimizer is not unique.

Specifically, let:

- **Source space:** $x_t = \{a\}$ with $p(x_t) \approx \delta_a$.

1242 • **Target space:** $x = \{b_1, b_2\}$ with $\rho(x) = \frac{1}{2}\delta_{b_1} + \frac{1}{2}\delta_{b_2}$.
 1243
 1244 • **Cost constant on pairs:** $\|a - b_1\|_2^2 = \|a - b_2\|_2^2 = K$ for some fixed $K \in \mathbb{R}$.
 1245

1246 The dual problem obtained from the primal with an additional term $\int f\left(\frac{\rho(x)}{p_T(x)}\right) p_T(x) dx$ acting
 1247 only on the target side admits multiple optimal solutions (u, w) ; in particular, uniqueness fails.
 1248

1249 The demonstration process can be summarized as follows:
 1250

1251 1) At the beginning, let us recall the feasibility for the multipliers u and w :

$$1252 \quad u(a) + w(b_j) \leq \|a - b_j\|_2^2 = K, \quad \forall j \in \{1, 2\}. \quad (\text{C.25})$$

1254 Based on this, we can define a shifted source potential $\tilde{u} := u - K$ and keep $\tilde{w} := w$.
 1255 Hence, the feasibility in Eq. (C.25) can be given as follows:

$$1256 \quad \tilde{u}(a) + \tilde{w}(b_j) \leq 0, \quad \forall j \in \{1, 2\}, \quad (\text{C.26})$$

1258 where the dual objective differs from the original by a global additive constant (independent
 1259 of (\tilde{u}, \tilde{w})), hence the set of maximizers is unaffected by this normalization. As such,
 1260 without loss of generality, it suffices to analyze the case $K = 0$. For notational simplicity
 1261 we drop tildes and write

$$1262 \quad u + w_j \leq 0, \quad \forall j \in \{1, 2\}. \quad (\text{C.27})$$

1263 2) Eliminating u and obtaining a piecewise-linear term Since $p(a) = 1$ and $\rho(b_1) = \rho(b_2) =$
 1264 $\frac{1}{2}$, the dual objective function (up to an additive constant) can be reformulated as follows:

$$1266 \quad \max_{u, w_1, w_2} \quad u + \frac{1}{2}w_1 + \frac{1}{2}w_2 - \frac{1}{2}f^*(-w_1) - \frac{1}{2}f^*(-w_2), \quad (\text{C.28})$$

1268 subject to $u \leq -w_1$ and $u \leq -w_2$. At optimum the constraint in u is tight, hence we have
 1269 the following result:

$$1270 \quad u = -\min\{w_1, w_2\}. \quad (\text{C.29})$$

1271 Substituting back yields an equivalent maximization over (w_1, w_2) :

$$1272 \quad \Phi(w_1, w_2) := -\min\{w_1, w_2\} + \frac{1}{2}w_1 + \frac{1}{2}w_2 - \frac{1}{2}f^*(-w_1) - \frac{1}{2}f^*(-w_2). \quad (\text{C.30})$$

1274 On this basis, we can define the “hinge” (V-shaped) linear part as follows:

$$1275 \quad L(w_1, w_2) := -\min\{w_1, w_2\} + \frac{1}{2}w_1 + \frac{1}{2}w_2 = \begin{cases} \frac{1}{2}(w_2 - w_1), & w_1 \leq w_2, \\ \frac{1}{2}(w_1 - w_2), & w_2 \leq w_1, \end{cases} \quad (\text{C.31})$$

1278 so that $L(w_1, w_2) = \frac{1}{2}|w_1 - w_2|$ and in particular $L(r, r) = 0$ for all r .

1280 Consequently, we have:

$$1281 \quad \Phi(w_1, w_2) = \frac{1}{2}|w_1 - w_2| - \frac{1}{2}f^*(-w_1) - \frac{1}{2}f^*(-w_2). \quad (\text{C.32})$$

1283 3) Notably, on the diagonal $w_1 = w_2 = r$, we have the following result:

$$1285 \quad \Phi(r, r) = -f^*(-r). \quad (\text{C.33})$$

1286 Since f^* is strictly convex, the one-dimensional problem $\max_r \Phi(r, r)$ has a unique maxi-
 1287 mizer r^* . Now let us consider antisymmetric perturbations around the diagonal:

$$1289 \quad w_1 = r^* + \delta, \quad w_2 = r^* - \delta, \quad \delta \in \mathbb{R}. \quad (\text{C.34})$$

1290 Then we obtain the following result:

$$1292 \quad \frac{1}{2}|w_1 - w_2| = \frac{1}{2}|2\delta| = |\delta|. \quad (\text{C.35})$$

1293 Using the second-order Taylor expansion of the strictly convex function f^* about $-r^*$, we
 1294 have for the following equality for small $|\delta|$:

$$1295 \quad -\frac{1}{2}f^*(-w_1) - \frac{1}{2}f^*(-w_2) = -f^*(-r^*) - \frac{1}{2}(f^{**}(-r^*))\delta^2 + \mathcal{O}(\delta^2). \quad (\text{C.36})$$

1296

Based on this, we get:

1297

$$\Phi(r^* + \delta, r^* - \delta) = |\delta| - \frac{1}{2} f^{**}(-r^*) \delta^2 - f^*(-r^*) + \mathcal{O}(\delta^2). \quad (\text{C.37})$$

1298

1299

1300 It should be pointed out that, for any sufficiently small but nonzero δ , the linear gain term
 1301 $|\delta|$ dominates the quadratic penalty term $\frac{1}{2} f^{**}(-r^*) \delta^2$, hence

1302

1303

$$\Phi(r^* + \delta, r^* - \delta) > \Phi(r^*, r^*).$$

1304

1305

1306

Consequently, the diagonal point $(w_1, w_2) = (r^*, r^*)$ is not uniquely optimal; in fact,
 there exists a continuum of distinct maximizers in a neighborhood along the antisymmetric
 direction. The corresponding u is

1307

1308

1309

$$u = -\min\{w_1, w_2\} = \begin{cases} -(r^* - \delta), & \delta \geq 0, \\ -(r^* + \delta), & \delta < 0, \end{cases}$$

1310

1311

yielding distinct optimal triples (u, w_1, w_2) for different $\delta \neq 0$.

1312

1313

1314

4) If the original cost is $\|a - b_j\|_2^2 = K$, recall $u = \tilde{u} + K$. Thus each optimal (\tilde{u}, w) constructed above gives an optimal (u, w) for the original problem by adding K to u . As the set of optimal w -pairs is already non-singleton, the full optimal dual variable pair (u, w) is non-unique.

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

In summary, our proof is based on the counter-example mentioned above. Specifically, in the symmetric two-target discrete setting, with the additional f -term acting only on the target potential w , the dual objective contains a V-shaped hinge $L(w_1, w_2) = \frac{1}{2}|w_1 - w_2|$ arising from eliminating u . This non-strict component competes with the strictly convex penalty $-\sum_j \rho(b_j) f^*(-w(b_j))$. Along antisymmetric perturbations, the first-order increase from the hinge dominates the second-order decrease from the convex penalty, producing a continuum of maximizers. Hence the optimal dual variable pair is not unique. Consequently, the dual problem defined in Eq. (7) admits non-unique optimal solutions. \square

C.4 DERIVATION OF PROPOSITION 3

1326

1327

1328

Proposition (3). Let $\kappa(x_t, x) := p(x_t) p_T(x)$ denote the reference joint PDF. The entropy-regularized primal problem is

1329

1330

1331

$$\begin{aligned} \mathcal{L}^{\text{E-Primal}} = \arg \min_{\rho \in \mathcal{P}_2(\mathbb{R}^D)} & \frac{1}{2\eta} \mathcal{W}_2^2(\rho(x), p(x_t)) + \mathbb{D}_f[\rho(x), p_T(x)] \\ & + \epsilon \iint \pi(x_t, x) [\log \frac{\pi(x_t, x)}{\kappa(x_t, x)} - 1] dx_t dx, \end{aligned} \quad (\text{C.38})$$

1332

1333

and is equivalent to the semi-dual optimization problem

1334

1335

1336

1337

1338

1339

$$\mathcal{L}^{\text{E-SemiDual}} = \sup_w -\epsilon \mathbb{E}_{p(x_t)} [\log \mathbb{E}_{p_T(x)} (\exp(\frac{w(x) - \frac{1}{2\eta} \|x - x_t\|_2^2}{\epsilon}))] - \mathbb{E}_{p_T(x)} [f^*(-w(x))], \quad (\text{C.39})$$

where f^* denotes the convex conjugate of f .

1340

1341

1342

1343

Proof. Define $c(x_t, x) := \frac{1}{2\eta} \|x_t - x\|_2^2$ as the quadratic transport cost. Introducing Lagrange multipliers $u(x_t) : \mathbb{R}^D \rightarrow \mathbb{R}$ (for the x_t -marginal) and $w(x) : \mathbb{R}^D \rightarrow \mathbb{R}$ (for the x -marginal). The Lagrangian of Eq. (C.38) is

1344

1345

1346

1347

1348

1349

$$\begin{aligned} \mathcal{L}(\pi, \rho; u, w) = & \iint c(x_t, x) \pi(x_t, x) dx_t dx + \epsilon \iint \pi(x_t, x) \left[\log \frac{\pi(x_t, x)}{\kappa(x_t, x)} - 1 \right] dx_t dx \\ & + \int f\left(\frac{\rho(x)}{p_T(x)}\right) p_T(x) dx + \int u(x_t) \left[p(x_t) - \int \pi(x_t, x) dx \right] dx_t \\ & + \int w(x) \left[\rho(x) - \int \pi(x_t, x) dx_t \right] dx. \end{aligned} \quad (\text{C.40})$$

1350 Grouping π -, ρ - and constant terms yields
 1351

$$\begin{aligned} 1352 \quad \mathcal{L} &= \iint (c(x_t, x) - u(x_t) - w(x)) \pi(x_t, x) dx_t dx \\ 1353 &+ \epsilon \iint \pi(x_t, x) \left[\log \frac{\pi(x_t, x)}{\kappa(x_t, x)} - 1 \right] dx_t dx \\ 1354 &+ \int \left(w(x) \rho(x) + f\left(\frac{\rho(x)}{p_T(x)}\right) p_T(x) \right) dx + \int u(x_t) p(x_t) dx_t. \end{aligned} \quad (C.41)$$

1359 Define $a(x_t, x) := c(x_t, x) - u(x_t) - w(x)$. For each fixed (x_t, x) , minimize
 1360

$$1361 \quad \phi(y) := a y + \epsilon \left(y \log \frac{y}{\kappa} - y \right), \quad y \geq 0. \\ 1362$$

1363 The first-order condition $a + \epsilon \log(y/\kappa) = 0$ gives
 1364

$$y^* = \kappa e^{-a/\epsilon} = \kappa e^{(u+w-c)/\epsilon}. \quad (C.42)$$

1366 Substituting back yields
 1367

$$\inf_{y \geq 0} \phi(y) = -\epsilon \kappa e^{-a/\epsilon} = -\epsilon \kappa \exp\left(\frac{u+w-c}{\epsilon}\right). \quad (C.43)$$

