Learning the eye of the beholder: Statistical modeling and estimation for personalized color perception

Xuanzhou Chen [1] **Austin Xu** [2 *] **Jingyan Wang** [3] **Ashwin Pananjady** [1 3]

Abstract

Color perception has long remained an intriguing topic in vision and cognitive science. It is a common practice to classify a person as either "color-normal" or "color-blind", and that there are a few prevalent types of color-blindness. However, empirical evidence has repeatedly suggested that at best, categories for color-blindness only serve as approximations to real manifestations of it. To better understanding individual-level color perception, we propose a color perception model that unifies existing theories for color-normal and color-blind people, which posits a low-dimensional structure in color space according to which any given user distinguishes colors. We design an algorithm to learn this low-dimensional structure from user queries, and prove statistical guarantees on its performance. Taking inspiration from these guarantees, we design a novel data collection paradigm based on perceptual adjustment queries (PAQs) that efficiently infers a user's color distinguishability profile from a small number of cognitively lightweight responses. In a host of simulations, PAQs offer significant advantages over the de facto method of collecting comparison-based similarity queries.

1. Introduction

A non-trivial fraction of the population, about 8% of men and 0.5% of women, have color-blindness conditions (also called color vision deficiency, or CVD). Color-blindness is conventionally classified into a few types, based on the person's level of difficulty in perceiving three primary colors (red, green, and blue). It also varies in severity, which can be quantified in color-blind test scores, from mild to complete deficiency in perceiving a primary color. Common tests for diagnosing color-blind conditions include the Ishihara test (naming numbers from colored dots) and the Farnsworth-Munsell test (arranging colors on a line to form a gradual hue change). However, empirical studies suggest that these tests often fail to accurately identify the type of color-blindness or extent of defect (Hardy et al., 1945; Dain et al., 2019; Miquilini et al., 2019; Van Staden et al., 2018).

In this work, we provide a principled framework towards fine-grained analysis of a person's color discrimination ability. We go beyond pre-defined categories and leverage confusion lines in a color space (Moreira et al., 2018) as the directions along which the person is not able to distinguish colors. These confusion lines intersect at one single point (termed the "copunctal point", which provides rich information about the impaired/missing cone); see Figure 1 for an example of the confusion lines and the copunctal point for the type of red-green color-blindness.

To locate the copunctal point, we draw on a second line of color vision literature that conducts experiments to study people's varying levels of difficulty in distinguishing similar colors, with or without color deficiency. This difficulty is characterized by ellipses in the color space, where colors within the same ellipse are hard for a person to distinguish from the color at the center of this ellipse. We term the center the "reference color" of the ellipse. The major axis of the ellipse represents the direction that the person has the most difficulty in distinguishing colors from the reference color. This major axis reduces to a confusion line for people with complete deficiency in perceiving a primary color.

We adopt a two-step procedure for computing the copunctal point. In the first step, we collect user data to estimate their ellipses and associated major axes. In the second step, we compute an "approximate" intersection point of these major axes and report it as our copunctal point estimate. We provide a theoretical result that relates the error in locating the copunctal point to the error in estimating the major axes (see Thm 3.2). This result enables us to use and assess different methods for estimating the ellipses in the first step.

ICML 2024 Workshop on Models of Human Feedback for AI Alignment. *Work done while at the School of Electrical and Computer Engineering, Georgia Institute of Technology. [1]School of Electrical and Computer Engineering, Georgia Institute of Technology. [2]Salesforce AI Research. [3]School of Industrial and Systems Engineering, Georgia Institute of Technology. Correspondence to: Xuanzhou Chen <xchen920@gatech.edu>.

ICML 2024 Workshop on Models of Human Feedback for AI Alignment, Vienna, Austria.

We exemplify this plug-in approach by using the perceptual adjustment queries (PAQ) (Xu et al., 2024) to estimate the ellipses in the first step. PAQ is designed for Mahalanobis metric learning, where the metric is represented by a PSD matrix. User responses to PAQs are collected from an interface that presents a reference item (in the metric space) to the user, and asks the user to move along a slider and stop when the item changes from being similar to dissimilar. PAQs are thus natural query mechanisms for our task: the user is presented a reference color, and instructed to identify the first color on the sliding bar that is the same as the reference color (see Figure 2). Furthermore, the problem of estimating the major axis of confusion ellipse is related to that of metric learning. We propose a least-squares estimator to estimate the ellipses and subsequently the major axes. We provide theoretical guarantees on the accuracy in estimating the ellipse (Thm 4.1), completing the first step in the plug-in approach.

