Efficient Stochastic Gradient Descent for Distributionally Robust Learning

Soumyadip Ghosh
IBM Research AI
T.J. Watson Research Center
Yorktown Heights, NY 10598
ghoshs@us.ibm.com

Mark S. Squillante
IBM Research AI
T.J. Watson Research Center
Yorktown Heights, NY 10598
mss@us.ibm.com

Ebisa D. Wollega
Department of Engineering
Colorado State University-Pueblo
Pueblo, CO 81001
ebisa.wollega@cspueblo.edu

Abstract

We consider a new stochastic gradient descent algorithm for efficiently solving general min-max optimization problems that arise naturally in distributionally robust learning. By focusing on the entire dataset, current approaches do not scale well. We address this issue by initially focusing on a subset of the data and progressively increasing this support to statistically cover the entire dataset.

1 Introduction

The problem of distributionally robust learning has been an area of great interest in the machine learning community over the past few years. This class of problems includes a fundamental tradeoff between bias and variance, or equivalently between approximation error and estimation error. An important issue with distributionally robust learning concerns the scalability of the learning algorithms for very large datasets, especially since existing approaches are based on operating on the entire collection of data samples in each iteration. To address these fundamental issues, we propose and investigate a new stochastic gradient descent (SGD) algorithm to efficiently solve general large-scale distributionally robust learning optimization problems by sub-sampling the support of the decision variables. Out method does progressively increase this support so as to eventually cover the dataset, and we do so by optimally, in a strong statistical sense, balancing the computational effort with the required level of accuracy. Our approach supports a general class of distance measures as part of the robust formulation. We derive and establish various theoretical results for our approach using a combination of methods from mathematical optimization and mathematical statistics. We also present empirical results that demonstrate and quantify the significant benefits of our approach over previous work in the area. All proofs and additional technical materials are provided in the supplement.

1.1 Distributionally Robust Learning

Consider a general formulation of the distributionally robust optimization problem of active interest. Let $\mathbb{X}$ denote a sample space, P a probability distribution on $\mathbb{X}$, and $\Theta \subseteq \mathbb{R}^d$ a parameter space. Define $L_P(\theta) := \mathbb{E}_P[l(\theta, \xi)]$ to be the expectation with respect to (w.r.t.) P of a loss function $l : \Theta \times \mathbb{X} \to \mathbb{R}$ representing the estimation error for a learning model with parameters $\theta \in \Theta$ over data $\xi \in \mathbb{X}$. Further define the expected worst-case loss function $R(\theta) := \mathbb{E}_{P^*(\theta)}[l(\theta, \xi)] = \sup_{P \in \mathcal{P}}\{L_P(\theta)\}$, which

maximizes the loss L_P over a well-defined set of measures $\mathcal{P}$ that typically takes the form

$$\mathcal{P} = \left\{ P \mid D(P, P_b) \leq \rho, \quad \int dP(\xi) = 1, \quad P(\xi) \geq 0 \right\}, \tag{1}$$

where $D(\cdot, \cdot)$ is a metric on the space of probability distributions on $\mathbb{X}$ and where the constraints limit the feasible candidates to be within a distance ρ of a base distribution, denoted by P_b. We then seek to find a $\theta \in \Theta$ that, for a given $\mathbb{X}$ and $\mathcal{P}$, solves the distributionally robust optimization problem

$$R(\theta^*_{rob}) \;=\; \min_{\theta \in \Theta} \left\{ R(\theta) \right\} \;=\; \min_{\theta \in \Theta} \left\{ \sup_{P \in \mathcal{P}} \left\{ L_P(\theta) \right\} \right\}. \tag{2}$$

The solution to the above min-max formulation renders an expected loss performance R^* that is robust w.r.t. $\xi \in \mathbb{X}$ taking any $P \in \mathcal{P}$. Hence, (2) explicitly treats the identity of the true (unknown) data distribution, denoted by P_0, as being ambiguous. Note that the likelihood of $P^*(\theta^*_{rob}) \neq P_0$ is generally quite high, and thus the loss at θ^*_{rob} is likely to be higher than the loss at the optimal θ^* for P_0, were it to be known; meanwhile, since this is rarely the case, θ^*_{rob} still hedges the performance of a model to the uncertainty in P_0. Note further that this entire approach is as opposed to only solving for θ^*_{erm}, the loss performance L_P under a fixed data distribution, often the empirical distribution.

The formulation in (2) w.r.t. the set $\mathcal{P}$ in (1) captures numerous use cases with different metrics D. In special cases of the definition of (1), the solution to the inner maximization problem (2) may be explicitly available. One example is based on specific instances of Wasserstein distance metrics, where the solution $P^*(\theta)$ of the inner problem can be explicitly characterized, the objective function value $\mathbb{E}_{P^*(\theta)}[l(\theta, \xi)]$ is available in closed form, and (2) reduces to a standard stochastic optimization problem; refer to [2, 14] for examples along these lines. Our primary interest herein lies in the general ϕ-divergence class of distance measures

$$D_\phi(P, P_b) = \mathbb{E}_{P_b}\left[\phi\left(\frac{dP}{dP_b}\right)\right], \tag{3}$$

where $\phi(t)$ is a non-negative convex function that takes a value of 0 only at $t = 1$. While not explicitly characterizable, formulation (2) with (3) constraints yield efficient solution procedures (see Section 2.2).

In the case of a χ^2-metric, corresponding to a ϕ-divergence with $\phi(t) = (t - 1)^2$, Namkoong and Duchi [11] analyze the formulation (2) and establish its equivalence to *variance regularization* of the empirical risk minimization problem. Specifically, for convex, bounded loss functions l with P_b as the empirical distribution $P_{b,N} = (\frac{1}{N})$ over a large dataset of size N, the following result is shown to hold with high probability

$$\mathbb{E}_{P^*(\theta)}[l(\theta, \xi)] = \mathbb{E}_{P_b}[l(\theta, \xi)] + \sqrt{\frac{\rho \mathrm{Var}_{P_b}(l(\theta, \xi))}{N}}, \;\; \theta \in \Theta. \tag{4}$$

Results in a similar vein have been obtained for other ϕ-divergence metrics [7, 4], most notably Kuhlback-Leibler (KL) divergence which uses $\phi(t) = t \log t + t - 1$. Namkoong and Duchi [11] consider specific instances of the loss function where an appropriate choice of ρ leads to an optimal solution θ^*_{rob} that has loss performance within $O(1/N)$ of the (unknown) true optimal θ^*; meanwhile, the θ^*_{erm} identified by minimizing $\mathbb{E}_{P_b}[l(\theta, \xi)]$ leads to a solution with $O(1/\sqrt{N})$ loss performance.

The formulation based on directly minimizing the variance regularized risk of the form (4) over $\theta \in \Theta$ is hard to solve because of the non-convexity of the second term, even if l is strongly convex. On the other hand, the formulation in (2) is a convex problem in θ. This, in combination with the (possibly) better statistical properties of θ^*_{rob}, makes it highly desirable to efficiently solve the general min-max formulation (2). The problem formulations of primary interest in this paper are such that the optimal solution and/or optimal function value of the inner maximization problem over $\mathcal{P}$ cannot be obtained in closed form, which appears to be the case for the important general class of ϕ-divergence distance metrics in (3). Define the vectors $P := (p_n)$ and $P_{b,N} := (1/N)$ of dimension N. We shall henceforth focus on the case where P_b is the empirical probability mass function (pmf) over a dataset of size N, and thus the loss function and constraint set $\mathcal{P}$ are given by

$$L_{P_b}(\theta) = \frac{1}{N} \sum_{n=1}^{N} l(\theta, \xi_n), \;\; \mathcal{P} = \left\{ P \mid \sum_{n=1}^{N} p_n = 1, \;\; p_n \geq 0, \forall n, \;\; D_\phi(P, P_{b,N}) = \frac{1}{N} \sum_{n=1}^{N} \phi(N p_n) \leq \rho \right\}.$$

Ben-Tal et al. [1] derive the dual of the inner objective of (2), obtaining

$$R^* = \min_{\theta \in \Theta} \max_{p_n \geq 0} \min_{\alpha \geq 0, \lambda} \left\{ L_{P_b}(\theta) + \alpha(\rho - D_\phi(P, P_{b,N})) + \lambda\left(1 - \sum_{n=1}^{N} p_n\right) \right\}$$

$$= \min_{\theta \in \Theta, \alpha \geq 0, \lambda} \alpha\rho + \lambda + \frac{\alpha}{N} \sum_{n=1}^{N} \phi^*\left(\frac{l(\theta, \xi_n) - \lambda}{\alpha}\right) \tag{5}$$

where the convex conjugate of ϕ, $\phi^*(s) = \max_{t \geq 0}\{st - \phi(t)\}$, is known in closed form for various ϕ, such as those corresponding to χ^2- and KL-divergence. Since (5) is in the form of a standard stochastic minimization problem, Ben-Tal et al. [1] propose to apply classical SGD methods to compute its solution. However, as Namkoong and Duchi [10] observe, the presence of α in the denominator of the argument of ϕ^* causes SGD to become unstable as $\alpha \to 0$, which our experiments show is likely. An alternative approach is proposed in [10] that interleaves one SGD step in the θ-space with a step in the P-space. Such primal-dual steps are a result of applying stochastic mirror-descent to each set of variables. This yield a method that applies SGD-type iterations to a formulation with a composite dimension of $d + N$. Each step requires solving convex proximal mapping optimization formulations, and the computational effort needed to do this makes it is desirable to avoid this significant expansion in dimension.

To this end, Namkoong and Duchi [11] propose to determine the optimal $P^*(\theta)$ that defines $R(\theta)$ directly, namely solving the problem (2) as a large deterministic gradient descent problem. This is a feasible approach for specific choices of ϕ-divergences. For the χ^2 case, Namkoong and Duchi [11] show that the inner maximization can be reduced to a one-dimensional root-finding problem, which can be solved via bisection search. The key issue is that this bisection search still requires an $O(N \log N)$ amount of effort (see Proposition 2) at each iteration, which can be expensive.

1.2 Our Contributions

We propose a new primal descent algorithm to solve (2) that is applicable for various ϕ-divergence distance measures (3). In Section 2.1, we (slightly) generalize the (exact) bisection-search result in [10, 11] for the inner maximization problem by utilizing similar results derived in [5] for KL-divergences and showing that this general approach can be successfully applied to other ϕ-divergence metrics. This still yields a computational effort of order $O(M \log M)$ for D_{χ^2}-constrained (2) and $O(M)$ for D_{KL}-constarints, where M is the dimension of the decision variables, in our case the size of the support of the pmf. To address this issue, instead of operating with the complete dataset $M_t = N$ for all iterations t of a gradient descent algorithm, we propose the following stochastic sub-gradient descent scheme

$$\theta_{t+1} = \theta_t - \gamma_t \nabla_\theta \hat{R}_{\mathcal{M}_t}(\theta_t), \tag{6}$$

where γ_t is variously called the step-size or gain sequence or learning rate, $\mathcal{M}_t$ is a relatively small subset of the full dataset having size $|\mathcal{M}_t| = M_t$, and $\hat{R}_{\mathcal{M}_t}(\cdot)$ is an approximation of $R(\cdot)$.

