
Quadratic Metric Elicitation for Fairness and Beyond (Supplementary material)

Gaurush Hiranandani¹

Jatin Mathur²

Harikrishna Narasimhan³

Oluwasanmi Koyejo^{2,3}

¹Amazon, Palo Alto, California, USA

²University of Illinois at Urbana-Champaign, Champaign, Illinois, USA

³Google Research, Mountain View, California, USA

A LINEAR PERFORMANCE METRIC ELICITATION (LPME)

In this section, we shed more light on the procedure from [Hiranandani et al., 2019b] that elicits a multiclass linear metric. We call it the Linear Performance Metric Elicitation (LPME) procedure. As discussed in Algorithm 1, we use this as a subroutine to elicit metrics in the quadratic family.

LPME exploits the enclosed sphere $\mathcal{S} \subset \mathcal{R}$ for eliciting linear multiclass metrics. Let the sphere \mathcal{S} 's radius be $\rho > 0$, and the oracle's scale invariant metric be $\phi^{\text{lin}}(\mathbf{r}) := \langle \mathbf{a}, \mathbf{r} \rangle$ such that $\|\mathbf{a}\|_2 = 1$. The oracle queries are $\Omega(\mathbf{r}_1, \mathbf{r}_2; \phi^{\text{lin}}) := \mathbb{1}[\phi^{\text{lin}}(\mathbf{r}_1) > \phi^{\text{lin}}(\mathbf{r}_2)]$. We first outline a trivial Lemma from [Hiranandani et al., 2019b].

Lemma 1. [Hiranandani et al., 2019b] *Let a normalized vector \mathbf{a} with $\|\mathbf{a}\|_2 = 1$ parametrize a linear metric $\phi^{\text{lin}} := \langle \mathbf{a}, \mathbf{r} \rangle$, then the unique optimal rate $\bar{\mathbf{r}}$ over \mathcal{S} is a rate on the boundary of \mathcal{S} given by $\bar{\mathbf{r}} = \rho\mathbf{a} + \mathbf{o}$, where \mathbf{o} is the center of \mathcal{S} .*

Algorithm 2 Linear Performance Metric Elicitation

1: **Input:** Query space $\mathcal{S} \subset \mathcal{R}$, binary-search tolerance $\epsilon > 0$, oracle $\Omega(\cdot, \cdot; \phi^{\text{lin}})$ with metric ϕ^{lin}

2: **for** $i = 1, 2, \dots, k$ **do**

3: Set $\mathbf{a} = \mathbf{a}' = (1/\sqrt{k}, \dots, 1/\sqrt{k})$.

4: Set $a'_i = -1/\sqrt{k}$.

5: Compute the optimal $\bar{\mathbf{s}}^{(\mathbf{a})}$ and $\bar{\mathbf{s}}^{(\mathbf{a}'')}$ over the sphere \mathcal{S} using Lemma 1

6: Query $\Omega(\bar{\mathbf{s}}^{(\mathbf{a})}, \bar{\mathbf{s}}^{(\mathbf{a}'')}$; $\phi^{\text{lin}})$
 {These queries reveal the search orthant}

7: Start with coordinate $j = 1$.

8: **Initialize:** $\boldsymbol{\theta} = \boldsymbol{\theta}^{(1)}$ { $\boldsymbol{\theta}^{(1)}$ is a point in the search orthant.}

9: **for** $t = 1, 2, \dots, T = 3(k-1)$ **do**

10: Set $\boldsymbol{\theta}^{(a)} = \boldsymbol{\theta}^{(c)} = \boldsymbol{\theta}^{(d)} = \boldsymbol{\theta}^{(e)} = \boldsymbol{\theta}^{(b)} = \boldsymbol{\theta}^{(t)}$.

11: Set $\theta_j^{(a)}$ and $\theta_j^{(b)}$ to be the min and max angle, respectively, based on the search orthant

12: **while** $|\theta_j^{(b)} - \theta_j^{(a)}| > \epsilon$ **do**

13: Set $\theta_j^{(c)} = \frac{3\theta_j^{(a)} + \theta_j^{(b)}}{4}$, $\theta_j^{(d)} = \frac{\theta_j^{(a)} + \theta_j^{(b)}}{2}$, and $\theta_j^{(e)} = \frac{\theta_j^{(a)} + 3\theta_j^{(b)}}{4}$.

14: Set $\bar{\mathbf{r}}^{(a)} = \mu(\boldsymbol{\theta}^{(a)})$ (i.e. parametrization of $\partial\mathcal{S}$). Similarly, set $\bar{\mathbf{r}}^{(c)}$, $\bar{\mathbf{r}}^{(d)}$, $\bar{\mathbf{r}}^{(e)}$, $\bar{\mathbf{r}}^{(b)}$

15: $[\theta_j^{(a)}, \theta_j^{(b)}] \leftarrow \text{ShrinkInterval}(\Omega, \bar{\mathbf{r}}^{(a)}, \bar{\mathbf{r}}^{(c)}, \bar{\mathbf{r}}^{(d)}, \bar{\mathbf{r}}^{(e)}, \bar{\mathbf{r}}^{(b)})$ {see Figure 1}

16: Set $\theta_j^{(d)} = \frac{1}{2}(\theta_j^{(a)} + \theta_j^{(b)})$

17: Set $\boldsymbol{\theta}^{(t)} = \boldsymbol{\theta}^{(d)}$.

18: Update coordinate $j \leftarrow j + 1$ cyclically.

19: **Output:** $\hat{a}_i = \prod_{j=1}^{i-1} \sin \theta_j^{(T)} \cos \theta_i^{(T)} \forall i \in [k-1]$, $\hat{a}_k = \prod_{j=1}^{k-1} \sin \theta_j^{(T)}$

Lemma 1 provides a way to define a one-to-one correspondence between a linear performance metric and its optimal rate over

Subroutine *ShrinkInterval*

Input: Oracle Ω and rate profiles $\bar{\mathbf{r}}^{(a)}, \bar{\mathbf{r}}^{(c)}, \bar{\mathbf{r}}^{(d)}, \bar{\mathbf{r}}^{(e)}, \bar{\mathbf{r}}^{(b)}$
 Query $\Omega(\bar{\mathbf{r}}^{(a)}, \bar{\mathbf{r}}^{(c)}; \phi^{\text{lin}})$.
If $(\bar{\mathbf{r}}^{(a)} \succ \bar{\mathbf{r}}^{(c)})$ Set $\theta_j^{(b)} = \theta_j^{(d)}$.
else Query $\Omega(\bar{\mathbf{r}}^{(c)}, \bar{\mathbf{r}}^{(d)}; \phi^{\text{lin}})$.
If $(\bar{\mathbf{r}}^{(c)} \succ \bar{\mathbf{r}}^{(d)})$ Set $\theta_j^{(b)} = \theta_j^{(d)}$.
else Query $\Omega(\bar{\mathbf{r}}^{(d)}, \bar{\mathbf{r}}^{(e)}; \phi^{\text{lin}})$.
If $(\bar{\mathbf{r}}^{(d)} \succ \bar{\mathbf{r}}^{(e)})$ Set $\theta_j^{(a)} = \theta_j^{(c)}$ and $\theta_j^{(b)} = \theta_j^{(e)}$.
else Set $\theta_j^{(a)} = \theta_j^{(c)}$.
Output: $[\theta_j^{(a)}, \theta_j^{(b)}]$.




Figure 1: (Left): The *ShrinkInterval* subroutine used in line 16 of Algorithm 2 (Right): Visual illustration of the subroutine *ShrinkInterval* [Hiranandani et al., 2019b]; *ShrinkInterval* shrinks the current search interval to half based on oracle responses to at most three queries.

the sphere. That is, given a linear performance metric, using Lemma 1, we may get a unique point in the query space lying on the boundary of the sphere $\partial\mathcal{S}$. Moreover, the converse is also true; i.e., given a feasible rate on the boundary of the sphere $\partial\mathcal{S}$, one may recover the linear metric for which the given rate is optimal. Thus, for eliciting a linear metric, Hiranandani et al. [2019b] essentially search for the optimal rate (over the sphere \mathcal{S}) using pairwise queries to the oracle. The optimal rate by virtue of Lemma 1 reveals the true metric. The LPME subroutine is summarized in Algorithm 2. Intuitively, Algorithm 2 minimizes a strongly convex function denoting distance of query points from a supporting hyperplane whose slope is the true metric (see Figure 2(c) in [Hiranandani et al., 2019b]). The procedure also uses the following standard parameterization for the surface of the sphere $\partial\mathcal{S}$.

Parameterizing the boundary of the enclosed sphere $\partial\mathcal{S}$. Let θ be a $(k-1)$ -dimensional vector of angles. In θ , all the angles except the primary angle are in $[0, \pi]$, and the primary angle is in $[0, 2\pi]$. A scale invariant linear performance metric with $\|\mathbf{a}\|_2 = 1$ can be constructed by assigning $a_i = \prod_{j=1}^{i-1} \sin \theta_j \cos \theta_i$ for $i \in [k-1]$ and $a_k = \prod_{j=1}^{k-1} \sin \theta_j$. Since we can easily compute the metric’s optimal rate over \mathcal{S} using Lemma 1, by varying θ in this procedure, we parametrize the surface of the sphere $\partial\mathcal{S}$. We denote this parametrization by $\mu(\theta)$, where $\mu: [0, \pi]^{k-2} \times [0, 2\pi] \rightarrow \partial\mathcal{S}$.

Description of Algorithm 2: Let the oracle’s metric be $\phi^{\text{lin}} = \langle \mathbf{a}, \mathbf{r} \rangle$ such that $\|\mathbf{a}\|_2 = 1$ (Section 2.2). Using the parametrization $\mu(\theta)$ for the boundary of the sphere $\partial\mathcal{S}$, Algorithm 2 returns an estimate $\hat{\mathbf{a}}$ with $\|\hat{\mathbf{a}}\|_2 = 1$. Line 2-6 recover the search orthant of the optimal rate over the sphere by posing k trivial queries. Once the search orthant of the optimal rate is known, the algorithm in each iteration of the for loop (line 9-18) updates one angle θ_j keeping other angles fixed by the *ShrinkInterval* subroutine. The *ShrinkInterval* subroutine (illustrated in Figure 1) is binary-search based routine that shrinks the interval $[\theta_j^a, \theta_j^b]$ by half based on the oracle responses to (at most) three queries.¹ Note that, depending on the oracle responses, one may reduce the search interval to half using less than three queries in some cases. Then the algorithm cyclically updates each angle until it converges to a metric sufficiently close to the true metric. We fix the number of cycles in coordinate-wise binary search to three. Therefore, in order to elicit a linear performance metric in k dimensions, the LPME subroutine requires at most $3 \times 3 \times k \log(\pi/2\epsilon)$ queries, where three is the number of cycles in coordinate wise binary search, three is the (maximum) number of queries to shrink the search interval into half, and the initial search interval for the angles is $\pi/2$.

B GEOMETRY OF THE FEASIBLE SPACE (PROOFS OF SECTION 2 AND SECTION 4)

Proof of Proposition 1 and Proposition 2. We prove Proposition 2. The proof of Proposition 1 is analogous where the probability measures (corresponding to classifiers and their rates) are not conditioned on any group.

The group-specific set of rates \mathcal{R}^g for a group g has the following properties [Hiranandani et al., 2020]:

- *Convex:* Consider two classifiers $h_1^g, h_2^g \in \mathcal{H}^g$ that achieve the rates $\mathbf{r}_1^g, \mathbf{r}_2^g \in \mathcal{R}^g$, respectively. Also, consider a classifier h^g that predicts what classifier h_1^g predicts with probability γ and predicts what classifier h_2^g predicts with probability $1 - \gamma$. Then the rate vector of the classifier h^g is:

¹The description of the binary search algorithm in Hiranandani et al. [2019b] always assumes getting responses to four queries that essentially correspond to the four intervals. In practice, the binary search can be adaptive and may only require at most three queries as we have discussed in this paper. The order of the queries in LPME, however, remains the same.