1369 Hence
 1370

$$1371 \quad \inf_{\pi \geq 0} \left\{ \pi\text{-terms of Eq. (C.41)} \right\} = -\epsilon \iint \kappa(x_t, x) \exp\left(\frac{u(x_t) + w(x) - c(x_t, x)}{\epsilon}\right) dx_t dx. \quad (C.44)$$

1373 For ρ , by Legendre–Fenchel conjugate (Touchette, 2005; Caluya & Halder, 2020), we have:
 1374

$$1375 \quad \inf_{\rho(x) \geq 0} \left\{ w(x) \rho(x) + f\left(\frac{\rho(x)}{p_T(x)}\right) p_T(x) \right\} = -p_T(x) f^*(-w(x)). \quad (C.45)$$

1377 Integrating over x gives
 1378

$$1379 \quad \inf_{\rho} \int [w \rho(x_t) + f\left(\frac{\rho(x_t)}{p_T(x)}\right) p_T(x)] dx = - \int p_T(x) f^*(-w(x)) dx.$$

1381 Combining Eq. (C.44) and Eq. (C.45), we obtain
 1382

$$\begin{aligned} 1383 \quad g(u, w) &= -\epsilon \iint \kappa(x_t, x) \exp\left(\frac{u(x_t) + w(x) - c(x_t, x)}{\epsilon}\right) dx_t dx \\ 1384 &- \int p_T(x) f^*(-w(x)) dx + \int u(x_t) p(x_t) dx_t. \end{aligned} \quad (C.46)$$

1388 Using $\kappa = p(x_t) \cdot p_T(x)$, define
 1389

$$1390 \quad A(x_t) := \int \exp\left(\frac{w(x) - c(x_t, x)}{\epsilon}\right) p_T(x) dx. \quad (C.47)$$

1392 Then
 1393

$$\iint \kappa \exp\left[\frac{(u(x_t) + w(x) - c(x_t, x))}{\epsilon}\right] dx_t dx = \int p(x_t) A(x_t) e^{\frac{u(x_t)}{\epsilon}} dx_t.$$

1395 Thus, Eq. (C.46) can be reformulated as follows:
 1396

$$1397 \quad g(u, w) = \int [p(x_t) u(x_t) - \epsilon p(x_t) A(x_t) e^{u(x_t)/\epsilon}] dx_t - \int p_T(x) f^*(-w(x)) dx. \quad (C.48)$$

1399 For each x_t , consider
 1400

$$1401 \quad \psi_{x_t}(u) := p(x_t) u - \epsilon p(x_t) A(x_t) e^{\frac{u(x_t)}{\epsilon}}.$$

1402 The first-order condition
 1403

$$\frac{d}{du} \psi_{x_t}(u) = p(x_t) - p(x_t) A(x_t) e^{\frac{u(x_t)}{\epsilon}} = 0$$

1404 gives

$$1405 \quad e^{u^*(x_t)/\epsilon} = \frac{1}{A(x_t)} \iff u^*(x_t) = -\epsilon \log A(x_t). \quad (C.49)$$

1406 Substituting back,

$$1407 \quad \sup_u \psi_{x_t}(u) = \epsilon p(x_t) (-\log A(x_t) - 1).$$

1408 Summing over x_t and discarding the constant $-\epsilon \int p(x_t) dx_t = -\epsilon$ (independent of $w(x)$), we
1409 obtain the semi-dual

$$1410 \quad \sup_w -\epsilon \mathbb{E}_{p(x_t)} [\log A(x_t)] - \mathbb{E}_{p_T(x)} [f^*(-w(x))], \quad (C.50)$$

1411 with $A(x_t)$ defined in Eq. (C.47). \square

1412 C.5 DERIVATION OF PROPOSITION 4

1413 **Proposition (4).** *The semi-dual formulation in Eq. (9) admits a unique optimal solution.*

1414 *Proof.* Let the entropy-regularized dual objective in Eq. (9) be

$$1415 \quad g(w) = -\epsilon \mathbb{E}_{p(x_t)} \{ \log \mathbb{E}_{p_T(x)} [\exp(\frac{w(x) - \|x - x_t\|_2^2}{\epsilon})] \} - \mathbb{E}_{p_T(x)} [f^*(-w(x))], \quad (C.51)$$

1416 where f^* is assumed to be strictly convex and proper, and $\epsilon > 0$.

1417 We seek to show that $g(w)$ is a strictly concave functional on an appropriate space of measurable
1418 functions w , thus its maximizer (if it exists) is unique.

1419 Our proof can be given by the following steps

1420 1) Define for fixed x_t :

$$1421 \quad \Phi_\epsilon(w; x_t) := -\epsilon \log \mathbb{E}_{p_T(x)} [\exp(\frac{w(x) - \|x - x_t\|_2^2}{\epsilon})], \quad (C.52)$$

1422 The mapping $w \mapsto \mathbb{E}_{p_T(x)} [\exp(\frac{w(x) - C(x, x_t)}{\epsilon})]$ is log-convex by Hölder's inequality, and
1423 therefore, $w \mapsto \Phi_\epsilon(w; x_t)$ is strictly concave, except in directions where w differs only
1424 by an additive constant almost everywhere. Taking the expectation over x_t preserves strict
1425 concavity unless w is constant almost everywhere.

1426 2) The term $-\mathbb{E}_x [f^*(-w(x))]$ is strictly concave with respect to w because f^* is strictly
1427 convex. Specifically, for any distinct $w_1 \neq w_2$, strict convexity of f^* gives for all $\lambda \in$
1428 $(0, 1)$,

$$1429 \quad -\mathbb{E}_x [f^*(-((1 - \lambda)w_1(x) + \lambda w_2(x)))] > -(1 - \lambda)\mathbb{E}_x [f^*(-w_1(x))] - \lambda\mathbb{E}_x [f^*(-w_2(x))]$$

1430 provided $w_1(x) \neq w_2(x)$ on a set of positive measure.

1431 3) Since the sum of a strictly concave function and a concave function is strictly concave,
1432 it follows that the full dual objective $g(w)$ is strictly concave on the set of admissible
1433 functions.

1434 As a result, $g(w)$ admits at most one maximizer, and the proposition is proved. \square

1435 C.6 DERIVATION OF THEOREM 5

1436 **Theorem (5).** *The optimal solution $\rho^*(x)$ to problem defined in Eq. (8) satisfies the following bound:*

$$1437 \quad \mathbb{D}_f[\rho^*(x), p_T(x)] \leq \mathcal{W}_2(p(x_t), p_T(x)). \quad (C.53)$$

1458 According to the definition of the JKO proximal recursion Eqs. (6) and (8), each update $p(x_{t+1})$
 1459 satisfies the following inequality (see Theorem 4.0.4 of reference (Ambrosio et al., 2005)):

$$1460 \quad \frac{1}{2\eta} \mathcal{W}_2^2(p(x_{t+1}), p(x_t)) + \mathbb{D}_f[p(x_{t+1}), p_T(x)] \leq \mathbb{D}_f[p(x_t), p_T(x)],$$

1463 where $\mathbb{D}_f[p(x_t), p_T(x)] \geq 0$ and attains its minimum at the target distribution $p_T(x)$. This inequality
 1464 implies that the total energy decreases at every step, thereby reducing the Wasserstein distance
 1465 to $p_T(x)$ when $\mathbb{D}_f[p(x_{t+1}), p_T(x)]$ is geodesically convex:

$$1466 \quad \mathcal{W}_2(p(x_{t+1}), p_T(x)) \leq \mathcal{W}_2(p(x_t), p_T(x)) - \Delta_t, \quad \Delta_t \geq 0. \quad (\text{C.54})$$

1467 Hence, as t increases, based on Eqs. (C.53) and (C.54) the upper bound on $\mathbb{D}_f[\rho^*(x_t), p_T(x)]$ is
 1468 gradually tightened:

$$1469 \quad \mathbb{D}_f[\rho^*(x_{t+1}), p_T(x)] \leq \mathbb{D}_f[\rho^*(x_t), p_T(x)] - \Delta_t,$$

1471 which shows that the transported probability density $\rho(x)$ progressively becomes more similar to
 1472 the target $p_T(x)$.

1473 *Proof.* To facilitate reading, we define the signal as follows:

$$1475 \quad \mathcal{W}_{2,\epsilon}^2(\rho, \xi) := \inf_{\pi \in \Pi(\rho, \xi)} \iint \|x - y\|_2^2 \pi(x, y) dx dy \\ 1476 \quad + \epsilon \iint \pi(\rho(x), p(x_t)) [\log \pi(\rho(x), p(x_t)) - 1] dx dx_t, \quad (\text{C.55})$$

1480 The dual representation of the f -divergence based on the Legendre–Fenchel conjugate is:

$$1482 \quad \mathbb{D}_f[\rho(x), p_T(x)] = \sup_{v(x)} \{ \mathbb{E}_{\rho(x)}[v(x)] - \mathbb{E}_{p_T(x)}[f^*(v(x))] \}. \quad (\text{C.56})$$

1484 Thus, the problem defined in Eq. (8) can be written as:

$$1485 \quad \inf_{\rho(x)} \mathcal{W}_{2,\epsilon}(\rho(x), p(x_t)) + \sup_{v(x)} \{ \mathbb{E}_{\rho(x)}[v(x)] - \mathbb{E}_{p_T(x)}[f^*(v(x))] \} \quad (\text{C.57})$$

1487 Interchanging $\min_{\rho(x)}$, $\sup_{v(x)}$ by the convexity-concavity and Sion’s theorem (Sion, 1958; Simons,
 1488 1995), we obtain the following result:

$$1490 \quad \sup_{v(x)} -\mathbb{E}_{p_T(x)}[f^*(v(x))] + \inf_{\rho(x)} \{ \mathcal{W}_{2,\epsilon}(\rho(x), p(x_t)) + \mathbb{E}_{\rho(x)}[v(x)] \} \quad (\text{C.58})$$

1492 The inner minimization with respect to $\rho(x)$ is precisely the entropic optimal transport problem in
 1493 the semi-dual form for PDFs $\rho(x)$ and $p(x_t)$:

$$1494 \quad \min_{\rho(x)} \mathcal{W}_{2,\epsilon}(\rho(x), p(x_t)) + \mathbb{E}_{\rho(x)}[v(x)] \quad (\text{C.59})$$

1496 whose optimal value equals

$$1498 \quad \mathbb{E}_{p(x_t)}[-\epsilon \log \int \exp(\frac{v(x) - c(x_t, x)}{\epsilon}) dy]. \quad (\text{C.60})$$

1500 This follows from standard duality in entropic optimal transport.

1501 Plug the expression above into the main problem:

$$1503 \quad \sup_{v(x)} \mathbb{E}_{p(x_t)}[-\epsilon \log \int \exp(\frac{v(x) - c(x_t, x)}{\epsilon}) dy] - \mathbb{E}_{p_T(x)}[f^*(v(x))]. \quad (\text{C.61})$$