The approximation $\hat{R}_{\mathcal{M}}$ is obtained by first *uniformly* sampling *without replacement* the subset $\mathcal{M}_t$ of the complete dataset of size N. This $\hat{R}_{\mathcal{M}}$ approximation solves the inner maximization over pmfs on this subset $\mathcal{M}$. Sampling without replacement differs from the standard with-replacement approach in the stochastic optimization literature, though it is preferred by practitioners in machine/deep learning. Remark 1 below describes why this strategy is needed here.

Defining $\hat{P} = (\hat{p}_m)$ of dimension M_t, we more precisely have the formulation

$$\hat{R}_{\mathcal{M}_t}(\theta) = \max_{\hat{P}=(\hat{p}_m)} \sum_{m \in \mathcal{M}_t} \hat{p}_m l(\theta, \xi_m) \text{ s.t. } \sum_{m \in \mathcal{M}_t} \phi(M\hat{p}_m) \leq M\rho_M, \sum_{m \in \mathcal{M}_t} \hat{p}_m = 1, \hat{p}_m \geq 0. \tag{7}$$

The cost of solving this problem via bisection search is $O(M_t \log M_t)$. Now suppose $\hat{P}^*(\theta) = (\hat{p}_m^*(\theta))$ is an optimal solution to (7). Then the vector

$$\nabla_\theta \hat{R}_{\mathcal{M}_t}(\theta_t) = \sum_{m \in \mathcal{M}_t} \hat{p}_m^*(\theta_t) \nabla_\theta l(\theta_t, \xi_m) \tag{8}$$

is a valid sub-gradient for $\hat{R}_{\mathcal{M}_t}(\theta_t)$ and thus we use it in (6).

3

With respect to the quality of the approximation of $\hat{R}^*$, or more particularly that of its sub-gradient, we provide a result in Section 2.3 on the rate at which the bias in the gradient estimation depends on the sample size M. Since this estimator is unbiased and the only control on it is via M, our method necessarily grows $M_t \nearrow N$ as $t \nearrow$. The result specifically depends on the $\mathcal{M}_t$ being sampled without replacement, and sampling with replacement yields a much slower bias dropoff that makes the method computationally burdensome.

We look at sample size growth rules where the maximum size N is hit after a (large but) finite number of iterations. In Section 2.4, we address the question of choosing a good sequence, and in particular balancing the added computational buden of each iteration against the expected reduction in optimality gap. We show for the strong-convex loss functions $l(\theta, \xi)$ that too slow a growth sequence is inefficient, while geometrically growing sequences are efficient in the sense that the expected optimality gap drops at a rate proporionate to the increase in computational budget.

This paper only treats strongly convex losses $l(\theta, \xi)$, but our analysis of bias and convergence and substantial aspects of the rate of convergence can be extended (in the spirit of [12, 6]) to the cases when $l(\theta, \xi)$ are convex but not c-strongly convex, or more importantly non-convex, e.g. training deep learning models. The algorithm proposed in [10] appears to be limited to convex $l(\theta, \xi)$. This subject is the focus of our ongoing research.

2 Algorithm and Analysis

2.1 SGD Algorithm

Our dynamically sampled subgradient descent algorithm for efficiently solving distributionally robust learning optimization problems is presented in Algorithm 1. Here we fix $\gamma_t = \gamma$ and increase the sub-sampling set $\mathcal{M}_t$ in a geometric manner so as to statistically cover the entire dataset. We will subsequently show that these parameter settings provide the desired statistical efficiency. The algorithm stops when $M_t \geq N$; in our experiments we proceed with the full gradient (deterministic) algorithm thereafter.

Algorithm 1 Dynamically Sampled Subgradient Descent

Given: Constant learning rate sequence γ; Initial sample size M_0 and Sample size growth factor $\nu < 1$; Initial iterate θ_0; Set $t_{\max} = \log(N/M_0)/(-\log(\nu))$

1: **for** $t = 1, 2, \ldots, t_{\max}$ **do**
2:      Sample M_t indices *without replacement* uniformly from $\{1, \ldots, N\}$, and gather them in $\mathcal{M}_t$
3:      Solve inner maximization to obtain optimal solution $\hat{P}^*_\mathcal{M}$ $\triangleright$ see Section 2.2
4:      Set $G_t \leftarrow \sum_{m \in \mathcal{M}} \hat{p}^*_m \nabla l(\theta_t, \xi_m)$
5:      Set $\theta_{t+1} \leftarrow \theta_t - \gamma G_t$ $\triangleright$ subgradient descent step
6:      Set $M_{t+1} \leftarrow M_t / \nu$
7:      Increment $t \leftarrow t + 1$
8: **end for**

Our detailed analysis in the remainder of this section starts with an exact solution to the inner maximization problem, generalizing the bisection-search result in [10, 11]. The final two subsections establish various mathematical properties for our approach w.r.t. bias and convergence, respectively.

2.2 Solving for $P^*(\theta)$ and $R(\theta)$

Recall the inner optimization problem expressed as $R(\theta) = \max_{P=(p_n)} \sum_n z_n p_n$ subject to $\sum_n (1/N) \phi(N p_n) \leq \rho/N$, $\sum_n p_n = 1$, $p_n \geq 0$. Following (7) and (8), we restrict the support of $\hat{P} = (p_m)$ in Algorithm 1 to a given set of indices $\mathcal{M}$, and only allow $p_n \geq 0$, $\forall n \in \mathcal{M}$, while setting the remaining elements as $p_n = 0$, $\forall n \notin \mathcal{M}$. We then define the restricted problem

$$\hat{R}_\mathcal{M}(\theta) = \max_{\hat{P}=(\hat{p}_m)} \sum_{m \in \mathcal{M}} z_m \hat{p}_m \text{ s.t. } \sum_{m \in \mathcal{M}} \frac{1}{M} \phi(M \hat{p}_m) \leq \rho_M, \sum_{m \in \mathcal{M}} \hat{p}_m = 1, \hat{p}_m \geq 0, \forall m \in \mathcal{M}.$$

$$(9)$$

The problem (9) states its target divergence value as ρ_M; in the subsequent sections we will prescribe specific values. Denote the optimal solution to (9) as $\hat{P}^*_{\mathcal{M}}$ and its objective value as $\hat{R}^*_{\mathcal{M}}(\theta)$; the latter is an approximation for the robust objective $R(\theta)$. Defining $\hat{P} = (\hat{p}_m)$ of dimension M and writing the Lagrangian objective of (9) as

$$\mathcal{L}(\alpha, \lambda, \hat{P}) = \sum_{m \in \mathcal{M}} z_m \hat{p}_m + \frac{\alpha}{M} \left(M \rho_M - \sum_{m \in \mathcal{M}} \Phi(M \hat{p}_m) \right) + \lambda \left(1 - \sum_{m \in \mathcal{M}} \hat{p}_m \right), \quad (10)$$

we then have $\hat{R}^*_{\mathcal{M}}(\theta) = \min_{\alpha \geq 0, \lambda} \max_{\hat{p}_m \geq 0} \mathcal{L}(\alpha, \lambda, \hat{P})$; refer to [9]. The equality constraint $\sum_{m \in \mathcal{M}} \hat{p}_m = 1$ will always be satisfied; but the ϕ-divergence inequality may not satisfied as an equality, given the optimality direction $z = (z_{m_1}, \dots, z_{m_M})$, the constraints that $1 \geq \hat{p}_m \geq 0$, and a ρ_M large enough so that the ϕ-divergence constraint allows the mass to accumulate at either of the bounds on $\hat{p}_m$. By complimentary slackness, we have the optimal $\alpha^* = 0$ in this case.

We will use the following general procedure to solve Lagrangian formulations in the proofs of Proposition 1 and Proposition 2. This has been followed by previous work[1, 5, 10, 11], either explicitly or in the same spirit:
Procedure 1.

1. *Case:* $\alpha^* = 0$ along with constraint $D_\phi(\hat{P}^*_{\mathcal{M}}, P_b) \leq \rho_M$.

   (a) Let $\mathcal{M}' = \{m \in \mathcal{M} : z_m = \max_{u \in \mathcal{M}} z_u\}$ and $M' = |\mathcal{M}'|$. Set $\alpha^* = 0$ in (10), and then observe that an optimal solution is $\hat{P}^*$ where $\hat{p}^*_m = \frac{1}{M'}$, $\forall m \in \mathcal{M}'$, and $\hat{p}^*_m = 0$, $\forall m \notin \mathcal{M}'$.

   (b) If $D_\phi(\hat{P}^*, P_b) \leq \rho_M$, then **stop** and return $\hat{P}^*$.

2. *Case:* constraint $D_\phi(\hat{P}^*_{\mathcal{M}}, P_b) = 0$ with $\alpha^* \geq 0$.

   (a) Keeping λ, α fixed, solve for the optimal $\hat{P}^*$ (as a function of λ, α) that maximizes $\mathcal{L}(\alpha, \lambda, \hat{P})$, applying the constraint $\hat{p}_m \geq 0$.

   (b) Keeping α fixed, solve for the optimal λ^* using the first order optimality condition on $\mathcal{L}(\alpha, \lambda, \hat{P}^*)$. Note that this is equivalent to satisfying the equation $\sum_{m \in \mathcal{M}} \hat{p}^*_m = 1$. This step usually leads to a λ^* available in closed form; see the results below.

   (c) Apply the first order optimality condition to the one-dimensional function $\mathcal{L}(\alpha, \lambda^*, \hat{P}^*)$ to obtain the optimal $\alpha^* \geq 0$. This is equivalent to requiring that α^* satisfy the equation $\sum_{m \in \mathcal{M}} \phi(\hat{p}^*_m) = \rho$. Substitute α^* in $\hat{P}^*$ and return it.

For the two results below, the last step of **Procedure 1** turns out to involve solving a root finding problem, where the left hand summation is a (strictly) monotonic function of α. We now apply this procedure to two specific ϕ-divergences, noting that the optimal value $\hat{R}^*_{\mathcal{M}}(\theta)$ for many other ϕ-divergences can be obtained in a similar manner. Algorithm 2 presents the solution to the χ^2-divergence constrained problem.

Proposition 1 *The optimal solution $\hat{P}^*$ to the problem (9) with a KL-divergence constraint (where $\phi(t) = t \log t - (t - 1)$) is given by*
1. Case $\alpha^* = 0$: $\hat{p}^*_m = \frac{1}{M'}$, *where* $M' = |\mathcal{M}'|$, $\mathcal{M}' = \{m \in \mathcal{M} | z_m = \max_u z_u\}$;

2. *Case* $D_{KL}(\hat{P}^*, P_b) = \rho_M$: $\hat{p}^*_m = \frac{e^{z_m \beta^*}}{\sum_j e^{z_j \beta^*}}$, *where* β^* *solves* $\beta \kappa'(\beta) - \kappa(\beta) = \rho_M$ *and* $\kappa(\beta) = \sum_j e^{z_j \beta} / M$. *The computational effort needed to solve this problem is* $O(M \log(1/\epsilon))$, *where* ϵ *is the desired accuracy.*

Proof of Proposition 1: We first handle the case when the KL-divergence constraint is not tight and $\alpha^* = 0$. Substituting this in (10) shows that any optimal solution $\hat{P}^*$ places mass only within the set $\mathcal{M}'$ as defined. Consider any such $\hat{P}^*$, and let $\hat{P}^*_b$ be the solution that assigns equal mass to the M'

support points in $\mathcal{M}'$. We then have

$$D_\phi(\hat{P}^*, P_b) = \sum_{j \in \mathcal{M}'} \frac{1}{M} \phi(M\hat{p}_j^*) + \sum_{j \notin \mathcal{M}'} \frac{1}{M} \phi(0) = \frac{M'}{M} \sum_{j \in \mathcal{M}'} \frac{1/M}{M'/M} \phi(M\hat{p}_j^*) + \sum_{j \notin \mathcal{M}'} \frac{1}{M} \phi(0)$$

$$\geq \frac{M'}{M} \phi \left(\sum_{j \in \mathcal{M}'} \frac{1/M}{M'/M} M\hat{p}_j^* \right) + \sum_{j \notin \mathcal{M}'} \frac{1}{M} \phi(0) = D_\phi(\hat{P}_b^*, P_b)$$

where we apply Jensen's inequality to the convex ϕ. Thus, among all optimal solutions, $\hat{P}_b^*$ obtains the smallest divergence, and hence is the best optimal candidate to meet the divergence constraint with slack. Note that this applies for any convex ϕ.

