
Optimizing Black-box Metrics with Adaptive Surrogates

Qijia Jiang¹ Olaoluwa Adigun² Harikrishna Narasimhan³ Mahdi Milani Fard³ Maya Gupta³

Abstract

We address the problem of training models with black-box and hard-to-optimize metrics by expressing the metric as a monotonic function of a small number of easy-to-optimize surrogates. We pose the training problem as an optimization over a relaxed surrogate space, which we solve by estimating local gradients for the metric and performing inexact convex projections. We analyze gradient estimates based on finite differences and local linear interpolations, and show convergence of our approach under smoothness assumptions with respect to the surrogates. Experimental results on classification and ranking problems verify the proposal performs on par with methods that know the mathematical formulation, and adds notable value when the form of the metric is unknown.

1. Introduction

We consider the problem of training a machine learning model when the true evaluation metric is difficult to optimize on the training set. This general problem arises with many flavors and in different scenarios. For example, we may have a black-box metric whose mathematical expression is unknown or difficult to approximate with a convex training loss. The latter is particularly true with non-decomposable evaluation metrics, such as the F-measure or ranking metrics like Precision@ K , where it is not straight-forward to construct a differentiable objective that closely approximates the metric.

Another example is when the training labels are only a proxy for the true label. This arises in problems where one has access to cheap-to-acquire noisy labels, such as clicks, but wishes to optimize for a more expensive label, such as whether users rate a result as good. If we have access to a small auxiliary validation set with true labels, how can this

information be used to influence the training loss? Similar examples also arise when the training data has noisy features and we have a small validation set with clean features, or in machine learning fairness problems where the training data contains group-dependent noise, but we may have access to a small set of auxiliary clean data.

In many of the above scenarios, one wishes to optimize a black-box metric M over d model parameters, but does not have access to explicit gradients for M , nor is it practical to obtain reliable gradient estimates when d is large. We provide a general solution to this problem by choosing $K \ll d$ convex surrogate losses, and expressing M as an *unknown* monotonic function $\psi : \mathbb{R}_+^K \rightarrow \mathbb{R}$ of the K surrogates. We then reformulate the original problem as an optimization of ψ over the K -dimensional surrogate space. The choice of surrogates can be as simple as the hinge losses on positive and negative samples, which should work well for metrics like the F-measure, or the surrogates can be chosen to be a family of different convex losses to handle robustness to training noise given a small set of clean validation samples.

Our strategy is to estimate gradients for the unknown function ψ with respect to its K inputs by measuring changes in the metric M and the K surrogates for different perturbations on the model, and use the estimates for $\nabla\psi$ to perform *projected gradient descent* over the K -dimensional surrogate space. We show how the projection step can be implemented inexactly but with convergence guarantees by solving a convex problem in the original d parameters. We are thus able to *adaptively* combine the K surrogates to align well with the target metric M .

The main contributions of this paper include:

1. A novel formulation that poses the problem of optimizing a black-box metric as a lower-dimensional problem in a surrogate space.
2. A projected gradient descent based training algorithm using finite-differences and local linear interpolations to estimate gradients.
3. Theoretical results showing convergence to a stationary point under smoothness assumptions on ψ .
4. Experiments showing that the proposed approach works as well as methods that take advantage of the form of the metric if known, but can give substantial gains when the metric truly is a black-box.

*Equal contribution ¹Stanford ²USC ³Google Research. Correspondence to: Harikrishna Narasimhan <hnarasimhan@google.com>.

2. Related Work

There has been much work on directly optimizing specialized classes of evaluation metrics during training. These include approaches that relax the metric using convex surrogates (Joachims et al., 2005; Kar et al., 2014; Narasimhan et al., 2015a; Kar et al., 2016), plug-in or post-shift methods that tune a threshold on estimates of class probabilities (Ye et al., 2012; Koyejo et al., 2014; Narasimhan et al., 2014; Yan et al., 2018), reduction approaches that formulate a sequence of cost-sensitive learning tasks (Parambath et al., 2014; Narasimhan et al., 2015b; Alabi et al., 2018; Narasimhan, 2018), and approaches that use constrained optimization and game-based formulations (Eban et al., 2017; Narasimhan et al., 2019).

However, all the above approaches require the evaluation metric to be available in closed-form. Of these, the closest to ours is the approach of Narasimhan et al. (2015b), which reformulates the learning problem as an optimization problem over the space of confusion matrices. To ensure the constraint set is convex, this approach requires the use of stochastic classifiers, and the theoretical guarantees assume that the metrics are convex or pseudo-convex in the confusion matrix. In contrast, we do not require stochastic classifiers, and can handle general metrics.

Recently, there has been some work on optimizing evaluation metrics that are only available as a black-box. Zhao et al. (2019) approximate black-box metrics with a weighted training loss where the weighting function acts on a low-dimensional embedding of each example, and a validation set is used to estimate the parameters of the example-weighting function. A related approach by Ren et al. (2018) uses meta-gradient descent to re-weight the training examples to handle training set biases and label noise. In contrast we model the unknown metric as a function of surrogate losses, and directly estimate the metric gradients, rather than estimating a weighting function on each example.

Huang et al. (2019) also propose jointly adaptively learning a metric with the model training. They use a parametric form for their learned metric, whereas we nonparametrically estimate the metric gradients. They use reinforcement learning to align the training objective’s optimum with that of the true metric, whereas we use gradient descent over a surrogate space. They do not provide any theoretical guarantees.

Grabocka et al. (2019) express the metric as a set function that maps each prediction to an embedding and maps the average embedding across all examples to the predicted metric. They jointly optimize the parameters for the loss and the model. This approach is similar to ours in that it expresses the metric as a function on surrogate losses, and attempts to learn that function. However, our approach is different in two key points. First, we take as *given* known-useful sur-

rogate losses, whereas they *learn* decomposable surrogate mappings from scratch. Second, they parameterize their surrogate functions and final mapping as neural networks, whereas we nonparametrically adaptively estimate the local gradients. They provide limited theoretical guarantees.

Similar to Grabocka et al. (2019), the work of Wu et al. (2018) also learns a parameterized metric (e.g. as a neural network). An auxiliary parametric “teacher” model is used to adaptively learn the parameters for the metric that will maximize performance on a validation set. They do not provide theoretical guarantees.

3. Problem Setup and High-level Approach

Let \mathcal{X} be some instance space and \mathcal{Y} be the label space. Let $f_\theta : \mathcal{X} \rightarrow \mathbb{R}$ be a model parametrized by $\theta \in \mathbb{R}^d$ that outputs a score $f_\theta(\mathbf{x})$ for instance $\mathbf{x} \in \mathcal{X}$. One can use this score to make a prediction; e.g. for binary classification problems, one predicts $\text{sign}(f_\theta(\mathbf{x}))$. We measure performance w.r.t. a test distribution D over $\mathcal{X} \times \mathcal{Y}$. We consider two scenarios, one where we are provided a training sample S of n examples directly drawn from D , and the other where the training sample S is drawn from a noisy distribution, and we are provided a smaller clean validation set from D .

The performance of f_θ is evaluated by a metric $M : \mathbb{R}^d \rightarrow [0, 1]$ computed on D , where M may be as simple as the error rate $M_{err}(\theta) = \mathbf{E}_{(\mathbf{x}, y) \sim D} [yf_\theta(\mathbf{x}) > 0]$ (or an estimate), or M may be a complex, non-decomposable metric such as Precision@ K that depends on the scores and the distribution in a more intricate manner. We consider settings where the form of M is unknown, and the metric is available only as a black-box, i.e., for a given $\theta \in \mathbb{R}^d$, we can evaluate $M(\theta)$. The goal is to learn a good f_θ by solving:

$$\min_{\theta \in \mathbb{R}^d} M(\theta). \quad (1)$$

3.1. Reformulation with Surrogates

To optimize (1), one could directly estimate gradients of M with respect to the d parameters, but d is usually too large for that to be practical. To relax (1) to a more tractable problem, we take as given K convex surrogate loss functions $\ell_1, \dots, \ell_K : \mathbb{R}^d \rightarrow \mathbb{R}_+$ where $K \ll d$, and express M as an *unknown* non-decreasing function of the K surrogates, with an *unknown* slack:

$$M(\theta) = \psi(\ell_1(\theta), \dots, \ell_K(\theta)) + \epsilon(\theta),$$

where $\psi : \mathbb{R}_+^K \rightarrow [0, 1]$ is *monotonic* but possibly non-convex, and the slack $\epsilon : \mathbb{R}^d \rightarrow [-1, 1]$ determines how well the metric can be approximated by the K surrogates. Note that this decomposition of M is not unique. Our results hold for any such decomposition, but to enable a tighter analysis we consider a ψ for which the associated worst-case slack over all θ , i.e., $\max_{\theta \in \mathbb{R}^d} |\epsilon(\theta)|$ is the minimum.

Here are examples of target metrics and convex surrogates.

Example 1 (Classification Metrics). Consider the task of minimizing the G-mean metric given by $1 - \sqrt{\text{TPR} \times \text{TNR}}$, where TPR is the true positive rate and TNR is the true negative rate. This metric promotes high accuracies on both the positive and negative class and is popular for classification tasks where there is class imbalance (Daskalaki et al., 2006). Possible surrogates for this metric include the average logistic or hinge losses on the positive and negatives examples as these serve as proxies for the TPR and TNR. It is reasonable to assume monotonic ψ here, since lower surrogate values tend to produce better TPR and TNR values, and in turn lower G-means. The F-measure is another popular metric that can be written as a monotonic function of the TPR and TNR (Koyejo et al., 2014), and there again the average positive and negative losses would make good surrogates.

Example 2 (Misaligned Training Data). Consider minimizing a metric using a training dataset that is noisy or misaligned with the test distribution, but we have access to a small validation set with clean data. The metric M here is evaluated on the clean validation set, and the surrogates ℓ_1, \dots, ℓ_K might be convex l_p losses on the training data with different values of $p > 1$ to tune the noise robustness. In this case, the precise mathematical relationship ψ between the validation metric and the surrogates is unknown.

Example 3 (ML Fairness Problems). For blackbox ML fairness metrics, good surrogates might be logistic losses on the positive and negative samples for different groups.

Example 4 (Ranking Metrics). Consider optimizing a ranking metric such as precision@ K . While there are different convex surrogates available for this metric (Joachims, 2005; Narasimhan et al., 2015a), the surrogate that performs the best can vary with the application. We have also observed in practice that sometimes setting a different value of K in the training loss produces a better precision@ K during evaluation time. The proposed set-up gives us a way to combine multiple available ranking surrogates (possibly with different K values) to align well with the test metric.

3.2. High-level Approach

Let $\mathcal{L} := \{(\ell_1(\theta), \dots, \ell_K(\theta)) \mid \theta \in \mathbb{R}^d\}$ be the set of feasible surrogate profiles. We then seek to approximate (1) by ignoring the slack ϵ and posing the problem as an optimization of ψ over the K -dimensional set \mathcal{L} :

$$\min_{\ell \in \mathcal{L}} \psi(\ell). \quad (2)$$

Our high-level idea is to solve this re-formulated problem by applying *projected gradient descent* over \mathcal{L} .

However, there are many challenges in implementing this idea. First, while each ℓ_k is convex, the space of feasible surrogates \mathcal{L} is not necessarily a convex set. Second, the

function ψ is unknown to us, and therefore we need to estimate gradients for ψ with only access to the metric M and the surrogates ℓ . Third, we would need to implement projections onto the K -dimensional surrogate space without explicitly constructing this set.

4. Surrogate Projected Gradient Descent

We now explain how we tackle the above challenges.

4.1. Convexifying the Surrogate Space

To turn (2) into a problem over a convex domain, we define the *epigraph* of the convex surrogate function profiles:

$$\mathcal{U} := \{\mathbf{u} \in \mathbb{R}_+^K \mid \mathbf{u} \geq \ell(\theta) \text{ for some } \theta \in \mathbb{R}^d\}.$$

Observation 1. \mathcal{U} is a convex superset of \mathcal{L} .

We then optimize ψ over this K -dimensional convex set:

$$\min_{\mathbf{u} \in \mathcal{U}} \psi(\mathbf{u}). \quad (3)$$

This relaxation preserves the optimizer for (2) because ψ is monotonic and \mathcal{U} consists of upper bounds on surrogate profiles in \mathcal{L} :

Observation 2. For any $\mathbf{u}^* \in \underset{\mathbf{u} \in \mathcal{U}}{\operatorname{argmin}} \psi(\mathbf{u})$, there exists $\ell^* \in \mathcal{L}$, $\ell^* \leq \mathbf{u}^*$, such that $\psi(\ell^*) = \psi(\mathbf{u}^*)$.

4.2. Projected Gradient Descent over \mathcal{U}

We then perform projected gradient descent over \mathcal{U} . We maintain iterates \mathbf{u}^t in \mathcal{U} , and at each step, (i) estimate the gradient of ψ w.r.t. the K -dimensional point \mathbf{u}^t , (ii) perform a descent step: $\tilde{\mathbf{u}}^{t+1} = \mathbf{u}^t - \eta \nabla \psi(\mathbf{u}^t)$, for some $\eta > 0$, and (iii) project $\tilde{\mathbf{u}}^{t+1}$ onto \mathcal{U} to get the next iterate \mathbf{u}^{t+1} .

In order to implement these steps without knowing ψ , or having direct access to the set \mathcal{U} , we simultaneously maintain iterates θ^t in the original parameter space that map to iterates $\mathbf{u}^t \in \mathcal{U}$, i.e., for which $\mathbf{u}^t = \ell(\theta^t)$.

Now to estimate gradients without direct access to ψ , we measure changes in the K surrogates $\ell(\cdot)$ and changes in the metric $M(\cdot)$ at different perturbations of θ^t and compute estimates of $\nabla \psi(\mathbf{u}^t)$ based on finite-differences or local linear interpolations. To compute projections without direct access to \mathcal{U} , we formulate a convex optimization problem over the original parameters θ , and show that this results in an over-constrained projection onto \mathcal{U} .

Thus we maintain iterates (\mathbf{u}^t, θ^t) such that $\mathbf{u}^t = \ell(\theta^t)$, and execute the following at every iteration:

$$\begin{aligned} \tilde{\mathbf{u}}^{t+1} &= \mathbf{u}^t - \eta \operatorname{gradient}_{\psi}(\theta^t; M, \ell) \\ (\mathbf{u}^{t+1}, \theta^{t+1}) &= \operatorname{project}_{\mathcal{U}}(\tilde{\mathbf{u}}^{t+1}; \ell). \end{aligned}$$

Figure 1 gives a schematic description of the updates. The gradient computation takes the current θ^t as input and




Figure 1. PGD over K -dimensional set \mathcal{U} . ‘project’ performs an over-constrained projection onto \mathcal{U} . ‘gradient’ probes M and ℓ and returns an estimate $\hat{\mathbf{g}}^{t+1} \in \mathbb{R}^K$ for $\nabla \psi$.




Figure 2. **Over-constrained projection.** The space of surrogate profiles $\mathcal{L} = \{(\ell_1(\theta), \ell_2(\theta)) \mid \theta \in \mathbb{R}^d\}$ is a non-convex set (solid line), and its epigraph $\mathcal{U} = \{\mathbf{u} \geq \ell \mid \ell \in \mathcal{L}\}$ is convex (shaded region). For the point $\tilde{\mathbf{u}}_a$ outside \mathcal{U} , the solution \mathbf{u}_a to (4) is the same as the exact projection $\Pi(\tilde{\mathbf{u}}_a)$ onto \mathcal{U} . For the point $\tilde{\mathbf{u}}_b$ inside the set, $\Pi(\tilde{\mathbf{u}}_b) = \tilde{\mathbf{u}}_b$, whereas \mathbf{u}_b is one of many solutions to (4) on the boundary and with $\mathbf{u}_b \leq \Pi(\tilde{\mathbf{u}}_b)$ in each coordinate.

probes M and ℓ to return an estimate of $\nabla \psi(\mathbf{u}^t)$. We elaborate on how we estimate gradients in Section 5. The projection computation takes the updated $\tilde{\mathbf{u}}^{t+1}$ as input and returns a point \mathbf{u}^{t+1} in \mathcal{U} and an associated θ^{t+1} such that $\mathbf{u}^{t+1} = \ell(\theta^{t+1})$. We explain this next.

4.3. Over-constrained Projection

To implement the projection without explicit access to \mathcal{U} , we set up an optimization over θ by penalizing a clipped L_2 -distance between the surrogate profile $\ell(\theta)$ and $\tilde{\mathbf{u}}^{t+1}$:

$$\begin{aligned} \theta^{t+1} &\in \operatorname{argmin}_{\theta \in \mathbb{R}^d} \|(\ell(\theta) - \tilde{\mathbf{u}}^{t+1})_+\|^2 \\ \mathbf{u}^{t+1} &= \ell(\theta^{t+1}), \end{aligned} \quad (4)$$

where $(z)_+ := \max\{0, z\}$ is applied element-wise and $\|\cdot\|$ is the L_2 -norm. Note that we penalize errors in only one direction (i.e. the errors where $\ell_k(\theta) \geq \tilde{u}_k^{t+1}$). This has the advantage of the optimization problem being convex. Moreover, as we show below, (4) results in an over-constrained projection: any solution \mathbf{u}^{t+1} to (4) is feasible (i.e. is in \mathcal{U}), and for a monotonic ψ , yields a ψ -value that is no worse than what we would get with an exact projection.

Lemma 1. *Let \mathbf{u}^+ be the exact projection of $\tilde{\mathbf{u}}^{t+1} \in \mathbb{R}_+^K$ onto \mathcal{U} . For any solution \mathbf{u}^{t+1} to (4), we have $\mathbf{u}^{t+1} \in \mathcal{U}$, $\mathbf{u}^{t+1} \leq \mathbf{u}^+$, and for a monotonic ψ , $\psi(\mathbf{u}^{t+1}) \leq \psi(\mathbf{u}^+)$.*

Problem (4) may not have a unique solution. For example,

Algorithm 1 Surrogate Projected Gradient Descent

- 1: **Input:** Black-box metric M , surrogate loss functions $\ell_1, \dots, \ell_K : \mathbb{R}^d \rightarrow \mathbb{R}_+^K$, hyper-parameters: T, η
- 2: Initialize $\theta^1 \in \mathbb{R}^d, \mathbf{u}^1 = \ell(\theta^1)$
- 3: **for** $t = 1$ **to** T **do**
- 4: **Gradient estimate:** Obtain an estimate $\hat{\mathbf{g}}^t$ for gradient $\nabla \psi(\mathbf{u}^t)$ by invoking Algorithms 2 or 3 with inputs θ^t, M and ℓ_1, \dots, ℓ_K
- 5: **Gradient update:** $\tilde{\mathbf{u}}^{t+1} = \mathbf{u}^t - \eta \hat{\mathbf{g}}^t$
- 6: **Over-constrained projection:** Solve:

$$\theta^{t+1} \in \operatorname{argmin}_{\theta \in \mathbb{R}^d} \|(\ell(\theta) - \tilde{\mathbf{u}}^{t+1})_+\|^2$$

to accuracy $\mathcal{O}(\frac{1}{\beta^2 T})$ and set $\mathbf{u}^{t+1} = \ell(\theta^{t+1})$

- 7: **end for**

when $\tilde{\mathbf{u}}^{t+1}$ is in the interior of \mathcal{U} , the exact projection \mathbf{u}^+ is the same as $\tilde{\mathbf{u}}^{t+1}$, whereas the solutions to (4) are the points \mathbf{u} on the boundary of \mathcal{U} with $\mathbf{u} \leq \mathbf{u}^+$ (see Figure 2). As ψ is monotonic, picking any of these solutions for the next iterate doesn’t hurt the convergence of the algorithm.

