
Analysis of Kernel Mirror Prox for Measure Optimization

Pavel Dvurechensky

Jia-Jie Zhu

{dvureche, zhu}@wias-berlin.de

Weierstrass Institute for Applied Analysis and Stochastics
Berlin, Germany

Abstract

By choosing a suitable function space as the dual to the non-negative measure cone, we study in a unified framework a class of functional saddle-point optimization problems, which we term the Mixed Functional Nash Equilibrium (MFNE), that underlies several existing machine learning algorithms, such as implicit generative models, distributionally robust optimization (DRO), and Wasserstein barycenters. We model the saddle-point optimization dynamics as an interacting Fisher-Rao-RKHS gradient flow when the function space is chosen as a reproducing kernel Hilbert space (RKHS). As a discrete time counterpart, we propose a primal-dual kernel mirror prox (KMP) algorithm, which uses a dual step in the RKHS, and a primal entropic mirror prox step. We then provide a unified convergence analysis of KMP in an infinite-dimensional setting for this class of MFNE problems, which establishes a convergence rate of $O(1/N)$ in the deterministic case and $O(1/\sqrt{N})$ in the stochastic case, where N is the iteration counter. As a case study, we apply our analysis to DRO, providing algorithmic guarantees for DRO robustness and convergence.

1 INTRODUCTION

Modern analysis of machine learning algorithms in high dimensions leverages the view of optimization over probability measures. For example, some state-of-the-art analysis of generative adversarial networks and

adversarial robustness formulates the training process as finding a mixed Nash Equilibrium (MNE) in a two-player zero-sum game (Hsieh et al., 2019; Domingo-Enrich et al., 2020; Wang & Chizat, 2022; Trillos & Trillos, 2023; Kim et al., 2024). The MNE problem seeks the solution, i.e., MNE, of the saddle-point optimization problem

$$\inf_{\mu \in \mathcal{M}} \sup_{\nu \in \mathcal{M}} G(\mu, \nu), \quad (1)$$

where $G(\mu, \nu)$ is a bi-variate objective function(al) of the probability measures μ, ν , and \mathcal{M} is the space of probability measures on a convex compact domain, e.g., a bounded closed convex subset of \mathbb{R}^d . Related to (1), we are also interested in treating (constrained) measure optimization problems of the form

$$\inf_{\mu \in \mathcal{B} \subset \mathcal{M}} \mathcal{E}(\mu). \quad (2)$$

In machine learning, there exists a commonly-used alternative paradigm to directly optimizing the measures, such as in the MNE problem (1) or in (2). This paradigm exploits the duality between probability measure space \mathcal{M} and a dual functional space, i.e., instead of optimizing w.r.t. a measure, we solve the (constrained) *functional optimization problem*

$$\inf_{f \in \mathcal{F}} \mathcal{V}(f), \quad (3)$$

where \mathcal{F} is some set of dual functions. Such methods often leverage scalable learning models such as reproducing kernel Hilbert spaces (RKHS) and deep neural networks for parameterizing and manipulating the dual function f , instead of directly searching in the space of probability distributions as in (1) or (2). We have witnessed success of such dual functional approaches in many domains of machine learning research, e.g., generative modeling (Nowozin et al.; Arjovsky et al., 2017; Gulrajani et al., 2017; Korotin et al., 2019) computing optimal transport (Genevay et al., 2016), (see also (Dvurechensky et al., 2018b;

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

Jambulapati et al., 2019; Guminov et al., 2021) for discrete transport setting), distributionally robust optimization (Zhu et al., 2021), Wasserstein barycenters (Li et al., 2020; Tiapkin et al., 2022; Korotin et al., 2021; Dvinskikh & Tiapkin, 2021). However, optimization over those functional spaces, e.g., RKHSs or neural networks, is inherently an infinite-dimensional problem that requires careful treatment. In particular, algorithms and their convergence analysis do not directly follow from the finite-dimensional setting.

In certain situations, the dual problem (3) may turn out to be non-smooth and from the algorithmic perspective, it is better to keep both functional and measure variables and consider a primal-dual problem formulation. This drives us to our main problem of interest in this paper. Namely, we formulate the *Mixed Functional Nash Equilibrium* (MFNE), whose outer optimization is an optimization problem over a set \mathcal{F} of functions

$$\inf_{f \in \mathcal{F}} \sup_{\mu \in \mathcal{M}} F(f, \mu), \quad (4)$$

which, like the MNE, is a special case of the infinite-dimensional pure NE problem. MFNE (4) has appeared in several cutting-edge algorithms in ML research, which we discuss in detail in Section 2.1. In these examples, algorithmic treatment is often made on a case-by-case basis and a unified convergence analysis for those applications is still missing. In this paper, we make a first step towards a unified treatment of such applications by considering a subclass of MFNE problems where \mathcal{F} is a subset of RKHS, which we call *Mixed Kernel Nash Equilibrium* (MKNE). A prominent example of such problems, as we show, is the distributionally robust optimization problem with the ambiguity set given by the kernel maximum mean discrepancy (MMD) distance.

We summarize our contributions as the following.

1. We model the infinite-dimensional continuous-time optimization dynamics of the general functional optimization problem (3) as RKHS gradient flows, which enjoys simple convexity structure. Then, motivated by an interacting gradient flow (19)-(20) that couples the Fisher-Rao and RKHS gradient flow, we derive a primal-dual kernel mirror prox algorithm for solving MFNE (4) with \mathcal{F} being a subset of Hilbert space.
2. We show that our specialized infinite-dimensional mirror prox algorithm, which we term primal-dual kernel mirror prox (KMP), can be applied to the MKNE problem to obtain a convergence rate $\mathcal{O}(\frac{1}{N})$ in the deterministic and $\mathcal{O}(\frac{1}{\sqrt{N}})$ in the stochastic settings, where N is the iteration counter. To the best of our knowledge, it is the first analysis with the kernel mirror prox steps in

the dual functional space, which differs from the typical mean-field analysis of measure optimization.

3. Our framework can be applied to existing learning tasks to derive convergence results. As a concrete case study, we apply our analysis to distributionally robust optimization (DRO) to establish a convergence rate via primal-dual kernel mirror prox in the case of MMD ambiguity sets. Importantly, our method does not require the loss to be convex w.r.t. the uncertain variable, i.e., $l(\theta; x)$ can be non-convex non-linear in x in (10), which is much more general than common assumptions in the Wasserstein DRO literature. We then provide robustness guarantees for the KMP solution both in the population risk and under distribution shift.

2 PRELIMINARIES

Notation. We use $\mathcal{M} = \text{Prob}(\bar{\Omega})$ to denote the space of probability measures defined on the closure of a bounded convex domain $\Omega \subset \mathbb{R}^d$. We say that (4) is convex-concave if the inner problem is concave maximization and the outer problem is convex minimization. The convexity notion, if not otherwise specified, refers to the regular convexity notion (defined in (43)). For PDE gradient flows, the states are functions of both time and space, for example, $u(t, x)$. When there is no ambiguity, we write $u(t) := u(t, \cdot)$ to denote the function at evolutionary time t . If not otherwise specified, proofs are deferred to the appendix.

2.1 Duality of Metrics on Probability Measures

One reason behind the ubiquity of the functional optimization problem (3) and MFNE (4) in machine learning problems is that common probability metrics admit a dual characterization. For example, the *optimal transport* distance (Santambrogio, 2015; Ambrosio et al., 2008), e.g., p -Wasserstein metric W_p with $p \geq 1$, can be characterized in the dual space via the dual Kantorovich problem

$$W_p^p(\mu, \nu) = \sup_{\psi_1, \psi_2} \int \psi_1 d\mu + \int \psi_2 d\nu \quad (5)$$

$$\text{s.t. } \psi_1(x) + \psi_2(y) \leq c(x, y), \quad (6)$$

$\forall x, y$, a.e. which is an infinite-dimensional optimization problem with an infinite constraint. Bounded continuous functions ψ_i are referred to as the Kantorovich potential functions. $c(x, y) = \|x - y\|_p^p$ is the transport cost function associated with the transport. In the machine learning literature, researchers have explored this dual formulation by directly parameterizing the Kantorovich potential ψ_i 's, e.g., using RKHS

functions (Genevay et al., 2016; Vacher et al., 2021) or input convex neural networks (ICNNs) (Li et al., 2020; Korotin et al., 2019).

Another commonly used metric is the integral probability metric (IPM), which is defined via the weak norm formulation

$$\text{IPM}(\mu, \nu) = \sup_{f \in \mathcal{F}} \int f d(\mu - \nu). \quad (7)$$

One particular choice of the test function family is the RKHS-norm-ball $\mathcal{F} = \{f : \|f\|_{\mathcal{H}} \leq 1\}$ (here \mathcal{H} is a RKHS), which yields the kernel maximum mean discrepancy (MMD) (Gretton et al., 2012). The choice $\mathcal{F} = \{f : \text{Lip}(f) \leq 1\}$ (i.e., class of 1-Lipschitz continuous functions) recovers the type-1 (Kantorovich) Wasserstein metric.

Such basic duality characterizes the relationship of the underlying geometry of the learning problem and the variational problem that optimizes w.r.t. functions, such as the IPM test functions and the Kantorovich potentials in OT. This has been exploited in several fields in machine learning.

Implicit generative models (IGM). We consider the following IGM formulation (Li et al., 2015; Dziugaite et al.; Goodfellow et al., 2014),

$$\inf_{G_\theta} \mathbb{E}_Z \mathcal{D}(P, G_\theta(Z)), \quad (8)$$

where G_θ is the generator function parameterized by $\theta \in \mathbb{R}^d$, P is a target distribution, Z is a random variable with (simple) latent distribution Q , \mathcal{D} can be chosen as a discrepancy measure between distributions such as the f -divergence family, optimal transport distance, or kernel MMD. In IGM, P is often taken to be the training data distribution. Recent theoretical analysis of the optimization for training (8) *lifts* the non-convex optimization problem to the space of probability measures, which is an instantiation of a MNE problem (1). The problem can then be cast into this paper’s general MFNE formulation (4), for example, by choosing the function \mathcal{D} from the IPM family (7)

$$\inf_{\mu \in \mathcal{M}} \sup_{f \in \mathcal{F}} \left\{ \int f(x) dP(x) - \mathbb{E}_{\theta \sim \mu} \int f(g_\theta(z)) dQ(z) \right\}, \quad (9)$$

where g_θ is the generator. If the function class \mathcal{F} is chosen to be the class of 1-Lipschitz functions, then the formulation is the Wasserstein GAN (Arjovsky et al., 2017; Gulrajani et al., 2017). On the other hand, if \mathcal{F} is the RKHS-norm ball, this is the MMD GAN (Li et al., 2017; Bińkowski et al., 2018). Note that this lifted problem is now *convex-concave* in the optimization variables.

Distributionally robust optimization (DRO).

We now consider the metric-ball-constrained DRO problem (Delage & Ye, 2010; Ben-Tal et al., 2013; Gao & Kleywegt, 2016; Zhao & Guan, 2018)

$$\inf_{\theta \in \Theta} \sup_{\mu: \mathcal{D}(\mu, \hat{P}_n) \leq \epsilon} \mathbb{E}_\mu[l(\theta; x)], \quad (10)$$

where the set of decision variables Θ is a closed convex set, e.g., $\Theta \subseteq \mathbb{R}^d$, \mathcal{D} is a discrepancy measure between distributions, \hat{P}_n is the empirical data distribution, $l(\theta; x)$ is the loss function given parameter θ and data point x . Below, we consider DRO in the two settings of probability metric \mathcal{D} : the p -Wasserstein DRO ($p \geq 1$) (Gao & Kleywegt, 2016) as well as the Kernel(-MMD) DRO (Zhu et al., 2021). Our scope is crucially broader than the common Wasserstein DRO reformulations using conic linear duality, which often omit the discussion of *non-linear-in- x losses* common in machine learning. We propose the following primal-dual reformulation of DRO.

Lemma 2.1 (Primal-dual reformulation of Wasserstein and Kernel DRO). *Suppose the probability metric is chosen to be the optimal transport metric (e.g., p -Wasserstein distance). Then, the DRO problem (10) admits the reformulation*

$$\inf_{\substack{\gamma > 0, \theta \in \Theta, \\ f \in \Psi_c}} \sup_{\mu \in \mathcal{M}} \mathbb{E}_\mu(l(\theta; x) - \gamma f(x)) - \frac{\gamma}{n} \sum_{i=1}^n f^c(x_i) + \gamma \epsilon. \quad (11)$$

Ψ_c is the set of c -concave functions (Santambrogio, 2015) and $f^c(y) := \inf_x c(x, y) - f(x)$ denotes the c -transform.

Suppose the probability metric \mathcal{D} is the MMD, then the DRO problem (10) is equivalent to the smooth saddle-point problem

$$\inf_{\theta \in \Theta, f \in \mathcal{H}} \sup_{\mu \in \mathcal{M}, h \in \mathcal{H}: \|h\|_{\mathcal{H}} \leq 1} \left(\frac{1}{n} \sum_{i=1}^n f(x_i) + \epsilon \langle h, f \rangle + \mathbb{E}_\mu(l(\theta; x) - f(x)) \right). \quad (12)$$

Note that c -concavity of a function f means that f is the c -transform of a bounded continuous function. Lemma 2.1 shows that those primal-dual DRO formulations have the *convex-concave* structure in μ, f as in the MFNE (4). (Although the convexity w.r.t. μ of (11) is lost along the Wasserstein geodesics (Ambrosio et al., 2008).) They are also convex in the learning model parameter θ when the loss l is, (12) is concave in the smoothing variable h , and (11) is trivially convex in the dual variable γ .

Wasserstein barycenter. Another important application is the Wasserstein barycenter problem (Agueh

& Carlier, 2011; Peyré & Cuturi, 2019; Li et al., 2020; Tiapkin et al., 2022; Kroshnin et al., 2019; Korotin et al., 2021; Dvurechensky et al., 2018a), which can be formulated as a saddle-point optimization problem.

$$\begin{aligned} & \min_{\mu \in \mathcal{M}} \sum_{i=1}^n \alpha_i [W_p^p(\mu, \nu_i)] \\ & = \min_{\mu \in \mathcal{M}} \sum_{i=1}^n \alpha_i \sup_{f_i \in \Psi_c} \left\{ \int f_i^c d\mu + \int f_i d\nu_i \right\}, \end{aligned} \quad (13)$$

where f_i^c again denotes the c -transform. $\nu_i \in \mathcal{M}$ are given probability measures, $f_i \in \Psi_c$ are the c -concave Kantorovich potentials associated with the Wasserstein distance.

As we have seen, several ML problems can be formulated in our framework of the infinite-dimensional MFNE problem (4), which gives a unified optimization view on these applications. An important role in these formulations is played by the dual function set \mathcal{F} . Note that in the aforementioned settings of p -Wasserstein ($p \neq 1$) metric, one may parameterize Kantorovich potential functions using tools such as the random Fourier features (Rahimi & Recht, 2007) and ICNN (Amos et al., 2017), such as done by (Genevay et al., 2016; Makkuva et al., 2020; Korotin et al., 2019). Therefore, as a first step towards the theoretical analysis of the MFNE problem and to simplify this analysis, we focus on the RKHS functions (and hence the kernel MMD in the DRO setting) setting in the rest of the paper. We underline that, even in this particular setting, we are not aware of algorithms and their convergence results.

2.2 Gradient Flow in Hilbert and Metric Spaces

Recent theoretical analysis of generative models via MNE, e.g., (Hsieh et al., 2019; Domingo-Enrich et al., 2020; Wang & Chizat, 2022; Trillos & Trillos, 2023), adopted the mean-field point of view closely related to PDE gradient flows of probability measures (Ambrosio et al., 2008; Otto, 1996; Jordan et al., 1998; Otto, 2001). Notably, works such as (Hsieh et al., 2019) modeled training dynamics of GAN as sampling using Langevin SDE, which is also equivalent to solving the Fokker-Planck PDE.