Algorithm 3 Obtaining the sphere $\mathcal{S} \subset \mathcal{R}$ (Figure 2(a)) of radius ρ centered at \mathbf{o}

- 1: **for** $j = 1, 2, \dots, k$ **do**
 - 2: Let α_j be the standard basis vector.
 - 3: Compute the maximum constant c_j such that $\mathbf{o} + c_j \alpha_j$ is feasible by solving (OP1).
 - 4: Let $CONV$ denote the convex hull of $\{\mathbf{o} \pm c_j \alpha_j\}_{j=1}^k$. It will be centered at \mathbf{o} .
 - 5: Compute the radius ρ of the largest ball that fits in $CONV$.
 - 6: **Output:** Sphere \mathcal{S} with radius ρ centered at \mathbf{o} .
-

$$\begin{aligned}
R_{ij}^g(h) &= \mathbb{P}(h^g = i | Y = i) \\
&= \mathbb{P}(h_1^g = i | h^g = h_1^g, Y = i) \mathbb{P}(h^g = h_1^g) + \mathbb{P}(h_2^g = i | h^g = h_2^g, Y = i) \mathbb{P}(h^g = h_2^g) \\
&= \gamma \mathbf{r}_1^g + (1 - \gamma) \mathbf{r}_2^g.
\end{aligned}$$

The above equations show that the convex combination of any two rates is feasible as well, i.e., one can construct a randomized classifier which will achieve the convex combination of rates. Hence, $\mathcal{R}^g \forall g \in [m]$ is convex. Since intersection of convex sets is convex, the intersection set $\mathcal{R}^1 \cap \dots \cap \mathcal{R}^m$ is convex as well.

- *Bounded:* Since $R_{ij}^g(h) = P[h = i | Y = i] \leq 1$ for all $i \in [k]$, $\mathcal{R}^g \subseteq [0, 1]^k$.
- *The rates \mathbf{o} and \mathbf{e}_i 's are always achievable:* A uniform random classifier, i.e, the classifier, which for any input, predicts all classes with probability $1/k$ achieves the rate profile \mathbf{o} . A classifier that always predicts class i achieves the rate \mathbf{e}_i . Thus, $\mathbf{e}_i \in \mathcal{R}^g \forall i \in [k], g \in [m]$ are always feasible.
- *\mathbf{e}_i 's are vertices:* Consider the supporting hyperplanes with the following slope: $a_i > a_j > 0$ and $a_l = 0$ for $l \in [k], l \neq i, j$. These hyperplanes will be supported by \mathbf{e}_i . Thus, \mathbf{e}_i 's are vertices of the convex set \mathcal{R}^g . From Assumption 1, one can construct a ball around the trivial rate \mathbf{o} and thus \mathbf{o} lies in the interior.

The above points apply to space of overall rates \mathcal{R} as well; thus, proving Proposition 1. □

B.1 FINDING THE SPHERE $\mathcal{S} \subset \mathcal{R}$

In this section, we provide details regarding how a sphere \mathcal{S} with sufficiently large radius ρ inside the feasible region \mathcal{R} may be found (see Figure 2(a)). The following discussion is borrowed from [Hiranandani et al., 2019b] and provided here for completeness.

The following optimization problem is a special case of OP2 in [Narasimhan, 2018]. The problem is associated with a feasibility check problem. Given a rate profile \mathbf{r}_0 , the optimization routine tries to construct a classifier that achieves the rate \mathbf{r}_0 within small error $\epsilon > 0$.

$$\min_{\mathbf{r} \in \mathcal{R}} 0 \quad s.t. \quad \|\mathbf{r} - \mathbf{r}_0\|_2 \leq \epsilon. \tag{OP1}$$

The above optimization problem checks the feasibility, and if there exists a solution to the above problem, then Algorithm 1 of [Narasimhan, 2018] returns it. Furthermore, Algorithm 3 computes a value of $\rho \geq \tilde{\rho}/k$, where $\tilde{\rho}$ is the radius of the largest ball contained in the set \mathcal{R} . Also, the approach in [Narasimhan, 2018] is consistent, thus we should get a good estimate of the sphere, provided we have sufficiently large number of samples. The algorithm is completely offline and does not impact oracle query complexity.

Lemma 2. [Hiranandani et al., 2019b] *Let $\tilde{\rho}$ denote the radius of the largest ball in \mathcal{R} centered at \mathbf{o} . Then Algorithm 3 returns a sphere with radius $\rho \geq \tilde{\rho}/k$, where k is the number of classes.*

The idea in Algorithm 3 can be trivially extended to finding a sphere $\bar{\mathcal{S}} \subset \mathcal{R}^1 \cap \dots \cap \mathcal{R}^m$ corresponding to Remark 3.

C QUADRATIC PERFORMANCE METRIC ELICITATION PROCEDURE

In this section, we describe how the subroutine calls to LPME in Algorithm 1 elicit a quadratic metric in Definition 3. We start with the shifted metric of Equation (5). Also, as explained in the main paper, we may assume $d_1 \neq 0$ due to Assumption 2. We can derive the following solution using any non-zero coordinate of \mathbf{d} , instead of d_1 . We can identify a non-zero coordinate using k trivial queries of the form $(\rho \alpha_i + \mathbf{o}, \mathbf{o}), \forall i \in [k]$.

1. From line 1 of Algorithm 1, we get local linear approximation at \mathbf{o} . Using Remark 2, we have (6) which is

$$d_i = \frac{f_{i0}}{f_{10}} d_1 \quad \forall i \in \{2, \dots, k\}. \quad (1)$$

2. Similarly, if we apply LPME on small balls around rate profiles \mathbf{z}_j , Remark 2 gives us:

$$\frac{d_i + (\rho - \varrho)B_{ij}}{d_1 + (\rho - \varrho)B_{1j}} = \frac{f_{ij}}{f_{1j}} \quad \forall i \in \{2, \dots, k\}, j \leq i. \quad (2)$$

$$\begin{aligned} \implies d_i + (\rho - \varrho)B_{ij} &= \frac{f_{ij}}{f_{1j}} (d_1 + (\rho - \varrho)B_{1j}) \\ \implies (\rho - \varrho)B_{ij} &= \frac{f_{ij}}{f_{1j}} (d_1 + (\rho - \varrho)B_{j1}) - d_i \\ \implies (\rho - \varrho)B_{ij} &= \frac{f_{ij}}{f_{1j}} (d_1 + \frac{f_{j1}}{f_{11}} (d_1 + (\rho - \varrho)B_{11}) - d_j) - \frac{f_{i0}}{f_{10}} d_1 \\ \implies (\rho - \varrho)B_{ij} &= \left(\frac{f_{ij}}{f_{1j}} - \frac{f_{i0}}{f_{10}} + \frac{f_{ij}}{f_{1j}} \left(\frac{f_{j1}}{f_{11}} - \frac{f_{j0}}{f_{10}} \right) \right) d_1 + (\rho - \varrho) \frac{f_{j1}}{f_{11}} B_{11}, \end{aligned} \quad (3)$$

where we have used that the matrix \mathbf{B} is symmetric in the second step, and (1) in the last two steps. We can represent each element in terms of B_{11} and d_1 . So, a relation between B_{11} and d_1 may allow us to represent each element of \mathbf{a} and \mathbf{B} in terms of d_1 .

3. Therefore, by applying LPME on small balls around rate profiles $-\mathbf{z}_1$, Remark 2 gives us (8):

$$\frac{d_2 - (\rho - \varrho)B_{21}}{d_1 - (\rho - \varrho)B_{11}} = \frac{f_{21}^-}{f_{11}^-}. \quad (4)$$

4. Using (2) and (4), we have:

$$(\rho - \varrho)B_{11} = \frac{\frac{f_{21}^-}{f_{11}^-} + \frac{f_{21}}{f_{11}} - 2\frac{f_{20}}{f_{10}}}{\frac{f_{21}^-}{f_{11}^-} - \frac{f_{21}}{f_{11}}} d_1. \quad (5)$$

Putting (5) in (3), we get:

$$\begin{aligned} B_{ij} &= \left[\frac{f_{ij}}{f_{1j}} \left(1 + \frac{f_{j1}}{f_{11}} \right) - \frac{f_{ij}}{f_{1j}} \frac{f_{j0}}{f_{10}} - \frac{f_{i0}}{f_{10}} + \frac{f_{ij}}{f_{1j}} \frac{f_{j1}}{f_{11}} \frac{\frac{f_{21}^-}{f_{11}^-} + \frac{f_{21}}{f_{11}} - 2\frac{f_{20}}{f_{10}}}{\frac{f_{21}^-}{f_{11}^-} - \frac{f_{21}}{f_{11}}} \right] d_1 \\ &= \left(F_{i,1,j} (1 + F_{j,1,1}) - F_{i,1,j} F_{j,1,0} - F_{i,1,0} + F_{i,1,j} \frac{F_{2,1,1}^- + F_{2,1,1} - 2F_{2,1,0}}{F_{2,1,1}^- - F_{2,1,1}} \right) d_1, \end{aligned} \quad (6)$$

where $F_{i,j,l} = \frac{f_{il}}{f_{jl}}$ and $F_{i,j,l}^- = \frac{f_{il}^-}{f_{jl}^-}$. As $\mathbf{a} = \mathbf{d} + \mathbf{B}\mathbf{o}$, we can represent each element of \mathbf{a} and \mathbf{B} using using (1) and (6) in terms of d_1 . We can then use the normalization condition $\|\mathbf{a}\|_2^2 + \|\mathbf{B}\|_F^2 = 1$ to get estimates of \mathbf{a} , \mathbf{B} which are independent of d_1 .

This completes the derivation of solution from QPME (section 3).

D FAIR (QUADRATIC) PERFORMANCE METRIC ELICITATION PROCEDURE

Algorithm 4: FPM Elicitation

Input: Query set \mathcal{S}' , search tolerance $\epsilon > 0$, oracle Ω'

1. Let $\mathcal{L} \leftarrow \emptyset$
 - 2: **For** $\sigma \in \mathcal{M}$ **do**
 - 3: $\beta^\sigma \leftarrow \text{QPME}(\mathcal{S}', \epsilon, \Omega')$
 - 4: Let ℓ^σ be Eq. (10), extend $\mathcal{L} \leftarrow \mathcal{L} \cup \{\ell^\sigma\}$
 - 5: $\hat{\mathbb{B}} \leftarrow$ normalized solution from (13) using \mathcal{L}
 - 6: $\hat{\lambda} \leftarrow$ trace back normalized solution from (10) for any σ
- Output:** $\hat{\mathbf{a}}, \hat{\mathbb{B}}, \hat{\lambda}$

We first discuss eliciting the fair (quadratic) metric in Definition 4, where all the parameters are unknown. We then provide an alternate procedure for eliciting just the trade-off parameter λ when the predictive performance and fairness violation coefficients are known. The latter is a separate application as discussed in [Zhang et al., 2020]. However, unlike Zhang et al. [2020], instead of ratio queries, we use simpler pairwise comparison queries.

In this section, we work with any number of groups $m \geq 2$. The idea, however, remains the same as described in the main paper for number of groups $m = 2$. We specifically select queries from the sphere $\bar{\mathcal{S}} \subset \mathcal{R}^1 \cap \dots \cap \mathcal{R}^m$, which is common to all the group-specific feasible region of rates, so to reduce the problem into multiple instances of the proposed QPME procedure of Section 3.

Suppose that the oracle's fair performance metric is ϕ^{fair} parametrized by $(\mathbf{a}, \mathbb{B}, \lambda)$ as in Definition 4. The overall fair metric elicitation procedure framework is summarized in Algorithm 4. The framework exploits the sphere $\bar{\mathcal{S}} \subset \mathcal{R}^1 \cap \dots \cap \mathcal{R}^m$ and uses the QPME procedure (Algorithm 1) as a subroutine multiple times.