Combining the theoretical results for the two steps, we derive end-to-end guarantees for estimating the copunctal point (Corollary 4.2). A sample experimental result for this end-to-end procedure is presented in Figure 3, which visualizes the ellipses and confusion lines for four individuals. These results provide new quantitative findings about color perception of normal and color-blind people (discussed in more detail in Section 5.2).

Our framework for copunctal point estimation leads to theoretical and practical implications. Theoretically, our result on ellipse estimation from the first step reveals the surprising phenomenon that the least-squares estimator is biased in estimating the ellipse, but unbiased in estimating its major axis. Practically, our framework has the potential to improve downstream applications. For example, the image recoloring task aims to enhance enhance contrast for color-blind people by performing a remapping of the colors in an image. Prior methods remap the colors by moving colors away from on the same confusion line (Meyer & Greenberg, 1988; Tsekouras et al., 2021b), where this line is determined by the type of color-blindness. Our approach yields an estimate for an *individual*'s copunctal point and hence an individualized confusion line along which to perform color correction. In summary, we envisage that our results lay the foundations towards using statistical and computational methods to quantitatively study and correct for color-blindness.

1.1. Our contributions and organization

Motivated by color theory, we formulate a forward model that characterizes color perception through a copunctal point and elliptical confusion regions. We set up the statistical estimation problem of computing the copunctal point from observations of confusion regions (Section 2), and present an algorithm for this task (Section 3). We prove that our

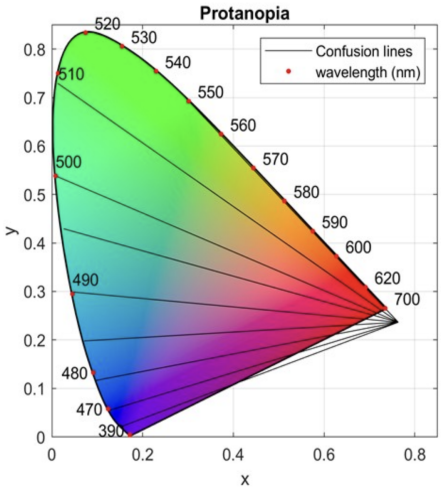

Figure 1. CIE xy color space and a set of confusion lines with intersection as the copunctal point for red-green color-blindness (protanopia) (Tsekouras et al., 2021a)

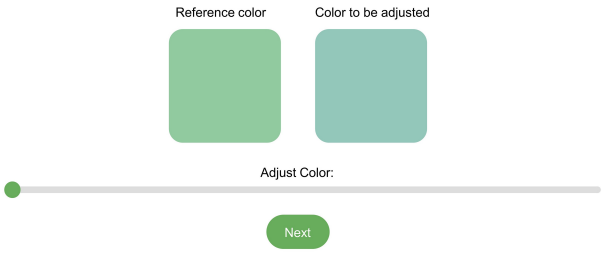

Figure 2. Mock user interface for color matching experiment

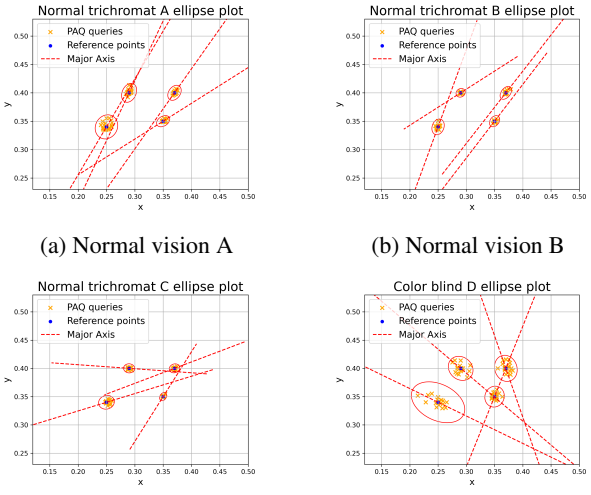

(a) Normal vision A (b) Normal vision B

(c) Normal vision C (d) Color-blind D

Figure 3. Ellipse plot and copunctal point estimation for normal trichromats A, B and C and color-blind D)

algorithm reliably estimates the copunctal point given good metric estimates, and complement this result with guarantees for metric learning (Section 4). We run some simulations to validate our algorithm with knowledge of the true copunctal point, and conclude with a small user study (Section 5).