Intuitively, a gradient flow describes a dynamical system that is driven towards the fastest dissipation of certain energy. This system is called a *gradient system*. For example, the dynamical system described by an ordinary differential equation in the Euclidean space that follows the negative gradient direction of some function f , $\dot{x}(t) = -\nabla f(x(t))$, $x(t) \in \mathbb{R}^d$, is a simple gradient system. Of particular interest to this

paper in the context of measure optimization (2) is the gradient flow in the Fisher-Rao metric space, which is the gradient system that generates the following *reaction equation* as its gradient flow equation

$$\begin{aligned} \partial_t \mu &= -\mu \cdot (\mathcal{E}'(\mu) - \int \mu \cdot \mathcal{E}'(\mu)), \\ \mu(0, x) &= \mu^0(x) \in \mathcal{M}, \end{aligned} \quad (14)$$

where $\mathcal{E}'(\mu)$ is the first variation. Alternatively, one can also view the whole r.h.s. as the Fisher-Rao metric gradient. More generally, gradient flows in metric spaces (Ambrosio et al., 2008) have gained significant attention in recent machine learning literature due to the study of Wasserstein gradient flow (WGF), originated from the seminal works of Otto and colleagues, e.g., (Otto, 1996; Jordan et al., 1998; Otto, 2001). Rigorous characterizations of general metric gradient systems have been carried out in PDE literature, for which we refer to (Ambrosio et al., 2008) for a complete treatment.

In practical optimization, this flow corresponds to the entropic mirror descent step

$$\mu_{k+1} \leftarrow \arg \inf_{\mu \in \mathcal{M}} \int \mathcal{E}'(\mu_k) d(\mu - \mu_k) + \frac{1}{\tau} D_{\text{KL}}(\mu, \mu_k), \quad (15)$$

giving multiplicative update $\mu_{k+1} = \frac{1}{Z} \mu_k \cdot e^{-\tau \mathcal{E}'(\mu_k)}$, where Z is a normalization constant.

3 RKHS GRADIENT FLOW FOR MODELING MEASURE OPTIMIZATION DYNAMICS

Our starting point is to model the infinite-dimensional continuous-time optimization dynamics of the functional optimization problem (3) using the RKHS gradient flows (RKHSGF)

$$\partial_t f = -\mathcal{V}'(f), \quad f(0, x) = f^0(x) \in \mathcal{H}, \quad (16)$$

where $\mathcal{V}'(f)$ is Fréchet (sub)differential in \mathcal{H} . The usual properties of this gradient flow, such as existence and uniqueness are provided in the appendix. In principle, other function spaces in the literature can also be considered in practice, such as the random Fourier feature functions, single hidden-layer neural networks (Bach, 2017), and ICNNs. To make the theoretical analysis tractable, we focus in the rest of the paper on the RKHS setting. A more general setting is left for future work.

RKHS has been used in machine learning applications of gradient flows. For example, in (Arbel et al.,

2019), the squared RKHS norm was used as the driving energy functional for the Wasserstein gradient flow, rather than its use as the dissipation geometry for the flow as in our formulation. A few other works such as (Chu et al., 2020; Chewi et al., 2020; Duncan et al., 2019; Liu & Wang, 2019; Korba et al., 2021) studied the kernelized Wasserstein gradient flow in the context of Stein geometry. Those particular cases do not exploit the simplicity of gradient flow structure in the Hilbert space. We outline some technical background regarding this gradient flow in the RKHS.

One important distinction between the functional optimization dynamics in RKHSGF and the measure optimization in MNE optimization dynamics is that we do not need advanced structures such as *generalized geodesic convexity* from the WGF setting (see, e.g., (Ambrosio et al., 2008)) — linearity or regular convexity is sufficient for RKHSGF. This can be seen in the following example.

Example 3.1. [Convexity] In a 2-Wasserstein space (\mathcal{M}, W_2) , linear energy functional $\mathcal{E}_0(\mu) = \int V_0 d\mu$ for $\mu \in (\mathcal{M}, W_2)$, is non-convex geodesically (Section A.3) if and only if the function $V_0(x)$ is non-convex (Ambrosio et al., 2008). On the other hand, in an RKHS \mathcal{H} , linear energy functional $\mathcal{E}_1(f) = \langle V_1, f \rangle_{\mathcal{H}}$ for $f \in \mathcal{H}$ is always convex, regardless of the non-convexity of V_1 .

Example 3.2. [Intractability of Wasserstein DRO with nonlinear losses] The Wasserstein-DRO problem

$$\inf_{\theta \in \Theta} \sup_{W_p(\mu, \hat{P}_n) \leq \epsilon} \mathbb{E}_{\mu}[l(\theta; x)] \quad (17)$$

where $l(\theta; x)$ is non-convex non-linear w.r.t. x is intractable. As discussed in Example 3.1, despite the linear-in-measure inner problem, the objective function $\mathbb{E}_{\mu}[l(\theta; x)]$ is not geodesically convex w.r.t. μ in the Wasserstein space. See Section A.3 for more details. Hence the popular approaches using linear duality reformulation, used in a large body of literature, are not tractable for this class of problems, unless strong restrictions are imposed. In other words, there is still no free lunch despite the lifting to the linear structure. The only principled treatment in the non-convex case is via Lagrangian relaxation under non-convex objectives (Sinha et al., 2020), which no longer preserves the ambiguity-ball structure. In contrast, we show convergence guarantees for Kernel-MMD DRO in Section 5 without placing restricting assumptions on $l(\theta; x)$ w.r.t. x .

In contrast, the RKHSGF in our formulation converges under the usual convexity in the Hilbert space. In the next section, we will use the explicit Euler discretization of this gradient flow, coupled with measure-update step, to derive the *primal-dual kernel mirror prox* algorithm, which is the discrete-time counterpart

to the interacting gradient flow.

Interacting Gradient Flow Dynamics Previously, (Domingo-Enrich et al., 2020) proposed the Interacting Wasserstein Gradient Flow as the infinite-dimensional continuous-time optimization dynamics for solving the MNE problem (1). We now model a gradient system that couples a gradient flow in the Fisher-Rao geometry and an *RKHS gradient flow* (RKHSGF). We term the mean-field dynamics *Interacting Gradient Flow*. Our perspective is to choose the (dual) function space in MFNE (4) as a RKHS \mathcal{H} (Wendland, 2004; Steinwart & Christmann, 2008; Berlinet & Thomas-Agnan, 2011). That is to say, letting $\mathcal{F} = \mathcal{H}$ in MFNE (4), we solve a specific variation of MFNE, which is the *Mixed Kernel Nash Equilibrium* (MKNE)

$$\inf_{f \in \mathcal{H}} \sup_{\mu \in \mathcal{M}} F(f, \mu). \quad (18)$$

The gradient flow equations governing the interacting system are

$$\partial_t \mu - \mu \cdot F'_\mu(f, \mu) + \int \mu \cdot F'_\mu(f, \mu) = 0, \quad (19)$$

$$\partial_t f + F'_f(f, \mu) = 0, \quad (20)$$

$$\mu(0, x) = \mu^0(x) \in \mathcal{M}, \quad f(0, x) = f^0(x) \in \mathcal{H}.$$

where the derivatives F'_μ, F'_f are taken in the sense of the Fréchet differential w.r.t. L^2 and RKHS norms respectively. While it is also possible to use other spaces to replace \mathcal{H} , such as neural networks to generate so-called neural flows, for the sake of analysis, this paper only focuses on the RKHS setting.

When viewed standalone, the properties of the RKHS gradient flow equation (20) and the *reaction equation* (19) have already been discussed above. Our choice of the Fisher-Rao type of flows is due to that its relation to the entropic mirror descent algorithm (15), which is a common tool for optimizing over probability measures in practical computation. Later, we further provide the convergence guarantee for a discrete-time mirror prox algorithm based on our choice of flows.

Motivated by the continuous dynamics, we now propose the discrete-time saddle-point optimization algorithm that mimics the continuous-time dynamics and prove its convergence rate bounds.

4 A PRIMAL-DUAL KERNEL MIRROR PROX ALGORITHM

To construct our mirror prox algorithm, following (Hsieh et al., 2019), we make some mild regularity assumptions. As noted by those authors, the following assumptions can be met by most practical appli-

cations. Namely, we restrict \mathcal{M} to the set of probability measures on $\tilde{\Omega}$ that admit densities w.r.t. the Lebesgue measure and have a density that is continuous and positive almost everywhere on $\tilde{\Omega}$. We note that this assumption is only made for the ease of the analysis. In practice, there are also numerous particle-based algorithms that can perform the mirror descent step, see Remark 4.6. We also assume that there is a Hilbert space \mathcal{H} and a convex and closed set $H \subseteq \mathcal{H}$. For the sake of generality, we consider a slight variation of MKNE (18) as the following general infinite-dimensional saddle-point problem on the spaces of measures and functions

$$\inf_{f \in H \subseteq \mathcal{H}} \sup_{\mu \in \mathcal{M}} F(f, \mu). \quad (21)$$

For shortness, we denote the set of all variables by $u := (f, \mu)$ and slightly abuse notation by defining $F(u) := F(f, \mu)$. We consider here the setting of two variables only for simplicity. An extension for a more general problem formulation covering the DRO problem (12) is provided in the appendix. Moreover, in the next section, we consider DRO problem (12) as a particular case study and provide technical details to check that the main assumptions of this section hold for that problem. Our first main assumption in this section is as follows.

Assumption 4.1. The functional $F(f, \mu)$ is convex in f for fixed μ and concave in μ for fixed f .

4.1 Preliminaries

To construct the mirror prox algorithm for problem (21) we need, first, to introduce proximal setup, which consists of norms, their dual, and Bregman divergences over the space of each of the variables.

For the space of the variable f , we use the norm of the Hilbert space $\|\cdot\|_{\mathcal{H}}$, distance generating function $d_f(f) = \frac{1}{2}\|f\|_{\mathcal{H}}^2$, which gives Bregman divergence $B_{\mathcal{H}}(f, \tilde{f}) = \frac{1}{2}\|f - \tilde{f}\|_{\mathcal{H}}^2$. This leads to the mirror descent step, given step size $\eta > 0$ and direction $\xi_f \in \mathcal{H}$,

$$f_+ = \text{Mirr}_{\eta}^{f, H}(f_0, \xi_f) = \arg \min_{\tilde{f} \in H} \left\{ \langle \tilde{f}, \eta \xi_f \rangle + \frac{1}{2} \|\tilde{f} - f_0\|_{\mathcal{H}}^2 \right\}. \quad (22)$$

Note that the explicit expression for the above step is the projection of $f_0 - \eta \xi_f$ onto the set H w.r.t. the RKHS norm.

For the space of the variable μ , we follow (Hsieh et al., 2019) and, first, introduce the Total Variation (TV) norm for the elements of \mathcal{M} $\|\mu\|_{TV} = \sup_{\|\xi\|_{L^\infty} \leq 1} \int \xi d\mu = \sup_{\|\xi\|_{L^\infty} \leq 1} \langle \xi, \mu \rangle$, where $\|\xi\|_{L^\infty}$ is the L^∞ -norm of functions. To define the mirror step, we use (negative) Shannon entropy and its Fenchel

dual defined respectively as

$$\Phi(\mu) = \int d\mu \ln \frac{d\mu}{dx}, \quad \Phi^*(\xi) = \ln \int e^{\xi} dx \quad (23)$$

defined for ξ from the space \mathcal{M}^* of all bounded integrable functions on $\tilde{\Omega}$. The corresponding Bregman divergence is the relative entropy given by

$$D_{\Phi}(\mu, \tilde{\mu}) = \int d\mu \ln \frac{d\mu}{d\tilde{\mu}}. \quad (24)$$

As discussed previously in (15), this leads to the mirror step

$$\begin{aligned} \mu_+ &= \text{Mirr}_{\eta}^{\mu}(\mu_0, \xi_{\mu}) = \arg \min_{\tilde{\mu} \in \mathcal{M}} \left\{ \langle \tilde{\mu}, \eta \xi_{\mu} \rangle + D_{\Phi}(\tilde{\mu}, \mu_0) \right\}. \\ &= d\Phi^*(d\Phi(\mu_0) - \eta \xi_{\mu}) \equiv d\mu_+ = \frac{e^{-\eta \xi_{\mu}} d\mu_0}{\int e^{-\eta \xi_{\mu}} d\mu_0}. \end{aligned} \quad (25)$$

Our second main assumption is as follows.

Assumption 4.2. The functional $F(u) := F(f, \mu)$ is Fréchet differentiable in f w.r.t. the RKHS norm and in μ , in L^2 , and the derivatives $F'_f(u) := F'_f(f, \mu)$, $F'_\mu(u) := F'_\mu(f, \mu)$ are Lipschitz continuous in the following sense

$$\|F'_f(u) - F'_f(\tilde{u})\|_{\mathcal{H}} \leq L_{ff} \|f - \tilde{f}\|_{\mathcal{H}} + L_{f\mu} \|\mu - \tilde{\mu}\|_{TV}, \quad (26)$$

$$\|F'_\mu(u) - F'_\mu(\tilde{u})\|_{L^\infty} \leq L_{\mu f} \|f - \tilde{f}\|_{\mathcal{H}} + L_{\mu\mu} \|\mu - \tilde{\mu}\|_{TV}. \quad (27)$$

We also denote

$$L = \max\{L_{f,f}, L_{f,\mu}, L_{\mu,f}, L_{\mu,\mu}\}. \quad (28)$$

4.2 Kernel Mirror Prox Algorithm and Its Analysis

The updates of the general infinite-dimensional mirror prox algorithm for problem (21) are given in Algorithm 1. Recall that, for shortness, we use the notation u for the pair (f, μ) and slightly abuse notation by setting $F(u) := F(f, \mu)$. For its analysis we need the following auxiliary results. The first one is used for the mirror steps applied to the variable f .

Lemma 4.3. Let \mathcal{H} be a Hilbert space and let $H \subset \mathcal{H}$ be convex and closed. Let $\tilde{h} \in H$ and $\xi, \tilde{\xi} \in \mathcal{H}^* = \mathcal{H}$, $\eta > 0$, and

$$\begin{aligned} h &= \arg \min_{\hat{h} \in H} \left\{ \langle \hat{h}, \eta \xi \rangle + \frac{1}{2} \|\tilde{h} - \hat{h}\|_{\mathcal{H}}^2 \right\} = \text{Mirr}_{\eta}^{h, H}(\tilde{h}, \xi), \\ \tilde{h}_+ &= \arg \min_{\hat{h} \in H} \left\{ \langle \hat{h}, \eta \tilde{\xi} \rangle + \frac{1}{2} \|\tilde{h} - \hat{h}\|_{\mathcal{H}}^2 \right\} = \text{Mirr}_{\eta}^{h, H}(\tilde{h}, \tilde{\xi}). \end{aligned}$$

Algorithm 1 General Mirror Prox

Require: Initial guess $\tilde{u}_0 := (\tilde{f}_0, \tilde{\mu}_0)$, step-sizes $\eta_f, \eta_\mu > 0$.

- 1: **for** $k = 0, 1, \dots, N - 1$ **do**
- 2: Compute

$$f_k = \text{Mirr}_{\eta_f}^{f,H}(\tilde{f}_k, F'_f(\tilde{u}_k)), \mu_k = \text{Mirr}_{\eta_\mu}^\mu(\tilde{\mu}_k, -F'_\mu(\tilde{u}_k)).$$

- 3: Compute

$$\tilde{f}_{k+1} = \text{Mirr}_{\eta_f}^{f,H}(\tilde{f}_k, F'_f(u_k)), \tilde{\mu}_{k+1} = \text{Mirr}_{\eta_\mu}^\mu(\tilde{\mu}_k, -F'_\mu(u_k)).$$

- 4: **end for**

- 5: Compute $(\bar{f}_N, \bar{\mu}_N) = \bar{u}_N = \frac{1}{N} \sum_{k=0}^{N-1} u_k$.
-

Then, for any $\hat{h} \in H$,

$$\begin{aligned} \langle h - \hat{h}, \eta \tilde{\xi} \rangle &\leq \frac{1}{2} \|\hat{h} - \tilde{h}\|_{\mathcal{H}}^2 - \frac{1}{2} \|\hat{h} - \tilde{h}_+\|_{\mathcal{H}}^2 \\ &\quad + \frac{\eta^2}{2} \|\tilde{\xi} - \xi\|_{\mathcal{H}}^2 - \frac{1}{2} \|h - \tilde{h}\|_{\mathcal{H}}^2. \end{aligned}$$

The second result characterizes the mirror step with respect to the measure μ .