Let us consider a non-empty set of sets $\mathcal{M} \subset 2^{[m]} \setminus \{\emptyset, [m]\}$. We will later discuss how to choose such a set \mathcal{M} . We partition the set of groups $[m]$ into two sets of groups. Let $\sigma \in \mathcal{M}$ and $[m] \setminus \sigma$ be one such partition of the m groups defined by the set of groups σ . For example, when $m = 3$, one may choose the set of groups $\sigma = \{1, 2\}$.

Now, consider a sphere \mathcal{S}' whose elements $\mathbf{r}^{1:m} \in \mathcal{S}'$ are given by:

$$\mathbf{r}^g = \begin{cases} \mathbf{s} & \text{if } g \in \sigma \\ \mathbf{o} & \text{o.w.} \end{cases} \quad (7)$$

This is an extension of the sphere \mathcal{S}' defined in the main paper for the $m > 2$ case. Elements in \mathcal{S}' have rate profiles $\mathbf{s} \in \bar{\mathcal{S}}$ to the groups in σ and trivial rate profile \mathbf{o} to the remaining groups in $[m] \setminus \sigma$. Analogously, the modified oracle is $\Omega'(\mathbf{r}_1, \mathbf{r}_2) = \Omega((\mathbf{r}_1^{1:m}), (\mathbf{r}_2^{1:m}))$, where $\mathbf{r}_1^{1:m}, \mathbf{r}_2^{1:m}$ are the elements of the spheres \mathcal{S}' above. Thus, for elements in \mathcal{S}' , the metric in Definition 4 reduces to:

$$\phi^{\text{fair}}(\mathbf{r}^{1:m} \in \mathcal{S}'; \mathbf{a}, \mathbb{B}, \lambda) = (1 - \lambda) \langle -\mathbf{a} \odot \boldsymbol{\tau}^\sigma, \mathbf{s} - \mathbf{o} \rangle + \lambda \frac{1}{2} (\mathbf{s} - \mathbf{o})^T \mathbf{W}^\sigma (\mathbf{s} - \mathbf{o}) + c^\sigma \quad (8)$$

where $\boldsymbol{\tau}^\sigma = \sum_{g \in \sigma} \boldsymbol{\tau}^g$, $\mathbf{W}^\sigma = \sum_{u \in \sigma, v \in [m] \setminus \sigma} B^{uv}$, and c^σ is a constant not affecting the oracle responses.

The above metric is a particular instance of $\bar{\phi}(\mathbf{s}; \mathbf{d}, \mathbf{B})$ in (4) with $\mathbf{d} := -(1 - \lambda)\mathbf{a} \odot \boldsymbol{\tau}^\sigma$ and $\mathbf{B} := \lambda \mathbf{W}^\sigma$; thus, we apply QPME procedure as a subroutine in Algorithm 4 to elicit the metric in (8).

The only change needed to be made to the algorithm is in line 5, where we need to take into account the changed relationship between \mathbf{d} and \mathbf{a} , and need to separately (not jointly) normalize the linear and quadratic coefficients. With this change, the output of the algorithm directly gives us the required estimates. Specifically, we have from step 1 of Algorithm 1 and (6) an estimate

$$\frac{d_i}{d_1} = \frac{\tau_i^\sigma a_i}{\tau_1^\sigma a_1} = \frac{f_{i0}}{f_{10}} \implies a_i = \frac{f_{i0}}{f_{10}} \frac{\tau_1^\sigma}{\tau_i^\sigma} a_1. \quad (9)$$

Using the normalization condition (i.e., $\|\mathbf{a}\|_2 = 1$), we directly get an estimate $\hat{\mathbf{a}}$ for the linear coefficients. Similarly, steps

2-4 of Algorithm 1 and (9) gives us: $\hat{B}_{ij} =$

$$\begin{aligned} \sum_{u \in \sigma, v \in [m] \setminus \sigma} \tilde{B}_{ij}^{uv} &= \left(F_{i,1,j}^\sigma (1 + F_{j,1,1}^\sigma) - F_{i,1,j}^\sigma F_{j,1,0}^\sigma d_1 - F_{i,1,0}^\sigma + F_{i,1,j}^\sigma \frac{F_{2,1,1}^{-,\sigma} + F_{2,1,1}^\sigma - 2F_{2,1,0}^\sigma}{F_{2,1,1}^{-,\sigma} - F_{2,1,1}^\sigma} \right) \tau_1^{-1} \hat{a}_1 \\ &= \beta^\sigma, \end{aligned} \quad (10)$$

where the above solution is similar to the two group case, but here it is corresponding to a partition of groups defined by σ , and $\tilde{\mathbf{B}}^{uv} := \lambda \mathbf{B}^{uv} / (1 - \lambda)$ is a scaled version of the true (unknown) \mathbf{B}^{uv} . Let equation (10) be denoted by ℓ^σ . Also, let the right hand side term of (10) be denoted by β^σ .

Since we want to elicit $\binom{m}{2}$ fairness violation weight matrices in \mathbb{B} , we require $\binom{m}{2}$ ways of partitioning the groups into two sets so that we construct $\binom{m}{2}$ independent matrix equations similar to (10). Let \mathcal{M} be those set of sets. Thus, running over all the choices of sets of groups $\sigma \in \mathcal{M}$ provides the system of equations $\mathcal{L} := \cup_{\sigma \in \mathcal{M}} \ell^\sigma$ (line 4 in Algorithm 4), which is:

$$\begin{bmatrix} \Xi & 0 & \dots & 0 \\ 0 & \Xi & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & \Xi \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{b}}_{(11)} \\ \tilde{\mathbf{b}}_{(12)} \\ \dots \\ \tilde{\mathbf{b}}_{(kk)} \end{bmatrix} = \begin{bmatrix} \beta_{(11)} \\ \beta_{(12)} \\ \dots \\ \beta_{(kk)} \end{bmatrix}, \quad (11)$$

where $\tilde{\mathbf{b}}_{(ij)} = (\tilde{b}_{ij}^1, \tilde{b}_{ij}^2, \dots, \tilde{b}_{ij}^{\binom{m}{2}})$ and $\gamma_{(ij)} = (\beta_{ij}^1, \beta_{ij}^2, \dots, \beta_{ij}^{\binom{m}{2}})$ are vectorized versions of the ij -th entry across groups for $i, j \in [k]$, and $\Xi \in \{0, 1\}^{\binom{m}{2} \times \binom{m}{2}}$ is a binary full-rank matrix denoting membership of groups in the set σ . For example, when one chooses $\mathcal{M} = \{\{1, 2\}, \{1, 3\}, \{2, 3\}\}$ for $m = 3$, Ξ is given by:

$$\Xi = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}.$$

One may choose any set of sets \mathcal{M} that allows the resulting group membership matrix Ξ to be non-singular. The solution of the system of equations \mathcal{L} is:

$$\begin{bmatrix} \tilde{\mathbf{b}}_{(11)} \\ \tilde{\mathbf{b}}_{(12)} \\ \dots \\ \tilde{\mathbf{b}}_{(kk)} \end{bmatrix} = \begin{bmatrix} \Xi & 0 & \dots & 0 \\ 0 & \Xi & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & \Xi \end{bmatrix}^{(-1)} \begin{bmatrix} \beta_{(11)} \\ \beta_{(12)} \\ \dots \\ \beta_{(kk)} \end{bmatrix}. \quad (12)$$

When all $\tilde{\mathbf{B}}^{uv}$'s are normalized, we have the estimated fairness violation weight matrices as:

$$\hat{\mathbf{B}}^{uv} = \frac{\tilde{\mathbf{B}}^{uv}}{\frac{1}{2} \sum_{u,v=1,v>u}^m \|\tilde{\mathbf{B}}^{uv}\|_F} \quad \text{for } u, v \in [m], v > u. \quad (13)$$

Due to the above normalization, the solution is again independent of the true trade-off λ .

Given estimates \hat{B}_{ij}^{uv} and \hat{a}_1 , we can now additionally estimate the trade-off parameter $\hat{\lambda}$ from ℓ^σ (10) for any $\sigma \in \mathcal{M}$. This completes the fair (quadratic) metric elicitation procedure.

D.1 ELICITING TRADE-OFF λ WHEN (LINEAR) PREDICTIVE PERFORMANCE AND (QUADRATIC) FAIRNESS VIOLATION COEFFICIENTS ARE KNOWN

We now provide an alternate binary search based method similar to [Hiranandani et al. \[2020\]](#) for eliciting the trade-off parameter λ when the linear predictive and quadratic fairness coefficients are already known. This is along similar lines to the application considered by [Zhang et al. \[2020\]](#), but unlike them, instead of ratio queries, we require simpler pairwise queries.

Here, the key insight is to approximate the non-linearity posed by the fairness violation in Definition 4, which then reduces the problem to a one-dimensional binary search. We have:

$$\phi^{\text{fair}}(\mathbf{r}^{1:m}; \mathbf{a}, \mathbb{B}, \lambda) := (1 - \lambda) \langle \mathbf{a}, \mathbf{1} - \mathbf{r} \rangle + \lambda \frac{1}{2} \left(\sum_{u,v=1,v>u}^m (\mathbf{r}^u - \mathbf{r}^v)^T \mathbf{B}^{uv} (\mathbf{r}^u - \mathbf{r}^v) \right). \quad (14)$$

Algorithm 5 Eliciting the trade-off λ when predictive performance and fairness violation are known

- 1: **Input:** Query space $\overline{\mathcal{S}}_{\mathbf{z}_1}'$, binary-search tolerance $\epsilon > 0$, oracle Ω
 - 2: **Initialize:** $\lambda^{(a)} = 0, \lambda^{(b)} = 1$.
 - 3: **while** $|\lambda^{(b)} - \lambda^{(a)}| > \epsilon$ **do**
 - 4: Set $\lambda^{(c)} = \frac{3\lambda^{(a)} + \lambda^{(b)}}{4}, \lambda^{(d)} = \frac{\lambda^{(a)} + \lambda^{(b)}}{2}, \lambda^{(e)} = \frac{\lambda^{(a)} + 3\lambda^{(b)}}{4}$
 - 5: Set $\mathbf{s}^{(a)} = \operatorname{argmax}_{\mathbf{s} \in \overline{\mathcal{S}}_{\mathbf{z}_1}'} \langle -(1 - \lambda^{(a)})\boldsymbol{\tau}^1 \odot \hat{\mathbf{a}} + \lambda^{(a)} \sum_{v=2}^m \hat{\mathbf{B}}^{1v}(\mathbf{z}_1 - \mathbf{o}), \mathbf{s} \rangle$ using Lemma 1
 - 6: Similarly, set $\mathbf{s}^{(c)}, \mathbf{s}^{(d)}, \mathbf{s}^{(e)}, \mathbf{s}^{(b)}$.
 - 7: $[\lambda^{(a)}, \lambda^{(b)}] \leftarrow \text{ShrinkInterval}(\Omega, \mathbf{s}^{(a)}, \mathbf{s}^{(c)}, \mathbf{s}^{(d)}, \mathbf{s}^{(e)}, \mathbf{s}^{(b)})$ using a subroutine analogous to the routine shown in Figure 1.
 - 8: **Output:** $\hat{\lambda} = \frac{\lambda^{(a)} + \lambda^{(b)}}{2}$.
-

To this end, we define a new sphere $\mathcal{S}' = \{(\mathbf{s}, \mathbf{o}, \dots, \mathbf{o}) | \mathbf{s} \in \overline{\mathcal{S}}\}$. The elements in \mathcal{S}' is the set of rate profiles whose first group achieves rates $\mathbf{s} \in \overline{\mathcal{S}}$ and rest of the groups achieve trivial rate \mathbf{o} (corresponding to uniform random classifier). For any element in \mathcal{S}' , the associated discrepancy terms $(\mathbf{r}^u - \mathbf{r}^v) = 0$ for $u, v \neq 1$. Thus for elements in \mathcal{S}' , the metric in Definition 4 reduces to:

$$\phi^{\text{fair}}((\mathbf{s}, \mathbf{o}, \dots, \mathbf{o}); \mathbf{a}, \mathbb{B}, \lambda) = (1 - \lambda) \langle -\boldsymbol{\tau}^1 \odot \mathbf{a}, \mathbf{s} - \mathbf{o} \rangle + \lambda \frac{1}{2} (\mathbf{s} - \mathbf{o})^T \sum_{v=2}^m \mathbf{B}^{1v} (\mathbf{s} - \mathbf{o}) + c. \quad (15)$$

Additionally, we consider a small sphere $\overline{\mathcal{S}}_{\mathbf{z}_1}'$, where $\mathbf{z}_1 := (\rho - \varrho)\boldsymbol{\alpha}_1 + \mathbf{o}$, similar to what is shown in Figure 2(a). We may approximate the quadratic term on the right hand side above by its first order Taylor approximation as follows:

$$\begin{aligned} \phi^{\text{fair}}((\mathbf{s}, \mathbf{o}, \dots, \mathbf{o}); \mathbf{a}, \mathbb{B}, \lambda) &\approx \phi^{\text{fair, apx}}((\mathbf{s}, \mathbf{o}, \dots, \mathbf{o}); \mathbf{a}, \mathbb{B}, \lambda) \\ &= \langle -(1 - \lambda)\boldsymbol{\tau}^1 \odot \mathbf{a} + \lambda \sum_{v=2}^m \mathbf{B}^{1v}(\mathbf{z}_1 - \mathbf{o}), \mathbf{s} \rangle \end{aligned} \quad (16)$$

for \mathbf{s} in a small neighbourhood around the rate profile \mathbf{z}_1 . Since the metric is essentially linear in \mathbf{s} , the following lemma from [Hiranandani et al., 2020] shows that the metric in (16) is quasiconcave in λ .