For the case when the KL-divergence constraint is tight using $\phi(t) = t \log t - t + 1$, we proceed according to the corresponding three steps in **Procedure 1** above.

Step 2(a). Setting to zero the gradient of $\mathcal{L}(\alpha, \lambda, \hat{P})$ with respect to $\hat{P}$, we obtain

$$\hat{p}_m^* = \frac{1}{M} e^{\frac{z_m - \lambda}{\alpha}}, \quad m \in \mathcal{M}.$$

This solution also satisfies the non-negativity constraint on $\hat{P}^*$.

Step 2(b). Setting $\sum \hat{p}_m^* = 1$ renders $e^{-\frac{\lambda}{\alpha}} = \frac{M}{\sum_j e^{z_j/\alpha}}$, which in turn yields $\hat{p}_m^* = \frac{e^{z_m/\alpha}}{\sum_j e^{z_j/\alpha}}$.

Step 2(c). To obtain α^*, substitute the $\hat{P}^*$ into the divergence constraint satisfied as an equality. Then, after some algebra, we conclude that α^* must satsify

$$\frac{\sum_m z_m e^{z_m/\alpha^*}}{\alpha^* \sum_j e^{z_j/\alpha^*}} - \log \sum_j \frac{e^{z_j/\alpha^*}}{M} = \rho_M.$$

Let $\beta = 1/\alpha$ and write $\kappa(\beta) = \log \sum_j e^{z_j \beta}/M$. Then, finding α^* is equivalent to obtaining the β^* that satisfies $\beta \kappa'(\beta) - \kappa(\beta) = \rho_M$. A unique root for this exists because the left hand expression is monotonic and takes on a value of 0 at $\beta = 0$, and $\kappa'(\beta) \to \infty$ as $\beta \to \infty$. Hence, a bisection search will render the optimal $\alpha^* = 1/\beta^*$. $\square$

Proposition 2 *An optimal solution to the problem (9) with a χ^2-divergence constraint (where $\phi(t) = (t-1)^2$) is given by*
1. *Case $\alpha^* = 0$: $\hat{p}_m^* = \frac{1}{M'}$, where $M' = |\mathcal{M}'|$, $\mathcal{M}' = \{m \in \mathcal{M} | z_m = \max_u z_u\}$;*

2. *Case $D_{KL}(\hat{P}^*, P_b) = \rho_M$: $\hat{p}_m^* = \begin{cases} \frac{z_m - z_{\min} - \lambda^*}{2\alpha^* M} + \frac{1}{M}, & z_m \geq z_{\min} + \lambda^* - 2\alpha, \\ 0, & z_m < z_{\min} + \lambda^* - 2\alpha, \end{cases}$ where* $z_{\min} = \min_m z_m$ *and* λ^*, α^* *jointly solve:* $\sum_m (z_m - z_{\min}) \mathbb{I}\{\hat{p}_m^* > 0\} - 2M\alpha^* = (\lambda^* - 2\alpha^*) \left(\sum_m \mathbb{I}\{\hat{p}_m^* > 0\} \right)$ *and* $\sum_m (z_m - z_{\min} - \lambda^*)^2 \mathbb{I}\{\hat{p}_m^* > 0\} = 4\rho_M M$. *Furthermore, the computational effort needed to obtain the primal-dual optimal solutions $\hat{P}^*, \alpha^*, \lambda^*$ is $O((M - \log \epsilon) \log M)$, where ϵ is the estimation precision required.*

Proof of Proposition 2: The case when $\alpha^* = 0$ is handled as in Proposition 1, and thus we only consider the case where $D_{\chi^2}(\hat{P}^*, P_b) = 0$.

First, order all the z_m into the increasing sequence $z_{(1)} \leq z_{(2)} \leq \ldots \leq z_{(M)}$, where the notation (i) denotes the index of the ith smallest z_m value. Additionally, define $v_i := z_{(i)} - z_{(1)}, \forall i = 1, \ldots, M$. Note that the objective function $\sum_m z_m \hat{p}_m = \sum_i (v_i + z_{(1)}) \hat{p}_i = \sum_i v_i \hat{p}_i + z_{(1)}$, and hence it is sufficient to maximize $\hat{p}_i$ with respect to the (non-negative) vector $v = (v_1, \ldots, v_M)$.

Step 2(a). Setting the gradient of $\mathcal{L}(\alpha, \gamma, \hat{P})$ with respect to $\hat{P}$ to zero componentwise for each $\hat{p}_i$, we obtain

$$\hat{p}_i^* = \begin{cases} \frac{v_i - \lambda}{2\alpha M} + \frac{1}{M}, & v_i \geq \lambda - 2\alpha, \\ 0, & v_i < \lambda - 2\alpha. \end{cases}$$

Let I represent the index for which the following condition holds:

$$v_{I+1} \geq \lambda - 2\alpha \quad \text{and} \quad v_I < \lambda - 2\alpha. \tag{11}$$

Step 2(b). Let $s_i = \sum_{j=i+1}^{M} v_j$, $\forall i = 0, \ldots, M-1$. The equality $\sum_i \hat{p}_i = 1$ can be rewritten as

$$\sum_{j \geq I+1} \left(\frac{v_j - \lambda}{2M\alpha} + \frac{1}{M} \right) = 1 \qquad \text{or} \qquad (s_I - 2M\alpha) = (\lambda - 2\alpha)(M - I), \qquad (12)$$

where λ and the index I satisfy the bounds in (11). The first term is a lower semi-continuous decreasing step function of λ, with steps at the λ where $\lambda - 2\alpha = v_i$ for each i; recall that $v_j \geq 0$. The right hand side is an increasing function of λ. Hence, a unique λ^* exists that satisfies (12); we only need to check the mismatch at the M breakpoints of the step-function to find this λ^*. A bisection search with computational effort of at most $O(\log M)$ (as described in Algorithm 2) yields this point.

Step 2(c). This last step requires the zero of the gradient of $\mathcal{L}(\alpha, \lambda, \hat{P})$ with respect to α, or equivalently the α^* that satisfies

$$\sum_{j >= I(\alpha)+1} (v_j - \lambda(\alpha))^2 - M\rho_M 4\alpha^2 = 0. \qquad (13)$$

The first term is a decreasing function of α since $I \to 0$ and $\lambda \to -\infty$ as $\alpha \to \infty$. Hence a unique root exists, which can again be found via a binary search (see Algorithm 2). From (12), we know that when α is large, the optimal $\lambda^*(\alpha) = s_0/M$ and $I^*(\alpha) = 0$. Let $\bar{\alpha} = (\sum_i (v_i - \sum_j v_j/M))/4\rho_M M$ and let the bisection algorithm search within $(0, \alpha_\infty]$ where $\alpha_\infty = K\bar{\alpha}$ for some large constant K.

The bisection for α^* involves $\log 1/\epsilon$ steps where ϵ is the precision required in solving (13), each of which takes $\log M$ steps to solve for the optimal $(\lambda^*(\alpha), I^*(\alpha))$ pair. The overall computational complexity of solving for $\hat{P}^*(\theta)$ is therefore $O((\log \frac{1}{\epsilon} + M) \cdot \log M)$, where the second $M \log M$ term arises from sorting the M values z_m into the vector v. $\square$

To summarize, the optimization procedure to obtain the solution to the χ^2-divergence constrained problem is presented in Algorithm 2.

2.3   Small-sample Approximation of $\nabla_\theta R(\theta)$

Algorithm 1 is proposed in the spirit of SGD methods, in that it is unnecessary to obtain precise values for the gradient especially for the initial iterations θ_t in (6). We therefore construct a sub-gradient approximation $\nabla_\theta \hat{R}_{\mathcal{M}}(\theta_t)$ in (8) to the full-gradient $\nabla_\theta R(\theta_t) = \sum_{n=1}^N p_n^* \nabla_\theta l(\theta_t, \xi_n)$, where $P^* = (p_1^*, \ldots, p_N^*)$ is the optimal solution to the full-data problem (2) and $\hat{P}_t^* = (\hat{p}_1^*, \ldots, \hat{p}_{M_t}^*)$ in (8) is the optimal solution to the restricted problem (7) based on uniformly sampling *without replacement* M_t data points from the full data set.

The primary concern with this approach is the bias induced by the subsampling of the full support, which we show in Theorem 3 to be of order $O(1/M_t - 1/N)^{1-\delta}$. We restrict our attention to ϕ-divergences that satisfy, for a small $\eta > 0$, the continuity condition

$$|\phi(t(1 + \eta)) - \phi(t)| \leq \kappa_1(\eta)\phi(t) + \kappa_2(\eta), \qquad (14)$$

where $\kappa_1(\eta)$ and $\kappa_2(\eta)$ are both $O(\eta)$. This continuity condition can be verified for many common ϕ-divergence measures of interest including the χ^2 and KL-divergence metrics. Let $\mathbb{E}_{\mathcal{M}}$ and $\mathbb{P}_{\mathcal{M}}$ be expectations and probabilities w.r.t. the uniform sampling without-replacement producing the random set $\mathcal{M}$.

Theorem 3 *Suppose the optimal solution P^* to (2) is unique and $\rho \ll 1$ in (1). Assume the ϕ-divergence satisfies (14) and define the D_ϕ-constraint target in (7) to be $\rho_M = \rho + \eta_M$, where $\eta_M = c \left(\frac{1}{M} - \frac{1}{N} \right)^{(1-\delta)/2}$ for constant $c > 0$ and small constant $\delta > 0$. Then, for all $M \geq M_0$ with M_0 sufficiently large, we have that the sub-gradient $\nabla_\theta \hat{R}_{\mathcal{M}}(\theta)$ and full-gradient $\nabla_\theta R(\theta)$ satisfy $\|\mathbb{E}_{\mathcal{M}}[\nabla_\theta \hat{R}_{\mathcal{M}}(\theta)] - \nabla_\theta R(\theta)\|_2^2 \to C\eta_M^2$ as $M \to N$, where C is a finite constant.*

We first provide a sketch of the proof of Theorem 3, with the full details to follow. First construct $\tilde{P}_{\mathcal{M}} = (\tilde{p}_1, \ldots, \tilde{p}_M)$, a restriction of the (unique) optimal solution P^* of the full-data problem (2) onto the (random) subset $\mathcal{M}$ of support points used in the restricted problem (9), where $\tilde{p}_m =$

Algorithm 2 Solve for $\hat{P}^*_{\mathcal{M}}$ under χ^2-divergence constraint

---

Given: Subset $\mathcal{M} \subset \mathcal{D}$ of full dataset $\mathcal{D}$, and optimization objective $z_m = l(\theta, \xi_m)$, $m \in \mathcal{M}$.