Lemma 4.4 ((Hsieh et al., 2019, Lemma 5)). *Let $\tilde{\mu} \in \mathcal{M}$ and $\xi, \tilde{\xi} \in \mathcal{M}^*$, $\eta > 0$, and*

$$\mu = \text{Mirr}_{\eta}^\mu(\tilde{\mu}, \xi), \quad \tilde{\mu}_+ = \text{Mirr}_{\eta}^\mu(\tilde{\mu}, \tilde{\xi}). \quad (29)$$

Then, for any $\hat{\mu} \in \mathcal{M}$

$$\begin{aligned} \langle \mu - \hat{\mu}, \eta \tilde{\xi} \rangle &\leq D_{\Phi}(\hat{\mu}, \tilde{\mu}) - D_{\Phi}(\hat{\mu}, \tilde{\mu}_+) \\ &\quad + \frac{\eta^2}{8} \|\tilde{\xi} - \xi\|_{L^\infty}^2 - 2\|\mu - \tilde{\mu}\|_{TV}^2. \end{aligned} \quad (30)$$

The following result gives the convergence rate of Algorithm 1.

Theorem 4.5. *Let Assumptions 4.1, 4.2 hold. Let also the step sizes in Algorithm 1 satisfy $\eta_f = \eta_\mu = \frac{1}{16L}$, where L is defined in (28). Then, for any compact set $U = U_f \times U_\mu \subseteq H \times \mathcal{M}$, the sequence $(\bar{f}_N, \bar{\mu}_N)$ generated by Algorithm 1 satisfies*

$$\begin{aligned} \max_{\mu \in U_\mu} F(\bar{f}_N, \mu) - \min_{f \in U_f} F(f, \bar{\mu}_N) \\ \leq \frac{8L}{N} \max_{u \in U} \left(\|f - \tilde{f}_0\|_{\mathcal{H}}^2 + 2D_{\Phi}(\mu, \tilde{\mu}_0) \right). \end{aligned}$$

Remark 4.6 (Implementation of mirror descent steps). Algorithm 1 requires two mirror steps in the functional variables μ, f , both have already been separately studied in the machine learning literature. The f -update, optimization w.r.t. RKHS functions, has already been implemented for various learning tasks, e.g., by (Dai et al., 2014; Genevay et al., 2016; Tiapkin et al., 2022;

Zhu et al., 2021; Kremer et al., 2023). For example, step (22) can be implemented as solving convex quadratic program when using the RKHS function parameterization in the form $f = \sum_j \alpha_j k(x_j, \cdot)$. See those references for more implementation details. For the μ -update, many algorithms have been proposed including many particle-update schemes such as (Mei et al., 2018; Wang et al., 2021; Chizat & Bach, 2018; Dai et al., 2016; Hsieh et al., 2019; Kim et al., 2024), as well as natural gradient descent (Amari, 1998; Khan & Nielsen, 2018). An orthogonal direction is to incorporate, e.g., based on inexact Mirror Prox schemes (Stonyakin et al., 2021, 2022), an additional error estimates incurred by adopting one of those specific implementation schemes, e.g., particle mirror descent, which is beyond the scope of this paper and left for future work. Our current scope is also similar to another recent work on infinite-dimensional mirror descent analysis by (Aubin-Frankowski et al., 2022).

4.3 Analysis of Stochastic Kernel Mirror Prox

To account for potential inexactness in the first-order information, we assume that, instead of exact derivatives, the algorithm uses their inexact counterparts $\tilde{F}'_f(u), \tilde{F}'_\mu(u)$ that may be random and are assumed to satisfy the following assumption.

Assumption 4.7.

$$\begin{aligned} F'_f(u) &= \mathbb{E} \tilde{F}'_f(u), \quad F'_\mu(u) = \mathbb{E} \tilde{F}'_\mu(u), \\ \mathbb{E} \|F'_f(u) - \tilde{F}'_f(u)\|_{\mathcal{H}}^2 &\leq \sigma_f^2, \\ \mathbb{E} \|F'_\mu(u) - \tilde{F}'_\mu(u)\|_{L^\infty}^2 &\leq \sigma_\mu^2. \end{aligned}$$

Theorem 4.8. *Let Assumptions 4.1, 4.2, 4.7 hold. Let also in Algorithm 1 the stochastic derivatives be used instead of the deterministic and the step sizes satisfy $\eta_f = \eta_\mu = \frac{1}{16L}$, where L is defined in (28). Then, for any compact set $U = U_f \times U_\mu \subseteq H \times \mathcal{M}$, the sequence $(\bar{f}_N, \bar{\mu}_N)$ generated by Algorithm 1 satisfies*

$$\begin{aligned} \mathbb{E} \left\{ \max_{\mu \in U_\mu} F(\bar{f}_N, \mu) - \min_{f \in U_f} F(f, \bar{\mu}_N) \right\} \\ \leq \frac{8L}{N} \max_{u \in U} \left(\|f - \tilde{f}_0\|_{\mathcal{H}}^2 + 2D_{\Phi}(\mu, \tilde{\mu}_0) \right) + \frac{3(\sigma_f^2 + \sigma_\mu^2)}{16L}. \end{aligned}$$

Let us denote $\sigma^2 = \sigma_f^2 + \sigma_\mu^2$. Theorem 4.8 guarantees the same convergence rate as in the exact case, but up to some vicinity which is governed by the level of noise. In most cases, the σ^2/L term can be made of the same order $1/N$ by using the mini-batching technique. Indeed, a mini-batch of size N allows us to reduce the variance from σ^2 to σ^2/N . Yet, we note that in this case, N iterations will require the number of samples $O(N^2)$. This corresponds to iteration

complexity $O(1/\varepsilon)$ and sample complexity $O(1/\varepsilon^2)$ to reach an accuracy ε .

An alternative would be to use the information about the diameter of the set U . Indeed, assume that

$$\max_{u \in U} \left(\|f - \tilde{f}_0\|_{\mathcal{H}}^2 + 2D_{\Phi}(\mu, \tilde{\mu}_0) \right) \leq \Omega_U^2.$$

Fixing the number of steps N and choosing $\eta_f = \eta_{\mu} = \min \left\{ \frac{1}{16L}, \frac{\Omega_U \sigma}{\sqrt{6N}} \right\}$, we obtain the following result

$$\mathbb{E} \left\{ \max_{\mu \in U_{\mu}} F(\bar{f}_N, \mu) - \min_{f \in U_f} F(f, \bar{\mu}_N) \right\} \leq \max \left\{ \frac{8L\Omega_U^2}{N}, \sqrt{\frac{3\sigma^2\Omega_U^2}{2N}} \right\}. \quad (31)$$

5 DRO Algorithmic Guarantees using Kernel Mirror Prox

In this section we particularize the elements of KMP in Algorithm 1 for the specific DRO problem (10) equipped with the kernel MMD constraint using its primal-dual formulation (12). In contrast with many reformulation techniques using simple linear duality formulation, we propose a principled *primal-dual convergence analysis* for DRO using our MKNE framework. Notably, we prove the finite-sample robustness *algorithmic guarantee* for the DRO solution generated by the KMP algorithm after N iterations under non-linear DRO losses.

Compared to problem (21), the DRO formulation (12) has two additional variables $\theta \in \mathbb{R}^d$ and $h \in H \subset \mathcal{H}$. For these variables, the proximal setup is introduced in the same way as for the variable f . We choose \mathcal{H} to be a reproducing kernel Hilbert space with kernel k . Our main assumptions for problem (12) are **1.** $L_0 = \sup_{x, \theta} \|\nabla_{\theta} l(\theta; x)\|_2 < +\infty$. **2.** $\nabla_{\theta} l(\theta; x)$ is $L(x)$ -Lipschitz w.r.t. θ and $L_1 = \sup_{\mu} \mathbb{E}_{x \sim \mu} L(x)^2 < +\infty$. **3.** $C = \sup_x k(x, x) < +\infty$. **4.** For any fixed x , $l(\theta; x)$ is convex w.r.t. θ (but not necessarily w.r.t. x).

Remark 5.1. Most importantly, we do not place restrictive assumptions on the loss l w.r.t. variable x , which is a common practice in the DRO literature. That is, we do not assume that $l(\theta; x)$ is convex, concave, affine, or quadratic in x . This is a significant improvement over the previous DRO analysis, see (Kuhn et al., 2019). We also note that the convexity assumption w.r.t. θ is only used to obtain the global convergence guarantee. When it does not hold in practice, we can still execute our primal-dual KMP algorithm for DRO. This is not possible with other existing Wasserstein or kernel DRO algorithms.

We provide more details regarding the set-up of the KMP algorithm for the DRO setting in the appendix. We now give the main result of this section that upper-bounds the solution (sub-)optimality of the KMP algorithm. To the best of our knowledge, this is the first general non-asymptotic algorithmic guarantees without strong assumptions on $l(\theta; x)$ in terms of its dependence on x . As a reminder, \hat{P}_n is the empirical data distribution, whose samples are generated from the (unknown) distribution μ_0 .

Proposition 5.2 (DRO Guarantee for KMP decision sub-optimality). *Suppose $\bar{\theta}_N = \frac{1}{N} \sum_{k=0}^{N-1} \theta_k$ is the averaged solution produced by the kernel mirror prox algorithm after N steps. Then, $\forall \varepsilon > 0$, the DRO risk associated with the decision $\bar{\theta}_N$ is bounded by*

$$\sup_{\text{MMD}(\mu, \hat{P}_n) \leq \varepsilon} \mathbb{E}_{\mu} l(\bar{\theta}_N; x) - \underbrace{\inf_{\theta \in \Theta} \sup_{\text{MMD}(\mu, \hat{P}_n) \leq \varepsilon} \mathbb{E}_{\mu} l(\theta; x)}_{(\text{Optimal DRO risk})} \leq O\left(\frac{1}{N}\right). \quad (32)$$

The bound becomes $O(\frac{1}{\sqrt{N}})$ in expectation in the stochastic case, which we provide in the appendix. The result above quantifies that the decision generated by the KMP algorithm after N iterations is close to the ideal true DRO solution θ_{DRO}^* in terms of the risk. This distinguishes our analysis from the ideal DRO bounds that do not incorporate optimization errors. The proof of this result, given in the appendix, is due to our analysis of the KMP algorithm in the previous section.

Using this main result, we can derive further guarantees such as the ones below. Let $\varepsilon_n(\delta) := \sqrt{\frac{C}{n}} + \sqrt{\frac{2C \log(1/\delta)}{n}}$ for $C := \sup_x |k(x, x)|$, which can be calculated for commonly-used kernels, e.g., $C = 1$ for the Gaussian kernel. Note that n is the sample size instead of the number of iterations N .

Corollary 5.3 (Robustness guarantee for population risk). *Suppose the ambiguity level ε is chosen such that $\varepsilon > \varepsilon_n(\delta)$. Then, with large probability, the population risk of the decision $\bar{\theta}_N$ output by the KMP algorithm after N steps is estimated from above by*

$$\mathbb{E}_{\mu_0} [l(\bar{\theta}_N; x)] - \inf_{\theta \in \Theta} \sup_{\text{MMD}(\mu, \hat{P}_n) \leq \varepsilon} \mathbb{E}_{\mu} l(\theta; x) \leq O\left(\frac{1}{N}\right). \quad (33)$$

Corollary 5.4 (Generalization guarantee under distribution shift). *The risk under the worst-case distribution shift from the population distribution μ_0 , of any radius $\varepsilon > 0$, of the decision $\bar{\theta}_N$ output by the KMP*

algorithm after N steps is upper-bounded with large probability by

$$\begin{aligned}
 & \sup_{\text{MMD}(\mu, \mu_0) \leq \epsilon} \mathbb{E}_\mu l(\bar{\theta}_N; x) \\
 - & \inf_{\theta \in \Theta} \sup_{\text{MMD}(\mu, \hat{P}_n) \leq \epsilon} \mathbb{E}_\mu l(\theta; x) \leq O\left(\max\left\{\frac{1}{N}, \frac{1}{\sqrt{n}}\right\}\right).
 \end{aligned}
 \tag{34}$$

6 DISCUSSION

In conclusion, we model the measure optimization dynamics as an interacting gradient flow that couples the Fisher-Rao and the RKHS flow. We then analyze its corresponding primal-dual kernel mirror prox algorithm. We provide the first unified convergence analysis for the MKNE problem class. Furthermore, by applying our analysis to the primal-dual reformulation of DRO with kernel MMD constraints, we established algorithmic guarantees for DRO under relaxed conditions such as nonlinear losses. In the future, we plan to investigate dual functional spaces \mathcal{F} other than RKHS as well as practical implementation and parameterization of the mirror prox algorithm. A related direction of future work is extension of the algorithms and analysis for problems lacking convexity based on techniques in (Beznosikov et al., 2022; Gorbunov et al., 2022).

Acknowledgements

The authors thank Anna Korba and Lénaïc Chizat for the helpful discussion and pointing out relevant references. This project has received funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689).

References

Agueh, M. and Carlier, G. Barycenters in the wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.

Amari, S.-I. Natural gradient works efficiently in learning. *Neural computation*, 10(2):251–276, 1998.

Ambrosio, L., Gigli, N., and Savare, G. *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*. Springer Science & Business Media, 2008.

Amos, B., Xu, L., and Kolter, J. Z. Input convex neural networks. In *International Conference on Machine Learning*, pp. 146–155. PMLR, 2017.

Antonakopoulos, K., Belmege, V., and Mertikopoulos, P. An adaptive mirror-prox method for variational

inequalities with singular operators. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8455–8465. Curran Associates, Inc., 2019.

Arbel, M., Korba, A., Salim, A., and Gretton, A. Maximum Mean Discrepancy Gradient Flow. *arXiv:1906.04370 [cs, stat]*, December 2019. URL <http://arxiv.org/abs/1906.04370>. arXiv:1906.04370.

Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein generative adversarial networks. In *International conference on machine learning*, pp. 214–223. PMLR, 2017.

Aubin-Frankowski, P.-C., Korba, A., and Léger, F. Mirror descent with relative smoothness in measure spaces, with application to sinkhorn and em. *Advances in Neural Information Processing Systems*, 35:17263–17275, 2022.

Bach, F. Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681, 2017.

Ben-Tal, A. and Nemirovski, A. *Lectures on Modern Convex Optimization (Lecture Notes)*. Personal web-page of A. Nemirovski, 2023.

Ben-Tal, A., Den Hertog, D., De Waegenaere, A., Melenberg, B., and Rennen, G. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013.

Berlinet, A. and Thomas-Agnan, C. *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media, 2011.

Beznosikov, A., Dvurechenskii, P., Koloskova, A., Samokhin, V., Stich, S. U., and Gasnikov, A. Decentralized local stochastic extra-gradient for variational inequalities. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 38116–38133. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/f9379afacdbabfdc6b060972b60f9ab8-Paper-Conference.pdf.

Bińkowski, M., Sutherland, D. J., Arbel, M., and Gretton, A. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*, 2018.

Boucheron, S., Bousquet, O., and Lugosi, G. Theory of Classification: a Survey of Some Recent Advances. *ESAIM: Probability and Statistics*, 9: 323–375, November 2005. ISSN 1292-8100, 1262-3318. doi: 10.1051/ps:2005018. URL <http://www.esaim-ps.org/10.1051/ps:2005018>.