1505 This is the desired semi-dual form.

1506 At optimality, plug in any variation $v = v^* + \delta\psi$ into $g(w)$ and take derivative w.r.t. δ at 0, then set
 1507 to zero. The calculation is:

$$1509 \quad 0 = \frac{\partial}{\partial \delta} g(v^* + \delta\psi) \Big|_{\delta=0} = \mathbb{E}_{p(x_t)} \left[\frac{\int \psi(x) \exp\left(\frac{v^*(x) - c(x_t, x)}{\epsilon}\right) dx}{\int \exp\left(\frac{v^*(x) - c(x_t, x)}{\epsilon}\right) dx} \right] - \mathbb{E}_{p_T(x)}[(f^*)'(v^*(x))\psi(x)], \quad (\text{C.62})$$

1512 which for all test functions $\psi(x)$ implies
 1513

$$1514 \quad \underbrace{\int p(x_t) \frac{\exp\left(\frac{v^*(x) - c(x_t, x)}{\epsilon}\right)}{\int \exp\left(\frac{v^*(x) - c(x_t, x)}{\epsilon}\right) dx} dx_t}_{1517 := \tilde{p}_T(x)} = p_T(x)(f^*)'(v^*(x)).$$

1519 That is, the pushforward of $p(x_t)$ under the mapping:
 1520

$$1521 \quad \mathbf{T}^*(x|x_t) = \frac{\exp\left(\frac{v^*(x) - c(x_t, x)}{\epsilon}\right)}{\int \exp\left(\frac{v^*(x) - c(x_t, x)}{\epsilon}\right) dx},$$

1524 which indicates that
 1525

$$\tilde{p}_T(x) = p_T(x)(f^*)'(v^*(x)). \quad (\text{C.63})$$

1526 So, $\rho^*(x) = \tilde{p}_T(x)$ is the marginal of the optimal transport π^* as claimed.
 1527

1528 Since the value of the primal objective at $\rho(x) = p_T(x)$ gives an upper bound:
 1529

$$\mathbb{D}_f[\rho^*(x)\|p_T(x)] + \mathcal{W}_{2,\epsilon}(\rho^*(x), p(x_t)) \leq \mathcal{W}_{2,\epsilon}(p_T, p(x_t)). \quad (\text{C.64})$$

1531 So in particular, we get:
 1532

$$\mathbb{D}_f[\rho^*(x)\|p_T(x)] \leq \mathcal{W}_{2,\epsilon}(p(x_t), p_T(x)). \quad (\text{C.65})$$

1533 In addition, we notice that the following inequality holds for $\epsilon > 0$:
 1534

$$1535 \quad \epsilon \iint \pi(\rho(x), p(x_t)) [\log \pi(\rho(x), p(x_t)) - 1] dx dx_t \leq 0. \quad (\text{C.66})$$

1537 Plugging Eq. (C.66) into Eq. (C.65), we arrive at the desired result. \square
 1538

1539 C.7 DERIVATION OF THEOREM 6

1541 **Theorem (6).** *Under mild assumptions, the E-SUOT-based GDA ensures that the target domain
 1542 generalization error is upper-bounded by the following inequality:*

$$1543 \quad \varepsilon_{p_T}(h_T) \leq \varepsilon_{p_0}(h_0) + \varepsilon_{p_0}(h_T^*) + \iota \zeta \mathcal{C} + \mathcal{S}_{\text{stat}}, \quad (\text{C.67})$$

1545 where ι is the Lipschitz constant of the loss function, ζ is the Lipschitz constant bound for hypotheses
 1546 in \mathcal{H} , \mathcal{C} aggregates the cumulative domain transportation and label continuity costs along the
 1547 adaptation path, and $\mathcal{S}_{\text{stat}}$ is the statistical error term.

1548 Before formally proving the theorem, we introduce the following assumptions, which are mild and
 1549 commonly satisfied in practical domain adaptation scenarios:
 1550

1551 **(A. 1)** The loss function $\mathcal{L}(\cdot, y)$ is ι -Lipschitz with respect to its first argument; that is, for any
 1552 a, a' and fixed y , we have:
 1553

$$1554 \quad |\mathcal{L}(a, y) - \mathcal{L}(a', y)| \leq \iota |a - a'|. \quad (\text{C.68})$$

1555 **(A. 2)** Each hypothesis $h \in \mathcal{H}$ is ζ -Lipschitz, i.e., for any x, x' , we have:
 1556

$$1557 \quad |h(x) - h(x')| \leq \zeta \|x - x'\|. \quad (\text{C.69})$$

1558 **(A. 3)** The labeling function q_t along the adaptation path is such that $|q_t(x) - q_{t-1}(x)|$ is small
 1559 for most x , to ensure local continuity.
 1560

1561 **(A. 4)** The sequence of domains (p_0, p_1, \dots, p_T) is induced by E-SUOT-based GDA transport, so
 1562 that the total cumulative cost \mathcal{C} as defined below is finite.
 1563

1564 **(A. 5)** At every step, empirical risk minimization over sufficient samples ensures a small
 1565 empirical-to-expected error gap, leading to a statistical error term $\mathcal{S}_{\text{stat}}$.
 1566

1567 **(A. 6)** The sample size for each domain is large enough to make $\mathcal{S}_{\text{stat}}$ negligible in the asymptotic
 1568 regime.
 1569

1566 Notably, Assumption (A.1) is standard and well-justified for classification tasks employing the categorical cross-entropy loss. More specifically, the gradient of \mathcal{L}_{CE} with respect to the logits a is
 1567 bounded as
 1568

$$1569 \left| \frac{\partial \mathcal{L}_{\text{CE}}}{\partial a_j} \right| = |[\text{softmax}(a)]_j - \mathbb{I}(j = y)| \leq 1, \\ 1570$$

1571 for any class index j , since each softmax component lies within $[0, 1]$, and the indicator function
 1572 $\mathbb{I}(j = y)$ takes values in $\{0, 1\}$. By the classical Lagrange mean value theorem (see, Theorem
 1573 4 in (Thomas et al., 2014)), this bounded-gradient property implies that \mathcal{L}_{CE} is globally Lipschitz
 1574 continuous with respect to its first argument on any bounded input domain, with a Lipschitz constant
 1575 of at most 1. In practice, this assumption can be further facilitated by applying weight normalization
 1576 or spectral-norm regularization, which help maintain bounded network outputs and thereby make
 1577 the Lipschitz condition more readily satisfied during optimization.

1578 Furthermore, Assumptions (A.2), (A.5) and (A.6) are standard and generally hold for commonly
 1579 used loss functions and hypothesis classes. Unless the loss or model is exceptionally non-standard,
 1580 these can be stated directly with the theorem and do not require additional justification.

1581 Assumption (A.3) holds in cases where the labeling function changes smoothly along the adaptation
 1582 path. For our construction, since the intermediate domains are generated by incremental, continuous
 1583 transformations (e.g., gradual style or environmental shifts), the underlying semantics of inputs
 1584 remain stable, and thus $\mathbb{E}_{p_{t-1}(x)}[q_t(x) - q_{t-1}(x)]$ is small for every t . This situation typically
 1585 occurs under covariate shift, where only the input distribution evolves while class definitions stay
 1586 fixed. However, we acknowledge that this assumption may not hold in fine-grained recognition
 1587 settings or tasks with rapidly changing label semantics, where small variations in features can lead
 1588 to distinct class assignments. In such cases, the theoretical guarantees derived under Assumption
 1589 (A.3) would apply only locally, within regions where the labeling function remains approximately
 1590 smooth. However, the assumption may not hold in tasks with abrupt semantic boundaries—such as
 1591 fine-grained classification—where visually similar samples can belong to distinct categories. Our
 1592 analysis therefore applies when the domain evolution does not induce significant concept shift.

1593 As for Assumption (A.4), in our E-SUOT-based GDA, each domain is generated via an iterative
 1594 unbalanced optimal transport step that progressively reduces the transport cost as we proved in The-
 1595 orem 6. This guarantees that the cumulative cost \mathcal{C} is finite, as can be bounded analytically. In
 1596 summary, all the above assumptions are justified in our setting. Based on these assumptions, we
 1597 now proceed with the formal proof.

1598 *Proof.* Our goal is to bound the target risk $\varepsilon_{p_T}(h_T)$. Consider the telescoping sum along the domain
 1599 adaptation path:

$$1600 \varepsilon_{p_T}(h_T) = \varepsilon_{p_0}(h_0) + [\varepsilon_{p_T}(h_T) - \varepsilon_{p_0}(h_0)]. \quad (\text{C.70})$$

1601 To make the recursion explicit, rewrite this as:

$$1602 \varepsilon_{p_T}(h_T) = \varepsilon_{p_0}(h_0) + \sum_{t=1}^T [\varepsilon_{p_t}(h_t) - \varepsilon_{p_{t-1}}(h_{t-1})]. \quad (\text{C.71})$$

1606 For each $t \in \{0, \dots, T-1\}$, we observe that

$$1607 \varepsilon_{p_t}(h_t) - \varepsilon_{p_{t-1}}(h_{t-1}) \\ 1608 = \underbrace{[\varepsilon_{p_t}(h_t) - \varepsilon_{p_t}(h_{t-1})]}_{\text{optimization error}} + \underbrace{[\varepsilon_{p_t}(h_{t-1}) - \varepsilon_{p_{t-1}}(h_{t-1})]}_{\text{domain shift term}} + \underbrace{[\varepsilon_{p_{t-1}}(h_{t-1}) - \varepsilon_{p_{t-1}}(h_t)]}_{\leq 0 \text{ by ERM}} \quad (\text{C.72})$$

1611 In practice, the last term is non-positive since ‘empirical risk minimization’ (Vapnik, 1999; Shalev-
 1612 Shwartz & Ben-David, 2014; Zhuang et al., 2024) ensures moving toward lower risk, so we can drop
 1613 it for an upper bound.