1: Form the non-negative (increasing) ordered $v_i = z_{(i)} - z_{(1)}$, where $z_{(i)}$ denotes the ith smallest value in $\{z_m\}$.
2: Form the sums $s_i = \sum_{j=i+1}^{M} v_j$, $i = 0, \ldots, M$.
3: Form the set $\mathcal{M}' \leftarrow \{i \in \mathcal{M} | v_i = v_1\}$. Set $M' \leftarrow |\mathcal{M}'|$.

Case $\alpha^* = 0$

4: Set $\hat{p}^*_m \leftarrow \frac{1}{M'}$, $\forall m \in \mathcal{M}'$, and $\hat{p}^*_m = 0$ otherwise. Let $\hat{P}^*_{\mathcal{M}} = (\hat{p}^*_1, \ldots, \hat{p}^*_M)$.
5: **if** $D_{\chi^2}(\hat{P}^*_{\mathcal{M}}, P_b) \leq M\rho_m$ **then**
6:      Return $\hat{P}^*_{\mathcal{M}}$ as an optimal solution.
7: **end if**

Case $D_{\chi^2}(\hat{P}^*_{\mathcal{M}}, P_b) = 0$

8: Set $\alpha_{\min} = 0$ and $\alpha_{\max} = K\sqrt{\frac{1}{4\rho_M M)} \sum_i \left(v_i - \frac{\sum_j v_j}{M} \right)}$ ▷ K is a large constant
9: **while** $|\alpha_{\max} - \alpha_{\min}| > \epsilon$ **do**
10:      Set $\alpha = \frac{\alpha_{\max} + \alpha_{\min}}{2}$
11:      Get $(\lambda, I) = \text{FIND\_OPTIMAL\_LAMBDA}(\alpha)$
12:      Set $\Delta = \sum_{j=I+1}^{M} (v_j - \lambda)^2 - 4M\rho_M \alpha^2$
13:      **if** $\eta > 0$ **then**
14:          $\alpha_{\min} = \alpha$
15:      **else**
16:          $\alpha_{\max} = \alpha$
17:      **end if**
18: **end while**
19: Set $\alpha^* = \frac{\alpha_{\max} + \alpha_{\min}}{2}$ and $(\lambda^*, I^*) = \text{FIND\_OPTIMAL\_LAMBDA}(\alpha)$
20: For all $j \geq I^* + 1$, set $\hat{p}^*_j = \frac{1}{M} + \frac{v_j - \lambda^*}{2\alpha^* M}$, else $\hat{p}^*_j = 0$. Return $\hat{P}^*_{\mathcal{M}}$.

Finding $(\lambda^*(\alpha), I^*(\alpha))$ **for a given** α **by bisection**

21: **function** FIND_OPTIMAL_LAMBDA(α)
22:      **if** $s_0 < 2M\alpha$ **then**
23:          Return $(\lambda = \frac{s_0}{2M}, I = 0)$.
24:      **end if**
25:      Set $I_{\min} = 1$ and $I_{\max} = M$.
26:      **while** $(I_{\max} - I_{\min}) > 0$ **do**
27:          Set $I \leftarrow \frac{I_{\min} + I_{\max}}{2}$
28:          **if** $s_I - 2M\alpha > v_{I+1}(M - I)$ **then**
29:              Set $I_{\min} = I + 1$
30:          **else if** $s_I - 2M\alpha < v_I(M - I)$ **then**
31:              Set $I_{\max} = I - 1$
32:          **else**
33:              Set $\lambda = \frac{s_I - 2M\alpha}{M - I}$
34:              Set Return (λ, I)
35:          **end if**
36:      **end while**
37: **end function**

$\frac{p_m^*}{\sum_{j \in \mathcal{M}} p_j^*}$, $\forall m \in \mathcal{M}$. The condition $\lfloor \rho \rfloor \ll N$ ensures that, with high probability, the summation in the denominator is greater than zero for a sufficiently large M. We then show that, with high probability (under the $\mathcal{M}$-sampling measure), the pmf $\tilde{P}_{\mathcal{M}}$ is a feasible solution to (9) when ρ_M is inflated as assumed. Next, we establish that $\mathbb{E}_{\mathcal{M}}[|z^T(\tilde{P}_{\mathcal{M}} - P^*)|]$ is of the order $O(\eta_M^{2/(1-\delta)})$, where z^T denotes the transpose of vector z. Since $\tilde{P}_{\mathcal{M}}$ is a feasible solution to (9), an appeal to the fundamental theorem of calculus yields the desired result. We extensively exploit the statistical properties of sampling a finite set without replacement, and therefore provide a brief summary here. Let $\{x_1, \ldots, x_N\}$ be a set of one-dimensional values with $\mu = \frac{1}{N} \sum_n x_n$ and $\sigma^2 = \frac{1}{N-1} \sum_n (x_n - \mu)^2$. Suppose we sample $M < N$ of these points uniformly without replacement to construct the set $\mathcal{M} = \{X_1, \ldots, X_M\}$. The probability that any particular set of M subsamples was chosen is $\left(\frac{(N-M)!}{N!}\right)$. Denote by $\mathbb{E}_{\mathcal{M}}$ the expectation under this probability measure, and let $\bar{X} = \frac{1}{M} \sum_{m=1}^{M} X_m$ and $\bar{S}^2 = \frac{1}{M-1} \sum_{m=1}^{M} (X_m - \bar{X})^2$ represent the sample mean and sample variance, respectively. We then know [15] that

$$\mathbb{E}_{\mathcal{M}}[\bar{X}] = \mu, \qquad \mathbb{E}_{\mathcal{M}}[\bar{S}^2] = \sigma^2, \qquad \mathbb{E}_{\mathcal{M}}[(\bar{X} - \mu)^2] = \left(\frac{1}{M} - \frac{1}{N}\right)\sigma^2.$$

The second term, i.e., the expectation of the sample variance, shows that the sample variance is an unbiased estimate of the true variance σ^2. Further note that the third term, i.e., the variance of the sample mean, reduces to zero as $M \to N$.

We now start by addressing the feasibility of the restriction $\tilde{P}_{\mathcal{M}}$ of the (unique) optimal solution P^* of the full-data problem onto the (randomly sampled) subset $\mathcal{M}$.

Lemma 4 *Suppose the ϕ-divergence function satisfies condition* (14). *Let the D_ϕ-constraint target ρ_M of the restricted problem* (9) *be set as*

$$\rho_M = \rho + c\eta_M, \qquad where \quad \eta_M = \left(\frac{1}{M} - \frac{1}{N}\right)^{(1-\delta)/2}, \quad c > 0, \text{ and } \delta > 0 \quad small.$$

Denote by $\mathcal{P}_{\mathcal{M}}$ the feasibility set of (9). *Then, we have*
$$\mathbb{P}_{\mathcal{M}}(\tilde{P}_{\mathcal{M}} \in \mathcal{P}_{\mathcal{M}}) \to 1 \qquad as \qquad M \to N.$$

Proof of Lemma 4: In the notation of sampling without-replacement introduced above, define a set of scalar values $x_n = Np_n^*$, $\forall n = 1, \ldots, N$. We then have

$$\mu = \frac{1}{N} \sum_n Np_n^* = 1 \quad and \quad \sigma^2 = \frac{1}{N-1} \sum_n (Np_n^* - 1)^2.$$

By Chebychev's inequality, the sample-average $\bar{X}$ of an M-subsample from this set satisfies

$$\mathbb{P}_{\mathcal{M}}\left(|\bar{X} - 1| > \eta_M\right) \leq \frac{1}{\eta_M^2}\mathbb{E}_{\mathcal{M}}[(\bar{X} - \mu)^2] \leq (\frac{1}{M} - \frac{1}{N})^\delta \sigma^2.$$

Hence, as $N \to M$, we have with probability at least $1 - \eta_M^{2\delta/(1-\delta)}$ that $|\bar{X} - 1| \leq \eta_M$.

The condition $\rho \ll 1$ ensures with high probability that the full data inner maximization (2) is tightly constrained by the D_ϕ constraint and a degenerate solution with $\alpha^* = 0$ (as in Case 1 of **Procedure 1**) does not apply. This lets us choose an M_0 such that $\mathbb{P}_{\mathcal{M}}(\sum_{j \in \mathcal{M}} NP_j^* > 0) \geq 1 - \delta'$ for all $|\mathcal{M}| \geq M_0$. Rearranging $|\bar{X} - 1| \leq \eta_M$, we obtain $(\frac{1}{M} \sum_j Np_j^*)^{-1} \leq 1 + (\eta_M/(1 - \eta_M))$. Let $\eta_M' = \eta_M/(1 - \eta_M) = O(\eta_M)$ as $M \to N$. Then the solution $\tilde{P}_{\mathcal{M}}$ is a pmf, and thus we only need to check whether $\mathbb{P}_{\mathcal{M}}(D_\phi(\tilde{P}_{\mathcal{M}}, P_{b,M}) > \rho_M)$ is small. For a sufficiently large $M \geq M_0$ such that η_M is small enough to satisfy the ϕ-continuity condition (14), we obtain

$$D_\phi(\tilde{P}_{\mathcal{M}}, P_{b,M}) = \frac{1}{M} \sum_m \phi\left(M\frac{p_m^*}{\sum_j p_j^*}\right)$$

$$= \frac{1}{M} \sum_m \phi\left(Np_m^* \frac{1}{\frac{1}{M}\sum_{j \in \mathcal{M}} Np_j^*}\right)$$

$$\leq \frac{1}{M}\left(\sum_m \phi(Np_m^*)(1 + O(\eta_M))\right) + O(\eta_M), \qquad (15)$$

9

where the last inequality follows from (14). Let $\{x_n = \phi(Np_n^*)\}_{n=1}^N$ be a vector from which we choose the (random) indices $\mathcal{M}$. Then $\mathbb{E}_\mathcal{M}\bar{X} = \mathbb{E}_\mathcal{M}[\frac{1}{M}\sum_{m\in\mathcal{M}}\phi(Np_m^*)] = \frac{1}{N}\sum_n \phi(Np_n^*) = D_\phi(P^*, P_b)$. From Section 2.2, when $\rho \ll 1$, the D_ϕ-constraint is tight at the optimal solution and we have that $D_\phi(P^*, P_b) = \rho$.

Therefore, taking expectations on both sides of the inequality (15) renders

$$\mathbb{E}_\mathcal{M}D_\phi(\tilde{P}_\mathcal{M}, P_{b,M}) \quad \le \quad (1 + O(\eta_M))D_\phi(P^*, P_b) + O(\eta_M) \quad = \quad \rho + O(\eta_M)(1 + \rho_M).$$

Define $\rho_M = \rho + c\eta_M$ by choosing a constant $c > 0$ such that last summand on the right hand side is strictly smaller than ρ_M. This then yields the desired high probability guarantee. $\square$

Lemma 4 shows that the specfic choice of ρ_M allows the restriction of the unique optimal $P_\mathcal{M}^*$ to be feasible for (9) with high probability as $M \nearrow N$. We next establish that the bias in the estimation of the optimal objective is $O(\eta_M)$.