- Chewi, S., Le Gouic, T., Lu, C., Maunu, T., and Rigollet, P. Sgvd as a kernelized wasserstein gradient flow of the chi-squared divergence. *Advances in Neural Information Processing Systems*, 33:2098–2109, 2020.
- Chizat, L. and Bach, F. On the global convergence of gradient descent for over-parameterized models using optimal transport. *Advances in neural information processing systems*, 31, 2018.
- Chu, C., Minami, K., and Fukumizu, K. The equivalence between stein variational gradient descent and black-box variational inference. *arXiv preprint arXiv:2004.01822*, 2020.
- Dai, B., Xie, B., He, N., Liang, Y., Raj, A., Balcan, M.-F. F., and Song, L. Scalable kernel methods via doubly stochastic gradients. *Advances in neural information processing systems*, 27, 2014.
- Dai, B., He, N., Dai, H., and Song, L. Provable bayesian inference via particle mirror descent. In *Artificial Intelligence and Statistics*, pp. 985–994. PMLR, 2016.
- Delage, E. and Ye, Y. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations research*, 58(3):595–612, 2010. Publisher: INFORMS.
- Domingo-Enrich, C., Jelassi, S., Mensch, A., Rotskoff, G., and Bruna, J. A mean-field analysis of two-player zero-sum games. In *Advances in Neural Information Processing Systems*, volume 33, pp. 20215–20226. Curran Associates, Inc., 2020.
- Duncan, A., Nüsken, N., and Szpruch, L. On the geometry of stein variational gradient descent. *arXiv preprint arXiv:1912.00894*, 2019.
- Dvinskikh, D. and Tiapkin, D. Improved complexity bounds in wasserstein barycenter problem. In Banerjee, A. and Fukumizu, K. (eds.), *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13–15, 2021, Virtual Event*, volume 130 of *Proceedings of Machine Learning Research*, pp. 1738–1746. PMLR, 2021. URL <http://proceedings.mlr.press/v130/dvinskikh21a.html>.
- Dvurechensky, P., Nesterov, Y., and Spokoiny, V. Primal-dual methods for solving infinite-dimensional games. *Journal of Optimization Theory and Applications*, 166(1):23–51, Jul 2015.
- Dvurechensky, P., Dvinskikh, D., Gasnikov, A., Uribe, C. A., and Nedić, A. Decentralize and randomize: Faster algorithm for Wasserstein barycenters. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems* 31, NeurIPS 2018, pp. 10783–10793. Curran Associates, Inc., 2018a.
- Dvurechensky, P., Gasnikov, A., and Kroshnin, A. Computational optimal transport: Complexity by accelerated gradient descent is better than by Sinkhorn’s algorithm. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1367–1376, 2018b.
- Dziugaite, G. K., Roy, D. M., and Ghahramani, Z. Training generative neural networks via maximum mean discrepancy optimization.
- Gao, R. and Kleywegt, A. J. Distributionally Robust Stochastic Optimization with Wasserstein Distance. *arXiv:1604.02199 [math]*, July 2016. URL <http://arxiv.org/abs/1604.02199>. arXiv: 1604.02199.
- Genevay, A., Cuturi, M., Peyré, G., and Bach, F. Stochastic optimization for large-scale optimal transport. *Advances in neural information processing systems*, 29, 2016.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative Adversarial Networks. *arXiv:1406.2661 [cs, stat]*, June 2014. URL <http://arxiv.org/abs/1406.2661>. arXiv: 1406.2661.
- Gorbunov, E., Danilova, M., Dobre, D., Dvurechenskii, P., Gasnikov, A., and Gidel, G. Clipped stochastic methods for variational inequalities with heavy-tailed noise. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 31319–31332. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/cb0ce861adaf6f8a93069c064733f402-Paper-Conference.pdf.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012. Publisher: JMLR. org.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. Improved Training of Wasserstein GANs. *arXiv:1704.00028 [cs, stat]*, December 2017. URL <http://arxiv.org/abs/1704.00028>. arXiv: 1704.00028.
- Guminov, S., Dvurechensky, P., Tupitsa, N., and Gasnikov, A. On a combination of alternating minimization and Nesterov’s momentum. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 3886–3898, Virtual, 18–24 Jul 2021.

- PMLR. URL <http://proceedings.mlr.press/v139/guminov21a.html>.
- Hsieh, Y.-P., Liu, C., and Cevher, V. Finding mixed Nash equilibria of generative adversarial networks. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2810–2819. PMLR, 09–15 Jun 2019.
- Jambulapati, A., Sidford, A., and Tian, K. A direct tilde $\{O\}(1/\epsilon)$ iteration parallel algorithm for optimal transport. *Advances in Neural Information Processing Systems*, 32, 2019.
- Jordan, R., Kinderlehrer, D., and Otto, F. The variational formulation of the Fokker–Planck equation. *SIAM journal on mathematical analysis*, 29(1):1–17, 1998. Publisher: SIAM.
- Khan, M. E. and Nielsen, D. Fast yet simple natural-gradient descent for variational inference in complex models. In *2018 International Symposium on Information Theory and Its Applications (ISITA)*, pp. 31–35. IEEE, 2018.
- Kim, J., Yamamoto, K., Oko, K., Yang, Z., and Suzuki, T. Symmetric Mean-field Langevin Dynamics for Distributional Minimax Problems, February 2024.
- Korba, A., Aubin-Frankowski, P.-C., Majewski, S., and Ablin, P. Kernel Stein Discrepancy Descent. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 5719–5730. PMLR, July 2021. ISSN: 2640-3498.
- Korotin, A., Egiazarian, V., Asadulaev, A., Safin, A., and Burnaev, E. Wasserstein-2 generative networks. *arXiv preprint arXiv:1909.13082*, 2019.
- Korotin, A., Li, L., Solomon, J., and Burnaev, E. Continuous Wasserstein-2 Barycenter Estimation without Minimax Optimization. *arXiv:2102.01752 [cs, stat]*, February 2021. URL <http://arxiv.org/abs/2102.01752>. arXiv: 2102.01752.
- Kremer, H., Nemmour, Y., Schölkopf, B., and Zhu, J.-J. Estimation beyond data reweighting: Kernel method of moments. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J. (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 17745–17783. PMLR, July 2023.
- Kroshnin, A., Tupitsa, N., Dvinskikh, D., Dvurechensky, P., Gasnikov, A., and Uribe, C. On the complexity of approximating Wasserstein barycenters. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3530–3540, Long Beach, California, USA, 09–15 Jun 2019. PMLR.
- Kuhn, D., Esfahani, P. M., Nguyen, V. A., and Shafieezadeh-Abadeh, S. Wasserstein Distributionally Robust Optimization: Theory and Applications in Machine Learning. In Netessine, S., Shier, D., and Greenberg, H. J. (eds.), *Operations Research & Management Science in the Age of Analytics*, pp. 130–166. INFORMS, October 2019. ISBN 978-0-9906153-3-0. doi: 10.1287/educ.2019.0198.
- Li, C.-L., Chang, W.-C., Cheng, Y., Yang, Y., and Póczos, B. Mmd gan: Towards deeper understanding of moment matching network. *Advances in neural information processing systems*, 30, 2017.
- Li, L., Genevay, A., Yurochkin, M., and Solomon, J. Continuous Regularized Wasserstein Barycenters. *arXiv:2008.12534 [cs, stat]*, October 2020. URL <http://arxiv.org/abs/2008.12534>. arXiv: 2008.12534.
- Li, Y., Swersky, K., and Zemel, R. Generative moment matching networks. In *International conference on machine learning*, pp. 1718–1727. PMLR, 2015.
- Liu, Q. and Wang, D. Stein Variational Gradient Descent: A General Purpose Bayesian Inference Algorithm. *arXiv:1608.04471 [cs, stat]*, September 2019. URL <http://arxiv.org/abs/1608.04471>. arXiv: 1608.04471.
- Makkuva, A., Taghvaei, A., Oh, S., and Lee, J. Optimal transport mapping via input convex neural networks. In *International Conference on Machine Learning*, pp. 6672–6681. PMLR, 2020.
- Mei, S., Montanari, A., and Nguyen, P.-M. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, August 2018. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1806579115. URL <http://www.pnas.org/lookup/doi/10.1073/pnas.1806579115>.
- Mielke, A. An introduction to the analysis of gradients systems, June 2023. URL <http://arxiv.org/abs/2306.05026>. arXiv:2306.05026 [math-ph].
- Nemirovski, A., Juditsky, A., Lan, G., and Shapiro, A. Robust Stochastic Approximation Approach to Stochastic Programming. *SIAM Journal on Optimization*, 19(4):1574–1609, January 2009. ISSN 1052-6234, 1095-7189. doi: 10.1137/070704277. URL <http://epubs.siam.org/doi/10.1137/070704277>.
- Nesterov, Y. Dual extrapolation and its applications to solving variational inequalities and related problems. *Mathematical Programming*, 109(2-3):319–344, 2007.

- Nowozin, S., Cseke, B., and Tomioka, R. f-GAN: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Otto, F. *Double degenerate diffusion equations as steepest descent*. Citeseer, 1996.
- Otto, F. The geometry of dissipative evolution equations: the porous medium equation. 2001. Publisher: Taylor & Francis.
- Peyré, G. and Cuturi, M. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019. ISSN 1935-8237. doi: 10.1561/2200000073. URL <http://dx.doi.org/10.1561/2200000073>.
- Rahimi, A. and Recht, B. Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20, 2007.
- Santambrogio, F. Optimal transport for applied mathematicians. *Birkäuser, NY*, 55(58-63):94, 2015. Publisher: Springer.
- Sinha, A., Namkoong, H., Volpi, R., and Duchi, J. Certifying Some Distributional Robustness with Principled Adversarial Training. *arXiv:1710.10571 [cs, stat]*, May 2020. URL <http://arxiv.org/abs/1710.10571>. arXiv: 1710.10571.
- Steinwart, I. and Christmann, A. *Support vector machines*. Springer Science & Business Media, 2008.
- Stonyakin, F., Tyurin, A., Gasnikov, A., Dvurechensky, P., Agafonov, A., Dvinskikh, D., Alkousa, M., Pasechnyuk, D., Artamonov, S., and Piskunova, V. Inexact model: a framework for optimization and variational inequalities. *Optimization Methods and Software*, 36(6):1155–1201, 2021. doi: 10.1080/10556788.2021.1924714. URL <https://doi.org/10.1080/10556788.2021.1924714>.
- Stonyakin, F., Gasnikov, A., Dvurechensky, P., Titov, A., and Alkousa, M. Generalized mirror prox algorithm for monotone variational inequalities: Universality and inexact oracle. *Journal of Optimization Theory and Applications*, 194(3):988–1013, Sep 2022. ISSN 1573-2878.
- Tiapkin, D., Gasnikov, A., and Dvurechensky, P. Stochastic saddle-point optimization for the wasserstein barycenter problem. *Optimization Letters*, 16(7):2145–2175, Sep 2022. ISSN 1862-4480. doi: 10.1007/s11590-021-01834-w. URL <https://doi.org/10.1007/s11590-021-01834-w>. arXiv:2006.06763.
- Tolstikhin, I., Sriperumbudur, B. K., and Muandet, K. Minimax Estimation of Kernel Mean Embeddings. *Journal of Machine Learning Research*, 18: 1–47, 2017.
- Trillos, C. G. and Trillos, N. G. On adversarial robustness and the use of wasserstein ascent-descent dynamics to enforce it. *arXiv preprint arXiv:2301.03662*, 2023.
- Vacher, A., Muzellec, B., Rudi, A., Bach, F., and Vialard, F.-X. A dimension-free computational upper-bound for smooth optimal transport estimation. In Belkin, M. and Kpotufe, S. (eds.), *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pp. 4143–4173. PMLR, 15–19 Aug 2021. URL <https://proceedings.mlr.press/v134/vacher21a.html>.
- Wainwright, M. J. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019.
- Wang, G. and Chizat, L. An exponentially converging particle method for the mixed nash equilibrium of continuous games. *arXiv preprint arXiv:2211.01280*, 2022.
- Wang, Y., Ma, X., Bailey, J., Yi, J., Zhou, B., and Gu, Q. On the Convergence and Robustness of Adversarial Training. *arXiv:2112.08304 [cs]*, December 2021. URL <http://arxiv.org/abs/2112.08304>. arXiv: 2112.08304.
- Wendland, H. *Scattered data approximation*, volume 17. Cambridge university press, 2004.
- Zhao, C. and Guan, Y. Data-driven risk-averse stochastic optimization with Wasserstein metric. *Operations Research Letters*, 46(2):262–267, March 2018. ISSN 01676377. doi: 10.1016/j.orl.2018.01.011.
- Zhu, J.-J., Jitkrittum, W., Diehl, M., and Schölkopf, B. Kernel Distributionally Robust Optimization: Generalized Duality Theorem and Stochastic Approximation. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, pp. 280–288. PMLR, March 2021. ISSN: 2640-3498.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]

- (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Not Applicable]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
 3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Not Applicable]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Not Applicable]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Not Applicable]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Not Applicable]
 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
 5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

A FURTHER TECHNICAL BACKGROUND

A.1 List of Acronyms

- DRO – Distributionally Robust Optimization
- $(\text{EVI})_{\lambda^-}$ – evolutionary variational inequality
- GAN – Generative Adversarial Network
- ICNN – input convex neural network
- IGM – implicit generative models
- IPM – integral probability metric
- KMP – Kernel Mirror Prox
- MMD – maximum mean discrepancy
- MMS – minimizing movement scheme
- MNE – Mixed Nash Equilibrium
- MFNE – Mixed Functional Nash Equilibrium
- MKNE – Mixed Kernel Nash Equilibrium
- RKHS – Reproducing kernel Hilbert space
- RKHSGF – RKHS gradient flow
- TV – Total Variation
- WGF – Wasserstein gradient flow

A.2 Proof of Lemma 2.1

We now prove Lemma 2.1, of which a more detailed version is stated below.

Lemma A.1 (Primal-dual reformulation of Wasserstein and kernel DRO). *Suppose the probability metric \mathcal{D} is chosen to be the MMD, then the DRO problem (10) admits the following equivalent primal-dual reformulations*

$$\inf_{\theta \in \Theta \subseteq \mathbb{R}^d, f \in \mathcal{H}, \tau > 0} \sup_{\mu \in \mathcal{M}} \mathbb{E}_{\mu}(l(\theta; x) - f(x)) + \frac{1}{n} \sum_{i=1}^n f(x_i) + \frac{\tau}{2} \|f\|_{\mathcal{H}}^2 + \frac{\epsilon^2}{2\tau}, \quad (35)$$

$$\inf_{\theta \in \Theta \subseteq \mathbb{R}^d, f \in \mathcal{H}} \sup_{\mu \in \mathcal{M}} \mathbb{E}_{\mu}(l(\theta; x) - f(x)) + \frac{1}{n} \sum_{i=1}^n f(x_i) + \epsilon \|f\|_{\mathcal{H}}. \quad (36)$$

Furthermore, it is equivalent to the smooth saddle-point problem

$$\inf_{\theta \in \Theta \subseteq \mathbb{R}^d, f \in \mathcal{H}} \sup_{\mu \in \mathcal{M}, h \in \mathcal{H}: \|h\|_{\mathcal{H}} \leq 1} \left\{ \frac{1}{n} \sum_{i=1}^n f(x_i) + \epsilon \langle h, f \rangle + \mathbb{E}_{\mu}(l(\theta; x) - f(x)) \right\}. \quad (37)$$

Suppose the probability metric \mathcal{D} is chosen to be the optimal transport metric, e.g., p -Wasserstein distance. Then, the DRO problem (10) admits the following equivalent reformulations

$$\inf_{\gamma > 0, \theta \in \Theta \subseteq \mathbb{R}^d, f \in \Psi_{\gamma \cdot c}} \sup_{\mu \in \mathcal{M}} \mathbb{E}_{\mu}(l(\theta; x) - f(x)) - \frac{1}{n} \sum_{i=1}^n f^{\gamma \cdot c}(x_i) + \gamma \cdot \epsilon, \quad (38)$$

$$\inf_{\gamma > 0, \theta \in \Theta \subseteq \mathbb{R}^d, f \in \Psi_c} \sup_{\mu \in \mathcal{M}} \mathbb{E}_{\mu}(l(\theta; x) - \gamma \cdot f(x)) - \frac{1}{n} \sum_{i=1}^n \gamma \cdot f^c(x_i) + \gamma \cdot \epsilon. \quad (39)$$

Here Ψ_c denotes the set of c -concave functions (Santambrogio, 2015) and $f^c(y) := \inf_x c(x, y) - f(x)$ denotes the c -transform.

Proof. For this proof, it suffices to consider the case where θ is fixed, since only the inner maximization is reformulated using duality. We first prove the result for the MMD setting.