An outline of the projected gradient descent with this inexact projection is presented in Algorithm 1. One can interpret the algorithm as *adaptively* combining the K surrogates ℓ_k ’s to optimize the metric M (see Appendix C for the details).

4.4. Convergence Guarantee

We show convergence of Algorithm 1 to a stationary point of $\psi(\ell(\cdot))$. Since we probe M to estimate gradients for ψ , the errors in the estimate would depend on how closely $\psi(\ell(\cdot))$ approximates M , and in turn on the magnitude of the slack term ϵ . We assume here that the gradient estimation error $\mathbf{E}[\|\hat{\mathbf{g}}^t - \nabla \psi(\ell(\theta^t))\|^2]$ at each step t is bounded by a $\kappa_\epsilon \in \mathbb{R}_+$ that depends on the slack ϵ . In Section 5, we present gradient estimates that satisfy this condition.

Theorem 2 (Convergence of Algorithm 1). *Let $M(\theta) = \psi(\ell(\theta)) + \epsilon(\theta)$, for a ψ that is monotonic, β -smooth and L -Lipschitz, and the worst-case slack $\max_{\theta \in \mathbb{R}^d} |\epsilon(\theta)|$ is the minimum among all such decompositions of M .*

Suppose each ℓ_k is γ -smooth and Φ -Lipschitz in θ with $\|\ell(\theta)\| \leq G, \forall \theta$. Suppose the gradient estimates $\hat{\mathbf{g}}^t$ satisfy $\mathbf{E}[\|\hat{\mathbf{g}}^t - \nabla \psi(\ell(\theta^t))\|^2] \leq \kappa_\epsilon, \forall t \in [T]$ and the projection step satisfies $\|(\ell(\theta^{t+1}) - \tilde{\mathbf{u}}^t)_+\|^2 \leq \min_{\theta \in \mathbb{R}^d} \|(\ell(\theta) - \tilde{\mathbf{u}}^t)_+\|^2 + \mathcal{O}(\frac{1}{\beta^2 T}), \forall t \in [T]$. Set stepsize $\eta = \frac{1}{\beta^2}$.

Then Algorithm 1 converges to an approximate stationary point of $\psi(\ell(\cdot))$:

$$\min_{1 \leq t \leq T} \mathbf{E}[\|\nabla \psi(\ell(\theta^t))\|^2] \leq C \left(\frac{\beta}{\sqrt{T}} + \sqrt{\kappa_\epsilon} + \sqrt{L \kappa_\epsilon^{1/4}} \right),$$

where the expectation is over the randomness in the gradient estimates, and $C = \mathcal{O}(KL(\gamma(G + \frac{L}{\beta^2}) + \Phi^2))$.

Remark 1 (Stationary point of M). When the gradient estimation error κ_ϵ is small and the number of steps $T \rightarrow \infty$, the algorithm reaches a model θ with a small gradient norm $\|\nabla\psi(\ell(\cdot))\|$. If additionally the slack term ϵ is Lipschitz in θ , then this implies that the algorithm also converges to an approximate stationary point of the metric M .

The proof of Theorem 2 proceeds in two parts. We first show that the algorithm converges to an approximate stationary point of ψ over \mathcal{U} . For this, we extend recent results (Ghadimi et al., 2016) on convergence of projected gradient descent for smooth non-convex objectives. We then exploit the smoothness of the surrogates ℓ to show that this result translates to the algorithm converging to an approximate stationary point of $\psi(\ell(\cdot))$ w.r.t. θ .

Remark 2 (Prior convergence results). A key difference between our analysis and prior works on zeroth-order gradient methods (Duchi et al., 2015; Ghadimi et al., 2016; Nesterov and Spokoiny, 2017) is that we do not directly optimize the given objective over the space of parameters θ , and instead perform an optimization over a relaxed surrogate space \mathcal{U} that is not directly specified, and do so using inexact projections and approximate gradient estimates.

5. Gradient Estimation Techniques

We now address the issue of estimating the gradient $\hat{\mathbf{g}}^t$ of ψ at a given $\ell(\theta^t)$ without explicit access to ψ . We provide an algorithm based on finite-differences, and another based on local linear interpolations. We also show error bounds for these algorithms, i.e. bound the errors κ_ϵ in Theorem 2.

5.1. Finite Differences

We first consider the case where both the surrogates ℓ and metric M are evaluated on the same sample. Let $\mathbf{f}_\theta := [f_\theta(\mathbf{x}_1), \dots, f_\theta(\mathbf{x}_n)]^\top \in \mathbb{R}^n$ denote the scores of the model θ computed on the n training examples. We overload notation and use $M(\mathbf{f}_\theta, \mathbf{y})$ to denote the value of the evaluation metric M on the model scores $\mathbf{f}_\theta \in \mathbb{R}^n$ and labels $\mathbf{y} \in \mathcal{Y}^n$. Similarly, we use $\ell_k(\mathbf{f}_\theta, \mathbf{y})$ to denote the value of surrogate loss ℓ_k on \mathbf{f}_θ and \mathbf{y} .

We present our method in Algorithm 2. We adopt a standard finite-difference gradient estimate (Nesterov and Spokoiny, 2017), which requires us to perturb the surrogates ℓ with random Gaussian vectors $Z^1, \dots, Z^m \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_K)$, evaluate ψ at the perturbed surrogate profiles, and calculate

$$\frac{1}{m} \sum_{j=1}^m \frac{\psi(\ell(\mathbf{f}_\theta, \mathbf{y}) + \sigma Z^j) - \psi(\ell(\mathbf{f}_\theta, \mathbf{y}))}{\sigma},$$

for $\sigma > 0$. In our case, we cannot directly perturb the surrogates ℓ and evaluate changes in ψ . Instead, we perturb the scores \mathbf{f}_θ so that the corresponding changes in ℓ follows a Gaussian distribution, and evaluate the difference between

Algorithm 2 Finite-difference Gradient Estimate

- 1: **Input:** $\theta' \in \mathbb{R}^d, M, \ell_1, \dots, \ell_K$
 - 2: **Hyper-parameters:** Num of perturbations m, σ
 - 3: Draw $Z^1, \dots, Z^m \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_K)$
 - 4: Find $\Delta^j \in \mathbb{R}^n$ s.t. $\ell(\mathbf{f}_{\theta'} + \Delta^j, \mathbf{y}) = \ell(\mathbf{f}_{\theta'}, \mathbf{y}) + \sigma Z^j$,
for $j = 1, \dots, m$
 - 5: $\hat{\mathbf{g}} = \frac{1}{m} \sum_{j=1}^m \frac{M(\mathbf{f}_{\theta'} + \Delta^j, \mathbf{y}) - M(\mathbf{f}_{\theta'}, \mathbf{y})}{\sigma} Z^j$
 - 6: **Output:** $\hat{\mathbf{g}}$
-

Algorithm 3 Linear Interpolation Gradient Estimate

- 1: **Input:** $\theta' \in \mathbb{R}^d, M, \ell_1, \dots, \ell_K$
 - 2: **Hyper-parameters:** Num of perturbations m, σ
 - 3: Draw $Z_1^1, \dots, Z_1^m, Z_2^1, \dots, Z_2^m \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$
 - 4: $\mathbf{H}_{j,:} = \ell(\theta' + \sigma Z_1^j) - \ell(\theta' + \sigma Z_2^j), j = 1, \dots, m$
 - 5: $\mathbf{M}_{j,:} = M(\theta' + \sigma Z_1^j) - M(\theta' + \sigma Z_2^j), j = 1, \dots, m$
 - 6: $\hat{\mathbf{g}} \in \operatorname{argmin}_{\hat{\mathbf{g}} \in \mathbb{R}^K} \|\mathbf{H}\hat{\mathbf{g}} - \mathbf{M}\|^2$
 - 7: **Output:** $\hat{\mathbf{g}}$
-

the metric M at the original and perturbed scores. This is possible, for example, when each $\ell_k(\mathbf{f}_\theta, \mathbf{y})$ is an average of point-wise losses $\phi_k(y_i f_\theta(x_i))$ on different subsets of the data, for some invertible function $\phi_k : \mathbb{R} \rightarrow \mathbb{R}$, in which case, it is easy to compute the right amount of perturbation to the scores \mathbf{f}_θ to produce the desired perturbation in ℓ_k .

Lemma 3 (Finite difference estimate). *Let M be as defined in Theorem 2 and $|\epsilon(\theta)| \leq \bar{\epsilon}, \forall \theta$. Let $\hat{\mathbf{g}}$ be returned by Algorithm 2 for a given θ', m perturbations and $\sigma = \frac{\sqrt{\bar{\epsilon}}}{\sqrt{K}\beta^2}$.*

$$\mathbf{E} [\|\hat{\mathbf{g}} - \nabla\psi(\ell(\theta'))\|^2] \leq \mathcal{O} \left(\frac{L^2 K}{m} + \bar{\epsilon} K^2 \beta^2 \right),$$

where the expectation is over the random perturbations.

This gives a bound on κ_ϵ in Theorem 2 when Algorithm 2 is used for gradient estimates. Note the error depends on the slack magnitude $\bar{\epsilon}$, and decreases with more perturbations.

5.2. Local Linear Interpolations

The finite-difference approach is not applicable to settings where the metric is evaluated on a validation sample but the surrogates are evaluated on training examples (as in Example 2), or where finding the right amount of perturbation on the scores is difficult. For such cases we present a local linear interpolation based approach in Algorithm 3, where we perturb the model parameters θ instead of the scores.

We use the fact that a smooth function ψ can be locally approximated by a linear function, and estimate the gradient of ψ of at $\ell(\theta)$ by perturbing θ , measuring the corresponding differences in the surrogates ℓ and the metric M , and fitting a linear function from the surrogate differences to

the metric differences. Specifically, for d model parameters, we draw two independent sets of d -dimensional Gaussian perturbations $Z_1^1, \dots, Z_1^m, Z_2^1, \dots, Z_2^m \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$, and return a linear fit from $\mathbf{H} = [\ell(\theta + \sigma Z_1^j) - \ell(\theta + \sigma Z_2^j)]_{j=1}^m$ to $\mathbf{M} = [M(\theta + \sigma Z_1^j) - M(\theta + \sigma Z_2^j)]_{j=1}^m$.

Lemma 4 (Linear interpolation estimate). *Let M be defined as in Theorem 2 and $|\epsilon(\theta)| \leq \bar{\epsilon}, \forall \theta$. Assume each ℓ_k is Φ -Lipschitz in θ w.r.t. the L_∞ -norm, and $\|\ell(\theta)\| \leq G \forall \theta$. Suppose for a given θ', σ and perturbation count m , the expected covariance matrix for the left-hand-side of the linear system \mathbf{H} is well-conditioned, and has the smallest singular value $\lambda_{\min}(\sum_{i=1}^m \mathbf{E}[\mathbf{H}_i \mathbf{H}_i^\top]) = \mathcal{O}(m\sigma^2\Phi^2)$. Then setting $\sigma = \tilde{\mathcal{O}}\left(\frac{G^{1/3}\bar{\epsilon}^{1/3}}{\Phi K^{3/2}\beta^{1/3}}\right)$ and $m = \tilde{\mathcal{O}}\left(\frac{G^4 K^9 \beta^2}{\bar{\epsilon}^2}\right)$, Algorithm 3 returns w.h.p. (over draws of random perturbations) a gradient estimate $\hat{\mathbf{g}}$ that satisfies:*

$$\|\hat{\mathbf{g}} - \nabla\psi(\ell(\theta'))\|^2 \leq \tilde{\mathcal{O}}\left(G^{1/3}\bar{\epsilon}^{1/3}K^3\beta^{2/3}\right).$$

We show in Appendix A.5 how this high probability statement can then be used to derive a bound on the expected errors κ_ϵ in Theorem 2. Prior works provide error bounds on a similar gradient estimate under an assumption that the perturbation matrix \mathbf{H} can be chosen to be invertible (Conn et al., 2008; 2009; Berahas et al., 2019). In our case, however, \mathbf{H} is not chosen explicitly, but instead contains measurements of changes in surrogates for random perturbations on θ . Hence to show an error bound, we need a slightly subtle condition on the correlation structure of the surrogates (that essentially says the variance of the perturbed surrogates are large enough and the rates are not strongly correlated with each other), which we express as a condition on the smallest singular value of the covariance of \mathbf{H} .

5.3. Handling Non-smooth Metrics

For ψ that is non-smooth and Lipschitz, we extend the finite difference gradient estimate in Section 5.1 with a two-step perturbation. We draw two sets of Gaussian random vectors $Z_1^1, \dots, Z_1^m, Z_2^1, \dots, Z_2^m \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_K)$ and calculate

$$\frac{1}{m} \sum_{j=1}^m \frac{\psi(\ell(\mathbf{f}_\theta, \mathbf{y}) + \sigma_1 Z_1^j + \sigma_2 Z_2^j) - \psi(\ell(\mathbf{f}_\theta, \mathbf{y}) + \sigma_1 Z_1^j)}{\sigma_2}$$

for $\sigma_1, \sigma_2 > 0$, where we perturb ℓ by perturbing the scores \mathbf{f}_θ . This approach computes a finite-difference gradient estimate for a smooth approximation to the original ψ , given by $\psi_{\sigma_1}(\mathbf{u}) := \mathbf{E}[\psi(\mathbf{u} + \sigma_1 Z_1)]$, where $Z_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_K)$. We provide error bounds in Appendix B by building on recent work by Duchi et al. (2015), and discuss asymptotic convergence of Algorithm 1 as $\sigma_1 \rightarrow 0$.

6. Experiments

We present experiments to show the proposed approach, Algorithm 1, is able to perform as well as methods that take

Table 1. Datasets used in our experiments.

Dataset	#instances	#features	Groups
Simulated	5000	2	-
COMPAS	4073	31	M/F
Adult	32561	122	M/F
Credit	30000	89	M/F
Business	11560	36	C/NC
KDD Cup 08	102294	117	-

Table 2. Test G-mean on sim. data. Lower is better.

	LogReg	PostShift	Proposed
Simulated	1.000	0.848	0.803

advantage of a metric’s form where available, and is also able to provide gains for metrics that are truly a black-box. We consider a simulated classification task, fair classification with noisy features, a ranking task and classification with proxy labels. The datasets we use are listed in Table 1.

We use the linear interpolation approach in Algorithm 3 for estimating gradients, as this is the most practical among the proposed estimation methods, and applicable when the surrogates and metrics are evaluated on different samples. We use linear models, and tune hyper-parameters such as step sizes and the perturbation parameter σ for gradient estimation using a held-out validation set. We run the projected gradient descent with 250 outer iterations and 1000 perturbations. For the projection step, we run 100 iterations of Adagrad. See Appendix D for more details and a discussion on perturbations. The code has been made available.

6.1. Optimizing G-mean on Simulated Data

We first apply our approach to maximize a non-black box evaluation metric: G-mean = $1 - \sqrt{\text{TPR} \times \text{TNR}}$, described in Example 1. We consider a simulated binary classification task in two dimensions, containing 10% positives and 90% negatives. The positive examples are drawn from a Gaussian with mean $[0, 0]$ and covariance matrix $0.2 \times \mathbf{I}_2$. The negative examples are drawn from a mixture of two Gaussians centered at $[-1, -1]$ and $[1, 1]$, with equal priors, and with a covariance matrix of $0.1 \times \mathbf{I}_2$.

We apply our method with two surrogate functions: the average hinge losses on the positive and negative examples. The results are shown in Table 2. We compare against two baselines: logistic regression that optimizes a standard cross-entropy loss, and a plug-in or post-shift approach that shifts the a threshold on the logistic regression model to optimize G-mean (Narasimhan et al., 2014). Because of the class imbalance, logistic regression learns to always predict the majority negative class and yields zero true positive rate and as a result a poor G-mean. Post-shift produces a better G-mean, but the proposed method performs the best. It is clear from the resulting decision boundaries shown in Figure 3

Table 3. Average test macro F-measure across groups with clean features. *Higher is better.* Despite having only black-box access to the metric, our approach performs comparable to methods that take advantage of the form of the metric.

	LogReg	PostShift	RelaxedFM	GenRates	Proposed
Business	0.793	0.789	0.794	0.793	0.796
COMPAS	0.560	0.631	0.614	0.620	0.629
Adult	0.668	0.664	0.665	0.654	0.665
Default	0.467	0.536	0.525	0.532	0.533




Figure 3. Hyperplanes learned by proposed method and PostShift on simulated data.

that the proposed method learns the better linear separator.

6.2. Macro F-measure with Noisy Features

For this experiment we consider training a classifier with fairness goals defined on binary protected attributes. We seek to maximize the average F-measure across the groups:

$$\text{Macro } F_1 = \frac{1}{2} \sum_{G \in \{0,1\}} \frac{2 \times \text{Precision}_G \times \text{Recall}_G}{\text{Precision}_G + \text{Recall}_G},$$

where Precision_G and Recall_G are the precision and recall on protected group G . Optimizing a sum of F-measures is harder than optimizing the binary F-measure because the summation destroys its pseudo-convexity property (Narasimhan et al., 2019).

We use four fairness datasets: (1) *COMPAS*, where the goal is to predict recidivism with *gender* as the protected attribute (Angwin et al., 2016); (2) *Adult*, where the goal is to predict if a person’s income is more than 50K/year, and we take *gender* as the protected group (Blake and Merz, 1998); (3) *Credit Default*, where the task is to predict whether a customer would default on his/her credit card payment, and we take *gender* as the protected group (Blake and Merz, 1998); (4) *Business Entity Resolution*, a proprietary dataset from a large internet services company, where the goal is to predict whether a pair of business descriptions refer to identical businesses, and we consider *non-chain* businesses as protected. In each case, we split the data into train-validation-test sets in the ratio 4/9 : 2/9 : 1/3.