Lemma 5 *Under the assumptions of Lemma 4, we have*

$$|\mathbb{E}_\mathcal{M}[\hat{R}_\mathcal{M}(\theta)] - R(\theta)| = O\left(\eta_M\right).$$

Proof of Lemma 5: We first estimate the gap between $z^T \tilde{P}_\mathcal{M}$ and $z^T P^*$, the optimal objective value of the full-data problem (2), as follows

$$\mathbb{E}_\mathcal{M}[z^T \tilde{P}_\mathcal{M}] - z^T P^* = \mathbb{E}_\mathcal{M}\left[\left(\sum_{m\in\mathcal{M}} z_m \frac{p_m^*}{\sum_j p_j^*} - \sum_{n=1}^N z_n p_n^*\right)\right],$$

$$= \mathbb{E}_\mathcal{M}\left[\left(\frac{\frac{1}{M}\sum_{m\in\mathcal{M}} z_m Np_m^*}{\frac{1}{M}\sum_j Np_j^*} - \frac{1}{N}\sum_{n=1}^N z_n Np_n^*\right)\right] = \mathbb{E}_\mathcal{M}\left[\left(\frac{\bar{X}(R)}{\bar{X}(P^*)} - \frac{\mu(R)}{\mu(P^*)}\right)\right], \quad (16)$$

where the last equality uses the sample and population means of the two N-dimensional vectors:

$$\{x_1(R),\ldots,x_n(R) = z_n Np_n^*,\ldots,x_N(R)\}, \qquad \bar{X}(R) = \frac{1}{M}\sum_{m\in\mathcal{M}} x_m(R), \quad \mu(R) = \frac{1}{N}\sum_{n=1}^N x_n(R);$$

$$\{x_1(P^*),\ldots,x_n(P^*) = Np_n^*,\ldots,x_N(P^*)\}, \qquad \bar{X}(P^*) = \frac{1}{M}\sum_{m\in\mathcal{M}} x_m(P^*), \quad \mu(P^*) = \frac{1}{N}\sum_{n=1}^N x_n(P^*).$$

The Taylor expansion of any smooth function $h(u,v)$ is given by

$$h(u,v) = h(u_o,v_o) + \nabla_\theta h(u_o,v_o)\begin{pmatrix}(u-u_o)\\(v-v_o)\end{pmatrix} + \begin{pmatrix}(u-u_o)\\(v-v_o)\end{pmatrix}^T \nabla_\theta^2 h(u_o,v_o)\begin{pmatrix}(u-u_o)\\(v-v_o)\end{pmatrix} + r(u,v,u_o,v_o),$$

where the higher order terms $r(u,v,u_o,v_o)$ are $o(\|u-u_o\|\cdot\|v-v_o\|)$. Applying this to $h(u,v) = u/v$ with $u = \bar{X}(R)$, $u_o = \mu(R)$, $v = \bar{X}(P^*)$, $v_o = \mu(P^*)$ and $Y = r(u,v,u_o,v_o)$, we obtain

$$h(u,v) - h(u_o,v_o) = \frac{\bar{X}(R)}{\bar{X}(P^*)} - \frac{\mu(R)}{\mu(P^*)} = \frac{1}{\mu(P^*)}(\bar{X}(R) - \mu(R)) - \frac{\mu(R)}{\mu(P^*)^2}(\bar{X}(P^*) - \mu(P^*))$$

$$+ \frac{2\mu(R)}{\mu(P^*)^3}(\bar{X}(P^*) - \mu(P^*))^2 - \frac{1}{\mu(P^*)^2}(\bar{X}(R) - \mu(R))(\bar{X}(P^*) - \mu(P^*)) + Y.$$

From [15], we have that the higher order terms Y involve $\mathbb{E}_\mathcal{M}[(\bar{X} - \mu)^3] = o(\eta^{2/(1-\delta)})$. Hence, the expectation of the last term in (16) is rendered as

$$\left|\mathbb{E}_\mathcal{M}\left[\frac{\bar{X}(R)}{\bar{X}(P^*)} - \frac{\mu(R)}{\mu(P^*)}\right]\right| \le 2\mu(R)\mathbb{E}_\mathcal{M}[(\bar{X}(P^*)-1)^2] + \sqrt{\mathbb{E}_\mathcal{M}[(\bar{X}(R) - \mu(R))^2]\mathbb{E}_\mathcal{M}[(\bar{X}(P^*) - 1)^2]} + Y,$$

where we apply the Cauchy-Schwartz inequality to the second term and Y comprises terms of order $o(\eta_M^{2/(1-\delta)})$. Each of the two summands on the right hand side are of order $\eta_M^{2/(1-\delta)}$.

Now, since $\tilde{P}_\mathcal{M}$ is a feasible solution to the restricted problem while $\hat{P}_\mathcal{M}^*$ is its optimal solution, $\mathbb{E}_\mathcal{M}z^T\tilde{P}_\mathcal{M} \le \mathbb{E}_\mathcal{M}z^T\hat{P}_\mathcal{M}^*$. On the other hand, the $\hat{P}_\mathcal{M}^*$ satisfies the D_ϕ-divergence constraint at

$\rho + \eta_M$. For large enough M, we note from the fact that the objective function (in the Lagrangian form) depends linearly on the parameter ρ that $z^T \hat{P}_{\mathcal{M}}^* \leq z^T P^* + \eta_M$. Thus,

$$\left| \mathbb{E}_{\mathcal{M}}[z^T \hat{P}_{\mathcal{M}}^*] - z^T P^* \right| \leq O\left(\max\{\eta_M, \ \eta_M^{2/(1-\delta)}\}\right),$$

which yields the final result. $\square$

Proof of Theorem 3: Given the robust loss function $R(\theta) = \sum_n l(\theta, \xi_n) p_n^*$ and our approximation $\hat{R}_{\mathcal{M}}(\theta) = \sum_{m \in \mathcal{M}} l(\theta, \xi_m) \hat{p}_m$ constructed from the subsampled $\mathcal{M}$, the mean-value theorem of calculus renders

$$(\nabla_\theta l(\theta, \xi_n))_u = \frac{\partial l(\theta, \xi_n)}{\partial \theta_u} = \frac{1}{h_{u,n}}(l(\theta + h_{u,n}\mathbf{e}_u) - l(\theta)),$$

where $h_{u,n}$ is a small positive value that depends on the component θ_u and on the sample ξ_n, with $\mathbf{e}_u$ the unit-vector in the uth coordinate. Let $\underline{h} = \min_{u,n} h_{u,n}$. From Lemma 5, we therefore have

$$\left| (\mathbb{E}_{\mathcal{M}} \nabla_\theta \hat{R}_{\mathcal{M}}(\theta) - \nabla_\theta R(\theta))_u \right| \leq \frac{1}{\underline{h}} \left| \mathbb{E}_{\mathcal{M}} \sum_n l(\theta + h_{u,n}\mathbf{e}_u, \xi_n)^T (p_n^* - \hat{p}_n) - l(\theta, \xi_n)^T (p_n^* - \hat{p}_n) \right|$$

$$\leq \frac{1}{\underline{h}} \left| \mathbb{E}_{\mathcal{M}} \left[\sum_n l(\theta + h\mathbf{e}_u, \xi_n)^T (p_n^* - \hat{p}_n) \right] \right| + \left| \mathbb{E}_{\mathcal{M}} \left[l(\theta, \xi_n)^T (p_n^* - \hat{p}_n) \right] \right| = O(\eta_M).$$

Squaring and combining these terms over all u yields the final desired result. $\square$

Remark 1 *The squared bias in Theorem 3 is more accurately stated as $O(\eta_{|\mathcal{M}_t|}^2)$, where $|\mathcal{M}_t|$ is the number of support points used in (9). We require sampling without replacement because M_t samples with replacement only produces a set $\mathcal{M}$ such that $|\mathcal{M}| = O(\log M_t)$. The resulting slow drop in bias makes the method inefficient in terms of the computational effort expended.*

2.4   Convergence of (6)

We now present an analysis of the convergence of Algorithm 1 under the following assumptions.

Assumption 1 *(i) For each ξ_n, $n = 1, \ldots, N$, the loss functions $l(\theta, \xi_n)$ are c-strongly convex and their gradients $\nabla_\theta l(\theta, \xi_n)$ are L-Lipschitz. Additionally, the Hessian $\nabla_\theta^2 l(\theta, \xi_n)$ exists.*

(ii) The robust loss function $R(\theta)$ has a unique minimizer θ_{rob} that satisfies (2).

(iii) The estimator $\nabla R(\theta)$ obeys a bound $\mathbb{E}_{\mathcal{M}} \left[\|\nabla \hat{R})\mathcal{M}(\theta) - \mathbb{E}_{\mathcal{M}}[\nabla \hat{R}]\mathcal{M}(\theta)]\|_2^2 \right] \leq C' \eta_M^{2/(1-\delta)}$.

Proposition 6 below shows that the properties in Assumption 1(i) translate over to the robust performance metric $R(\theta)$ as defined in (2). We can relax the assumption to have sample-dependent constants $c(\xi)$ and $L(\xi)$. Since the number of samples is finite, the values $\bar{L} = \max_\xi L(\xi)$ and $\underline{c} = \min_\xi c(\xi)$ are sample-independent values that can be used in place of c, L in Proposition 6 to obtain the same properties.

Proposition 6 *With Assumption 1, the function $R(\theta) = \max_{P \in \mathcal{P}} L_P(\theta)$ is c-strongly convex, and its gradient $\nabla_\theta R(\theta)$ is L-Lipschitz.*

Proof of Proposition 6: Since each $l(\theta, \xi_n)$ is c-strongly convex, we have

$$l(\theta_1, \xi_n) + \nabla_\theta l(\theta_1, \xi_n)^T (\theta_2 - \theta_1) + \frac{c}{2} \|\theta_2 - \theta_2\|_2^2 \leq l(\theta_2, \xi_n).$$

Take any pmf P with components p_n and sum up each side to obtain

$$\sum_n p_n \left(l(\theta_1, \xi_n) + \nabla_\theta l(\theta_1, \xi_n)^T (\theta_2 - \theta_1) + \frac{c}{2} \|\theta_2 - \theta_2\|_2^2 \right) \leq \sum_n p_n l(\theta_2, \xi_n).$$

Since the above applies to any P, apply this for $P^*(\theta_1)$, the optimal pmf for the inner maximization that defines $R(\theta_1)$, with components $p_n^*(\theta_1)$. As discussed in Section 2.2, if the D_ϕ-constraint is tight

11

enough, i.e., $\rho \ll 1$, then $P^*(\theta_1)$ is, with high probability, unique for θ_1, and thus the subgradient $\nabla R(\theta)$ is the gradient. We then derive

$$\left(\sum_n p_n^*(\theta_1)l(\theta_1,\xi_n)\right) + \left(\sum_n p_n^*(\theta_1)\nabla_\theta l(\theta_1,\xi_n)\right)^T (\theta_2 - \theta_1) + \frac{c}{2}\|\theta_2 - \theta_2\|_2^2 \le \sum_n p_n^*(\theta_1)l(\theta_2,\xi_n)$$

$$R(\theta_1) + \nabla_\theta R(\theta_1)^T(\theta_2 - \theta_1) + \frac{c}{2}\|\theta_2 - \theta_1\|_2^2 \le \max_{P \in \mathcal{P}} \sum_n p_n l(\theta_2,\xi_n) = R(\theta_2),$$

which verifies that $R(\theta)$ is c-strongly convex.

