
Randomized Stochastic Gradient Descent Ascent

Othmane Sebbouh
ENS, PSL, CNRS
CREST-ENSAE

Marco Cuturi
CREST-ENSAE

Gabriel Peyré
ENS, PSL, CNRS

Abstract

An increasing number of machine learning problems, such as robust or adversarial variants of existing algorithms, require minimizing a loss function that is itself defined as a maximum. Carrying a loop of stochastic gradient ascent (SGA) steps on the (inner) maximization problem, followed by an SGD step on the (outer) minimization, is known as Epoch Stochastic Gradient *Descent Ascent* (ESGDA). While successful in practice, the theoretical analysis of ESGDA remains challenging, with no clear guidance on choices for the inner loop size nor on the interplay between inner/outer step sizes. We propose RSGDA (Randomized SGDA), a variant of ESGDA with stochastic loop size with a simpler theoretical analysis. RSGDA comes with the first (among SGDA algorithms) almost sure convergence rates when used on nonconvex min/strongly-concave max settings. RSGDA can be parameterized using optimal loop sizes that guarantee the best convergence rates known to hold for SGDA. We test RSGDA on toy and larger scale problems, using distributionally robust optimization and single-cell data matching using optimal transport as a testbed.

1 INTRODUCTION

Min-Max problems in ML. Consider the following stochastic min-max optimization problem:

$$\min_{\theta \in \mathbb{R}^d} \phi(\theta), \quad \text{where } \phi(\theta) \stackrel{\text{def}}{=} \max_{v \in \mathcal{V}} F(\theta, v), \quad (1)$$
$$\text{and } F(\theta, v) = \mathbb{E}_z [f(\theta, v; z)],$$

with $\mathcal{V} \subseteq \mathbb{R}^n$. Problems such as (1) appear in the estimation of generative models (Goodfellow et al., 2014), reinforcement learning (Dai et al., 2018), online learning (Cesa-Bianchi and Lugosi, 2006), and in many other domains including mathematics and economics (Von Neumann and Morgenstern, 2007, Başar and Olsder, 1998) (see Nouiehed et al. (2019a) and references therein).

Gradient Descent-Ascent algorithms. Most machine learning applications resort to stochastic gradient methods to solve (1). These methods, which we refer to as *Gradient Descent Ascent (GDA) algorithms*, alternate between possibly many ascent steps in v to increase F , with a descent step along a stochastic gradient direction to decrease ϕ . Within that space alone, several algorithms have been proposed, some of them taking advantage of assumptions on the properties of F . The setting that has received the most attention is by far that where F is convex-concave (Sion, 1958, Korpelevich, 1976, Nemirovski, 2004, Nedić and Ozdaglar, 2009, Azizian et al., 2020). Only recently has the application of GDA to non-convex functions been thoroughly analyzed from a theoretical perspective, resulting in a flurry of works (Namkoong and Duchi, 2016, Sinha et al., 2018, Rafique et al., 2018, Grnarova et al., 2017, Lu et al., 2020, Nouiehed et al., 2019b, Thekumparampil et al., 2019, Kong and Monteiro, 2019, Jin et al., 2020, Zhang et al., 2020). These recent works are particularly relevant to modern machine learning problems, where the *min* problem in (1) often requires optimizing the parameters of models using a non-convex loss ϕ .

Non-convex strongly concave optimization. We consider in this work the setting where F is non-convex in θ , yet strongly concave and smooth in v . This setting is practically relevant since it appears in Temporal Difference learning (Dai et al., 2018), robust optimization (Sinha et al., 2018), or entropic optimal transport (Cuturi, 2013). It is also theoretically appealing, since it has been shown that GDA algorithms achieve a complexity similar to single-variable minimization, up to a factor depending on the conditioning of the problem (Sanjabi et al., 2018b, Nouiehed et al., 2019b, Lin et al., 2019, Qiu et al., 2020, Huang et al., 2020). For example, compared to the nonconvex-concave and nonconvex-nonconcave settings, one does not need to assume that the step size in (1) is bounded in order to ensure convergence.

In practice, one of the most widely used algorithms to solve (1) is epoch stochastic gradient descent ascent (Goodfellow et al., 2014, Sanjabi et al., 2018a, Lin et al., 2019, Jin et al., 2020, Nouiehed et al., 2019b, Sanjabi et al., 2018b, Houdard et al., 2021), where we make an arbitrary fixed number of stochastic gradient ascent steps followed by a gradient descent step. Unfortunately, there is very little understanding of the theoretical justifications behind this method, leaving practitioners in the dark as to what guarantees they might expect from their parameter settings.

Our Contributions. In this work, we aim to quantify in the form of convergence rates and suggested parameter settings how much the practical choices are justified by theory. To this end, we propose Randomized Stochastic Gradient Descent Ascent (RSGDA), a randomization of Epoch Stochastic Gradient descent ascent (ESGDA) which is more amenable to theoretical analysis and performs very similarly to ESGDA in practice. From a theoretical perspective, we show that RSGDA enjoys the best known convergence rates, which are also verified by one-step stochastic gradient descent ascent (Lin et al., 2019) and other variants (Huang et al., 2020, Qiu et al., 2020), and we demonstrate how the step sizes and number of gradient ascent steps should be set in order to retain these convergence guarantees. We evaluate our suggested parameter settings on problems from robust optimization and optimal transport: we consider (i) the problem of distributionally robust optimization (Sinha et al., 2018), where we aim to learn a classifier which is robust to adversarial inputs, and (ii) the problem of single-cell data matching using regularized optimal transport (Schiebinger et al., 2019, Stark et al., 2020, Cuturi, 2013).

2 BACKGROUND

We present in this section relevant assumptions for analysis, and briefly review state of the art results.

2.1 Assumptions

Smoothness and strong concavity. Throughout the paper, we assume that F is smooth in both variables, and strongly concave in the second variable.

Assumption 2.1. We assume that F is L -smooth on \mathbb{R}^d and μ -strongly concave on V .

The smoothness and strong concavity of F ensure that the function defined in (1) is smooth as well.

Lemma 2.2 (Lemma 4.3 in Lin et al. (2019)). Let Assumption 2.1 hold. Let $\theta^* \stackrel{\text{def}}{=} \arg\max_{\theta \in \mathbb{R}^d} F(\theta; v)$ for all $v \in V$. Then, θ^* is μ -lipschitz and F is $2L$ -smooth.

Assumptions on noise. We assume that the stochastic gradient in (1) has bounded variance, which is standard in non-convex stochastic optimization.

Assumption 2.3. There exists $\sigma^2 > 0$ such that for all $(\theta; v) \in \mathbb{R}^d \times V$,

$$\mathbb{E}_z \|\text{grad}_\theta f(\theta; v; z) - \text{grad}_\theta F(\theta; v)\|^2 \leq \sigma^2$$

For the stochastic gradient in v , we only assume that the noise is finite at the maximizer.

Assumption 2.4. Define $v^*(\theta) \stackrel{\text{def}}{=} \arg\max_{v \in V} f(\theta; v)$. Let $\sigma^2 \stackrel{\text{def}}{=} \mathbb{E}_z \|\text{grad}_v f(\theta; v^*(\theta); z)\|^2 < 1$.

2.2 Review of Gradient Descent Ascent Methods

Definitions. Let $\epsilon > 0$. We call $\theta \in \mathbb{R}^d$ an ϵ -approximate stationary point if $\|\text{grad}_\theta F(\theta)\| \leq \epsilon$. We call $v \in V$ a ϵ -approximate maximizer if for some $\theta \in \mathbb{R}^d$ we have $F(\theta; v) \geq \max_{v' \in V} F(\theta; v') - \epsilon$.

SGDmax, SGDA and ESGDA. One of the most analyzed algorithms for solving min-max problems such as (1) is (S)GDmax (Sanjabi et al., 2018a, Lin et al., 2019, Jin et al., 2020, Nouiehed et al., 2019b, Sanjabi et al., 2018b, Houdard et al., 2021), where at each iteration we make the number of (stochastic) gradient ascent steps necessary to reach an approximate maximizer v_{k+1} , before making a descent step using the gradient $\text{grad}_\theta f(\theta_k; v_k)$. See Alg. 1, SGDmax.

To reach an ϵ -stationary point v_k , this algorithm requires $O(\log(1/\epsilon)^2)$ (resp. $O(\log(1/\epsilon)^3)$) total gradient computations in the deterministic (resp. stochastic) setting (Lin et al., 2019). In practice, however, because SGDmax involves a subroutine where we need to ensure that we reach an approximate maximizer, this algorithm is rarely implemented. It is instead approximated by ESGDA.

Algorithm 1 SGDmax/SGDA/ESGDA

Inputs: step sizes α and β , loop size m , max-oracle accuracy ϵ

for $k = 0; 1; 2; \dots;$ do

SGDmax :

Find v_{k+1} s.t. $E_k [F(v_k; v_{k+1})] \leq \epsilon_k + \beta$

SGDA :

Sample $z_k^0 \sim D$

$v_{k+1} = v_k + \alpha \nabla v f(v_k; z_k^0)$

ESGDA :

for $t = 0; 1; 2; \dots; m - 1$ do

Sample $z_k^t \sim D$

$v_k^{t+1} = v_k^t + \alpha \nabla v f(v_k^t; z_k^t)$

end for

$v_{k+1} = v_{k+1}^0 = v_k^m$

Sample $z_k \sim D$

$v_{k+1} = v_k + \alpha \nabla v f(v_k; z_k)$

end for

At the other end of the spectrum is one-step Gradient Descent Ascent (Lin et al., 2019, Chen et al., 2020) (also referred to as GDA in the literature), in which one ascent step is followed by one descent step. See Alg. 1, SGDA. This algorithm has two important advantages compared with (S)GDmax: (a) it is simple: it doesn't require any inner ascent loop or stopping criteria, (b) it has better convergence rates: $O(\log(1/\epsilon)^2)$ (resp. $O(\log(1/\epsilon)^3)$) in the deterministic (resp. stochastic) setting.

In practice, instead of choosing a precision or running SGDA, one popular choice (Goodfellow et al., 2014, Sinha et al., 2018, Houdard et al., 2021) is using epoch stochastic gradient descent ascent (ESGDA), where we make a fixed number of ascent steps m followed by a descent step on v . See Alg. 1, ESGDA. The goal of the ascent steps is to have a good enough approximation of $v(v_k)$, and hence of the gradient $\nabla v f(v_k)$, in order to make a descent step on v . Despite its popularity, we know little about the theoretical properties of ESGDA in the nonconvex-strongly concave setting. Yan et al. (2020) studies a version of ESGDA with an iteration dependent number of gradient ascent steps. The problem they consider F

weakly convex and strongly concave is inherently harder than the smooth nonconvex-strongly concave problem. The complexity of their method for finding a nearly stationary point is $O(\log(1/\epsilon)^4)$, but to reach a stationary point, they still need $O(\log(1/\epsilon)^6)$ iterations, even when their results are specialized to the smooth setting (See Prop. 4.11 in Lin et al. (2019) for the relation between stationarity and near stationarity). Chen et al. (2021) consider the harder problem of stochastic nested optimization (which includes bilevel and min-max optimization). They devise an algorithm (ALSET) with $O(\log(1/\epsilon)^4)$ complexity when the epoch size is m . But they assume that their function and its hessian are both lipschitz, whereas the analysis of SGDA only requires that the gradients are lipschitz.

Despite the popularity of ESGDA, to the best of our knowledge, there is no theoretical analysis ensuring that this algorithm converges under the same conditions as SGDA and SGDmax. In this work, we aim to support practical implementation choices of ESGDA with solid theory, and in turn suggest new parameter settings to further improve how ESGDA is implemented.

Motivating RSGDA. RSGDA is a randomized version of ESGDA with a stochastic loop size. Empirically, RSGDA performs similarly to ESGDA (Fig. 1), but its theoretical analysis is simpler, thanks to its inner-loop free structure. This is showcased in Prop. 3.1, our central result, a descent inequality (up to additional noise terms), from which we draw several conclusions. First, we derive almost sure convergence rates for RSGDA (Cor. 3.2). Then, we show that, like SGDA, RSGDA enjoys the best known convergence rates in expectation among stochastic gradient descent ascent algorithms (Cor. 3.8). Importantly, for each of our convergence results, we determine (i) a range of descent step probabilities p that guarantees the best possible convergence rate, and (ii) how the step sizes should be set depending on p .

3 RANDOMIZED STOCHASTIC GRADIENT DESCENT ASCENT

We now introduce RSGDA (Alg. 2) in detail. At each iteration k , we toss a coin: if it lands heads (with probability $1 - p$), we keep v_k fixed and make an ascent step on v_k along the stochastic gradient $\nabla v f(v_k; z_k)$, where $z_k \sim D$ (or a projected ascent step if $V \notin \mathbb{R}^n$); if it lands tails (with probability p), we keep v_k fixed and make a descent step along the

stochastic gradient $r_k = \nabla f_k(v_k; z_k)$. The algorithm can also be seen as a version of ESGDA where the size of the inner loop is stochastic and equal to $1 = p$ in expectation.

Algorithm 2 Randomized SGDA

Inputs: step sizes $(\alpha_k)_k$ and $(\beta_k)_k$, $p \in (0; 1)$.
 Initialisation: $v_0 \in \mathbb{R}^n$; $z_0 \in \mathbb{R}^d$
 for $k = 0; 1; 2; \dots$; do
 Sample $z_k \in D$
 $v_k^+ = v_k + \alpha_k \nabla f(v_k; z_k)$
 $v_k^* = v_k + \beta_k \nabla f(v_k; z_k)$
 $(v_{k+1}; z_{k+1}) = \begin{cases} (v_k^+; z_k) & \text{w. p. } p \\ (v_k^*; z_k) & \text{w. p. } 1 - p \end{cases}$
 end for

From a stochastic optimization perspective, this randomization trick is reminiscent of the way Loopless SVRG (Stochastic Variance Reduced Gradient) (Hofmann et al., 2015, Kovalev et al., 2020) avoids using the inner loop of the original SVRG method (Johnson and Zhang, 2013). Like for Loopless SVRG and SVRG, RSGDA results in a much simpler analysis than for ESGDA.