MMD setting. We consider the kernel mean embedding map as a linear constraint

$$\int \phi(x) d\mu(x) = h, \quad (40)$$

where h is a function in \mathcal{H} and $\phi(x) := k(x, \cdot)$ is a canonical feature function of \mathcal{H} . Using the Lagrange multipliers $f \in \mathcal{H}$ for the above linear constraint and $1/(2\tau)$ for the inequality constraint $\text{MMD}(\mu, \hat{\mu})^2 \leq \epsilon^2$, we obtain the Lagrangian

$$\int l d\mu - \frac{1}{2\tau} \text{MMD}^2(\mu, \hat{\mu}) + \frac{\epsilon^2}{2\tau} + \langle f, \int \phi d\mu - h \rangle_{\mathcal{H}}. \quad (41)$$

Since $\text{MMD}(\mu, \hat{\mu}) = \|h - \hat{h}\|_{\mathcal{H}}$, we arrive at the Lagrange saddle-point problem

$$\inf_{\tau > 0, f \in \mathcal{H}} \sup_{\mu \in \mathcal{M}, h \in \mathcal{H}} \left\{ \int l d\mu - \frac{1}{2\tau} \|h - \hat{h}\|_{\mathcal{H}}^2 + \frac{\epsilon^2}{2\tau} - \langle f, h \rangle_{\mathcal{H}} + \int \langle f, \phi \rangle_{\mathcal{H}} d\mu \right\} \quad (42)$$

where \hat{h} is the kernel mean embedding of the empirical measure $\hat{P}_n = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}$. Carrying out the quadratic optimization problem w.r.t. h in closed form, we eliminate the variable h . Finally, using the reproducing property, and rearranging the terms, we obtain the result in (35). An optimal choice of the dual variable τ yields the equivalent reformulation (36). Smoothing via the definition of the dual norm in the Hilbert spaces, i.e., using that $\|f\|_{\mathcal{H}} = \max_{h \in \mathcal{H}: \|h\|_{\mathcal{H}} \leq 1} \langle h, f \rangle$, we obtain (37). Note that the strong duality holds due to the infinite-dimensional linear program duality; see (Zhu et al., 2021) for an elementary proof of strong duality.

Wasserstein setting. The derivaiton is similar to the MMD setting by combining the dual Kantorovich representation of OT (5) and Lagrange duality. \square

A.3 Background of Gradient Flow and Geodesic Convexity

Recall that a functional \mathcal{E} defined on a Hilbert space \mathcal{H} is λ -convex if $\forall s \in [0, 1], \forall u_0, u_1 \in \mathcal{H}$,

$$\mathcal{E}((1-s)u_0 + su_1) \leq (1-s)\mathcal{E}(u_0) + s\mathcal{E}(u_1) - \frac{\lambda}{2}s(1-s)\|u_0 - u_1\|_{\mathcal{H}}^2. \quad (43)$$

If $\lambda > 0$, \mathcal{E} is strongly convex. This notion of convexity does not make sense in the Wasserstein space.

Recent machine learning literature has explored *Wasserstein gradient flow* (WGF). The 2-Wasserstein space (\mathcal{M}, W_2) is geodesically convex, as defined below.

In a metric space $(\mathcal{M}, \mathcal{D})$ a curve $\gamma : [0, 1] \rightarrow \mathcal{M}$ is a (constant speed) geodesic if

$$\forall r, s \in [0, 1]: \quad \mathcal{D}(\gamma(r), \gamma(s)) = |s - r| \mathcal{D}(\gamma(0), \gamma(1)).$$

We refer to that as the geodesic γ connects the points $\gamma(0)$ and $\gamma(1)$ and write $\text{Geod}(\gamma(0), \gamma(1))$ for the set of all such geodesics.

Definition A.2 (Geodesic metric spaces). The metric space $(\mathcal{M}, \mathcal{D})$ is a geodesic space, if for all $u_0, u_1 \in \mathcal{M}$ there exists a geodesic connecting u_0 and u_1 .

Definition A.3 (Geodesic convexity). A functional $\mathcal{E} : \mathcal{M} \rightarrow \mathbb{R}_{\infty}$ is geodesically λ -convex if $\forall u_0, u_1 \in \mathcal{M} \exists \gamma \in \text{Geod}(u_0, u_1)$ s.t.

$$\mathcal{E}(\gamma(s)) \leq (1-s)\mathcal{E}(\gamma(0)) + s\mathcal{E}(\gamma(1)) - \frac{\lambda}{2}s(1-s)\mathcal{D}(\gamma(0), \gamma(1))^2, \quad \forall s \in [0, 1].$$

As we have seen in the main text, e.g., discussions around Example 3.2, this complication of convexity structure makes optimization in the Wasserstein space more difficult than in general Hilbert spaces, which motivates our approach to work in the RKHS.

Next, we mention standard results for RKHSGF in the context of this paper for completeness. First, we outline the existence, uniqueness, and a powerful result known as the *evolutionary variational inequalities* $(\text{EVI})_\lambda$. Since an RKHS is a Hilbert space, Lemma A.4 is simply the $(\text{EVI})_\lambda$ in a Hilbert space, whose proof is standard (Ambrosio et al., 2008; Santambrogio, 2015; Mielke, 2023).

Lemma A.4 (Characterizations of RKHS gradient flow). *Suppose the energy functional \mathcal{V} is proper, upper semicontinuous, λ -convex for some $\lambda \in \mathbb{R}$, and has compact sublevel sets. Then for any initial condition in the RKHS $f(0, x) \in \mathcal{H}$, there exists a unique solution to (16) at time t , $f(t) \in \mathcal{H}$.*

Furthermore, the gradient flow solution $f(t, x)$ satisfies the following $(\text{EVI})_\lambda$, for $t, s \in [0, T]$:

$$\frac{1}{2} \|f(t) - h\|_{\mathcal{H}}^2 \leq \frac{1}{2} e^{-\lambda(t-s)} \|f(s) - h\|_{\mathcal{H}}^2 + M_\lambda(t-s)(\mathcal{V}(h) - \mathcal{V}(f(t))), \quad \forall h \in \text{dom}(\mathcal{V}) \subset \mathcal{H} \quad (44)$$

where $M_\lambda(\tau) = \int_0^\tau e^{-\lambda(\tau-s)} ds$. Using $(\text{EVI})_\lambda$, we can effortlessly extract convergence results. Suppose a minimizer $f^* \in \inf_{f \in \mathcal{H}} \mathcal{V}(f)$ of the energy exists, we set $h = f^*$, $s = 0$ in $(\text{EVI})_\lambda$ and obtain

$$\|f(t) - f^*\|_{\mathcal{H}}^2 \leq e^{-\lambda t} \|f(0) - f^*\|_{\mathcal{H}}^2 + 2M_\lambda t \left(\inf_{f \in \mathcal{H}} \mathcal{V}(f) - \mathcal{V}(f(t)) \right) \leq e^{-\lambda t} \|f(0) - f^*\|_{\mathcal{H}}^2, \quad (45)$$

yielding an *exponential convergence in time* if the energy \mathcal{V} is strongly convex (in the usual sense) w.r.t. the f variable, i.e., $\lambda > 0$. Note that the convexity condition can be further weakened using functional inequalities such as the Łojasiewicz inequality.

A time-discretization of the RKHS gradient flow (16) using the standard *minimizing movement scheme* (MMS) in the gradient flow literature, is as follows

$$\hat{f}^{k+1} \in \arg \inf_{f \in \mathcal{H}} \left\{ \mathcal{V}(f) + \frac{1}{2\tau} \|f - \hat{f}^k\|_{\mathcal{H}}^2 \right\}, \quad (46)$$

$$\hat{f}^0(x) = f(0, x) \in \mathcal{H}, \quad (47)$$

where $\tau > 0$ is the step size for time-discretization. Defining the piecewise constant function $\bar{f}_\tau(t, x) := \hat{f}^k(x)$ for $t \in [k\tau, (k+1)\tau]$, standard PDE proofs (see, e.g., (Ambrosio et al., 2008)) guarantee that \bar{f}_τ converges to the continuous-time RKHSGF solution, i.e., $\bar{f}_\tau \rightarrow f(t)$ as $\tau \rightarrow 0$. The fully-implicit Euler discretization is difficult to implement in practice. In the next section, we will use the explicit step, coupled with measure-update step.

B THEORETICAL ANALYSIS OF PRIMAL-DUAL KERNEL MIRROR PROX IN SECTION 4

For the sake of generality, in particular, to cover the DRO problem (12), we consider a more general saddle-point problem. Theorems 4.5 and 4.8 are obtained as corollaries of the results obtained in this section. To make this section self-contained and for the reader's convenience we repeat some definitions and results given in the main text.

We consider generic variable $x \in \mathbb{R}^n$ with domain \mathcal{X} . We denote by \mathcal{M} the set of all probability measures on \mathcal{X} that admit densities w.r.t. the Lebesgue measure and the density is continuous and positive almost everywhere on \mathcal{X} . We also assume that there are two Hilbert spaces $\mathcal{H}_f, \mathcal{H}_h$, a convex set $\Theta \subseteq \mathbb{R}^d$, and convex a compact $H \in \mathcal{H}_h$. We consider the following general infinite-dimensional saddle-point problem

$$\inf_{\theta \in \Theta \subseteq \mathbb{R}^d, f(x) \in \mathcal{H}_f} \sup_{\mu \in \mathcal{M}, h(x) \in H \subseteq \mathcal{H}_h} F(\theta, f, \mu, h). \quad (48)$$

For shortness, we denote the set of all variables by $u = (\theta, f, \mu, h)$ and slightly abuse notation to define $F(u) = F(\theta, f, \mu, h)$.

Our first main assumption is as follows.

Assumption B.1. The functional $F(\theta, f, \mu, h)$ is convex in (θ, f) for fixed (μ, h) and concave in (μ, h) for fixed (θ, f) .

B.1 Preliminaries

To construct the mirror prox algorithm for problem (48) we need to first introduce proximal setup, which consists of norms, their dual, and Bregman divergences on each space of the variables.

For the space of the variable θ , we introduce the standard proximal setup with the Euclidean norm $\|\cdot\|_2$, distance-generating function $d_\theta(\theta) = \frac{1}{2}\|\theta\|_2^2$, which gives Bregman divergence $B_\theta(\theta, \check{\theta}) = \frac{1}{2}\|\theta - \check{\theta}\|_2^2$. This leads to the mirror step defined as

$$\theta_+ = \text{Mirr}_\eta^{\theta, \Theta}(\theta_0, \xi_\theta) = \arg \min_{\check{\theta} \in \Theta} \{ \langle \check{\theta}, \eta \xi_\theta \rangle + \frac{1}{2} \|\check{\theta} - \theta_0\|_2^2 \}. \quad (49)$$

We note that our choice of the Euclidean proximal setup is made for simplicity and that other standard proximal setups as in (Nemirovski et al., 2009) are possible.

For the space of the variable f , we use the norm of the Hilbert space $\|\cdot\|_{\mathcal{H}_f}$, distance generating function $d_f(f) = \frac{1}{2}\|f\|_{\mathcal{H}_f}^2$, which gives Bregman divergence $B_{\mathcal{H}_f}(f, \check{f}) = \frac{1}{2}\|f - \check{f}\|_{\mathcal{H}_f}^2$. This leads to the mirror step

$$f_+ = \text{Mirr}_\eta^{f, \mathcal{H}_f}(f_0, \xi_f) = \arg \min_{\check{f}} \{ \langle \check{f}, \eta \xi_f \rangle + \frac{1}{2} \|\check{f} - f_0\|_{\mathcal{H}_f}^2 \} = f_0 - \eta \xi_f. \quad (50)$$

Note that if $\xi_f = \mathcal{V}'(f)$ (the Fréchet differential of the energy \mathcal{V}), the above equation is equivalent to $\frac{1}{\tau}(f_+ - f) = -\mathcal{V}'(f)$, which can be seen as a time discretization of (20).

For the space of the variable μ , we follow Hsieh et al. (2019) and, first, introduce the Total Variation norm for the elements of \mathcal{M}

$$\|\mu\|_{TV} = \sup_{\|\xi\|_{L^\infty} \leq 1} \int \xi d\mu = \sup_{\|\xi\|_{L^\infty} \leq 1} \langle \xi, \mu \rangle,$$

where $\|\xi\|_{L^\infty}$ is the L^∞ -norm of functions. To define the mirror step, we use (negative) Shannon entropy

$$\Phi(\mu) = \int d\mu \ln \frac{d\mu}{dx} \quad (51)$$

and its Fenchel dual

$$\Phi^*(\xi) = \ln \int e^{\xi} dx \quad (52)$$

defined for ξ from the space \mathcal{M}^* of all bounded integrable functions on \mathcal{X} . The corresponding Bregman divergence is the relative entropy given by

$$D_\Phi(\mu, \check{\mu}) = \int d\mu \ln \frac{d\mu}{d\check{\mu}}. \quad (53)$$

This leads to the mirror step (Hsieh et al., 2019, Theorem 1)

$$\mu_+ = \text{Mirr}_\eta^\mu(\mu_0, \xi_\mu) = d\Phi^*(d\Phi(\mu_0) - \eta \xi_\mu) \quad \equiv \quad d\mu_+ = \frac{e^{-\eta \xi_\mu} d\mu_0}{\int e^{-\eta \xi_\mu} d\mu_0}. \quad (54)$$

This step can be seen as a time discretization of the dynamics (19) as outlined in the discussion around (15).

Finally, for the space of the variable h , we do the same as for the variable f . Namely, we use the distance generating function $d_h(h) = \frac{1}{2}\|h\|_{\mathcal{H}_h}^2$, which gives Bregman divergence $B_{\mathcal{H}_h}(h, \check{h}) = \frac{1}{2}\|h - \check{h}\|_{\mathcal{H}_h}^2$. This leads to the mirror step

$$h_+ = \text{Mirr}_\eta^{h, \mathcal{H}_h}(h_0, \xi_h) = \arg \min_{\check{h} \in H} \{ \langle \check{h}, \eta \xi_h \rangle + \frac{1}{2} \|\check{h} - h_0\|_{\mathcal{H}_h}^2 \}. \quad (55)$$

Our second main assumption is as follows

Assumption B.2. The functional $F(\theta, f, \mu, h)$ is (Fréchet) differentiable in θ w.r.t. the Euclidean norm, in f w.r.t. the RKHS norm, in μ w.r.t. L^2 , and in h w.r.t. the RKHS norm. Furthermore, we assume that the

derivatives are Lipschitz continuous in the following sense

$$\|F'_\theta(u) - F'_\theta(\tilde{u})\|_2 \leq L_{\theta\theta}\|\theta - \tilde{\theta}\|_2 + L_{\theta f}\|f - \tilde{f}\|_{\mathcal{H}_f} + L_{\theta\mu}\|\mu - \tilde{\mu}\|_{TV} + L_{\theta h}\|h - \tilde{h}\|_{\mathcal{H}_h}, \quad (56)$$

$$\|F'_f(u) - F'_f(\tilde{u})\|_{\mathcal{H}_f} \leq L_{f\theta}\|\theta - \tilde{\theta}\|_2 + L_{ff}\|f - \tilde{f}\|_{\mathcal{H}_f} + L_{f\mu}\|\mu - \tilde{\mu}\|_{TV} + L_{fh}\|h - \tilde{h}\|_{\mathcal{H}_h}, \quad (57)$$

$$\|F'_\mu(u) - F'_\mu(\tilde{u})\|_{L^\infty} \leq L_{\mu\theta}\|\theta - \tilde{\theta}\|_2 + L_{\mu f}\|f - \tilde{f}\|_{\mathcal{H}_f} + L_{\mu\mu}\|\mu - \tilde{\mu}\|_{TV} + L_{\mu h}\|h - \tilde{h}\|_{\mathcal{H}_h}, \quad (58)$$

$$\|F'_h(u) - F'_h(\tilde{u})\|_{\mathcal{H}_h} \leq L_{h\theta}\|\theta - \tilde{\theta}\|_2 + L_{hf}\|f - \tilde{f}\|_{\mathcal{H}_f} + L_{h\mu}\|\mu - \tilde{\mu}\|_{TV} + L_{hh}\|h - \tilde{h}\|_{\mathcal{H}_h}. \quad (59)$$

We also denote

$$L = \max_{\kappa_1, \kappa_2 \in \{\theta, f, \mu, h\}} \{L_{\kappa_1 \kappa_2}\}. \quad (60)$$

B.2 Mirror Prox Algorithm and Its Analysis

The updates of the general infinite-dimensional Mirror Prox algorithm¹ for problem (48) are given in Algorithm 2.