Training with no noise. The first set of experiments tests if the proposed approach is able to match the performance of existing methods that are customized to optimize the macro F-measure. We compare against (i) plain logistic regression method, (ii) a plug-in or post-shift method that tunes a

threshold on the logistic regression model to maximize the F-measure (Koyejo et al., 2014; Narasimhan et al., 2014), (iii) an approach that optimizes a continuous relaxation to the F-measure that replaces the indicators with the hinge loss, and (iv) the recent “generalized rates” approach of Narasimhan et al. (2019) for optimizing metrics that are a sum of ratios. We apply our approach using four surrogate losses, each one is the hinge loss averaged over either the positive or negative examples, calculated separately for each of the two groups. As seen in Table 3, despite having only black-box access to the metric, the proposed approach performs comparable to the other methods that are directly tailored to optimize the macro F-measure.

Training with noisy features. The second set of experiments evaluates the performance of these methods when the training set has noisy features for just one of the groups, while the smaller validation set contains clean features. We use our approach to adaptively combine the same four surrogate losses computed on the noisy training set to best optimize the macro F-measure on the clean validation set.

We chose a certain fraction of the examples at random from one of the groups, which we refer to as group 0, and for these examples, we add Gaussian noise to the real features (with mean 0 and the same standard deviation as the feature), and flip the binary features with probability 0.9. Figure 4 shows the test F-measure for the different methods with varying fraction of noisy examples in group 0. Except for logistic regression, all other methods have access to the validation set: post-shift uses the validation set to tune a threshold on the logistic regression model; the RelaxedFM and GenRates method optimize their loss on the training set, but pick the best model iterate using the validation set. The proposed approach is able to make the best use of the validation set, and consistently performs the best across most noise levels.

6.3. Ranking to Optimize PRBEP

We next consider a ranking task, where the goal is to learn a scoring function f that maximizes the precision-recall break-even point (PRBEP), i.e. yields maximum precision at the threshold where precision and recall are equal. PRBEP is a special case of Precision@ K when K is set to the number of positive examples in the dataset. For this task, we experiment with the KDD Cup 2008 breast cancer detection data set (Rao et al., 2008) popularly used in this literature

Optimizing Black-box Metrics with Adaptive Surrogates

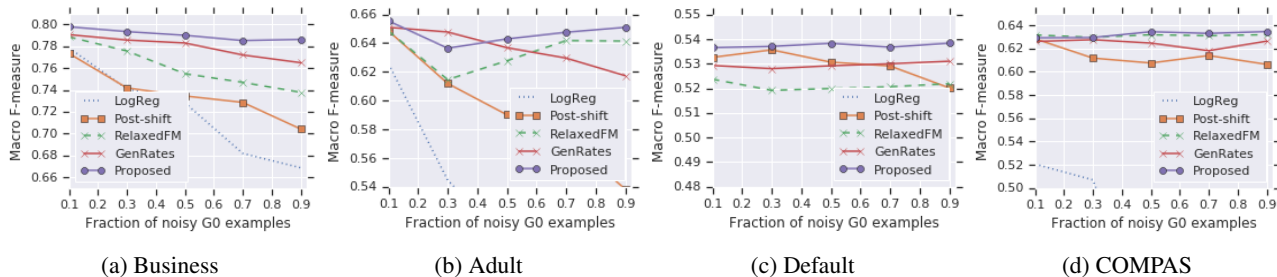


Figure 4. Test macro F-measure across groups for varying noise levels, averaged over 5 trials. Higher is better.

Table 4. Train and test PRBEP on KDD Cup 2008 data. Higher is better.

	LogReg	Kar et al. (2015)	Proposed
Train	0.480	0.473	0.546
Test	0.472	0.441	0.480

Table 5. Test classification error where the training labels are only proxy labels with unknown relationship to the true labels. The proposed method was run with both hinge and sigmoid surrogates. Lower is better.

	LogReg	PostShift	Hinge	Sigmoid
Adult	0.333	0.322	0.314	0.314
Business	0.340	0.251	0.256	0.236

(Kar et al., 2015; Mackey et al., 2018). We randomly split this dataset 60/20/20 for training, validation, and test.

Since the break-even point for a dataset is not known beforehand, we use surrogates that approximate precision at different recall thresholds τ . We use the quantile-based surrogate losses of Mackey et al. (2018) with $\tau = 0.25, 0.5, 0.75$. As a comparison, we optimize the avg-precision@ K surrogate provided by Kar et al. (2015). As seen in Table 4, the proposed approach is able to learn a better training loss by combining the three quantile surrogates, and yields the best PRBEP on the both the training and test sets.

6.4. Classification with Proxy Labels

Next, we consider classification tasks where the training labels are proxies for the true labels, but the validation data has the true labels. We seek to minimize the classification error on the validation set by combining hinge loss surrogates evaluated separately on the positive and negative training examples. While the theory requires convex losses, we experiment with also running the algorithm with non-convex sigmoid losses as surrogates.

For the Adult data, we predict whether a candidate’s gender is female, and take the marital-status-wife feature as the proxy label. For the Business Entity Resolution data, we predict whether a pair of business descriptions refer to the same business, and use the has-same-phone-number feature

as a proxy label.

We compare with a logistic regression model trained with the proxy labels and a post-shift method that corrects the logistic regression threshold to minimize classification error on the validation data. As expected logistic regression yields the highest test error. On Adult, both variants of the proposed method are better than PostShift. On Business, the proposed method performs slightly worse than PostShift when run with hinge surrogates, but yields notable improvements when run with sigmoid surrogates, which are tighter relaxations to the true errors.

7. Discussion

There is currently a lot of interest in training models with better alignment with evaluation metrics. Here, we have investigated a simple method that directly estimates only the needed gradients for gradient descent training, and does not require assuming a parametric form. This simplicity enabled us to provide rigorous theoretical guarantees.

Experimentally, our approach was as good as strategies that take advantage of a metric’s form (where available), gave notable gains over baselines for black-box ranking, and was significantly better than post-shifting for experiments with group-dependent noise. For the proxy label experiments, however, the results were mixed, with the proposed method requiring a tighter surrogate relaxation to perform better than post-shift. Post-shift is a strong baseline – in theory, for many metrics it is optimal to simply post-shift the Bayes class probability model $\mathbf{P}(y = 1|x)$ with a suitable threshold β (Koyejo et al., 2014; Yan et al., 2018). Post-shift only has one degree of freedom, which limits it, but also enables choosing β to directly optimize the true metric. In contrast, our method acts through surrogate losses to optimize the target metric. We argue that post-shift should be a required baseline for experiments on custom metric optimization.

We look forward to seeing further theoretical analysis for handling black-box metrics, and further experimentation comparing methods with fewer but smarter parameters to those with more flexible modeling.

References

- Alabi, D., Immorlica, N., and Kalai, A. (2018). Unleashing linear optimizers for group-fair learning and optimization. In *COLT*.
- Angwin, J., Larson, J., Mattu, S., and Kirchner, L. (2016). Machine bias. *ProPublica*, May, 23.
- Berahas, A. S., Cao, L., Choromanski, K., and Scheinberg, K. (2019). A theoretical and empirical comparison of gradient approximations in derivative-free optimization. *arXiv preprint arXiv:1905.01332*.
- Blake, C. and Merz, C. J. (1998). UCI repository of machine learning databases.
- Conn, A. R., Scheinberg, K., and Vicente, L. N. (2008). Geometry of interpolation sets in derivative free optimization. *Mathematical programming*, 111(1-2):141–172.
- Conn, A. R., Scheinberg, K., and Vicente, L. N. (2009). *Introduction to derivative-free optimization*, volume 8. Siam.
- Daskalaki, S., Kopanas, I., and Avouris, N. (2006). Evaluation of classifiers for an uneven class distribution problem. *Applied Artificial Intelligence*, 20:381–417.
- Duchi, J. C., Jordan, M. I., Wainwright, M. J., and Wibisono, A. (2015). Optimal rates for zero-order convex optimization: The power of two function evaluations. *IEEE Transactions on Information Theory*, 61(5):2788–2806.
- Eban, E., Schain, M., Mackey, A., Gordon, A., Saurous, R. A., and Elidan, G. (2017). Scalable learning of non-decomposable objectives. In *AISTATS*.
- Garmanjani, R. and Vicente, L. N. (2013). Smoothing and worst-case complexity for direct-search methods in nonsmooth optimization. *IMA Journal of Numerical Analysis*, 33(3):1008–1028.
- Ghadimi, S., Lan, G., and Zhang, H. (2016). Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. *Mathematical Programming*, 155(1):267–305.
- Grabocka, J., Scholz, R., and Schmidt-Thieme, L. (2019). Learning surrogate losses. *arXiv preprint arXiv:1905.10108*.
- Huang, C., Zhai, S., Talbott, W., Bautista, M. A., Sun, S.-Y., Guestrin, C., and Susskind, J. (2019). Addressing the loss-metric mismatch with adaptive loss alignment. In *ICML*.
- Jin, C., Netrapalli, P., Ge, R., Kakade, S. M., and Jordan, M. I. (2019). A short note on concentration inequalities for random vectors with subgaussian norm. *arXiv preprint arXiv:1902.03736*.
- Joachims, T. (2005). A support vector method for multivariate performance measures. In *ICML*.
- Joachims, T., Granka, L., Pan, B., Hembrooke, H., and Gay, G. (2005). Accurately interpreting clickthrough data as implicit feedback. *Proc. SIGIR*.
- Kar, P., Li, S., Narasimhan, H., Chawla, S., and Sebastiani, F. (2016). Online optimization methods for the quantification problem. In *KDD*.
- Kar, P., Narasimhan, H., and Jain, P. (2014). Online and stochastic gradient methods for non-decomposable loss functions. In *NIPS*.
- Kar, P., Narasimhan, H., and Jain, P. (2015). Surrogate functions for maximizing precision at the top. In *ICML*.
- Koyejo, O., Natarajan, N., Ravikumar, P., and Dhillon, I. (2014). Consistent binary classification with generalized performance metrics. In *NIPS*.
- Mackey, A., Luo, X., and Eban, E. (2018). Constrained classification and ranking via quantiles. *arXiv preprint arXiv:1803.00067*.
- Narasimhan, H. (2018). Learning with complex loss functions and constraints. In *AISTATS*.
- Narasimhan, H., Cotter, A., and Gupta, M. (2019). Optimizing generalized rate metrics through game equilibrium. In *NeurIPS*.
- Narasimhan, H., Kar, P., and Jain, P. (2015a). Optimizing non-decomposable performance measures: A tale of two classes. In *ICML*.
- Narasimhan, H., Ramaswamy, H., Saha, A., and Agarwal, S. (2015b). Consistent multiclass algorithms for complex performance measures. In *ICML*.
- Narasimhan, H., Vaish, R., and Agarwal, S. (2014). On the statistical consistency of plug-in classifiers for non-decomposable performance measures. In *NIPS*.
- Nesterov, Y. and Spokoiny, V. (2017). Random gradient-free minimization of convex functions. *Found. Comput. Math.*, 17(2):527–566.
- Parambath, S., Usunier, N., and Grandvalet, Y. (2014). Optimizing F-measures by cost-sensitive classification. In *NIPS*.

- Rao, R. B., Yakhnenko, O., and Krishnapuram, B. (2008). Kdd cup 2008 and the workshop on mining medical data. *ACM SIGKDD Explorations Newsletter*, 10(2):34–38.
- Ren, M., Zeng, W., Yang, B., and Urtasun, R. (2018). Learning to reweight examples for robust deep learning. In *ICML*.
- Tropp, J. A. (2015). An introduction to matrix concentration inequalities. *Found. Trends Mach. Learn.*, 8(1-2):1–230.
- Wainwright, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press.
- Wu, L., Tian, F., Xia, Y., Fan, Y., Qin, T., Jian-Huang, L., and Liu, T.-Y. (2018). Learning to teach with dynamic loss functions. In *NeurIPS*, pages 6466–6477.
- Yan, B., Koyejo, O., Zhong, K., and Ravikumar, P. (2018). Binary classification with karmic, threshold-quasi-concave metrics. In *ICML*.
- Ye, N., Chai, K., Lee, W., and Chieu, H. (2012). Optimizing F-measures: A tale of two approaches. In *ICML*.
- Zhao, S., Milani Fard, M., Narasimhan, H., and Gupta, M. R. (2019). Metric-optimized example weights. In *ICML*.

Notations. We use $[K]$ to denote $\{1, \dots, K\}$. We use $\|\cdot\|$ to denote the L_2 -norm. Unless specified otherwise, all smoothness and Lipschitz definitions are with respect to the L_2 -norm.

A. Proofs for Lemmas and Theorems

A.1. Proof of Observation 1

Proof. To see that the vector $[u_1, \dots, u_K]$ belongs to a convex set, since by assumption $\{\ell_i\}_{i=1}^K$ are convex functions, therefore the set of constraints $\ell_i(\theta) \leq u_i$ defines a convex set in $[\theta, u_1, \dots, u_K]$ as intersection of sublevel sets of convex functions are convex. \square

A.2. Proof of Lemma 1

Lemma 1 (Restated). *Let \mathbf{u}^+ be the exact projection of $\tilde{\mathbf{u}}^{t+1} \in \mathbb{R}_+^K$ onto \mathcal{U} . For any solution \mathbf{u}^{t+1} to (4), we have $\mathbf{u}^{t+1} \in \mathcal{U}$, $\mathbf{u}^{t+1} \leq \mathbf{u}^+$, and for a monotonic ψ , $\psi(\mathbf{u}^{t+1}) \leq \psi(\mathbf{u}^+)$.*

We first show how one can compute an exact projection onto \mathcal{U} , and show that the projection described in Lemma 1 implements this approximately.

Lemma 5 (Exact projection). *The projection \mathbf{u}^+ of $\tilde{\mathbf{u}}^{t+1} \in \mathbb{R}_+^K$ onto \mathcal{U} is given by:*

- (i) $\theta^{t+1} \in \operatorname{argmin}_{\theta \in \mathbb{R}^d} \frac{1}{2} \|(\ell(\theta) - \tilde{\mathbf{u}}^{t+1})_+\|^2$; $\mathbf{u}^{t+1} = \ell(\theta^{t+1})$
- (ii) $u_k^+ = \max\{\tilde{u}_k, u_k^{t+1}\}$, $\forall k \in [K]$,

where $(z)_+ = \max\{0, z\}$, applied element-wise.

Proof. It is easy to see that step (i) is a convex problem because each ℓ_k is convex in θ , and both $(\cdot)_+$ and $\|\cdot\|^2$ are convex and monotonic in their arguments, making the composition $\|(\ell(\theta) - \tilde{\mathbf{u}})_+\|^2$ also convex in θ .

To perform the projection, since step (i) above is a convex problem, the optimality condition gives

$$\sum_{i=1}^K (\ell_i(\theta^+) - \tilde{u}_i)_+ \cdot \mathbb{1}\{\ell_i(\theta^+) - \tilde{u}_i > 0\} \cdot \nabla_{\theta} \ell_i(\theta^+) = \mathbf{0}_d$$

which is the same as

$$\sum_{i=1}^K (u_i^+ - \tilde{u}_i) \cdot \nabla_{\theta} \ell_i(\theta^+) = \mathbf{0}_d \quad (5)$$

by the second step of the procedure. We shall use (5) to show that u_i^+ is the projection in the \mathcal{U} -space.

The projection in the \mathcal{U} -space can equivalently be written as the following convex problem

$$\begin{aligned} & \underset{u_1, \dots, u_K, \theta}{\text{minimize}} && \frac{1}{2} \sum_{i=1}^K (u_i - \tilde{u}_i)^2 \\ & \text{subject to} && u_i - \ell_i(\theta) \geq 0 \quad \forall i \in [K]. \end{aligned}$$

Introducing the dual variable $\lambda \in \mathbb{R}^K$ and the KKT condition of the problem becomes

$$\begin{aligned} \sum_{i=1}^K (u_i - \tilde{u}_i) - \sum_{i=1}^K \lambda_i &= 0 & \sum_{i=1}^K \lambda_i \cdot \nabla_{\theta} \ell_i(\theta) &= 0 \\ u_i - \ell_i(\theta) \geq 0 & \quad \lambda_i \geq 0 & \lambda_i (u_i - \ell_i(\theta)) &= 0 \quad \forall i \in [K] \end{aligned}$$

if $(u_1, \dots, u_K, \theta)$ and λ are optimal.

Taking $\lambda_i = u_i^+ - \tilde{u}_i$ and $\theta = \theta^+$ with $u_i = u_i^+$, one can easily verify using (5) that all the conditions hold. Since the optimization problem satisfies Slater's constraint qualification and therefore we can conclude that the primal optimal solution is \mathbf{u}^+ , as defined in the lemma statement. \square

We go on to prove Lemma 1.

Proof of Lemma 1. Because \mathbf{u}^{t+1} is the surrogate loss at θ^{t+1} , it clearly lies in \mathcal{L} and hence in the superset $\mathcal{U} \supseteq \mathcal{L}$. Next, notice that the over-constrained projection \mathbf{u}^{t+1} in Lemma 1 is the same as step (i) in the exact projection in Lemma 1, with step (ii) giving us that the exact projection $u_k^+ = \max\{\tilde{u}_k, u_k^{t+1}\}$, $\forall k \in [K]$. It follows that: $u_k^{t+1} \leq u_k^+$, $\forall k \in [K]$. So for a monotonic ψ , we have $\psi(\mathbf{u}^{t+1}) \leq \psi(\mathbf{u}^+)$. \square

A.3. Proof of Theorem 2

Theorem 2 (Restated). Let $M(\theta) = \psi(\ell(\theta)) + \epsilon(\theta)$, for a ψ that is monotonic, β -smooth and L -Lipschitz, and the worst-case slack $\max_{\theta \in \mathbb{R}^d} |\epsilon(\theta)|$ is the minimum among all such decompositions of M .

Suppose each ℓ_k is γ -smooth and Φ -Lipschitz in θ with $\|\ell(\theta)\| \leq G$, $\forall \theta$. Suppose the gradient estimates $\hat{\mathbf{g}}^t$ satisfy $\mathbf{E} [\|\hat{\mathbf{g}}^t - \nabla \psi(\ell(\theta^t))\|^2] \leq \kappa_\epsilon$, $\forall t \in [T]$ and the projection step satisfies $\|(\ell(\theta^{t+1}) - \tilde{\mathbf{u}}^t)_+\|^2 \leq \min_{\theta \in \mathbb{R}^d} \|\ell(\theta) - \tilde{\mathbf{u}}^t\|_+^2 + \mathcal{O}(\frac{1}{\beta^2 T})$, $\forall t \in [T]$. Set stepsize $\eta = \frac{1}{\beta^2}$.