This is apparent through Prop. 3.1, in which we establish a one-step recurrence inequality which is central to deriving our convergence results.

Proposition 3.1. Consider the iterates of Alg. 2. Let Assumption 2.1 hold. Let $p \in (0; 1)$ and define for all $k \geq N$,

$$D_k = \alpha_k \min_{z \in D} \|\nabla f(v_k; z)\|^2; \quad r_k = \alpha_k \|\nabla f(v_k; z_k)\|^2;$$

$$\text{and } E_k = D_k + L \frac{p \alpha_k}{(1 - p) \alpha_k} r_k;$$

Let $(\alpha_k)_k$ and $(\beta_k)_k$ be two positive decreasing sequences such that $\frac{\alpha_k}{k}$ is decreasing as well, with

$$\alpha_k \leq \frac{1}{2L} \text{ and } \beta_k \leq \frac{p(1-p)\alpha_k}{4 + p(2p+(1-p)\alpha_k)}.$$

$$p \alpha_k \|\nabla f(v_k; z_k)\|^2 + 2E_k [E_{k+1}] \leq 2E_k + 4 \alpha_k p \alpha_k L \alpha_k^{-2} + 2 \alpha_k^2 p \frac{2}{k} L + \frac{p^2 (2p+(1-p)\alpha_k)^3}{(1-p)^2 k^2} \alpha_k^4; \quad (2)$$

Almost sure convergence rates. Establishing the recurrence inequality of Prop. 3.1 is key to deriving almost sure convergence results in stochastic optimization (Bertsekas and Tsitsiklis, 2000, Gadat et al., 2018, Sebbouh et al., 2021), where one can directly

Figure 1: We use RSGDA (Alg. 2) and ESGDA (Alg. 1) on the minmax problem described in §4.1. This figure shows that ESGDA with a given loop size m and RSGDA with $p = 1/(m+1)$ perform similarly. The objective is defined in (5).

apply the Robbins-Siegmund theorem (Robbins and Siegmund, 1971). In contrast, it is unclear how almost sure convergence rates can be derived for ESGDA.

Similar to what was done in Sebbouh et al. (2021) for SGD, we can use this lemma and Prop. 3.1 to derive small- α almost sure convergence rates for RSGDA. To the best of our knowledge, these are the first such rates for a gradient descent ascent algorithm.

Corollary 3.2. Consider the setting of Prop. 3.1. We have that

$$\min_{t=0; \dots; k} \sum_{t=0}^k \alpha_t \|\nabla f(v_t; z_t)\|^2 = o\left(\frac{1}{k}\right) \text{ a.s.} \quad (3)$$

as long as the step sizes α_k and β_k verify

$$\sum_{k=1}^{\infty} \alpha_k = 1; \quad \sum_{k=1}^{\infty} \alpha_k^2 < 1;$$

$$\sum_{k=1}^{\infty} \beta_k \alpha_k^{-2} < 1; \quad \sum_{k=1}^{\infty} \frac{\beta_k^3}{k} < 1; \quad (4)$$

Let $\epsilon > 0$ and take for all $k \geq N$, $\alpha_k = \frac{1}{2L(k+1)^{2+\epsilon}}$ and $\beta_k = \frac{p}{2} \frac{1-p}{p(2p+(1-p)\alpha_k)} \frac{\alpha_k}{(k+1)^{1+\epsilon}}$, we have that the step sizes verify (4) and

$$\min_{t=0; \dots; k} \sum_{t=0}^k \alpha_t \|\nabla f(v_t; z_t)\|^2 = o\left(k^{-\frac{2}{5}+\epsilon}\right) \text{ almost surely.}$$

Remark 3.3. Using the same analysis, it is also possible to show that using the exact gradients, Randomized GDA converges almost surely at $\mathcal{O}(k^{-1})$ rate. Indeed, in that setting, $\alpha_k^2 = \alpha_k^{-2} = 0$ and any constant step sizes $\alpha_k > 0$ verify conditions (4). Substituting in (3) gives the desired rate.

Convergence rates in expectation. In the remainder of this section, we establish convergence rates in expectation for Randomized GDA (where we

use the exact gradient at each iteration) and RSGDA (where we use stochastic gradients). The rates we recover are similar to the best that can be derived for GDA and SGDA (Lin et al., 2019). The interest of the results we present is that (i) we show that it is possible to do more gradient ascent steps and retain the same complexity as SGDA, and we quantify this by determining a range of values of p for which we have this complexity, (ii) we highlight the effect of changing the probability p on the step sizes, and (iii) we show that taking $p = 1/2$ doesn't result in a diverging algorithm; instead the convergence rate is worse by a factor which is an increasing function of p , so that the strategy of making moderately more descent steps than ascent steps is worth exploring depending on the application.

3.1 Randomized GDA (RGDA)

We first consider RGDA, a version of Alg. 2 where we use the exact gradients, i.e. where the only randomness in the algorithm comes from the coin tosses governed by the probability p . This is the randomized equivalent of Epoch GDA.

Corollary 3.4. Consider the setting of Prop. 3.1. Let for all $k \geq 2 \in \mathbb{N}$, $\alpha_k = \frac{1}{2L}$ and $\beta_k = \frac{(1-p)}{4} \frac{p}{2p+(1-p)\alpha_k}$. Then, for all $p \in [1/2, 1]$, we have that

$$\min_{t=0, \dots, k-1} \mathbb{E} [kr - (t)k^2] = O(k^{-1}) :$$

Hence, finding an ϵ -stationary point requires at most $O(k^{-2})$ iterations for all $p \in [1/2, 1]$.

Remark. The value $p=1/2$ of the upper bound on p is arbitrary, and any constant higher than $p=1/2$ but independent of L still ensures the same complexity.

We draw two important insights from this corollary:

(a) Flexibility to the choice of p . RGDA admits a range of values of p for which we have the same order of complexity. In particular, with $p = 1/2$, we recover the same convergence rate for GDA which was derived in Lin et al. (2019), with better constants, as highlighted in the appendix.

(b) Suggested choice of step sizes. It is known from Lin et al. (2019) that the step sizes α_k and β_k need to verify the following relation: $\alpha_k = \beta_k^2$. How should the step sizes change when we make more ascent steps? Since $p \in [1/2, 1]$ gives the tightest convergence rate, let us consider the case where p is in this range. In this setting, Cor. 3.4 shows that the relation should be $\alpha_k = \beta_k^2$. Hence, there is a

linear relation between the descent step size and p : the lower p , i.e. the more ascent steps we make, the higher α_k should be.

We can draw similar conclusions in the stochastic setting as well, but with different ranges for p and different choices for the step sizes.

3.2 Randomized SGDA (RSGDA)

We now consider RSGDA, where we use stochastic gradients instead of deterministic ones. This is the randomized version of Epoch SGDA.

Decreasing Step Sizes. As is the case for SGD, without additional structure on the objective F , we can only guarantee the anytime convergence (without knowledge of the last iteration or the required precision) of RSGDA if we use decreasing step sizes.

Corollary 3.5. Consider the setting of Prop. 3.1. Let $\alpha_k = \frac{1}{2L(k+1)^{2-5\epsilon}}$ and $\beta_k = \frac{p}{2} \frac{1-p}{p(2p+(1-p)\alpha_k)}$. Then,

$$\min_{t=0, \dots, k-1} \mathbb{E} [kr - (t)k^2] = O(\log(k)k^{2-5\epsilon}) :$$

The previous corollary presented the parameter settings that resulted in the best rate when using decreasing step sizes. We have used the Θ notation for brevity even though the result is non-asymptotic. A complete version of this corollary, including the explicit bound, can be found in the appendix.

RSGDA With Fixed Step Sizes. In practice, it is often the case that stochastic gradient algorithms are implemented using a fixed step size.

Corollary 3.6. Consider the setting of Prop. 3.1. Let $\epsilon > 0$. With suitable choices of step sizes, if $\alpha_k = \frac{1}{2L(k+1)^{2-5\epsilon}}$, then $\min_{t=0, \dots, k-1} \mathbb{E} [kr - (t)k] = O(k^{3-5\epsilon})$. Moreover, for all $p \in [1/2, 1]$, the complexity is still $O(k^{3-5\epsilon})$.

We give the explicit values for the step sizes and the complexity in the appendix. In this corollary, we recover the complexity given without proof in Lin et al. (2019) when using $p = 1/2$, but we extend their result to a larger number of ascent steps (instead of a single one). Note that the step sizes rely on some quantities that are generally impossible to obtain, like α_k . As is the case for SGD, without knowledge of this constant and using constant step sizes, it is possible to guarantee only sublinear convergence towards the neighborhood of a solution (see e.g. (Gower et al., 2019)).

RSGDA With Large Batch Sizes. The best known complexity for SGDA is $O(k^{-3/4})$ (Lin et al., 2019), and it is obtained using fixed step sizes and large minibatch sizes. In the following corollary, we show that the same rate can be recovered for RSGDA using large minibatch sizes. As done in the previous corollaries, we also give the range of values of p which give the same rate. Following Lin et al. (2019), we assume that all stochastic gradients have bounded variance in this setting.

Assumption 3.7. There exists $\sigma^2 > 0$ such that for all $(; v) \in \mathbb{R}^d \times V$,

$$\begin{aligned} \mathbb{E}_z \|\nabla_{;v} f(; v; z) - \nabla_{;v} F(; v)\|^2 &\leq \sigma^2 \\ \mathbb{E}_z \|\nabla_{;v} f(; v; z) - \nabla_{;v} F(; v)\|^2 &\leq \sigma^2 \end{aligned}$$

Corollary 3.8 (Large minibatch sizes). Consider the setting of Prop. 3.1. Let Assumption 3.7 hold and choose $p \in [1, \frac{1}{2}]$. Using the step sizes of Cor. 3.4 and a sufficiently large minibatch size, the total number of stochastic gradient evaluations to reach an ϵ -stationary point is $O(k^{3/4})$.

Discussion about other values of p . Note that Prop. 3.1 and the subsequent corollaries allow for any value of p in $(0, 1)$ and still guarantee the convergence of RSGDA, albeit with an additional factor in the convergence rates (See Appendix). This suggests tuning p to values that are moderately larger than $1/2$, and we would still expect RSGDA to perform almost as well as with values of p in the range which gives the best convergence rates. In applications where descent steps are cheaper than ascent steps, this strategy is sound.

3.2.1 Interpolation Setting

We consider using RSGDA in a more favorable setting, where the maximization problem is easier, as in the interpolation case, when F is a finite-sum and $\sigma^2 = 0$.

Assumption 3.9 (Interpolation). For all $(; v) \in \mathbb{R}^d \times \mathbb{R}^n$; $F(; v) = \frac{1}{n} \sum_{i=1}^n F_i(; v)$, where for all $i \in [n]$, F_i verifies Assumption 2.1 and for all $; \in \mathbb{R}^d$, there exists $v(;) \in \mathbb{R}^n$ such that for all $i \in [n]$, $\nabla_{;v} F_i(; v(;)) = 0$.

Note that we do not assume that $\sigma^2 = 0$, which, by Assumption 2.3, would have implied, for $i \in [n]$ and $(; v) \in \mathbb{R}^d \times \mathbb{R}^n$ that $\nabla_{;v} F_i(; v) = \nabla_{;v} F(; v)$.

From an optimization perspective, this setting has been explored in many works on SGD (Vaswani et al., 2019a,b, Loizou et al., 2021, Sebbouh et al., 2021),

where it was shown that if the two previous assumptions are verified (for the single-variable objective), SGD has the same convergence rate as Gradient Descent. This setting has also recently been studied for bilinear minimax optimization (Li et al., 2021),

The next result shows that in that setting, RSGDA converges at the improved rate of $O(k^{-2/4})$ (versus $O(k^{-3/4})$), without requiring large batch sizes.

Corollary 3.10. Consider the setting of Proposition 3.1 and let Assumption 3.9 hold. Let $k = 1 + (2L)$ for all $k \geq 2$.

- Almost sure convergence. Let $\epsilon > 0$ and $k = O(\frac{1}{\epsilon^{2(k+1)^{1+2p}}})$. Then, $\min_{t=0, \dots, k-1} \|\nabla_{;v} f(;) - \nabla_{;v} F(;)\|^2 = o(k^{-1-2p})$ a.s.
- Anytime convergence in expectation. Let $k = O(\frac{1}{\epsilon^{2p/(k+1)}})$. Then, $\min_{t=0, \dots, k-1} \mathbb{E} \|\nabla_{;v} f(;) - \nabla_{;v} F(;)\|^2 = O(\frac{\log(k)}{k^{1+2p}})$.
- Convergence in expectation for a given precision. With a suitable choice of k , if $p \in [1/2, 1]$ and $k = O(\frac{1}{\epsilon^{2/4}})$, then $\min_{t=0, \dots, k-1} \mathbb{E} \|\nabla_{;v} f(;) - \nabla_{;v} F(;)\|^2 = O(\epsilon^{1+2p})$.

By setting $\sigma^2 = 0$, all these results can be derived from Prop. 3.1 in a similar fashion to the previous corollaries. As an illustration of these results, we apply our method to distributionally robust optimization in §4.1, where Assumption 3.9 holds for the problem we consider (Sinha et al., 2018).

4 APPLICATIONS

To illustrate our results, we consider two nonconvex strongly concave problems: distributionally robust optimization (Shah et al., 2015, Kuhn et al., 2019) and learning with a Sinkhorn loss (Genevay et al., 2018).