Algorithm 2 Ideal General Mirror-Prox

Require: Initial guess $\tilde{u}_0 = (\tilde{\theta}_0, \tilde{f}_0, \tilde{\mu}_0, \tilde{h}_0)$, step-sizes $\eta_\theta, \eta_f, \eta_\mu, \eta_h > 0$.

1: **for** $k = 0, 1, \dots, N - 1$ **do**

2: Compute for $\tilde{u}_k = (\tilde{\theta}_k, \tilde{f}_k, \tilde{\mu}_k, \tilde{h}_k)$

$$\begin{aligned} \theta_k &= \text{Mirr}_{\eta_\theta}^{\theta, \Theta}(\tilde{\theta}_k, F'_\theta(\tilde{u}_k)), & f_k &= \text{Mirr}_{\eta_f}^{f, \mathcal{H}_f}(\tilde{f}_k, F'_f(\tilde{u}_k)), \\ \mu_k &= \text{Mirr}_{\eta_\mu}^\mu(\tilde{\mu}_k, -F'_\mu(\tilde{u}_k)), & h_k &= \text{Mirr}_{\eta_h}^{h, H}(\tilde{h}_k, -F'_h(\tilde{u}_k)). \end{aligned}$$

3: Compute for $u_k = (\theta_k, f_k, \mu_k, h_k)$

$$\begin{aligned} \tilde{\theta}_{k+1} &= \text{Mirr}_{\eta_\theta}^{\theta, \Theta}(\tilde{\theta}_k, F'_\theta(u_k)), & \tilde{f}_{k+1} &= \text{Mirr}_{\eta_f}^{f, \mathcal{H}_f}(\tilde{f}_k, F'_f(u_k)), \\ \tilde{\mu}_{k+1} &= \text{Mirr}_{\eta_\mu}^\mu(\tilde{\mu}_k, -F'_\mu(u_k)), & \tilde{h}_{k+1} &= \text{Mirr}_{\eta_h}^{h, H}(\tilde{h}_k, -F'_h(u_k)). \end{aligned}$$

4: **end for**

5: Compute $(\bar{\theta}_N, \bar{f}_N, \bar{\mu}_N, \bar{h}_N) = \bar{u}_N = \frac{1}{N} \sum_{k=0}^{N-1} u_k = \frac{1}{N} \sum_{k=0}^{N-1} (\theta_k, f_k, \mu_k, h_k)$.

For the analysis of the mirror prox algorithm, we need the following auxiliary results. The first one is used for the mirror steps applied to the variables θ, f, h .

Lemma B.3. *Let \mathcal{H} be (possibly finite-dimensional) Hilbert space and let $H \subset \mathcal{H}$ be convex and closed. Let $\tilde{h} \in H$ and $\xi, \tilde{\xi} \in \mathcal{H}^* = \mathcal{H}$, $\eta > 0$, and*

$$h = \arg \min_{\hat{h} \in H} \left\{ \langle \hat{h}, \eta \xi \rangle + \frac{1}{2} \|\tilde{h} - \hat{h}\|_{\mathcal{H}}^2 \right\} = \text{Mirr}_\eta^{h, H}(\tilde{h}, \xi), \quad (61)$$

$$\tilde{h}_+ = \arg \min_{\hat{h} \in H} \left\{ \langle \hat{h}, \eta \tilde{\xi} \rangle + \frac{1}{2} \|\tilde{h} - \hat{h}\|_{\mathcal{H}}^2 \right\} = \text{Mirr}_\eta^{h, H}(\tilde{h}, \tilde{\xi}). \quad (62)$$

Then, for any $\hat{h} \in H$

$$\langle h - \hat{h}, \eta \tilde{\xi} \rangle \leq \frac{1}{2} \|\hat{h} - \tilde{h}\|_{\mathcal{H}}^2 - \frac{1}{2} \|\hat{h} - \tilde{h}_+\|_{\mathcal{H}}^2 + \frac{\eta^2}{2} \|\tilde{\xi} - \xi\|_{\mathcal{H}}^2 - \frac{1}{2} \|h - \tilde{h}\|_{\mathcal{H}}^2. \quad (63)$$

Proof. By the optimality condition in (62), we have for all $\hat{h} \in H$

$$\langle \eta \tilde{\xi} - (\tilde{h} - \tilde{h}_+), \hat{h} - \tilde{h}_+ \rangle \geq 0. \quad (64)$$

¹We believe that the Dual Averaging algorithm (Nesterov, 2007) can be extended to our setting in a similar fashion as it was done in (Dvurechensky et al., 2015) for saddle-point problems in Hilbert spaces.

Rearranging, we obtain

$$\langle \tilde{h}_+ - \hat{h}, \eta \tilde{\xi} \rangle \leq \langle \tilde{h}_+ - \hat{h}, \tilde{h} - \tilde{h}_+ \rangle = -\frac{1}{2} \|\tilde{h}_+ - \tilde{h}\|_{\mathcal{H}}^2 - \frac{1}{2} \|\hat{h} - \tilde{h}_+\|_{\mathcal{H}}^2 + \frac{1}{2} \|\hat{h} - \tilde{h}\|_{\mathcal{H}}^2. \quad (65)$$

In the same way, by the optimality condition in (61), we have for all $\hat{h} \in H$, and, in particular for \tilde{h}_+

$$\langle \eta \xi - (\tilde{h} - h), \tilde{h}_+ - h \rangle \geq 0. \quad (66)$$

Rearranging, we obtain

$$\langle h - \tilde{h}_+, \eta \xi \rangle \leq \langle h - \tilde{h}_+, \tilde{h} - h \rangle = -\frac{1}{2} \|h - \tilde{h}\|_{\mathcal{H}}^2 - \frac{1}{2} \|\tilde{h}_+ - h\|_{\mathcal{H}}^2 + \frac{1}{2} \|\tilde{h}_+ - \tilde{h}\|_{\mathcal{H}}^2. \quad (67)$$

Combining the last inequality with (65) and using the Fenchel inequality, we obtain

$$\langle h - \hat{h}, \eta \tilde{\xi} \rangle = \langle \tilde{h}_+ - \hat{h}, \eta \tilde{\xi} \rangle + \langle h - \tilde{h}_+, \eta \xi \rangle + \langle h - \tilde{h}_+, \eta(\tilde{\xi} - \xi) \rangle \quad (68)$$

$$\leq -\frac{1}{2} \|\tilde{h}_+ - \tilde{h}\|_{\mathcal{H}}^2 - \frac{1}{2} \|\hat{h} - \tilde{h}_+\|_{\mathcal{H}}^2 + \frac{1}{2} \|\hat{h} - \tilde{h}\|_{\mathcal{H}}^2 \quad (69)$$

$$- \frac{1}{2} \|h - \tilde{h}\|_{\mathcal{H}}^2 - \frac{1}{2} \|\tilde{h}_+ - h\|_{\mathcal{H}}^2 + \frac{1}{2} \|\tilde{h}_+ - \tilde{h}\|_{\mathcal{H}}^2 \quad (70)$$

$$+ \frac{\eta^2}{2} \|\tilde{\xi} - \xi\|_{\mathcal{H}}^2 + \frac{1}{2} \|h - \tilde{h}\|_{\mathcal{H}}^2, \quad (71)$$

which gives the result of the Lemma. \square

The second result characterizes the mirror step with respect to the measure μ .

Lemma B.4 ((Hsieh et al., 2019) [Lemma 5]). *Let $\tilde{\mu} \in \mathcal{M}$ and $\xi, \tilde{\xi} \in \mathcal{M}^*$, $\eta > 0$, and*

$$\mu = \text{Mirr}_\eta^\mu(\tilde{\mu}, \xi), \quad (72)$$

$$\tilde{\mu}_+ = \text{Mirr}_\eta^\mu(\tilde{\mu}, \tilde{\xi}). \quad (73)$$

Then, for any $\hat{\mu} \in \mathcal{M}$

$$\langle \mu - \hat{\mu}, \eta \tilde{\xi} \rangle \leq D_\Phi(\hat{\mu}, \tilde{\mu}) - D_\Phi(\hat{\mu}, \tilde{\mu}_+) + \frac{\eta^2}{8} \|\tilde{\xi} - \xi\|_{L^\infty}^2 - 2\|\mu - \tilde{\mu}\|_{TV}^2. \quad (74)$$

The following result gives the convergence rate of Algorithm 2.

Theorem B.5. *Let Assumptions B.1, B.2 hold. Let also the step sizes in Algorithm 2 satisfy $\eta_\theta = \eta_f = \eta_\mu = \eta_h = \frac{1}{16L}$, where L is defined in (60). Then, for any compact set $U = U_\theta \times U_f \times U_\mu \times U_h \subseteq \Theta \times \mathcal{H}_f \times \mathcal{M} \times H$, the sequence $(\bar{\theta}_N, \bar{f}_N, \bar{\mu}_N, \bar{h}_N)$ generated by Algorithm 2 satisfies*

$$\begin{aligned} & \max_{\mu \in U_\mu, h \in U_h} F(\bar{\theta}_N, \bar{f}_N, \mu, h) - \min_{\theta \in U_\theta, f \in U_f} F(\theta, f, \bar{\mu}_N, \bar{h}_N) \\ & \leq \frac{8L}{N} \max_{u \in U} \left(\|\theta - \tilde{\theta}_0\|_2^2 + \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + 2D_\Phi(\mu, \tilde{\mu}_0) + \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 \right). \end{aligned}$$

Before we prove the theorem, we would like to underline that the l.h.s. of the estimate is an appropriate measure of the suboptimality. This notion of the duality gap is quite standard for algorithms for saddle-point problems, see, e.g., Sect 5.3.6.1 of (Ben-Tal & Nemirovski, 2023). The motivation is that the saddle-point problem can be considered as a primal-dual pair of optimization problems $\text{Opt}(P) = \min_{\theta, f} \{\bar{F}(\theta, f) := \max_{\mu \in U_\mu, h \in U_h} F(\theta, f, \mu, h)\}$, $\text{Opt}(D) = \max_{\mu, h} \{\underline{F}(\mu, h) := \min_{\theta \in U_\theta, f \in U_f} F(\theta, f, \mu, h)\}$. By the duality we have that

$$\begin{aligned} & \bar{F}(\bar{\theta}_N, \bar{f}_N) - \text{Opt}(P) + (\text{Opt}(D) - \underline{F}(\bar{\mu}_N, \bar{h}_N)) \\ & = \max_{\mu \in U_\mu, h \in U_h} F(\bar{\theta}_N, \bar{f}_N, \mu, h) - \text{Opt}(P) + \text{Opt}(D) - \min_{\theta \in U_\theta, f \in U_f} F(\theta, f, \bar{\mu}_N, \bar{h}_N) \\ & = \max_{\mu \in U_\mu, h \in U_h} F(\bar{\theta}_N, \bar{f}_N, \mu, h) - \min_{\theta \in U_\theta, f \in U_f} F(\theta, f, \bar{\mu}_N, \bar{h}_N). \end{aligned}$$

Thus, our convergence rate implies simultaneously the convergence rate in the primal and in the dual problem.

We also underline that the algorithm does not need the set $U = U_\theta \times U_f \times U_\mu \times U_h$ as input and that it solves the original problem on the set $\Theta \times \mathcal{H}_f \times \mathcal{M} \times H$. The set U appears only in the convergence result and this can be any compact set. In particular, if U contains a neighborhood of the solution to the original problem, then the original problem is equivalent to the same problem but on the set U . Additionally, saddle-point problems can be equivalently reformulated as variational inequalities, see Example 2.2 in (Antonakopoulos et al., 2019). Then, our notion of duality gap in the l.h.s. of the bound in the above Theorem can be bounded by the restricted gap function defined in (2.5a) of (Antonakopoulos et al., 2019), see also Sect. 6 of (Stonyakin et al., 2022). Then, Lemma 1 in (Antonakopoulos et al., 2019) implies that U can be an arbitrary compact set that contains a neighborhood of the solution.

Proof. Applying Lemma B.3 to the step in θ , we obtain for any $\theta \in \Theta$ and for $k = 0, \dots, N-1$

$$\langle \theta_k - \theta, \eta_\theta F'_\theta(u_k) \rangle \leq \frac{1}{2} \|\theta - \tilde{\theta}_k\|_2^2 - \frac{1}{2} \|\theta - \tilde{\theta}_{k+1}\|_2^2 - \frac{1}{2} \|\theta_k - \tilde{\theta}_k\|_2^2 + \frac{\eta_\theta^2}{2} \|F'_\theta(u_k) - F'_\theta(\tilde{u}_k)\|_2^2. \quad (75)$$

Summing these inequalities for $k = 0, \dots, N-1$, we obtain

$$\begin{aligned} \sum_{k=0}^{N-1} \langle \theta_k - \theta, F'_\theta(u_k) \rangle &\leq \frac{1}{2\eta_\theta} \|\theta - \tilde{\theta}_0\|_2^2 + \frac{R_\theta}{2\eta_\theta}, \\ R_\theta &= \sum_{k=0}^{N-1} \left(-\|\theta_k - \tilde{\theta}_k\|_2^2 + \eta_\theta^2 \|F'_\theta(u_k) - F'_\theta(\tilde{u}_k)\|_2^2 \right). \end{aligned}$$

In the same way, we obtain for all $f \in \mathcal{H}_f$

$$\begin{aligned} \sum_{k=0}^{N-1} \langle f_k - f, F'_f(u_k) \rangle &\leq \frac{1}{2\eta_f} \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + \frac{R_f}{2\eta_f}, \\ R_f &= \sum_{k=0}^{N-1} \left(-\|f_k - \tilde{f}_k\|_{\mathcal{H}_f}^2 + \eta_f^2 \|F'_f(u_k) - F'_f(\tilde{u}_k)\|_{\mathcal{H}_f}^2 \right) \end{aligned}$$

and for all $h \in H$

$$\begin{aligned} \sum_{k=0}^{N-1} \langle h_k - h, -F'_h(u_k) \rangle &\leq \frac{1}{2\eta_h} \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 + \frac{R_h}{2\eta_h}, \\ R_h &= \sum_{k=0}^{N-1} \left(-\|h_k - \tilde{h}_k\|_{\mathcal{H}_h}^2 + \eta_h^2 \|F'_h(u_k) - F'_h(\tilde{u}_k)\|_{\mathcal{H}_h}^2 \right). \end{aligned}$$

Finally, applying Lemma B.4 to the step in μ , we obtain for any $\mu \in \mathcal{M}$

$$\begin{aligned} \sum_{k=0}^{N-1} \langle \mu_k - \mu, -F'_\mu(u_k) \rangle &\leq \frac{1}{\eta_\mu} D_\Phi(\mu, \tilde{\mu}_0) + \frac{R_\mu}{2\eta_\mu}, \\ R_\mu &= \sum_{k=0}^{N-1} \left(-\|\mu_k - \tilde{\mu}_k\|_{TV}^2 + \eta_\mu^2 \|F'_\mu(u_k) - F'_\mu(\tilde{u}_k)\|_{L^\infty}^2 \right). \end{aligned}$$

By convexity of F in (θ, f) and concavity of F in (μ, h) , we have, for all $(\theta, f) \in \Theta \times \mathcal{H}_f$

$$\begin{aligned} \frac{1}{N} \sum_{k=0}^{N-1} F(u_k) - F(\theta, f, \bar{\mu}_N, \bar{h}_N) &\leq \frac{1}{N} \sum_{k=0}^{N-1} (F(u_k) - F(\theta, f, \mu_k, h_k)) \\ &\leq \frac{1}{N} \sum_{k=0}^{N-1} (\langle \theta_k - \theta, F'_\theta(u_k) \rangle + \langle f_k - f, F'_f(u_k) \rangle) \\ &\leq \frac{1}{2N\eta_\theta} \|\theta - \tilde{\theta}_0\|_2^2 + \frac{R_\theta}{2N\eta_\theta} + \frac{1}{2N\eta_f} \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + \frac{R_f}{2N\eta_f}. \end{aligned}$$

