



HAL
open science

Regularization for Wasserstein Distributionally Robust Optimization

Waïss Azizian, Franck Iutzeler, Jérôme Malick

► **To cite this version:**

Waïss Azizian, Franck Iutzeler, Jérôme Malick. Regularization for Wasserstein Distributionally Robust Optimization. 2022. hal-03837681v1

HAL Id: hal-03837681

<https://hal.science/hal-03837681v1>

Preprint submitted on 3 Nov 2022 (v1), last revised 23 Mar 2023 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Regularization for Wasserstein Distributionally Robust Optimization

Waïss Azizian^{*†}

Franck Iutzeler[†]

Jérôme Malick[‡]

Abstract

Optimal transport has recently proved to be a useful tool in various machine learning applications needing comparisons of probability measures. Among these, applications of distributionally robust optimization naturally involve Wasserstein distances in their models of uncertainty, capturing data shifts or worst-case scenarios. Inspired by the success of the regularization of Wasserstein distances in optimal transport, we study in this paper the regularization of Wasserstein distributionally robust optimization. First, we derive a general strong duality result of regularized Wasserstein distributionally robust problems. Second, we refine this duality result in the case of entropic regularization and provide an approximation result when the regularization parameters vanish.

1 Introduction

Optimal transport (OT) has a long history and exciting recent developments, notably around applications in machine learning and data science; we refer to the monographs [Villani \(2003\)](#), [Santambrogio \(2015\)](#), [Peyré et al. \(2019\)](#), and [Merigot and Thibert \(2021\)](#). One of the key technical properties at the core of recent success of OT in these applications is the use of regularization, and specifically entropic regularization, opening the way to efficient computational schemes (see e.g., [Cuturi \(2013\)](#)) to get theoretically-grounded approximations of the Wasserstein distances.

Distributionally robust optimization (DRO) has recently been formulated using OT metrics and has proven to be useful in machine learning (see e.g., [Kuhn et al. \(2019\)](#)). But regularization has still to be studied and used in this context. In this paper, we propose a study of regularization in Wasserstein distributionally robust optimization (WDRO), inspired from several recent developments in OT, namely [Genevay et al. \(2016\)](#), [Carlier et al. \(2017\)](#), [Genevay et al. \(2019\)](#), and [Paty and Cuturi \(2020\)](#).

1.1 Distributionally robust optimization with Wasserstein neighborhoods

DRO is a popular approach in optimization under uncertainty. We briefly present here the ideas and the notation that we will use; we refer to the celebrated paper [Esfahani and Kuhn \(2018\)](#) and the survey paper [Kuhn et al. \(2019\)](#) for more details and for applications in machine learning.

^{*}DI, ENS, Univ. PSL, 75005, Paris, France

[†]Univ. Grenoble Alpes, 38000, Grenoble, France

[‡]CNRS & LJK, 38000, Grenoble, France

Standard approaches in stochastic optimization consider minimizing the expectation of a random loss with respect to some input distribution or available data points: For an objective $f_\theta : \Xi \rightarrow \mathbb{R}$ defined on a sample space Ξ and depending on parameters $\theta \in \Theta$, this consists in considering

$$\min_{\theta \in \Theta} \mathbb{E}_{\xi \sim P}[f_\theta(\xi)].$$

Here P is a fixed probability distribution on Ξ ; in practice, it is typically an empirical distribution $P = \frac{1}{n} \sum_{i=1}^n \delta_{\xi_i}$ stemming from data samples $(\xi_i)_{i=1}^n$.

A DRO counterpart of this problem is to minimize the expectation of the loss with respect to a set of probability distributions close to P . More precisely, we choose a neighborhood $\mathcal{U}(P)$ of P (called the ambiguity set or the distributional uncertainty region) in the set of probabilities measures on Ξ , denoted by $\mathcal{P}(\Xi)$ and consider the worst possible expectation of the objective in this neighborhood. The resulting problem is thus of the form

$$\min_{\theta \in \Theta} F(\theta) \quad \text{where} \quad F(\theta) := \sup_{Q \in \mathcal{U}(P)} \mathbb{E}_{\xi \sim Q}[f_\theta(\xi)]. \quad (1)$$

A natural way to define the ambiguity set $\mathcal{U}(P)$ is to consider a ball centered at P with radius $\rho > 0$ controlling the required level of robustness. When using the Wasserstein distance to define the ball, this gives so-called Wasserstein DRO problems (WDRO).

For a cost function $c : \Xi \times \Xi \rightarrow \mathbb{R}_+$, the Wasserstein distance between two probability distributions $P, Q \in \mathcal{P}(\Xi)$ is defined as

$$W_c(P, Q) = \inf \left\{ \mathbb{E}_{(\xi, \zeta) \sim \pi} c(\xi, \zeta) : \pi \in \mathcal{P}(\Xi \times \Xi), \pi_1 = P, \pi_2 = Q \right\},$$

where π_1 and π_2 denote the first and second marginals of the coupling probability, or transport plan, π defined on $\Xi \times \Xi$. A WDRO problem thus has the form (1) with the ambiguity set

$$\mathcal{U}(P) = \{Q \in \mathcal{P}(\Xi) : W_c(P, Q) \leq \rho\}. \quad (2)$$

When the the objective f_θ exhibits a simple structure, this problem reformulates as a tractable convex optimization problem; see e.g., [Kuhn et al. \(2019\)](#). This is exploited in several applications, for instance: logistic regression (see e.g., [Yu et al. \(2021\)](#)), support vector machines (see e.g., [Shafieezadeh-Abadeh et al. \(2019\)](#)), or ℓ^1 -regression (see e.g., [\(Chen and Paschalidis, 2020\)](#)). Another argument supporting a WDRO approach for machine learning applications is that it provides generalization guarantees, see e.g., [Esfahani and Kuhn \(2018\)](#); [An and Gao \(2021\)](#).

1.2 Contributions, related works, and outline

In this paper we study regularization in the context of Wasserstein distributionally robust optimization. First, we propose a unified framework for double regularization of the WDRO objective function (both in the objective and in the constraint). We then provide a strong duality result with general convex regularization functions. This result can be seen as the analogue for WDRO of the general result of [Paty and Cuturi \(2020\)](#) for OT. Second, we refine our analysis in the case of the entropic regularization and obtain an explicit expression for the dual problem. Furthermore, we provide approximation guarantees when the regularization parameters are driven to 0, adapting the reasoning from [Carlier et al. \(2017\)](#). These results can be seen as analogues for WDRO of results of [Genevay et al. \(2019\)](#) for the entropic-regularized OT. Note however, that the reasoning

and results used OT do not directly apply in the context of WDRO. Indeed, in OT, the entropic regularization depends on (the product of) the two marginal distributions whereas, in WDRO, the second marginal is not fixed but rather an optimization variable, which makes the analysis different and more involved. This distinction is further discussed in the beginning of [Section 3](#).

Up to our knowledge, regularization in the context of WDRO has not been investigated on its own, as we do in this paper.¹ Nevertheless, let us mention the two recent papers related to our developments: [Blanchet and Kang \(2020\)](#) and [Wang et al. \(2021\)](#). In [Blanchet and Kang \(2020\)](#), an entropic smoothing of a specific WDRO dual function is introduced and used for computational purposes. Such dual smoothing implicitly corresponds to a regularization of the associated primal problem, but this link is not formally made. In contrast, the preprint [Wang et al. \(2021\)](#) (which appeared while we were preparing this manuscript) shares similar spirit as our work. The entropic regularization of WDRO is proposed and analyzed, with a special focus on computational aspects. This is complementary to our work which provides a theoretical study of general regularizations as well as approximation guarantees for the entropic regularization. We will come back more precisely on the connections between our results and those of [Blanchet and Kang \(2020\)](#) and [Wang et al. \(2021\)](#) in [Remark 1](#).

The outline of this paper is the following. The introduction ends below with the definition of the framework of this paper. In [Section 2](#), we provide a duality result for a general double regularization, together with an illustration in the case when the transport cost is used a regularization function. In [Section 3](#), we focus our analysis to the entropic regularization to get refined expressions and an explicit control of the quality of the approximation of the underlying WDRO problem.