1614 By the Lipschitz property of \mathcal{L} and h ,

$$1616 |\varepsilon_{p_t}(h) - \varepsilon_{p_{t-1}}(h)| \leq \zeta \cdot \mathcal{W}_1(p_{t-1}, p_t). \quad (\text{C.73})$$

1617 Suppose the true label function q_t changes along the path. Following standard analysis, this gives
 1618 an additional cost due to the label discrepancy:

$$1619 \zeta \mathbb{E}_{p_t(x)} |q_t(x) - q_{t-1}(x)|. \quad (\text{C.74})$$

1620 Therefore, each step can be bounded by
 1621

$$|\varepsilon_{p_t}(h_t) - \varepsilon_{p_{t-1}}(h_{t-1})| \leq \iota\zeta\mathcal{W}_1(p_{t-1}, p_t) + \iota\mathbb{E}_{p_t(x)}|f_t(x) - f_{t-1}(x)| + s_t \quad (\text{C.75})$$

1622 where s_t denotes the statistical error at step t .
 1623

1624 Let

$$\mathcal{C} := \sum_{t=0}^{T-1} \left[\mathcal{W}_1(p_{t-1}, p_t) + \frac{1}{\zeta}\mathbb{E}_{p_t(x)}|q_t(x) - q_{t-1}(x)| \right] \quad (\text{C.76})$$

1625 and

$$\mathcal{S}_{\text{stat}} := \sum_{t=1}^{T-1} s_t. \quad (\text{C.77})$$

1626 Sum these bounds for all $t \in \{0, \dots, T-1\}$, we get:
 1627

$$\sum_{t=0}^{T-1} |\varepsilon_{p_t}(h_t) - \varepsilon_{p_{t-1}}(h_{t-1})| \leq \iota\zeta\mathcal{C} + \mathcal{S}_{\text{stat}}. \quad (\text{C.78})$$

1628 As the final classifier h_T may not be optimally trained with respect to p_0 , include the approximation
 1629 gap:
 1630

$$\varepsilon_{p_0}(h_0) + \varepsilon_{p_0}(h_T^*) - \varepsilon_{p_0}(h_0) \quad (\text{C.79})$$

1631 where h_T^* is the risk minimizer in \mathcal{H} for p_0 .
 1632

1633 Finally,

$$\varepsilon_{p_T}(h_T) \leq \varepsilon_{p_0}(h_0) + \varepsilon_{p_0}(h_T^*) + \iota\zeta\mathcal{C} + \mathcal{S}_{\text{stat}},$$

1634 as desired. \square
 1635

1636 While Theorem 6 provides a clean decomposition, the last two terms are not directly computable
 1637 from data. To bridge this gap between theory and practice, we attempt to estimate them using the
 1638 following strategies:
 1639

- **Loss Lipschitz constant ι :** According to the analysis of Assumption (A.1), we conclude that $\iota < 1$ with the help of applying weight normalization or spectral-norm regularization.
- **Hypothesis Lipschitz constant ζ .** This bounds how sensitively hypotheses $h \in \mathcal{H}$ react to input perturbations. In our implementation process, we use the multi-layer-perceptron with ReLU activated function. Thus, the Lipschitz constant ζ for the classifier can be estimated as follows:

$$\zeta \leq \prod_{\ell=1}^L \|W_\ell\|_2, \quad (\text{C.80})$$

1640 where W_ℓ is the weight of the linear layer, $\|W_\ell\|_2$ is the spectral norm of W_ℓ . Similarly, ζ can be
 1641 controlled by enforcing spectral normalization or weight normalization on layers.
 1642

- **Cumulative cost \mathcal{C} :** Recall Eq. (C.76), we observe that the 1-Wasserstein distance term $\mathcal{W}_1(p_{t-1}, p_t)$ can be approximated using sample-based optimal transport distances, for example, Sinkhorn distance (Cuturi, 2013), computed on intermediate feature representations. Besides, the inter-step labeling-function shift term $\mathbb{E}_{p_t(x)}|q_t(x) - q_{t-1}(x)|$ in Eq. (C.76) can be approximated in practice via pseudo-labels predicted by the models at steps $t-1$ and t , effectively quantifying how much pseudo-labels change along the adaptation path.

- **Statistical error $\mathcal{S}_{\text{stat}}$:** The term $\mathcal{S}_{\text{stat}} = \sum_{t=1}^{T-1} s_t$ collects the statistical deviations between empirical and population risks at each adaptation step. Under mild assumptions, including G -Lipschitz losses and norm-constrained neural networks, standard uniform-convergence bounds based on Rademacher complexity yield the following result, adapted from Section 9.4 of Bach (2024):

$$s_t = \mathcal{O}\left(\frac{1}{\sqrt{N}}\right), \quad (\text{C.81})$$

1669 Consequently, for all $t \in \{0, 1, \dots, T-1\}$, we can estimate $\mathcal{S}_{\text{stat}}$ as follows:
 1670

$$\mathcal{S}_{\text{stat}} = \mathcal{O}\left(\frac{T}{\sqrt{N}}\right). \quad (\text{C.82})$$

1674 C.8 DISCUSSIONS ON THE SELECTION OF STEP SIZE η
1675

1676 Let us recall Eqs. (B.6) and (6) as follows:

1677
$$\begin{cases} \mathcal{L}^{\text{Primal}} = \arg \min_{\rho(x) \in \mathcal{P}_2(\mathbb{R}^D)} \frac{1}{2\eta} \mathcal{W}_2^2(\rho(x), p(x_t)) + \mathbb{D}_f[\rho(x), p_T(x)], \\ \min_{\pi \geq 0} \iint c(x, y) \pi(x, y) dy dx + \lambda_1 \mathbb{D}_f(\tilde{\rho}(x), \rho(x)) + \lambda_2 \mathbb{D}_f(\tilde{\xi}(y), \xi(y)). \end{cases} \quad (\text{C.83})$$

1678

1680 Since Eq. (6) is a variant of Eq. (B.6), where $\lambda_1 \equiv 0$. Based on this, we may raise one question:
1681 How to select the λ_2 , i.e. η ?1683 Since our target is decreasing the functional $\mathbb{D}_f[\rho(x), p_T(x)]$ along the simulation process, thus one
1684 of the key factor is that the selection of the η can decrease the functional $\mathbb{D}_f[\rho(x), p_T(x)]$. Take
1685 the KL divergence, the f -divergence we consider in the proposed E-SUOT approach, we have the
1686 following proposition for selecting the η :1687 **Proposition 7.** Suppose that the $\|\frac{\delta \mathbb{D}_{\text{KL}}[\rho(x), p_T(x)]}{\delta \rho(x)}\| \leq \mathcal{A}$ and $\|\nabla \frac{\delta \mathbb{D}_{\text{KL}}[\rho(x), p_T(x)]}{\delta \rho(x)}\| \leq \mathcal{B}$, there
1688 exists a constant \mathcal{H}_0 control the tailness of $p_T(x)$, and let $\{\rho_t(x) | t = 1, \dots, T\}$ denote the sequence
1689 of empirical PDF of the intermediate domain generated by the JKO recursion. When η satisfies the
1690 following condition, the sequence of KL divergence $\{\mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)] | t = 1, \dots, T\}$ converges
1691 to a finite value as $t \rightarrow \infty$:

1692
$$0 < \eta < \min\left(\frac{1}{\mathcal{B}}, \frac{\mathcal{H}_0}{\mathcal{A}}\right). \quad (\text{C.84})$$

1693

1694 Before proposing the proof, we should introduce the light-tailness property on the target distribution
1695 $p_T(x)$ in order to ensure the validity of our Taylor expansion and to control higher-order discretiza-
1696 tion errors during the proof process. Specifically, we say that $p_T(x)$ is light-tailed (Ambrosio et al.,
1697 2005; Johnson & Zhang, 2018; 2021) if there exists a universal constant $\mathcal{H}_0 < \infty$ such that

1698
$$\int \|\nabla \log p_T(x)\| p_T(x) dx < \mathcal{H}_0. \quad (\text{C.85})$$

1699

1700 We call \mathcal{H}_0 the ‘‘light-tail constant’’ of $p_T(x)$. This condition requires that the expectation (under
1701 $p_T(x)$) of the norm of the score function $\nabla p_T(x)$ is finite. Intuitively, this ensures that $p_T(x)$ decays
1702 sufficiently rapidly in the tails so that the gradients do not blow up at infinity. On this basis, the proof
1703 is articulated as follows:1704 *Proof.* Suppose at time t and time $t + \eta$, we have:

1705
$$x_{t+\eta} = \mathbf{T}(x) := x_t + \eta v_t(x_t). \quad (\text{C.86})$$

1706

1707 We denote the probability distributions, before and after applying Eq. (C.86) as $\rho_t(x)$ and $\rho_{t+\eta}(x)$,
1708 respectively. The aim is to Taylor expand the evolution of

1709
$$\mathbb{D}_{\text{KL}}[\rho_{t+\eta}(x), p_T(x)] = \int \rho_{t+\eta}(x) \log \frac{\rho_{t+\eta}(x)}{p_T(x)} dx, \quad (\text{C.87})$$

1710

1711 with respect to η around $\eta = 0$. The new probability distribution, for small η , can be given as
1712 follows according to the Liouville’s theorem:

1713
$$\rho_{t+\eta}(x) = \rho_t(\mathbf{T}^{-1}(x)) \cdot |\det \mathbf{J}_{\mathbf{T}^{-1}}(z)| \quad (\text{C.88})$$

1714

1715 where $\mathbf{J}_{\mathbf{T}^{-1}}(z)$ is the Jacobian matrix of the inverse map, and when η is small enough, the inverse
1716 function $\mathbf{T}^{-1}(x)$ of function $\mathbf{T}(x)$ can be given as follows:

1717
$$\mathbf{T}^{-1}(x) \approx x - \eta v_t(x). \quad (\text{C.89})$$

1718

1719 Hence, expanding to the first order in η can be given as follows:

1720
$$\rho_{t+\eta}(x) \approx \rho_t(x - \eta v_t(x)) [1 - \eta \nabla \cdot \phi(z)] \approx \rho_t(x) - \eta \nabla [\rho_t(x) v_t(x)]. \quad (\text{C.90})$$

1721

1722 Define $F(\eta)$:

1723
$$F(\eta) := \mathbb{D}_{\text{KL}}[\rho_{t+\eta}(x), p_T(x)] = \int \rho_{t+\eta}(x) \log \frac{\rho_{t+\eta}(x)}{p_T(x)} dx. \quad (\text{C.91})$$

1724

1725 Applying Taylor’s expansion at $\eta = 0$, we can obtain the following result:

1726
$$F(\eta) = F(0) + \eta F'(0) + \mathcal{O}(\eta^2). \quad (\text{C.92})$$

1727

Now, we start deriving $F'(0)$. When we take the derivative inside the integral, we have:

$$F'(\eta) = \frac{d}{d\eta} \int \rho_{t+\eta}(x) \log \frac{\rho_{t+\eta}(x)}{p_T(x)} dx = \int \frac{d}{d\eta} \rho_{t+\eta}(x) [1 + \log \frac{\rho_{t+\eta}(x)}{p_T(x)}] dx \quad (\text{C.93})$$

At $\eta = 0$, $\rho_{t+\eta}(x) = \rho_t(x)$:

$$F'(0) = \int \frac{d}{d\eta} \rho_{t+\eta}(x) \Big|_{\eta=0} [1 + \log \frac{\rho_t(x)}{p_T(x)}] dx. \quad (\text{C.94})$$

Now, using the result from the calculus of variations:

$$\frac{d}{d\eta} \rho_{t+\eta}(x) \Big|_{\eta=0} = -\nabla \cdot (\rho_t(x) v_t(x)) \quad (\text{C.95})$$

Thus,

$$F'(0) = - \int \nabla \cdot (\rho_t(x) v_t(x)) [1 + \log \frac{\rho_t(x)}{p_T(x)}] dx. \quad (\text{C.96})$$

Now, use integration by parts:

$$\int -\nabla \cdot [\rho_t(x) v_t(x)] g(x) dx = \int [\rho_t(x) v_t(x)]^\top \nabla g(x) dx. \quad (\text{C.97})$$

Set $g(x) = 1 + \log \frac{\rho_t(x)}{p_T(x)}$. Its gradient is:

$$\nabla g(x) = \nabla \log \rho_t(x) - \nabla \log p_T(x). \quad (\text{C.98})$$