Lemma 3. *Under the regularity assumption that $\langle -\boldsymbol{\tau}^1 \odot \mathbf{a}, \sum_{v=2}^m \mathbf{B}^{1v}(\mathbf{z}_1 - \mathbf{o}) \rangle \neq 1$, the function*

$$\vartheta(\lambda) := \max_{\mathbf{s} \in \overline{\mathcal{S}}_{\mathbf{z}_1}'} \phi^{\text{fair, apx}}((\mathbf{s}, \mathbf{o}, \dots, \mathbf{o}); \mathbf{a}, \mathbb{B}, \lambda) \quad (17)$$

is strictly quasiconcave (and therefore unimodal) in λ .

The unimodality of $\vartheta(\lambda)$ allows us to perform the one-dimensional binary search in Algorithm 5 using the query space $\overline{\mathcal{S}}_{\mathbf{z}_1}'$, tolerance ϵ , and the oracle Ω . The binary search algorithm is same as Algorithm 4 in [Hiranandani et al., 2020] and provided here for completeness.

E EXTENSION TO ELICITING GENERAL QUADRATIC METRICS

In this section, we discuss how the entire setup including the proposed procedure and the guarantees of the main paper described in terms of the *diagonal* entries of the predictive rate matrix extends to a setup where the metric is defined in all the terms of the rate matrix. For this section, we use an additional notation. For a matrix \mathbf{A} , let *off-diag*(\mathbf{A}) returns a vector of off-diagonal elements of \mathbf{A} .

Just like the main paper, we consider a k -class classification setting with $X \in \mathcal{X}$ and $Y \in [k]$ denoting the input and output random variables, respectively. We assume access to an n -sized sample $\{(\mathbf{x}, y)_i\}_{i=1}^n$ generated *iid* from a distribution $\mathbb{P}(X, Y)$. We work with randomized classifiers $h : \mathcal{X} \rightarrow \Delta_k$ that for any \mathbf{x} gives a distribution $h(\mathbf{x})$ over the k classes and use $\mathcal{H} = \{h : \mathcal{X} \rightarrow \Delta_k\}$ to denote the set of all classifiers.

General Predictive rates: We define the predictive rate matrix for a classifier h by $\mathbf{R}(h, \mathbb{P}) \in \mathbb{R}^{k \times k}$, where the ij -th entry is the fraction of label- i examples for which the randomized classifier h predicts j :

$$R_{ij}(h, \mathbb{P}) := \mathbb{P}(h(X) = j | Y = i) \quad \text{for } i, j \in [k], \quad (18)$$

where the probability is over draw of $(X, Y) \sim \mathbb{P}$ and the randomness in h .

Notice that each diagonal entry of \mathbf{R} can be written in terms of its off-diagonal elements:

$$R_{ii}(h, \mathbb{P}) = 1 - \sum_{j=1, j \neq i}^k R_{ij}(h, \mathbb{P}).$$

Thus, we can represent a rate matrix with its $q := (k^2 - k)$ off-diagonal elements, write it as a vector $\mathbf{r}(h, \mathbb{P}) = \text{off-diag}(\mathbf{R}(h, \mathbb{P}))$, and interchangeably refer to it as the ‘vector of general rates’ or ‘off-diagonal rates’. To distinguish from rates considered in the main paper, we will call the rates entries corresponding to the diagonals of the rate matrix, i.e., $\mathbb{P}(h(X) = i|Y = i)$ as discussed in Equation(1), as the ‘diagonal rates’.

Feasible general rates: The set of all feasible general rates is given by:

$$\mathcal{R} = \{\mathbf{r}(h, \mathbb{P}) \in [0, 1]^q : h \in \mathcal{H}\}.$$

The quadratic metric in general rates is defined in the same way as Definition 3 as follows:

Definition 1 (Quadratic Metric in General Rates). *For a vector $\mathbf{a} \in \mathbb{R}^q$ and a positive semi-definite symmetric matrix $\mathbf{B} \in \mathbb{R}^{q \times q}$ with $\|\mathbf{a}\|_2^2 + \|\mathbf{B}\|_F^2 = 1$ (w.l.o.g. due to scale invariance):*

$$\phi^{\text{quad}}(\mathbf{r}; \mathbf{a}, \mathbf{B}) = \langle \mathbf{a}, \mathbf{r} \rangle + \frac{1}{2} \mathbf{r}^T \mathbf{B} \mathbf{r}. \quad (19)$$

Example 1 (Distribution matching). We can extend Example 2 in the multiclass case as follows. In certain applications, one needs the proportion of predictions for each class (i.e., the coverage) to match a target distribution $\boldsymbol{\pi} \in \Delta_k$ [Goh et al., 2016, Narasimhan, 2018]. A measure often used for this task is the squared difference between the per-class coverage and the target distribution: $\phi^{\text{cov}}(\mathbf{r}) = \sum_{i=1}^k (\text{cov}_i(\mathbf{r}) - \pi_i)^2$, where $\text{cov}_i(\mathbf{r}) = 1 - \sum_{j=1}^{k-1} r_{(i-1)(k-1)+j} + \sum_{j>i} r_{(j-1)(k-1)+i} + \sum_{j<i} r_{(j-1)(k-1)+i-1}$. Similar metrics can be found in the quantification literature where the target is set to the class prior $\mathbb{P}(Y = i)$ [Esuli and Sebastiani, 2015, Kar et al., 2016]. We capture more general quadratic distance measures for distributions, e.g. $(\text{cov}(\mathbf{r}) - \boldsymbol{\pi})^T \mathbf{Q} (\text{cov}(\mathbf{r}) - \boldsymbol{\pi})$ for a positive semi-definite matrix $\mathbf{Q} \in PSD_k$ [Lindsay et al., 2008].

The definition of metric elicitation and oracle query remain the same except that the vector \mathbf{r} now represents the vector of general rates, and not just the diagonal rates.

Consider Appendix B, where we discuss proof of Proposition 1, Proposition 2, and a procedure to construct a feasible sphere of appropriate radius in the convex set of diagonal rates. The entire methodology applies to general set of rates by just replacing diagonal rates in the proofs with the general rates. Thus, all the geometrical properties discussed in Proposition 1 for the set of diagonal rates applies to the set of general rates. The exact geometry of the set of diagonal rates, as shown in Figures 2(a) and 2(b) may differ from the geometry of the set of general rates; however, the geometric properties including \mathbf{e}_i being the vertices remains the same. Therefore, under the same Assumption 1, we can guarantee an existence of a sphere in the set of general rates similar to Remark 1.

Once we guarantee a sphere in the set of general rates, we can follow LPME for eliciting linear metrics in general rates or QPME for eliciting linear metrics in general rates. The computational and query complexity will depend on the number of unknowns, which in the case of general rates, will be $\tilde{O}(q)$ for LPME and $\tilde{O}(q^2)$ for QPME.

F PRACTICALITY OF QUERYING ORACLE

Recall that in our setup, any query posed to the oracle needs to be feasible, i.e., should be achievable by some classifier (see definition of feasible rates in Section 2). Therefore the oracle we query can be a human expert or a group of experts who compare intuitive visualizations of rates, or can be an entire population of users (as would be the case with A/B testing).

An important practical concern in employing the proposed QPME procedure is the number of queries needed to be posed to the oracle. We note that (i) the number of queries needed by our proposal is optimal (i.e. matches the lower bound for the problem in Theorem 2), (ii) has only a *linear* dependence on the number of unknowns (Theorem 1), and (iii) can be considerably reduced by making reasonable practical structural assumptions about the metric to reduce the number

of unknowns. While our procedure’s query complexity for the most general setup (with $O(k^2)$ unknowns) is $\tilde{O}(k^2)$, the quadratic dependence on k is merely an artifact of there being $O(k^2)$ unknowns in this setup. For example, when the number of classes is large, one may just cluster the classes from error perspective. For example, one may assume same error costs for similar classes. This will reduce the number of unknowns to $O(c^2)$, where $c \ll k$ is the number of cluster of classes.

We also stress that in many internet-based settings, one can deploy A/B tests to obtain preferences by aggregating feedback from a large group of participants. In this case, the entire user population serves as an oracle. Most internet-based companies run thousands of A/B tests daily making it practical to get preferences for our metric elicitation procedure. Moreover, one can employ practical improvements such as running A/B tests with fewer participants in the initial rounds (when the rates are far apart) and switch to running A/B tests with more precision in later rounds. Note that because the queries posed by our method always corresponds to a feasible classifier (see Section 3), one can easily run comparisons between classifiers as a part of an A/B test.

If needed, our algorithms can also work with queries that compare classification statistics directly, instead of classifiers. There has been growing work on visualization of confusion matrices (predictive rates) for non-expert users. For example, see [Beauxis-Aussalet and Hardman, 2014] and [Shen et al., 2020]. With the aid of such intuitive visualizations, it is reasonable to expect human practitioners to comprehend the queries posed to them and provide us with pairwise comparisons. Moreover, we have shown that our approach is resilient to noisy responses, which enhances our confidence in their ability to handle human feedback. In practice, the most viable option will depend on the target population.

Finally, we would like to emphasize that because our query complexity has only a linear dependence on the number of unknowns, and the number of unknowns can be reduced with practical structural assumptions, our proposal is as practical as the prior methods [Hiranandani et al., 2019a,b] for linear metric elicitation. In fact, despite eliciting from a more flexible class, our proposal has the same dependence on the number of unknowns as those prior methods.

G ELICITATION GUARANTEE FOR THE QPME PROCEDURE

G.1 SAMPLE COMPLEXITY BOUNDS

Recall from Definition 5 that the oracle responds correctly as long as $|\phi(\mathbf{r}_1) - \phi(\mathbf{r}_2)| > \epsilon_\Omega$. For simplicity, we assume that our algorithm has access to the population rates \mathbf{r} defined in Eq. (1). In practice, we expect to estimate the rates using a sample $D := \{\mathbf{x}, y\}_{i=1}^n$ drawn from the distribution \mathbb{P} , and to query classifiers from a hypothesis class \mathcal{H} with finite capacity. Standard generalization bounds (e.g. Daniely et al. [2015]) give us that with high probability over draw of D , the estimates $\hat{\mathbf{r}}$ are close to the population rates \mathbf{r} , up to the desired tolerance ϵ_Ω , as long as we have sufficient samples. Further, since the metrics ϕ are Lipschitz w.r.t. rates, with high probability, we thus gather correct oracle feedback from querying with finite sample estimates $\Omega(\hat{\mathbf{r}}_1, \hat{\mathbf{r}}_2)$.