4.1 Distributionally Robust Optimization

The goal of distributionally robust optimization is to learn machine learning models which are robust to changes in the distribution of test data compared to training data. Consider a training dataset $D \stackrel{\text{def}}{=} \{(x_1, y_1), \dots, (x_n, y_n)\} \in \mathbb{R}^{p+q}; p, q \geq 1$, and suppose we want to learn a robust classifier from a parametric family $f : \mathbb{R}^p \rightarrow \mathbb{R}^q; \theta \in \mathbb{R}^d$. Sinha et al. (2018) showed that one way to do so is to solve

Figure 2: Training loss (5) when using RSGDA. $\eta = 10$. We study the effect of the descent step size and the descent probability p .

the following optimization problem:

$$\min_{\theta \in \mathbb{R}^d} F(\theta; v); \tag{5}$$

$$F(\theta; v) = \max_{v \in \mathbb{R}^{pn}} \frac{1}{n} \sum_{j=1}^n \ell(f(v_j); y_j) + \kappa \|v - x\|_2^2;$$

where $\ell : \mathbb{R}^p \times \mathbb{R}^q \rightarrow \mathbb{R}$ is the loss function. Intuitively, the inner maximization problem requires that we find the (adversarial) inputs $f(v_j; g_j)$ that maximize the loss, while minimizing the average distance to the original inputs $f(x_j; g_j)$. The hyperparameter κ controls the trade-off between these two objectives. A low κ allows the classifier to be robust to adversarial perturbations, but sacrifices model accuracy on training data. A high κ ensures that the model performs well on training data, but sacrifices robustness. [Sinha et al. \(2018\)](#) showed that for κ large enough, solving this problem is nonconvex strongly concave. Moreover, we also have that Assumption 3.9 is verified.

Lemma 4.1 ([Sinha et al. \(2018\)](#)). Consider F defined in (5). Assume that $\ell(\cdot; v) : \mathbb{R}^p \times \mathbb{R}^q \rightarrow \mathbb{R}$ is smooth for all $(\cdot; v; y) \in \mathbb{R}^d \times \mathbb{R}^{pn} \times \mathbb{R}^q$, and that the noise assumption 2.3 holds. Then for a large enough κ , F verifies Assumption 2.1. Moreover, since F is separable in each coordinate $\theta_j \in \mathbb{R}^d; j \in [n]$, we have that Assumption 3.9 is verified.

Application To MNIST. We reproduce the setting of [Sinha et al. \(2018\)](#) for the MNIST dataset and concentrate on the optimization aspect of their procedure. For the experiment to still be meaningful, we only considered the models which resulted in higher than 98% validation accuracy. As done by [Sinha et al. \(2018\)](#), we set $\kappa = 1:3$ and choose ℓ to be a CNN with smooth ELU activations. More details on the experimental settings can be found in the Appendix. Contrary to [Lin et al. \(2019\)](#), we report the loss of interest $F(\cdot)$ in order to better assess the effect of the parameter settings.

Fig. 2 shows the performance of RSGDA (Alg. 2) on Problem (5) with various parameter settings. We use minibatch RSGDA and set the ascent step size $\eta = 10$. We noticed that further tuning this hyperparameter had less impact than κ and p on the performance of RSGDA. For each probability p , we display the step size η that resulted in the lowest training loss. For all experiments with $p \in [0:5]$, choosing $\eta > 0:25$ led the model to diverge or result in a low validation accuracy. As predicted by theory, decreasing descent probability allows to take larger step sizes ($\eta = 0:25; \kappa = 0:3$) and results in faster optimization, but decreasing p too much ($p = 0:1$) leads to a slow algorithm, as θ is updated less often. Further increasing the step size when $p = 0:1$ makes RSGDA diverge.

4.2 Learning with a Sinkhorn Loss

The second application we consider is that of learning with a semi-dual Sinkhorn loss ([Cuturi and Peyré, 2018](#), [Kitagawa et al., 2019](#)). We establish that this problem is indeed nonconvex-strongly concave. Then, inspired by [Schiebinger et al. \(2019\)](#), [Stark et al. \(2020\)](#), we apply semi-dual optimal transport to a single-cell data integration problem using RSGDA. But first, let us recall the definition of the regularized OT loss between two measures ([Cuturi, 2013](#)).

Regularized OT Loss. Let X and Y be two metric spaces. Consider two probability measures $(\mu; \nu) \in \mathcal{P}(X) \times \mathcal{P}(Y)$ and let $\epsilon > 0$. The regularized OT metric between these two measures is given by

$$W(\mu; \nu) \stackrel{\text{def}}{=} \min_{\gamma \in \mathcal{Z}(\mu; \nu)} \int c(x; y) d\gamma(x; y) + \epsilon \log \frac{d(\mu; \nu)}{d(\mu) d(\nu)} d(\mu; \nu);$$

where $c : X \times Y \rightarrow \mathbb{R}_+$ is the cost to move a unit of mass from x to y , $(\mu; \nu) = (\sum_{i=1}^n \mu_i \delta_{x_i}; \sum_{j=1}^n \nu_j \delta_{y_j})$ and $P_1(x; y) = x$ and $P_2(x; y) = y$ are projection operators. With $X = Y$, d_x a metric on X , and $c = d_x^p$, we have that $W_0^{1=p}$ defines a distance on $\mathcal{P}(X)$. Unfortunately, computing W_0 is too costly in most applications, which justifies using the regularized OT loss with $\epsilon > 0$. Fortunately, when one of the measures is discrete, we can express $W(\mu; \nu)$ as a finite-dimensional stochastic maximization problem ([Genevay et al., 2016](#)). Consider a dataset $Y_n = \{y_1; \dots; y_n\} \subset Y$ and let $\mu = \frac{1}{n} \sum_{j=1}^n \delta_{y_j}$.

Figure 3: t-SNE embeddings of the point cloud X_k (7) along the iterations of RSGDA ($\rho = 0.9$).

Figure 4: Loss $W(\cdot; \cdot)$ (7) when using RSGDA vs. Sinkhorn's algorithm, depending on the descent probability p (for RSGDA) or the number of iterations m_{sin} of Sinkhorn's algorithm. $\epsilon = 5$.

Then,

$$W(\cdot; \cdot) = \max_{\nu \in \mathcal{P}(Z)} E_x [h(x; \nu)]; \quad \text{where} \quad (6)$$

$$h(x; \nu) \stackrel{\text{def}}{=} \sum_{j=1}^n \frac{v_j}{n} \log \sum_{j=1}^n \frac{\exp \frac{v_j c(x; y_j)}{n}}{n} A :$$

Learning With a Semi-discrete Sinkhorn Loss is a Nonconvex Strongly Concave Problem. Consider the task of learning a parametric map from observed data $Y_n = \{y_1, \dots, y_n\} \subset Y$. Given a probability space Z and a possibly continuous distribution $\nu \in \mathcal{P}(Z)$, we may want to solve one of two problems:

$$\min_{\nu \in \mathcal{P}(Z)} W(\cdot; \nu); \quad \frac{1}{n} \sum_{j=1}^n \log \sum_{j=1}^n \frac{\exp \frac{v_j c(x; y_j)}{n}}{n} A \quad \max_{\nu \in \mathcal{P}(Z)} E_z [h(f(z); \nu)];$$

$$\text{or} \quad \min_{\nu \in \mathcal{P}(Z)} W(\cdot; \nu); \quad \frac{1}{n} \sum_{j=1}^n \log \sum_{j=1}^n \frac{\exp \frac{v_j c(x; y_j)}{n}}{n} A =: W(\cdot; \cdot); \quad (7)$$

where $f: Z \rightarrow X$; $\nu \in \mathcal{P}(Z)$. An instance of the first problem is learning generative models (Genevay et al., 2018, Houdard et al., 2021), where one wants to fit to the dataset Y_n . There, ν is

typically a Gaussian in a low-dimensional space, and Y_n is for example a dataset of images. An instance of the second problem is single-cell data matching, which we present later in this section. Using the formulation (6) of W , we can express both problems as finite dimensional stochastic min-max problems. We now show that these problems are indeed nonconvex strongly concave.

Proposition 4.2. Define for all $(\cdot; \nu) \in \mathcal{P}(Z) \times \mathcal{P}(X)$; $F(\cdot; \nu) = E_z [h(f(z); \nu)]$. Assume that for all $z \in Z$, $\nu \in \mathcal{P}(X)$ $c(f(z); \cdot)$ is Lipschitz continuous and smooth, and that $\nu \in \mathcal{P}(X)$ $c(f(z); \cdot)$ is Lipschitz-continuous almost surely for all $(\cdot; \nu) \in \mathcal{P}(Z) \times \mathcal{P}(X)$. Then, there exists $V \subset \mathcal{P}(X)$ such that F verifies Assumptions 2.1, 2.3 and 2.4 on $\mathcal{P}(Z) \times V$ and $(\cdot) = \max_{\nu \in V} F(\cdot; \nu)$ for all $\cdot \in \mathcal{P}(Z)$.

A similar result can be proved for (7) (see App. for proof of Lem. 4.2). This result is of independent interest for OT practitioners since combining Lem. 4.2 with our convergence results in §3, we extend (Sanjabi et al., 2018a) to the semi-discrete setting, and we give stronger theoretical guarantees than those in (Houdard et al., 2021) when using GDA algorithms for learning with a semi-discrete Sinkhorn loss.

Single-cell Genomic Data Matching. We consider the single cell melanoma tumor dataset from (Stark et al., 2020). Patient data is analyzed using two different technologies, scRNA and CyTOF, resulting in two point clouds of different sample sizes and dimensions. The goal is to integrate both datasets in order to understand the correspondances between technologies, and obtain a unified analysis on larger sample sizes. We denote the CyTOF point cloud as $\{x_1, \dots, x_m\} \subset \mathbb{R}^{41}$; $m = 135334$, and define $\nu = \frac{1}{m} \sum_{i=1}^m \delta_{x_i}$, and the scRNA one as $\{y_1, \dots, y_n\} \subset \mathbb{R}^{1024}$; $n = 4683$, and define $\mu = \frac{1}{n} \sum_{j=1}^n \delta_{y_j}$. Instead of embedding both point clouds in another low-dimensional space as done in Stark et al. (2020), we aim to map the point cloud directly into \mathbb{R}^d via a multilayer perceptron (MLP), namely

solve (7), where f is an MLP. We compare minibatch RSGDA (where we sample from \mathcal{D}) against Sinkhorn's algorithm (the most widely used approach to compute W) with minibatching (Cuturi, 2013) in Fig. 4. Details on the architecture and on Sinkhorn's algorithm are provided in the appendix (see also Genevay et al. (2018)). Using a single iteration of Sinkhorn's algorithm resulted in the best performance for this algorithm: due to the bias introduced by using Sinkhorn's algorithm on minibatches (Fratras et al., 2020), increasing the number of iterations resulted in an increasingly biased gradient direction for the descent step. For RSGDA, we found that increasing the descent probability resulted in a faster algorithm up to the value $p = 0.95$, after which the algorithm is slower than $p = 0.5$. On this problem, decreasing p below 0.5 was not beneficial, despite increasing the descent step size. Using t-SNE embeddings (Van der Maaten and Hinton, 2008), we also display in Fig. 3 the evolution of $\|x_k\|$ for the best run of RSGDA with $p = 0.9$.

Conclusion. We have presented RSGDA, a randomized version of Epoch Stochastic Gradient Descent Ascent, a popular method for solving min-max optimization problems in machine learning. We showed empirically that RSGDA performs similarly to ESGDA, and demonstrated theoretically that RSGDA enjoys the same theoretical properties as Stochastic Gradient Descent Ascent (Lin et al., 2019), a method which is well grounded in theory. We also provided practical parameter settings, which we tested numerically on a distributionally robust optimization problem and single-cell data integration using optimal transport.

Acknowledgements

The work of Gabriel Peyré and Othmane Sebbouh was supported in part by the French government under management of Agence Nationale de la Recherche as part of the Investissements d'avenir program, reference ANR19-P3IA-0001 (PRAIRIE 3IA Institute). Gabriel Peyré also acknowledges support from the European Research Council (ERC project NORIA). Othmane Sebbouh also acknowledges the support of a Chaire d'excellence de l'IDEX Paris Saclay. The authors thank Pierre Ablin, Michael Sander and Mathieu Blondel for valuable feedback and discussions.

References

- Azizian, W., Mitliagkas, I., Lacoste-Julien, S., and Gidel, G. (2020). A tight and unified analysis of gradient-based methods for a whole spectrum of differentiable games. In *International Conference on Artificial Intelligence and Statistics*, pages 2863–2873. PMLR. 1
- Bandura, D. R., Baranov, V. I., Ornatsky, O. I., Antonov, A., Kinach, R., Lou, X., Pavlov, S., Vorobiev, S., Dick, J. E., and Tanner, S. D. (2009). Mass cytometry: technique for real time single cell multitarget immunoassay based on inductively coupled plasma time-of-flight mass spectrometry. *Analytical chemistry*, 81(16):6813–6822. 28
- Ba³ar, T. and Olsder, G. J. (1998). *Dynamic noncooperative game theory* SIAM. 1
- Bertsekas, D. P. and Tsitsiklis, J. N. (2000). Gradient convergence in gradient methods with errors. *SIAM Journal on Optimization*, 10(3):627–642. 4
- Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, learning, and games* Cambridge university press. 1
- Chen, T., Sun, Y., and Yin, W. (2021). Tighter analysis of alternating stochastic gradient method for stochastic nested problems. *arXiv preprint arXiv:2106.13781*. 3
- Chen, Z., Zhou, Y., Xu, T., and Liang, Y. (2020). Proximal gradient descent-ascent: Variable convergence under k^a geometry. In *International Conference on Learning Representations* 3
- Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems* 26:2292–2300. 2, 7, 9
- Cuturi, M. and Peyré, G. (2018). Semidual regularized optimal transport. *SIAM Review*, 60(4):941–965. 7
- Dai, B., Shaw, A., Li, L., Xiao, L., He, N., Liu, Z., Chen, J., and Song, L. (2018). Sbeed: Convergent reinforcement learning with nonlinear function approximation. In *International Conference on Machine Learning*, pages 1125–1134. PMLR. 1, 2
- Fratras, K., Zine, Y., Flamary, R., Gribonval, R., and Courty, N. (2020). Learning with minibatch wasserstein: asymptotic and gradient properties. In *International Conference on Artificial Intelligence and Statistics*, pages 2131–2141. PMLR. 9, 29
- Gadat, S., Panloup, F., and Saadane, S. (2018). Stochastic heavy ball. *Electronic Journal of Statistics*, 12(1):461–529. 4