Then Algorithm 1 converges to an approximate stationary point of $\psi(\ell(\cdot))$:

$$\min_{1 \leq t \leq T} \mathbf{E} [\|\nabla \psi(\ell(\theta^t))\|^2] \leq C \left(\frac{\beta}{\sqrt{T}} + \sqrt{\kappa_\epsilon} + \sqrt{L\kappa_\epsilon^{1/4}} \right),$$

where the expectation is over the randomness in the gradient estimates, and $C = \mathcal{O}(KL(\gamma(G + \frac{L}{\beta^2}) + \Phi^2))$.

While the above theorem prescribes a specific learning rate η for the projected gradient descent, in our experiments, we tune η using a held-out validation set.

The proof proceeds in two parts. In Section A.3.1, we first show that the algorithm converges to an approximate stationary point of ψ over \mathcal{U} . In Section A.3.2, we then translate this a guarantee in θ , i.e. we show that the algorithm converges to an approximate stationary point of $\psi(\ell(\cdot))$ over θ .

A.3.1. CONVERGENCE IN \mathcal{U} -SPACE

Lemma 6. Define the gradient mapping at $\mathbf{u} \in \mathcal{U}$ for a vector $g \in \mathbb{R}^K$ as $P(\mathbf{u}, g) := \frac{1}{\eta}(\mathbf{u} - \Pi_{\mathcal{U}}(\mathbf{u} - \eta \cdot g))$, where $\Pi_{\mathcal{U}}(z)$ denotes the projection of z onto \mathcal{U} . Then under the assumptions of Theorem 2,

$$\min_{1 \leq t \leq T} \mathbf{E} [\|P(\mathbf{u}^t, \nabla \psi(\mathbf{u}^t))\|^2] \leq \mathcal{O} \left(\frac{\beta^2}{T} + \kappa_\epsilon + L\sqrt{\kappa_\epsilon} \right).$$

Before we prove this result, we will find it useful to state the following lemma.

Lemma 7 (Properties of inexact projection). Fix $\mathbf{u} \in \mathcal{U}$ where \mathcal{U} is a convex set and arbitrary vectors $g_1, g_2 \in \mathbb{R}^K$. Let

$$\theta_1^+ \in \operatorname{argmin}_{\theta \in \mathbb{R}^d} \frac{1}{2} \|(\ell(\theta) - (\mathbf{u} - \eta g_1))_+\|^2 \quad \text{and} \quad \theta_2^+ \in \operatorname{argmin}_{\theta \in \mathbb{R}^d} \frac{1}{2} \|(\ell(\theta) - (\mathbf{u} - \eta g_2))_+\|^2,$$

and let $\mathbf{u}_1^+ = \max\{\ell(\theta_1^+), \mathbf{u} - \eta g_1\}$ and $\mathbf{u}_2^+ = \max\{\ell(\theta_2^+), \mathbf{u} - \eta g_2\}$. Define the gradient mapping $P(\mathbf{u}, g_1) := \frac{1}{\eta}(\mathbf{u} - \mathbf{u}_1^+)$ and $P(\mathbf{u}, g_2) := \frac{1}{\eta}(\mathbf{u} - \mathbf{u}_2^+)$. Denote $\tilde{\theta}_1^+, \tilde{\theta}_2^+$ as approximate minimizers such that

$$\frac{1}{2} \|(\ell(\tilde{\theta}_1^+) - (\mathbf{u} - \eta g_1))_+\|^2 \leq \frac{1}{2} \|(\ell(\theta_1^+) - (\mathbf{u} - \eta g_1))_+\|^2 + \alpha \quad (6)$$

and

$$\frac{1}{2} \|(\ell(\tilde{\theta}_2^+) - (\mathbf{u} - \eta g_2))_+\|^2 \leq \frac{1}{2} \|(\ell(\theta_2^+) - (\mathbf{u} - \eta g_2))_+\|^2 + \alpha, \quad (7)$$

and let $\tilde{\mathbf{u}}_1^+ = \max\{\ell(\tilde{\theta}_1^+), \mathbf{u} - \eta g_1\}$ and $\tilde{\mathbf{u}}_2^+ = \max\{\ell(\tilde{\theta}_2^+), \mathbf{u} - \eta g_2\}$. Define the corresponding gradient mapping $\tilde{P}(\mathbf{u}, g_1) := \frac{1}{\eta}(\mathbf{u} - \tilde{\mathbf{u}}_1^+)$ and $\tilde{P}(\mathbf{u}, g_2) := \frac{1}{\eta}(\mathbf{u} - \tilde{\mathbf{u}}_2^+)$. Then the following holds:

1. $\|\tilde{P}(\mathbf{u}, g_1) - P(\mathbf{u}, g_1)\| \leq \frac{\sqrt{2\alpha}}{\eta}$.

2. $\langle g_1, \tilde{P}(\mathbf{u}, g_1) \rangle \geq \frac{3}{4} \|\tilde{P}(\mathbf{u}, g_1)\|^2 - \frac{2\alpha}{\eta^2}$.
3. $\|\tilde{P}(\mathbf{u}, g_1)\| \leq \|g_1\| + \frac{\sqrt{2\alpha}}{\eta}$.
4. $\|P(\mathbf{u}, g_1) - P(\mathbf{u}, g_2)\| \leq \|g_1 - g_2\|$.
5. $\|\tilde{P}(\mathbf{u}, g_1) - \tilde{P}(\mathbf{u}, g_2)\| \leq \|g_1 - g_2\| + 2\frac{\sqrt{2\alpha}}{\eta}$.

Proof. We have:

$$\begin{aligned}
 \frac{1}{2} \|\tilde{\mathbf{u}}_1^+ - (\mathbf{u} - \eta g_1)\|^2 &= \frac{1}{2} \|\max\{\mathbf{u} - \eta g_1, \ell(\tilde{\theta}_1^+)\} - (\mathbf{u} - \eta g_1)\|^2 \\
 &= \frac{1}{2} \|(\ell(\tilde{\theta}_1^+) - (\mathbf{u} - \eta g_1))_+\|^2 \\
 &\leq \frac{1}{2} \|(\ell(\theta_1^+) - (\mathbf{u} - \eta g_1))_+\|^2 + \alpha \quad (\text{Assumption (6)}) \\
 &= \frac{1}{2} \|\max\{\mathbf{u} - \eta g_1, \ell(\theta_1^+)\} - (\mathbf{u} - \eta g_1)\|^2 + \alpha \\
 &= \frac{1}{2} \|\mathbf{u}_1^+ - (\mathbf{u} - \eta g_1)\|^2 + \alpha,
 \end{aligned}$$

which implies that

$$g_1^\top \tilde{\mathbf{u}}_1^+ + \frac{1}{2\eta} \|\tilde{\mathbf{u}}_1^+ - \mathbf{u}\|^2 - g_1^\top \mathbf{u}_1^+ - \frac{1}{2\eta} \|\mathbf{u}_1^+ - \mathbf{u}\|^2 \leq \frac{\alpha}{\eta}. \quad (8)$$

Part (1) now follows from

$$\begin{aligned}
 \|\tilde{P}(\mathbf{u}, g_1) - P(\mathbf{u}, g_1)\| &= \frac{1}{\eta} \|\mathbf{u}_1^+ - \tilde{\mathbf{u}}_1^+\| \\
 &\leq \frac{\sqrt{2\eta}}{\eta} \sqrt{F_{g_1}(\tilde{\mathbf{u}}_1^+) - F_{g_1}(\mathbf{u}_1^+) - \nabla F_{g_1}(\mathbf{u}_1^+)^\top (\tilde{\mathbf{u}}_1^+ - \mathbf{u}_1^+)} \\
 &\leq \frac{\sqrt{2\eta}}{\eta} \sqrt{\frac{\alpha}{\eta}} \leq \frac{\sqrt{2\alpha}}{\eta},
 \end{aligned}$$

where we used $\frac{1}{\eta}$ -strong convexity of the objective $F_{g_1}(\mathbf{z}) := g_1^\top \mathbf{z} + \frac{1}{2\eta} \|\mathbf{z} - \mathbf{u}\|^2$ for $\mathbf{z}, \mathbf{u} \in \mathcal{U}$ and the fact that \mathbf{u}_1^+ is the exact minimizer over the convex set \mathcal{U} , implying $\nabla F_{g_1}(\mathbf{u}_1^+)^\top (\mathbf{z} - \mathbf{u}_1^+) \geq 0 \forall \mathbf{z} \in \mathcal{U}$.

For part (4), since \mathbf{u}_1^+ and \mathbf{u}_2^+ are optimal points of function $F_{g_1}(\cdot)$ and $F_{g_2}(\cdot)$ over convex set \mathcal{U} respectively, from optimality condition we have

$$\left(g_1 + \frac{1}{\eta}(\mathbf{u}_1^+ - \mathbf{u})\right)^\top (\mathbf{z} - \mathbf{u}_1^+) \geq 0 \quad \text{and} \quad \left(g_2 + \frac{1}{\eta}(\mathbf{u}_2^+ - \mathbf{u})\right)^\top (\mathbf{z} - \mathbf{u}_2^+) \geq 0 \quad \text{for all } \mathbf{z} \in \mathcal{U}. \quad (9)$$

Setting $\mathbf{z} = \mathbf{u}_2^+$ in the first and $\mathbf{z} = \mathbf{u}_1^+$ in the second equation and summing up we have

$$(g_1 - g_2)^\top (\mathbf{u}_2^+ - \mathbf{u}_1^+) \geq \frac{1}{\eta} \|\mathbf{u}_2^+ - \mathbf{u}_1^+\|^2.$$

Therefore using Cauchy-Schwarz

$$\|P(\mathbf{u}, g_1) - P(\mathbf{u}, g_2)\| = \frac{1}{\eta} \|\mathbf{u}_2^+ - \mathbf{u}_1^+\| \leq \|g_1 - g_2\|.$$

Part (5) now follows immediately from part (1) and (4) by

$$\begin{aligned}
 \|\tilde{P}(\mathbf{u}, g_1) - \tilde{P}(\mathbf{u}, g_2)\| &\leq \|P(\mathbf{u}, g_1) - P(\mathbf{u}, g_2)\| + \|\tilde{P}(\mathbf{u}, g_1) - P(\mathbf{u}, g_1) + P(\mathbf{u}, g_2) - \tilde{P}(\mathbf{u}, g_2)\| \\
 &\leq \|g_1 - g_2\| + 2\|\tilde{P}(\mathbf{u}, g_1) - P(\mathbf{u}, g_1)\| \\
 &\leq \|g_1 - g_2\| + \frac{2\sqrt{2\alpha}}{\eta}.
 \end{aligned}$$

To see part (2), we plug in $\mathbf{z} = \mathbf{u}$ in the first equation of display (9), giving $g_1^\top(\mathbf{u} - \mathbf{u}_1^+) \geq \frac{1}{\eta} \|\mathbf{u} - \mathbf{u}_1^+\|^2$. Moreover from equation (8) we know

$$g_1^\top(\mathbf{u}_1^+ - \tilde{\mathbf{u}}_1^+) \geq -\frac{\alpha}{\eta} + \frac{1}{2\eta} \|\tilde{\mathbf{u}}_1^+ - \mathbf{u}\|^2 - \frac{1}{2\eta} \|\mathbf{u}_1^+ - \mathbf{u}\|^2.$$

Consequently,

$$g_1^\top(\mathbf{u} - \tilde{\mathbf{u}}_1^+) = g_1^\top(\mathbf{u} - \mathbf{u}_1^+) + g_1^\top(\mathbf{u}_1^+ - \tilde{\mathbf{u}}_1^+) \geq \frac{1}{\eta} \|\mathbf{u} - \mathbf{u}_1^+\|^2 - \frac{\alpha}{\eta} + \frac{1}{2\eta} \|\tilde{\mathbf{u}}_1^+ - \mathbf{u}\|^2 - \frac{1}{2\eta} \|\mathbf{u}_1^+ - \mathbf{u}\|^2.$$

Now to relate $\|\mathbf{u} - \mathbf{u}_1^+\|$ to $\|\mathbf{u}_1^+ - \tilde{\mathbf{u}}_1^+\|$, we have

$$\begin{aligned} \frac{1}{2\eta} \|\mathbf{u} - \tilde{\mathbf{u}}_1^+\|^2 &\leq \frac{1}{\eta} \|\mathbf{u} - \mathbf{u}_1^+\|^2 + \frac{1}{\eta} \|\mathbf{u}_1^+ - \tilde{\mathbf{u}}_1^+\|^2 \\ &\leq \frac{1}{\eta} \|\mathbf{u} - \mathbf{u}_1^+\|^2 + 2[F_{g_1}(\tilde{\mathbf{u}}_1^+) - F_{g_1}(\mathbf{u}_1^+) - \nabla F_{g_1}(\mathbf{u}_1^+)^\top(\tilde{\mathbf{u}}_1^+ - \mathbf{u}_1^+)] \\ &\leq \frac{1}{\eta} \|\mathbf{u} - \mathbf{u}_1^+\|^2 + \frac{2\alpha}{\eta}. \end{aligned}$$

Putting things together $g_1^\top \tilde{P}(\mathbf{u}, g_1) = \frac{1}{\eta} g_1^\top(\mathbf{u} - \tilde{\mathbf{u}}_1^+) \geq \frac{3}{4} \|\tilde{P}(\mathbf{u}, g_1)\|^2 - \frac{2\alpha}{\eta^2}$, as claimed.

Finally, for part (3) since $\|g_1\| \cdot \|\mathbf{u} - \mathbf{u}_1^+\| \geq g_1^\top(\mathbf{u} - \mathbf{u}_1^+) \geq \frac{1}{\eta} \|\mathbf{u} - \mathbf{u}_1^+\|^2$ and using part (2),

$$\begin{aligned} \|\tilde{P}(\mathbf{u}, g_1)\| &= \frac{1}{\eta} \|\mathbf{u} - \tilde{\mathbf{u}}_1^+\| \leq \frac{1}{\eta} \|\mathbf{u} - \mathbf{u}_1^+\| + \frac{1}{\eta} \|\mathbf{u}_1^+ - \tilde{\mathbf{u}}_1^+\| \\ &\leq \|g_1\| + \frac{\sqrt{2\alpha}}{\eta}, \end{aligned}$$

where we used part (1) for the last step. This concludes the proof of the lemma. \square

Equipped with the above results, we move on to prove Lemma 6, i.e. to show that the algorithm converges to an approximate stationary point of ψ over \mathcal{U} .

Proof of Lemma 6. We will assume that the gradient estimates $\hat{\mathbf{g}}^t$ satisfy $\mathbf{E} [\|\hat{\mathbf{g}}^t - \nabla\psi(\ell(\theta^t))\|^2] \leq \kappa_\epsilon$, $\forall t \in [T]$ and the projection step satisfies $\frac{1}{2} \|(\ell(\theta^{t+1}) - \tilde{\mathbf{u}}^t)_+\|^2 \leq \min_{\theta \in \mathbb{R}^d} \frac{1}{2} \|(\ell(\theta) - \tilde{\mathbf{u}}^t)_+\|^2 + \alpha$, $\forall t \in [T]$.

Let $\mathbf{u}^{t+1} = \ell(\theta^{t+1})$ and $\tilde{\mathbf{u}}^{t+1} = \max\{\mathbf{u}^{t+1}, \mathbf{u}^t - \eta\hat{\mathbf{g}}^t\}$ be the next iterate had we executed step (ii) of the projection given Lemma 1. Define $\delta^t := \hat{\mathbf{g}}^t - \nabla\psi(\mathbf{u}^t)$. For any $g \in \mathbb{R}^K$, let the gradient mapping $P(\mathbf{u}, g)$ and approximate gradient mapping $\tilde{P}(\mathbf{u}, g)$ be defined as in Lemma 7. Note that $\tilde{\mathbf{u}}^{t+1} = \mathbf{u}^t - \eta\tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t)$.

$$\begin{aligned} \psi(\mathbf{u}^{t+1}) &\leq \psi(\tilde{\mathbf{u}}^{t+1}) \quad (\text{from monotonicity of } \psi) \\ &\leq \psi(\mathbf{u}^t) - \eta \langle \nabla\psi(\mathbf{u}^t), \tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t) \rangle + \frac{\beta^2}{2} \eta^2 \|\tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2 \quad (\text{using smoothness of } \psi) \\ &= \psi(\mathbf{u}^t) - \eta \langle \hat{\mathbf{g}}^t, \tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t) \rangle + \eta \langle \hat{\mathbf{g}}^t - \nabla\psi(\mathbf{u}^t), \tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t) \rangle + \frac{\beta^2}{2} \eta^2 \|\tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2 \\ &= \psi(\mathbf{u}^t) - \eta \langle \hat{\mathbf{g}}^t, \tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t) \rangle + \eta \langle \delta^t, \tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t) \rangle + \frac{\beta^2}{2} \eta^2 \|\tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2 \\ &\leq \psi(\mathbf{u}^t) - \left(\frac{3}{4} \eta - \frac{\beta^2}{2} \eta^2 \right) \|\tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2 + \eta \langle \delta^t, \tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t) \rangle + \frac{2\alpha}{\eta} \quad (\text{from Lemma 7, statement 2}) \\ &= \psi(\mathbf{u}^t) - \left(\frac{3}{4} \eta - \frac{\beta^2}{2} \eta^2 \right) \|\tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2 + \eta \langle \delta^t, \tilde{P}(\mathbf{u}^t, \nabla\psi(\mathbf{u}^t)) \rangle + \eta \langle \delta^t, \tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t) - \tilde{P}(\mathbf{u}^t, \nabla\psi(\mathbf{u}^t)) \rangle + \frac{2\alpha}{\eta} \\ &\leq \psi(\mathbf{u}^t) - \left(\frac{3}{4} \eta - \frac{\beta^2}{2} \eta^2 \right) \|\tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2 + \eta \langle \delta^t, \tilde{P}(\mathbf{u}^t, \nabla\psi(\mathbf{u}^t)) \rangle + \eta \|\delta^t\|^2 + 2\sqrt{2\alpha} \|\delta^t\| + \frac{2\alpha}{\eta} \\ &\leq \psi(\mathbf{u}^t) - \left(\frac{3}{4} \eta - \frac{\beta^2}{2} \eta^2 \right) \|\tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2 + \eta \|\delta^t\| \left(\|\nabla\psi(\mathbf{u}^t)\| + \frac{\sqrt{2\alpha}}{\eta} \right) + \eta \|\delta^t\|^2 + 2\sqrt{2\alpha} \|\delta^t\| + \frac{2\alpha}{\eta} \\ &\leq \psi(\mathbf{u}^t) - \left(\frac{3}{4} \eta - \frac{\beta^2}{2} \eta^2 \right) \|\tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2 + (\eta L + \sqrt{2\alpha}) \|\delta^t\| + \eta \|\delta^t\|^2 + 2\sqrt{2\alpha} \|\delta^t\| + \frac{2\alpha}{\eta}, \end{aligned}$$

where the third-last inequality uses Lemma 7, statement 5 together with Cauchy-Schwarz and the second-last inequality uses Lemma 7, statement 3, and the fact that ψ is L -Lipschitz. Summing up over $t = 1, \dots, T$,

$$\left(\frac{3}{4}\eta - \frac{\beta^2}{2}\eta^2\right) \sum_{t=1}^T \|\tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2 \leq \psi(\mathbf{u}^1) - \psi(\mathbf{u}^{T+1}) + \sum_{t=1}^T \left((\eta L + 3\sqrt{2\alpha})\|\delta^t\| + \eta\|\delta^t\|^2 + \frac{2\alpha}{\eta} \right).$$