1.2. Related work

We utilize two salient concepts from color vision theory, both of which utilize the CIE 1931 xy chromaticity diagram (Smith & Guild, 1931). The concept of confusion lines is typically defined for different types of color-blindness, and 2D color discrimination ellipses are commonly used to reason about people with normal color vision. We unify these concepts in presenting our forward model for color vision. Our work is also related to human data elicitation for metric learning. For example, pairwise comparisons (Ying et al., 2009) can be used to compare colors in pairs to judge if they are (in)distinguishable, and triplet comparisons (Mason et al., 2017) can be used to assess degrees of distinguishability from a reference color. We compare both these mechanisms with PAQ in simulation. Our theoretical analysis also touches upon popular themes in functional estimation. Our algorithm is inspired by robust optimization and a delicate use of the Davis-Kahan theorem (Yu et al., 2015).

1.3. Notation and Definition

For two real numbers a and b, let $a \wedge b = \min\{a, b\}$ and $a \vee b = \max\{a, b\}$. Given a matrix $\boldsymbol{A} \in \mathbb{R}^{d_1 \times d_2}$, denote $\|\boldsymbol{A}\|_F$ and $\|\boldsymbol{A}\|_{\text{op}}$ as its Frobenius norm and operator norm, respectively.

2. Forward model and copunctal point estimation problem formulation

In this section, we present a formal mathematical model and problem statement for personalized color perception estimation. We begin with a brief overview of color theory in Section 2.1 before presenting a formal mathematical model in Section 2.2 and problem statement in Section 2.3.

2.1. Color theory and a unified model for color perception

In this work, we operate primarily in the CIE 1931 xy chromaticity diagram (Smith & Guild, 1931), a two-dimensional color space that is a projected version of CIE xyY. We refer the chromaticity diagram as our color space. The two non-negative coordinates in this space characterize a color's *hue* and *colorfulness*, two properties that collectively make up a color's *chromaticity*. The range of visible colors forms a horseshoe shaped subset of the non-negative 2D quadrant known as a *gamut*. For color-blind people, distinguishing colors is difficult along a particular set of directions in the color space, known as *confusion lines*. Though these confusion lines vary based on the type of color-blindness (e.g., red-green or blue-yellow) as well as the particular colors considered, the orientation of these lines is not random in that they intersect at a point called the *copunctal point*. Difficulty distinguishing colors is not specific to color-blind people. For any color in the color space, nearby colors are naturally indistinguishable, forming local regions around colors known as *color discrimination ellipse (i.e., confusion regions)*.

However, the relationship between the confusion lines and the confusion regions are seldom discussed in previous research. In this paper, we posit a single model of human color perception that unifies the copunctal point structure for color-blind people and the elliptical confusion regions that arise for all people.

Model 1 (A unified model for color perception). Each user, regardless of color-blindness status, has a unique copunctal point $\boldsymbol{w}^\star$, from which confusion lines are emitted. These confusion lines define the major axes of the region of confusion centered at any color in the color space.

Estimating the copunctal point characterizes how they perceive differences with respect to any color in the color space.

2.2. A mathematical model for color perception

We now describe a mathematical model that captures various facets of color perception discussed in Section 2.1. In what follows, we explain how the problem of identifying the copunctal point for individuals is naturally captured by understanding a metric space and its geometry. We consider the CIE xy chromaticity color space in $\mathbb{R}^2$. The ellipse at each color $\boldsymbol{z} \in \mathbb{R}^2$ can then be equivalently represented by a matrix $\boldsymbol{\Sigma}_{\boldsymbol{z}}^\star \in \mathbb{R}^{2\times 2}$. We use the shorthand notation $\boldsymbol{\Sigma}^\star$ if its associated color $\boldsymbol{z}$ is clear from the context. All points within the ellipse centered are hence captured by the set $\{\boldsymbol{v} \in \mathbb{R}^2 : (\boldsymbol{v}-\boldsymbol{z})^T \boldsymbol{\Sigma}_{\boldsymbol{z}}^\star (\boldsymbol{v}-\boldsymbol{z}) \le y\}$. In the metric learning literature, the matrix $\boldsymbol{\Sigma}_{\boldsymbol{z}}^\star$ is called a Mahalanobis metric, and the product $(\boldsymbol{v}-\boldsymbol{z})^T \boldsymbol{\Sigma}_{\boldsymbol{z}}^\star (\boldsymbol{v}-\boldsymbol{z})$ is the distance (amount of color difference perceived by an individual) between the two points $\boldsymbol{v}$ and $\boldsymbol{z}$ in this metric space. The value of y is the minimum distance such that the user is able to perceive the color difference, and hence the ellipse $\{\boldsymbol{v} \in \mathbb{R}^2 : (\boldsymbol{v}-\boldsymbol{z})^T \boldsymbol{\Sigma}_{\boldsymbol{z}}^\star (\boldsymbol{v}-\boldsymbol{z}) \le y\}$ represents all colors that cannot be distinguished from the reference color $\boldsymbol{z}$. Note that we are free to scale $\boldsymbol{\Sigma}^\star$, so we assume $y = 1$ without loss of generality for the rest of the paper.