- Genevay, A., Chizat, L., Bach, F., Cuturi, M., and Peyré, G. (2019). Sample complexity of sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1574–1583. [25](#)
- Genevay, A., Cuturi, M., Peyré, G., and Bach, F. (2016). Stochastic optimization for large-scale optimal transport. *Advances in Neural Information Processing Systems* pages 3440–3448. [7](#), [22](#)
- Genevay, A., Peyré, G., and Cuturi, M. (2018). Learning generative models with sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics, AISTATS 2018*, pages 1608–1617. [6](#), [8](#), [9](#)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27. [1](#), [2](#), [3](#)
- Gower, R. M., Loizou, N., Qian, X., Sailanbayev, A., Shulgin, E., and Richtárik, P. (2019). Sgd: General analysis and improved rates. In *International Conference on Machine Learning* pages 5200–5209. [5](#), [16](#)
- Gnarnova, P., Levy, K. Y., Lucchi, A., Hofmann, T., and Krause, A. (2017). An online learning approach to generative adversarial networks. *arXiv preprint arXiv:1706.03269* [1](#)
- Hendrycks, D. and Gimpel, K. (2016). Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415* [28](#)
- Hofmann, T., Lucchi, A., Lacoste-Julien, S., and McWilliams, B. (2015). Variance reduced stochastic gradient descent with neighbors. *Advances in Neural Information Processing Systems* 28:2305–2313. [4](#)
- Houdard, A., Leclaire, A., Papadakis, N., and Rabin, J. (2021). On the existence of optimal transport gradient for learning generative models. *arXiv preprint arXiv:2102.05542*. [2](#), [3](#), [8](#)
- Huang, F., Gao, S., Pei, J., and Huang, H. (2020). Accelerated zeroth-order and first-order momentum methods from mini to minimax optimization. *arXiv preprint arXiv:2008.08170*. [2](#)
- Jin, C., Netrapalli, P., and Jordan, M. (2020). What is local optimality in nonconvex-nonconcave minimax optimization? In *International Conference on Machine Learning*, pages 4880–4889. PMLR. [1](#), [2](#)
- Johnson, R. and Zhang, T. (2013). Accelerating stochastic gradient descent using predictive variance reduction. *Advances in neural information processing systems* 26:315–323. [4](#)
- Kitagawa, J., Mérigot, Q., and Thibert, B. (2019). Convergence of a newton algorithm for semi-discrete optimal transport. *Journal of the European Mathematical Society*, 21(9):2603–2651. [7](#)
- Kong, W. and Monteiro, R. D. (2019). An accelerated inexact proximal point method for solving nonconvex-concave min-max problems. *arXiv preprint arXiv:1905.13433*. [1](#)
- Korpelevich, G. M. (1976). The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756. [1](#)
- Kovalev, D., Horváth, S., and Richtárik, P. (2020). Don't jump through hoops and remove those loops: SvrG and katyusha are better without the outer loop. In *Algorithmic Learning Theory*, pages 451–467. PMLR. [4](#)
- Kuhn, D., Esfahani, P. M., Nguyen, V. A., and Shafieezadeh-Abadeh, S. (2019). Wasserstein distributionally robust optimization: Theory and applications in machine learning. In *Operations Research & Management Science in the Age of Analytics*, pages 130–166. INFORMS. [6](#)
- Li, C. J., Yu, Y., Loizou, N., Gidel, G., Ma, Y., Roux, N. L., and Jordan, M. I. (2021). On the convergence of stochastic extragradient for bilinear games with restarted iteration averaging. *arXiv preprint arXiv:2107.00464*. [6](#)
- Lin, T., Jin, C., and Jordan, M. I. (2019). On Gradient Descent Ascent for Nonconvex-Concave Minimax Problems. *arXiv preprint arXiv:1906.00331*. [2](#), [3](#), [5](#), [6](#), [7](#), [9](#), [20](#)
- Loizou, N., Vaswani, S., Laradji, I. H., and Lacoste-Julien, S. (2021). Stochastic polyak step-size for sgd: An adaptive learning rate for fast convergence. In *International Conference on Artificial Intelligence and Statistics* pages 1306–1314. PMLR. [6](#)
- Lu, S., Tsaknakis, I., Hong, M., and Chen, Y. (2020). Hybrid block successive approximation for one-sided non-convex min-max problems: algorithms and applications. *IEEE Transactions on Signal Processing* 68:3676–3691. [1](#)
- Namkoong, H. and Duchi, J. C. (2016). Stochastic gradient methods for distributionally robust optimization with f-divergences. In *NIPS*, volume 29, pages 2208–2216. [1](#)
- Nedić, A. and Ozdaglar, A. (2009). Subgradient methods for saddle-point problems. *Journal of optimization theory and applications*, 142(1):205–228. [1](#)

- Nemirovski, A. (2004). Prox-method with rate of convergence $O(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251. [1](#)
- Nouiehed, M., Sanjabi, M., Huang, T., Lee, J. D., and Razaviyayn, M. (2019a). Solving a class of non-convex min-max games using iterative first order methods. *Advances in Neural Information Processing Systems* 32(NeurlPS). [1](#)
- Nouiehed, M., Sanjabi, M., Huang, T., Lee, J. D., and Razaviyayn, M. (2019b). Solving a class of non-convex min-max games using iterative first order methods. In *Advances in Neural Information Processing Systems*, volume 32. [1](#), [2](#)
- Qiu, S., Yang, Z., Wei, X., Ye, J., and Wang, Z. (2020). Single-timescale stochastic nonconvex-concave optimization for smooth nonlinear reinforcement learning. *arXiv preprint arXiv:2008.10103*. [2](#)
- Raque, H., Liu, M., Lin, Q., and Yang, T. (2018). Non-convex min max optimization: provable algorithms and applications in machine learning (2018). *arXiv preprint arXiv:1810.02060*. [1](#)
- Robbins, H. and Siegmund, D. (1971). A convergence theorem for nonnegative almost supermartingales and some applications. *Optimizing methods in Statistics*, pages 233–257. [4](#), [18](#)
- Sanjabi, M., Ba, J., Razaviyayn, M., and Lee, J. D. (2018a). On the convergence and robustness of training GANs with regularized optimal transport. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 7091–7101. [2](#), [8](#)
- Sanjabi, M., Razaviyayn, M., and Lee, J. D. (2018b). Solving non-convex non-concave min-max games under Polyak-Frostig condition. *arXiv preprint arXiv:1812.02878*. [2](#)
- Schiebinger, G., Shu, J., Tabaka, M., Cleary, B., Subramanian, V., Solomon, A., Gould, J., Liu, S., Lin, S., Berube, P., et al. (2019). Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943. [2](#), [7](#)
- Sebbouh, O., Gower, R. M., and Defazio, A. (2021). Almost sure convergence rates for stochastic gradient descent and stochastic heavy ball. In *Conference on Learning Theory*, pages 3935–3971. PMLR. [4](#), [6](#), [18](#)
- Shaeezadeh-Abadeh, S., Esfahani, P. M., and Kuhn, D. (2015). Distributionally robust logistic regression. In *Proceedings of the 28th International Conference on Neural Information Processing Systems*-Volume 1, pages 1576–1584. [6](#)
- Sinha, A., Namkoong, H., and Duchi, J. (2018). Certifying some distributional robustness with principled adversarial training. In *International Conference on Learning Representations* [1](#), [2](#), [3](#), [6](#), [7](#)
- Sion, M. (1958). On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176. [1](#)
- Stark, S. G., Ficek, J., Locatello, F., Bonilla, X., Chevrier, S., Singer, F., Rätsch, G., Lehmann, K.-V., Consortium, T. P., et al. (2020). Scim: universal single-cell matching with unpaired feature sets. *Bioinformatics (Oxford, England)*, 36(Supplement_2):i919–i927. [2](#), [7](#), [8](#), [27](#)
- Tang, F., Barbacioru, C., Wang, Y., Nordman, E., Lee, C., Xu, N., Wang, X., Bodeau, J., Tuch, B. B., Siddiqui, A., et al. (2009). mRNA-seq whole-transcriptome analysis of a single cell. *Nature methods* 6(5):377–382. [28](#)
- Thekumparampil, K. K., Jain, P., Netrapalli, P., and Oh, S. (2019). Efficient algorithms for smooth minimax optimization. *arXiv preprint arXiv:1907.01543* [1](#)
- Van der Maaten, L. and Hinton, G. (2008). Visualizing data using t-sne. *Journal of machine learning research* 9(11). [9](#)
- Vaswani, S., Bach, F., and Schmidt, M. (2019a). Fast and faster convergence of SGD for over-parameterized models and an accelerated perceptron. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1195–1204. PMLR. [6](#)
- Vaswani, S., Mishkin, A., Laradji, I., Schmidt, M., Gidel, G., and Lacoste-Julien, S. (2019b). Painless stochastic gradient: Interpolation, line-search, and convergence rates. *Advances in neural information processing systems* 32:3732–3745. [6](#)
- Von Neumann, J. and Morgenstern, O. (2007). *Theory of games and economic behavior*. Princeton university press. [1](#)
- Yan, Y., Xu, Y., Lin, Q., Liu, W., and Yang, T. (2020). Optimal epoch stochastic gradient descent ascent methods for min-max optimization. *Advances in Neural Information Processing Systems* 33. [3](#)
- Zhang, J., Xiao, P., Sun, R., and Luo, Z.-Q. (2020). A single-loop smoothed gradient descent-ascent algorithm for nonconvex-concave min-max problems. *arXiv preprint arXiv:2010.15768*. [1](#)

APPENDIX

Contents

1	INTRODUCTION	1
2	BACKGROUND	2
2.1	Assumptions	2
2.2	Review of Gradient Descent Ascent Methods	2
3	RANDOMIZED STOCHASTIC GRADIENT DESCENT ASCENT	3
3.1	Randomized GDA (RGDA)	5
3.2	Randomized SGDA (RSGDA)	5
3.2.1	Interpolation Setting	6
4	APPLICATIONS	6
4.1	Distributionally Robust Optimization	6
4.2	Learning with a Sinkhorn Loss	7
A	Complete results	14
A.1	Corollary 3.4	14
A.2	Corollary 3.5	14
A.3	Corollary 3.6	15
A.4	Corollary 3.8	15
B	Proofs	15
B.1	Proof of Proposition 3.1	15
B.2	Proof of Corollary 3.2	18
B.3	Proofs for convergence rates in expectation	19
B.3.1	Proof of Corollary 3.4	19
B.3.2	Proof of Corollary 3.5	19
B.3.3	Proof of Corollary 3.6	19
B.3.4	Proof of Corollary 3.8	20
B.4	Proofs for the interpolation results of Section 3.2.1 (Corollary 3.10)	21
B.4.1	Almost sure convergence	21
B.4.2	Anytime convergence in expectation	21
B.4.3	Convergence in expectation for a given precision.	21
C	Learning with a semi-discrete Sinkhorn loss is a nonconvex-strongly concave problem	22
C.1	Problem setting	22

C.2 Assumptions and consequences	22
C.3 Restricting the maximization problem	25
C.4 The semi-dual objective is strongly concave on the subset	26
D Details about the experimental setting of Section 4.2	27
D.1 Dataset, architecture and OT hyperparameters	27
D.2 Using Sinkhorn's algorithm	28
D.2.1 Impact of the choice of the number of iterations of Sinkhorn's algorithm	28

The appendix is organized as follows:

- ^ In appendix A, we present the detailed statements of the corollaries we presented in §3.
- ^ In appendix B, we present the proofs of all the statements of §3.
- ^ In appendix C, we present the proof of Proposition 4.2.
- ^ In appendix D, we present details on the experimental setting of §4.2.

A Complete results

In this section, we give the detailed statements of the corollaries we presented in §3. The proofs of these results are presented in appendix B.

A.1 Corollary 3.4

Corollary 3.4 presented the convergence rate in expectation of Randomized Gradient Descent Ascent (RGA) (Alg. 2) when using deterministic gradients to update \mathbf{k} and \mathbf{v}_k .

Corollary A.1. Consider the setting of Proposition 3.1. Let $\eta_k = \frac{1}{L}$ and $\rho = \frac{(1-p)}{2} \frac{p}{p(2p+1-p)}$. Then,

$$\min_{t=0, \dots, k} \mathbb{E} \sum_{i=1}^h \|\mathbf{k}^i(t)\|^2 \leq \frac{4L}{(1-p)k} \frac{\rho}{2 + \frac{1-p}{p}} D_0 + \frac{L^2}{(1-p)k} \rho_0.$$

In particular, for all $p \in [\frac{1}{2}, 1]$, we have that

$$\min_{t=0, \dots, k} \mathbb{E} \sum_{i=1}^h \|\mathbf{k}^i(t)\|^2 = O\left(\frac{1}{k}\right):$$

This implies that finding an ϵ -stationary point requires at most $O\left(\frac{1}{\epsilon^2}\right)$ iterations for all $p \in [\frac{1}{2}, 1]$.

A.2 Corollary 3.5

Corollary 3.5 presented the anytime convergence rate in expectation of RSGDA (Alg. 2) using decreasing step sizes.

Corollary A.2. Consider the setting of Proposition 3.1. Then,

$$\begin{aligned} \min_{t=0, \dots, k} \mathbb{E} \sum_{i=1}^h \|\mathbf{k}^i(t)\|^2 &\leq \frac{2E_0}{\sum_{t=0}^{k-1} p_t} + \frac{4L}{\sum_{t=0}^{k-1} p_t} \sum_{t=0}^{k-1} p_t^2 \\ &+ \frac{2L}{\sum_{t=0}^{k-1} p_t} \sum_{t=0}^{k-1} p_t^2 \\ &+ \frac{2}{\sum_{t=0}^{k-1} p_t} \sum_{t=0}^{k-1} \frac{2p_t^3}{(1-p_t)^2} + \frac{p_t^3}{(1-p_t)} : \end{aligned}$$

Let $p = p_2(0; 1)$, $\eta_k = \frac{1}{2L(k+1)^{2.5}}$ and $\rho_k = \frac{(1-p)k}{4p^2(k+1)^{1.5}}$. Then,

$$\min_{t=0, \dots, k} \mathbb{E} \sum_{i=1}^h \|\mathbf{k}^i(t)\|^2 = O\left(\frac{\log(k)}{k^{2.5}}\right):$$

A.3 Corollary 3.6

Corollary 3.6 presented the convergence rate in expectation of RSGDA (Alg.2) given a specified precision $\epsilon > 0$ for an arbitrary minibatch size.