Taking expectations on both sides and using the assumption $0 \leq \psi(\mathbf{u}) \leq 1 \forall \mathbf{u} \in \mathcal{U}$,

$$\begin{aligned} \left(\frac{3}{4}\eta - \frac{\beta^2}{2}\eta^2\right) \sum_{t=1}^T \mathbf{E} \left[\|\tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2 \right] &\leq 1 + \sum_{t=1}^T \left((\eta L + 3\sqrt{2\alpha})\mathbf{E} [\|\delta^t\|] + \eta\mathbf{E} [\|\delta^t\|^2] + \frac{2\alpha}{\eta} \right) \\ &\leq 1 + \sum_{t=1}^T \left((\eta L + 3\sqrt{2\alpha})\sqrt{\mathbf{E} [\|\delta^t\|^2]} + \eta\mathbf{E} [\|\delta^t\|^2] + \frac{2\alpha}{\eta} \right) \\ &\leq 1 + T \left((\eta L + 3\sqrt{2\alpha})\sqrt{\kappa_\epsilon} + \eta\kappa_\epsilon + \frac{2\alpha}{\eta} \right), \end{aligned}$$

where we used the assumption on the gradient estimate error $\mathbf{E} [\|\delta^t\|^2]$ in the last step. Rearranging we have

$$\frac{1}{T} \sum_{t=1}^T \mathbf{E} \left[\|\tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2 \right] \leq \frac{1/T + (\eta L + 3\sqrt{2\alpha})\sqrt{\kappa_\epsilon} + \eta\kappa_\epsilon + \frac{2\alpha}{\eta}}{\frac{3}{4}\eta - \frac{\beta^2}{2}\eta^2}.$$

Using Lemma 7, statement 1,

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbf{E} [\|P(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2] &\leq \frac{2}{T} \sum_{t=1}^T \mathbf{E} [\|\tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2] + \frac{2}{T} \sum_{t=1}^T \mathbf{E} [\|\tilde{P}(\mathbf{u}^t, \hat{\mathbf{g}}^t) - P(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2] \\ &\leq \frac{2/T + 2(\eta L + 3\sqrt{2\alpha})\sqrt{\kappa_\epsilon} + 2\eta\kappa_\epsilon + \frac{4\alpha}{\eta}}{\frac{3}{4}\eta - \frac{\beta^2}{2}\eta^2} + \frac{4\alpha}{\eta^2}. \end{aligned}$$

Setting stepsize $\eta = \frac{1}{\beta^2}$:

$$\frac{1}{T} \sum_{t=1}^T \mathbf{E} [\|P(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2] \leq \frac{8\beta^2}{T} + 8L\sqrt{\kappa_\epsilon} + 8\kappa_\epsilon + 24\beta^2\sqrt{2\alpha\kappa_\epsilon} + 20\alpha\beta^4.$$

We can now bound the average gradient map norm across iterations:

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbf{E} [\|P(\mathbf{u}^t, \nabla\psi(\mathbf{u}^t))\|^2] &\leq \frac{2}{T} \sum_{t=1}^T \mathbf{E} [\|P(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2] + \frac{2}{T} \sum_{t=1}^T \mathbf{E} [\|P(\mathbf{u}^t, \nabla\psi(\mathbf{u}^t)) - P(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2] \\ &\leq \frac{2}{T} \sum_{t=1}^T \mathbf{E} [\|P(\mathbf{u}^t, \hat{\mathbf{g}}^t)\|^2] + \frac{2}{T} \sum_{t=1}^T \mathbf{E} [\|\nabla\psi(\mathbf{u}^t) - \hat{\mathbf{g}}^t\|^2] \\ &\leq \frac{16\beta^2}{T} + 16L\sqrt{\kappa_\epsilon} + 16\kappa_\epsilon + 48\beta^2\sqrt{2\alpha\kappa_\epsilon} + 40\alpha\beta^4 + 2\kappa_\epsilon \end{aligned}$$

where we used Lemma 7, statement 4 for the second inequality and the assumption on the gradient estimation error for the last inequality. Thus:

$$\min_{1 \leq t \leq T} \mathbf{E} [\|P(\mathbf{u}^t, \nabla\psi(\mathbf{u}^t))\|^2] \leq \frac{16\beta^2}{T} + 16L\sqrt{\kappa_\epsilon} + 18\kappa_\epsilon + 48\beta^2\sqrt{2\alpha\kappa_\epsilon} + 40\alpha\beta^4.$$

Now picking $\alpha = \frac{1}{\beta^2 T}$ completes the proof. \square

A.3.2. CONVERGENCE IN θ -SPACE

We are now ready to prove Theorem 2. We translate the near-stationarity result in Lemma from \mathbf{u} -space to θ -space.

Proof of Theorem 2. For a given T , let $t^* \in \operatorname{argmin}_{1 \leq t \leq T} \|P(\mathbf{u}^t, \nabla\psi(\mathbf{u}^t))\|^2$. Pick iterates θ^{t^*} and θ^{t^*+1} of Algorithm 1. The corresponding iterates in the \mathcal{U} -space are $\mathbf{u}^{t^*} = \ell(\theta^{t^*})$ and $\mathbf{u}^{t^*+1} = \ell(\theta^{t^*+1})$.

Further, let $\tilde{\mathbf{u}}^{t^*+1} = \mathbf{u}^{t^*} - \eta \nabla\psi(\mathbf{u}^{t^*})$ be the un-projected next iterate, and $\hat{\mathbf{u}}^{t^*+1} = \mathbf{u}^{t^*} - \eta \cdot P(\mathbf{u}^{t^*}, \nabla\psi(\mathbf{u}^{t^*}))$ be the one obtained after an exact projection, both using exact gradient $\nabla\psi(\mathbf{u}^{t^*})$.

We start with the assumption that (as promised by Lemma 6):

$$\mathbf{E}[\|P(\mathbf{u}^{t^*}, \nabla\psi(\mathbf{u}^{t^*}))\|^2] = \frac{1}{\eta^2} \mathbf{E}[\|\mathbf{u}^{t^*} - \eta \cdot P(\mathbf{u}^{t^*}, \nabla\psi(\mathbf{u}^{t^*})) - \mathbf{u}^{t^*}\|^2] = \frac{1}{\eta^2} \mathbf{E}[\|\hat{\mathbf{u}}^{t^*+1} - \mathbf{u}^{t^*}\|^2] \leq \epsilon^2$$

or equivalently,

$$\mathbf{E}[\|\hat{\mathbf{u}}^{t^*+1} - \mathbf{u}^{t^*}\|^2] \leq \eta^2 \epsilon^2 \quad (10)$$

and would like to bound the gradient norm of $\psi(\ell(\cdot))$ at θ^{t^*} .

We start by translating (10) to a guarantee in the θ -space. We know that

$$\hat{\mathbf{u}}^{t^*+1} \in \arg \min_{\mathbf{u} \in \mathcal{U}} \|\mathbf{u} - \tilde{\mathbf{u}}^{t^*+1}\|^2. \quad (11)$$

Put together (10) and (11), and take expectation over randomness in \mathbf{u}^{t^*} ,

$$\begin{aligned} \mathbf{E}[\|\mathbf{u}^{t^*} - \tilde{\mathbf{u}}^{t^*+1}\|^2] &\leq \mathbf{E}[\|\hat{\mathbf{u}}^{t^*+1} - \tilde{\mathbf{u}}^{t^*+1}\|^2] + \mathbf{E}[\|\hat{\mathbf{u}}^{t^*+1} - \mathbf{u}^{t^*}\|^2] + 2\mathbf{E}[\|\hat{\mathbf{u}}^{t^*+1} - \tilde{\mathbf{u}}^{t^*+1}\| \|\hat{\mathbf{u}}^{t^*+1} - \mathbf{u}^{t^*}\|] \\ &\leq \mathbf{E}[\|\hat{\mathbf{u}}^{t^*+1} - \tilde{\mathbf{u}}^{t^*+1}\|^2] + \eta^2 \epsilon^2 + 2\eta\epsilon \sqrt{\mathbf{E}[\|\hat{\mathbf{u}}^{t^*+1} - \tilde{\mathbf{u}}^{t^*+1}\|^2]} \\ &\leq \mathbf{E}[\|\hat{\mathbf{u}}^{t^*+1} - \tilde{\mathbf{u}}^{t^*+1}\|^2] + \eta^2 \epsilon^2 + 2\eta\epsilon \sqrt{\mathbf{E}[\|\mathbf{u}^{t^*} - \tilde{\mathbf{u}}^{t^*+1}\|^2]} \\ &= \mathbf{E}[\|\hat{\mathbf{u}}^{t^*+1} - \tilde{\mathbf{u}}^{t^*+1}\|^2] + \eta^2 \epsilon^2 + 2\eta^2 \epsilon \sqrt{\mathbf{E}[\|\nabla\psi(\mathbf{u}^{t^*})\|^2]}, \end{aligned}$$

where we used Cauchy-Schwarz for the second step. Using the fact that ψ is L -Lipschitz:

$$\mathbf{E}[\|\mathbf{u}^{t^*} - \tilde{\mathbf{u}}^{t^*+1}\|^2] \leq \mathbf{E}[\|\hat{\mathbf{u}}^{t^*+1} - \tilde{\mathbf{u}}^{t^*+1}\|^2] + \epsilon', \quad (12)$$

where $\epsilon' = \eta^2(\epsilon^2 + 2L\epsilon)$.

We also know that $\hat{\mathbf{u}}^{t^*+1}$ can be equivalently obtained by performing an optimization in the θ -space as follows:

$$\hat{\theta}^{t^*+1} \in \arg \min_{\theta \in \mathbb{R}^d} \|\max\{\ell(\theta), \tilde{\mathbf{u}}^{t^*+1}\} - \tilde{\mathbf{u}}^{t^*+1}\|^2$$

and setting $\hat{\mathbf{u}}^{t^*+1} = \max\{\ell(\hat{\theta}^{t^*+1}), \tilde{\mathbf{u}}^{t^*+1}\}$. So (12) translates to the following guarantee in the θ -space:

$$\mathbf{E}[\|\ell(\theta^{t^*}) - \tilde{\mathbf{u}}^{t^*+1}\|^2] \leq \mathbf{E}[\min_{\theta \in \mathbb{R}^d} \|\max\{\ell(\theta), \tilde{\mathbf{u}}^{t^*+1}\} - \tilde{\mathbf{u}}^{t^*+1}\|^2] + \epsilon', \quad (13)$$

where we have used $\mathbf{u}^{t^*} = \ell(\theta^{t^*})$. Now since

$$\|\max\{\ell(\theta^{t^*}), \tilde{\mathbf{u}}^{t^*+1}\} - \tilde{\mathbf{u}}^{t^*+1}\|^2 = \|(\ell(\theta^{t^*}) - \tilde{\mathbf{u}}^{t^*+1})_+\|^2 \leq \|\ell(\theta^{t^*}) - \tilde{\mathbf{u}}^{t^*+1}\|^2,$$

together with (13) we have

$$\mathbf{E}[\|\max\{\ell(\theta^{t^*}), \tilde{\mathbf{u}}^{t^*+1}\} - \tilde{\mathbf{u}}^{t^*+1}\|^2] \leq \mathbf{E}[\min_{\theta \in \mathbb{R}^d} \|\max\{\ell(\theta), \tilde{\mathbf{u}}^{t^*+1}\} - \tilde{\mathbf{u}}^{t^*+1}\|^2] + \epsilon'. \quad (14)$$

Having translated our initial assumption on the gradient mapping to θ -space, we can now provide a guarantee on the gradient of $\psi(\ell(\cdot))$. Let $Q(\theta) := \|\max\{\ell(\theta), \tilde{\mathbf{u}}^{t^*+1}\} - \tilde{\mathbf{u}}^{t^*+1}\|^2 = \|(\ell(\theta) - \tilde{\mathbf{u}}^{t^*+1})_+\|^2$.

Taking as given that Q is smooth in θ with smoothness parameter ω for now, by standard properties of smooth functions, we have for any θ' :

$$\|\nabla Q(\theta')\|^2 \leq 2\omega \cdot (Q(\theta') - \min_{\theta \in \mathbb{R}^d} Q(\theta)).$$

Using the above property and (14), taking expectation on both sides, we have:

$$\mathbf{E}[\|\nabla Q(\theta^{t^*})\|^2] \leq 2\omega\epsilon',$$

or equivalently,

$$\mathbf{E}\left[\left\|2\sum_{k=1}^K(\ell_k(\theta^{t^*}) - \tilde{u}_k^{t^*+1})_+ \nabla_{\theta} \ell_k(\theta^{t^*})\right\|^2\right] \leq 2\omega\epsilon',$$

therefore

$$4\eta^2 \mathbf{E}\left[\left\|\sum_{k=1}^K(\nabla \psi_k(\ell^{t^*}))_+ \nabla_{\theta} \ell_k(\theta^{t^*})\right\|^2\right] \leq 2\omega\epsilon',$$

where we use the short-hand $\ell^{t^*} = \ell(\theta^{t^*})$. By monotonicity of ψ , the gradient of ψ is always non-negative, and the above becomes:

$$4\eta^2 \mathbf{E}\left[\left\|\sum_{k=1}^K \nabla \psi_k(\ell^{t^*}) \nabla_{\theta} \ell_k(\theta^{t^*})\right\|^2\right] \leq 2\omega\epsilon',$$

and we have:

$$\mathbf{E}[\|\nabla_{\theta} \psi(\ell(\theta^{t^*}))\|^2] \leq \omega\epsilon'/2\eta^2 = \omega(\epsilon^2 + 2L\epsilon)/2,$$

as desired. It remains to justify the smoothness of $Q(\theta)$. For any $\theta_1, \theta_2 \in \mathbb{R}^d$,

$$\begin{aligned} & \|\nabla Q(\theta_1) - \nabla Q(\theta_2)\| \\ &= \left\|2\sum_{k=1}^K(\ell_k(\theta_1) - \tilde{u}_k^{t^*+1})_+ \cdot \nabla_{\theta} \ell_k(\theta_1) - 2\sum_{k=1}^K(\ell_k(\theta_2) - \tilde{u}_k^{t^*+1})_+ \cdot \nabla_{\theta} \ell_k(\theta_2)\right\| \\ &\leq 2\sum_{k=1}^K \left\|(\ell_k(\theta_1) - \tilde{u}_k^{t^*+1})_+ \cdot (\nabla_{\theta} \ell_k(\theta_1) - \nabla_{\theta} \ell_k(\theta_2))\right\| + \left\|[(\ell_k(\theta_1) - \tilde{u}_k^{t^*+1})_+ - (\ell_k(\theta_2) - \tilde{u}_k^{t^*+1})_+] \cdot \nabla_{\theta} \ell_k(\theta_2)\right\| \\ &\leq 2\sum_{k=1}^K |\ell_k(\theta_1) - \tilde{u}_k^{t^*+1}| \cdot \gamma \|\theta_1 - \theta_2\| + |\ell_k(\theta_1) - \ell_k(\theta_2)| \cdot \|\nabla_{\theta} \ell_k(\theta_2)\| \\ &= 2\sum_{k=1}^K |\ell_k(\theta_1) - \ell_k(\theta^{t^*}) + \eta \nabla \psi_k(u^{t^*})| \cdot \gamma \|\theta_1 - \theta_2\| + \Phi^2 \|\theta_1 - \theta_2\| \\ &\leq 2K[(G + \eta L) \cdot \gamma + \Phi^2] \cdot \|\theta_1 - \theta_2\| = 2K\left[\left(G + \frac{L}{\beta^2}\right) \cdot \gamma + \Phi^2\right] \cdot \|\theta_1 - \theta_2\| \end{aligned}$$

where we used γ -smoothness and Φ -lipschitz property of ℓ_k and $\|\ell(\theta)\| \leq G$, together with $(a)_+ - (b)_+ \leq |a - b|$, therefore $\omega = 2K[(G + \frac{L}{\beta^2}) \cdot \gamma + \Phi^2]$. \square

A.4. Proof of Lemma 3

Recall from Algorithm 2 that the finite difference estimate of the gradient of ψ at θ' is given by:

$$\hat{\mathbf{g}} = \frac{1}{m} \sum_{j=1}^m \frac{M(\mathbf{f}_{\theta'} + \Delta^j, \mathbf{y}) - M(\mathbf{f}_{\theta'}, \mathbf{y})}{\sigma} Z^j.$$

Lemma 3 (Restated). *Let M be as defined in Theorem 2 and $|\epsilon(\theta)| \leq \bar{\epsilon}, \forall \theta$. Let $\hat{\mathbf{g}}$ be returned by Algorithm 2 for a given θ' , m perturbations and $\sigma = \frac{\sqrt{\bar{\epsilon}}}{\sqrt{K}\beta^2}$.*

$$\mathbf{E}[\|\hat{\mathbf{g}} - \nabla \psi(\ell(\theta'))\|^2] \leq \mathcal{O}\left(\frac{L^2 K}{m} + \bar{\epsilon} K^2 \beta^2\right),$$

where the expectation is over the random perturbations.

We will find it useful to re-state results from Nesterov and Spokoiny (2017), extended to our setting.