Hence,

$$F'(0) = \int [\rho_t(x) v_t(x)]^\top [\nabla \log \rho_t(x) - \nabla \log p_T(x)] dx = \mathbb{E}_{\rho_t(x)} [v_t^\top(x) (\nabla \log \rho_t(x) - \nabla \log p_T(x))]. \quad (\text{C.99})$$

But with a negative sign because the original derivative is minus divergence:

$$F'(0) = -\mathbb{E}_{\rho_t(x)} [v_t^\top(x) (\nabla \log \rho_t(x) - \nabla \log p_T(x))]. \quad (\text{C.100})$$

Putting all together, we get:

$$\begin{aligned} & \mathbb{D}_{\text{KL}}[\rho_{t+\eta}(x), p_T(x)] \\ &= \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)] - \eta \mathbb{E}_{\rho_t(x)} [v_t^\top(x) (\nabla \log p_T(x) - \nabla \log \rho_t(x))] + \mathcal{O}(\eta^2). \end{aligned} \quad (\text{C.101})$$

Notably, the optimal velocity field $v_t^*(x)$ for KL divergence is $-\nabla \frac{\delta \mathbb{D}_{\text{KL}}[\rho_t(x)]}{\delta \rho_t(x)}$. Thus, we have:

$$\begin{aligned} & \mathbb{D}_{\text{KL}}[\rho_{t+\eta}(x), p_T(x)] \\ &= \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)] - \underbrace{\eta \mathbb{E}_{\rho_t(x)} [(\nabla \log p_T(x) - \nabla \log \rho_t(x))^\top v_t^*(x)]}_{\leq 0} + \mathcal{O}(\eta^2). \end{aligned} \quad (\text{C.102})$$

Since $\|\nabla \frac{\delta \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)]}{\delta \rho_t(x)}\| \leq \mathcal{B}$, there exists a positive constant C such that:

$$\begin{aligned} & \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)] - \eta \mathbb{E}_{\rho_t(x)} [(\nabla \log p_T(x) - \nabla \log \rho_t(x))^\top v_t^*(x)] + \mathcal{O}(\eta^2) \\ & \leq \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)] - \eta \mathbb{E}_{\rho_t(x)} [(\nabla \log p_T(x) - \nabla \log \rho_t(x))^\top v_t^*(x)] + C\eta^2, \end{aligned} \quad (\text{C.103})$$

where constant C satisfies the following condition:

$$C \propto \mathcal{B}^2. \quad (\text{C.104})$$

To avoid $C\eta^2$ dominating the right-hand-side of Eq. (C.103), we should satisfy the following condition:

$$\frac{1}{\eta^2} \gg \mathcal{B}^2 \Rightarrow \eta \ll \frac{1}{\mathcal{B}} \Rightarrow \eta < \frac{1}{\mathcal{B}}. \quad (\text{C.105})$$

According to the log-Sobolev inequality (Ambrosio et al., 2005; Villani et al., 2009), for the target distribution $p_T(x)$ satisfying the curvature condition controlled by $\mathcal{H}_0 > 0$ (i.e., strong

1782 log-concavity or equivalent tailness control), there exists a log-Sobolev constant \mathcal{H}_0 such that for
 1783 any smooth density $\rho_t(x)$,

$$1785 \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)] \leq \frac{\mathcal{H}_0}{2} \mathbb{E}_{\rho_t(x)} [\|\nabla \log \frac{\rho_t(x)}{p_T(x)}\|^2]. \quad (\text{C.106})$$

1787 Equivalently, the following lower-bound form holds:

$$1789 \mathbb{E}_{\rho_t(x)} [\|v_t^*(x)\|^2] \geq \frac{1}{\mathcal{H}_0} \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)], \quad (\text{C.107})$$

1791 where we used $v_t^*(x) = \nabla \log p_T(x) - \nabla \log \rho_t(x)$.

1792 When the variational derivative $\frac{\delta \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)]}{\delta \rho_t(x)}$ is bounded by \mathcal{A} , and the spatial gradient of this
 1793 functional derivative is bounded by \mathcal{B} , the log-Sobolev inequality admits a perturbation-corrected
 1794 version:

$$1796 \mathbb{E}_{\rho_t(x)} [\|v_t^*(x)\|^2] \geq \frac{1}{\mathcal{H}_0} \{ \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)] - \frac{\mathcal{A}}{\mathcal{H}_0} \}, \quad (\text{C.108})$$

1798 where the correction term $\frac{\mathcal{A}}{\mathcal{H}_0}$ compensates for the bounded $\|\nabla \frac{\delta \mathbb{D}_{\text{KL}}}{\delta \rho}\|$ and ensures dimensional
 1799 consistency of the energy inequality. Plugging Eq. (C.108) into Eq. (C.103) gives:

$$1800 \begin{aligned} & \mathbb{D}_{\text{KL}}[\rho_{t+\eta}(x), p_T(x)] \\ 1801 & \leq \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)] - \eta \mathbb{E}_{\rho_t} [\|v_t^*(x)\|^2] + C\eta^2 \\ 1803 & \leq \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)] - \frac{\eta}{\mathcal{H}_0} \{ \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)] - \frac{\mathcal{A}}{\mathcal{H}_0} \} + C\eta^2, \end{aligned} \quad (\text{C.109})$$

1805 Rearranging Eq. (C.109), we have:

$$1807 \mathbb{D}_{\text{KL}}[\rho_{t+\eta}(x), p_T(x)] \leq (1 - \frac{\eta}{\mathcal{H}_0}) \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)] + \frac{\eta \mathcal{A}}{\mathcal{H}_0^2} + C\eta^2. \quad (\text{C.110})$$

1809 To promise the iteration will gradually reduce $\mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)]$, we should require:

$$1811 (1 - \frac{\eta}{\mathcal{H}_0}) \in (0, 1) \Rightarrow 0 < \eta < \mathcal{H}_0. \quad (\text{C.111})$$

1812 According to Eq. (C.105), ignoring $C\eta^2$, we obtain the equilibrium point, corresponding to the
 1813 steady state as $t \rightarrow \infty$:

$$1815 \mathbb{D}_{\text{KL}}[\rho_\infty(x), p_T(x)] = \frac{\mathcal{A}}{\mathcal{H}_0}. \quad (\text{C.112})$$

1817 Next, to ensure that each discrete update indeed decreases the KL divergence, we impose

$$1819 \mathbb{D}_{\text{KL}}[\rho_{t+\eta}(x), p_T(x)] < \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)]. \quad (\text{C.113})$$

1820 Substituting Eq. (C.110) into Eq. (C.113) yields the following result:

$$1822 -\frac{\eta}{\mathcal{H}_0} \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)] + \frac{\eta \mathcal{A}}{\mathcal{H}_0^2} + C\eta^2 < 0. \quad (\text{C.114})$$

1824 Dividing both sides by $\eta > 0$ and rearranging terms gives the following equation:

$$1826 \mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)] > \frac{\mathcal{A}}{\mathcal{H}_0} + C\eta \mathcal{H}_0. \quad (\text{C.115})$$

1828 In the late stage of GDA task, $\mathbb{D}_{\text{KL}}[\rho_t(x), p_T(x)]$ approaches its equilibrium value
 1829 $\mathbb{D}_{\text{KL}}[\rho_\infty(x), p_T(x)] = \frac{\mathcal{A}}{\mathcal{H}_0}$, so that:

$$1831 \mathbb{D}_{\text{KL}}[\rho_t, p_T] - \mathbb{D}_{\text{KL}}[\rho_\infty, p_T] \approx \mathbb{D}_{\text{KL}}[\rho_\infty, p_T] = \frac{\mathcal{A}}{\mathcal{H}_0}. \quad (\text{C.116})$$

1833 Hence, the typical contraction strength per update is of order:

$$1835 \frac{\eta}{\mathcal{H}_0} (\mathbb{D}_{\text{KL}}[\rho_t, p_T] - \mathbb{D}_{\text{KL}}[\rho_\infty, p_T]) \approx \frac{\eta}{\mathcal{H}_0} \frac{\mathcal{A}}{\mathcal{H}_0}. \quad (\text{C.117})$$

1836 To guarantee that the quadratic residual $C\eta^2$ does not dominate the contraction term in Eq. (C.117),
 1837 we require:

$$1838 \quad C\eta^2 \ll \frac{\eta}{\mathcal{H}_0} \frac{\mathcal{A}}{\mathcal{H}_0} \implies \eta \ll \frac{\mathcal{H}_0}{\mathcal{A}}. \quad (C.118)$$

1840 That is, the discretization step must satisfy the stability condition since $\mathcal{A} > 0$:

$$1842 \quad 0 < \eta < \frac{\mathcal{H}_0}{\mathcal{A}} < \mathcal{H}_0, \quad (C.119)$$

1844 which ensures that the numerical update is dominated by the contraction term rather than by the
 1845 additive bias or high-order error. Based on Eqs. (C.105) and (C.119), we arrive at the desired result.

□

1849 D DETAILED ALGORITHM OF THE E-SUOT FRAMEWORK

1851 While Algorithm 1 outlines the general workflow for generating the intermediate domain, it does
 1852 not specify how E-SUOT can be applied to the GDA task. To bridge this gap, we first present the
 1853 complete workflow for E-SUOT-based GDA in Algorithm 2.

1855 Building on this foundation, the complete workflow for E-SUOT-based gradual domain adaptation
 1856 is summarized in Algorithm 2 based on Algorithm 1. Notably, our algorithm decouples the training
 1857 of the transport function \mathbf{T}_θ from the fine-tuning of the classifier h_ω . This separation allows the
 1858 intermediate domain to be generated offline and subsequently used for online inference, potentially
 1859 reducing overall computation time comparable to traditional GDA approaches.

1860 **Algorithm 2** Overall Workflow for Constraining E-SUOT-based Gradual Domain Adaptation

1861 **Input:** Source domain samples: $\{(x_0^{(i)}, y_0^{(i)})\}_{i=1}^N$, target domain samples: $\{(x_T^{(i)}, y_T^{(i)})\}_{i=1}^N$, entropy
 1862 regularization strength: ϵ , step size: η , number of intermediate domain $T - 1$, neural network batch
 1863 size \mathcal{B} , and neural network training epochs: \mathcal{E} .

1864 **Output:** Classifier in target domain $h_{\omega,T}$.