Corollary A.3. Consider the setting of Proposition 3.1. Choose following step sizes

$$\eta = \frac{2}{24L^2} \quad \text{and} \quad \beta = \frac{2}{12L^2} \min \left\{ \frac{2}{12L^2}; \frac{1-p}{48} \frac{p^3}{3^3 L^3}; \frac{r}{p} \frac{1-p}{12} \frac{p^2}{2^2 L^2}; \frac{r}{8^2 L^2} \frac{1-p}{p(2p + \frac{(1-p)^2}{24} \frac{p^2}{L^2})} \right\};$$

and a number of iterations k which verifies

$$k \geq \frac{12}{\epsilon^2} \left(\frac{12D_0 L^2}{p^2}; \frac{48^3 D_0^3 L^3}{(1-p)^3}; \frac{12^2 D_0^2 L^2}{p(1-p)^2}; \frac{8^2 L^2}{1-p} \frac{2 + \frac{(1-p)^2}{24} \frac{p^2}{L^2}}{p}; \frac{2L^2 r_0}{(1-p)^2} \right);$$

Then, by choosing $\beta \geq \frac{1}{2}$, we have that finding an ϵ -stationary point requires at most $O(k^3)$ iterations.

A.4 Corollary 3.8

Corollary 3.8 presented the convergence rate in expectation of RSGDA (Alg.2) given a specified precision $\epsilon > 0$ for a large enough minibatch size.

Corollary A.4 (Large minibatch sizes) Consider the setting of proposition 3.1. Let Assumption 3.7 hold. Choose the step sizes

$$\eta = \frac{1}{L} \quad \text{and} \quad \beta = \frac{1-p}{2^2 L^2} \frac{1-p}{p(2p + \frac{1-p}{L})};$$

There exists a minibatch size $M(\epsilon; p)$ such that if the total number of stochastic gradient evaluations is larger than

$$O \left(\frac{1}{\epsilon^2} \frac{2L D_0}{1-p} \frac{2 + \frac{1-p}{p} + \frac{L^2 r_0}{1-p}}{2} \max\{1; M(\epsilon; p)\} \right);$$

then, $\min_{t=0, \dots, k-1} E[kr_t - f^*] \leq \epsilon$. Moreover, choosing $\beta \geq \frac{1}{2}$ ensures that the total number stochastic gradient computations is $O(k^3)$.

B Proofs

B.1 Proof of Proposition 3.1

Proof. In this proof, we define $\mathbf{x}_k^+ = \mathbf{x}_k - \beta_k \mathbf{r}_k f(\mathbf{x}_k; \mathbf{v}_k; \mathbf{z}_k)$ and $\mathbf{v}_k^+ = \mathbf{v}_k + \beta_k \mathbf{r}_k \nabla f(\mathbf{x}_k; \mathbf{v}_k; \mathbf{z}_k)$. We denote by $E_{\mathbf{z}_k}[\cdot]$ the expectation conditioned on the random variable \mathbf{z}_k , and by $E_k[\cdot]$ the expectation conditioned on all past random variables.

Recall that

$$(\mathbf{x}_{k+1}; \mathbf{v}_{k+1}) = \begin{cases} (\mathbf{x}_k^+; \mathbf{v}_k) & \text{w. p. } p \\ (\mathbf{x}_k; \mathbf{v}_k^+) & \text{w. p. } 1-p \end{cases}$$

We have

$$E_k[(\mathbf{x}_{k+1})] = p E_{\mathbf{z}_k}(\mathbf{x}_k^+) + (1-p)(\mathbf{x}_k); \tag{8}$$

From Lemma 2.2, we have that \tilde{v}_k is $2L$ -smooth. Hence,

$$\langle \tilde{v}_k^+, \tilde{v}_k \rangle - \langle \tilde{v}_k, \tilde{v}_k \rangle \leq \frac{1}{2} L \|\tilde{v}_k - \tilde{v}_k^+\|^2 + \frac{1}{2} L \|\tilde{v}_k + \tilde{v}_k^-\|^2.$$

Thus,

$$\begin{aligned} E_{z_k} \langle \tilde{v}_k^+, \tilde{v}_k \rangle &= \langle \tilde{v}_k, \tilde{v}_k \rangle + \frac{1}{2} L E_{z_k} \|\tilde{v}_k - \tilde{v}_k^+\|^2 \\ &= \langle \tilde{v}_k, \tilde{v}_k \rangle + \frac{1}{2} L \|\tilde{v}_k - \tilde{v}_k^+\|^2 + \frac{1}{2} L \|\tilde{v}_k + \tilde{v}_k^-\|^2 \\ &\quad + \frac{1}{2} L E_{z_k} \|\tilde{v}_k - \tilde{v}_k^+\|^2 + \frac{1}{2} L \|\tilde{v}_k + \tilde{v}_k^-\|^2 \\ &= \langle \tilde{v}_k, \tilde{v}_k \rangle + \frac{1}{2} L \|\tilde{v}_k - \tilde{v}_k^+\|^2 + \frac{1}{2} L \|\tilde{v}_k + \tilde{v}_k^-\|^2 \\ &= \frac{1}{2} (1 - 2L) \|\tilde{v}_k - \tilde{v}_k^+\|^2 + \frac{1}{2} L \|\tilde{v}_k + \tilde{v}_k^-\|^2 \\ &= \langle \tilde{v}_k, \tilde{v}_k \rangle + \frac{1}{2} L \|\tilde{v}_k - \tilde{v}_k^+\|^2 + \frac{1}{2} L \|\tilde{v}_k + \tilde{v}_k^-\|^2 \\ &= \frac{1}{4} L \|\tilde{v}_k - \tilde{v}_k^+\|^2; \end{aligned}$$

where we used the fact that $\|\tilde{v}_k - \tilde{v}_k^+\| \leq \frac{1}{4L}$ in the last inequality.

Let $\tilde{v}_k \stackrel{\text{def}}{=} \tilde{v}_k - \tilde{v}_k^+$. From assumption 2.1, F is L -smooth. Hence,

$$E_{z_k} \langle \tilde{v}_k^+, \tilde{v}_k \rangle \leq \langle \tilde{v}_k, \tilde{v}_k \rangle + \frac{1}{2} L \|\tilde{v}_k - \tilde{v}_k^+\|^2 + \frac{1}{4} L \|\tilde{v}_k + \tilde{v}_k^-\|^2.$$

Using this inequality in (8), we have

$$E_k [\langle \tilde{v}_{k+1}, \tilde{v}_k \rangle] \leq \langle \tilde{v}_k, \tilde{v}_k \rangle + \frac{1}{2} L \|\tilde{v}_k - \tilde{v}_k^+\|^2 + \frac{1}{4} L \|\tilde{v}_k + \tilde{v}_k^-\|^2.$$

Besides, we have

$$E_k [\langle \tilde{v}_{k+1}^+, \tilde{v}_k \rangle] = (1 - p) E_{z_k} \langle \tilde{v}_k^+, \tilde{v}_k \rangle + p E_{z_k} \langle \tilde{v}_k, \tilde{v}_k^+ \rangle.$$

First, note that since $\tilde{v}_k^+ = \tilde{v}_k - \tilde{v}_k^+$ and $\tilde{v}_k \in V$, where V is a convex set, we have that

$$\langle \tilde{v}_k^+, \tilde{v}_k \rangle = \langle \tilde{v}_k - \tilde{v}_k^+, \tilde{v}_k \rangle = \|\tilde{v}_k\|^2 - \langle \tilde{v}_k^+, \tilde{v}_k \rangle,$$

where we used the fact that \tilde{v}_k is contractive. Hence, using the classical analysis for SGD in the strongly convex and smooth setting (see for example Gower et al. (2019)), we have that

$$E_{z_k} \langle \tilde{v}_k^+, \tilde{v}_k \rangle \leq (1 - p) \|\tilde{v}_k\|^2 + 2p(1 - 2L) \|\tilde{v}_k - \tilde{v}_k^+\|^2 + 2 \frac{1}{k} \|\tilde{v}_k - \tilde{v}_k^+\|^2.$$

Hence, with $\|\tilde{v}_k - \tilde{v}_k^+\| \leq \frac{1}{2L}$, we have

$$E_k [\langle \tilde{v}_{k+1}^+, \tilde{v}_k \rangle] \leq (1 - p) \|\tilde{v}_k\|^2 + p E_{z_k} \langle \tilde{v}_k, \tilde{v}_k^+ \rangle + 2(1 - p) \frac{1}{k} \|\tilde{v}_k - \tilde{v}_k^+\|^2.$$

Let $\tilde{v}_k > 0$. Then,

$$E_k [\langle \tilde{v}_{k+1}, \tilde{v}_k \rangle] \leq ((1 - p)(1 - \tilde{v}_k) + p(1 + \tilde{v}_k)) \|\tilde{v}_k\|^2 + \frac{1}{k} E_{z_k} \langle \tilde{v}_k, \tilde{v}_k^+ \rangle + 2(1 - p) \frac{1}{k} \|\tilde{v}_k - \tilde{v}_k^+\|^2.$$

From Lemma 2.2, we have that \tilde{v}_k is L -lipschitz. Hence, by also using Assumption 2.3,

$$E_{z_k} \langle \tilde{v}_k, \tilde{v}_k^+ \rangle \leq \frac{1}{2} L \|\tilde{v}_k - \tilde{v}_k^+\|^2 + \frac{1}{2} L \|\tilde{v}_k + \tilde{v}_k^-\|^2 + \frac{1}{k} \|\tilde{v}_k - \tilde{v}_k^+\|^2.$$

Hence,

$$E_k [k_{k+1}] \leq ((1-p)(1-k) + p(1+k))k + p \left(1 + \frac{1}{k}\right) \frac{2}{k} \sum_{k'}^2 \text{kr} F(k; v_k) k^2 + 2(1-p) \frac{2}{k} \sum_{k'}^2 + p \left(1 + \frac{1}{k}\right) \frac{2}{k} \sum_{k'}^2.$$

With $k = \frac{(1-p)k}{2p}$, this inequality becomes

$$E_k [k_{k+1}] \leq \left(1 + \frac{(1-p)k}{2}\right)k + \frac{p(2p+(1-p)k)}{(1-p)k} \frac{2}{k} \sum_{k'}^2 \text{kr} F(k; v_k) k^2 + 2(1-p) \frac{2}{k} \sum_{k'}^2 + \frac{p(2p+(1-p)k)}{(1-p)k} \frac{2}{k} \sum_{k'}^2.$$

Rearranging, we have

$$k \frac{2}{(1-p)k} k \frac{2}{(1-p)k} E_k [k_{k+1}] + \frac{2p(2p+(1-p)k)}{(1-p)^2 \frac{2}{k} \sum_{k'}^2} \sum_{k'}^2 \text{kr} F(k; v_k) k^2 + \frac{4}{k} \sum_{k'}^2 + \frac{2p(2p+(1-p)k)}{(1-p)^2 \frac{2}{k} \sum_{k'}^2} \sum_{k'}^2.$$

Hence,

$$\frac{p}{2} \frac{k L^2}{k} \frac{p}{(1-p)k} k \frac{p}{(1-p)k} E_k [k_{k+1}] + \frac{p^2(2p+(1-p)k)}{(1-p)^2 \frac{2}{k} \sum_{k'}^2} \frac{3}{k} \sum_{k'}^4 \text{kr} F(k; v_k) k^2 + 2 \frac{p}{k} \frac{k L}{k} \sum_{k'}^2 + \frac{p^2(2p+(1-p)k)}{(1-p)^2 \frac{2}{k} \sum_{k'}^2} \frac{3}{k} \sum_{k'}^4 + \frac{p}{(1-p)k} k \frac{p}{(1-p)_{k+1}} E_k [k_{k+1}] + \frac{p^2(2p+(1-p)k)}{(1-p)^2 \frac{2}{k} \sum_{k'}^2} \frac{3}{k} \sum_{k'}^4 \text{kr} F(k; v_k) k^2 + 2 \frac{p}{k} \frac{k L}{k} \sum_{k'}^2 + \frac{p^2(2p+(1-p)k)}{(1-p)^2 \frac{2}{k} \sum_{k'}^2} \frac{3}{k} \sum_{k'}^4;$$

where we used in the last inequality that $\frac{k+1}{k+1} = \frac{k}{k}$. Using this inequality in (B.1) and rearranging gives

$$\frac{p}{2} \frac{k}{k} \text{kr} F(k) k^2 + E_k F(k_{k+1}) + \frac{p}{(1-p)_{k+1}} \frac{L}{k+1} F(k) + \frac{p}{(1-p)k} k \frac{k p}{4} \left(1 + \frac{4p(2p+(1-p)k)}{(1-p)^2 \frac{2}{k} \sum_{k'}^2}\right) \sum_{k'}^2 \text{kr} F(k; v_k) k^2 + p \frac{2}{k} L \sum_{k'}^2 + 2 \frac{p}{k} \frac{k L}{k} \sum_{k'}^2 + \frac{p^2(2p+(1-p)k)}{(1-p)^2 \frac{2}{k} \sum_{k'}^2} \frac{3}{k} \sum_{k'}^4.$$

Hence, since $k = \frac{p}{4} \frac{(1-p)k}{p(2p+(1-p)k)}$,

$$p \frac{k}{k} \text{kr} F(k) k^2 + E_k F(k_{k+1}) + \frac{2p}{(1-p)_{k+1}} \frac{L}{k+1} F(k) + 2 F(k) + \frac{2p}{(1-p)k} \frac{k L}{k} + 2p \frac{2}{k} L \sum_{k'}^2 + 4 \frac{p}{k} \frac{k L}{k} \sum_{k'}^2 + \frac{2p^2(2p+(1-p)k)}{(1-p)^2 \frac{2}{k} \sum_{k'}^2} \frac{3}{k} \sum_{k'}^4.$$

Hence,

$$p \frac{k}{k} \text{kr} F(k) k^2 + 2E_k [E_{k+1}] - 2E_k + 2 \sum_{k'}^2 + p \frac{2}{k} L + \frac{p^2(2p+(1-p)k)}{(1-p)^2 \frac{2}{k} \sum_{k'}^2} \frac{3}{k} \sum_{k'}^4 + 4 \frac{p}{k} \frac{k L}{k} \sum_{k'}^2:$$

□

B.2 Proof of Corollary 3.2

Before the proof of Corollary 3.2, we first present a simplified version of the Robbins-Siegmund theorem.