Lemma 8. *Suppose ψ is L -Lipschitz and β -smooth. Define $\psi_\sigma(\mathbf{u}) := \mathbf{E}_{Z \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_K)} [\psi(\mathbf{u} + \sigma Z)]$. Let $\hat{\mathbf{g}}_1 = \frac{1}{m} \sum_{j=1}^m \frac{\psi(\ell(\mathbf{f}_\theta + \Delta^j, \mathbf{y})) - \psi(\ell(\mathbf{f}_\theta, \mathbf{y}))}{\sigma} Z^j$, where Δ^j is as defined in Algorithm 2. Then:*

1. $\hat{\mathbf{g}}_1$ is an unbiased estimate of the gradient of ψ_σ at $\ell(\theta)$, i.e., $\mathbf{E}[\hat{\mathbf{g}}_1] = \nabla \psi_\sigma(\ell(\theta))$.
2. $\mathbf{E} [\|\hat{\mathbf{g}}_1 - \mathbf{E}[\hat{\mathbf{g}}_1]\|^2] \leq \frac{\sigma^2 \beta^2}{m} (K+6)^3 + \frac{4L^2}{m} (K+4)$.
3. $\|\nabla \psi_\sigma(\ell(\theta)) - \nabla \psi(\ell(\theta))\| \leq \frac{\sigma \beta^2}{2} (K+3)^{3/2}$.

Proof. See Eq. (21) in Nesterov et al. (2017) for part 1. Theorem 4 of Nesterov et al. together with the fact that $\text{Var}(X) \leq \mathbf{E}[X^2]$ implies part 2. See Lemma 3 of Nesterov et al. for part 3. \square

Proof of Lemma 3. We can write out the gradient estimate as:

$$\begin{aligned} \hat{\mathbf{g}} &= \frac{1}{m} \sum_{j=1}^m \frac{M(\mathbf{f}_\theta + \Delta^j, \mathbf{y}) - M(\mathbf{f}_\theta, \mathbf{y})}{\sigma} Z^j \\ &= \frac{1}{m} \sum_{j=1}^m \frac{\psi(\ell(\mathbf{f}_\theta + \Delta^j, \mathbf{y})) - \psi(\ell(\mathbf{f}_\theta, \mathbf{y}))}{\sigma} Z^j + \frac{1}{m} \sum_{j=1}^m \frac{\epsilon(\mathbf{f}_\theta + \Delta^j, \mathbf{y}) - \epsilon(\mathbf{f}_\theta, \mathbf{y})}{\sigma} Z^j \\ &= \frac{1}{m} \sum_{j=1}^m \frac{\psi(\ell(\theta) + \sigma Z^j) - \psi(\ell(\theta))}{\sigma} Z^j + \frac{1}{m} \sum_{j=1}^m \frac{\epsilon(\mathbf{f}_\theta + \Delta^j, \mathbf{y}) - \epsilon(\mathbf{f}_\theta, \mathbf{y})}{\sigma} Z^j \\ &:= \hat{\mathbf{g}}_1 + \hat{\mathbf{g}}_2, \end{aligned}$$

where $\epsilon(\mathbf{f}_\theta, \mathbf{y})$ is the unknown slack function in Section 3.1, re-written in terms of the scores \mathbf{f}_θ and labels \mathbf{y} .

Let ψ_σ be defined as in Lemma 8. Then the gradient estimate error can be expanded as:

$$\begin{aligned} \mathbf{E} [\|\hat{\mathbf{g}} - \nabla \psi(\ell(\theta))\|^2] &\leq 2\mathbf{E} [\|\hat{\mathbf{g}} - \nabla \psi_\sigma(\ell(\theta))\|^2] + 2\|\nabla \psi_\sigma(\ell(\theta)) - \nabla \psi(\ell(\theta))\|^2 \\ &\leq 4\mathbf{E} [\|\hat{\mathbf{g}}_1 - \nabla \psi_\sigma(\ell(\theta))\|^2] + 4\mathbf{E} [\|\hat{\mathbf{g}}_2\|^2] + 2\|\nabla \psi_\sigma(\ell(\theta)) - \nabla \psi(\ell(\theta))\|^2 \\ &\leq 4\mathbf{E} [\|\hat{\mathbf{g}}_1 - \nabla \psi_\sigma(\ell(\theta))\|^2] + \frac{16\bar{\epsilon}^2}{\sigma^2 m} \sum_{j=1}^m \mathbf{E} [\|Z^j\|^2] + 2\|\nabla \psi_\sigma(\ell(\theta)) - \nabla \psi(\ell(\theta))\|^2 \\ &\leq \frac{4\sigma^2}{m} \beta^2 (K+6)^3 + \frac{16}{m} L^2 (K+4) + \frac{16\bar{\epsilon}^2 K}{\sigma^2} + \frac{\sigma^2}{2} \beta^4 (K+3)^3, \end{aligned}$$

where we used the fact that (1) $\hat{\mathbf{g}}_1$ is an unbiased estimate of $\nabla \psi_\sigma(\ell(\theta))$ (see part 1 of Lemma 8); (2) the assumption that $|\epsilon(\theta)| \leq \bar{\epsilon}$; (3) $\|a_1 + \dots + a_m\|^2 \leq m(\|a_1\|^2 + \dots + \|a_m\|^2)$, and the last step follows from Parts 2–3 of Lemma 8.

Setting $\sigma = \frac{\sqrt{\bar{\epsilon}}}{\sqrt{K}\beta^2}$ completes the proof. \square

A.5. Proofs and Discussion for Linear Interpolation Gradient Estimates

Lemma 4 (Restated). *Let M be defined as in Theorem 2 and $|\epsilon(\theta)| \leq \bar{\epsilon}, \forall \theta$. Assume each ℓ_k is Φ -Lipschitz in θ w.r.t. the L_∞ -norm, and $\|\ell(\theta)\| \leq G \forall \theta$. Suppose for a given θ^l , σ and perturbation count m , the expected covariance matrix for the left-hand-side of the linear system \mathbf{H} is well-conditioned with the smallest singular value $\lambda_{\min}(\sum_{i=1}^m \mathbf{E}[\mathbf{H}_i \mathbf{H}_i^\top]) \geq \mu_{\min} = \mathcal{O}(m\sigma^2\Phi^2)$. Then for any $\delta > 0$, setting $\sigma = \frac{G^{1/3}\bar{\epsilon}^{1/3}}{\Phi K^{3/2} \log(d)^{2/3} \beta^{1/3}}$ and $m = \frac{G^4 K^9 \log(d)^4 \beta^2 \log(K/\delta)}{\bar{\epsilon}^2}$, Algorithm 3 returns w.p. $\geq 1 - \delta$ (over draws of random perturbations) a gradient estimate $\hat{\mathbf{g}}$ that satisfies:*

$$\|\hat{\mathbf{g}} - \nabla \psi(\ell(\theta^l))\|^2 \leq \tilde{\mathcal{O}} \left(G^{1/3} \bar{\epsilon}^{1/3} K^3 \beta^{2/3} \right).$$

We first discuss the assumptions in Lemma 4 in Section A.5.1. We then provide the proof for the high probability statement in the lemma in Section A.5.2. We then show how this can be translated to a bound on the expected gradient error via truncation in Section A.5.3.

A.5.1. ASSUMPTIONS IN LEMMA 4

We discuss example settings where the assumptions in the lemma hold.

Correlation Assumption on \mathbf{H} . One of the key assumptions we make is that the matrix \mathbf{H} is well-conditioned. Recall that \mathbf{H} is a $m \times K$ matrix, where each row corresponds to a perturbation of the surrogates, and contains differences in the K surrogates ℓ_1, \dots, ℓ_K at two independent perturbations to the model parameters θ . We assume that the smallest singular value of \mathbf{H} 's covariance matrix $\sum_{i=1}^m \mathbf{E}[\mathbf{H}_i \mathbf{H}_i^\top]$ scales as $m\sigma^2\Phi^2$. This assumption essentially states that the perturbations on the K surrogates are weakly correlated. The scaling factors σ and Φ come from the fact that Gaussian perturbations on the model parameters θ have standard deviation σ and the surrogates ℓ_k are Φ -Lipschitz.

As an example scenario where this assumption holds, consider a ML fairness task where the instances belong to K non-overlapping protected groups. Further, assume that the group membership attribute is included in the feature vector, i.e., the d -dimensional feature vector $\mathbf{x} = [g_1, \dots, g_K, \tilde{x}_1, \dots, \tilde{x}_{d-K}]$, where g_k is a Boolean indicating if the instance belongs to group k , and $\tilde{x}_1, \dots, \tilde{x}_{d-K}$ are group-independent features. A natural choice of surrogates for this application would be average losses computed on the K individual groups. For example, with a linear model θ , we could choose ℓ_k to be the average squared loss conditioned on examples from group k , i.e., $\ell_k(\theta) = \mathbf{E}_{(x,y)|x_k=1}[(\theta^\top x - y)^2]$.

Note that the first K coordinates of the model vector θ correspond to weights on the K Boolean group attributes. So adding noise $Z_k \in \mathbb{R}$ to the k -th coordinate of θ only affects scores on examples from the k -th group (i.e., examples for which $x_k = 1$), and hence only perturbs surrogate ℓ_k . Specifically, adding $Z_k \in \mathbb{R}$ to the k -th coordinate of θ would perturb $\ell_k(\theta)$ to $\ell_k(\theta) + C_k Z_k + Z_k^2$, where $C_k = 2\mathbf{E}_{(x,y)|x_k=1}[\theta^\top x - y]$, and leave the other surrogates $\ell_j, j \neq k$ unchanged.

Now suppose we add independent σ -Gaussian noise to only the first K coordinates of θ . The expected covariance matrix as defined in the lemma statement then takes the form:

$$\sum_{i=1}^m \mathbf{E}[\mathbf{H}_i \mathbf{H}_i^\top] = \begin{bmatrix} \mathcal{O}(m(C_1^2\sigma^2 + \sigma^4)) & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & \mathcal{O}(m(C_K^2\sigma^2 + \sigma^4)) \end{bmatrix} = \begin{bmatrix} \Omega(m\sigma^2) & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & \Omega(m\sigma^2) \end{bmatrix},$$

where recall that the k -th column of \mathbf{H} contains the differences of $\ell_k(\theta)$ at two different σ -Gaussian perturbations on the first K coordinates of θ , and C_k 's are constants that are independent of the random perturbations.

In the more general case, where we perturb all coordinates of θ , the assumption on \mathbf{H} would still hold if there exists a subset of coordinates for each surrogate ℓ_k that when perturbed produce larger changes to ℓ_k than to the other surrogates.

Lipschitz Assumption on $\ell(\theta)$ Another key assumption we make is that the surrogates ℓ_k are Φ -Lipschitz w.r.t. the L_∞ -norm. This allows us to produce perturbations in the K surrogates by perturbing the model parameters θ , and do so without a strong dependence on the dimension of θ in the error bound. Note that the choice of the infinity norm results in a mild logarithmic dependence on the dimension d in the bound. When the surrogates are not L_∞ -Lipschitz, but are instead Lipschitz w.r.t. the L_2 -norm, we prescribe perturbing only a small number of $d' \ll d$ coordinates of θ that are most closely related to the surrogate (such as e.g. the group attribute coordinates in the fairness example above), and this would result in a bound that has a polynomial dependence on d' .

A.5.2. PROOF OF LEMMA 4

We will make use of the fact that because we perturb the model parameters θ with Gaussian random noise, the resulting perturbations on the surrogates ℓ follow a sub-Gaussian distribution. We first state a few well-known facts about sub-Gaussian random vectors.

Lemma 9 (Properties of sub-Gaussian distribution).

- (i) Let (Z_1, \dots, Z_d) be a vector of i.i.d standard gaussian variables and $f : \mathbb{R}^N \rightarrow \mathbb{R}$ be Φ -Lipschitz w.r.t. L_2 -norm. Then the random variable $f(\sigma Z) - \mathbf{E}[f(\sigma Z)]$ is sub-Gaussian with parameter at most $\sigma\Phi$.
- (ii) Let Z_1, \dots, Z_K be K (not necessarily independent) sub-Gaussian random variables with parameters at most σ . Then the random vector (Z_1, \dots, Z_K) is a sub-Gaussian random vector with parameter σK .
- (iii) For a sub-Gaussian random vector $Z \in \mathbb{R}^K$ with parameter at most σ , we have for any $p \in \mathbb{N}$:

$$(\mathbf{E}[\|Z - \mathbf{E}[Z]\|_2^p])^{1/p} \leq 2\sqrt{2}\sigma\sqrt{K}\sqrt{p}.$$

Proof. For a proof of (1), see e.g. [Wainwright \(2019\)](#), Chapter 2. For a proof of (3), see [Jin et al. \(2019\)](#). We now prove (2).

For a random vector (Z_1, \dots, Z_K) where the coordinates Z_k 's are σ -sub-Gaussian and not necessarily independent, we have that for any $v \in \mathbb{S}^{K-1}$ and $\lambda \in \mathbb{R}$,

$$\begin{aligned} \mathbf{E}[\exp(\lambda v^\top (Z - \mathbf{E}[Z]))] &= \mathbf{E}\left[\prod_{k=1}^K \exp\left(\lambda v_k (Z_k - \mathbf{E}[Z_k])\right)\right] \\ &\leq \prod_{k=1}^K \mathbf{E}\left[\left(\exp(\lambda v_k (Z_k - \mathbf{E}[Z_k]))\right)^K\right]^{1/K} \\ &\leq \prod_{k=1}^K \exp\left(\frac{1}{2}\lambda^2 K^2 \sigma^2\right)^{1/K} = \prod_{k=1}^K \exp\left(\frac{1}{2}\lambda^2 \sigma^2 K\right) \leq \exp\left(\frac{1}{2}\lambda^2 \sigma^2 K^2\right), \end{aligned}$$

where we have used Hölder's inequality for the second step. \square

We can write the optimization problem in [Algorithm 3](#) as solving the following linear system

$$\begin{bmatrix} h'_{11} - h''_{11} & \cdots & h'_{1K} - h''_{1K} \\ \vdots & \cdots & \vdots \\ h'_{m1} - h''_{m1} & \cdots & h'_{mK} - h''_{mK} \end{bmatrix} \cdot \hat{\mathbf{g}} = \begin{bmatrix} \psi(\ell(\theta') + \mathbf{h}'_1) - \psi(\ell(\theta') + \mathbf{h}''_1) + \epsilon_{11} - \epsilon_{12} \\ \vdots \\ \psi(\ell(\theta') + \mathbf{h}'_m) - \psi(\ell(\theta') + \mathbf{h}''_m) + \epsilon_{m1} - \epsilon_{m1} \end{bmatrix},$$

and use the resulting $\hat{\mathbf{g}} \in \mathbb{R}^K$ as the gradient estimate, where we denote $\mathbf{h}'_j := \ell(\theta' + \sigma Z_1^j) - \ell(\theta') \in \mathbb{R}^K$ and $\mathbf{h}''_j := \ell(\theta' + \sigma Z_2^j) - \ell(\theta') \in \mathbb{R}^K$ for $j \in [m]$, and $\epsilon_{j1} = \epsilon(\theta' + \sigma Z_1^j)$ and $\epsilon_{j2} = \epsilon(\theta' + \sigma Z_2^j)$. We further denote

$$\mathbf{L} = [\ell(\theta'); \dots; \ell(\theta')] \in \mathbb{R}^{m \times K}$$

$$\mathbf{H}' = [\mathbf{h}'_1; \dots; \mathbf{h}'_m] \in \mathbb{R}^{m \times K}, \quad \mathbf{H}'' = [\mathbf{h}''_1; \dots; \mathbf{h}''_m] \in \mathbb{R}^{m \times K}$$

$$\epsilon_1 = [\epsilon_{11}; \dots; \epsilon_{m1}] \in \mathbb{R}^m, \quad \epsilon_2 = [\epsilon_{12}; \dots; \epsilon_{m2}] \in \mathbb{R}^m,$$

and equivalently re-write the above linear system as:

$$(\mathbf{H}' - \mathbf{H}'') \cdot \hat{\mathbf{g}} = \psi(\mathbf{L} + \mathbf{H}') - \psi(\mathbf{L} + \mathbf{H}'') + \epsilon_1 - \epsilon_2, \quad (15)$$

where the matrix \mathbf{H} that we defined in the lemma statement is the same as $\mathbf{H}' - \mathbf{H}''$.

Below we state a lemma involving implications of our assumptions on the left-hand-side perturbation matrices \mathbf{H}' and \mathbf{H}'' .

Lemma 10 (Properties of \mathbf{H}' and \mathbf{H}''). *Suppose each $\ell_k(\theta)$ is Φ -Lipschitz w.r.t. the L_∞ -norm and $\|\ell(\theta)\| \leq G, \forall \theta$. Then each \mathbf{h}'_i and each \mathbf{h}''_i is a sub-Gaussian vector with parameter at most $\sigma\Phi K$. The differences $\mathbf{h}'_i - \mathbf{h}''_i$ are also sub-Gaussian random vectors with parameter at most $2\sigma\Phi K$, and have mean zero. Moreover, $\|\mathbf{h}'_i\| \leq 2G$, $\|\mathbf{h}''_i\| \leq 2G$, $\|\mathbf{h}'_i - \mathbf{h}''_i\| \leq 2G, \forall i$.*

The proof follows directly from [Lemma 9](#)(i)–(ii) and the fact that a function ℓ_k that is Φ -Lipschitz w.r.t. the L_∞ -norm is also Φ -Lipschitz w.r.t. the L_2 -norm. We also have from the smoothness of ψ that

$$|\psi(\ell(\theta') + \mathbf{h}'_i) - [\psi(\ell(\theta')) + \nabla\psi(\ell(\theta'))^\top \mathbf{h}'_i]| \leq \frac{\beta}{2} \|\mathbf{h}'_i\|_2^2. \quad (16)$$

With this in hand, we are ready to bound the error in the gradient estimate $\hat{\mathbf{g}}$ compared to $\nabla\psi(\ell(\theta))$.