1.3 Set-up, notation, and assumptions

The framework of this paper is the following. With Ξ a subset of \mathbb{R}^d , P a reference probability distribution over Ξ , and $f: \Xi \rightarrow \mathbb{R}$ the underlying objective function (we drop the dependence in θ to simplify the notation), we consider the sup problem in the objective function of [\(1\)](#) with the Wasserstein ball of radius ρ as an ambiguity set [\(2\)](#). Our objective thus writes:

$$\sup\{\mathbb{E}_{\xi \sim Q} f(\xi) : Q \in \mathcal{P}(\Xi), W_c(P, Q) \leq \rho\}.$$

We reformulate the above problem, in a concise way, using only couplings as

$$\sup_{\pi \in \mathcal{P}_P(\Xi \times \Xi): \mathbb{E}_{\pi} c \leq \rho} \mathbb{E}_{\pi_2} f \tag{WDRO}$$

where $\mathcal{P}_P(\Xi \times \Xi)$ is the set of probability distributions on $\Xi \times \Xi$ having P as a first marginal

$$\mathcal{P}_P(\Xi \times \Xi) := \{\pi \in \mathcal{P}(\Xi \times \Xi) : \pi_1 = P\}.$$

When the space Ξ is compact, we have that the topological dual of $\mathcal{C}(\Xi \times \Xi)$, the set of continuous functions on $\Xi \times \Xi$, is exactly $\mathcal{M}(\Xi \times \Xi)$, the set of finite signed measures over $\Xi \times \Xi$ by the Riesz representation ([Rudin, 1987](#), Th. 2.14). We denote by $\langle \cdot, \cdot \rangle$ the duality pairing between $\mathcal{C}(\Xi \times \Xi)$ and $\mathcal{M}(\Xi \times \Xi)$:

$$\begin{cases} \mathcal{M}(\Xi \times \Xi) \times \mathcal{C}(\Xi \times \Xi) & \longrightarrow \mathbb{R} \\ (\pi, \varphi) & \longmapsto \langle \pi, \varphi \rangle := \int \varphi \, d\pi . \end{cases}$$

¹Let us mention that studying “regularization for WDRO” as we propose here should not be confused with studying “the regularizing effect of WDRO on learning problems”, which is a separate field of study (see e.g., [Blanchet et al. \(2019\)](#); [Shafieezadeh-Abadeh et al. \(2019\)](#))

When establishing general duality results, we will also make a constant use of the convex conjugate of a function $F: \mathcal{C}(\Xi \times \Xi) \rightarrow \mathbb{R} \cup \{+\infty\}$ defined as

$$F^*: \begin{cases} \mathcal{M}(\Xi \times \Xi) & \longrightarrow \mathbb{R} \cup \{+\infty\} \\ \pi & \longmapsto \sup_{\varphi \in \mathcal{C}(\Xi)} \langle \pi, \varphi \rangle - F(\varphi), \end{cases}$$

as well as the preconjugate of a function $G: \mathcal{M}(\Xi \times \Xi) \rightarrow \mathbb{R} \cup \{+\infty\}$ defined as

$$G_*: \begin{cases} \mathcal{C}(\Xi \times \Xi) & \longrightarrow \mathbb{R} \cup \{+\infty\} \\ \varphi & \longmapsto \sup_{\pi \in \mathcal{M}(\Xi)} \langle \pi, \varphi \rangle - G(\pi). \end{cases}$$

In presence of convexity, these two operations are dual one another; see e.g., [Clason and Valkonen \(2020, Rem. 5.2\)](#). More precisely, we have that $(F^*)_* = F$ when F is lower semi-continuous (l.s.c.), convex, and proper, and $(G_*)^* = G$ when G is weakly- \star l.s.c., convex, and proper. Furthermore, the following duality result will be instrumental in our developments; it is a reformulation ([Bot et al., 2009, Th. 3.2.6](#)), adapted to our purposes.

Lemma 1.1 (General Fenchel duality). *Consider a compact subset $\mathcal{X} \subset \mathbb{R}^d$, a function $h \in \mathcal{C}(\mathcal{X})$, and two functions $F, G: \mathcal{M}(\mathcal{X}) \rightarrow \mathbb{R} \cup \{+\infty\}$ convex weakly- \star l.s.c. proper. If there exists $\varphi \in \text{dom } F_* \cap (h - \text{dom } G_*)$ such that F_* is continuous at φ , then*

$$\sup_{\pi \in \mathcal{M}(\mathcal{X})} \langle \pi, h \rangle - F(\pi) - G(\pi) = \inf_{\varphi, \psi \in \mathcal{C}(\mathcal{X}): \varphi + \psi = h} F_*(\varphi) + G_*(\psi).$$

Proof. First, the right-hand side (RHS) can be rewritten as

$$\inf_{\varphi, \psi \in \mathcal{C}(\mathcal{X}): \varphi + \psi = h} F_*(\varphi) + G_*(\psi) = \inf_{\varphi \in \mathcal{C}(\mathcal{X})} F_*(\varphi) + G_*(h - \varphi).$$

Then, since F is convex weakly- \star l.s.c., we get that $F = (F_*)^*$ and that F_* is proper, convex, and l.s.c.; see ([Clason and Valkonen, 2020, Rem. 5.2](#)). The same holds for G .

We can thus apply the duality result from ([Bot et al., 2009, Th. 3.2.6](#)) with F_* and $G_*(h - \cdot)$ as primal functions that are proper convex functions, and $\mathcal{C}(\mathcal{X})$ as primal space. Indeed, the regularity assumption of the lemma exactly gives the regularity condition (RC_1^{id}) of this result. Hence,

$$\begin{aligned} \inf_{\varphi \in \mathcal{C}(\mathcal{X})} F_*(\varphi) + G_*(h - \varphi) &= \sup_{\pi \in \mathcal{M}(\mathcal{X})} -(F_*)^*(-\pi) - (G_*(h - \cdot))^*(\pi) \\ &= \sup_{\pi \in \mathcal{M}(\mathcal{X})} -F(-\pi) - G(-\pi) - \langle \pi, h \rangle. \end{aligned}$$

Carrying out the change of variable $\pi \leftarrow -\pi$ then concludes the proof. ■

2 Regularization of the WDRO objective function

In this section, we study (WDRO) objectives with additional regularization functions both in the constraints $\mathbb{E}_{\pi} c \leq \rho$ and in the objective $\mathbb{E}_{\pi_2} f$. For two arbitrary convex functions $R, S: \mathcal{M}(\Xi \times \Xi) \rightarrow \mathbb{R} \cup \{+\infty\}$, the regularized objective we consider is

$$\sup_{\pi \in \mathcal{P}_{\mathbb{P}}(\Xi \times \Xi): \mathbb{E}_{\pi} c + S(\pi) \leq \rho} \mathbb{E}_{\pi_2} f - R(\pi). \quad (\text{R-WDRO})$$

We give in [Theorem 2.1](#) the expression of the dual of this problem, under the following assumptions.

- Assumption 1.** (i) $\Xi \subset \mathbb{R}^d$ is convex and compact;
(ii) $f: \Xi \rightarrow \mathbb{R}$ and $c: \Xi \times \Xi \rightarrow \mathbb{R}_+$ are both continuous;
(iii) For all ξ in Ξ , $c(\xi, \xi) = 0$.

The dual expression for (R-WDRO) with general R and S is the analogue for WDRO of the main result of [Paty and Cuturi \(2020\)](#) which gives the dual of the OT problem regularized by a general convex function.

The formulation of the dual of the regularized problem can also be seen as a generalization of the existing one for the non-regularized case ($R = S = 0$), see e.g., [Blanchet and Murthy \(2019\)](#); [Gao and Kleywegt \(2016\)](#). Note however that the duality results in these two papers rely on weaker assumptions; in particular, Ξ is not assumed to be compact and f, c are only upper- and lower-semicontinuous respectively. In the following result, these assumptions are needed to handle general regularizations.

Theorem 2.1 (Strong duality for doubly-regularized WDRO). *Let [Assumption 1](#) hold and take two convex, proper, and weakly- \star l.s.c. functions $R, S: \mathcal{M}(\Xi \times \Xi) \rightarrow \mathbb{R} \cup \{+\infty\}$, such that $R + S$ is also proper. If the primal problem (R-WDRO) is strictly feasible (i.e., if there exists $\pi \in \mathcal{P}_{\mathbb{P}}(\Xi \times \Xi) \cap \text{dom } R$ such that $\mathbb{E}_{\pi} c + S(\pi) < \rho$), then we have*

$$\text{val (R-WDRO)} = \inf_{\lambda \geq 0} \inf_{\varphi \in \mathcal{C}(\Xi \times \Xi)} \lambda \rho + \mathbb{E}_{\xi \sim \mathbb{P}} \left[\sup_{\zeta \in \Xi} f(\zeta) - \lambda c(\xi, \zeta) - \varphi(\xi, \zeta) \right] + (R + \lambda S)_{\star}(\varphi), \quad (3)$$

and there exists a primal optimal solution $\pi^{\star} \in \mathcal{P}_{\mathbb{P}}(\Xi \times \Xi)$ and an optimal dual parameter $\lambda^{\star} \geq 0$.

We prove this theorem by carefully combining two standard duality results (the Lagrangian duality theorem in Banach spaces [Peypouquet \(2015\)](#) and the Fenchel duality theorem recalled in [Lemma 1.1](#)) with a powerful theorem for exchanging minimization/integration of [Rockafellar and Wets \(1998\)](#).