For the L-Lipschitz gradient condition, we check a relatively lesser used condition, namely that the L-Lipschitz-gradient condition is equivalent to establishing the convexity of the function

$$g(\theta,\xi_n) := \frac{L}{2}\theta^T\theta - l(\theta,\xi_n). \tag{17}$$

Use the convexity definition of $g(\theta,\xi_n)$ that $g(\theta_2,\xi_n) \ge g(\theta_1,\xi_n) + \nabla_\theta g(\theta_1,\xi_n)^T(\theta_2 - \theta_1)$ to verify its equivalence to the more commonly used L-Lipschitz-gradient condition that

$$l(\theta_1,\xi_n) + \nabla_\theta l(\theta_1,\xi_n)^T(\theta_2 - \theta_1) + \frac{L}{2}\|\theta_2 - \theta_2\|_2^2 \ge l(\theta_2,\xi_n).$$

By assumption, the Hessians $\nabla_\theta^2 l(\theta,\xi_n)$ exist, and so (17) can be equivalently represented as the positive semi-definiteness of $LI - \nabla_\theta^2 l(\theta,\xi_n)$.

Now, for any two postive semi-definite matrices A and B, we have that $\det(A + B) \ge \det(A) + \det(B) \ge 0$. To see this, factorize $A = C^T C$ for some matrix C whose (generalized) inverse C^{-1} exists. Then, the matrix $(C^{-1})^T B C^{-1}$ is positive semi-definite, and we have that

$$\det(A + B) = \det(C^T(I + (C^{-1})^T B C^{-1})C) = \det(C)^2 \det(I + (C^{-1})^T B C^{-1})$$
$$\ge \det(C)^2(1 + \det(C)^{-2}\det(B)) = \det(A) + \det(B).$$

The first inequality above uses the fact that for a semidefinite D with eigenvalues e_i, $\det(I + D) = \Pi(1 + e_i) \ge 1 + \Pi e_i = 1 + \det(D)$.

We use this sum-of-determinants identity to check the condition (17) for $R(\theta) = \sum_n p_n^*(\theta)l(\theta,\xi_n)$ to obtain the desired result (recalling that $\theta \in \mathbb{R}^d$):

$$\det(LI - \sum_n p_n^*(\theta)\nabla_\theta^2 l(\theta,\xi_n)) = \det(\sum_n p_n^*(\theta)(LI - \nabla_\theta^2 l(\theta,\xi_n)))$$
$$\ge \sum_n (p_n^*(\theta))^d \det((LI - \nabla_\theta^2 l(\theta,\xi_n)) \ge 0.$$

$\square$

The Assumption 1(ii) is likely to hold given the strong-convexity of $R(\theta)$ from Proposition 6. The variance experienced by the average of a set sampled (without replacement) from a larger dataset is $O(\eta_M^{-2/(1-\delta)})$ [15], and we expect this to hold as in Assumption 1(iii).

The standard prescription from SGD algorithms is that the sample size be maintained at a constant $M_t = M$ throughout the iterations. However, this would lead to biased sampling in the iterates given Theorem 3, which provides only $M_t \nearrow N$ as a control. Fixed bias violates a basic requirement for SGD that the gradient estimator $\nabla \hat{R}(\theta) = \Theta(\nabla R(\theta))$ (see, e.g., 4.3 in [3]). Hence, convergence of (6) cannot be guaranteed when $M_t = M$, $\forall t$, and this is amply demonstrated in the experimental setup in Section 3.

Moreover, $M_t \nearrow M$ as $t \to \infty$ reduces both the variance and bias in the gradient estimation. We no longer require $\gamma_t \searrow 0$ and thus Algorithm 1 chooses to take fixed step sizes. Since N is finite, any scheme to increase M_t will eventually end with $M_{t_{\max}} = N$ for some $t_{\max} < \infty$, at which point it is advisable to switch to a deterministic optimization algorithm. The key consideration then is that of "optimally" increasing the M_t. Here the tradeoff is between the reduction in stochastic error, which includes the bias and variance in $\nabla \hat{R}_{\mathcal{M}_t}(\theta_t)$, and the increased computational effort in each iteration as M_t increases.

In the remainder of this section we will argue that the choice in Algorithm 1 of geometrically increasing M_t is efficient. Our notion of efficiency will be developed w.r.t. the total computational budget W_t that is expended up till iterate t, which is the sum of the amount of individual work w_t in each iterate. From Proposition 1, we have that $w_t = O(M_t)$ for the D_{KL}-constraint case, while from Proposition 2 we have $w_t = O(M_t \log M_t)$ for the D_{χ^2}-constrained formulation.

Defining the ratio $\nu_t := M_t/M_{t+1}$ as the *growth factor* of the sequence $\{M_t\}$, we will consider the cases: (1) *Sub-geometric*: if $\nu_t \nearrow 1$ as $t \to \infty$, e.g., polynomial growth $\nu_t = 1 - \frac{1}{t}$; (2) *Geometric*: if $\nu_t = \nu < 1$, $\forall t$.

Lemma 7 *From the definition of the sample growth rates, we have: (1) If M_t grows geometrically, then $w_t = \Theta(W_t)$ for the two cases $w_t = O(M_t)$ and $w_t = O(M_t \log M_t)$. ; (2) If M_t grows sub-geometrically, then $w_t = o(W_t)$ when $w_t = O(M_t)$.*

Proof of Lemma 7: For (1), $\{M_t\}$ is geometric with rate $\nu < 1$. First consider the case when $w_t = O(M_t)$. In this case, writing $S_t := \frac{W_t}{w_t} = \sum_{s=1}^{t} \frac{M_s}{M_t} = \sum_{s=1}^{t} \prod_{s+1}^{t} \nu_i$, we have

$$S_t = \sum_{s=1}^{t} \nu^{(t-s-1)} = \sum_{s=1}^{t} \nu^{(s-1)} = \frac{1-\nu^t}{1-\nu} \to \frac{1}{1-\nu},$$

thus proving (1). When $w_t = O(M_t \log M_t)$, the same follows for super-geometric sequences because $c_k \le c < 1$ for some c and all $k \ge k_0 = k_0(c)$. When $w_t = O(M_t \log M_t)$ and $M_t = \nu^{-t}$, we obtain

$$\sum_{s \le t} M_s \log M_s = \log \nu \sum_{s \le t} -s\nu^{-s}$$

$$= \frac{\log \nu}{\nu} \sum_{s \le t} (-s)\nu^{-(s-1)} = \frac{\log \nu}{\nu} \frac{\partial}{\partial \nu} \left(\sum_{s \le t} \nu^{-s} \right)$$

$$= \frac{\log \nu}{\nu} \frac{t\nu^{-(t+1)} - \nu^{-t} + 1}{(\nu^{-1} - 1)^2}.$$

Dividing the last expression by $w_t = O(-t\nu^{-t} \log \nu)$ renders a slighlty different limit as $t \to \infty$, but $W_t = O(w_t)$ holds.

For (2), start with the case when $w_t = O(M_t)$. Consider any small $\epsilon > 0$. Since we have that $\nu_t \nearrow 1$ for sub-geometric growth of $\{M_t\}$, then for a sufficiently large t there exists $t_0(\epsilon)$ such that, $\forall t \ge s \ge t_0(\epsilon)$,

$$1 \ge \nu_s \ge \left(\frac{1}{t^{\epsilon/2}} \right)^{1/t^\epsilon} \ge 0.$$

Then, using the definition of S_t as before, we obtain

$$S_t \ge \sum_{s=t-t^\epsilon}^{t} \prod_{u=s+1}^{t} \nu_u \ge t^\epsilon \prod_{u=t-t^\epsilon}^{t} \nu_u \ge t^{\epsilon/2}.$$

Hence, as $t \to \infty$, we have that $S_t \to \infty$, thus proving the result. $\square$

Hence, geometrically growing sequences are sufficiently fast that the work done in the last iterate w_t is of the same order as the cumulative computational effort W_t expended up until t, while this is not the case for slower sequences. Our final result that characterizes the rate at which the expected optimality gap $O_{t+1} := \mathbb{E}_{\mathcal{M}}[R(\theta_{t+1})] - R(\theta_{\rm rob})$ drops as $M_t \nearrow N$.

Theorem 8 *Suppose the constant step-size $\gamma_t = \gamma$ satisfies $\gamma \le \min\{\frac{1}{4L}, 4c\}$, $\forall t$. Let $r = 1 - \frac{\gamma}{4c}$. We then have: (1) If M_t grows geometrically with parameter $\nu < r$, then for $t \le t_{\max}$, $O_{t+1} = O(r^t)$. Further, if we use D_{KL}-constraints, then $O_{t+1} = O(W_t^{-1}(r/\nu)^t)$, and if D_{χ^2}-constraints are used then $O_{t+1} = O(W_t^1 (r/\nu)^t/t)$; (2) If M_t grows sub-geometrically, then $O_{t+1} = O(w_t^{-1})$.*

Theorem 8 establishes that any sub-geometric rate of growth will lead to sub-optimal reduction in the optimality gap w.r.t. the total computational effort. Intuitively, this can be understood to happen

13

because the stochastic error drops to zero much slower than the deterministic error that can be attained for strongly convex optimization objective, and thus the stochastic error dominates. This together with the scaling implications of Lemma 7 yields that sub-geometric rates are suboptimal in the sense that $W_t O_{t+1} \to \infty$ as $t \nearrow$, indicating that the error O_t is unable to drop fast enough compared to the rate at which W_t grows.

Geometrically increasing the sampling will, on the other hand, attain a balance between the rate of convergence of the stochastic error and the deterministic improvement possible for strongly convex functions, thus attaining a better balance between the optimality gap and the level of computational effort. Note that the fastest convergence is attained when $\nu = r$, eliminating the $(r/\nu)^t$ inflation factor. However, r depends on c and L through γ and so is hard to obtain in practice.

The following lemma on the intermixing of sequences of real numbers will be useful in the proof of Theorem 8.