Proof of Lemma 4. The least squares estimate for the linear system in (15) is given by:

$$\begin{aligned}
 \hat{\mathbf{g}} &= \left((\mathbf{H}' - \mathbf{H}'')^\top (\mathbf{H}' - \mathbf{H}'') \right)^{-1} (\mathbf{H}' - \mathbf{H}'')^\top [\psi(\mathbf{L} + \mathbf{H}') + \boldsymbol{\epsilon}_1 - \psi(\mathbf{L} + \mathbf{H}'') - \boldsymbol{\epsilon}_2] \\
 &= \left((\mathbf{H}' - \mathbf{H}'')^\top (\mathbf{H}' - \mathbf{H}'') \right)^{-1} (\mathbf{H}' - \mathbf{H}'')^\top \left[(\mathbf{H}' - \mathbf{H}'') \nabla \psi(\boldsymbol{\ell}) + \psi(\mathbf{L} + \mathbf{H}') - \psi(\mathbf{L} + \mathbf{H}'') \right. \\
 &\quad \left. - (\mathbf{H}' - \mathbf{H}'') \nabla \psi(\boldsymbol{\ell}) + \boldsymbol{\epsilon}_1 - \boldsymbol{\epsilon}_2 \right] \\
 &= \nabla \psi(\boldsymbol{\ell}(\theta')) + \left((\mathbf{H}' - \mathbf{H}'')^\top (\mathbf{H}' - \mathbf{H}'') \right)^{-1} (\mathbf{H}' - \mathbf{H}'')^\top \left[\psi(\mathbf{L} + \mathbf{H}') - \psi(\mathbf{L} + \mathbf{H}'') \right. \\
 &\quad \left. - (\mathbf{H}' - \mathbf{H}'') \nabla \psi(\boldsymbol{\ell}) + \boldsymbol{\epsilon}_1 - \boldsymbol{\epsilon}_2 \right].
 \end{aligned}$$

The error in the least squares based gradient estimate is then:

$$\begin{aligned}
 &\|\hat{\mathbf{g}} - \nabla \psi(\boldsymbol{\ell}(\theta'))\| \\
 &\leq \underbrace{\left\| \left((\mathbf{H}' - \mathbf{H}'')^\top (\mathbf{H}' - \mathbf{H}'') \right)^{-1} \right\|_{\text{op}}}_{\text{term}_1} \underbrace{\left\| (\mathbf{H}' - \mathbf{H}'')^\top \left[\psi(\mathbf{L} + \mathbf{H}') - \psi(\mathbf{L} + \mathbf{H}'') - (\mathbf{H}' - \mathbf{H}'') \nabla \psi(\boldsymbol{\ell}) + \boldsymbol{\epsilon}_1 - \boldsymbol{\epsilon}_2 \right] \right\|_2}_{\text{term}_2}.
 \end{aligned} \tag{17}$$

Bounding the second term in (17). We first bound the second term in (17). We have:

$$\begin{aligned}
 &\|\psi(\mathbf{L} + \mathbf{H}') - \psi(\mathbf{L} + \mathbf{H}'') - (\mathbf{H}' - \mathbf{H}'') \nabla \psi(\boldsymbol{\ell}) + \boldsymbol{\epsilon}_1 - \boldsymbol{\epsilon}_2\|_2 \\
 &= \left\| \psi(\mathbf{L} + \mathbf{H}') - \mathbf{H}' \nabla \psi(\boldsymbol{\ell}) - \psi(\mathbf{L}) + \psi(\mathbf{L}) + \mathbf{H}'' \nabla \psi(\boldsymbol{\ell}) - \psi(\mathbf{L} + \mathbf{H}'') + \boldsymbol{\epsilon}_1 - \boldsymbol{\epsilon}_2 \right\|_2 \\
 &\leq \left\| \psi(\mathbf{L} + \mathbf{H}') - \mathbf{H}' \nabla \psi(\boldsymbol{\ell}) - \psi(\mathbf{L}) \right\|_2 + \left\| \psi(\mathbf{L}) + \mathbf{H}'' \nabla \psi(\boldsymbol{\ell}) - \psi(\mathbf{L} + \mathbf{H}'') \right\|_2 + \|\boldsymbol{\epsilon}_1\|_2 + \|\boldsymbol{\epsilon}_2\|_2 \\
 &\leq \frac{\beta}{2} \|\mathbf{H}'\|_{\text{F}}^2 + \frac{\beta}{2} \|\mathbf{H}''\|_{\text{F}}^2 + 2\sqrt{m}\bar{\epsilon},
 \end{aligned}$$

where we used (16) and the assumption $|\epsilon(\theta)| \leq \bar{\epsilon} \forall \theta$. This in turn gives

$$\begin{aligned}
 \text{term}_2 &= \left\| (\mathbf{H}' - \mathbf{H}'')^\top \left[\psi(\mathbf{L} + \mathbf{H}') - \psi(\mathbf{L} + \mathbf{H}'') - (\mathbf{H}' - \mathbf{H}'') \nabla \psi(\boldsymbol{\ell}) + \boldsymbol{\epsilon}_1 - \boldsymbol{\epsilon}_2 \right] \right\|_2 \\
 &\leq \sum_{j=1}^m \|\mathbf{h}'_j - \mathbf{h}''_j\| \cdot \frac{\beta}{2} (\|\mathbf{h}'_j\|_2^2 + \|\mathbf{h}''_j\|_2^2) + 2\sqrt{m}G\sqrt{m}\bar{\epsilon}
 \end{aligned}$$

where each \mathbf{h}'_j is of length K with (correlated) subgaussian coordinates. Therefore using Cauchy-Schwarz,

$$\mathbf{E}[\|\mathbf{h}'_j - \mathbf{h}''_j\| \cdot \|\mathbf{h}'_j\|_2^2] \leq \sqrt{\mathbf{E}[\|\mathbf{h}'_j - \mathbf{h}''_j\|_2^2]} \cdot \sqrt{\mathbf{E}[\|\mathbf{h}'_j\|_2^4]}.$$

Note that $\mathbf{E}[h'_{ij}] = \mathbf{E}[\ell_j(\theta' + \sigma Z^i)] - \ell_j(\theta') \leq \sigma \Phi \mathbf{E}[\|Z^i\|_\infty] \leq \mathcal{O}(\sigma \Phi \sqrt{\log(d)})$, where we've used that the max of d independent standard normal random variables scales as $\sqrt{\log(d)}$. Similarly, $\mathbf{E}[h''_{ij}] \leq \mathcal{O}(\sigma \Phi \sqrt{\log(d)})$. Together with these facts and Lemma 10 and Lemma 9(iii) we have

$$\sqrt{\mathbf{E}[\|\mathbf{h}'_j\|_2^4]} \leq \sqrt{8(4\sqrt{2}\sigma\Phi K^{3/2})^4 + \mathcal{O}(\sigma^2\Phi^2 \log(d)K)^2} \leq \mathcal{O}(\sigma^2\Phi^2 K^3 (\log(d))^2)$$

where we used triangle inequality and $(a+b)^p \leq 2^{p-1}(a^p + b^p)$. Similarly, we have:

$$\sqrt{\mathbf{E}[\|\mathbf{h}'_j - \mathbf{h}''_j\|_2^2]} \leq 4\sigma\Phi K^{3/2}.$$

Now since $\|\mathbf{h}'_j\|_2 \leq G$, we can apply Hoeffding's inequality to these bounded random variables to get

$$\mathbf{P} \left(\sum_{j=1}^m \|\mathbf{h}'_j - \mathbf{h}''_j\|_2 \cdot \|\mathbf{h}'_j\|_2^2 \geq \mathcal{O}(\sigma^3\Phi^3 K^{9/2} (\log(d))^2 m) + mt \right) \leq 2 \exp \left(-\frac{2mt^2}{G^6} \right),$$

which further gives us:

$$\mathbf{P} \left(\text{term}_2 \geq \mathcal{O}(\sigma^3 \Phi^3 K^{9/2} (\log(d))^2 m \beta) + m \beta t + 2mG\bar{\epsilon} \right) \leq 2 \exp \left(-\frac{2mt^2}{G^6} \right), \quad (18)$$

Bounding the first term in (17). Now the first term in (17) is simply

$$\text{term}_1 = \left\| \left((\mathbf{H}' - \mathbf{H}'')^\top (\mathbf{H}' - \mathbf{H}'') \right)^{-1} \right\|_{\text{op}} = \lambda_{\min}^{-1} \left((\mathbf{H}' - \mathbf{H}'')^\top (\mathbf{H}' - \mathbf{H}'') \right).$$

Let us denote $\hat{\Sigma} := \sum_{i=1}^m (\mathbf{h}'_i - \mathbf{h}''_i)(\mathbf{h}'_i - \mathbf{h}''_i)^\top$ as the empirical covariance matrix. We now apply a matrix Chernoff inequality (see e.g. [Tropp \(2015\)](#)) to lower bound the smallest eigenvalue of $\hat{\Sigma}$. We first note that the largest eigenvalue of this matrix is bounded above:

$$\lambda_{\max}(\hat{\Sigma}) = \max_{\|u\|=1} \frac{1}{m} \sum_{i=1}^m \left((\mathbf{h}'_i - \mathbf{h}''_i)^\top u \right)^2 \leq 4G^2,$$

This together with the matrix Chernoff bound gives us for $\mu_{\min} \leq \lambda_{\min}(\hat{\Sigma})$, we have

$$\mathbf{P} \left(\lambda_{\min}(\hat{\Sigma}) \leq \frac{\mu_{\min}}{2} \right) \leq K \cdot \exp \left(-\frac{\mu_{\min}}{32G^2} \right).$$

The assumption $\mu_{\min} = \mathcal{O}(m\sigma^2\Phi^2)$ then yields:

$$\mathbf{P} \left(\text{term}_1 \leq \mathcal{O}(m\sigma^2\Phi^2) \right) \leq K \cdot \exp \left(-\frac{m\sigma^2\Phi^2}{G^2} \right). \quad (19)$$

Combining the above bound (19) with the bound on the second term (18) (picking $t = \sigma^3\Phi^3$), we get the following tail bound:

$$\mathbf{P} \left(\|\hat{\mathbf{g}} - \nabla\psi(\ell(\theta'))\| \geq \mathcal{O} \left(\sigma\Phi K^{9/2} \log(d)^2 \beta + \frac{G\bar{\epsilon}}{\sigma^2\Phi^2} \right) \right) \leq K \cdot \exp \left(-\frac{m\sigma^2\Phi^2}{G^2} \right) + 4 \exp \left(-\frac{2m\sigma^6\Phi^6}{G^6} \right).$$

Then for any $\delta > 0$, setting $\sigma = \frac{G^{1/3}\bar{\epsilon}^{1/3}}{\Phi K^{3/2} \log(d)^{2/3} \beta^{1/3}}$ and $m = \frac{G^4 K^9 \log(d)^4 \beta^2 \log(K/\delta)}{\bar{\epsilon}^2}$, Algorithm 3 returns w.p. $\geq 1 - \delta$ (over draws of random perturbations) a gradient estimate $\hat{\mathbf{g}}$ that satisfies:

$$\|\hat{\mathbf{g}} - \nabla\psi(\ell(\theta'))\|^2 \leq \mathcal{O} \left(G^{1/3} \bar{\epsilon}^{1/3} K^3 (\log(d))^{4/3} \beta^{2/3} \right),$$

which completes the proof. \square

A.5.3. TRANSLATING TO A BOUND ON THE EXPECTED ERROR

Lemma 4 provides a high probability bound on the gradient estimation error. This means that with a small probability the gradient estimation error may not be bounded. To translate this high probability bound into a bound on the expected gradient error, we first truncate the estimated gradients to be in a bounded range:

$$\text{trunc}(\hat{\mathbf{g}}) = \begin{cases} \hat{\mathbf{g}} & \text{if } \|\hat{\mathbf{g}}\| \leq 2\sqrt{KL} \\ \mathbf{0} & \text{otherwise} \end{cases},$$

where L is the Lipschitz constant for ψ .

Corollary 1. *Under the assumptions in Lemma 4, for any $\delta \in (0, 1)$, setting $\sigma = \frac{G^{1/3}\bar{\epsilon}^{1/3}}{\Phi K^{3/2} \log(d)^{2/3} \beta^{1/3}}$ and $m = \frac{G^4 K^9 \log(d)^4 \beta^2 \log(K/\delta)}{\bar{\epsilon}^2}$, Algorithm 3 returns a gradient estimate $\hat{\mathbf{g}}$ that satisfies:*

$$\mathbf{E} \left[\|\text{trunc}(\hat{\mathbf{g}}) - \nabla\psi(\ell(\theta'))\|^2 \right] \leq \tilde{\mathcal{O}} \left(G^{1/3} \bar{\epsilon}^{1/3} K^3 \beta^{2/3} \right) + 10KL^2\delta.$$

Proof. Because both the truncated gradient estimates and the true gradients are bounded, the gradient error is trivially bounded by:

$$\|\text{trunc}(\hat{\mathbf{g}}) - \nabla\psi(\boldsymbol{\ell}(\theta'))\|^2 \leq 2(\|\text{trunc}(\hat{\mathbf{g}})\|^2 + \|\nabla\psi(\boldsymbol{\ell}(\theta'))\|^2) \leq 2(4KL^2 + L^2) \leq 10KL^2. \quad (20)$$

In the case where $\|\hat{\mathbf{g}}\| \leq 2\sqrt{KL}$, the gradient error for the truncated $\hat{\mathbf{g}}$ is the same as that for $\hat{\mathbf{g}}$:

$$\|\text{trunc}(\hat{\mathbf{g}}) - \nabla\psi(\boldsymbol{\ell}(\theta'))\|^2 = \|\hat{\mathbf{g}} - \nabla\psi(\boldsymbol{\ell}(\theta'))\|^2. \quad (21)$$

When $\|\hat{\mathbf{g}}\| > 2\sqrt{KL}$, the gradient error for the truncated estimates $\text{trunc}(\hat{\mathbf{g}})$ is upper bounded by:

$$\|\text{trunc}(\hat{\mathbf{g}}) - \nabla\psi(\boldsymbol{\ell}(\theta'))\|^2 = \|\nabla\psi(\boldsymbol{\ell}(\theta'))\|^2 \leq L^2,$$

whereas the the gradient error for the original estimates $\hat{\mathbf{g}}$ is lower bounded by:

$$\begin{aligned} \|\hat{\mathbf{g}} - \nabla\psi(\boldsymbol{\ell}(\theta'))\|^2 &\geq \max_{k \in [K]} (\hat{g}_k - \nabla_k \psi(\boldsymbol{\ell}(\theta')))^2 \geq \left(\max_{k \in [K]} |\hat{g}_k| - \max_{k \in [K]} |\nabla_k \psi(\boldsymbol{\ell}(\theta'))| \right)^2 \\ &\geq \left(\frac{1}{\sqrt{K}}(2\sqrt{KL}) - L \right)^2 = L^2. \end{aligned}$$

Therefore even in this case, the gradient error for $\text{trunc}(\hat{\mathbf{g}})$ is bounded by that for $\hat{\mathbf{g}}$:

$$\|\text{trunc}(\hat{\mathbf{g}}) - \nabla\psi(\boldsymbol{\ell}(\theta'))\|^2 \leq L^2 \leq \|\hat{\mathbf{g}} - \nabla\psi(\boldsymbol{\ell}(\theta'))\|^2. \quad (22)$$

Combining (21) and (22) with the trivial upper bound in (20) allows us to convert the the high probability result in Lemma 4 to the following bound on the expected error. For any $\delta \in (0, 1)$, setting $\sigma = \frac{G^{1/3}\bar{\epsilon}^{1/3}}{\Phi K^{3/2} \log(d)^{2/3} \beta^{1/3}}$ and $m = \frac{G^4 K^9 \log(d)^4 \beta^2 \log(K/\delta)}{\bar{\epsilon}^2}$, we have:

$$\mathbf{E} [\|\text{trunc}(\hat{\mathbf{g}}) - \nabla\psi(\boldsymbol{\ell}(\theta'))\|^2] \leq \tilde{\mathcal{O}} \left((1 - \delta) G^{1/3} \bar{\epsilon}^{1/3} K^3 \beta^{2/3} \right) + 10\delta KL^2,$$

as desired. \square

B. Handling Non-smooth Metrics

For ψ that is only L -Lipschitz and non-smooth, we extend the finite difference gradient estimate in Section 5.1 with a two-step perturbation method, as detailed in Algorithm 4. This approach can be seen as computing a finite-difference gradient estimate for a smooth approximation to the original ψ , given by $\psi_{\sigma_1}(\mathbf{u}) := \mathbf{E}[\psi(\mathbf{u} + \sigma_1 Z_1)]$, where $Z_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_K)$. Since ψ_{σ_1} is a convolution of ψ with a Gaussian density kernel, it is always smooth. For this setting, we build on recent work by Duchi et al. (2015), and show that the two-step perturbation approach provides a gradient estimate for ψ_{σ_1} .

Lemma 11 (Two-step finite difference gradient estimate). *Let $M(\theta) = \psi(\boldsymbol{\ell}(\theta)) + \epsilon(\theta)$, for a ψ that is L -Lipschitz, and the worst-case slack $\max_{\theta \in \mathbb{R}^d} |\epsilon(\theta)|$ is the minimum among all such decompositions of M . Suppose $|\epsilon(\theta)| \leq \bar{\epsilon}$, $\forall \theta$. Let $\hat{\mathbf{g}}$ be returned by Algorithm 4 for a fixed $\sigma_1 > 0$ and $\sigma_2 = \sqrt{\frac{\sigma_1}{K^{3/2}L}}$. Then:*

$$\mathbf{E} [\|\hat{\mathbf{g}} - \nabla\psi_{\sigma_1}(\boldsymbol{\ell}(\theta))\|^2] \leq \tilde{\mathcal{O}} \left(\frac{L^{7/4} K^{13/8}}{m\sigma_1^{1/4}} + \frac{LK^{5/2}\bar{\epsilon}^2}{\sigma_1} \right).$$

Drawing upon the result of Theorem 2, we can repeat the analysis on the smooth function $\psi_{\sigma_1}(\cdot)$ to get the following convergence guarantee for Algorithm 1.

Corollary 2 (Convergence of Algorithm 1 for non-smooth ψ). *Let $M(\theta) = \psi(\boldsymbol{\ell}(\theta)) + \epsilon(\theta)$, for a ψ that is monotonic, and L -Lipschitz, and the worst-case slack $\max_{\theta \in \mathbb{R}^d} |\epsilon(\theta)|$ is the minimum among all such decompositions of M .*

Suppose each ℓ_k is γ -smooth and Φ -Lipschitz in θ with $\|\boldsymbol{\ell}(\theta)\| \leq G$, $\forall \theta$. Suppose the gradient $\hat{\mathbf{g}}^t$ are estimated with Algorithm 1 for a choice $\sigma_1 > 0$, number of perturbation m , and $\sigma_2 = \sqrt{\frac{\sigma_1}{K^{3/2}L}}$. Suppose the projection step satisfies $\|(\boldsymbol{\ell}(\theta^{t+1}) - \hat{\mathbf{u}}^t)_+\|^2 \leq \min_{\theta \in \mathbb{R}^d} \|(\boldsymbol{\ell}(\theta) - \hat{\mathbf{u}}^t)_+\|^2 + \mathcal{O}(\frac{\sigma_1^2}{TKL^2})$, $\forall t \in [T]$. Set stepsize $\eta = \frac{\sigma_1^2}{KL^2}$.