Proof. The first step of the proof consists in applying the Lagrangian duality theorem of [Peypouquet \(2015, Thm. 3.68\)](#). Let us carefully check its assumptions. First note that Slater's condition holds by assumption, so we only need to check that the primal problem has at least a solution. This is the case by the following arguments: (i) the problem is feasible by assumption; (ii) $\mathcal{P}(\Xi \times \Xi)$ is weakly- \star sequentially compact (see e.g., [Brezis \(2010, Cor. 3.30\)](#)); (iii) the constraint set $\{\pi \in \mathcal{P}_{\mathbb{P}}(\Xi \times \Xi) : \mathbb{E}_{\pi} c + S(\pi) \leq \rho\}$ is weakly- \star closed (since S is weakly- \star l.s.c. and the constraint $\pi_1 = \mathbb{P}$ is weakly- \star closed); and (iv) the objective $\pi \mapsto \mathbb{E}_{(\xi, \zeta) \sim \pi} [f(\zeta) - \lambda c(\xi, \zeta)] - R(\pi)$ is weakly- \star upper semi-continuous (u.s.c.) by assumption. As a result, we have Lagrangian duality and existence of a dual solution

$$\text{val (R-WDRO)} = \inf_{\lambda \geq 0} \sup_{\pi \in \mathcal{P}(\Xi \times \Xi) : \pi_1 = \mathbb{P}} \mathbb{E}_{(\xi, \zeta) \sim \pi} [f(\zeta) - \lambda c(\xi, \zeta)] - (R + \lambda S)(\pi) + \lambda \rho \quad (4)$$

The next step is to write the inner sup as an inf. For concision, let us introduce $T_{\lambda} := R + \lambda S$ and $F_{\lambda}(\xi, \zeta) := f(\zeta) - \lambda c(\xi, \zeta)$. The sup over π , for a fixed $\lambda \geq 0$, thus writes

$$\sup_{\pi \in \mathcal{P}(\Xi \times \Xi) : \pi_1 = \mathbb{P}} \mathbb{E}_{(\xi, \zeta) \sim \pi} [f(\zeta) - \lambda c(\xi, \zeta)] - T_{\lambda}(\pi) = \sup_{\pi \in \mathcal{M}(\Xi \times \Xi)} \langle \pi, F_{\lambda} \rangle - \iota_{\mathcal{P}_{\mathbb{P}}(\Xi \times \Xi)}(\pi) - T_{\lambda}(\pi)$$

where $\iota_{\mathcal{P}_P(\Xi \times \Xi)}$ is the indicator function in the sense of convex analysis, i.e., for $\pi \in \mathcal{M}(\Xi \times \Xi)$, $\iota_{\mathcal{P}_P(\Xi \times \Xi)}(\pi) = 0$ if $\pi \in \mathcal{P}(\Xi \times \Xi)$ and $\pi_1 = P$, and $+\infty$ otherwise.

Now, we want to apply the duality result of [Lemma 1.1](#) with $F \leftarrow \iota_{\mathcal{P}_P(\Xi \times \Xi)}$, $G \leftarrow T_\lambda$, and $h \leftarrow F_\lambda$. To do so, we need to derive the (pre)conjugates of $\iota_{\mathcal{P}_P(\Xi \times \Xi)}$ and T_λ .

By the disintegration theorem, any coupling $\pi(d\xi, d\zeta)$ can be written as $P(d\xi)Q(d\zeta|\xi)$ with Q conditional probability on Ξ . Therefore,

$$\begin{aligned} (\iota_{\mathcal{P}_P(\Xi \times \Xi)})_*(\varphi) &= \sup \left\{ \mathbb{E}_{(\xi, \zeta) \sim \pi} \varphi(\xi, \zeta) : \pi \in \mathcal{P}(\Xi \times \Xi), \pi_1 = P \right\} \\ &= \sup \left\{ \mathbb{E}_{\xi \sim P} \mathbb{E}_{\zeta \sim Q(\cdot|\xi)} \varphi(\xi, \zeta) : Q(\cdot|\cdot) \text{ conditional probability on } \Xi \right\} \\ &\leq \mathbb{E}_{\xi \sim P} \left[\sup_{\zeta \in \Xi} \varphi(\xi, \zeta) \right]. \end{aligned}$$

We now proceed to the reverse inequality, noting that a measurable map $\zeta: \Xi \rightarrow \Xi$ induces a conditional probability $Q(\cdot|\xi) = \delta_{\zeta(\xi)}$. Hence,

$$\begin{aligned} (\iota_{\mathcal{P}_P(\Xi \times \Xi)})_*(\varphi) &= \sup \left\{ \mathbb{E}_{\xi \sim P} \mathbb{E}_{\zeta \sim Q(\cdot|\xi)} \varphi(\xi, \zeta) : Q(\cdot|\cdot) \text{ conditional probability on } \Xi \right\} \\ &\geq \sup \left\{ \mathbb{E}_{\xi \sim P} \varphi(\xi, \zeta(\xi)) : \zeta: \Xi \rightarrow \Xi \text{ measurable} \right\} \end{aligned}$$

where the final sup is finite since Ξ is assumed to be compact, and φ to be continuous. Now, let us define N as the $\Xi \times \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ mapping defined as $N(\xi, \zeta) = -\varphi(\xi, \zeta)$ if $\zeta \in \Xi$ and $+\infty$ otherwise. Since φ is continuous and Ξ is closed, N is jointly l.s.c., and, as a consequence, it is a normal integrand by ([Rockafellar and Wets, 1998](#), Ex. 14.31). So we can apply ([Rockafellar and Wets, 1998](#), Thm. 14.60) to get that

$$\inf \left\{ \mathbb{E}_{\xi \sim P} -\varphi(\xi, \zeta(\xi)) : \zeta: \Xi \rightarrow \Xi \text{ measurable} \right\} = \mathbb{E}_{\xi \sim P} \left[\inf_{\zeta \in \Xi} -\varphi(\xi, \zeta) \right].$$

Inverting the signs, we have showed both upper and lower-inequalities, which means that

$$(\iota_{\mathcal{P}_P(\Xi \times \Xi)})_*(\varphi) = \mathbb{E}_{\xi \sim P} \left[\sup_{\zeta \in \Xi} \varphi(\xi, \zeta) \right] \quad \text{for any } \varphi \in \mathcal{C}(\Xi \times \Xi),$$

and thus $\text{dom}(\iota_{\mathcal{P}_P(\Xi \times \Xi)})_* = \mathcal{C}(\Xi \times \Xi)$. Also, $\varphi \in \mathcal{C}(\Xi \times \Xi) \mapsto \sup_{\zeta \in \Xi} \varphi(\cdot, \zeta) \in \mathcal{C}(\Xi)$ is 1-Lipschitz w.r.t. the norm of the uniform convergence so $(\iota_{\mathcal{P}_P(\Xi \times \Xi)})_*$ is continuous on its domain.

Moreover, since T_λ is convex, proper and weakly- \star l.s.c., $T_{\lambda*}$ is proper and therefore $\text{dom } T_* \neq \emptyset$. Thus, [Lemma 1.1](#) can be used and gives that

$$\sup_{\pi \in \mathcal{P}(\Xi \times \Xi): \pi_1 = P} \mathbb{E}_{(\xi, \zeta) \sim \pi} [f(\zeta) - \lambda c(\xi, \zeta)] - T(\pi) = \inf_{\varphi \in \mathcal{C}(\Xi \times \Xi)} \mathbb{E}_{\xi \sim P} \sup_{\zeta \in \Xi} f(\zeta) - \lambda c(\xi, \zeta) - \varphi(\xi, \zeta) + T_{\lambda*}(\varphi),$$

which, combined with (4), leads to the claimed result. \blacksquare

As an illustration of this duality result, let us consider the case when the transport cost itself is used a regularization. In this case, the expression of the dual simplifies as follows. This expression will be used in the analysis of the next section.

Corollary 2.2 (Duality for cost-regularized WDRO). *Let [Assumption 1](#) hold and take $\varepsilon, \delta > 0$. Then we have*

$$\sup_{\pi \in \mathcal{P}_P(\Xi \times \Xi): \mathbb{E}_{\pi} c + \delta \mathbb{E}_{\pi} c \leq \rho} \mathbb{E}_{\pi_2} f - \varepsilon \mathbb{E}_{\pi} c = \inf_{\lambda \geq 0} \lambda \rho + \mathbb{E}_{\xi \sim P} \sup_{\zeta \in \Xi} f(\zeta) - (\varepsilon + (1 + \delta)\lambda)c(\xi, \zeta).$$

Proof. Define $R : \pi \in \mathcal{M}(\Xi \times \Xi) \mapsto \varepsilon \langle \pi, c \rangle$, $S : \pi \in \mathcal{M}(\Xi \times \Xi) \mapsto \delta \langle \pi, c \rangle$ which are convex, proper and weakly- \star continuous by construction. The pre-conjugate of their sum is $(R + \lambda S)_* = \iota_{\{(\varepsilon + \lambda \delta)c\}}$. Moreover, the primal is strictly feasible thanks to the transport plan $\pi(d\xi, d\zeta) = P(d\xi) \delta_\xi(d\zeta)$. Thus, we can apply [Theorem 2.1](#) to get the expression. \blacksquare

3 Entropic regularization

In this section, we specialize and refine the study of the previous section in the case entropic regularization, i.e., when the Kullback-Leibler (KL) divergence is used a regularizing function. This regularization is defined for two signed measures with finite variations μ, ν as

$$\text{KL}(\mu|\nu) = \begin{cases} \int \log \frac{d\mu}{d\nu} d\mu & \text{if } \mu \text{ and } \nu \text{ are non-negative and } \mu \ll \nu \\ +\infty & \text{otherwise} \end{cases}.$$

This kind of regularization is very popular in computational OT, where it enables the derivation useful approximations of the Wasserstein distance, as, for example the so-called Sinkhorn distance:

$$\inf_{\pi \in \mathcal{P}(\Xi \times \Xi): \pi_1 = P, \pi_2 = Q} \mathbb{E}_\pi c + \varepsilon \text{KL}(\pi | P \otimes Q) \quad \text{for } P, Q \in \mathcal{P}(\Xi) \text{ given.} \quad (5)$$

When Ξ is compact, the dual of this problem is given in [Genevay et al. \(2016, Prop. 2.1\)](#)

$$\sup_{f \in \mathcal{C}(\Xi)} \mathbb{E}_P f - \varepsilon \mathbb{E}_{\xi \sim P} \log \left(\mathbb{E}_{\zeta \sim Q} e^{\frac{f(\xi) - c(\xi, \zeta)}{\varepsilon}} \right). \quad (6)$$

In [Section 3.1](#), we establish a similar result for WDRO. But let us point out here a technical difficulty arising from the WDRO framework compared to OT. The KL divergence (5) is taken w.r.t. the measure $P \otimes Q$, which does not restrict the set of feasible transport plans (if $\pi \in \mathcal{P}(\Xi \times \Xi)$ satisfies $\pi_1 = P$ and $\pi_2 = Q$, π is indeed absolutely continuous w.r.t. $P \otimes Q$). In WDRO however, only one marginal of the transport plans π is fixed and the support of the optimal coupling can be arbitrary, so that the same regularization as in (5) cannot directly be used.