- 1865 1: Initialize the classifier $h_{\omega,0}$: $h_{\omega,0} \leftarrow \arg \min_{\omega} \mathcal{L}_{\text{CE}}(x_0, h_{\omega,0}, y_0)$.
- 1866 2: Train $\mathcal{T} = \{\mathbf{T}_{\theta,t}\}_{t=1}^{T-1}$: $\mathcal{T} \leftarrow$ Algorithm 1.
- 1867 3: **for** $t = 0$ **to** $T - 1$ **do**
- 1868 4: Obtain the intermediate domain data $\{(x_{t+1}^{(i)}, y_{t+1}^{(i)})\}$: $x_{t+1}^{(i)} \leftarrow \mathbf{T}_{\theta,t}(x_t^{(i)})$ and $y_{t+1}^{(i)} \leftarrow y_t^{(i)}$ for
 1869 all $i \in \{1, \dots, N\}$.
- 1870 5: Finetune the classifier $h_{\omega,t+1}$: $h_{\omega,t+1} \leftarrow \arg \min_{\omega} \mathcal{L}_{\text{CE}}(x_{t+1}, h_{\omega,t}, y_{t+1})$.
- 1871 6: **end for**

1874 E DETAILED INFORMATION FOR EXPERIMENTS

1876 E.1 DATASET DESCRIPTIONS

- 1878 • **Portraits:** Portraits is a binary gender classification dataset comprising 37,921 front-facing
 1879 portrait images collected between 1905 and 2013. Following the chronological split protocol of (Kumar et al., 2020), we divide the data into a source domain (the earliest 2,000 images),
 1880 intermediate domains (14,000 images not utilized in this work), and a target domain (the subsequent 2,000 images), similar to the setting in reference (Zhuang et al., 2024).
- 1882 • **Rotated MNIST:** Rotated MNIST is a variant of the standard MNIST dataset Deng (2012)
 1883 in which images are rotated to create domain adaptation challenges. As described in He
 1884 et al. (2024); Kumar et al. (2020), we use 4,000 source images and 4,000 target images,
 1885 with the target images rotated by 45° to 60° .
- 1887 • **Office-Home:** Office-Home is a domain adaptation benchmark dataset consisting of ap-
 1888 proximately 15,500 images categorized into 65 object classes commonly found in office
 1889 and home environments (Venkateswara et al., 2017). The dataset encompasses four vi-
 sually distinct domains—*Artistic (Ar)*, *Clipart (Cl)*, *Product (Pr)*, and *Real-World (Rw)*.

1890 Following common domain adaptation protocols, one domain is selected as the source do-
 1891 main while another serves as the target domain, resulting in a total of 12 domain transfer
 1892 tasks (e.g., Ar→Rw, Cl→Pr, etc.).
 1893

1894 **E.2 EXPERIMENTAL SETTINGS**

1895 **E.2.1 PRELIMINARIES OF SELF-TRAINING METHOD**

1896 Self-training is a classical semi-supervised learning strategy that leverages the model’s own pre-
 1897 dictions on unlabeled data to iteratively improve its performance. Given a model h_ω trained on labeled
 1898 source data, we use it to generate pseudo-labels for unlabeled samples in a target or auxiliary dataset
 1899 \mathcal{D}_{aux} . Each unlabeled input $x_i \in \mathcal{D}_{\text{aux}}$ is assigned a pseudo-label $\tilde{y}_i = \text{sign}(h_\omega(x_i))$, indicating a
 1900 positive or negative prediction. A new model $h'_{\omega'}$ is then trained to minimize the empirical loss on
 1901 this pseudo-labeled dataset:
 1902

$$1903 \text{ST}(\theta, \mathcal{D}_{\text{aux}}) = \arg \min_{h'_{\omega'} \in \mathcal{H}} \frac{1}{N_{\text{aux}}} \sum_{x_i \in \mathcal{D}_{\text{aux}}} \mathcal{L}(h'_{\omega'}(x_i), \text{sign}(h_\omega(x_i))), \quad (E.1)$$

1904 where ST is the abbreviation of self-training, N_{aux} denotes the sample size of auxiliary dataset \mathcal{D}_{aux} .
 1905 This procedure can be iteratively repeated, replacing θ with the newly optimized θ' to refine the
 1906 pseudo-labels over time.
 1907

1908 Intuitively, self-training alternates between
 1909

- 1910 1) Producing pseudo-labels using the current classifier.
- 1911 2) Retraining the model on these pseudo-labels, thereby progressively refining the decision
 1912 boundary.

1913 In practice, confidence thresholding or pseudo-label sharpening can be incorporated to reduce noise
 1914 accumulation and improve stability.
 1915

1916 **E.2.2 TRAINING & EVALUATION PROTOCOLS**

1917 For GDA and UDA task, we follow the standard domain adaptation protocol, where model training
 1918 and hyperparameter tuning are performed using labeled source data merely since the validation on
 1919 the target domain is infeasible in unsupervised adaptation setting. All results are reported on the
 1920 target domain dataset without using target labels for validation or early stopping.
 1921

1922 For classifier training under this protocol, let $\hat{h}_{\omega,t}(x_t)$ denote the logits produced by the classifier
 1923 parameterized by ω at time step t . We define
 1924

$$1925 h_{\omega,t}(x_t) = \text{softmax}(\hat{h}_{\omega,t}(x_t)) \quad (E.2)$$

1926 as the corresponding class-probability vector. Given the ground-truth label y_t , the categorical cross-
 1927 entropy loss is formulated as follows:
 1928

$$1929 \mathcal{L}_{\text{CE}}(x_t, \omega, y_t) = - \sum_{i=1}^B y_t^{(i)} \log h_{\omega,t}^{(i)}(x_t) = - \sum_{i=1}^B y_t^{(i)} \log [\text{softmax}(\hat{h}_{\omega,t}(x_t))]^{(i)}. \quad (E.3)$$

1930 On this basis, we use the classification accuracy (denoted as “Accuracy”) as the evaluation metric,
 1931 defined as the proportion of correctly predicted target samples:
 1932

$$1933 \text{Accuracy} = \left[\frac{1}{N_{\text{tgt}}} \sum_{i=1}^{N_{\text{tgt}}} \mathbb{I}(\hat{y}_i = y_i) \right] \times 100\%, \quad (E.4)$$

1934 where N_{tgt} is the number of target test samples, \hat{y}_i denotes the predicted label of the i -th sample, y_i
 1935 is the corresponding ground-truth label, and $\mathbb{I}(\cdot)$ is the indicator function that equals 1 when the con-
 1936 dition holds and 0 otherwise. A larger accuracy value corresponds to better adaptation performance.
 1937

1944
1945

E.2.3 GDA TASK

1946
1947
1948
1949
1950
1951
1952
1953
1954

The official implementations of GOAT (He et al., 2024) and CNF (Sagawa & Hino, 2025) are used in our experiments. Additionally, we employ UMAP (McInnes et al., 2018) to reduce the dimensionality of the three GDA datasets to 8 in order to align with the experimental tuple provided by (Zhuang et al., 2024). The experiments are conducted on a workstation equipped with two NVIDIA RTX 4090 GPUs under five different random seeds at least three times. The overall hyper-parameters we use in our GDA task are summarized in Table E.1.

1955
1956
1957
1958
1959
1960
1961
1962

In all experiments, we parameterize the classifier h_ϕ as a three-layer multi-layer perceptron (MLP) at each step, utilizing ReLU activation functions and a hidden dimension of 100 for each layer. For both T_θ and w_ϕ , we employ a two-layer MLP with the SiLU activation function and incorporate a skip connection to enable a residual structure (He et al., 2016). All models are optimized by the Adam optimizer (Kingma & Ba, 2015) with learning rate at 0.0001. For all three GDA datasets, we apply UMAP (McInnes et al., 2018) to reduce their dimensionality to eight. For the classifier h_ω , we use a two-layer MLP with ReLU activations and 128 hidden units. All baseline models are trained on features embedded by the UMAP.

1963
1964
1965
1966
1967
1968
1969

E.2.4 UDA TASK

1970
1971
1972
1973
1974
1975
1976
1977

For the UDA task, we adopt the Office-Home dataset (Venkateswara et al., 2017) as the benchmark to evaluate the performance of the proposed E-SUOT framework. Following the standard unsupervised domain adaptation (UDA) protocol, model training and hyperparameter tuning are performed solely using the labeled source data, without access to target labels for validation or early stopping. All results are reported on the target domain dataset.

1978
1979
1980
1981
1982
1983

We compare E-SUOT with a diverse set of representative UDA approaches, including DANN (Ganin & Lempitsky, 2015), MSTN (Xie et al., 2018), GVB-GD (Cui et al., 2020), RSDA (Gu et al., 2020; 2022), LAMBDA (Le et al., 2021), SENTRY (Prabhu et al., 2021), FixBi (Na et al., 2021), CST (Liu et al., 2021a), CoVi (Na et al., 2022), and GGF (Zhuang et al., 2024). For baselines including DANN, MSTN, GVB-GD, RSDA, LAMBDA, SENTRY, FixBi, CST, and CoVi, we directly report the publicly available results from their original papers under identical experimental settings (i.e., the same dataset and evaluation protocol). For the GGF method, as the public release of GGF provides only code for GDA task, we re-implemented the missing components according to the paper’s description and ran the experiments locally to maintain consistency with the original experimental protocol.

1984
1985
1986
1987
1988
1989

In addition, similar to GGF, our E-SUOT framework builds upon CoVi as the backbone feature extractor. The extracted features are embedded into an eight-dimensional space using UMAP. The classifiers h_ω for GGF and E-SUOT are implemented as two-layer ReLU MLPs with 256 hidden units. We set the η , \mathcal{B} , ϵ , and T as 0.5, 1024, 0.001, and 4, respectively. Both GGF and E-SUOT models are trained under the same conditions for fair comparison.

1990
1991
1992
1993
1994
1995
1996
1997

E.3 DETAILED INFORMATION FOR ABLATION STUDIES

1998
1999

For the ablation study, we ablate two module namely the training strategy of T_θ and the objective functional. The detailed information are elaborated in this part.

For “Training Strategy”, the detailed experimental protocols are given as follows:

- **Adversarial Training:** In our adversarial training scheme, we optimize Eq. (7). Building on (Korotin et al., 2021; 2023; Choi et al., 2023; 2024), the training of T_θ is formulated adversarially, as summarized in Algorithm 3. In Line 5, the penalty term $\frac{1}{2\eta} \|x_{t-1} - T_{\theta,t-1}(x_{t-1})\|_2^2$ is omitted since it is constant with respect to $w_{\phi,t-1}$.
- **Barycentric-based Training:** We propose the algorithm for barycentric-based training in Algorithm 4. For barycentric-based training, rather than first compute the transport map, we attempt to compute the optimal transport map π^* between $\rho(x_{t-1})$ and $p_T(x)$ as we demonstrate in Line 4. Based on this, we make barycentric projection (Courty et al., 2017b; Perrot et al., 2016) using this π^* to obtain the proxy points (Liu et al., 2021b; 2023)

Table E.1: Hyperparameters for E-SUOT on GDA task.

Datasets	η	\mathcal{B}	ϵ	T
Portraits	0.5	1024	0.1	5
MNIST 45°	0.5	1024	0.01	5
MNIST 60°	0.5	2048	0.005	5

1998 for transport map learning as we demonstrate in *Line 5*. Finally, the transport map $\mathbf{T}_{\theta, t-1}$
 1999 is constructed based on these points, similar to the flow matching (Lipman et al., 2023), as
 2000 we demonstrate in *Line 6*.
 2001

Algorithm 3 Adversarial Training for $\{\mathbf{T}_{\theta, t}\}_{t=1}^{T-1}$.