Lemma B.1 (Robbins and Siegmund (1971)) Consider a filtration $(\mathcal{F}_k)_k$, the nonnegative sequences of $(F_k)_k$ adapted processes $(V_k)_k$, $(U_k)_k$ and $(Z_k)_k$ such that $\sum_k Z_k < \infty$ almost surely, and

$$\forall k \geq N; E[V_{k+1} | \mathcal{F}_k] + U_{k+1} \leq V_k + Z_k;$$

Then $(V_k)_k$ converges and $\sum_k U_k < \infty$ almost surely.

We now move on to the proof of Corollary 3.2.

Proof. In this proof, we use a similar proof technique to Sebbouh et al. (2021). From Proposition 3.1, we have

$$E_k - E_{k+1} \leq \alpha_k^2 k^2 + \frac{2E_k[E_{k+1}]}{p} - \frac{2E_k}{p} + 4\alpha_k k L^2 + 2\alpha_k^2 L^2 + \frac{p(2p+(1-\alpha_k)k)^3}{(1-\alpha_k)^2 k^4};$$

Using Lemma B.1 and the fact that $\sum_k \alpha_k^2 k^2 < \infty$, $\sum_k \alpha_k k L^2 < \infty$, $\sum_k \frac{3}{k^2} < \infty$, $\sum_k \frac{3}{k^4} < \infty$, we have that $(E_k)_k$ converges almost surely. Now define for all $k \geq N$,

$$w_k = \frac{2\alpha_k}{k}; \quad g_0 = \alpha_0 k^2; \quad g_{k+1} = (1 - w_k)g_k + w_k \alpha_k k^2;$$

Notice that since $(\alpha_k)_k$ is decreasing, we have $w_k \in [0, 1]$. Hence, using the convexity of the squared norm, we have

$$\frac{1}{2} \sum_{j=0}^k g_{k+1} + \frac{2E_k[E_{k+1}]}{p} + \frac{\alpha_k}{2} g_k - \frac{1}{2} \sum_{j=0}^{k-1} g_k + \frac{2E_k}{p} + 4\alpha_k k L^2 + 2\alpha_k^2 L^2 + \frac{p(2p+(1-\alpha_k)k)^3}{(1-\alpha_k)^2 k^4};$$

Using Lemma B.1 and the step size conditions again, and the fact that $(E_k)_k$ converges almost surely, gives that $\sum_{j=0}^k g_{k+1}$ converges almost surely and that $\sum_k \alpha_k g_k < \infty$ almost surely. In particular, this implies that $\lim_k \alpha_k g_k = 0$. Notice that $\alpha_k g_k = \frac{1}{k} \sum_{j=0}^{k-1} g_k$. Hence, since we have that $\sum_{j=0}^k g_{k+1}$ converges almost surely and $\sum_k \frac{1}{k} = \infty$ (which is a consequence of the fact that $\sum_k \alpha_k = 1$), then $\lim_k \frac{1}{k} \sum_{j=0}^{k-1} g_k = 0$, i.e.

$$g_k = o\left(\frac{1}{k}\right)$$

Finally, since g_k is a weighted average of $\alpha_0 k^2; \dots; \alpha_{k-1} k^2$, we have that $g_k \geq \min_{t=0, \dots, k-1} \alpha_t k^2$. Hence,

$$\min_{t=0, \dots, k-1} \alpha_t k^2 = o\left(\frac{1}{k}\right);$$

□

B.3 Proofs for convergence rates in expectation

All the convergence in expectation proofs follow by telescopic cancellation in (2). Indeed, summing Inequality (2) between $t = 0$ and $k - 1$, and using the fact that $\sum_{t=0}^{k-1} \mathbb{E} \| \mathbf{g}_t \|^2 \leq \min_{j=0, \dots, k-1} \mathbb{E} \| \mathbf{g}_j \|^2$, we have

$$\begin{aligned} \min_{t=0, \dots, k-1} \mathbb{E} \| \mathbf{g}_t \|^2 &\leq \frac{2(\|\mathbf{v}_0\|)^2}{p^{k-1}} + \frac{2D_0 L \|\mathbf{v}_0\|}{(1-p)^0} \frac{1}{p^{k-1}} + \frac{4L^2 \tilde{\rho}^{k-1}}{p^{k-1}} \\ &+ \frac{2L^2 \tilde{\rho}^{k-1}}{p^{k-1}} + \frac{2^2 4p}{(1-p)^2} \frac{1}{p^{k-1}} \sum_{t=0}^{k-1} 2p \frac{1}{t^3} + (1-p) \sum_{t=0}^{k-1} \frac{1}{t^3} \end{aligned} \quad (9)$$

B.3.1 Proof of Corollary 3.4

Using constant step sizes in (9), we have

$$\min_{t=0, \dots, k-1} \mathbb{E} \| \mathbf{g}_t \|^2 \leq \frac{2D_0}{pk} + \frac{2Lr_0}{(1-p)k} + 4L^2 \tilde{\rho}^{k-1} + 2L^2 + \frac{2^2 4p^2}{(1-p)^2} + \frac{2^2 3Lp}{(1-p)} \quad (10)$$

Proof. When using the exact gradients, we have that $\tilde{\rho} = \rho^2 = 0$. Using the constant step sizes of Corollary 3.4 in (10) directly gives the desired result. \square

B.3.2 Proof of Corollary 3.5

Proof. Using Inequality (9) and the parameter settings of Corollary 3.5, we have

$$\begin{aligned} \min_{t=0, \dots, k-1} \mathbb{E} \| \mathbf{g}_t \|^2 &\leq \frac{2(\|\mathbf{v}_0\|)^2}{p^0 (k+1)^{2=5}} + \frac{2L \|\mathbf{v}_0\|}{(1-p)^0 (k+1)^{2=5}} \\ &+ \frac{4L^2 \tilde{\rho}^0 \log(k+1)}{(k+1)^{2=5}} + \frac{2L^2 \tilde{\rho}^0}{(k+1)^{2=5}} \\ &+ \frac{2^2 4p^2}{o(1-p)^2 (k+1)^{2=5}} \frac{2p \log(k+1)}{o} + (1-p) \quad ; \end{aligned}$$

where $o(\cdot)$ indicates that we omit the absolute constants arising in the summations. The asymptotically dominant term is $\frac{4L^2 \tilde{\rho}^0 \log(k+1)}{(k+1)^{2=5}}$. Hence

$$\min_{t=0, \dots, k-1} \mathbb{E} \| \mathbf{g}_t \|^2 = O\left(\frac{\log(k)}{k}\right)$$

\square

B.3.3 Proof of Corollary 3.6

Proof. For simplicity, we consider that $\tilde{\rho} \leq \frac{1}{2}$. Otherwise, we can simply replace $\tilde{\rho}$ and $\tilde{\rho}^2$ by $\max\{\tilde{\rho}, \frac{1}{2}\}$, and the proof will still hold with $\max\{\tilde{\rho}, \frac{1}{2}\}$ instead of $\tilde{\rho}$.

Let $\epsilon > 0$. The proof follows simply from forcing each term of the LHS of (10) to be smaller than $\frac{\epsilon}{6}$. This

results in the following step sizes

$$= \frac{2}{24L^2} \quad \text{and} \quad = \frac{2}{12L^2} \min \left\{ \frac{2}{12L^2}; \frac{1-p}{p} \frac{8^2 L^3}{48^3}; \frac{1-p}{p} \frac{2}{12L^2}; \frac{1-p}{8^2 L^2 p^2 + \frac{(1-p)^2}{24L^2}} \right\};$$

and the following lower bound on k

$$k \geq \frac{12^2 D_0 L^2}{p^2}; \frac{48^3 D_0^3 L^3}{(1-p)^3}; \frac{12^2 D_0^2 L^2}{p(1-p)^2}; \frac{8^2 L^2 + \frac{(1-p)^2}{24L^2}}{1-p}; \frac{2L^2 r_0}{(1-p)^2};$$

By choosing $p \geq \frac{1}{2}$, we have that the RHS is at most of the order $\frac{3}{5}$. □

B.3.4 Proof of Corollary 3.8

To prove Corollary 3.8, we need to use the following lemma, which ensures that the variance of the stochastic gradients decreases linearly with the minibatch size.

Lemma B.2 (Lemma A.2 in Lin et al. (2019)). Let Assumption 3.7 hold. Then, is $G_z(\cdot; v) = \frac{1}{M} \sum_{i=1}^M r_i f(\cdot; v; z^i)$ where z_1, \dots, z_M are sampled i.i.d, then

$$E \|k G_z(\cdot; v) - r F(\cdot; v)\|^2 \leq \frac{2}{M};$$

And the same holds with the gradient with respect to θ .

Proof. Using (10) with the constant step sizes of Corollary 3.8 and Lemma

$$\min_{t=0, \dots, k-1} E \|k r_{(t)} - k^2\|^2 \leq \frac{8^2 L^2}{(1-p)k} \frac{1-p}{2p} D_0 + \frac{2L^2}{(1-p)k} r_0 + \frac{2^2}{M} \frac{1-p}{2} \frac{1-p}{2p + \frac{1-p}{2}} + 2 + \frac{p}{2^2} \frac{1-p}{2p + \frac{1-p}{2}} + \frac{1-p}{4} \frac{1-p}{2p + \frac{1-p}{2}} A;$$

Hence, since $\frac{p}{2^2(2p + \frac{1-p}{2})} + \frac{1-p}{4(2p + \frac{1-p}{2})} \leq 2$, we have

$$\min_{t=0, \dots, k-1} E \|k r_{(t)} - k^2\|^2 \leq \frac{8^2 L^2}{(1-p)k} \frac{1-p}{2p} D_0 + \frac{2L^2}{(1-p)k} r_0 + \frac{2^2}{M} \frac{1-p}{2} \frac{1-p}{2p + \frac{1-p}{2}} + 4A;$$

Choosing $k \geq \frac{2}{1-p} \frac{1-p}{2p} D_0 + \frac{2L^2}{(1-p)k} r_0$ and $M = M(\cdot; \cdot)$ ensures that

$$\frac{8^2 L^2}{(1-p)k} \frac{1-p}{2p} D_0 + \frac{2L^2}{(1-p)k} r_0 \leq \frac{2}{2} \quad \text{and} \quad \frac{2^2}{M} \frac{1-p}{2} \frac{1-p}{2p + \frac{1-p}{2}} + 4A \leq \frac{2}{2};$$

Hence, the total number of samples required to guarantee that $\min_{t=0, \dots, k-1} E \|k r_{(t)} - k^2\| \leq \epsilon$ is

$$\max_{k \in \mathbb{N}} \mathbb{E} \left[\sum_{t=0}^k \text{kr} \left(\frac{1}{k} \right) k^2 \right] = O \left(\frac{2L D_0}{1-p} \frac{1}{2} + \frac{L^2 r_0}{1-p} \frac{1}{2} \right) \max_{k \in \mathbb{N}} \mathbb{E} \left[\sum_{t=0}^k \text{kr} \left(\frac{1}{k} \right) k^2 \right] A :$$

□

B.4 Proofs for the interpolation results of Section 3.2.1 (Corollary 3.10)

B.4.1 Almost sure convergence

Proof. Using $\tilde{\alpha}^2 = 0$ and $k = \frac{1}{2L}$ in (2), we have that

$$\mathbb{E} \left[\sum_{t=0}^k \text{kr} \left(\frac{1}{k} \right) k^2 + 2 \mathbb{E}_k [E_{k+1}] \right] \leq 2 \mathbb{E}_k + 2 \frac{2}{k} \frac{1}{2} \frac{1}{p} L + \frac{4p^2 - 2p + (1-p)}{(1-p)^2} k^4 L^2 \frac{1}{A} :$$

Now note that by choosing $k = \frac{p}{4(k+1)^{1-2p}}$, we have that

$$\sum_{k=1}^{\infty} \frac{1}{k} < \infty \quad \text{and} \quad \sum_{k=1}^{\infty} \frac{1}{k^2} < \infty \quad \text{and} \quad \sum_{k=1}^{\infty} \frac{1}{k^3} < \infty :$$

Thus, proceeding as in the proof of Corollary 3.2, but with different choices of step sizes, we have that

$$\min_{t=0, \dots, k-1} \text{kr} \left(\frac{1}{k} \right) k^2 = o(k^{-1-2p}) :$$

□

B.4.2 Anytime convergence in expectation

Proof. Using $\tilde{\alpha}^2 = 0$ and $k = \frac{1}{2L}$ in (2), we have that

$$\begin{aligned} \min_{t=0, \dots, k-1} \mathbb{E} \left[\sum_{t=0}^k \text{kr} \left(\frac{1}{k} \right) k^2 \right] &\leq \frac{2 \left(\frac{D_0}{p} \right)}{p \frac{1}{k-1} \frac{1}{t}} + \frac{4 D_0 L^2 k v_0}{(1-p) p \frac{1}{k-1} \frac{1}{t}} \\ &+ \frac{2L^2}{p \frac{1}{k-1} \frac{1}{t}} + \frac{2^2 4p}{(1-p)^2 p \frac{1}{k-1} \frac{1}{t}} 8pL^2 \sum_{t=0}^k \frac{1}{t^3} + 2(1-p)L \sum_{t=0}^k \frac{1}{t^3} : \end{aligned}$$

Hence, choosing $k = \frac{p}{4(k+1)^{1-2p}}$, we have, omitting absolute constants. As is the case in the proof of Corollary 3.5, the asymptotically dominant term in the RHS is $\frac{1}{p} \frac{1}{k-1} \frac{1}{t}$. Hence,

$$\min_{t=0, \dots, k-1} \mathbb{E} \left[\sum_{t=0}^k \text{kr} \left(\frac{1}{k} \right) k^2 \right] = O \left(\frac{\log(k)}{k+1} \right) :$$

□

B.4.3 Convergence in expectation for a given precision.

Proof. In this proof, we proceed as in the proof of Corollary 3.6. For simplicity, we consider that $\tilde{\alpha}^2 = 1$. Otherwise, we can simply replace $\tilde{\alpha}^2$ by 1, and the proof will still hold with $\tilde{\alpha}^2$ instead of 1.