We thus propose to regularize (WDRO) with KL using a base coupling π_0 with first marginal $(\pi_0)_1 = P$ and consider

$$\sup_{\pi \in \mathcal{P}_P(\Xi \times \Xi): \mathbb{E}_\pi c + \delta \text{KL}(\pi | \pi_0) \leq \rho} \mathbb{E}_{\pi_2} f - \varepsilon \text{KL}(\pi | \pi_0). \quad (\text{E-WDRO})$$

The choice of π_0 restricts the set of transport plans to those that are absolutely continuous w.r.t. π_0 . We see in [Section 3.2](#) that this restriction has a very limited impact since a natural choice of π_0 still provides a good approximations of the original problem (WDRO) when the regularization parameters ε, δ vanish.

3.1 Duality for the entropy-regularized problem

We derive here a duality theorem for (E-WDRO) involving the KL regularization with an arbitrary base coupling $\pi_0 \in \mathcal{P}_P(\Xi \times \Xi)$. The obtained result naturally involves the same features as in (6). Note also that a similar result for WDRO (in the case of the KL-regularization in constraints only)

appears in (Wang et al., 2021, Thm. 1), based on arguments specific to the KL divergence. Here we derive our duality result with double KL regularizations, by building on Theorem 2.1 with

$$R(\pi) = \varepsilon \text{KL}(\pi | \pi_0) \quad \text{and} \quad S(\pi) = \delta \text{KL}(\pi | \pi_0). \quad (7)$$

Theorem 3.1 (Strong duality for entropy-regularized problems). *Let Assumption 1 hold, take $\varepsilon, \delta > 0$, and fix an arbitrary $\pi_0 \in \mathcal{P}_{\mathbb{P}}(\Xi \times \Xi)$. If the primal problem (E-WDRO) is strictly feasible (i.e., if there exists π such that $\mathbb{E}_{\pi_0} c + \delta \text{KL}(\pi | \pi_0) < \rho$), then*

$$\text{val}(\text{E-WDRO}) = \inf_{\lambda \geq 0} \lambda \rho + (\varepsilon + \lambda \delta) \mathbb{E}_{\xi \sim \mathbb{P}} \log \left(\mathbb{E}_{\zeta \sim \pi_0(\cdot | \xi)} e^{\frac{f(\zeta) - \lambda c(\xi, \zeta)}{\varepsilon + \lambda \delta}} \right), \quad (8)$$

and there exists a primal optimal solution $\pi^* \in \mathcal{P}_{\mathbb{P}}(\Xi \times \Xi)$ and an optimal dual parameter $\lambda^* \geq 0$.

The interest of using the entropy as a regularization appears when comparing (8) to the general dual (3). We see that there is no inf on φ and moreover the inner sup is replaced by a smoothed approximation (of the log-sum-exp type).

Remark 1 (Two closely related results). We observe that the expression (8) is proposed in Blanchet and Kang (2020), independently from KL/regularization considerations, in a specific context of semi-supervised learning with a finite set Ξ and the uniform distribution over Ξ for π_0 .

The most closely related work is Wang et al. (2021) which proposes and analyses a KL regularization of WDRO. In particular, a very similar dual expression is given. The main difference is that we study here a double regularization with an arbitrary π_0 , while the paper considers only a regularization of the Wasserstein distance ($R = 0$) with a specific π_0 . Also, our duality result holds under stronger assumptions (including a Slater-like assumption), which additionally provides existence of primal and dual optimal solutions.

Proof. First, we can check that R and S as defined in (7) do satisfy the regularity conditions of Theorem 2.1: R and S are proper (since they are non-negative and $R(\pi_0) = S(\pi_0) = 0$); R and S are convex by (Feydy et al., 2019, Prop. 6); R and S are weakly- \star l.s.c. by (Feydy et al., 2019, Prop. 8). Moreover, (Feydy et al., 2019, Prop. 7) tells us that $\text{KL}(\cdot | \pi_0)$ is the conjugate of $\varphi \mapsto \langle \pi_0, e^\varphi - 1 \rangle$ which is convex, proper and l.s.c., so that we get

$$(R + \lambda S)_* : \varphi \mapsto (\varepsilon + \lambda \delta) \langle \pi_0, e^{\frac{\varphi}{\varepsilon + \lambda \delta}} - 1 \rangle.$$

We can then apply Theorem 2.1 with the regularizations defined in (7) to get

$$\text{val}(\text{E-WDRO}) = \inf_{\lambda \geq 0} \lambda \rho + \underbrace{\inf_{\varphi \in \mathcal{C}(\Xi \times \Xi)} \mathbb{E}_{\xi \sim \mathbb{P}} \left[\sup_{\zeta \in \Xi} f(\zeta) - \lambda c(\xi, \zeta) - \varphi(\xi, \zeta) \right]}_{(a)} + (\varepsilon + \lambda \delta) \langle \pi_0, e^{\frac{\varphi}{\varepsilon + \lambda \delta}} - 1 \rangle.$$

For a fixed $\lambda \geq 0$, we can simplify the expression of the term (a) above by introducing $\tau := \varepsilon + \lambda \delta$, defining $F_\lambda : (\xi, \zeta) \mapsto f(\zeta) - \lambda c(\xi, \zeta) \in \mathcal{C}(\Xi \times \Xi)$, and carrying out the change of variable $\varphi \leftarrow F_\lambda - \varphi$. We thus obtain

$$\begin{aligned} (a) &= \inf_{\varphi \in \mathcal{C}(\Xi \times \Xi)} \mathbb{E}_{\xi \sim \mathbb{P}} \left[\sup_{\zeta \in \Xi} f(\zeta) - \lambda c(\xi, \zeta) - \varphi(\xi, \zeta) \right] + \tau \mathbb{E}_{\pi_0} \left[e^{\frac{\varphi}{\tau}} - 1 \right] \\ &= \inf_{\varphi \in \mathcal{C}(\Xi \times \Xi)} \mathbb{E}_{\xi \sim \mathbb{P}} \left[\sup_{\zeta \in \Xi} \varphi(\xi, \zeta) \right] + \tau \mathbb{E}_{\pi_0} \left[e^{\frac{F_\lambda - \varphi}{\tau}} - 1 \right]. \end{aligned} \quad (9)$$

The rest of the proof is devoted to the reformulation of the term (a) expressed as (9). In order to get rid of the supremum in this expression, we restrict the minimization to $\mathcal{C}(\Xi)$ instead of $\mathcal{C}(\Xi \times \Xi)$:

$$\text{val (9)} \leq \inf_{g \in \mathcal{C}(\Xi)} \mathbb{E}_{\xi \sim P} [g(\xi)] + \tau \mathbb{E}_{(\xi, \zeta) \sim \pi_0} \left[e^{\frac{F_\lambda(\xi, \zeta) - g(\xi)}{\tau}} - 1 \right]. \quad (10)$$

Let us prove that the inequality above is actually an equality. Fix a bivariate function $\varphi \in \mathcal{C}(\Xi \times \Xi)$ and consider the univariate function $g^\varphi := \sup_{\zeta \in \Xi} \varphi(\cdot, \zeta)$. Then, since $-\varphi(\xi, \zeta) \geq -g^\varphi(\xi)$, we have

$$\mathbb{E}_{\xi \sim P} \left[\sup_{\zeta \in \Xi} \varphi(\xi, \zeta) \right] + \tau \mathbb{E}_{\pi_0} \left[e^{\frac{F_\lambda - \varphi}{\tau}} - 1 \right] \geq \mathbb{E}_{\xi \sim P} [g^\varphi(\xi)] + \tau \mathbb{E}_{(\xi, \zeta) \sim \pi_0} \left[e^{\frac{F_\lambda(\xi, \zeta) - g^\varphi(\xi)}{\tau}} - 1 \right]. \quad (11)$$

However, g^φ is not continuous in general but only l.s.c., hence we cannot lower bound the RHS of (11) by the RHS of (10). To remedy this issue, we approximate g^φ with continuous functions defined for $k \geq 1$ by

$$g_k^\varphi(\xi) := \inf_{\zeta \in \Xi} g^\varphi(\zeta) + k \|\xi - \zeta\|.$$