More formally, for $\delta \in (0, 1)$, as long as the sample size n is greater than $O(\log(|\mathcal{H}|/\delta)/\epsilon_\Omega^2)$, the guarantee in Theorem 1 hold with probability at least $1 - \delta$ (over draw of D), where $|\mathcal{H}|$ can in turn be replaced by a measure of capacity of the hypothesis class \mathcal{H} . For example, one can show the following corollary to Theorem 1 for a hypothesis class \mathcal{H} in which each classifier is a randomized combination of a finite number of deterministic classifiers chosen from a set $\bar{\mathcal{H}}$, and whose capacity is measured in terms of the Natarajan dimension [Natarajan, 1989] of $\bar{\mathcal{H}}$.

Corollary 1. *Suppose the hypothesis class \mathcal{H} of randomized classifiers used to choose queries to the oracle is of the form:*

$$\mathcal{H} = \left\{ x \mapsto \sum_{t=1}^T \alpha_t h_t(x) \mid T \in \mathbb{Z}_+, \alpha \in \Delta_T, h_1, \dots, h_T \in \bar{\mathcal{H}} \right\},$$

for some class $\bar{\mathcal{H}}$ of deterministic multiclass classifiers $h : \mathcal{X} \rightarrow \{0, 1\}^k$. Suppose the deterministic hypothesis class $\bar{\mathcal{H}}$ has Natarajan dimension $d > 0$, and ϕ is 1-Lipschitz. Then for any $\delta \in (0, 1)$, as long as the sample size $n \geq O\left(\frac{d \log(k) + \log(1/\delta)}{\epsilon_\Omega^2}\right)$, the guarantee in Theorem 1 hold with probability at least $1 - \delta$ (over draw of $D = \{\mathbf{x}_i, y_i\}_{i=1}^n$ from \mathbb{P}).

The proof adapts generalization bounds from Daniely et al. [2015], and uses the fact that the predictive rate for any randomized classifier in \mathcal{H} is a convex combination of rates for deterministic classifiers in $\bar{\mathcal{H}}$ (due to linearity of expectation).

G.2 PROOFS

Before presenting the proof of Theorem 1, we re-write the LPME guarantees from [Hiranandani et al., 2019b] for linear metrics in the presence of an oracle noise parameter ϵ_Ω from Definition 5.

Lemma 4 (LPME guarantees with oracle noise [Hiranandani et al., 2019b]). *Let the oracle Ω 's metric be $\phi^{\text{lin}} = \langle \mathbf{a}, \mathbf{r} \rangle$ and its feedback noise parameter from Definition 5 be ϵ_Ω . Then, if the LPME procedure (Algorithm 2) is run using a sphere $\mathcal{S} \subset \mathcal{R}$ of radius ϱ and the binary-search tolerance ϵ , then by posing $O(k \log(1/\epsilon))$ queries it recovers coefficients $\hat{\mathbf{a}}$ with $\|\mathbf{a} - \hat{\mathbf{a}}\|_2 \leq O\left(\sqrt{k}(\epsilon + \sqrt{\epsilon_\Omega/\varrho})\right)$.*

Proof of Theorem 1. We first find the smoothness coefficient of the metric in Definition 3.

A function ϕ is said to be L -smooth if for some bounded constant L , we have:

$$\|\nabla\phi(x) - \nabla\phi(y)\|_2 \leq L\|x - y\|_2.$$

For the metric in Definition 3, we have:

$$\begin{aligned} \|\nabla\phi^{\text{quad}}(x) - \nabla\phi^{\text{quad}}(y)\|_2 &= \|\mathbf{a} + \mathbf{B}x - (\mathbf{a} + \mathbf{B}y)\|_2 \\ &\leq \|\mathbf{B}\|_2\|x - y\|_2 \\ &\leq \|\mathbf{B}\|_F\|x - y\|_2, \\ &\leq 1 \cdot \|x - y\|_2, \end{aligned}$$

where in the last step, we have used the scale invariance condition from Definition 3, i.e., $\|\mathbf{a}\|_2^2 + \|\mathbf{B}\|_F^2 = 1$, which implies that $\|\mathbf{B}\|_F^2 = 1 - \|\mathbf{a}\|_2^2 \leq 1$. Hence, the metrics in Definition 3 are 1-smooth.

Now, we look at the error in Taylor series approximation when we approximate the metric ϕ^{quad} in Definition 3 with a linear approximation. Our metric is

$$\phi^{\text{quad}}(\mathbf{r}) = \langle \mathbf{a}, \mathbf{r} \rangle + \frac{1}{2}\mathbf{r}^T \mathbf{B} \mathbf{r}.$$

We approximate it with the first order Taylor polynomial around a point \mathbf{z} , which we define as follows:

$$T_1(\mathbf{r}) = \langle \mathbf{a}, \mathbf{z} \rangle + \frac{1}{2}\mathbf{z}^T \mathbf{B} \mathbf{z} + \langle \mathbf{a} + \mathbf{B} \mathbf{z}, \mathbf{r} \rangle$$

The bound on the error in this approximation is:

$$\begin{aligned} |E(\mathbf{r})| &= |\phi^{\text{quad}}(\mathbf{r}) - T_1(\mathbf{r})| \\ &= \frac{1}{2}|(\mathbf{r} - \mathbf{z})^T \Delta\phi^{\text{quad}}|_{\mathbf{c}}(\mathbf{r} - \mathbf{z})| && \text{(First-order Taylor approximation error)} \\ &= \frac{1}{2}|(\mathbf{r} - \mathbf{z})^T \mathbf{B}(\mathbf{r} - \mathbf{z})| && \text{(Hessian at any point } \mathbf{c} \text{ is the matrix } \mathbf{B}) \\ &\leq \frac{1}{2}\|\mathbf{B}\|_2\|\mathbf{r} - \mathbf{z}\|_2^2 \\ &\leq \frac{1}{2}\|\mathbf{B}\|_F\varrho^2 \\ &\leq \frac{1}{2}\varrho^2 && \text{(Due to the scale invariance condition)} \end{aligned}$$

So when the oracle is asked $\Omega(\mathbf{r}_1, \mathbf{r}_2) = \mathbb{1}[\phi^{\text{quad}}(\mathbf{r}_1) > \phi^{\text{quad}}(\mathbf{r}_2)]$, the approximation error can be treated as feedback error from the oracle with feedback noise $2 \times \frac{1}{2}\varrho^2$. Thus, the overall feedback noise by the oracle is $\epsilon_\Omega + \varrho^2$ for the purposes of using Lemma 4 later.

We first prove guarantees for the matrix \mathbf{B} and then for the vector \mathbf{a} . We write Equation (9) in the following form assuming $d_1 = 1$ (since we normalize the coefficients at the end due to scale invariance):

$$\begin{aligned}
B_{ij} &= F_{ij} = \left[\frac{f_{ij}}{f_{1j}} \left(1 + \frac{f_{j1}}{f_{11}} \right) - \frac{f_{ij}}{f_{1j}} \frac{f_{j0}}{f_{10}} - \frac{f_{i0}}{f_{10}} + \frac{f_{ij}}{f_{1j}} \frac{f_{j1}}{f_{11}} \frac{\frac{f_{21}^-}{f_{11}} + \frac{f_{21}}{f_{11}} - 2\frac{f_{20}}{f_{10}}}{\frac{f_{21}^-}{f_{11}} - \frac{f_{21}}{f_{11}}} \right] \\
\Rightarrow \mathbf{B}[:, j] &= \mathbf{f}_j \left(\frac{1}{f_{1j}} + \frac{f_{j1}}{f_{1j}f_{11}} + \frac{f_{j0}}{f_{1j}f_{10}} + \frac{f_{j1}}{f_{1j}f_{11}} \left(\frac{\frac{f_{21}^-}{f_{11}} + \frac{f_{21}}{f_{11}} - 2\frac{f_{20}}{f_{10}}}{\frac{f_{21}^-}{f_{11}} - \frac{f_{21}}{f_{11}}} \right) \right) + \mathbf{f}_0 \frac{1}{f_{10}} \\
&= c_j \mathbf{f}_j + c_0 \mathbf{f}_0,
\end{aligned} \tag{20}$$

where $\mathbf{B}[:, j]$ is the j -th column of the matrix \mathbf{B} , and the constants c_j and c_0 are well-defined due to the regularity Assumption 4. Notice that,

$$\frac{\partial \mathbf{B}[:, j]}{\partial \mathbf{f}_j} = \text{diag}(\mathbf{c}'_j) \odot \mathbf{I} \quad , \text{ and } \quad \frac{\partial \mathbf{B}[:, j]}{\partial \mathbf{f}_0} = \text{diag}(\mathbf{c}'_0) \odot \mathbf{I},$$

where $\mathbf{c}'_j, \mathbf{c}'_0$ are vector of Lipschitz constants (bounded due to Assumption 4). This implies

$$\begin{aligned}
\|\bar{\mathbf{B}}[:, j] - \hat{\mathbf{B}}[:, j]\|_2 &\leq c'_j \|\bar{\mathbf{f}}_j - \hat{\mathbf{f}}_j\|_2 + c'_0 \|\bar{\mathbf{f}}_0 - \hat{\mathbf{f}}_0\|_2 \\
&\leq c'_j \sqrt{k} \left(\epsilon + \sqrt{\varrho + \epsilon_\Omega / \varrho} \right) + c'_0 \sqrt{k} \left(\epsilon + \sqrt{\varrho + \epsilon_\Omega / \varrho} \right) \\
&= O \left(\sqrt{k} \left(\epsilon + \sqrt{\varrho + \epsilon_\Omega / \varrho} \right) \right),
\end{aligned}$$

where we have used LPME guarantees from Lemma 4 under the oracle-feedback noise parameter $\epsilon_\Omega + \varrho^2$.

The above inequality provides bounds on each column of \mathbf{B} . Since $\|\mathbf{x}\|_\infty \leq \|\mathbf{x}\|_2$, we have $\max_{ij} |B_{ij} - \hat{B}_{ij}| \leq O \left(\sqrt{k} \left(\epsilon + \sqrt{\varrho + \epsilon_\Omega / \varrho} \right) \right)$, and consequentially, $\|\mathbf{B} - \hat{\mathbf{B}}\|_F \leq O \left(k \sqrt{k} \left(\epsilon + \sqrt{\varrho + \epsilon_\Omega / \varrho} \right) \right)$.

Now let us look at guarantees for \mathbf{a} . Since $\mathbf{a} = \mathbf{d} - \mathbf{B}\mathbf{o}$ from (4), we can write

$$\mathbf{a} = c_0 \mathbf{f}_0 - \sum_{j=1}^k o_j \mathbf{B}[:, j],$$

where $c_0 = 1/f_{10}$. Since \mathbf{o} is the rate achieved by random classifier, $o_j = 1/k \forall j \in [k]$, and thus we have

$$\frac{\partial \mathbf{a}}{\partial \mathbf{f}_0} = c_0 \mathbf{I} \quad \text{and} \quad \frac{\partial \mathbf{a}}{\partial \mathbf{B}[:, j]} = \frac{1}{k} \mathbf{I}.$$

Thus,

$$\begin{aligned}
\|\mathbf{a} - \hat{\mathbf{a}}\|_2 &\leq c'_0 \sqrt{k} \left(\epsilon + \sqrt{\varrho + \epsilon_\Omega / \varrho} \right) + \frac{1}{k} \sum_{j=1}^k \sqrt{k} \left(\epsilon + \sqrt{\varrho + \epsilon_\Omega / \varrho} \right) \\
&= c'_0 \sqrt{k} \left(\epsilon + \sqrt{\varrho + \epsilon_\Omega / \varrho} \right) + \frac{1}{k} \sum_{j=1}^k c'_j \sqrt{k} \left(\epsilon + \sqrt{\varrho + \epsilon_\Omega / \varrho} \right) \\
&= O \left(\sqrt{k} \left(\epsilon + \sqrt{\varrho + \epsilon_\Omega / \varrho} \right) \right),
\end{aligned}$$

where c'_0, c'_j 's are some Lipschitz constants (bounded due to Assumption 4). □

Notice the trade-off in the elicitation error that depends on the size of the sphere. As expected, when the radius of the sphere ϱ increases, the error due to approximation increases, but at the same time, error due to feedback reduces because we get better responses from the oracle. In contrast, when the radius of the sphere ϱ decreases, the error due to approximation decreases, but the error due to feedback increases.