Lemma 9 *Let $\{a_{t,s}\}$, $t \geq s \geq 1$, be a triangular array of positive-values real numbers. Assume that:*

1. *There exists s^* and $\beta > 1$ such that $\frac{a_{t,s+1}}{a_{t,s}} \geq \beta$, $\forall s \in [s^*, t-1]$ and all $t \geq 1$. Moreover, $\limsup_t \frac{a_{t,s}}{a_{t,t}} = l_s < \infty$ for each $s \in [1, s*1]$. Then*

$$S_t = \sum_{s=1}^t a_{t,s} = O(a_{t,t}).$$

2. *There exists s^* and $\beta > 1$ such that $\frac{a_{t,s}}{a_{t,s+1}} \geq \beta$, $\forall s \in [1, s^*-1]$ and all $t \geq 1$. Moreover, $\limsup_t \frac{a_{t,s}}{a_{t,t}} = l_s < \infty$ for each $s \in [s*, t]$. Then*

$$S_t = \sum_{s=1}^t a_{t,s} = O(a_{t,1}).$$

Proof of Lemma 9: We prove (1), and (2) follows similarly. For large enough t and any $\epsilon > 0$, we have

$$
\begin{aligned}
S_t &= a_{t,t} \left(\sum_{s=0}^{s^*-1} \frac{a_{t,s}}{a_{t,t}} + \sum_{s=s^*}^{t-1} \frac{a_{t,s}}{a_{t,t}} \right) \\
&\leq a_{t,t} \left(s^* \epsilon + \sum_{s=0}^{s^*-1} \ell_s + \sum_{s=s^*}^{t-1} \beta^{s-t} \right),
\end{aligned} \tag{18}
$$

where the inequality follows from assumptions in (1). Since $\beta > 1$, $s^* < \infty$, and $\ell_s < \infty$, the term within parentheses on the right-hand side is finite and the assertion holds. $\square$

Proof of Theorem 8: Theorem 3 and Assumption 1(iii) shows that, for any θ and a set $\mathcal{M}$ sampled to have M support points,

$$
\begin{aligned}
\mathbb{E}_{\mathcal{M}} \left[\|\nabla\hat{R})\mathcal{M}(\theta) - \nabla R(\theta)\|_2^2 \right] &\leq \mathbb{E}_{\mathcal{M}} \left[\|\nabla\hat{R})\mathcal{M}(\theta) - \mathbb{E}_{\mathcal{M}}[\nabla\hat{R})\mathcal{M}(\theta)]\|_2^2 \right] + \|\mathbb{E}_{\mathcal{M}}[\nabla\hat{R})\mathcal{M}(\theta)] - \nabla R(\theta)\|_2^2 \\
&\leq O(\eta_M^{2/(1-\delta)}) + O(\eta_M^2) = O(\eta_M^2).
\end{aligned}
$$

Hence, the slower rate of drop in bias prevails as the rate at which the mean squared error drops to zero. Elementary algebraic manipulations yield the following two implications:

$$\mathbb{E}_{\mathcal{M}} \left[\|\nabla\hat{R})\mathcal{M}(\theta)\|_2^2 \right] \leq C\eta_M^2 + \|\nabla R(\theta)\|_2^2 \tag{19}$$

$$-\mathbb{E}_{\mathcal{M}} \left[\left(\nabla\hat{R})\mathcal{M}(\theta)\right)^T \nabla R(\theta) \right] \leq C\eta_M^2 - \|\nabla R(\theta)\|_2^2 - \mathbb{E}_{\mathcal{M}} \left[\|\nabla\hat{R})\mathcal{M}(\theta)\|_2^2 \right]. \tag{20}$$

The change in the optimality gap $R(\theta_t) - R(\theta_{\mathrm{rob}})$ after one iteration of (6) can be bounded as

$$\mathbb{E}_{\mathcal{M}}[R(\theta_{t+1})] - R(\theta_{\mathrm{rob}}) \leq \left(\mathbb{E}_{\mathcal{M}}[R(\theta_t)] - R(\theta_{\mathrm{rob}})\right) - \gamma \mathbb{E}_{\mathcal{M}}[\nabla R(\theta_k)^T \nabla \hat{R}_{\mathcal{M}}(\theta)] + \frac{L\gamma^2}{2}\mathbb{E}_{\mathcal{M}}\left[\|\nabla \hat{R}(\theta_t)\|_2^2\right]$$

$$\leq \left(\mathbb{E}_{\mathcal{M}}[R(\theta_t)] - R(\theta_{\mathrm{rob}})\right) + \frac{C\gamma\eta_M^2}{2} - \frac{\gamma}{2}\|\nabla R(\theta_t)\|_2^2 + \left(\frac{L\gamma^2 - \gamma}{2}\right)\left(\|\nabla R(\theta_t)\|_2^2 + C\eta_M^2\right)$$

$$\leq \left(\mathbb{E}_{\mathcal{M}}[R(\theta_k)] - R(\theta_{\mathrm{rob}})\right)\left(1 - \frac{2\gamma - L\gamma^2}{4c}\right) + \frac{CL\gamma^2\eta_M^2}{2}$$

$$\leq \left(1 - \frac{\gamma}{4c}\right)\left(\mathbb{E}_{\mathcal{M}}[R(\theta_k)] - R(\theta_{\mathrm{rob}})\right) + \frac{CL\gamma^2\eta_{M_t}^2}{2}. \tag{21}$$

The first inequality starts with the L-Lipschitzness of $\nabla R(\cdot)$, and the second inequality substitutes the relations in (19) and (20). The third inequality uses the c-strong convexity of $R(\theta)$, specifically the implication that $\|\nabla R(\theta)\|_2^2 \leq (R(\theta) - R(\theta_{\mathrm{rob}}))/2c$. The final inequality utilitzes the conditions imposed on γ.

Let $r = 1 - \gamma/4c < 1$. The form of (21) is quite informative, in that it clearly displays the tradeoff being made by the algorithm : the first summand provides a geometric reduction in the optimality gap, which is to be balanced with the stochastic error in the second summand. Defining $O_t := (\mathbb{E}_{\mathcal{M}}[R(\theta_t)] - R(\theta_{\mathrm{rob}}))$, then telescoping the optimality gap in the first summand renders

$$O_{t+1} \leq rO_t + \frac{CL\gamma^2\eta_{M_t}^2}{2}$$

$$\leq r^t O_0 + \frac{CL\gamma^2}{2}\sum_{s=0}^{t}\eta_{M_s}^2 r^{t-s+1}$$

$$\leq r^t O_0 + \frac{CL\gamma^2}{2}\sum_{s=0}^{t}\left(\frac{1}{M_s} + \frac{1-\delta}{N}\right)r^{t-s+1}$$

$$= r^t O_0 + \frac{CL\gamma^2}{2}\sum_{s=0}^{t}r^{t-s+1}\left(\prod_{u=0}^{s}\nu_u\right) + \frac{(1-\delta)CL\gamma^2}{N}\sum_{s=0}^{t}r^s$$

$$= r^t O_0 + \frac{(1-\delta)CL\gamma^2}{N}\frac{1-r^t}{1-r} + \frac{CL\gamma^2}{2}\sum_{s=0}^{t}r^{t-s+1}\left(\prod_{u=0}^{s}\nu_u\right). \tag{22}$$

The third inequality is obtained by noting that for an x, $0 < x < 1$, we have $(1-x)^{(1-\delta)} \leq 1 + (1-\delta)x$. In (22), we observe that the the error after t steps is as a result of the balance between r and ν_t.

Denote the sth term in the summation in the last term of (22) as $a_{t,s}$ (as per the notation in Lemma 9), and take the ratio of successive terms in the summation:

$$\frac{a_{t,s+1}}{a_{t,s}} = \frac{\nu_{s+1}}{r}.$$

This ratio determines the value to which the second term in (22) converges. If $\nu_{s+1} < r$, then by Theorem 9(ii) it is of size $O(r^t)$, exactly of the same size as the first term. This is indeed the case for when M_t are sampled geometrically with a $\nu < r$ as assumed in the theorem statement, and thus $O_{t+1} = O(r^t)$. For this case, we also have from Lemma 7(1) that $W_t = O(w_t)$. When D_{KL}-constraints are considered, we obtain $w_t = O(M_t) = O(\nu^{-t})$, which leads to the desired result. When D_{χ^2}-constraints are used, $w_t = O(M_t \log M_t) = O(\nu^{-t}t)$, rendering the desired result.

On the other hand, if $\nu_{s+1} > r$ for sufficiently many s, the stochastic noise dominates and by Lemma 9(i), the second term is of size $O\left(\prod_{s=0}^{t}\nu_s\right) = O(M_t)$. When M_t are sampled at sub-geometric rates, the $\nu_t \nearrow 1$ and so this condition holds. Hence, for sub-geometrically growing sample sizes, $O_{t+1} = O(M_t)$. When the computational effort is $w_t = O(M_t)$, we obtain the desired result by noting that, for this case, Lemma 7(ii) implies $w_t = o(W_t)$. $\square$

3   Experimental Results

Numerous experiments were conducted to empirically evaluate our new SGD algorithm in comparison with the full-gradient algorithm based on two main datasets from [8]: HIV-1 Protease Cleavage; and Adult Income.
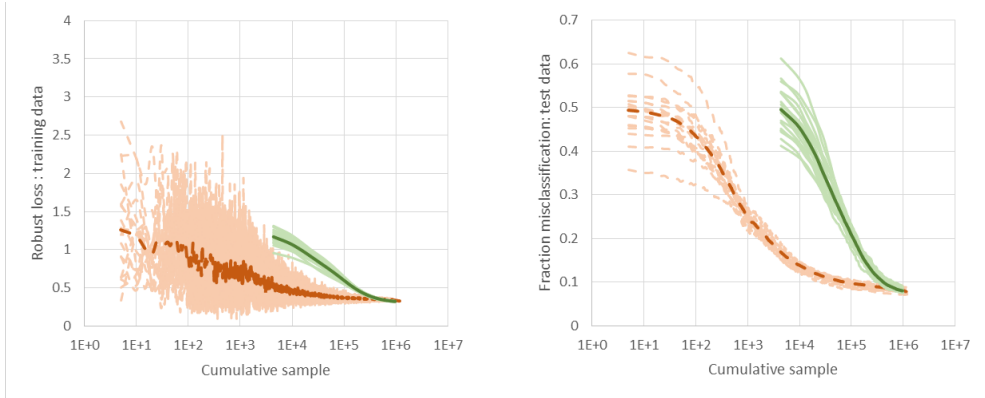
(a) $\rho = 0.1$, HIV-1 training data

(b) $\rho = 0.1$, HIV-1 testing data

Figure 1: Comparisons of sub-gradient (orange, dashed) and full-gradient (green, solid) algorithms.
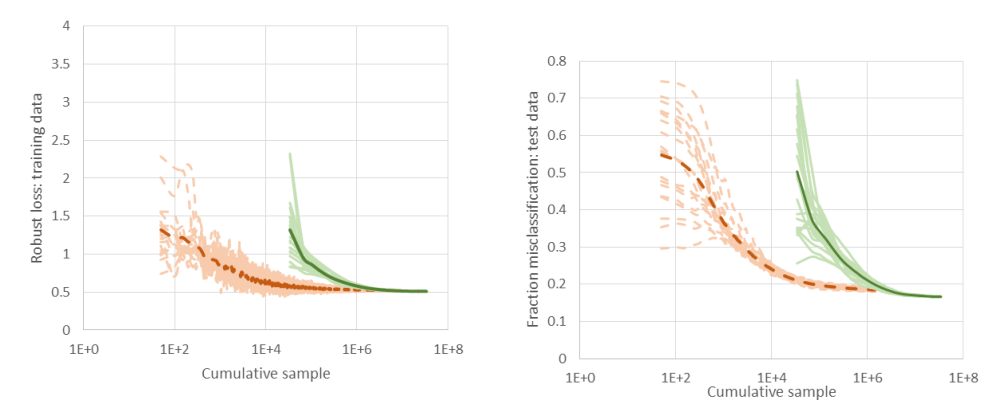
(a) $\rho = 0.1$, adult income training data

(b) $\rho = 0.1$, adult income testing data

Figure 2: Comparisons of sub-gradient (orange, dashed) and full-gradient (green, solid) algorithms.

3.1   HIV-1 Protease Cleavage

The HIV-1 protease cleavage dataset is compiled from four data source files, with the primary purpose to develop effective protease cleavage inhibitors by predicting whether the HIV-1 protease will cleave the protein sequence in its central position. We preprocessed the data to remove conflicting and overlapping samples, following the information in [13]. A logistic regression loss function $l(\theta; (x, y)) = \log(1 + \exp(-y\theta^t x))$ is used in this experiment, where x represents an $m \times n$ matrix with m the total number of samples (5830), n the binary 160-dimensional feature vector using orthogonal binary representation. The y label is coded as 1 if the HIV-1 protease cleaves at the center of an octamer, otherwise the y label coded as -1. The sample has 991 cleaved and 4839 non-cleaved labels. We initialized the values of θ to $U[-1, 1]$ in every experimental run.