Since g^φ is l.s.c., the functions (g_k^φ) converge pointwise to g^φ when k goes to $+\infty$, see e.g., (Rockafellar and Wets, 1998, Ex. 9.11). Moreover, these functions are uniformly bounded by $\sup_{\xi \in \Xi} |g^\varphi(\xi)|$, thus Lebesgue's dominated convergence implies that

$$\begin{aligned} \mathbb{E}_{\xi \sim P} \left[\sup_{\zeta \in \Xi} \varphi(\xi, \zeta) \right] + \tau \mathbb{E}_{\pi_0} \left[e^{\frac{F_\lambda - \varphi}{\tau}} - 1 \right] &\geq \mathbb{E}_{\xi \sim P} [g^\varphi(\xi)] + \tau \mathbb{E}_{(\xi, \zeta) \sim \pi_0} \left[e^{\frac{F_\lambda(\xi, \zeta) - g^\varphi(\xi)}{\tau}} - 1 \right] \\ &= \lim_{k \rightarrow +\infty} \mathbb{E}_{\xi \sim P} [g_k^\varphi(\xi)] + \tau \mathbb{E}_{(\xi, \zeta) \sim \pi_0} \left[e^{\frac{F_\lambda(\xi, \zeta) - g_k^\varphi(\xi)}{\tau}} - 1 \right] \\ &\geq \inf_{g \in \mathcal{C}(\Xi)} \mathbb{E}_{\xi \sim P} [g(\xi)] + \tau \mathbb{E}_{(\xi, \zeta) \sim \pi_0} \left[e^{\frac{F_\lambda(\xi, \zeta) - g(\xi)}{\tau}} - 1 \right]. \end{aligned} \quad (12)$$

Combining (10) and (12) gives that

$$(a) = \inf_{g \in \mathcal{C}(\Xi)} \mathbb{E}_{\xi \sim P} [g(\xi)] + \tau \mathbb{E}_{(\xi, \zeta) \sim \pi_0} \left[e^{\frac{F_\lambda(\xi, \zeta) - g(\xi)}{\tau}} - 1 \right] \quad (13)$$

The final step of the proof consists in solving the above minimum over g . Indeed, the objective

$$g \mapsto \mathbb{E}_{\xi \sim P} [g(\xi)] + \tau \mathbb{E}_{(\xi, \zeta) \sim \pi_0} \left[e^{\frac{F_\lambda(\xi, \zeta) - g(\xi)}{\tau}} - 1 \right]$$

is convex and differentiable on $\mathcal{C}(\Xi)$. As a consequence, the critical points are minimizers. The gradient at a continuous function $g \in \mathcal{C}(\Xi)$ lives in $\mathcal{M}(\Xi)$ and is given by

$$P(d\xi) - \left(\int_{\Xi} e^{\frac{F_\lambda(\xi, \zeta) - g(\xi)}{\tau}} \pi_0(d\zeta|\xi) \right) P(d\xi). \quad (14)$$

Thus the continuous function

$$g^* : \xi \mapsto \tau \log \left(\int_{\Xi} e^{\frac{F_\lambda(\xi, \zeta)}{\tau}} \pi_0(d\zeta|\xi) \right)$$

is a solution of (13). We get

$$(a) = \mathbb{E}_{\xi \sim \mathcal{P}} \left[\tau \log \left(\int_{\Xi} e^{\frac{F_{\lambda}(\xi, \zeta)}{\tau}} \pi_0(d\zeta|\xi) \right) \right] + \underbrace{\tau \mathbb{E}_{\xi \sim \mathcal{P}} \left[\int_{\Xi} e^{\frac{F_{\lambda}(\xi, \zeta) - g^*(\xi)}{\tau}} \pi_0(d\zeta|\xi) - 1 \right]}_{=0 \text{ (by nullity of the gradient in (14) for } g^*)}$$

which in turn gives the desired expression (8). \blacksquare

The proof of this result thus relies on basic optimization rationale and explicit calculus, from the expression of [Theorem 2.1](#). Let us note that an alternative proof could be obtained from duality formulas for variational inference; see ([Lee, 2021](#), Thm. 2.1) which is itself inspired by [Boucheron et al. \(2005, Lem. 1\)](#).

3.2 Approximation error of entropy-regularized problems

In this section, we study the behavior of the approximation error of (E-WDRO) as the regularization parameters vanish to 0. To quantify this approximation, we specify the cost c and the reference measure π_0 . Specifically, we consider that the cost c is a norm to some power $p \geq 1$

$$c(\xi, \zeta) = \|\xi - \zeta\|^p \tag{15}$$

and that the reference measure $\pi_0 \in \mathcal{P}_{\mathbb{P}}(\Xi \times \Xi)$ is taken, for some $\sigma > 0$, as

$$\pi_0(d\xi, d\zeta) \propto \mathbb{P}(d\xi) \mathbb{1}_{\zeta \in \Xi} e^{-\frac{c(\xi, \zeta)}{2^{p-1}\sigma}} d\zeta. \tag{16}$$

For example, when $c(\xi, \zeta) = \|\xi - \zeta\|_2^p$ with $p \in \{1, 2\}$, $\pi_0(\cdot|\xi)$ is a Laplace or a Gaussian distribution (which is easy to sample from for any $\xi \in \Xi$). We also slightly strengthen [Assumption 1](#) by assuming that Ξ is a convex body and that the functions are Lipschitz continuous.

Theorem 3.2 (Approximation for entropic regularization). *Let the following conditions hold:*

- (i) *the objective $f: \Xi \rightarrow \mathbb{R}$ and the cost $c: \Xi \times \Xi \rightarrow \mathbb{R}_+$ are Lipschitz continuous;*
- (ii) *the cost c and the coupling π_0 are taken as (15) and (16) with $\sigma > 0$ such that $\mathbb{E}_{\pi_0} c < \rho$;*
- (iii) *the set $\Xi \subset \mathbb{R}^d$ is compact, convex, with nonempty interior.*

Then, as the regularization parameters $\varepsilon, \delta > 0$ go to zero, we have

$$0 \leq \text{val}(\text{WDRO}) - \text{val}(\text{E-WDRO}) \leq \mathcal{O} \left(d(\varepsilon + \bar{\lambda}\delta) \log \frac{1}{\varepsilon + \bar{\lambda}\delta} \right)$$

where $\bar{\lambda} := \frac{2 \sup_{\Xi} |f|}{\rho - \mathbb{E}_{\pi_0} c}$ is an explicit dual bound.

This result is the analogue for (WDRO) of quantitative bounds for (5) established in [Genevay et al. \(2019\)](#) in the context of OT. The core of the proof consists in introducing a block approximation of the optimal transport plan, following [Carlier et al. \(2017\)](#). In our situation, we modify this approximation scheme to take into account the fact that the second marginals of the transport plans π are not fixed. We also need to introduce an auxiliary regularized problem. Thus the proof of [Theorem 3.2](#) requires several original steps, as described in the next section.

3.3 Proof of the approximation theorem

This section is devoted to the proof of [Theorem 3.2](#). In fact, we state and prove a slightly more detailed result, formalized in the next theorem, and we show afterwards how [Theorem 3.2](#) can be derived as a consequence. The following theorem can indeed be seen as a global version of [Theorem 3.2](#), with explicit constants and slightly more general assumptions. To simplify the reading, we denote by the optimal solution of the *entropy*-regularized problem.

$$F_\rho^{\varepsilon, \delta}(f) = \sup_{\pi \in \mathcal{P}_P(\Xi \times \Xi): \mathbb{E}_\pi c + \delta \text{KL}(\pi | \pi_0) \leq \rho} \mathbb{E}_{\pi_2} f - \varepsilon \text{KL}(\pi | \pi_0).$$

Theorem 3.3 (Extended approximation theorem). *Take a radius $\rho > 0$, regularization parameters $\varepsilon, \delta > 0$, and suppose that the following conditions hold:*

- (i) *The objective $f: \Xi \rightarrow \mathbb{R}$ and the cost $c: \Xi \times \Xi \rightarrow \mathbb{R}_+$ are Lipschitz continuous, and that their respective Lipschitz constants satisfy $\varepsilon \leq \text{Lip}(f)$ and $\delta \leq \text{Lip}(c)$;*
- (ii) *The cost c and the coupling π_0 are taken as [\(15\)](#) and [\(16\)](#) with $\sigma > 0$ such that $\mathbb{E}_{\pi_0} c < \rho$;*
- (iii) *The set $\Xi \subset \mathbb{R}^d$ is compact, convex, and satisfies (for $\mathbb{B}(\xi, \Delta)$ the ball for $\|\cdot\|$)*

$$V := \inf_{\xi \in \Xi, 0 < \Delta \leq d} \frac{\text{vol}(\Xi \cap \mathbb{B}(\xi, \Delta))}{\Delta^d} > 0. \quad (17)$$

Then, we have,

$$\begin{aligned} F_{\frac{\rho}{1+\delta/\sigma}}^{0,0}(f) - (\varepsilon + \bar{\lambda}\delta) \left(d + d \log \left(\frac{L}{(\varepsilon + \bar{\lambda}\delta)d} \right) + C + \frac{1}{\sigma} \left(\frac{(\varepsilon + \bar{\lambda}\delta)d}{L} \right)^p \right) - \frac{\varepsilon\rho}{\sigma + \delta} \\ \leq F_\rho^{\varepsilon, \delta}(f) \leq F_\rho^{0,0}(f). \end{aligned}$$

where $\bar{\lambda} = \frac{2 \sup_{\Xi} |f|}{\rho - \mathbb{E}_{\pi_0} c}$, $L = \text{Lip}(f) + \bar{\lambda} \text{Lip}(c)$, and $C = \min \left\{ \log \frac{\text{vol}(\Xi)}{V}, \log \frac{I_\sigma}{V} \right\}$ with $I_\sigma = \sigma^{\frac{d}{p}} \int_{\mathbb{R}^d} e^{-\frac{\|\zeta\|^p}{2^{p-1}}} d\zeta$.