Input: Intermediate domain samples: $\{(x_{t-1}^{(i)}, y_{t-1}^{(i)})\}_{i=1}^N$ for all $t \in \{1, \dots, T\}$, target domain samples: $\{(x_T^{(i)}, y_T^{(i)})\}_{i=1}^N$, entropy regularization strength: ϵ , step size: η , neural network batch size \mathcal{B} , and neural network training epochs: \mathcal{E} .

Output: The transportation map at $t-1$: $\mathbf{T}_{\theta, t-1}$.

```

1: Initialize
2: for  $e = 1$  to  $\mathcal{E}$  do
3:   Sample a batch  $\{x_{t-1}^{(i)}\}_{i=1}^{\mathcal{B}} \sim \{(x_{t-1}^{(i)}, y_{t-1}^{(i)})\}_{i=1}^N$  and  $\{x_T^{(i)}\}_{i=1}^{\mathcal{B}} \sim \{(x_T^{(i)}, y_T^{(i)})\}_{i=1}^N$ .
4:   Update  $w_{\phi, t-1}$  by:  $\phi \leftarrow \arg \min_{\phi} \frac{1}{\mathcal{B}} \sum_{i=1}^{\mathcal{B}} \frac{1}{2\eta} \|x_{t-1}^{(i)} - \mathbf{T}_{\theta, t-1}(x_{t-1}^{(i)})\|_2^2 +$ 
5:    $w_{\phi, t-1}(\mathbf{T}_{\theta, t-1}(x_t^{(i)})) + \frac{1}{\mathcal{B}} \sum_{j=1}^{\mathcal{B}} f^*(-w_{\phi, t-1}(x_T^{(j)}))$ .
6:   Sample a batch  $\{x_{t-1}^{(i)}\}_{i=1}^{\mathcal{B}} \sim \{(x_{t-1}^{(i)}, y_{t-1}^{(i)})\}_{i=1}^N$ .
7:   Update  $\mathbf{T}_{\theta, t-1}$  by:  $\theta \leftarrow \arg \min_{\theta} \frac{1}{\mathcal{B}} \sum_{i=1}^{\mathcal{B}} \frac{1}{2\eta} \|x_{t-1}^{(i)} - \mathbf{T}_{\theta, t-1}(x_{t-1}^{(i)})\|_2^2 -$ 
8:    $w_{\phi, t-1}(\mathbf{T}_{\theta, t-1}(x_{t-1}^{(i)}))$ .
9: end for

```

Algorithm 4 Barycentric-based training for $\{\mathbf{T}_{\theta, t}\}_{t=1}^{T-1}$.

Input: Intermediate domain samples: $\{(x_{t-1}^{(i)}, y_{t-1}^{(i)})\}_{i=1}^N$ for all $t \in \{1, \dots, T\}$, target domain samples: $\{(x_T^{(i)}, y_T^{(i)})\}_{i=1}^N$, entropy regularization strength: ϵ , step size: η , neural network batch size \mathcal{B} , and neural network training epochs: \mathcal{E} .

Output: The transportation map at $t-1$: $\mathbf{T}_{\theta, t-1}$.

```

1: Initialize
2: for  $e = 1$  to  $\mathcal{E}$  do
3:   Sample a batch  $\{x_{t-1}^{(i)}\}_{i=1}^{\mathcal{B}} \sim \{(x_{t-1}^{(i)}, y_{t-1}^{(i)})\}_{i=1}^N$  and  $\{x_T^{(i)}\}_{i=1}^{\mathcal{B}} \sim \{(x_T^{(i)}, y_T^{(i)})\}_{i=1}^N$ .
4:   Obtain the optimal transport map  $\pi^*(x_{t-1}, x_T)$  by:  $\pi^*(x_{t-1}, x_T) \leftarrow \inf_{\pi} \frac{1}{2\eta} \mathcal{W}_2^2(\rho(x_{t-1}), p_T(x)) + \epsilon \iint \pi(x_{t-1}, x_T) [\log \pi(x_{t-1}, x_T) - 1] dx_{t-1} dx_T +$ 
5:    $\mathbb{D}_f[\rho(x_{t-1}), p_T(x)]$ .
6:   Obtain the projected samples  $\tilde{x}_t$  via  $\pi^*(x_{t-1}, x_T)$ :  $\tilde{x}_t = x_{t-1} \pi^*(x_{t-1}, x_T)$ .
7:   Update  $\mathbf{T}_{\theta, t-1}$  by:  $\theta \leftarrow \frac{1}{\mathcal{B}} \sum_{i=1}^{\mathcal{B}} \|\tilde{x}_t^{(i)} - \mathbf{T}_{\theta, t-1}(x_{t-1}^{(i)})\|_2^2$ 
8: end for

```

For “Objective Functional”, the detailed experimental protocols are given as follows:

- χ^2 Divergence: The expression for χ^2 divergence can be given as follows:

$$\mathbb{D}_{\chi^2}[\rho(x_t), p_T(x)] \int p_T(x) \left[\frac{\rho(x_t)}{p_T(x)} - 1 \right]^2 dx_t, \quad \text{where } f(x) = (x-1)^2. \quad (\text{E.5})$$

Based on this, the corresponding conjugate function f^* can be given as follows:

$$f^*(x) = \begin{cases} \frac{1}{4}x^2 + x, & \text{if } x \geq -2 \\ -1, & \text{if } x < -2 \end{cases}. \quad (\text{E.6})$$

- Identity: For the identity function, we remove the f -divergence-based regularization term during the construction of E-SUOT framework. Based on this, the training objective for w_ϕ is reformulated as follows:

$$\mathcal{L}_{\text{Identity}}^{\text{E-SemiDual}} = \sup_w -\epsilon \mathbb{E}_{p(x_t)} \{ \log \mathbb{E}_{p_T(x)} [\exp(\frac{w(x) - \frac{1}{2\eta} \|x - x_t\|_2^2}{\epsilon})] \} + \mathbb{E}_{p_T(x)} [w(x)], \quad (\text{E.7})$$

2052
 2053 • **Softplus:** We directly parameterize the $f^*(x)$ using the smooth, convex, and non-decreasing softplus function as follows:
 2054

2055
$$f^*(x) = \log(1 + \exp(x)). \quad (\text{E.8})$$

 2056

2057 **F ADDITIONAL EXPERIMENTAL RESULTS**
 2058

2059 **F.1 IMPACT OF EARLY TRANSPORT STEPS ON ADAPTATION PERFORMANCE**
 2060

2061 The multi-step structure of E-SUOT involves a series of learned transport maps $\{T_{\theta,0}, T_{\theta,1}, \dots\}$,
 2062 where each step aims to progressively align intermediate feature distributions between domains.
 2063 While this design promotes smooth domain alignment, it also raises an important question: how
 2064 sensitive the overall adaptation performance is to the quality of early transport maps, and whether
 2065 suboptimal early mappings introduce cumulative errors that affect subsequent steps.

2066 To investigate this, we conduct an experiment on the Portraits dataset, where we selectively disable
 2067 the training of certain transport maps to simulate incomplete or inaccurate early-stage optimization.
 2068 Two complementary training strategies are designed:

2069 • **Forward strategy:** Progressively remove the training of early transport maps $(T_{\theta,0}, T_{\theta,1}, \dots)$
 2070 while keeping the later ones active, thereby testing whether missing early steps hinder later GDA
 2071 performance.
 2072 • **Backward strategy:** progressively disable the training of later maps $(T_{\theta,4}, T_{\theta,3}, \dots)$ while retaining
 2073 the trained early steps, examining whether well-trained initial stages are sufficient to sustain
 2074 strong performance.
 2075

2076 Table F.1 summarizes the results. We observe that in the “Forward” direction, excluding early trans-
 2077 port steps causes a substantial accuracy drop (up to 17.7%), indicating that early-stage mappings
 2078 are crucial for forming a reliable transport foundation. In contrast, the “Backward” experiments
 2079 show that once these early steps are properly optimized, subsequent refinements yield consistent
 2080 performance gains, suggesting that later transport maps mainly provide fine adjustments on top of
 2081 an already well-aligned feature space.

2082
 2083 **Table F.1: Performance of the E-SUOT vary different training stage on the Portraits dataset.**
 2084

Direction	Forward					Backward									
	Time Index	$t = 0$	$t = 1$	$t = 2$	$t = 3$	$t = 4$	Accuracy (%)	Δ	$t = 0$	$t = 1$	$t = 2$	$t = 3$	$t = 4$	Accuracy (%)	Δ
Training Status	$t = 0$	✗	✓	✓	✓	✓	76.4 \pm 2.06E-2	\uparrow 7.2%	✓	✗	✗	✗	✗	81.5 \pm 1.70E-2	\uparrow 14.4%
	$t = 1$	✗	✗	✓	✓	✓	75.0 \pm 2.48E-2	\uparrow 5.3%	✓	✓	✗	✗	✗	83.2 \pm 1.13E-2	\uparrow 16.8%
	$t = 2$	✗	✗	✗	✓	✓	74.4 \pm 1.92E-2	\uparrow 4.4%	✓	✓	✓	✗	✗	83.9 \pm 8.10E-3	\uparrow 17.8%
	$t = 3$	✗	✗	✗	✗	✓	74.0 \pm 1.64E-2	\uparrow 3.9%	✓	✓	✓	✓	✗	84.0 \pm 6.56E-3	\uparrow 17.9%
	$t = 4$	✗	✗	✗	✗	✗	58.6 \pm 1.87E-2	\downarrow 17.7%	✓	✓	✓	✓	✓	86.4 \pm 8.72E-2	\uparrow 21.5%

2091 *Kindly Note:* Δ denotes performance change percentage of the initial classifier accuracy.
 2092

2093 Overall, the results highlight that the early transport steps are the key drivers of successful adap-
 2094 tation. When the early mappings are well trained, the rest of the chain benefits from a stabilized
 2095 feature representation, leading to larger and more consistent improvements. This behavior parallels
 2096 diffusion-like processes, where the early transport transformations largely determine the shape and
 2097 quality of the final distribution, as also illustrated in Fig. 3 in reference (Caluya & Halder, 2020)
 2098 and Fig. 1 in reference (Liu & Wang, 2016).
 2099

2100 **F.2 PERFORMANCE IMPROVEMENT VARYING DIFFERENT UDA FEATURE EXTRACTORS**
 2101

2102 In addition to the results reported in Table 2, we further assess the effectiveness of the proposed E-
 2103 SUOT framework by integrating it with four representative UDA backbones—MSTN, RSDA, FixBi,
 2104 and CoVi—on the Office-Home dataset. To facilitate a fair and consistent comparison, we apply
 2105 UMAP-based dimensionality reduction to obtain an 8-dimensional embedding. The corresponding
 results are summarized in Table F.2.