The dataset was split by randomly selecting 25% of the data selected for testing, with the remaining data used for training. Figure 1 presents our empirical results from 20 experimental runs under our sub-gradient algorithm and under the full-gradient algorithm with $\rho = 0.1$, where the dashed and

bold lines respectively illustrate the average of the twenty runs; the corresponding figures for $\rho = 0.5$ and $\rho = 2.5$ are presented in Figure 3 and Figure 4 respectively.

(a) $\rho = 0.5$, HIV-1 training data

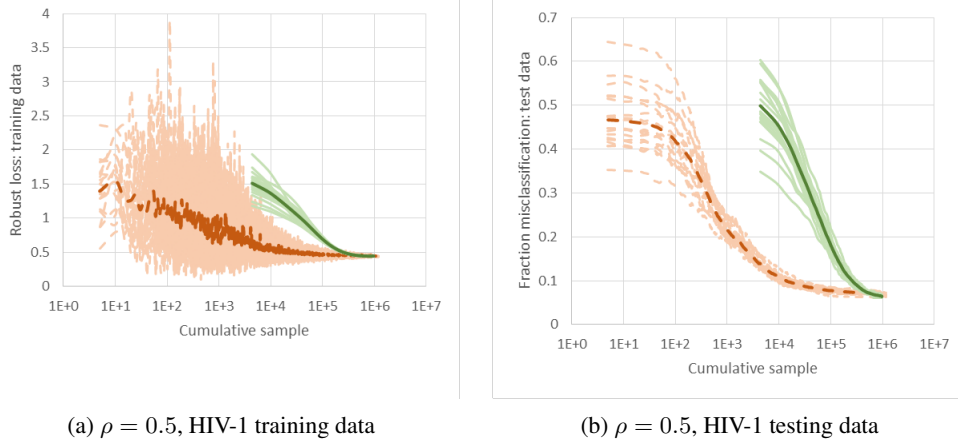(b) $\rho = 0.5$, HIV-1 testing data

Figure 3: Comparisons of the dynamically sampled stochastic sub-gradient algorithm (orange broken lines) and the full-sampled (solid green lines) algorithm.
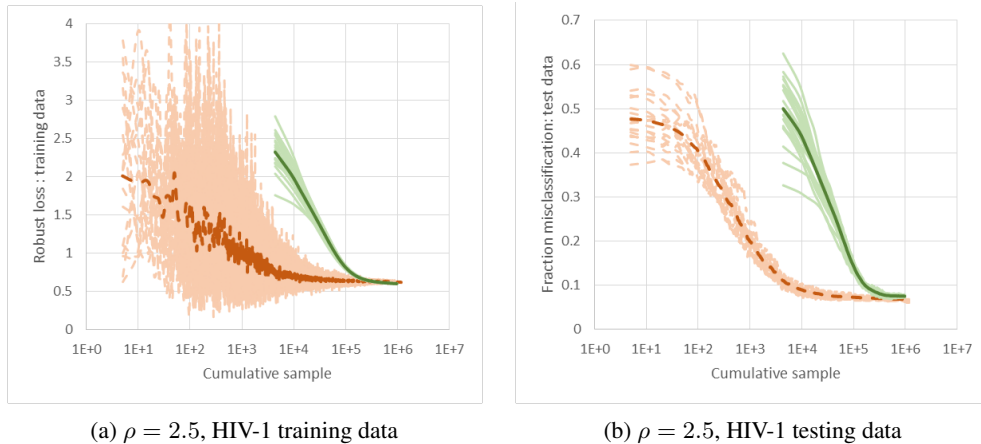

(a) $\rho = 2.5$, HIV-1 training data

(b) $\rho = 2.5$, HIV-1 testing data

Figure 4: Comparisons of the dynamically sampled stochastic sub-gradient algorithm (orange broken lines) and the N sampled (solid green lines) algorithm.

The leftmost plots compare the robust loss performance objective of the two algorithms based on the training testing data, and the rightmost plots compare the fractional misclassification performance of the two algorithms based on the testing data. It is readily apparent that our proposed SGD algorithm outperforms the full-gradient method for each value of ρ considered, with the best results obtained for $\rho = 0.1$. We note that our method can become somewhat unstable when the ρ value increases, which is consistent with our above results prescribing $\rho < 1$.

3.2   Adult Income

The adult income dataset is comprised of 32561 and 16281 samples for training and testing, respectively. Each sample has 15 attributes where the primary purpose to predict adult annual income w.r.t. the remaining 14 attributes. A logistic regression loss function $l(\theta; (x, y)) = \log(1 + \exp(-y\theta^t x))$ is used, where x represents an $m \times n$ matrix with m the total number of samples, n the binary 119-dimensional feature vector including binary encoding of categorical attributes, and the y label is coded as 1 if the annual income is above 50K, otherwise the y label coded as -1.

Figure 2 presents our empirical results from 20 experimental runs under our sub-gradient algorithm and under the full-gradient algorithm with $\rho = 0.1$, where the dashed and bold lines respectively illustrate the average of the twenty runs; the corresponding figures for $\rho = 0.5$ and $\rho = 2.5$ are presented in Figure 5 and Figure 6 respectively.

(a) $\rho = 0.5$, adult income training data

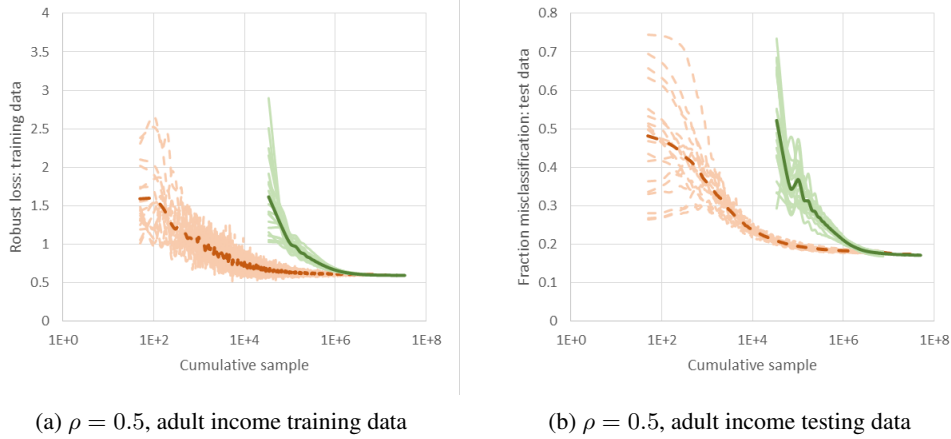(b) $\rho = 0.5$, adult income testing data

Figure 5: Comparisons of the dynamically sampled stochastic sub-gradient algorithm (orange broken lines) and the N sampled (solid green lines) algorithm.
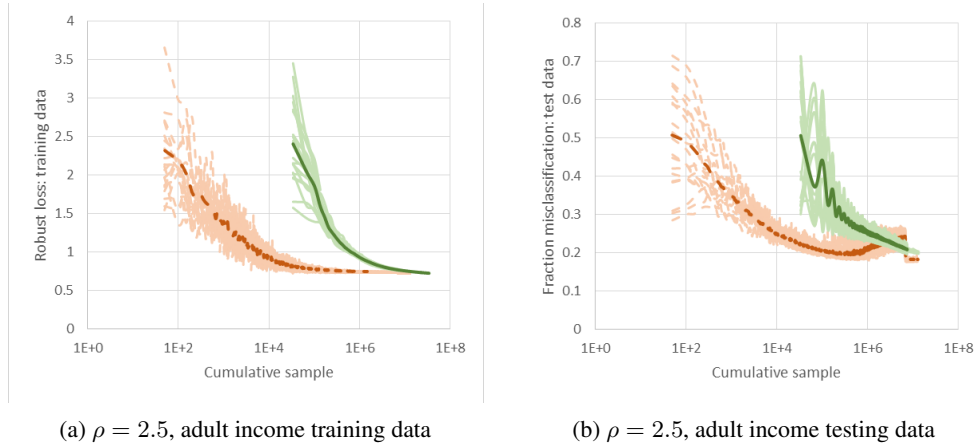

(a) $\rho = 2.5$, adult income training data

(b) $\rho = 2.5$, adult income testing data

Figure 6: Comparisons of the dynamically sampled stochastic sub-gradient algorithm (orange broken lines) and the N sampled (solid green lines) algorithm.

The leftmost plots compare the robust loss performance of the two algorithms based on the training testing data, and the rightmost plots compare the fractional misclassification performance of the two algorithms based on the testing data. Once again, it is readily apparent that our proposed SGD algorithm outperforms the full-gradient method for each value of ρ considered, with the best results obtained for $\rho = 0.1$. We note again that our method can become somewhat unstable when the ρ value increases, which is consistent with our above results prescribing $\rho < 1$.

References

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

[2] J. Blanchet, Y. Kang, and K. Murthy. Robust Wasserstein Profile Inference and Applications to Machine Learning. *ArXiv e-prints*, Oct. 2016.

[3] L. Bottou, F. E. Curtis, and J. Nocedal. Optimization Methods for Large-Scale Machine Learning. *ArXiv e-prints*, June 2016.

[4] J. Duchi, P. Glynn, and H. Namkoong. Statistics of Robust Optimization: A Generalized Empirical Likelihood Approach. *ArXiv e-prints*, Oct. 2016.

[5] S. Ghosh and H. Lam. Robust analysis in stochastic simulation: Computation and performance guarantees. *Operations Research*, pages To–appear, 2018.

[6] F. S. Hashemi, R. Pasupathy, and M. R. Taaffe. The adaptive sampling gradient method: Optimizing smooth functions with an inexact oracle. 2017. Under review.

[7] H. Lam. Robust sensitivity analysis for stochastic systems. *Mathematics of Operations Research*, 41(4):1248–1275, 2016.

[8] M. Lichman. Uci machine learning repository, 2013. URL `https://archive.ics.uci.edu/ml/datasets`.

[9] D. G. Luenberger. *Optimization by Vector Space Methods*. John Wiley & Sons, 1969.

[10] H. Namkoong and J. C. Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 2208–2216. Curran Associates, Inc., 2016.

[11] H. Namkoong and J. C. Duchi. Variance-based regularization with convex objectives. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 2971–2980. Curran Associates, Inc., 2017. URL `http://papers.nips.cc/paper/6890-variance-based-regularization-with-convex-objectives.pdf`.

[12] R. Pasupathy, P. Glynn, S. Ghosh, and F. S. Hashemi. On sampling rates in simulation-based recursions. *SIAM Journal on Optimization*, 28(1):45–73, 2018. doi: 10.1137/140951679. URL `https://doi.org/10.1137/140951679`.

[13] T. Rognvaldsson, L. You, and D. Garwicz. State of the art prediction of hiv-1 protease cleavage sites. *Bioinformatics*, 31(8):1204–1210, 2014.

[14] A. Sinha, H. Namkoong, and J. Duchi. Certifying Some Distributional Robustness with Principled Adversarial Training. In *ICLR 2018*. Oct. 2017.

[15] S. S. Wilks. *Mathematical Statistics*. John Wiley & Sons, 1962.