Algorithm 4 Two-step Finite-difference Gradient Estimate

- 1: **Input:** $\theta \in \mathbb{R}^d, M, \ell_1, \dots, \ell_k$, estimation accuracy ϵ
- 2: Draw $Z_1^1, \dots, Z_1^m, Z_2^1, \dots, Z_2^m \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_K)$
- 3: Find $\Delta_1^j \in \mathbb{R}^n$ s.t. $\ell(\mathbf{f}_\theta + \Delta_1^j, \mathbf{y}) = \ell(\mathbf{f}_\theta, \mathbf{y}) + \sigma_1 Z_1^j$ for $j = 1, \dots, m$
- 4: Find $\Delta_2^j \in \mathbb{R}^n$ s.t. $\ell(\mathbf{f}_\theta + \Delta_2^j, \mathbf{y}) = \ell(\mathbf{f}_\theta, \mathbf{y}) + \sigma_1 Z_1^j + \sigma_2 Z_2^j$ for $j = 1, \dots, m$
- 5: $\hat{\mathbf{g}} = \frac{1}{m} \sum_{j=1}^m \frac{M(\mathbf{f}_\theta + \Delta_2^j, \mathbf{y}) - M(\mathbf{f}_\theta + \Delta_1^j, \mathbf{y})}{\sigma_2} Z_2^j$
- 6: **Output:** $\hat{\mathbf{g}}$

Then Algorithm 1 converges to an approximate stationary point of the smooth approximation $\psi_{\sigma_1}(\ell(\cdot))$:

$$\min_{1 \leq t \leq T} \mathbf{E} [\|\nabla \psi_{\sigma_1}(\ell(\theta^t))\|^2] \leq C \left(\frac{\sqrt{KL}}{\sigma_1 \sqrt{T}} + \sqrt{\kappa} + \sqrt{L\kappa^{1/4}} \right),$$

where the expectation is over the randomness in the gradient estimates, and $C = \mathcal{O}(KL(\gamma(G + \frac{\sigma_1^2}{KL}) + \Phi^2))$ and $\kappa = \tilde{\mathcal{O}}\left(\frac{L^{7/4}K^{13/8}}{m\sigma_1^{1/4}} + \frac{LK^{5/2}\epsilon^2}{\sigma_1}\right)$.

The above result guarantees convergence to the stationary point of the smoothed metric $\psi_{\sigma_1}(\ell(\cdot))$ and not the original metric $\psi(\ell(\cdot))$. However, as long as the surrogate functions ℓ are continuously differentiable, by taking $\sigma_1 \rightarrow 0$ and allowing T to increase as σ_1 decreases, the algorithm can be made to converge to a stationary point of the original metric $\psi(\ell(\cdot))$, in the sense of Clark-subdifferential (see e.g. Garmanjani and Vicente (2013)).

B.1. Proof of Lemma 11

We will find it useful to re-state results from Duchi et al. (2015) and Nesterov and Spokoiny (2017), extended to our setting.

Lemma 12. Suppose ψ is L -Lipschitz. Define $\psi_{\sigma_1}(\mathbf{u}) := \mathbf{E}_{Z_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_K)} [\psi(\mathbf{u} + \sigma_1 Z_1)]$ and $\psi_{\sigma_1, \sigma_2}(\mathbf{u}) := \mathbf{E}_{Z_2 \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_K)} [\psi_{\sigma_1}(\mathbf{u} + \sigma_2 Z_2)]$. Let $\hat{\mathbf{g}}_1 = \frac{1}{m} \sum_{j=1}^m \frac{\psi(\ell(\mathbf{f}_\theta + \Delta_2^j, \mathbf{y})) - \psi(\ell(\mathbf{f}_\theta + \Delta_1^j, \mathbf{y}))}{\sigma_2} Z_2^j$, where Δ_1^j, Δ_2^j are as defined in Algorithm 4. Then:

1. $\hat{\mathbf{g}}_1$ is an unbiased estimate of the gradient of $\psi_{\sigma_1, \sigma_2}$ at $\ell(\theta)$, i.e., $\mathbf{E}[\hat{\mathbf{g}}_1] = \nabla \psi_{\sigma_1, \sigma_2}(\ell(\theta))$.
2. $\psi_{\sigma_1}(\cdot)$ is smooth with smoothness parameter $\frac{\sqrt{KL}}{\sigma_1}$ and Lipschitz with constant L .
3. $\mathbf{E} [\|\hat{\mathbf{g}}_1 - \mathbf{E}[\hat{\mathbf{g}}_1]\|^2] \leq \frac{CL^2K}{m} \left(\sqrt{\frac{\sigma_2}{\sigma_1}} K + \log K + 1 \right)$ for some constant C .
4. $\|\nabla \psi_{\sigma_1, \sigma_2}(\ell(\theta)) - \nabla \psi_{\sigma_1}(\ell(\theta))\| \leq \frac{\sigma_2}{2} \frac{\sqrt{KL}}{\sigma_1} (K + 3)^{\frac{3}{2}}$.

Proof. Part 1 follows by trivially observing

$$\mathbf{E}_{Z_1, Z_2} [\hat{\mathbf{g}}_1] = \mathbf{E}_{Z_2} \left[\frac{\psi_{\sigma_1}(\mathbf{u} + \sigma_2 Z_2) - \psi_{\sigma_1}(\mathbf{u})}{\sigma_2} Z_2 \right] = \nabla \psi_{\sigma_1, \sigma_2}(\mathbf{u})$$

where we invoked part 1 of Lemma 8. See Lemma 2 of Nesterov and Spokoiny (2017) for part 2. Part 2 together with Lemma 2 in Duchi et al. (2015) give the result in part 3. See Lemma 3 of Nesterov and Spokoiny (2017) for part 4. \square

Now we are ready to bound the MSE in gradient estimate.

Proof of Lemma 11. We can write out the gradient estimate as:

$$\begin{aligned}
 \hat{\mathbf{g}} &= \frac{1}{m} \sum_{j=1}^m \frac{M(\mathbf{f}_\theta + \Delta_2^j, \mathbf{y}) - M(\mathbf{f}_\theta + \Delta_1^j, \mathbf{y})}{\sigma_2} Z_2^j \\
 &= \frac{1}{m} \sum_{j=1}^m \frac{\psi(\ell(\mathbf{f}_\theta + \Delta_2^j, \mathbf{y})) - \psi(\ell(\mathbf{f}_\theta + \Delta_1^j, \mathbf{y}))}{\sigma_2} Z_2^j + \frac{1}{m} \sum_{j=1}^m \frac{\epsilon(\mathbf{f}_\theta + \Delta_2^j, \mathbf{y}) - \epsilon(\mathbf{f}_\theta + \Delta_1^j, \mathbf{y})}{\sigma_2} Z_2^j \\
 &= \frac{1}{m} \sum_{j=1}^m \frac{\psi(\ell(\theta) + \sigma_1 Z_1^j + \sigma_2 Z_2^j) - \psi(\ell(\theta) + \sigma_1 Z_1^j)}{\sigma_2} Z_2^j + \frac{1}{m} \sum_{j=1}^m \frac{\epsilon(\mathbf{f}_\theta + \Delta_2^j, \mathbf{y}) - \epsilon(\mathbf{f}_\theta + \Delta_1^j, \mathbf{y})}{\sigma_2} Z_2^j \\
 &:= \hat{\mathbf{g}}_1 + \hat{\mathbf{g}}_2,
 \end{aligned}$$

where $\epsilon(\mathbf{f}_\theta, \mathbf{y})$ is the unknown slack function in Section 3.1, re-written in terms of the scores \mathbf{f}_θ and labels \mathbf{y} .

Let ψ_{σ_1} and $\psi_{\sigma_1, \sigma_2}$ be defined as in Lemma 12. Then the gradient estimate error can be expanded as:

$$\begin{aligned}
 \mathbf{E} [\|\hat{\mathbf{g}} - \nabla \psi_{\sigma_1}(\ell(\theta))\|^2] &\leq 2\mathbf{E} [\|\hat{\mathbf{g}}_1 - \nabla \psi_{\sigma_1, \sigma_2}(\ell(\theta))\|^2] + 2\|\nabla \psi_{\sigma_1, \sigma_2}(\ell(\theta)) - \nabla \psi_{\sigma_1}(\ell(\theta))\|^2 \\
 &\leq 4\mathbf{E} [\|\hat{\mathbf{g}}_1 - \nabla \psi_{\sigma_1, \sigma_2}(\ell(\theta))\|^2] + 4\mathbf{E} [\|\hat{\mathbf{g}}_2\|^2] + 2\|\nabla \psi_{\sigma_1, \sigma_2}(\ell(\theta)) - \nabla \psi_{\sigma_1}(\ell(\theta))\|^2 \\
 &\leq 4\mathbf{E} [\|\hat{\mathbf{g}}_1 - \nabla \psi_{\sigma_1, \sigma_2}(\ell(\theta))\|^2] + \frac{16\bar{\epsilon}^2}{m\sigma_2^2} \sum_{j=1}^m \mathbf{E} [\|Z_2^j\|^2] + 2\|\nabla \psi_{\sigma_1, \sigma_2}(\ell(\theta)) - \nabla \psi_{\sigma_1}(\ell(\theta))\|^2 \\
 &\leq \frac{CL^2K}{m} \left(\sqrt{\frac{\sigma_2}{\sigma_1}} K + \log K + 1 \right) + \frac{16\bar{\epsilon}^2K}{\sigma_2^2} + \frac{\sigma_2^2 KL^2}{2\sigma_1^2} (K+3)^3,
 \end{aligned}$$

where we used that (1) $\hat{\mathbf{g}}_1$ is an unbiased estimate of $\nabla \psi_{\sigma_1, \sigma_2}(\ell(\theta))$ (see part 1 of Lemma 12); (2) boundness assumption $|\epsilon(\theta)| \leq \bar{\epsilon}$; (3) $\|a_1 + \dots + a_m\|^2 \leq m(\|a_1\|^2 + \dots + \|a_m\|^2)$, and the last step follows from Parts 3–4 of Lemma 12.

Setting $\sigma_2 = \sqrt{\frac{\sigma_1}{K^{3/2}L}}$ completes the proof. \square

B.2. Proof of Corollary 2

Proof. We begin by observing that convolution operation preserves monotonicity, convexity, and range of the function. Let $g_{\sigma_1}(\cdot)$ denotes Gaussian density function with variance σ_1^2 , since $\psi_{\sigma_1}(\mathbf{u})$ is a positively-weighted linear combination of shifted $\psi(\cdot)$, i.e.,

$$\psi_{\sigma_1}(\mathbf{u}) = \int_{\mathbb{R}^K} \psi(\mathbf{u} - \mathbf{z}) \cdot g_{\sigma_1}(\mathbf{z}) d\mathbf{z} = \int_{\mathbb{R}^K} \psi(\mathbf{z}) \cdot g_{\sigma_1}(\mathbf{u} - \mathbf{z}) d\mathbf{z},$$

Lipschitz property and convexity follows immediately from those on $\psi(\cdot)$. Moreover, since g_{σ_1} is a probability distribution, we always have $\max |\psi_{\sigma_1}(\mathbf{u})| \leq \max |\psi(\mathbf{u})|$. Taking derivatives, we have if $\psi(\cdot)$ is monotonic,

$$\frac{\partial \psi_{\sigma_1}(\mathbf{u})}{\partial u_i} = \nabla \psi_{\sigma_1}(\mathbf{u})^\top \mathbf{e}_i = \int_{\mathbb{R}^K} \nabla \psi(\mathbf{z})^\top \mathbf{e}_i \cdot g_{\sigma_1}(\mathbf{u} - \mathbf{z}) d\mathbf{z} > 0$$

therefore $\psi_{\sigma_1}(\cdot)$ is also monotonic. Moreover, from Lemma 12 we know $\psi_{\sigma_1}(\cdot)$ is smooth with parameter $\beta = \frac{\sqrt{KL}}{\sigma_1}$ and is L -Lipschitz, and that the mean-squared-error in gradient estimate $\hat{\mathbf{g}}$ is bounded by $\kappa = \tilde{O}\left(\frac{L^{7/4}K^{13/8}}{m\sigma_1^{1/4}} + \frac{LK^{5/2}\bar{\epsilon}^2}{\sigma_1}\right)$ from Lemma 11. Applying Theorem 2 on the smoothed metric $\psi_{\sigma_1}(\cdot)$ with $\eta = \frac{1}{\beta^2} = \frac{\sigma_1^2}{KL^2}$ then completes the proof. \square

C. Surrogate PGD as Optimizing a Linear Combination of Surrogates

In this section, we provide an interpretation of Algorithm 1 as optimizing an *adaptively chosen* linear combination of the surrogates $\ell(\theta)$ with an additional proximal penalty like term. Recall that Step 6 of the surrogate projected gradient descent algorithm in Algorithm 1 solves the following optimization problem:

$$\theta^{t+1} \in \operatorname{argmin}_{\theta \in \mathbb{R}^d} \|\ell(\theta) - \tilde{\mathbf{u}}^{t+1}\|_+^2. \quad (23)$$

Lemma 13. *The optimization problem in (23) is equivalent to:*

$$\theta^{t+1} \in \operatorname{argmin}_{\theta \in \mathbb{R}^d} \langle \hat{\mathbf{g}}^t, \ell(\theta) \rangle + \mathbb{D}(\theta, \theta^t),$$

where $\mathbb{D}(\theta, \theta^t) = \frac{1}{2\eta} \|\ell(\theta) - \ell(\theta^t)\|^2 + \frac{1}{2\eta} \|(\ell(\theta) - \ell(\theta^t) + \eta \hat{\mathbf{g}}^t)_+\|^2 - \frac{1}{2\eta} \|(\ell(\theta^t) - \eta \hat{\mathbf{g}}^t - \ell(\theta))_+\|^2$.

Thus (23) can be seen as minimizing a sum of linear combination of the surrogates and (roughly speaking) a term penalizing some form of distance between the current iterate θ^{t+1} and the previous iterate θ^t .

Proof. Expanding the optimization problem in (23):

$$\theta^{t+1} \in \operatorname{argmin}_{\theta \in \mathbb{R}^d} \|(\ell(\theta) - (\ell(\theta^t) - \eta \hat{\mathbf{g}}^t))_+\|^2.$$

Using the identity $(x)_+ = \frac{x+|x|}{2}$, we can write the objective in the above problem as

$$\begin{aligned} & \frac{1}{4} \left\| \ell(\theta) - \ell(\theta^t) + \eta \hat{\mathbf{g}}^t + |\ell(\theta) - \ell(\theta^t) + \eta \hat{\mathbf{g}}^t| \right\|^2 \\ &= \frac{1}{2} \left\| \ell(\theta) - \ell(\theta^t) + \eta \hat{\mathbf{g}}^t \right\|^2 + \frac{1}{2} \left\langle \ell(\theta) - \ell(\theta^t) + \eta \hat{\mathbf{g}}^t, |\ell(\theta) - \ell(\theta^t) + \eta \hat{\mathbf{g}}^t| \right\rangle \end{aligned}$$

which by ignoring constant terms and noticing that the second term is positive for the coordinates for which $\ell_k(\theta) > \ell_k(\theta^t) - \eta \hat{\mathbf{g}}_k^t$ and negative otherwise, we have that

$$\theta^{t+1} \in \operatorname{argmin}_{\theta \in \mathbb{R}^d} \langle \hat{\mathbf{g}}^t, \ell(\theta) \rangle + \frac{1}{2\eta} \left\| \ell(\theta) - \ell(\theta^t) \right\|^2 + \frac{1}{2\eta} \left\| (\ell(\theta) - \ell(\theta^t) + \eta \hat{\mathbf{g}}^t)_+ \right\|^2 - \frac{1}{2\eta} \left\| (\ell(\theta^t) - \eta \hat{\mathbf{g}}^t - \ell(\theta))_+ \right\|^2,$$

as desired. \square

D. Additional Experimental Details

D.1. Choice of Hyper-parameter

For the inner projection step in Algorithm 1, we run Adagrad with a fixed step-size of 1.0 for 100 iterations. We used Adagrad as the optimization method for each of the baselines (including logistic regression, and the Relaxed F-measure approach and the Generalized Rates approach in Section 6.2). We tuned the hyper-parameters such as the step size η for the proposed surrogate PGD algorithm and for the baseline Adagrad solvers, and the perturbation parameter σ for gradient estimation in Algorithm 3 using a held-out validation set.

For the F-measure experiments in Section 6.2, we chose the step sizes from the range $\{0.05, 0.1, 0.5, 1.0, 5.0\}$ and σ from the range $\{0.05, 0.1, 0.5\}$. For the ranking experiments in Section 6.3, we chose the step sizes from $\{0.001, 0.005, 0.01\}$ and found a fixed σ of 1.5 to work well across all runs. For the proxy label experiments in Section 6.4, we chose the step sizes from the range $\{0.01, 0.05, 0.1, 0.5, 1.0\}$ and σ from the range $\{0.01, 0.05, 0.1, 0.5, 1.0\}$.

D.2. Choice of Number of Perturbations

In our experiments, we chose to use 1000 perturbations to estimate gradients as this was a sufficiently large number that worked well for all experiments. But for many experiments, we could get comparable results with fewer perturbations. For example for the experiments in Sec 6.1, with as few as 10 perturbations, our approach achieved a test G-mean of 0.801, a comparable value to what we report in Table 2 for the proposed method (0.803). Similarly, for the macro F-measure experiments in Table 3, we got comparable results with 10 perturbations, as shown in Table 6. For the larger KDD Cup 2008 dataset in the ranking experiments in Section 6.3, we used minibatches of size 100 and only perturb examples within each batch to estimate the gradients.

D.3. Dependence of the Gradient Estimation Error on K

While the error bound for the linear interpolation based gradient estimation approach in Lemma 4 has a strong dependence on the number of surrogates K , we find that in our simulations, the dependence is less severe. This is evident from the

Table 6. Average test macro F-measure across groups with clean features. *Higher* is better. We compare the results for the proposed method with 10 and 1000 perturbations to estimate gradients.

	#perturbations = 10	#perturbations = 1000
Business	0.796	0.796
COMPAS	0.630	0.629
Adult	0.661	0.665
Default	0.532	0.533




Figure 5. Mean squared estimation error for gradients estimated by the local linear interpolation approach in Algorithm 3 for a synthetic K -dimensional gradient estimation problem, as K varies.

plot shown in Figure 5, where we consider the toy problem of estimating the gradient of the function $f(z) = \left(\prod_{k=1}^K z_k\right)^{1/K}$, where $z \in \mathbb{R}_+^K$, and we draw each coordinate z_k from $0.1 + \text{Unif}(0, 0.9)$. We use the local linear interpolation based approach in Algorithm 3 to estimate gradients for f and evaluate the mean squared error for the gradient estimates w.r.t. the true gradient of f as K varies. We use 100 perturbations, and report the average estimation errors over 100 random draws of z and over 100 random trials for each draw of z .