The following corollary translates our guarantees on the elicited metric to the guarantees on the optimal rate of the elicited metric. This is useful in practice, because the optimal classifier (rate) obtained by optimizing a certain metric is often the key entity for many applications.

Corollary 2. *Let ϕ^{quad} be the original quadratic metric of the oracle and $\hat{\phi}^{\text{quad}}$ be its estimate obtained by the QPME procedure (Algorithm 1). Moreover, let \mathbf{r}^* and $\hat{\mathbf{r}}^*$ be the maximizers of ϕ^{quad} and $\hat{\phi}^{\text{quad}}$, respectively. Then, $\phi^{\text{quad}}(\mathbf{r}^*) \leq \phi^{\text{quad}}(\hat{\mathbf{r}}^*) + O\left(k^2\sqrt{k}\left(\epsilon + \sqrt{\varrho + \epsilon\Omega/\varrho}\right)\right)$.*

Proof. We first show that if $|\phi^{\text{quad}}(r) - \hat{\phi}^{\text{quad}}(r)| \leq \epsilon$ for all rates r and some slack ϵ , then it follows that $\phi^{\text{quad}}(\hat{\mathbf{r}}^*) \geq \phi^{\text{quad}}(\mathbf{r}^*) - 2\epsilon$. This is because:

$$\begin{aligned} \phi^{\text{quad}}(\hat{\mathbf{r}}^*) &\geq \hat{\phi}^{\text{quad}}(\hat{\mathbf{r}}^*) - \epsilon && \left(\text{as } \hat{\phi}^{\text{quad}} \text{ approximates } \phi^{\text{quad}}\right) \\ &\geq \hat{\phi}^{\text{quad}}(\mathbf{r}^*) - \epsilon && \left(\text{as } \hat{\mathbf{r}}^* \text{ maximizes } \hat{\phi}^{\text{quad}}\right) \\ &\geq \phi^{\text{quad}}(\mathbf{r}^*) - 2\epsilon && \left(\text{as } \hat{\phi}^{\text{quad}} \text{ approximates } \phi^{\text{quad}}\right) \end{aligned} \quad (21)$$

Now, let us derive the trivial bound $|\phi^{\text{quad}}(r) - \hat{\phi}^{\text{quad}}(r)|$ for any rate \mathbf{r} .

$$\begin{aligned} |\phi^{\text{quad}}(r) - \hat{\phi}^{\text{quad}}(r)| &= |\langle \mathbf{a} - \hat{\mathbf{a}}, \mathbf{r} \rangle + \frac{1}{2} \mathbf{r}^T (\mathbf{B} - \hat{\mathbf{B}}) \mathbf{r}| \\ &\leq |\langle \mathbf{a} - \hat{\mathbf{a}}, \mathbf{r} \rangle| + \frac{1}{2} |\mathbf{r}^T (\mathbf{B} - \hat{\mathbf{B}}) \mathbf{r}| \\ &\leq \|\mathbf{a} - \hat{\mathbf{a}}\|_2 \|\mathbf{r}\|_2 + \frac{1}{2} \|\mathbf{B} - \hat{\mathbf{B}}\|_2 \|\mathbf{r}\|_2^2 \\ &\leq \|\mathbf{a} - \hat{\mathbf{a}}\|_2 \sqrt{k} + \frac{1}{2} \|\mathbf{B} - \hat{\mathbf{B}}\|_F k \\ &\leq O\left(k^2\sqrt{k}\left(\epsilon + \sqrt{\varrho + \epsilon\Omega/\varrho}\right)\right), \end{aligned} \quad (22)$$

where in the fourth step, we have used the fact that the rates are bounded in $[0, 1]$; hence $\|\mathbf{r}\|_2 \leq \sqrt{k}$, and in the fifth step, we have used the guarantees from Theorem 1. Combining (21) and (22) gives us the desired result. \square

Proof of Theorem 2. For the purpose of this proof, let us replace $\left(\epsilon + \sqrt{\varrho + \epsilon\Omega/\varrho}\right)$ by some slack ϵ . Theorem 1 guarantees that after running the QPME procedure for $O(k^2 \log(1/\epsilon))$ queries, we have

- $\|\mathbf{a} - \hat{\mathbf{a}}\|_2 \leq O(\sqrt{k}\epsilon)$
- $\left\| \mathbf{B} - \hat{\mathbf{B}} \right\|_F \leq O(k\sqrt{k}\epsilon)$.

If we vectorize the tuple (\mathbf{a}, \mathbf{B}) and denote it by w , we have $\|w - \hat{w}\|_2 \leq O(k\sqrt{k}\epsilon)$, where both $\|w\|_2, \|\hat{w}\|_2 = 1$, due to the scale invariance condition from Definition 3. Note that w is $\frac{k^2+3k}{2}$ -dimensional vector and defines the scale-invariant quadratic metric elicitation problem. Now, we have to count the minimum number of \hat{w} that are possible such that $\|w - \hat{w}\|_2 \leq O(k\sqrt{k}\epsilon)$.

This translates to finding the covering number of a ball in $\|\cdot\|_2$ norm with radius 1, where the covering balls have radius $k\sqrt{k}\epsilon$. Let us denote the cover by $\{u_i\}_{i=1}^N$ and the ball with radius 1 as \mathbb{B} . We then have:

$$\begin{aligned} \text{Vol}(\mathbb{B}) &\leq \sum_{i=1}^N \text{Vol}(k\sqrt{k}\epsilon\mathbb{B} + u_i) \\ &= N \text{Vol}(k\sqrt{k}\epsilon\mathbb{B}) \\ &= (k\sqrt{k}\epsilon)^{\frac{k^2+3k}{2}}. \end{aligned}$$

Number of queries for QPME		Number of queries for Fair-QPME				
k		$m \backslash k$	2	3	4	5
2	265.43	2	332.10	867.65	1663.73	2738.59
3	669.29	3	796.37	2127.44	4094.31	6734.15
4	1205.91	4	1398.14	3808.67	7363.82	12180.84
5	1879.74	5	2130.92	5887.99	11454.18	18999.71

Table 1: Number of queries required for eliciting regular quadratic metrics (Def. 3) and fairness quadratic metrics (Def. 4) in Section 6. The number of predictive rates and sensitive groups are denoted by k and m , respectively. Recall that a quadratic metric has $O(k^2)$ unknowns. We see that the number of queries is of order $O(k^2)$ for the quadratic metric in rates, and additionally, $O(m^2k^2)$ for the fair (quadratic) metrics. Theorem 2 shows that one cannot improve on this query complexity. Nonetheless, one may make more structural assumptions on the metric to bring down the number of queries in practice.

Thus the number of \hat{w} that are possible are at least

$$c \left(\frac{1}{k\sqrt{k\epsilon}} \right)^{\frac{k^2+3k}{2}-1} \leq N,$$

where c is a constant. Since each pairwise comparison provides at most one bit, at least $O(k^2) \log(\frac{1}{k\sqrt{k\epsilon}})$ bits are required to get a possible \hat{w} . We require $O(k^2) \log(\frac{1}{\epsilon})$ queries, which is near-optimal barring log terms. \square

H EXTENDED EXPERIMENTS

The source code is provided along with the supplementary material. The experiments in this paper were conducted on a machine with the following configuration: *2.6 GHz 6 core Intel i7 processor with 16GB RAM.*

H.1 MORE DETAILS ON SIMULATED EXPERIMENTS ON QUADRATIC METRIC ELICITATION

Number of queries. First, we look at the number of queries that were actually required to elicit quadratic and fair (quadratic) metrics in Section 6. Recall that the QPME procedure (Algorithm 1) requires running the LPME subroutine (Algorithm 2) $k + 2$ times. As discussed in Appendix A, each run of LPME requires at most $3 \times 3 \times k \log(\pi/2\epsilon)$ queries. So, the maximum number of queries required in eliciting a quadratic metric is $(k + 2) \times 3 \times 3 \times k \times 8$ for a binary search tolerance $\epsilon = 10^{-2}$, where we vary $k \in \{2, 3, 4, 5\}$ in the experiments.

However, note that, the elicitation error shown in Figure 4 is averaged over 100 simulated oracles, each one with its own simulated quadratic metric. Due to the nature of the binary search involved in the LPME subroutine (see Algorithm 2 and Figure 1), not every reduction of the search interval requires three queries. Many times the interval can be shrunk in less than three queries. The actual number of queries may vary across the oracles. The number of queries averaged over the 100 oracles corresponding to experiments in Section 6 is shown in Table 1.

Theorem 2 shows that our query complexity (which is linear in the number of unknowns) matches the lower bound for the problem, which means that it is theoretically impossible to obtain a better complexity order for our problem setup. In practice, it can be considerably reduced by making reasonable assumptions on the metric. For example, when the number of classes is large, one may just cluster the classes from error perspective. For example, one may assume same error costs for similar classes. This will reduce the number of unknowns to $O(c^2)$, where $c \ll k$ is the number of cluster of classes.

Comparison to a baseline. In Figures 4(a)–4(b), we show box plots of the ℓ_2 (Frobenius) norm between the true and elicited linear (quadratic) coefficients. We generally find that QPME is able to elicit metrics close to the true ones.

To reinforce this point, we also compare the elicitation error of the QPME procedure and the elicitation error of a baseline which assigns equal coefficients to \mathbf{a} and \mathbf{B} in Figure 2. We see that the elicitation error of the baseline is order of magnitude higher than the elicitation error of the QPME procedure. This holds for varying k showing that the QPME procedure is able to elicit oracle’s multiclass quadratic metrics very well.

Effect of Assumption 4. The larger standard deviation for $k = 5$ in Figure 4 is due to Assumption 4 failing to hold with sufficiently large constants $c_0, c_{-1}, c_1 \dots, c_q$ in a small number of trials and the resulting estimates not being as accurate.




Figure 2: Elicitation error in comparison to a baseline which assigns equal coefficients.

Table 2: Dataset statistics

Dataset	k	#samples	#features
default	2	30000	33
adult	2	43156	74
sensIT Vehicle	3	98528	50
covtype	7	581012	54

We now analyze in greater detail the effect of this regularity assumption in eliciting quadratic metrics and understand how the lower bounding constants impact the elicitation error. Assumption 4 effectively ensures that the ratios computed in (9) are well-defined. To this end, we generate two sets of 100 quadratic metrics. One set is generated following Assumption 4 with one coordinate in the gradient being greater than 10^{-2} , and the other is generated randomly without any regularity condition. For both sets, we run QPME and elicit the corresponding metrics.




Figure 3: Elicitation error for metrics following Assumption 4 vs elicitation error for completely random metrics.

In Figure 3, we see that the elicitation error is much higher when the regularity Assumption 4 is not followed, owing to the fact that the ratio computation in (9) is more susceptible to errors when gradient coordinates approach zero in some cases of randomly generated metrics. The dash-dotted curve (in red color) shows the trajectory of the theoretical bounds with increasing k (within a constant factor). In Figure 3, we see that the mean of ℓ_2 (analogously, Frobenius) norm better follow the theoretical bound trajectory in the case when regularity Assumption 4 is followed by the metrics.