In the same way, we obtain that, for all $(\mu, h) \in \mathcal{M} \times H$

$$\begin{aligned}
 & -\frac{1}{N} \sum_{k=0}^{N-1} F(u_k) + F(\bar{\theta}_N, \bar{f}_N, \mu, h) \leq \frac{1}{N} \sum_{k=0}^{N-1} (-F(u_k) + F(\theta_k, f_k, \mu, h)) \\
 & \leq \frac{1}{N} \sum_{k=0}^{N-1} (\langle \mu_k - \mu, -F'_\mu(u_k) \rangle + \langle h_k - h, -F'_h(u_k) \rangle) \\
 & \leq \frac{1}{N\eta_\mu} D_\Phi(\mu, \tilde{\mu}_0) + \frac{R_\mu}{2N\eta_\mu} + \frac{1}{2N\eta_h} \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 + \frac{R_h}{2N\eta_h}.
 \end{aligned}$$

Combining the last two bounds, we obtain that for all $\theta \in \Theta, f \in \mathcal{H}_f, \mu \in \mathcal{M}, h \in H$ it holds that

$$\begin{aligned}
 & F(\bar{\theta}_N, \bar{f}_N, \mu, h) - F(\theta, f, \bar{\mu}_N, \bar{h}_N) \\
 & \leq \frac{1}{2N\eta_\theta} \|\theta - \tilde{\theta}_0\|_2^2 + \frac{1}{2N\eta_f} \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + \frac{1}{N\eta_\mu} D_\Phi(\mu, \tilde{\mu}_0) + \frac{1}{2N\eta_h} \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 \\
 & + \frac{1}{2N} \left(\frac{R_\theta}{\eta_\theta} + \frac{R_f}{\eta_f} + \frac{R_\mu}{\eta_\mu} + \frac{R_h}{\eta_h} \right).
 \end{aligned}$$

Our next goal is to show that

$$\frac{R_\theta}{\eta_\theta} + \frac{R_f}{\eta_f} + \frac{R_\mu}{\eta_\mu} + \frac{R_h}{\eta_h} \leq 0.$$

Using the Lipschitz condition in Assumption B.2, we obtain

$$\begin{aligned}
 R_\theta & = \sum_{k=0}^{N-1} \left(-\|\theta_k - \tilde{\theta}_k\|_2^2 + \eta_\theta^2 \|F'_\theta(u_k) - F'_\theta(\tilde{u}_k)\|_2^2 \right) \\
 & \leq \sum_{k=0}^{N-1} \left(-\|\theta_k - \tilde{\theta}_k\|_2^2 + 4\eta_\theta^2 (L_{\theta\theta}^2 \|\theta_k - \tilde{\theta}_k\|_2^2 + L_{\theta f}^2 \|f_k - \tilde{f}_k\|_{\mathcal{H}_f}^2 + L_{\theta\mu}^2 \|\mu_k - \tilde{\mu}_k\|_{TV}^2 + L_{\theta h}^2 \|h_k - \tilde{h}_k\|_{\mathcal{H}_h}^2) \right).
 \end{aligned}$$

Combining this with the similar estimates for R_f, R_μ, R_h and rearranging the terms, we obtain

$$\begin{aligned}
 \frac{R_\theta}{\eta_\theta} + \frac{R_f}{\eta_f} + \frac{R_\mu}{\eta_\mu} + \frac{R_h}{\eta_h} & \leq \sum_{k=0}^{N-1} \left(\|\theta_k - \tilde{\theta}_k\|_2^2 (-1/\eta_\theta + 4(L_{\theta\theta}^2 \eta_\theta + L_{f\theta}^2 \eta_f + L_{\mu\theta}^2 \eta_\mu + L_{h\theta}^2 \eta_h)) \right. \\
 & \quad + \|f_k - \tilde{f}_k\|_{\mathcal{H}_f}^2 (-1/\eta_f + 4(L_{\theta f}^2 \eta_\theta + L_{f f}^2 \eta_f + L_{\mu f}^2 \eta_\mu + L_{h f}^2 \eta_h)) \\
 & \quad + \|\mu_k - \tilde{\mu}_k\|_{TV}^2 (-1/\eta_\mu + 4(L_{\theta\mu}^2 \eta_\theta + L_{f\mu}^2 \eta_f + L_{\mu\mu}^2 \eta_\mu + L_{h\mu}^2 \eta_h)) \\
 & \quad \left. + \|h_k - \tilde{h}_k\|_{\mathcal{H}_h}^2 (-1/\eta_h + 4(L_{\theta h}^2 \eta_\theta + L_{f h}^2 \eta_f + L_{\mu h}^2 \eta_\mu + L_{h h}^2 \eta_h)) \right) \leq 0,
 \end{aligned}$$

where we used that $\eta_\theta, \eta_f, \eta_\mu, \eta_h \leq \frac{1}{16L}$ for L defined in (60).

Thus, we finally obtain that for any compact $U = U_\theta \times U_f \times U_\mu \times U_h \subset \Theta \times \mathcal{H}_f \times \mathcal{M} \times H$

$$\begin{aligned}
 & \max_{\mu \in U_\mu, h \in U_h} F(\bar{\theta}_N, \bar{f}_N, \mu, h) - \min_{\theta \in U_\theta, f \in U_f} F(\theta, f, \bar{\mu}_N, \bar{h}_N) \\
 & \leq \frac{8L}{N} \max_{u \in U} \left(\|\theta - \tilde{\theta}_0\|_2^2 + \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + 2D_\Phi(\mu, \tilde{\mu}_0) + \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 \right).
 \end{aligned}$$

□

B.3 Analysis in the Stochastic Case

To account for potential inexactness in the first-order information, we assume that instead of exact derivatives, the algorithm uses their inexact counterparts $\tilde{F}'_\theta(u), \tilde{F}'_f(u), \tilde{F}'_\mu(u), \tilde{F}'_h(u)$, that may be random and are assumed to satisfy the following assumption.

Assumption B.6.

$$F'_\theta(u) = \mathbb{E}\tilde{F}'_\theta(u), \quad (76)$$

$$F'_f(u) = \mathbb{E}\tilde{F}'_f(u), \quad (77)$$

$$F'_\mu(u) = \mathbb{E}\tilde{F}'_\mu(u), \quad (78)$$

$$F'_h(u) = \mathbb{E}\tilde{F}'_h(u), \quad (79)$$

$$\mathbb{E}\|F'_\theta(u) - \tilde{F}'_\theta(u)\|_2^2 \leq \sigma_\theta^2, \quad (80)$$

$$\mathbb{E}\|F'_f(u) - \tilde{F}'_f(u)\|_{\mathcal{H}_f}^2 \leq \sigma_f^2, \quad (81)$$

$$\mathbb{E}\|F'_\mu(u) - \tilde{F}'_\mu(u)\|_{L^\infty}^2 \leq \sigma_\mu^2, \quad (82)$$

$$\mathbb{E}\|F'_h(u) - \tilde{F}'_h(u)\|_{\mathcal{H}_h}^2 \leq \sigma_h^2. \quad (83)$$

Theorem B.7. *Let Assumptions B.1, B.2, and B.6 hold. Let also in Algorithm 2 the stochastic derivatives be used instead of the deterministic and the step sizes satisfy $\eta_\theta = \eta_f = \eta_\mu = \eta_h = \frac{1}{16L}$, where L is defined in (60). Then, for any compact set $U = U_\theta \times U_f \times U_\mu \times U_h \subseteq \Theta \times \mathcal{H}_f \times \mathcal{M} \times H$, the sequence $(\bar{\theta}_N, \bar{f}_N, \bar{\mu}_N, \bar{h}_N)$ generated by Algorithm 2 satisfies*

$$\begin{aligned} & \mathbb{E} \left\{ \max_{\mu \in U_\mu, h \in U_h} F(\bar{\theta}_N, \bar{f}_N, \mu, h) - \min_{\theta \in U_\theta, f \in U_f} F(\theta, f, \bar{\mu}_N, \bar{h}_N) \right\} \\ & \leq \frac{8L}{N} \max_{u \in U} \left(\|\theta - \tilde{\theta}_0\|_2^2 + \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + 2D_\Phi(\mu, \tilde{\mu}_0) + \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 \right) + \frac{3}{16L} (\sigma_\theta^2 + \sigma_f^2 + \sigma_\mu^2 + \sigma_h^2). \end{aligned}$$

Proof. We proceed as in the proof of Theorem B.5 changing in Algorithm 2 the exact first-order information to its inexact counterpart. In this way, we obtain the following counterpart of (75)

$$\langle \theta_k - \theta, \eta_\theta \tilde{F}'_\theta(u_k) \rangle \leq \frac{1}{2} \|\theta - \tilde{\theta}_k\|_2^2 - \frac{1}{2} \|\theta - \tilde{\theta}_{k+1}\|_2^2 - \frac{1}{2} \|\theta_k - \tilde{\theta}_k\|_2^2 + \frac{\eta_\theta^2}{2} \|\tilde{F}'_\theta(u_k) - \tilde{F}'_\theta(\tilde{u}_k)\|_2^2. \quad (84)$$

Using the inequality

$$\mathbb{E}\|\tilde{F}'_\theta(u_k) - \tilde{F}'_\theta(\tilde{u}_k)\|_2^2 \leq 3\mathbb{E} \left(\|\tilde{F}'_\theta(u_k) - F'_\theta(u_k)\|_2^2 + \|\tilde{F}'_\theta(\tilde{u}_k) - F'_\theta(\tilde{u}_k)\|_2^2 + \|F'_\theta(u_k) - F'_\theta(\tilde{u}_k)\|_2^2 \right) \quad (85)$$

$$\stackrel{\text{Assumpt. B.6}}{\leq} 6\sigma_\theta^2 + 3\mathbb{E}\|F'_\theta(u_k) - F'_\theta(\tilde{u}_k)\|_2^2 \quad (86)$$

and taking the expectation in the previous inequality, we obtain the following counterpart of (75)

$$\langle \theta_k - \theta, \eta_\theta F'_\theta(u_k) \rangle \leq \frac{1}{2} \mathbb{E}\|\theta - \tilde{\theta}_k\|_2^2 - \frac{1}{2} \mathbb{E}\|\theta - \tilde{\theta}_{k+1}\|_2^2 - \frac{1}{2} \mathbb{E}\|\theta_k - \tilde{\theta}_k\|_2^2 \quad (87)$$

$$+ \frac{3\eta_\theta^2}{2} \mathbb{E}\|F'_\theta(u_k) - F'_\theta(\tilde{u}_k)\|_2^2 + 3\eta_\theta^2 \sigma_\theta^2. \quad (88)$$

Repeating the same steps as in the proof of Theorem B.5, we obtain that for any compact $U = U_\theta \times U_f \times U_\mu \times U_h \subset \Theta \times \mathcal{H}_f \times \mathcal{M} \times H$

$$\begin{aligned} & \mathbb{E} \left\{ \max_{\mu \in U_\mu, h \in U_h} F(\bar{\theta}_N, \bar{f}_N, \mu, h) - \min_{\theta \in U_\theta, f \in U_f} F(\theta, f, \bar{\mu}_N, \bar{h}_N) \right\} \\ & \leq \frac{8L}{N} \max_{u \in U} \left(\|\theta - \tilde{\theta}_0\|_2^2 + \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + 2D_\Phi(\mu, \tilde{\mu}_0) + \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 \right) + \frac{3}{16L} (\sigma_\theta^2 + \sigma_f^2 + \sigma_\mu^2 + \sigma_h^2). \quad (89) \end{aligned}$$

□

Let us denote $\sigma^2 = \sigma_\theta^2 + \sigma_f^2 + \sigma_\mu^2 + \sigma_h^2$. As we see, Theorem B.7 guarantees the same convergence rate as in the exact case, but up to some vicinity which is governed by the level of noise. In most cases, the σ^2/L term can be made of the same order $1/N$ by using mini-batching technique. Indeed, a mini-batch of size N allows to change

the variance from σ^2 to σ^2/N . Yet, we note that in this case, N iterations will require the number of samples $O(N^2)$.

An alternative would be to use the information about the diameter of the set U . Indeed, assume that

$$\max_{u \in U} \left(\|\theta - \tilde{\theta}_0\|_2^2 + \|f - \tilde{f}_0\|_{\mathcal{H}_f}^2 + 2D_\Phi(\mu, \tilde{\mu}_0) + \|h - \tilde{h}_0\|_{\mathcal{H}_h}^2 \right) \leq \Omega_U^2.$$

Then, we obtain the following counterpart of the r.h.s. of (89) substituting $\eta_\theta = \eta_f = \eta_\mu = \eta_h = \eta$

$$\frac{\Omega_U^2}{2N\eta} + 3\sigma^2\eta.$$

Fixing the number of steps N and choosing

$$\eta_\theta = \eta_f = \eta_\mu = \eta_h = \eta = \min \left\{ \frac{1}{16L}, \frac{\Omega_U\sigma}{\sqrt{6N}} \right\},$$

we obtain the following result

$$\begin{aligned} & \mathbb{E} \left\{ \max_{\mu \in U_\mu, h \in U_h} F(\bar{\theta}_N, \bar{f}_N, \mu, h) - \min_{\theta \in U_\theta, f \in U_f} F(\theta, f, \bar{\mu}_N, \bar{h}_N) \right\} \\ & \leq \max \left\{ \frac{8L\Omega_U^2}{N}, \sqrt{\frac{3\sigma^2\Omega_U^2}{2N}} \right\}. \end{aligned} \quad (90)$$

C DERIVATIONS OF RESULTS FOR DISTRIBUTIONALLY ROBUST OPTIMIZATION

C.1 Set-up and Adaptation of the General KMP Algorithm

In this subsection we particularize the elements of Algorithm 2 for the specific DRO problem (12). We choose $\mathcal{H}_f = \mathcal{H}_h = \mathcal{H}$ to be a reproducing kernel Hilbert space with kernel k .

Our main assumptions for this problem are

- l is convex w.r.t. θ .
- $L_0 = \sup_{x, \theta} \|\nabla_\theta l(\theta; x)\|_2 < +\infty$.
- $\nabla_\theta l(\theta; x)$ is $L(x)$ -Lipschitz w.r.t. θ and $L_1 = \sup_\mu \mathbb{E}_{x \sim \mu} L(x)^2 < +\infty$.
- $C = \sup_x k(x, x) < +\infty$.

Clearly, then the objective F is convex in (θ, f) for fixed (μ, h) and concave in (μ, h) for fixed (θ, f) .

The Frechet derivatives of F with respect to the variables (θ, f, μ, h) are given by

$$F'_\theta = \mathbb{E}_{x \sim \mu} \nabla_\theta l(\theta; x) \quad (91)$$

$$F'_f = \int k(x, x') d\hat{\mu}(x') + \epsilon h(x) - \int k(x, x') d\mu(x') = \mathbb{E}_{x \sim \hat{\mu}} k(\cdot, x) + \epsilon h(\cdot) - \mathbb{E}_{x \sim \mu} k(\cdot, x) \quad (92)$$

$$-F'_\mu = f(\cdot) - l(\theta; \cdot) \quad (93)$$

$$-F'_h = -\epsilon f(\cdot). \quad (94)$$

Since the derivatives w.r.t. θ and f have the form of expectation, we can use the following stochastic counterparts. We can take a sample of X_i 's from μ to construct an unbiased stochastic derivative

$$\tilde{F}'_\theta = \frac{1}{N_\theta} \sum_{i=1}^{N_\theta} \nabla_\theta l(\theta; X_i). \quad (95)$$

Similarly, we can take a sample of X_i 's from μ and \hat{X}_i from $\hat{\mu}$ to construct an unbiased stochastic derivative

$$\tilde{F}'_f = \epsilon h(\cdot) + \frac{1}{N_f} \sum_{i=1}^{N_f} (k(\cdot, \hat{X}_i) + k(\cdot, X_i)). \quad (96)$$

The Lipschitz constants of the derivatives are estimated in the following way. The derivative F'_θ depends only on μ and θ . Thus, $L_{\theta f} = L_{\theta h} = 0$. Further, we have

$$\|\mathbb{E}_{x \sim \mu} \nabla_\theta l(\theta_1; x) - \mathbb{E}_{x \sim \mu} \nabla_\theta l(\theta_2; x)\|_2 \leq \mathbb{E}_{x \sim \mu} L(x) \|\theta_1 - \theta_2\|_2 \quad (97)$$

and $L_{\theta\theta} = L_1$.

$$\|\mathbb{E}_{x \sim \mu_1} \nabla_\theta l(\theta; x) - \mathbb{E}_{x \sim \mu_2} \nabla_\theta l(\theta; x)\|_2 \leq L_0 \|\mu_1 - \mu_2\|_{TV} \quad (98)$$

and $L_{\theta\mu} = L_0$.