The proof of this result requires a few preliminary steps. First, we provide in [Lemma 3.4](#) a simple approximation result for the *cost*-regularized problem

$$G_\rho^{\varepsilon, \delta}(f) = \sup_{\pi \in \mathcal{P}_P(\Xi \times \Xi): \mathbb{E}_\pi c + \delta \mathbb{E}_\pi c \leq \rho} \mathbb{E}_{\pi_2} f - \varepsilon \mathbb{E}_\pi c. \quad (18)$$

Next, we bound in [Lemma 3.5](#) the dual optimal solution of the *entropy*-regularized problem $F_\rho^{\varepsilon, \delta}(f)$. Finally, for a fixed dual variable, we compare in [Lemma 3.6](#) the values of the Lagrangians of the *entropy*-regularized problem and the *cost*-regularized one, which is the most technical part of the proof. After these three lemmas, we prove the approximation result in the extended version ([Theorem 3.3](#)), and show how the initial version ([Theorem 3.2](#)) can be derived from it.

Lemma 3.4 (Approximation for cost-regularization). *Let [Assumption 1](#) hold and take $\varepsilon, \delta > 0$. Then, the following bound hold*

$$G_{\frac{\rho}{1+\delta}}^{0,0}(f) - \frac{\varepsilon\rho}{1+\delta} \leq G_\rho^{\varepsilon, \delta}(f) \leq G_\rho^{0,0}(f).$$

Proof. Since the cost function is non-negative, we directly have $G_\rho^{\varepsilon,\delta}(f) \leq G_\rho^{0,0}(f)$. From [Corollary 2.2](#), we write the cost-regularized function $G_\rho^{\varepsilon,\delta}(f)$ of (18) as follows

$$\begin{aligned}
G_\rho^{\varepsilon,\delta}(f) &= \inf_{\lambda \geq 0} \lambda \rho + \mathbb{E}_{\xi \sim \mathbb{P}} \left[\sup_{\zeta \in \Xi} f(\zeta) - (\varepsilon + (1 + \delta)\lambda)c(\xi, \zeta) \right] \\
&= \inf_{\lambda' \geq \varepsilon} \frac{\lambda' - \varepsilon}{1 + \delta} \rho + \mathbb{E}_{\xi \sim \mathbb{P}} \left[\sup_{\zeta \in \Xi} f(\zeta) - \lambda' c(\xi, \zeta) \right] \quad (\text{with } \lambda' = \varepsilon + (1 + \delta)\lambda) \\
&\geq \inf_{\lambda' \geq 0} \frac{\lambda' - \varepsilon}{1 + \delta} \rho + \mathbb{E}_{\xi \sim \mathbb{P}} \left[\sup_{\zeta \in \Xi} f(\zeta) - \lambda' c(\xi, \zeta) \right] \\
&\geq \frac{-\varepsilon \rho}{1 + \delta} + \inf_{\lambda' \geq 0} \lambda' \frac{\rho}{1 + \delta} + \mathbb{E}_{\xi \sim \mathbb{P}} \left[\sup_{\zeta \in \Xi} f(\zeta) - \lambda' c(\xi, \zeta) \right] \\
&= \sup_{\pi \in \mathcal{P}_{\mathbb{P}}(\Xi \times \Xi) : \mathbb{E}_\pi c \leq \frac{\rho}{1 + \delta}} \mathbb{E}_{\pi_2} f - \frac{\varepsilon \rho}{1 + \delta} = G_{\frac{\rho}{1 + \delta}}^{0,0}(f) - \frac{\varepsilon \rho}{1 + \delta},
\end{aligned}$$

where the last line follows again from [Corollary 2.2](#) with $\varepsilon = \delta = 0$. \blacksquare

Lemma 3.5 (Upper-bound on dual solutions). *Under the assumptions of [Theorem 3.1](#), the optimal solution λ^* of the dual problem of $F_\rho^{\varepsilon,\delta}(f)$ is bounded as follows*

$$\lambda^* \leq \bar{\lambda} = \frac{2 \sup_{\Xi} |f|}{\rho - \mathbb{E}_{\pi_0} c}. \quad (19)$$

Proof. [Theorem 3.1](#) gives the existence of λ^* , which, by definition, minimizes

$$g: \lambda \mapsto \lambda \rho + (\varepsilon + \lambda \delta) \mathbb{E}_{\xi \sim \mathbb{P}} \log \left(\mathbb{E}_{\zeta \sim \pi_0(\cdot|\xi)} e^{\frac{f(\zeta) - \lambda c(\xi, \zeta)}{\varepsilon + \lambda \delta}} \right).$$

On the one hand, $g(\lambda^*)$ is upper bounded as

$$g(\lambda^*) \leq g(0) = \varepsilon \mathbb{E}_{\xi \sim \mathbb{P}} \log \left(\mathbb{E}_{\zeta \sim \pi_0(\cdot|\xi)} e^{\frac{f(\zeta)}{\varepsilon}} \right) \leq \sup_{\Xi} |f|.$$

On the other hand, thanks to Jensen's inequality, $g(\lambda^*)$ is lower-bounded as

$$g(\lambda^*) \geq \lambda^* \rho + \mathbb{E}_{(\xi, \zeta) \sim \pi_0} [f(\zeta) - \lambda^* c(\xi, \zeta)] \geq \lambda^* (\rho - \mathbb{E}_{\pi_0} c) - \sup_{\Xi} |f|.$$

Combining the two inequalities gives (19). \blacksquare

Lemma 3.6 (Approximation bound for the Lagrangians). *Under the assumptions of [Theorem 3.3](#), consider*

$$\begin{aligned}
F_\rho^{\varepsilon,\delta}(\lambda, f) &= \sup_{\pi \in \mathcal{P}_{\mathbb{P}}(\Xi \times \Xi)} \mathbb{E}_{(\xi, \zeta) \sim \pi} [f(\zeta) - \lambda c(\xi, \zeta)] - (\varepsilon + \lambda \delta) \text{KL}(\pi | \pi_0), \\
\text{and } G_{\frac{\varepsilon}{\sigma}, \frac{\delta}{\sigma}}(\lambda, f) &= \sup_{\pi \in \mathcal{P}_{\mathbb{P}}(\Xi \times \Xi)} \mathbb{E}_{(\xi, \zeta) \sim \pi} [f(\zeta) - \lambda c(\xi, \zeta)] - \left(\frac{\varepsilon + \lambda \delta}{\sigma} \right) c(\xi, \zeta).
\end{aligned}$$

Then we have, for a fixed $\Delta \in (0, d]$ and with $I_\sigma(\xi) := \int_{\Xi} e^{-\frac{c(\xi, \zeta)}{2^{p-1}\sigma}} d\zeta$,

$$G_{\frac{\varepsilon}{\sigma}, \frac{\delta}{\sigma}}(\lambda, f) \leq F_\rho^{\varepsilon,\delta}(\lambda, f) + (\text{Lip}(f) + \lambda \text{Lip}(c))\Delta + (\varepsilon + \lambda \delta) \left(\frac{\Delta^p}{\sigma} - \log(V\Delta^d) + \mathbb{E}_{\xi \sim \mathbb{P}} \log I_\sigma(\xi) \right).$$

Proof. We start by reformulating $G_{\rho}^{\frac{\varepsilon}{\sigma}, \frac{\delta}{\sigma}}(\lambda, f)$. By continuity and compactness, the function $\zeta \mapsto f(\zeta) - (\lambda + (\varepsilon + \lambda\delta)/\sigma)c(\xi, \zeta)$ has a maximizer on Ξ for every ξ . By (Rockafellar and Wets, 1998, Thm. 14.37), we get that there exists a measurable map $\zeta^*: \Xi \rightarrow \Xi$ such that $\zeta^*(\xi) \in \arg \max_{\zeta \in \Xi} f(\zeta) - (\lambda + (\varepsilon + \lambda\delta)/\sigma)c(\xi, \zeta)$ for any $\xi \in \Xi$. Then,

$$\pi^*(d\xi, d\zeta) := P(d\xi) \delta_{\zeta^*(\xi)}(d\zeta)$$

is an optimal solution, and therefore

$$G_{\rho}^{\frac{\varepsilon}{\sigma}, \frac{\delta}{\sigma}}(\lambda, f) = \mathbb{E}_{\pi^*} F_{\lambda} - \frac{\varepsilon + \lambda\delta}{\sigma} \mathbb{E}_{\pi^*} c. \quad (20)$$

Now define $\mathbb{B}^{\Delta}(\xi) := \mathbb{B}(\zeta^*(\xi), \Delta)$ and $\pi^{\Delta} \in \mathcal{P}_{\mathbb{P}}(\Xi \times \Xi)$ such that

$$\pi^{\Delta} \propto \mathbf{1}_{\zeta \in \mathbb{B}^{\Delta}(\xi)} \pi_0(d\xi, d\zeta).$$