We next analyze the ratio of estimated fractions to the true fractions used in (9) over 1000 simulated runs. Ideally, this ratio should be 1, but as we see in Figure 4, these estimated ratios can be off by a significant amount for a few trials when the metrics are generated randomly. The estimated ratios, however, are more stable under Assumption 4. Since we multiply fractions in (9), even then we may observe the compounding effect of fraction estimation errors in the final estimates. Hence, we see for $k = 5$ in Figure 4(a)-4(b), the standard deviation is high due to few trials where the lower bound of 10^{-2} on the constants in Assumption 4 may not be enough. However, majority of the trials as shown in Figure 4(a)-4(b) and Figure 2 incur low elicitation error.

H.2 RANKING OF REAL-WORLD CLASSIFIERS

Performance metrics provide quantifiable scores to classifiers. This score is then often used to rank classifiers and select the best set of classifiers in practice. In this section, we discuss the benefits of elicited metrics in comparison to some default metrics while ranking real-world classifiers.




Figure 4: Ratio of estimated to true fractions over 1000 simulated runs with and without Assumption 4.

Ranking in case of quadratic metrics: For this experiment, we work with four real world datasets with varying number of classes $k \in \{2, 3, 7\}$. See Table 2 for details of the datasets. We use 60% of each dataset to train classifiers. The rest of the data is used to compute (testing) predictive rates. For each dataset, we create a pool of 80 classifiers by tweaking hyper-parameters in some famous machine learning models that are routinely used in practice. Specifically, we create 20 classifiers each from logistic regression models [Kleinbaum et al., 2002], multi-layer perceptron models [Pal and Mitra, 1992], LightGBM models [Ke et al., 2017], and support vector machines [Joachims, 1999]. We compare ranking of these 80 classifiers provided by competing baseline metrics with respect to the ground truth ranking, which is provided by the oracle’s true metric.

We generate a random quadratic metric ϕ^{quad} following Definition 3. We treat the true ϕ^{quad} as oracle’s metric. It provides us the ground truth ranking of the classifiers in the pool. We then use our proposed procedure QPME (Algorithm 1) to recover the oracle’s metric. For comparison in ranking of real-world classifiers, we choose two linear metrics that are routinely employed by practitioners as baselines. The first is accuracy $\phi^{\text{acc}} = 1/\sqrt{k}\langle \mathbf{1}, \mathbf{r} \rangle$, and the second is weighted accuracy, where we just use the linear part $\langle \mathbf{a}, \mathbf{r} \rangle$ of the oracle’s true quadratic metric $\langle \mathbf{a}, \mathbf{r} \rangle + \frac{1}{2}\mathbf{r}^T \mathbf{B} \mathbf{r}$. We repeat this experiment over 100 trials.




Figure 5: Performance of competing metrics while ranking real-world classifiers. ‘elicited’ is the metric elicited by QPME, ‘linear’ is the metric that comprises only the linear part of the oracle’s true quadratic metric, and ‘accuracy’ is the linear metric which weigh all classification errors equally (often used in practice).

We report NDCG (with exponential gain) [Valizadegan et al., 2009] and Kendall-tau coefficient [Shieh, 1998] averaged over the 100 trials in Figure 5. We observe consistently for all the datasets that the elicited metrics using the QPME procedure achieve the highest possible NDCG and Kendall-tau coefficient of 1. As we saw in Section 5, QPME may incur elicitation error, and thus the elicited metrics may not be very accurate; however, Figure 5 shows that the elicited metrics may still achieve near-optimal ranking results. This implies that when given a set of classifiers, ranking based on elicited metric scores align most closely to true ranking in comparison to ranking based on default metric scores. Consequentially, the elicited metrics may allow us to select or discard classifiers for a given task. This is advantageous in practice. For the *covtype* dataset, we see that the *linear* metric also achieves high NDCG values, so perhaps ranking at the top is quite accurate; however Kendall-tau coefficient is low suggesting that the overall ranking of classifiers is poor. We also observe that, in general, the weighted version (*linear* metric) is better than *accuracy* while ranking classifiers.

Ranking in case of fair (quadratic) metrics: With regards to fairness, we performed a similar experiment as above for comparing fair-classifiers’ ranking on Adult and Default datasets with gender as the protected group. There are two genders provided in the datasets, i.e., $m = 2$. We simulate fairness metrics as given in Definition 4 that gives ground-truth ranking of classifiers and evaluate the ranking by the elicited (fair-quadratic) metric using the procedure described in Section 4 (also depicted in Figure 3). In Table 3, we show the NDCG and KD-Tau values for our method and for three baselines: (a) ‘Linear with no fairness’, which is the metric that comprises only the linear part of the oracle’s true quadratic fair metric from Definition 4 without the fairness violation, (b) ‘Accuracy with eq. odds’ is the metric which weigh all classification errors and fairness violations equally, and (c) Fair Performance Metric Elicitation (FPME) procedure from [Hiranandani et al., 2020].² We again see that the ranking by the metric elicited using the proposed fair-QPME procedure (Section 4) is closest to the ground-truth ranking. The metric elicited by FPME [Hiranandani et al., 2020] ranks classifiers better than ‘Linear with no fairness’ and ‘Accuracy with equalized odds’; however, it is beaten by the proposed fair-QPME procedure.

Dataset →	Adult		Default	
Method ↓, Ranking Measures →	NDCG	KD-TAU	NDCG	KD-TAU
Linear with no fairness	0.9875	0.5918	0.9994	0.9057
Accuracy with equalized odds	0.9857	0.3763	0.9889	0.4953
Elicited via FPME [Hiranandani et al., 2020]	0.9989	0.9611	0.9974	0.9650
Elicited via Fair-QPME (Proposed)	1.0000	0.9972	1.0000	0.9981

Table 3: Performance of competing metrics while ranking real-world classifiers for fairness. ‘Linear with no fairness’ is the metric that comprises only the linear part of the oracle’s true quadratic fair metric from Definition 4 without the fairness violation, ‘Accuracy with eq. odds’ is the metric which weigh all classification errors and fairness violations equally (often used in practice), ‘Elicited via FPME [Hiranandani et al., 2020]’ is the metric elicited using the procedure from Hiranandani et al. [2020], ‘Elicited via Fair-QPME’ is the metric elicited by the proposed (quadratic) fairness metric elicitation procedure from Section 4 (also depicted in Figure 3),

Ranking in case of added structural assumptions on the metrics: Lastly, we discuss an experiment where we show how one may make structural assumptions on the metric when the *actual* number of unknowns is large and still get comparable results in practical settings. For this experiment, we assume that the oracle’s true metric is quadratic in general rate entries as explained in Appendix E. Thus, the number of unknowns is $O(q^2)$, where $q = k^2 - k$ and is the number of off-diagonal entries of the rate matrix. We can apply the QPME procedure as it is and elicit a quadratic metric in general rates with $O(q^2)$ queries (see Appendix E), since there are $O(q^2)$ unknowns.

Note that, even if the oracle’s original metric is a quadratic metric in off-diagonal entries, as a heuristic, we could still use our procedure to elicit a quadratic metric in diagonal rate entries. Moreover, we can use LPME procedure (Appendix A), too, to elicit a linear metric in off-diagonals and diagonal rate entries depending on the assumption we make on the metric.

Thus, for the ranking based experiments explained in Figure 5, we additionally ran (a) linear elicitation with diagonal rates, (b) linear elicitation with general rates, and (c) quadratic elicitation with diagonal rates, and compare their ranking with the elicited quadratic metric in general rates. As seen in Table 4, the quadratic approximation in the diagonal rates performs significantly better than eliciting a linear approximation in the general rates, while requiring the same query complexity ($\tilde{O}(k^2)$), and is close to the elicited quadratic metric in general rates, which require ($\tilde{O}(k^4)$) queries. Hence, one can make structural assumptions on the metric to reduce the query complexity and still get comparable results in practice.

Dataset →	Adult dataset		Default dataset	
Elicited Metric ↓, Rank- Measure →	NDCG	KD-TAU	NDCG	KD-TAU
Linear-diagonal ($\tilde{O}(k)$)	0.9783	0.6053	0.9790	0.4536
Linear-general ($\tilde{O}(k^2)$)	0.9908	0.7713	0.9863	0.6216
Quadratic-diagonal ($\tilde{O}(k^2)$)	0.9968	0.9611	1.0000	0.9979
Quadratic-general ($\tilde{O}(k^4)$)	1.0000	0.9986	1.0000	0.9979

Table 4: Performance on ranking real-world classifiers when the oracle’s true metric is quadratic in general rate entries: the elicited *quadratic metric in diagonal entries perform better than elicited linear metric* in general rate entries (while requiring same no. of queries), and close to the elicited quadratic metric in general rates.

²While FPME [Hiranandani et al., 2020] does not elicit a quadratic metric, one can still compare the elicited metrics based on how they rank candidate classifiers on real-world data.

I EXTENDED RELATED WORK

We discuss how the area of metric elicitation, in general, and our quadratic elicitation proposal, in particular, differ from two related fields: (i) inverse reinforcement learning and (ii) ranking from pairwise comparisons using choice models.

I.1 INVERSE REINFORCEMENT LEARNING (RL)

The idea of learning a reward/cost function in the inverse RL problems is conceptually similar to metric elicitation. However, there are many key differences. Studies such as [Ng et al., 2000, Wu et al., 2020, Abbeel and Ng, 2004] try to learn a linear reward function either by knowing the optimal policy or expert demonstrations. Not only is the type of feedback in these studies different from the pairwise feedback we handle, but the studies are focused on linear rewards. In contrast, our goal is to use pairwise feedback to elicit quadratic metrics which are important for classification problems, especially, fairness. Indeed nonlinear reward estimation in inverse RL problems has been tackled before [Levine et al., 2011, Fu et al., 2017], but these are passive learning approaches and do not come with query complexity guarantees like we do. Because of the use of a complex function class, these methods are not easy to analyze. There has been some work on actively estimating the reward function in the inverse RL problems [Lopes et al., 2009]; this work involves discretizing the feature space and using maximum entropy based ideas to elicit a distribution over rewards, which clearly uses a different set of modeling assumptions than us. A recent work [Sadigh et al., 2017] elicits reward functions through active learning, but is again tied to eliciting linear functions and provides limited theoretical guarantees; whereas, we specifically focus on quadratic elicitation with rigorous guarantees.

In summary, our work is significantly different from inverse RL methods, in that unlike them, we are tied to a particular geometry of the query space (the space of error statistics achieved by feasible classifiers), and elicit quadratic (or polynomial) functions from pairwise comparisons, specifically, in an active learning manner.

I.2 RANKING FROM PAIRWISE COMPARISONS

Our work is also quite different from the use of choice models such the Bradley-Terry-Luce (BTL) model for rank aggregation. (i) Firstly, choice models such as BTL are commonly used to learn an aggregate global ranking of a finite set of N items from pairwise comparisons. The underlying problem involves estimating a N -dimensional quality score vector for the items [Shah et al., 2015]. In contrast, metric elicitation estimates an oracle’s classification metric, a function of a classifier’s error statistics. The applications for the two problems are very different: while ranking aggregation strategies using BTL are often prescribed for aggregating user opinions on a restaurant or product, metric elicitation seeks to find the right objective to optimize for a classification task. (ii) Secondly, the noise model in BTL is stochastic and depends on the distance between the quality of items; whereas, our noise model in Definition 5 is not stochastic and is oblivious to the distance of rates unless the rates are very close. (iii) Thirdly, while there is some work on the extended BTL model where items are represented by feature vectors [Niranjan and Rajkumar, 2017], and the goal is to learn weights on the features to complete the ranking of the items, most of the work in this area considers a passive setting, where pairwise comparisons are assumed to be iid. In contrast, our work involves actively learning nonlinear utilities with theoretical bounds on query complexity. (iv) Lastly, the closest active learning work we could find with BTL models [Mohajer et al., 2017] does not generalize to feature-dependent utilities and is proposed for finding the top- k items, which is entirely different from metric elicitation.

Other fields that are less closely related to our work include learning scoring functions for supervised label ranking problems [Fürnkranz and Hüllermeier, 2010], and the more traditional metric learning literature, where the task is to learn a distance metric that captures similarities between data points, with the goal of using it for downstream learning tasks [Kulis et al., 2013].