2106

2107 Table F.2: Accuracy (%) improvement over different UDA feature extractors.

Method	Ar→Cl	Ar→Pr	Ar→Rw	Cl→Ar	Cl→Pr	Cl→Rw	Pr→Ar	Pr→Cl	Pr→Rw	Rw→Ar	Rw→Cl	Rw→Pr	Avg.
MSTN	49.8	70.3	76.3	60.4	68.5	69.6	61.4	48.9	75.7	70.9	55	81.1	65.7
E-SUOT+MSTN	57.8	75.9	79.6	65.5	75.9	74.8	64.5	58.5	81.4	73.7	59.5	84.4	71
Δ	↑16.1%	↑8.0%	↑4.3%	↑8.4%	↑10.8%	↑7.5%	↑5.0%	↑19.6%	↑7.5%	↑3.9%	↑8.2%	↑4.1%	↑8.1%
RSDA	53.2	77.7	81.3	66.4	74	76.5	67.9	53	82	75.8	57.8	85.4	70.9
E-SUOT+RSDA	61.5	78.8	81.7	67.6	77.3	77.6	67.2	61	82.7	76	62.4	85.3	73.3
Δ	↑15.7%	↑1.4%	↑0.5%	↑1.8%	↑4.4%	↑1.5%	↓1.0%	↑15.2%	↑0.9%	↑0.3%	↑7.9%	↓0.1%	↑3.3%
FixBi	58.1	77.3	80.4	67.7	79.5	78.1	65.8	57.9	81.7	76.4	62.9	86.7	72.7
E-SUOT+FixBi	61.7	79.1	81.7	67.6	77.6	78.2	67.3	61.3	82.7	76	62.5	85.3	73.4
Δ	↑6.2%	↑2.3%	↑1.6%	↓0.1%	↓2.4%	↑0.1%	↑2.3%	↑5.9%	↑1.2%	↓0.5%	↓0.6%	↓1.6%	↑1.0%
CoVi	58.5	78.1	80	68.1	80	77	66.4	60.2	82.1	76.6	63.6	86.5	73.1
E-SUOT+CoVi	61.6	79.3	81.8	67.6	77.7	78.1	67.4	61.2	82.9	76.3	62.5	85.2	73.5
Δ	↑5.3%	↑1.5%	↑2.2%	↓0.7%	↓2.9%	↑1.4%	↑1.5%	↑1.7%	↑1.0%	↓0.4%	↓1.7%	↓1.5%	↑0.5%

2118 *Kindly Note:* Δ denotes performance change percentage of the vanilla UDA feature extractor compared to E-SUOT.

2119

2120

2121 As shown in Table F.2, E-SUOT consistently enhances the performance of all baseline methods, 2122 achieving an average improvement of 8.1%, 3.3%, 1.0%, and 0.5% for MSTN, RSDA, FixBi, and 2123 CoVi, respectively. The largest relative gain (+16.1%) is observed on the challenging “Ar→Cl” 2124 transfer, indicating that E-SUOT is particularly effective when the domain gap is large. Even when 2125 combined with more recent and competitive adaptation approaches (e.g., FixBi, CoVi), E-SUOT still 2126 yields moderate yet consistent improvements, demonstrating its strong complementary capability 2127 rather than competing nature. Overall, the consistent gains across different backbones highlight the 2128 adaptability, scalability, and robustness of the E-SUOT framework in handling UDA task.

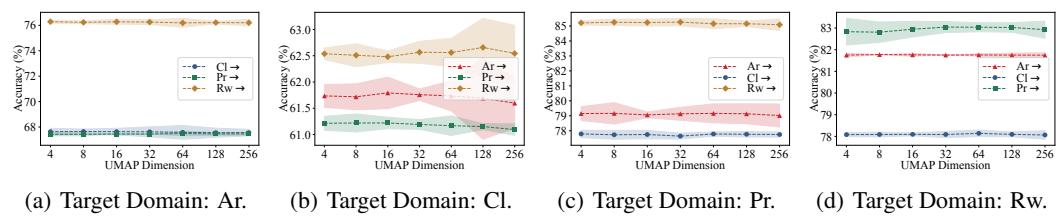
2129

2130 F.3 INVESTIGATION ON THE UMAP’S DIMENSION

2131

2132 We further conduct experiments on the sensitivity of the UMAP dimension to investigate the performance 2133 of E-SUOT under different input dimensions. Specifically, we vary the UMAP’s reduction 2134 dimension and investigate the model performance. Since the backbone dimension is 256, setting the 2135 dimension to 256 indicates that we did not use the UMAP embedding feature.

2136 From Fig. F.1, we observe that as the embedding dimension changes, the classification accuracy 2137 tends to fluctuate within a relatively small range across all target domains. In detail, for each target 2138 domain, the performance remains quite stable as we vary the UMAP dimension from 4 to 256—no 2139 drastic drops are observed. In some domains (e.g., Ar in Fig. F.1(a), Pr in Fig. F.1(c) and Rw 2140 in Fig. F.1(d)), the accuracy curves are almost flat, indicating the proposed method is insensitive 2141 to UMAP dimension choices in these cases. For domain Cl in Fig. F.1(b), although the standard 2142 deviation is higher, the main trend is still relatively stable. It suggests that E-SUOT retains strong 2143 robustness to the UMAP embedding dimension and does not rely heavily on fine-tuning this hyper- 2144 parameter. In summary, the sensitivity analysis shows that E-SUOT remains robust to variations in 2145 the UMAP embedding dimension.



2155 Figure F.1: Sensitivity of E-SUOT performance to the UMAP embedding dimension on the Office- 2156 Home dataset under the UDA setting. For dimension 256, the vanilla backbone features are used 2157 without UMAP. The feature extractor is pre-trained using the CoVI method. The shaded area indicates 2158 the ± 5.0 times standard deviation error.

2159

2160 **G LIMITATIONS & FUTURE DIRECTIONS AND BROADER IMPACT**
21612162 **G.1 LIMITATIONS & FUTURE DIRECTIONS**
21632164 The limitations and future research directions of this work can be summarized as follows:
2165

- **Consideration of Label Information:** In this work, we focused primarily on feature adaptation and did not explicitly incorporate label or discriminator information into the adaptation process. As a result, the performance of the proposed E-SUOT framework may degrade under scenarios involving significant covariate shift (Sugiyama et al., 2007; Sugiyama & Kawanabe, 2012). An important direction for future research is to integrate label information into the transportation process, for example, classifier guidance approach (Courty et al., 2017a; Dhariwal & Nichol, 2021; Bonet et al., 2025; Zhuang et al., 2024), which could further enhance model robustness and adaptation performance.
- **Regularization for Transport Plan:** To facilitate computation, we introduced entropy regularization on the transport plan; however, this may introduce potential instability or blur sparsity in the map (Yin et al., 2025). Future work may explore alternative regularization strategies (Courty et al., 2014; 2017b), such as group sparsity (to better incorporate label priors) or Laplacian regularization (to preserve local relationships), in order to further stabilize training and improve the properties of the learned potential function w .
- **Exploration of Other Discrepancy:** In this work, we adopted the Wasserstein distance as the primary metric for measuring domain discrepancy. However, other discrepancy measures, such as the Fisher-Rao distance (Zhang et al., 2022; Wang et al., 2023; Zhu, 2025), could also be explored to enable more flexible or principled adaptation approaches. Future work may investigate the use of alternative metrics (Neklyudov et al., 2023; Skreta et al., 2025) to further improve the effectiveness of the quality of intermediate domain thereby improving the performance of GDA.
- **Assumption of Label Invariance along the Transport Path:** The current formulation assumes that labels remain invariant during adaptation, i.e., $y_{t+1} \leftarrow y_t$, and thus primarily focuses on aligning the marginal feature distributions $p(x_t)$. This assumption may limit performance under pronounced label shift scenarios, where the conditional relationship $p(y|x)$ varies across domains, a case often encountered in unbalanced or fine-grained settings. Although our framework can be extended by incorporating classifier uncertainty or pseudo-label refinement (drawing inspiration from self-training schemes such as Eq. (3)–(4) in reference (Kumar et al., 2020)), handling substantial concept drift remains an open challenge. Future work may consider integrating adaptive label transport (Courty et al., 2017a) or uncertainty-aware pseudo-labeling (Kumar et al., 2020; Zhuang et al., 2024) to explicitly account for label-shift dynamics along the adaptation trajectory.
- **Higher Efficiency Utilization of Neural Networks:** In our current design, each stage requires training three separate networks, namely w_ϕ , T_θ , and h_ω , to generate each intermediate domain. Although this strategy can save computation time compared to existing approaches that perform intermediate-domain generation online during the domain adaptation stage, it may still be suboptimal in terms of overall training efficiency in the offline stage. A promising future direction is to reformulate the training of T_θ into a more parameter-efficient form, such as adopting a LoRA-style adaptation (Hu et al., 2022) or using reparameterization trick to parameterize the difference between different stage (Choi et al., 2024). For w_ϕ and h_ω , one possible improvement is to fine-tune only the last layer (Harrison et al., 2024; Branzema et al., 2025), which could further reduce the offline training cost.

2205 **G.2 BROADER IMPACT STATEMENT**
22062207 GDA addresses a critical challenge in machine learning: transferring knowledge from a labeled
2208 source domain to an unlabeled target domain when there is a substantial gap between the two.
2209 Rather than relying on abrupt, one-shot shifts—which are often brittle in the face of large distri-
2210 butional discrepancies—GDA interpolates through a series of intermediate domains, allowing for
2211 a smoother and more effective adaptation process. This paradigm has direct implications for many
2212 real-world applications. For example, in recommender systems, GDA enables knowledge transfer to
2213 serve cold-start users or to integrate new items, and in language processing it allows models trained
on high-resource languages to adapt more robustly to low-resource languages. By constructing and

2214 navigating intermediate distributions, GDA provides a principled foundation for bridging domain
2215 gaps and ensuring stable model performance under challenging conditions. Our work advances
2216 the field of GDA by unifying flow-based methods and optimal transport within the semi-dual for-
2217 mulation, identifying fundamental issues of stability and generalization that have limited previous
2218 approaches. We further propose theoretically-grounded regularization strategies that improve the
2219 robustness and reliability of the adaptation process. These advances not only deepen the theoretical
2220 understanding of GDA but also offer practical benefits for deploying adaptable machine learning
2221 systems in diverse settings. We believe our findings will help catalyze the development of more
2222 general, stable, and information-preserving domain adaptation methods, with impact across fields
2223 ranging from recommendation and computational linguistics to broader AI applications.

2224

2225 H LLM USAGE STATEMENT

2226

2227 In accordance with the conference guidelines, we disclose our use of Large Language Models
2228 (LLMs) in the preparation of this paper as follows:

2229

2230 We used LLMs (specifically, OpenAI GPT-4.1, GPT-5 and Google Gemini 2.5) *solely for checking*
2231 *grammar errors and improving the readability of the manuscript*. The LLMs *were not involved in*
2232 *research ideation, the development of research contributions, experiment design, data analysis, or*
2233 *interpretation of results*. All substantive content and scientific claims were created entirely by the
2234 authors. The authors have reviewed all LLM-assisted text to ensure accuracy and originality, and
2235 take full responsibility for the contents of the paper. The LLMs are not listed as an author.

2236

2237

2238

2239

2240

2241

2242

2243

2244

2245

2246

2247

2248

2249

2250

2251

2252

2253

2254

2255

2256

2257

2258

2259

2260

2261

2262

2263

2264

2265

2266

2267