Note first that we have

$$\mathbb{E}_{\pi^*} F_{\lambda} - \mathbb{E}_{\pi^{\Delta}} F_{\lambda} = \mathbb{E}_{\xi \sim P} \mathbb{E}_{\zeta \sim \pi^{\Delta}(\cdot|\xi)} [F_{\lambda}(\xi, \zeta^*(\xi)) - F_{\lambda}(\xi, \zeta)] \leq (\text{Lip}(f) + \lambda \text{Lip}(c))\Delta, \quad (21)$$

since F_{λ} is $(\text{Lip}(f) + \lambda \text{Lip}(c))$ -Lipschitz continuous and the support of $\pi^{\Delta}(\cdot|\xi)$ is $\mathbb{B}^{\Delta}(\xi)$. Now, we proceed to bound $\text{KL}(\pi^{\Delta} | \pi_0)$ by first noticing that

$$\begin{aligned} \text{KL}(\pi^{\Delta} | \pi_0) &= \mathbb{E}_{\xi \sim P} \mathbb{E}_{\zeta \sim \pi^{\Delta}(\cdot|\xi)} \log \left(\frac{d\pi^{\Delta}(\zeta|\xi)}{d\pi_0(\zeta|\xi)} \right) \\ &= -\mathbb{E}_{\xi \sim P} \log \left(\int_{\Xi \cap \mathbb{B}^{\Delta}(\xi)} e^{-\frac{c(\xi, \zeta)}{2^{p-1}\sigma}} d\zeta \right) + \mathbb{E}_{\xi \sim P} \log \left(\int_{\Xi} e^{-\frac{c(\xi, \zeta)}{2^{p-1}\sigma}} d\zeta \right) \end{aligned} \quad (22)$$

We focus on lower-bounding $\int_{\Xi \cap \mathbb{B}^{\Delta}(\xi)} e^{-\frac{c(\xi, \zeta)}{2^{p-1}\sigma}} d\zeta$. First, note that by the triangular inequality (and since $p \geq 1$), for any $\xi, \zeta \in \Xi$,

$$\frac{c(\xi, \zeta)}{2^p} \leq \frac{c(\xi, \zeta^*(\xi)) + c(\zeta^*(\xi), \zeta)}{2}.$$

If, in addition, ζ is in $\mathbb{B}^{\Delta}(\xi)$, this bound becomes $\frac{c(\xi, \zeta)}{2^{p-1}} \leq c(\xi, \zeta^*(\xi)) + \Delta^p$. Hence, we have

$$\int_{\Xi \cap \mathbb{B}^{\Delta}(\xi)} e^{-\frac{c(\xi, \zeta)}{2^{p-1}\sigma}} d\zeta \geq e^{-\frac{c(\xi, \zeta^*(\xi)) + \Delta^p}{\sigma}} \text{vol}(\Xi \cap \mathbb{B}^{\Delta}(\xi)) \geq e^{-\frac{c(\xi, \zeta^*(\xi)) + \Delta^p}{\sigma}} V \Delta^d,$$

where V is defined in (17). Plugging the above lower and upper bounds into (22) yields

$$\begin{aligned} \text{KL}(\pi^{\Delta} | \pi_0) &\leq \frac{\mathbb{E}_{\xi \sim P} c(\xi, \zeta^*(\xi)) + \Delta^p}{\sigma} - \log(V \Delta^d) + \mathbb{E}_{\xi \sim P} \log I_{\sigma}(\xi) \\ &= \frac{\mathbb{E}_{\pi^*} c + \Delta^p}{\sigma} - \log(V \Delta^d) + \mathbb{E}_{\xi \sim P} \log I_{\sigma}(\xi). \end{aligned} \quad (23)$$

Finally, putting (20), (21), and (23) together gives,

$$\begin{aligned}
G_{\rho^{\frac{\varepsilon}{\sigma}, \frac{\delta}{\sigma}}}(\lambda, f) &= \mathbb{E}_{\pi^*} F_{\lambda} - \frac{\varepsilon + \lambda\delta}{\sigma} \mathbb{E}_{\pi^*} c \\
&= \mathbb{E}_{\pi^{\Delta}} F_{\lambda} - (\varepsilon + \lambda\delta) \text{KL}(\pi^{\Delta} | \pi_0) + (\mathbb{E}_{\pi^*} F_{\lambda} - \mathbb{E}_{\pi^{\Delta}} F_{\lambda}) + \left((\varepsilon + \lambda\delta) \text{KL}(\pi^{\Delta} | \pi_0) - \frac{\varepsilon + \lambda\delta}{\sigma} \mathbb{E}_{\pi^*} c \right) \\
&\leq F_{\rho^{\varepsilon, \delta}}(\lambda, f) + (\text{Lip}(f) + \lambda \text{Lip}(c))\Delta + (\varepsilon + \lambda\delta) \left(\frac{\Delta^p}{\sigma} - \log(V\Delta^d) + \mathbb{E}_{\xi \sim P} \log I_{\sigma}(\xi) \right)
\end{aligned}$$

which is the claimed inequality. \blacksquare

We have now all the ingredients to establish the extended version of the approximation result.

Proof of Theorem 3.3. First, notice that by Lemma 3.4, we have that

$$F_{\frac{\rho}{1+\delta/\sigma}}^{0,0}(f) - \frac{\varepsilon\rho}{\sigma + \delta} = G_{\frac{\rho}{1+\delta/\sigma}}^{0,0}(f) - \frac{\varepsilon\rho}{\sigma + \delta} \leq G_{\rho^{\frac{\varepsilon}{\sigma}, \frac{\delta}{\sigma}}}(f).$$

Thus, using the bound at λ fixed, given by Lemma 3.6 and the upper-bound (19), we get

$$\begin{aligned}
G_{\rho^{\frac{\varepsilon}{\sigma}, \frac{\delta}{\sigma}}}(f) &\leq \inf_{0 \leq \lambda \leq \bar{\lambda}} G_{\rho^{\frac{\varepsilon}{\sigma}, \frac{\delta}{\sigma}}}(\lambda, f) \\
&\leq \inf_{0 \leq \lambda \leq \bar{\lambda}} F_{\rho^{\varepsilon, \delta}}(\lambda, f) + (\text{Lip}(f) + \lambda \text{Lip}(c))\Delta + (\varepsilon + \lambda\delta) \left(\frac{\Delta^p}{\sigma} - \log(V\Delta^d) + \mathbb{E}_{\xi \sim P} \log I_{\sigma}(\xi) \right) \\
&\leq F_{\rho^{\varepsilon, \delta}}(f) + (\text{Lip}(f) + \bar{\lambda} \text{Lip}(c)) \Delta + (\varepsilon + \bar{\lambda}\delta) \left(\frac{\Delta^p}{\sigma} - \log(V\Delta^d) + \mathbb{E}_{\xi \sim P} \log I_{\sigma}(\xi) \right).
\end{aligned}$$

Minimizing the above bound over $\Delta > 0$, we get that the optimal value Δ^* is of the form $\Delta^* = \frac{(\varepsilon + \bar{\lambda}\delta)d}{L} + \mathcal{O}((\varepsilon + \bar{\lambda}\delta)^{p+1})$, with $L = \text{Lip}(f) + \bar{\lambda} \text{Lip}(c)$, as introduced in the statement of the theorem.

As a consequence, we set $\Delta = \frac{(\varepsilon + \bar{\lambda}\delta)d}{L}$ in the bound above, which becomes

$$G_{\rho^{\frac{\varepsilon}{\sigma}, \frac{\delta}{\sigma}}}(f) \leq F_{\rho^{\varepsilon, \delta}}(f) + (\varepsilon + \bar{\lambda}\delta) \left(d + d \log \left(\frac{L}{(\varepsilon + \bar{\lambda}\delta)d} \right) + \mathbb{E}_{\xi \sim P} \log I_{\sigma}(\xi) + \log \frac{1}{V} + \frac{1}{\sigma} \left(\frac{(\varepsilon + \bar{\lambda}\delta)d}{L} \right)^p \right).$$

There is only left to bound the term in $\mathbb{E}_{\xi \sim P} \log I_{\sigma}(\xi)$. On one hand, we have $e^{-\frac{c(\xi, \zeta)}{2^{p-1}\sigma}} \leq 1$ so that

$$\int_{\Xi} e^{-\frac{c(\xi, \zeta)}{2^{p-1}\sigma}} d\zeta \leq \text{vol}(\Xi).$$

On the other hand, we also have (using the change of variable $\zeta' = \sigma^{-\frac{1}{p}}(\zeta - \xi)$ to get I_{σ})

$$\int_{\Xi} e^{-\frac{c(\xi, \zeta)}{2^{p-1}\sigma}} d\zeta \leq \int_{\mathbb{R}^d} e^{-\frac{\|\xi - \zeta\|^p}{2^{p-1}\sigma}} d\zeta = I_{\sigma}$$

This makes the constant C appear in the bound and thus ends the proof. \blacksquare

We finish by explaining how the main theorem, Theorem 3.2, stems from Theorem 3.3. On the left-hand side (LHS) of the inequality in Theorem 3.3, the unregularized objective has radius $\frac{\rho}{1+\delta/\sigma}$, instead of simply ρ in $F_{\rho}^{0,0}(f) = \text{val}(\text{WDRO})$. Thus, we compare in the next lemma the optimal values for these two parameters.