Using $\tilde{\alpha}^2 = 0$ and $k = \frac{1}{2L}$ in (10), we have

$$\min_{t=0, \dots, k-1} \mathbb{E} \left[\sum_{t=0}^k \text{kr} \left(\frac{1}{k} \right) k^2 \right] \leq \frac{2D_0}{pk} + \frac{4L^2 r_0}{(1-p)k} + 2L^2 + \frac{8^2 4L^2 p^2}{(1-p)^2} + \frac{4^2 3L^2 p^2}{(1-p)^2} :$$

Let $\epsilon > 0$. The proof follows simply from forcing each term of the LHS of the previous inequality to be smaller than $\frac{\epsilon}{5}$. This results in the following choice of ϵ ,

$$\epsilon = \min \left\{ \frac{8}{10L^2}; \frac{p(1-p)}{2^{10}Lp}; \frac{r}{p} \frac{1-p}{2^5 p^2 L}; \frac{1-p}{8^2 L p^{2p+1-p}}; \frac{9}{8^2 L p^{2p+1-p}} \right\};$$

and the following lower bound on k :

$$k \geq \max \left\{ \frac{8}{p^2}; \frac{100D_0 L^2}{(1-p)^3}; \frac{20^p D_0^p L}{p(1-p)^4}; \frac{80D_0^2 L}{(1-p)^2 \left(2 + \frac{1-p}{p}\right)}; \frac{20L^2 r_0}{(1-p)^2}; \frac{9}{(1-p)^2} \right\};$$

By choosing $p \geq \frac{1}{2}$, we have that the RHS is at most of the order ϵ^2 . □

C Learning with a semi-discrete Sinkhorn loss is a nonconvex-strongly concave problem

The goal of this Section is to present the details and proof of Lemma 4.2. Let us first recall the problem setting presented in §4.2 in more detail.

C.1 Problem setting

Let $X \subseteq \mathbb{R}^p$. For all $n \geq 2$, define $p[n] \stackrel{\text{def}}{=} p_1, \dots, p_n$. Assume that we are given a fixed dataset $(y_1, \dots, y_n) \in X^n$ and that $\sum_{j=1}^n p_j = 1$, where $p_j = \frac{1}{n}$ (we generalize the setting of §4.2 to non-uniform probabilities). Then, from Genevay et al. (2016), we have that for any $\mu \in \mathcal{P}(X)$,

$$W(\mu; \nu) = \max_{\nu \in \mathcal{P}(X)} E_{z \sim \nu} [h(z; \mu)];$$

where

$$h(x; \nu) = \sum_{j=1}^n p_j \log \sum_{i=1}^n \exp \left(\frac{p_i}{p_j} c(x; y_i) \right) + \nu_j;$$

For a mapping $g(\cdot; \mu) : Z \rightarrow X$, we have

$$W(g(\cdot; \mu); \nu) = \max_{\nu \in \mathcal{P}(X)} E_{z \sim \nu} [h(g(z; \mu); \mu)];$$

Thus, the problem of learning with a Sinkhorn loss can be formulated as

$$\min_{\mu \in \mathcal{P}(X)} \max_{\nu \in \mathcal{P}(X)} F(\mu; \nu) \stackrel{\text{def}}{=} E_{z \sim \nu} [h(g_z(\mu); \nu)]; \tag{11}$$

Assuming that a minimum exists, our goal is to find

$$\mu^* \stackrel{\text{def}}{=} \arg \min_{\mu \in \mathcal{P}(X)} F(\mu; \nu^*); \quad \text{where } \nu^* \stackrel{\text{def}}{=} \max_{\nu \in \mathcal{P}(X)} F(\mu^*; \nu);$$

C.2 Assumptions and consequences

We now detail the assumptions of Lemma 4.2.

Assumption C.1. We have

- for all $z \in Z$; $\forall! c(g_z(\mu); y)$ and $y \in \mathcal{Y}$; $c(g_z(\mu); y)$ are L_c -lipschitz for k_2 and k_1 respectively a.s.

2. for all $z \in Z$ and $y \in Y$, $\nabla^2 c(g_z(\cdot); y)$ is twice differentiable and L_c -smooth a.s.

Assumption C.1 has the following consequences on the function defined in (11).

Lemma C.2. $\forall \nu \nabla^2 F(\cdot; \nu)$ is L -smooth for all $\nu \in \mathbb{R}^d$. Moreover, if assumption C.1 holds, then

1. $\nabla^2 F(\cdot; \nu)$ is L_c -lipschitz and L_F -smooth, where $L_F \stackrel{\text{def}}{=} L_c + 2L_c^2$.
2. $\forall \nu \nabla^2 F(\cdot; \nu)$ and $\nabla^2 \log F(\cdot; \nu)$ are $2L_c$ -lipschitz.

Before proving Lemma C.2, we present some notations and preliminary calculations that will be used throughout the proofs.

Notations. We define for all $\nu \in \mathbb{R}^d, z \in Z, x, y \in X$

$$c_j(x) \stackrel{\text{def}}{=} c(x; y_j)$$

$$w_j(x) \stackrel{\text{def}}{=} \frac{\exp(\nu_j c_j(x))}{\sum_{k=1}^n \exp(\nu_k c_k(x))}$$

$$c(g_z(\cdot); y) = [c_1(g_z(\cdot)); \dots; c_n(g_z(\cdot))]^\top$$

Closed form gradients. Let $\nu \in \mathbb{R}^d, z \in Z$ and $y \in X$. We have

$$\nabla_\nu h(g_z(\cdot); \nu) = \frac{\exp(\nu c(g_z(\cdot); y))}{\sum_{k=1}^n \exp(\nu_k c_k(g_z(\cdot); y))} \in \mathbb{R}^n; \tag{12}$$

and

$$\nabla_x^2 h(g_z(\cdot); \nu) = \sum_{j=1}^n \frac{\exp(\nu_j c_j(g_z(\cdot)))}{\sum_{k=1}^n \exp(\nu_k c_k(g_z(\cdot)))} \nabla_x^2 c_j(g_z(\cdot)) \in \mathbb{R}^{d \times d};$$

Proof of Lemma C.2. Let Assumption C.1 hold. Let $y \in X$.

$\forall \nu \nabla^2 F(\cdot; \nu)$ is L -smooth.

Let $\nu \in \mathbb{R}^d, z \in Z$. Define $a \stackrel{\text{def}}{=} \exp(\nu c(g_z(\cdot); y))$. Differentiating $\nabla_\nu h(g_z(\cdot); \nu)$ twice, we have

$$\nabla_\nu^2 h(g_z(\cdot); \nu) = \frac{1}{(a^\top \mathbf{1}_n)^2} \frac{\text{diag}(a)}{a^\top \mathbf{1}_n}$$

Using the Cauchy-Schwartz inequality, we can show that $\nabla_\nu^2 h(x; \nu) \succeq 0$. Moreover, we have for all $b \in \mathbb{R}^n$,

$$\frac{b^\top \text{diag}(a) b}{a^\top \mathbf{1}_n} = \sum_{i=1}^n \frac{a_i}{a_j} b_i^2 \leq \sum_{j=1}^n b_j^2;$$

that is, $\frac{\text{diag}(a)}{a^\top \mathbf{1}_n} \preceq I$. This, together with the fact that $\frac{aa^\top}{(a^\top \mathbf{1}_n)^2} \succeq 0$, implies that $\mathbf{1}_n^\top \nabla_\nu^2 h(x; \nu) \succeq 0$. This in turn implies that $\forall \nu \nabla^2 F(\cdot; \nu)$ is L -smooth.

7! $F(\cdot; v)$ is $\frac{L_c + 2L_c^2}{2}$ -smooth.

Let $z \in \mathbb{R}^d, v \in \mathbb{R}^n, z \in \mathcal{Z}$. We have

$$\begin{aligned} r^2 h(g_z(\cdot); v) &= \frac{1}{2} \sum_{j=1}^n w_j(g_z(\cdot)) \left(r^2 c_j(g_z(\cdot)) + (r c_j(g_z(\cdot)) - c_j(g_z(\cdot)))^2 \right) + \sum_{k=1}^n w_k(g_z(\cdot)) (r c_k(g_z(\cdot)) - c_k(g_z(\cdot)))^2 \\ &= \frac{1}{2} \sum_{k=1}^n w_k(g_z(\cdot)) \left(w_j(g_z(\cdot)) r^2 c_j(g_z(\cdot)) + (r c_k(g_z(\cdot)) - c_k(g_z(\cdot)))^2 \right) \end{aligned}$$

Let $b \in \mathbb{R}^d$. Then

$$\begin{aligned} r^2 h(g_z(\cdot); v) &\leq \frac{1}{2} \sum_{k=1}^n w_k(g_z(\cdot)) \left(\sum_{j=1}^n w_j(g_z(\cdot)) r^2 c_j(g_z(\cdot)) + (r c_k(g_z(\cdot)) - c_k(g_z(\cdot)))^2 \right) \end{aligned}$$

Using Cauchy-Schwarz and the fact that $c_j(\cdot)$ is L_c -lipschitz, we have

$$(r c_j(g_z(\cdot)) - c_j(g_z(\cdot)))^2 \leq (r c_j(g_z(\cdot)) - c_j(g_z(\cdot)))^2 \leq L_c^2 \|b\|^2$$

Moreover, since $c_j(\cdot)$ is L_c -smooth, we have $r^2 c_j(g_z(\cdot)) \leq L_c \|b\|^2$. Hence,

$$\begin{aligned} r^2 h(g_z(\cdot); v) &\leq \frac{2L_c^2 + L_c}{2} \|b\|^2 \sum_{k=1}^n w_k(g_z(\cdot)) \sum_{j=1}^n w_j(g_z(\cdot)) \\ &= \frac{2L_c^2 + L_c}{2} \|b\|^2 \end{aligned}$$

Taking the expectation and using Jensen's inequality gives

$$r^2 F(\cdot; v) \leq \mathbb{E}_z [r^2 h(g_z(\cdot); v)] \leq \frac{2L_c^2 + L_c}{2} \|b\|^2$$

This shows that $F(\cdot; v)$ is $\frac{2L_c^2 + L_c}{2}$ -smooth.

7! $r F(\cdot; v)$ and $r \nabla F(\cdot; v)$ are $2L_c$ -lipschitz.

Let $a \in \mathbb{R}^d, v \in \mathbb{R}^n, z \in \mathcal{Z}$. Let $a \in \mathbb{R}^d$. Differentiating $r \nabla h(g_z(\cdot); v)$, we have

$$\begin{aligned} r^2 \nabla h(g_z(\cdot); v) &= \frac{1}{2} \sum_{i=1}^n w_i(g_z(\cdot)) \sum_{k=1}^n w_k(g_z(\cdot)) (r c_i(g_z(\cdot)) - c_i(g_z(\cdot))) (r c_k(g_z(\cdot)) - c_k(g_z(\cdot))) \\ &= \frac{1}{2} \sum_{i=1}^n w_i(g_z(\cdot))^2 \sum_{k=1}^n w_k(g_z(\cdot)) (r c_i(g_z(\cdot)) - c_i(g_z(\cdot))) (r c_k(g_z(\cdot)) - c_k(g_z(\cdot))) \\ &\leq \frac{1}{2} \sum_{i=1}^n w_i(g_z(\cdot))^2 \sum_{k=1}^n w_k(g_z(\cdot)) (r c_i(g_z(\cdot)) - c_i(g_z(\cdot)))^2 + \frac{1}{2} \sum_{i=1}^n w_i(g_z(\cdot))^2 \sum_{k=1}^n w_k(g_z(\cdot)) (r c_k(g_z(\cdot)) - c_k(g_z(\cdot)))^2 \\ &\leq \frac{4L_c^2}{2} \|a\|^2 \sum_{i=1}^n w_i(g_z(\cdot))^2 \\ &\leq \frac{4L_c^2}{2} \|a\|^2 \end{aligned}$$

where we used Jensen's inequality in the first inequality, Cauchy-Schwarz in the second, the fact that $\theta \mapsto c_i(\theta)$ is L_c -lipschitz *a.s.* in the third, and $w_i(\theta)^2 \leq w_i(\theta)$ in the fourth. Note that the norm on the left hand side is in \mathbb{R}^n , while that on the right hand side is in \mathbb{R}^d . Since the squared norm is convex, using Jensen's inequality,

$$\nabla_{\theta v}^2 F(\theta, v) a^2 \leq \mathbb{E}_{z \sim \mu} \nabla_{\theta v}^2 h(g_z(\theta), v) a^2 \leq \frac{4L_c^2}{\epsilon^2} \|a\|^2.$$

We conclude that $\theta \mapsto \nabla_v F(\theta, v)$ is $\frac{2L_c}{\epsilon}$ -lipschitz for all $v \in \mathbb{R}^n$. Similarly, we have that $v \mapsto \nabla_{\theta} F(\theta, v)$ is $\frac{2L_c}{\epsilon}$ -lipschitz for all $\theta \in \mathbb{R}^d$.

□

C.3 Restricting the maximization problem

In this section, we give the explicit form of the set \mathcal{V} in Lemma 4.2. Indeed, we show that the problem $\max_{v \in \mathbb{R}^n} F(\theta, v)$ can be restricted to a smaller bounded set on which $v \mapsto F(\theta, v)$ is strongly-concave for all $\theta \in \mathbb{R}^d$.

Lemma C.3. *Let Assumption C.1 hold. Define*

$$\mathcal{V} \stackrel{\text{def}}{=} \left\{ v \in \mathbb{R}^n : \|v\|_2 \leq \frac{L_c}{n} \sum_{i,k=1}^n \|y_k - y_i\|_{\infty} \text{ and } \sum_{i=1}^n v_i = 0 \right\}. \quad (13)$$

Then, for all $\theta \in \mathbb{R}^d$,

$$\max_{v \in \mathbb{R}^n} F(\theta, v) = \max_{v \in \mathcal{V}} F(\theta, v), \quad (14)$$

and F admits a unique maximizer on \mathcal{V} .