J PRELIMINARY USER STUDY

We are actively conducting user studies for eliciting performance metrics. In this section, we provide a peek into our future work. The goal of this preliminary study is to check workflow of the practical implementation of the metric elicitation framework with real data, and to a certain extent, support or reject the hypothesis that the implicit user preferences can be quantified using the pairwise comparison queries over confusion matrices or predictive rates. In addition, the goal includes




Figure 6: Flow-chart and bar-chart based visualizations for (binary classification) confusion matrices in the recidivism prediction task from Shen et al. [2020].

testing certain assumptions regarding the noise in the subject’s (oracle’s) responses, work around with finite samples, eliciting actual performance metrics in real-life scenarios, and evaluating the quality of the recovered metric.

The following user study works with the space of confusion matrices, i.e., entries of the type $\mathbb{P}(Y = i, h = j)$ for $i, j \in [k]$, instead of the predictive rates. In the future, we plan on incorporating rates, i.e., entries of the type $\mathbb{P}(h = j|Y = i)$ for $i, j \in [k]$, in the visualizations as well. Our contributions are summarized as follows:

- We create a web UI that uses existing visualizations of confusion matrices (predictive rates) that are refined to capture preferences over pairwise comparisons.
- The UI implements the binary-search procedure from Algorithm 2 at the back end that make use of the real-time responses over confusion matrices to elicit a linear performance metric for a binary classification task.
- We perform a user study with ten subjects and elicit their linear performance metrics using the proposed web UI. We compare the quality of the recovered metric by comparing their responses to the elicited metric’s responses over a set of randomly chosen pairwise comparison queries.

J.1 CHOICE OF TASK AND DATASET USED

Our choice of task is *cancer diagnosis* [Yang and Naiman, 2014] for which we use the Breast Cancer Wisconsin (Original) dataset from the UCI repository.³ The dataset has been extensively used in the literature for binary classification, where the label 1 denotes *malignant* cancer and label 0 denotes *benign* cancer. There are 699 samples in total, wherein each sample has 9 features. The task for any classifier is to take the 9 features of a patient as input and predict whether or not the patient has cancer. We divide this data into two equally sized parts – the training and the test data. Using the training data, we learn a logistic regression model to obtain an estimate of the class-conditional probability, i.e., $\hat{\eta}(x) = \hat{\mathbb{P}}(y = 1|X)$. We then create a sphere using Algorithm 3 inside the space of confusion matrices computed on the test data .

J.2 CHOICE OF VISUALIZATION

In modern times, ensuring effective public understanding of algorithmic decisions, especially, machine learning models has become an imperative task. With this view in mind, we borrow the visualizations of confusion matrices for the binary classifications setup from Shen et al. [2020]. The authors provide a concrete step towards the above goal by redesigning confusion matrices to support non-experts in understanding the performance of machine learning models. The final visualizations that we use from Shen et al. [2020] are created over multiple iterative user-studies. The visualizations are shown in Figure 6 in the context of a recidivism prediction task. One is the *flow-chart*, which helps users in understanding the direction of the data, and the other is the *bar chart*, which helps users in understanding the quantities involved.

However, in light of our preliminary discussions with Human-Computer Interaction (HCI) and machine learning researchers, we make/recommend the following changes in the visualization for pairwise comparison purposes in the metric elicitation framework.

1. Based on the observation that multiple visualizations of the information help in better user understanding, we choose to use both *flow-chart* and *bar-chart*, together to depict a confusion matrix.
2. We transform the data statistics so that the numbers denote out-of-100 samples.

³The dataset can be downloaded from <https://tinyurl.com/dn2esyvw>.



Figure 7: A sample of a pairwise comparison query from a run of the binary-search based procedure Algorithm 2

3. We found that the total number of positive and negative labels along with total number of positive and negative predictions are very helpful in comparing two confusion matrices. Therefore, we add the total numbers in the flow-chart boxes and on axes in the bar-charts.

A sample of a pairwise comparison query with modified visualizations incorporating the points above is shown in Figure 7. We next discuss the web user interface.

J.3 USER INTERFACE

We discuss our proposed web User Interface (UI) in detail and discuss our rationale behind its components. The UI has two parts to it as explained below.

The first phase of the UI is where we actually ask subjects for pairwise preferences over confusion matrices, and implement our binary-search procedure from Algorithm 2. The subjects have to make a choice reflecting on the trade-off between false positives and false negatives. Algorithm 2 takes in real-time preferences of the subjects, generates next set of queries based on the current responses, and converge to a linear performance metric at the back end. We save this (linear) performance metric for each subject. We stop the binary-search when the search interval becomes less than or equal to 0.05 (ϵ in line 1 of Algorithm 2).

In order to evaluate the quality of the recovered metric, in the second phase, we ask the subjects fifteen pairwise comparison queries, each on a separate web page, right after the binary search algorithm has converged, and we have elicited the metric. The subjects do not know this information and are shown evaluation queries in continuation to the previous phase (i.e., the binary search). The query comprises of two randomly selected confusion matrices that lie inside the feasible region. This set of queries is used to evaluate the effectiveness of the elicited metric.

J.4 STUDY RESULTS

We compute the fraction of times our elicited metric’s preferences matches with the subject’s preferences on the fifteen queries, i.e.,

$$\mathcal{M} := \frac{\sum_{i=1}^{15} \mathbb{1}[\text{subject's prefer. for query } i == \text{metric's prefer. for query } i]}{15} \times 100. \quad (23)$$

We show the elicited metric for the fifteen subjects and the measure \mathcal{M} values in Table 5. We see for nine out of ten subjects that more than 85% of the time our elicited metric’s preferences matches with the subject’s preferences on the fifteen evaluation queries. For three subjects, our metric’s preference matches exactly for all the evaluation queries.

The absolute numbers for the \mathcal{M} measure look good; however, how good they are is still a missing piece in this study because of the lack of a baseline. In future, we plan to devise ways to develop a baseline for the metric elicitation task and compare to that baseline on the measure \mathcal{M} .

Table 5: The elicited linear performance metrics for the ten subjects along with the fraction of times (in %) the elicited metric’s preferences matches with the subject’s preferences over the fifteen evaluation queries.

Subjects	Linear Performance Metric	\mathcal{M}
S1	0.125 TN + 0.875 TP	87
S2	0.141 TN + 0.859 TP	100
S3	0.125 TN + 0.875 TP	93
S4	0.141 TN + 0.859 TP	100
S5	0.328 TN + 0.672 TP	73
S6	0.031 TN + 0.969 TP	87
S7	0.031 TN + 0.969 TP	100
S8	0.359 TN + 0.641 TP	87
S9	0.125 TN + 0.875 TP	93
S10	0.141 TN + 0.859 TP	87

References

- Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *ICML*, 2004.
- Emma Beauxis-Aussalet and Lynda Hardman. Visualization of confusion matrix for non-expert users. In *IEEE Conference on Visual Analytics Science and Technology, Poster Proceedings*, 2014.
- Amit Daniely, Sivan Sabato, Shai Ben-David, and Shai Shalev-Shwartz. Multiclass learnability and the ERM principle. *JMLR*, 16(1):2377–2404, January 2015.
- A. Esuli and F. Sebastiani. Optimizing text quantifiers for multivariate loss functions. *ACM Transactions on Knowledge Discovery and Data*, 9(4), 2015.
- Justin Fu, Katie Luo, and Sergey Levine. Learning robust rewards with adversarial inverse reinforcement learning. *ArXiv preprint, arXiv:1710.11248*, 2017.
- Johannes Fürnkranz and Eyke Hüllermeier. Preference learning and ranking by pairwise comparison. In *Preference Learning*, pages 65–82. Springer, 2010.
- Gabriel Goh, Andrew Cotter, Maya Gupta, and Michael P Friedlander. Satisfying real-world goals with dataset constraints. In *NeurIPS*, 2016.
- Gaurush Hiranandani, Shant Boodaghians, Ruta Mehta, and Oluwasanmi Koyejo. Performance metric elicitation from pairwise classifier comparisons. In *AISTATS*, 2019a.
- Gaurush Hiranandani, Shant Boodaghians, Ruta Mehta, and Oluwasanmi O Koyejo. Multiclass performance metric elicitation. In *NeurIPS*, 2019b.

- Gaurush Hiranandani, Harikrishna Narasimhan, and Oluwasanmi Koyejo. Fair performance metric elicitation. In *NeurIPS*, 2020.
- Thorsten Joachims. Svm-light: Support vector machine. URL: <http://svmlight.joachims.org/>, University of Dortmund, 1999.
- Purushottam Kar, Shuai Li, Harikrishna Narasimhan, Sanjay Chawla, and Fabrizio Sebastiani. Online optimization methods for the quantification problem. In *ACM SIGKDD*, 2016.
- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. In *NeurIPS*, 2017.
- David G Kleinbaum, K Dietz, M Gail, Mitchel Klein, and Mitchell Klein. *Logistic regression*. Springer, 2002.
- Brian Kulis et al. Metric learning: A survey. *Foundations and Trends® in Machine Learning*, 5(4):287–364, 2013.
- Sergey Levine, Zoran Popovic, and Vladlen Koltun. Nonlinear inverse reinforcement learning with gaussian processes. *NeurIPS*, 24:19–27, 2011.
- Bruce G Lindsay, Marianthi Markatou, Surajit Ray, Ke Yang, Shu-Chuan Chen, et al. Quadratic distances on probabilities: A unified foundation. *The Annals of Statistics*, 36(2):983–1006, 2008.
- Manuel Lopes, Francisco Melo, and Luis Montesano. Active learning for reward estimation in inverse reinforcement learning. In *ECML-PKDD*, pages 31–46. Springer, 2009.
- Soheil Mohajer, Changho Suh, and Adel Elmahdy. Active learning for top- k rank aggregation from noisy comparisons. In *ICML*, 2017.
- Harikrishna Narasimhan. Learning with complex loss functions and constraints. In *AISTATS*, 2018.
- B. K. Natarajan. On learning sets and functions. *Machine Learning*, 4(1):67–97, 1989.
- Andrew Y Ng, Stuart J Russell, et al. Algorithms for inverse reinforcement learning. In *ICML*, 2000.
- UN Niranjan and Arun Rajkumar. Inductive pairwise ranking: going beyond the $n \log(n)$ barrier. In *AAAI*, volume 31, 2017.
- SK Pal and S Mitra. Multilayer perceptron, fuzzy sets, and classification. *IEEE Transactions on Neural Networks*, 3(5): 683–697, 1992.
- Dorsa Sadigh, Anca D Dragan, Shankar Sastry, and Sanjit A Seshia. Active preference-based learning of reward functions. In *RSS*, 2017.
- Nihar Shah, S. Balakrishnan, J. Bradley, A. Parekh, K. Ramchandran, and Martin Wainwright. Estimation from pairwise comparisons: Sharp minimax bounds with topology dependence. In *AISTATS*, 2015.
- Hong Shen, H. Jin, A. A. Cabrera, A. Perer, H. Zhu, and Jason I Hong. Designing alternative representations of confusion matrices to support non-expert public understanding of algorithm performance. *CSCW2*, 2020.
- Grace S Shieh. A weighted kendall’s tau statistic. *Statistics & probability letters*, 39(1):17–24, 1998.
- H. Valizadegan, R. Jin, R. Zhang, and J. Mao. Learning to rank by optimizing NDCG measure. In *NeurIPS*, 2009.
- Zheng Wu, Liting Sun, Wei Zhan, Chenyu Yang, and Masayoshi Tomizuka. Efficient sampling-based maximum entropy inverse reinforcement learning with application to autonomous driving. *IEEE Robotics and Automation Letters*, 5(4): 5355–5362, 2020.
- Sitan Yang and Daniel Q Naiman. Multiclass cancer classification based on gene expression comparison. *Statistical Applications in Genetics and Molecular Biology*, 13(4):477–496, 2014.
- Yunfeng Zhang, Rachel Bellamy, and Kush Varshney. Joint optimization of ai fairness and utility: A human-centered approach. In *AIES*, pages 400–406, 2020.