Lemma 3.7 (Comparing optimal values). *Under the assumptions of [Theorem 3.3](#),*

$$F_\rho^{0,0}(f) \leq F_{\frac{\rho}{1+\delta/\sigma}}^{0,0}(f) + \mathcal{O}(\delta).$$

Proof. We first apply ([Santambrogio, 2015](#), Theorem 5.27) to get a constant-speed geodesic for the p -Wasserstein distance connecting P and Q , which is $W_c(P, Q)^{\frac{1}{p}}$ with our notation. This means that there exists a family of probability distributions $(Q_t)_{t \in [0,1]}$ such that $Q_0 = Q$, $Q_1 = P$ and, for any $t \in [0, 1]$,

$$W_c(P, Q_t)^{\frac{1}{p}} = (1-t)W_c(P, Q)^{\frac{1}{p}} \quad \text{and} \quad W_c(Q_t, Q)^{\frac{1}{p}} = tW_c(P, Q)^{\frac{1}{p}}.$$

We apply these equations with Q such that $W_c(P, Q) \leq \rho$ and $t = 1 - (1 + \delta/\sigma)^{-\frac{1}{p}}$ to obtain

$$W_c(P, Q_t) \leq \frac{\rho}{1 + \delta/\sigma} \quad \text{and} \quad W_c(Q, Q_t) \leq t^p \rho = \mathcal{O}(\delta).$$

Note that the first inequality above yields

$$\mathbb{E}_{Q_t} f \leq \sup_{\pi \in \mathcal{P}_P(\Xi \times \Xi) : \mathbb{E} \pi c \leq \frac{\rho}{1+\delta/\sigma}} \mathbb{E}_{\pi_2} f = F_{\frac{\rho}{1+\delta/\sigma}}^{0,0}(f), \quad (24)$$

We now use the Kantorovich-Rubinstein inequality (e.g., [Villani \(2003, Thm. 1.14\)](#)) to write

$$\mathbb{E}_Q f - \mathbb{E}_{Q_t} f \leq W_1(Q, Q_t) \text{Lip}(f) \leq W_c(Q, Q_t)^{\frac{1}{p}} \text{Lip}(f),$$

where for the second inequality we used that the p -Wasserstein distance is always greater than or equal to the 1-Wasserstein distance (e.g., [Villani \(2003, §7.1.2\)](#)). Together with (24), this yields:

$$\mathbb{E}_Q f \leq \mathbb{E}_{Q_t} f + \mathcal{O}(\delta) \leq F_{\frac{\rho}{1+\delta/\sigma}}^{0,0}(f) + \mathcal{O}(\delta).$$

Taking the supremum over all Q such that $W_c(P, Q) \leq \rho$ allows us to conclude. ■

With the help of the previous lemma, the proof of [Theorem 3.2](#) comes easily from [Theorem 3.3](#).

Proof of [Theorem 3.2](#). We start with checking that the fact that Ξ is a compact convex body (condition (iv) in [Theorem 3.2](#)) implies that $V > 0$ (condition (iv) in [Theorem 3.3](#)). Introduce, for $\xi \in \Xi$, the function

$$\nu_\xi : \Delta \mapsto \frac{\text{vol}(\Xi \cap \mathbb{B}(\xi, \Delta))}{\Delta^d} = \text{vol}\left(\frac{1}{\Delta}(\Xi - \xi) \cap \mathbb{B}(0, 1)\right).$$

Since Ξ is convex, we easily get that ν_ξ is non-increasing. Thus we can lower-bound the constant V as follows, using $\text{diam}(\Xi) := \sup_{\xi, \zeta \in \Xi} \|\xi - \zeta\|$ the diameter of Ξ .

$$\begin{aligned} V &= \inf_{\xi \in \Xi, 0 < \Delta \leq d} \nu_\xi(\Delta) \geq \inf_{\xi \in \Xi, 0 < \Delta \leq \max(d, \text{diam}(\Xi))} \nu_\xi(\Delta) \\ &= \inf_{\xi \in \Xi} \frac{\text{vol}(\Xi \cap \mathbb{B}(\xi, \max(d, \text{diam}(\Xi))))}{(\max(d, \text{diam}(\Xi)))^d} \geq \frac{\text{vol}(\Xi)}{(\max(d, \text{diam}(\Xi)))^d} > 0. \end{aligned}$$

So we get $V > 0$ which is condition (iv) in [Theorem 3.3](#). Thus we can apply [Theorem 3.3](#) and [Lemma 3.7](#). Noting that $F_\rho^{\varepsilon, \delta}(f) = \text{val}(\text{E-WDRO})$ and $F_\rho^{0,0}(f) = \text{val}(\text{WDRO})$, and combining the obtained bound with [Lemma 3.7](#) gives the result. ■

4 Conclusion, perspectives

Inspired by the success of regularization in OT, we proposed and studied a regularization scheme for WDRO problems. We derived the expression of the dual objective function in the general case as well as a refined one in the particular setting of the entropic regularization. In addition, we showed that the difference between the original WDRO problem and the entropic one is properly controlled by the regularization parameters.

Finally, since regularization in OT has shown attractive computational advantages and statistical benefits, an interesting direction would be to investigate whether similar gains hold for regularization in WDRO.

Acknowledgments

This work has been supported by MIAI Grenoble Alpes (ANR-19-P3IA-0003).

References

- Yang An and Rui Gao. Generalization bounds for (wasserstein) robust optimization. *Advances in Neural Information Processing Systems*, 34, 2021.
- Jose Blanchet and Yang Kang. Semi-supervised learning based on distributionally robust optimization. *Data Analysis and Applications 3: Computational, Classification, Financial, Statistical and Stochastic Methods*, 5:1–33, 2020.
- Jose Blanchet and Karthyek Murthy. Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research*, 44(2):565–600, 2019.
- Jose Blanchet, Yang Kang, and Karthyek Murthy. Robust wasserstein profile inference and applications to machine learning. *Journal of Applied Probability*, 56(3):830–857, 2019.
- Radu Ioan Bot, Sorin-Mihai Grad, and Gert Wanka. *Duality in Vector Optimization*. Vector Optimization. Springer Berlin Heidelberg, 2009.
- Stéphane Boucheron, Olivier Bousquet, Gábor Lugosi, and Pascal Massart. Moment inequalities for functions of independent random variables. *The Annals of Probability*, 33(2):514 – 560, 2005.
- Haim Brezis. *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Springer New York, 2010.
- Guillaume Carlier, Vincent Duval, Gabriel Peyré, and Bernhard Schmitzer. Convergence of entropic schemes for optimal transport and gradient flows. *SIAM Journal on Mathematical Analysis*, 49(2):1385–1418, 2017.
- Ruidi Chen and Ioannis Ch Paschalidis. Distributionally Robust Learning. *Found. Trends Optim.*, 4:1–243, 2020.
- Christian Clason and Tuomo Valkonen. Introduction to nonsmooth analysis and optimization. *arXiv preprint arXiv:2001.00216*, 2020.

- Marco Cuturi. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
- Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1):115–166, 2018.
- Jean Feydy, Thibault Séjourné, François-Xavier Vialard, Shun-ichi Amari, Alain Trounev, and Gabriel Peyré. Interpolating between optimal transport and mmd using sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 2019.
- Rui Gao and Anton J Kleywegt. Distributionally robust stochastic optimization with wasserstein distance. *arXiv preprint arXiv:1604.02199*, 2016.
- Aude Genevay, Marco Cuturi, Gabriel Peyré, and Francis Bach. Stochastic Optimization for Large-scale Optimal Transport. In *Thirtieth Annual Conference on Neural Information Processing System*, NIPS, 2016.
- Aude Genevay, Lénaïc Chizat, Francis Bach, Marco Cuturi, and Gabriel Peyré. Sample complexity of sinkhorn divergences. In *22nd International Conference on Artificial Intelligence and Statistics*, pages 1574–1583. PMLR, 2019.
- Daniel Kuhn, Peyman Mohajerin Esfahani, Viet Anh Nguyen, and Soroosh Shafieezadeh-Abadeh. Wasserstein distributionally robust optimization: Theory and applications in machine learning. In *Operations Research & Management Science in the Age of Analytics*. INFORMS, 2019.
- Se Yoon Lee. Gibbs sampler and coordinate ascent variational inference: A set-theoretical review. *Communications in Statistics - Theory and Methods*, 0(0):1–21, 2021.
- Quentin Merigot and Boris Thibert. Optimal transport: discretization and algorithms. In *Handbook of Numerical Analysis*, volume 22, pages 133–212. Elsevier, 2021.
- François-Pierre Paty and Marco Cuturi. Regularized optimal transport is ground cost adversarial. In *International Conference on Machine Learning*, pages 7532–7542. PMLR, 2020.
- Juan Peypouquet. *Convex Optimization in Normed Spaces: Theory, Methods and Examples*. SpringerBriefs in Optimization. Springer International Publishing, 2015.
- 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.
- R. Tyrrell Rockafellar and Roger J.-B. Wets. *Variational Analysis*. Grundlehren Der Mathematischen Wissenschaften. Springer-Verlag, 1998.
- Walter Rudin. *Real and Complex Analysis*. McGraw-Hill, 1987.
- Filippo Santambrogio. *Optimal Transport for Applied Mathematicians*, volume 87 of *Progress in Nonlinear Differential Equations and Their Applications*. Springer International Publishing, 2015.
- Soroosh Shafieezadeh-Abadeh, Daniel Kuhn, and Peyman Mohajerin Esfahani. Regularization via Mass Transportation. *Journal of Machine Learning Research*, 20:1–68, 2019.

Cédric Villani. *Topics in Optimal Transportation*. American Mathematical Soc., 2003.

Jie Wang, Rui Gao, and Yao Xie. Sinkhorn distributionally robust optimization. *arXiv preprint arXiv:2109.11926*, 2021.

Yaodong Yu, Tianyi Lin, Eric Mazumdar, and Michael I. Jordan. Fast Distributionally Robust Learning with Variance Reduced Min-Max Optimization. *arXiv:2104.13326 [cs, math, stat]*, 2021.