Proof. Let $v^* \in \operatorname{argmax}_{v \in \mathbb{R}^n} F(\theta, v)$. Since $\theta \mapsto c(g_z(\theta), y)$ is L_c -lipschitz, we have that the (Sinkhorn) Kantorvitch potentials are L_c -Lipschitz on \mathcal{Y} as well (Genevay et al., 2019, Proposition 1). This means, in the discrete setting, that for any $v^* \in \operatorname{argmax}_{v \in \mathbb{R}^n} F(\theta, v)$ and $(k, i) \in [n]^2$,

$$|v_k^* - v_i^*| \leq L_c \|y_k - y_i\|_{\infty}.$$

Hence,

$$v_i^* - L_c \|y_k - y_i\|_{\infty} \leq v_k^* \leq v_i^* + L_c \|y_k - y_i\|_{\infty}.$$

Further, since the Kantorovitch potentials are uniquely defined upto a constant, we can restrict our attention to the *unique* vector $v^* \in \mathbb{R}^n$ which verifies $\sum_{i=1}^n v_i^* = 0$. Thus, summing the previous inequality from $i = 1$ to n and rearranging, we have

$$|v_k^*| \leq \frac{L_c}{n} \sum_{i=1}^n \|y_k - y_i\|_{\infty}.$$

Finally, summing between $k = 1$ and n and using the fact that $\|v\|_2 \leq \|v\|_1$, we have that

$$\|v^*\|_2 \leq \frac{L_c}{n} \sum_{i,k=1}^n \|y_k - y_i\|_{\infty}.$$

That is, the unique solution v^* to the maximization problem $\max_{v \in \mathbb{R}^n} F(\theta, v)$ which verifies $\sum_{i=1}^n v_i^* = 0$ also verifies the previous inequality. Thus, we can restrict the maximization to the set \mathcal{V} defined in (13) □

Remark C.4. Suppose we want to solve Problem (14) using Projected Stochastic Gradient Ascent. For simplicity, suppose here that $g(z, \theta) = z$ for all $z, \theta \in \mathcal{Z} \times \mathbb{R}^d$ almost surely. Suppose that we initialize at $v_0 = 0$ and update for all $k \in \mathbb{N}$, $v_{k+1} = \Pi_{\mathcal{V}}(v_k + \eta \nabla_v h(z_k, v_k))$, where $z_k \sim \mu$. From (12), we have that for all $z \in \mathcal{Z}$ and $v \in \mathbb{R}^n$, $\sum_{j=1}^n \partial_{v_j} h(z, v) = 0$, so that, since we initialize at $v_0 = 0$, we will have $\sum_{j=1}^n v_k^j = 0$ for all $k \in \mathbb{N}$. Hence, we can replace the set \mathcal{V} by $\tilde{\mathcal{V}} \stackrel{\text{def}}{=} \{v \in \mathbb{R}^n : \|v\|_2 \leq \frac{L_c}{n} \sum_{i,k=1}^n \|y_k - y_i\|_\infty\}$. Let $\beta = \frac{L_c}{n} \sum_{i,k=1}^n \|y_k - y_i\|_\infty$. The projection operator is then simply

$$\Pi_{\tilde{\mathcal{V}}}(v) = \frac{\beta v}{\max\{\|v\|, \beta\}}$$

for all $v \in \mathbb{R}^n$.

C.4 The semi-dual objective is strongly concave on the subset \mathcal{V}

We now show that $v \mapsto F(\theta, v)$ is strongly concave on \mathcal{V} for all $\theta \in \mathbb{R}^d$.

Proposition C.5. Let Assumption C.1 hold. Define

$$\mathcal{V} \stackrel{\text{def}}{=} \left\{ v \in \mathbb{R}^n : \|v\|_2 \leq \frac{L_c}{n} \sum_{i,k=1}^n \|y_k - y_i\|_\infty \text{ and } \sum_{i=1}^n v_i = 0 \right\}.$$

Let $\Delta_y^n = \max_{i,k \in [n]} \|y_k - y_i\|_\infty$. Then, the function $v \mapsto F(\theta, v)$, where F is defined in (11), is ξ -strongly concave on \mathcal{V} , with

$$\xi \stackrel{\text{def}}{=} \frac{\exp\left(\frac{-2(n+2)L_c \Delta_y^n}{\epsilon} \min_{k \in [n]} \|\cdot\|_k\right)}{2n\epsilon}.$$

Proof. Let $v, d \in \mathcal{V}$ and define $a = \exp\left(\frac{v-c(g_z(\theta), y)}{\epsilon}\right) \odot \cdot$. By differentiating $v \mapsto h(g_z(\theta), v)$ twice, we have, using standard computations for the log-sum-exp function:

$$\nabla_v^2 h(x, v) = \frac{1}{\epsilon} \frac{aa^\top}{(\mathbf{1}_n^\top a)^2} - \frac{\text{diag}(a)}{\mathbf{1}_n^\top a}.$$

Hence,

$$\begin{aligned} d^\top \nabla_v^2 h(x, v) d &= \frac{1}{\epsilon} \left(\frac{\sum_{i=1}^n a_i d_i^2}{(\mathbf{1}_n^\top a)^2} - \sum_{i=1}^n a_i d_i^2 \right) \\ &= \frac{1}{\epsilon} \left(\frac{\sum_{i=1}^n a_i d_i^2}{(\mathbf{1}_n^\top a)^2} - \frac{\sum_{i=1}^n a_i d_i^2}{\mathbf{1}_n^\top a} \right) \end{aligned}$$

Define $w_i = \frac{a_i}{\sum_{k=1}^n a_k}$. Then, we can rewrite the previous equality as

$$d^\top \nabla_v^2 h(x, v) d = \frac{1}{\epsilon} \left(\sum_{i=1}^n w_i d_i^2 - \sum_{i=1}^n w_i d_i^2 \right)$$

Since $\sum_{j=1}^n w_j = 1$, it is easy to show that

$$\sum_{i=1}^n w_i d_i^2 - \left(\sum_{i=1}^n w_i d_i \right)^2 = \frac{1}{2} \sum_{ij} w_i w_j (d_i - d_j)^2.$$

Hence,

$$d^\top \nabla_v^2 h(x, v) d = -\frac{1}{2\epsilon} \sum_{ij} w_i w_j (d_i - d_j)^2.$$

But we have for all $i \in [n]$, $w_i = \frac{\exp(\frac{v_i}{\epsilon}) \nu_i}{\sum_{k=1}^n \exp(\frac{v_k + c(x, y_i) - c(x, y_k)}{\epsilon}) \nu_k}$. And since $y \mapsto c(x, y)$ is L_c -lipschitz for all $x \in \mathcal{X}$, we have

$$|c(x, y_i) - c(x, y_k)| \leq L_c \|y_i - y_k\|_\infty \leq L_c \Delta_y^n, \quad \text{where } \Delta_y^n = \max_{i, k \in [n]} \|y_i - y_k\|_\infty.$$

Since $v \in \mathcal{V}$, we also have that

$$\|v\|_2 \leq \frac{L_c}{n} \sum_{i, k} \|y_k - y_i\|_\infty \leq n L_c \Delta_y^n.$$

Using these two inequalities to lower bound w_i , we have that

$$w_i \geq \frac{\exp\left(\frac{-(n+2)L_c \Delta_y^n}{\epsilon}\right) \min_{k \in [n]} \nu_k}{n}.$$

Hence,

$$d^\top \nabla_v^2 h(x, v) d \leq - \frac{\exp\left(\frac{-2(n+2)L_c \Delta_y^n}{\epsilon}\right) \min_{k \in [n]} \nu_k}{2n^2 \epsilon} \sum_{i, j} (d_i - d_j)^2.$$

But since $d \in \mathcal{V}$, we have that $\sum_{i=1}^n d_i = 0$. Hence; $\sum_{i, j} (d_i - d_j)^2 = 2n \sum_{i=1}^n d_i^2 = 2n \|d\|^2$. Hence,

$$d^\top \nabla_v^2 h(g_z(\theta), v) d \leq - \frac{\exp\left(\frac{-2(n+2)L_c \Delta_y^n}{\epsilon}\right) \min_{k \in [n]} \nu_k}{2n \epsilon} \|d\|^2.$$

This inequality holds for all $z \sim \mathcal{D}$ almost surely. Taking the expectation shows that $v \mapsto F(\theta, v)$ is strongly concave on \mathcal{V} for all $\theta \in \mathbb{R}^d$. \square

D Details about the experimental setting of Section 4.2

For ease of exposition, we redefine here the semi-dual function used in §4.2. Let $x \in \mathcal{X}$, and consider a point cloud $(y_j)_{j=1}^n \subseteq \mathcal{Y}$. Then we define for all $v \in \mathbb{R}^n$,

$$h(x, (y_j)_{j=1}^n; v) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{j=1}^n v_j - \epsilon \log \frac{1}{n} \sum_{j=1}^n \exp\left(\frac{v_j - c(x, y_j)}{\epsilon}\right) - \epsilon. \quad (15)$$

The goal is to solve the following problem

$$\min_{\theta \in \mathbb{R}^d} \mathcal{W}_\epsilon(\mu, \sum_{j=1}^n \delta_{f_\theta(y_j)}) \equiv \max_{v \in \mathbb{R}^n} \mathbb{E}_{x \sim \mu} h(x, (f_\theta(y_j))_{j=1}^n; v).$$

Thus, if $\mu = \frac{1}{m} \sum_{i=1}^m \delta_{x_i}$ for $(x_i)_{i=1}^m \subseteq \mathcal{X}$, we want to solve

$$\min_{\theta \in \mathbb{R}^d} \max_{v \in \mathbb{R}^n} \frac{1}{m} \sum_{i=1}^m h(x_i, (f_\theta(y_j))_{j=1}^n; v).$$

D.1 Dataset, architecture and OT hyperparameters

Dataset. The dataset we used was first considered in Stark et al. (2020). It consists of a single-cell profile of a metastatic melanoma sample from the Tumor Profiler Consortium¹. It contains single-cell data from a

¹<https://tpreports.nexus.ethz.ch/download/scim/data/tupro/>

cohort of patients which is analyzed using two technologies: Cytometry by Time of Flight (CyTOF, [Bandura et al. \(2009\)](#)) and scRNA-sequencing ([Tang et al., 2009](#)). The number of cells analyzed using CyTOF was $m = 135334$ and the dimension of the resulting points is $d = 41$, and the number of cells analyzed using scRNA-sequencing was $n = 4683$, with a dimension $d' = 1024$. We denote the CyTOF point cloud by $(x_i)_{i=1}^m$ and $\mu = \frac{1}{m} \sum_{i=1}^m \delta_{x_i}$, and the scRNA point cloud by $(y_j)_{j=1}^n$ and $\nu = \frac{1}{n} \sum_{j=1}^n \delta_{y_j}$.

Architecture. Instead of mapping both point clouds to a lower dimensional space, we instead map the smaller point cloud ν to the larger one μ . This strategy has two benefits: (i) it reduces the possible error due to learning the parametric map since we only incur the error for one point cloud, (ii) it is computationally less expensive, since we need to displace a smaller point cloud instead of both point clouds. The map we used was a 3-layer MLP with a number of hidden units 128-96-64, and we used GELU activation functions ([Hendrycks and Gimpel, 2016](#)). We used a minibatch size $b = 128$.

OT hyperparameters. We used a regularization parameter $\epsilon = 0.1$ and a quadratic cost function $c(x, y) = \|x - y\|^2$ for all $x, y \in \mathcal{X}$.

D.2 Using Sinkhorn’s algorithm

We now present Sinkhorn’s algorithm.

Algorithm 3 Sinkhorn($(x_i)_{i=1}^m, (y_j)_{j=1}^n, c, \epsilon, m_{sin}, a_0, b_0$)

Inputs: point clouds $(x_i)_{i=1}^m, (y_j)_{j=1}^n$, cost function c , regularization ϵ , number of iterations m . Optional: initialization $(a_0, b_0) \in \mathbb{R}_+^m \times \mathbb{R}_+^n$.

Compute $K \in \mathbb{R}^{n \times m}$, where $K_{i,j} = \exp -\frac{c(x_i, y_j)}{\epsilon}$ for all (i, j)

Initialisation: by default $a_0 = \mathbf{1}_m, b_0 = \mathbf{1}_n$

for $\ell = 0, \dots, m_{sin} - 1$ **do**

$$a_{\ell+1} = \frac{\mathbf{1}_m}{K b_\ell}, b_{k+1} = \frac{\mathbf{1}_n}{K a_{\ell+1}}$$

end for

return $(a_{m_{sin}}, b_{m_{sin}})$

In §4.2, we use Sinkhorn’s algorithm (Alg. 3) as a subroutine instead of the gradient ascent step in RSGDA (Alg. 2). We present the resulting algorithm explicitly in Alg. 4.

Algorithm 4 Learning using Sinkhorn’s algorithm

Inputs: $(x_i)_{i=1}^m, (y_j)_{j=1}^n$, cost function c , regularization ϵ , number of iterations m , parametric map f_θ .

$a_0 = \mathbf{1}_m, b_0 = \mathbf{1}_n, \theta_0 \in \mathbb{R}^d$, minibatch size b , step size α

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

Sample a minibatch $B \subset [n]$ of size $|B| = b$.

$$a_{k+1}, b_{k+1} = \text{Sinkhorn}((x_i)_{i \in B}, (y_j)_{j=1}^n, c, \epsilon, m_{sin}, a_k, b_k) \quad \triangleright \text{Using Alg. 3}$$

$$v_{k+1} = -\epsilon \log(b_{k+1})$$

$$\theta_{k+1} = \theta_k - \frac{\alpha}{b} \sum_{i \in B} h(x_i, (f_\theta(y_j))_{j=1}^n); v_{k+1} \quad \triangleright h \text{ defined in (15)}$$

end for

Another way to use Sinkhorn’s algorithm would be to initialize each algorithm with $a_0 = \mathbf{1}_m$ and $b_0 = \mathbf{1}_n$, but the resulting algorithm was not competitive.

D.2.1 Impact of the choice of the number of iterations of Sinkhorn’s algorithm

Here, we examine how the number of iterations of Sinkhorn’s algorithm m_{sin} should be set in Alg. 4. We find that we should only use one step, and that the more steps we use, the slower the optimization. This is due to the bias introduced by solving a minibatch version of the true transport problem $\mathcal{W}_\epsilon(\mu, \nu_{\theta_k})$ (7)