The derivative F'_f depends only on μ and h . Thus, $L_{ff} = L_{f\theta} = 0$. Further, we have

$$\|\mathbb{E}_{x \sim \hat{\mu}} k(\cdot, x) + \epsilon h_1(\cdot) + \mathbb{E}_{x \sim \mu_1} k(\cdot, x) - (\mathbb{E}_{x \sim \hat{\mu}} k(\cdot, x) + \epsilon h_2(\cdot) + \mathbb{E}_{x \sim \mu_2} k(\cdot, x))\|_{\mathcal{H}} \quad (99)$$

$$\leq \epsilon \|h_1 - h_2\|_{\mathcal{H}} + \|\mathbb{E}_{x \sim \mu_1} k(\cdot, x) - \mathbb{E}_{x \sim \mu_2} k(\cdot, x)\|_{\mathcal{H}} \leq \epsilon \|h_1 - h_2\|_{\mathcal{H}} + \sqrt{C} \|\mu_1 - \mu_2\|_{TV}, \quad (100)$$

i.e., $L_{f\mu} = \sqrt{C}$, $L_{fh} = \epsilon$. Here we used that

$$\|\mathbb{E}_{x \sim \mu_1} k(\cdot, x) - \mathbb{E}_{x \sim \mu_2} k(\cdot, x)\|_{\mathcal{H}} \leq \sqrt{C} \|\mu_1 - \mu_2\|_{TV}.$$

The derivative F'_μ depends only on f and θ . Thus, $L_{\mu\mu} = L_{\mu h} = 0$. Further, we have

$$\| -f_1(\cdot) + l(\theta_1; \cdot) - (-f_2(\cdot) + l(\theta_2; \cdot)) \|_{L^\infty} \quad (101)$$

$$\leq \sqrt{C} \|f_1 - f_2\|_{\mathcal{H}} + L_0 \|\theta_1 - \theta_2\|_{\mathcal{H}}, \quad (102)$$

i.e., $L_{f\mu} = \sqrt{C}$, $L_{f\theta} = L_0$. Here we used that

$$\begin{aligned} \|f_2(\cdot) - f_1(\cdot)\|_{L^\infty} &= \sup_x |f_2(x) - f_1(x)| = \sup_x \langle f_2 - f_1, \phi(x) \rangle_{\mathcal{H}} \\ &\leq \|f_2 - f_1\|_{\mathcal{H}} \cdot \sup_x \|\phi(x)\|_{\mathcal{H}} \leq \sqrt{C} \cdot \|f_2 - f_1\|_{\mathcal{H}}. \end{aligned} \quad (103)$$

Finally, the derivative F'_h depends only on f . Thus, $L_{h\theta} = L_{h\mu} = L_{hh} = 0$. Further, we have

$$\| -\epsilon f_1(\cdot) - (-\epsilon f_2(\cdot)) \|_{\mathcal{H}} \leq \epsilon \|f_1 - f_2\|_{\mathcal{H}}. \quad (104)$$

Thus, $L_{hf} = \epsilon$.

As we see, our main assumptions in Theorem B.5 hold for the DRO problem (12). Moreover, stochastic derivatives in (95) and (96) satisfy Assumption B.6. Indeed, we have

$$\mathbb{E}_{X \sim \mu} \|\mathbb{E}_{x \sim \mu} \nabla_\theta l(\theta; x) - \nabla_\theta l(\theta; X)\|_2^2 \leq \mathbb{E}_{X \sim \mu} \|\nabla_\theta l(\theta; X)\|_2^2 \leq L_0^2, \quad (105)$$

$$\mathbb{E}_{\hat{X} \sim \hat{\mu}, X \sim \mu} \left\| \epsilon h(\cdot) + k(\cdot, \hat{X}) + k(\cdot, X) - (\mathbb{E}_{x \sim \hat{\mu}} k(\cdot, x) + \epsilon h(\cdot) - \mathbb{E}_{x \sim \mu} k(\cdot, x)) \right\|_{\mathcal{H}}^2 \quad (106)$$

$$\leq 2\mathbb{E}_{X \sim \mu} \|k(\cdot, X)\|_{\mathcal{H}}^2 + 2\mathbb{E}_{\hat{X} \sim \hat{\mu}} \|k(\cdot, \hat{X})\|_{\mathcal{H}}^2 \leq 4C. \quad (107)$$

This allows us to apply also Theorem B.7 to the DRO problem (12). This leads to the following lemma, which summarizes the applicability of our analysis to the DRO setting.

Lemma C.1. *Assumptions B.2, B.6 hold for the smoothed DRO problem (12). Consequently, the results of Theorems B.5 and B.7 hold for the DRO problem (12).*

C.2 Proof of Guarantees for Distributionally Robust Optimization

For notational conciseness, we denote the DRO risk associated with the decision θ as

$$Z_{\epsilon, \hat{P}_n}(\theta) := \sup_{\mu: \mathcal{D}(\mu, \hat{P}_n) \leq \epsilon} \mathbb{E}_\mu[l(\theta; x)]. \quad (108)$$

Note that $Z_{\epsilon, \hat{P}_n}(\theta)$ is a random variable since \hat{P}_n is a random sample from the generating distribution. Let us denote the ideal DRO decision

$$\theta_{\text{DRO}}^* := \arg \inf_{\theta} \sup_{\mu: \mathcal{D}(\mu, \hat{P}_n) \leq \epsilon} \mathbb{E}_\mu[l(\theta; x)],$$

which is often not realizable in practice.

C.2.1 Proof of Proposition 5.2

We restate the results and provide more detailed statements.

Proposition C.2 (DRO Guarantee for KMP decision sub-optimality). *Suppose $\bar{\theta}_N = \frac{1}{N} \sum_{k=0}^{N-1} \theta_k$ is the averaged solution produced by the KMP algorithm after N steps. Then, $\forall \epsilon > 0$, the DRO risk associated with the decision $\bar{\theta}_N$ is bounded by*

$$\sup_{\text{MMD}(\mu, \hat{P}_n) \leq \epsilon} \mathbb{E}_\mu l(\bar{\theta}_N; x) - \underbrace{\inf_{\theta \in \Theta} \sup_{\text{MMD}(\mu, \hat{P}_n) \leq \epsilon} \mathbb{E}_\mu l(\theta; x)}_{(\text{Optimal DRO risk})} \leq O\left(\frac{1}{N}\right).$$

In the stochastic setting, we find

$$\mathbb{E} \left[\sup_{\text{MMD}(\mu, \hat{P}_n) \leq \epsilon} \mathbb{E}_\mu l(\bar{\theta}_N; x) - \underbrace{\inf_{\theta \in \Theta} \sup_{\text{MMD}(\mu, \hat{P}_n) \leq \epsilon} \mathbb{E}_\mu l(\theta; x)}_{(\text{Optimal DRO risk})} \right] \leq O\left(\frac{1}{\sqrt{N}}\right).$$

Proof. We note the relation

$$\begin{aligned} Z_{\epsilon, \hat{P}_n}(\theta_{\text{DRO}}^*) &= \inf_{\theta \in \Theta, f \in \mathcal{H}} \sup_{\mu \in \mathcal{M}, h \in \mathcal{H}: \|h\|_{\mathcal{H}} \leq 1} F(\theta, f, \mu, h) \\ &\geq \inf_{\theta \in \Theta, f \in \mathcal{H}} F(\theta, f, \bar{\mu}_N, \bar{h}_N) \end{aligned}$$

and

$$\begin{aligned} Z_{\epsilon, \hat{P}_n}(\bar{\theta}_N) &= \inf_{f \in \mathcal{H}} \sup_{\mu \in \mathcal{M}, h \in \mathcal{H}: \|h\|_{\mathcal{H}} \leq 1} F(\bar{\theta}_N, f, \mu, h) \\ &\leq \sup_{\mu \in \mathcal{M}, h \in \mathcal{H}: \|h\|_{\mathcal{H}} \leq 1} F(\bar{\theta}_N, \bar{f}_N, \mu, h). \end{aligned}$$

Taking the difference of the above expressions,

$$Z_{\epsilon, \hat{P}_n}(\bar{\theta}_N) - Z_{\epsilon, \hat{P}_n}(\theta_{\text{DRO}}^*) \leq \sup_{\mu \in \mathcal{M}, h \in \mathcal{H}: \|h\|_{\mathcal{H}} \leq 1} F(\bar{\theta}_N, \bar{f}_N, \mu, h) - \inf_{\theta \in \Theta, f \in \mathcal{H}} F(\theta, f, \bar{\mu}_N, \bar{h}_N). \quad (109)$$

Applying the statement of Theorem B.5 completes the proof for the deterministic case. The stochastic case follows from Theorem B.7. \square

C.2.2 Proof of Corollary 5.3

Corollary C.3 (Robustness guarantee for population risk). *Suppose the ambiguity level ϵ is chosen such that $\epsilon > \epsilon_n(\delta)$. Then, with probability $1 - \delta$, the population risk of the decision $\bar{\theta}_N$ output by the KMP algorithm after N steps is estimated from above by*

$$\mathbb{E}_{\mu_0}[l(\bar{\theta}_N; x)] - \inf_{\theta \in \Theta} \sup_{\text{MMD}(\mu, \hat{P}_n) \leq \epsilon} \mathbb{E}_\mu l(\theta; x) \leq O\left(\frac{1}{N}\right).$$

The bound is $O(\frac{1}{\sqrt{N}})$ in expectation in the stochastic case.

Proof. With probability $1 - \delta$, the empirical estimation error of MMD can be upper bounded by (Tolstikhin et al., 2017)

$$\text{MMD}(\widehat{P}_n, \mu_0) \leq \epsilon_n = \sqrt{\frac{C}{n}} + \sqrt{\frac{2C \log(1/\delta)}{n}}. \quad (110)$$

Since, $\epsilon > \epsilon_n(\delta)$, we trivially have

$$\mathbb{E}_{\mu_0}[l(\bar{\theta}_N; x)] \leq Z_{\epsilon, \widehat{P}_n}(\bar{\theta}_N). \quad (111)$$

Applying Proposition 5.2, we find the desired estimate

$$\mathbb{E}_{\mu_0}[l(\bar{\theta}_N; x)] \leq Z_{\epsilon, \widehat{P}_n}(\bar{\theta}_N) \leq Z_{\epsilon, \widehat{P}_n}(\theta_{\text{DRO}}^*) + O\left(\frac{1}{N}\right). \quad (112)$$

□

C.2.3 Proof of Corollary 5.4

Corollary C.4 (Robustness guarantee under distribution shift). *The risk under the worst-case distribution shift from the population distribution μ_0 , of any radius $\epsilon > 0$, of the decision θ_N output by the KMP algorithm after N steps is upper-bounded with large probability by*

$$\sup_{\text{MMD}(\mu, \mu_0) \leq \epsilon} \mathbb{E}_{\mu} l(\bar{\theta}_N; x) - \inf_{\theta \in \Theta} \sup_{\text{MMD}(\mu, \widehat{P}_n) \leq \epsilon} \mathbb{E}_{\mu} l(\theta; x) \leq O\left(\max\left\{\frac{1}{N}, \frac{1}{\sqrt{n}}\right\}\right).$$

In the stochastic setting, we find

$$\mathbb{E}\left[\sup_{\text{MMD}(\mu, \mu_0) \leq \epsilon} \mathbb{E}_{\mu} l(\bar{\theta}_N; x) - \inf_{\theta \in \Theta} \sup_{\text{MMD}(\mu, \widehat{P}_n) \leq \epsilon} \mathbb{E}_{\mu} l(\theta; x)\right] \leq O\left(\max\left\{\frac{1}{\sqrt{N}}, \frac{1}{\sqrt{n}}\right\}\right).$$

Proof. The starting point of the proof is the decomposition

$$Z_{\epsilon, \mu_0}(\bar{\theta}_N) - Z_{\epsilon, \widehat{P}_n}(\theta_{\text{DRO}}^*) = Z_{\epsilon, \mu_0}(\bar{\theta}_N) - Z_{\epsilon, \widehat{P}_n}(\bar{\theta}_N) + Z_{\epsilon, \widehat{P}_n}(\bar{\theta}_N) - Z_{\epsilon, \widehat{P}_n}(\theta_{\text{DRO}}^*). \quad (113)$$

We already have the upper bound for the last two terms due to Proposition 5.2. We now estimate the first two terms by standard concentration inequalities and the uniform law of large numbers; see, e.g., (Boucheron et al., 2005; Wainwright, 2019). First, for any θ in the domain and RKHS function $f \in \mathcal{H}$, with probability at least $1 - \delta$,

$$\int f d\mu_0 \leq \int f d\widehat{P}_n + O\left(\sqrt{\frac{\log(1/\delta)}{n}}\right).$$

We consider the standard setting of an RKHS of functions with finite norms. Using uniform convergence, the term $O\left(\sqrt{\frac{\log(1/\delta)}{n}}\right)$ does not depend on the specific function f within the class \mathcal{F} , and the term is finite. Taking infimum on both sides,

$$\begin{aligned} \inf_f \left\{ \int f d\mu_0 + \epsilon \|f\|_{\mathcal{H}} + \int (l(\theta; \cdot) - f) d\mu \right\} \leq \\ \inf_f \left\{ \int f d\widehat{P}_n + \epsilon \|f\|_{\mathcal{H}} + \int (l(\theta; \cdot) - f) d\mu \right\} + O\left(\sqrt{\frac{\log(1/\delta)}{n}}\right). \end{aligned} \quad (114)$$

We use the shorthand notation

$$\tilde{\mu}_0 := \arg \sup_{\mu} \inf_f \int f d\mu_0 + \epsilon \|f\|_{\mathcal{H}} + \int (l(\theta; \cdot) - f) d\mu. \quad (115)$$

Then, by (strong) linear duality,

$$\begin{aligned}
 Z_{\epsilon, \mu_0}(\theta) &= \sup_{\mu} \inf_f \left\{ \int f d\mu_0 + \epsilon \|f\|_{\mathcal{H}} + \int (l(\theta; \cdot) - f) d\mu \right\} \\
 &\leq \inf_f \left\{ \int f d\widehat{P}_n + \epsilon \|f\|_{\mathcal{H}} + \int (l(\theta; \cdot) - f) d\tilde{\mu}_0 \right\} + O\left(\sqrt{\frac{\log(1/\delta)}{n}}\right) \\
 &\leq \sup_{\mu} \inf_f \left\{ \int f d\widehat{P}_n + \epsilon \|f\|_{\mathcal{H}} + \int (l(\theta; \cdot) - f) d\mu \right\} + O\left(\sqrt{\frac{\log(1/\delta)}{n}}\right) \\
 &= Z_{\epsilon, \widehat{P}_n}(\theta) + O\left(\sqrt{\frac{\log(1/\delta)}{n}}\right) \quad (116)
 \end{aligned}$$

holds with probability at least $1 - \delta$.

In summary, plugging in $\theta = \bar{\theta}_N$,

$$\mathbb{P}\left(Z_{\epsilon, \mu_0}(\bar{\theta}_N) - Z_{\epsilon, \widehat{P}_n}(\bar{\theta}_N) \geq O\left(\sqrt{\frac{\log(1/\delta)}{n}}\right)\right) \leq \delta.$$

Furthermore, to derive the expected bound, we use the Gaussian integral to integrate the tail bound, combining with (113) and Proposition 5.2. We find the desired statement in expectation

$$\mathbb{E}\left[Z_{\epsilon, \mu_0}(\bar{\theta}_N) - Z_{\epsilon, \widehat{P}_n}(\theta_{\text{DRO}}^*)\right] \leq O\left(\max\left\{\frac{1}{N}, \frac{1}{\sqrt{n}}\right\}\right). \quad (117)$$

The stochastic case follows from the previous settings. □