The major axis of the ellipse is the direction that the individual has the most difficulty distinguishing colors, namely the direction that the individual's color perception changes at the slowest speed. It is the direction that the value $\boldsymbol{a}^T \boldsymbol{\Sigma}_{\boldsymbol{z}}^\star \boldsymbol{a}$

changes the slowest, namely the direction of the last eigenvector (which reduces to the second eigenvector in 2D space). The copunctal point is hence the intersection of the lines corresponding to these eigenvectors.

2.3. Problem statement

Under the mathematical model for color perception, our goal is to learn a user's copunctal point from human responses to queries that gauge color similarity to a variety of reference colors. Formally, we assume that we have N reference colors $\{\boldsymbol{z}_n\}_{n=1}^N$, against which we query for perceived similarity and differences in color. For each reference color $\boldsymbol{z}_n$, we obtain M_n responses of the form $\{(\gamma_m^{(n)}, \mathcal{A}_m^{(n)})\}_{m=1}^{M_n}$, where $\gamma_m^{(n)}$ is the human response to a query comprised of the reference color $\boldsymbol{z}_n$ and the query items in the query set $\mathcal{A}_m^{(n)}$. From these responses and query sets, our goal is to estimate the copunctal point $\boldsymbol{w}^\star$. Two natural stages arise in the copunctal point estimation problem. First, human responses are used to estimate a metric. Then the estimated metrics can be used to estimate the copunctal point by leveraging their major axes.

3. A copunctal point estimator and guarantees

In this section, we present an estimator that takes a set of estimated metrics corresponding to different reference points, and estimates the copunctal point $\boldsymbol{w}^\star$. For this estimator, we present our main theoretical results on copunctal point estimation, a deterministic result that bounds copunctal point estimation error given noisy estimates of the ellipses.

3.1. An error cone-based estimator

We now present our copunctal point estimator, which takes as input a set of estimates of distance metrics at various reference points. We only require that the estimates have bounded operator norm error. Specifically, we assume that we have metric estimates $\{\widehat{\boldsymbol{\Sigma}}_n\}_{n=1}^N$ that each satisfy the error bound

$$\|\widehat{\boldsymbol{\Sigma}}_n - \boldsymbol{\Sigma}_n^\star\|_{\text{op}} \leq \tau_n. \tag{1}$$

As discussed in Section 2.2, the copunctal point is at the intersection of confusion lines. At each reference point, the confusion line aligns with the major axes of the level-set lines induced by the ground truth metric. This major axis is determined by the second eigenvector of the true metric. The estimator we present makes use of this fact and proceeds in three steps below, see Figure 4. This general procedure is outlined in Algorithm 1.

Step 1: Computing the major axes We assume that for each of the N reference points $\{\boldsymbol{z}_n\}_{n=1}^N$, we have estimated a distance metric $\{\widehat{\boldsymbol{\Sigma}}_n\}_{n=1}^N$ from collected human responses. For every estimated metric $\widehat{\boldsymbol{\Sigma}}_n$ at reference points $\boldsymbol{z}_n$, we obtain its eigenvectors $\widehat{\boldsymbol{u}}_1^{(n)}$ and $\widehat{\boldsymbol{u}}_2^{(n)}$ and corresponding eigenvalues $\widehat{\lambda}_1^{(n)}$ and $\widehat{\lambda}_2^{(n)}$ via its eigenvalue decomposition. Here, $\widehat{\lambda}_1^{(n)}$ and $\widehat{\lambda}_2^{(n)}$ are the largest and smallest eigenvalues, respectively, and $\widehat{\boldsymbol{u}}_1^{(n)}$ and $\widehat{\boldsymbol{u}}_2^{(n)}$ are their respective eigenvectors. In particular, we are interested in the second eigenvector $\widehat{\boldsymbol{u}}_2^{(n)}$, as discussed in Section 2.2, aligns with the confusion line for this reference point. If we were guaranteed exact estimation of the metrics, then the intersection of the second eigenvectors of any two metrics yields the copunctal point exactly. However, because our measurements are noisy and our estimation of the metrics is imprecise, our estimate of the major axes is also imprecise. As such, the estimated eigenvectors play a critical role in how we estimate the copunctal point.

Step 2: Constructing an error cone Since our estimate of the metric is imperfect, our estimated eigenvectors are not exactly aligned with the true confusion lines. However, if we have a reasonable estimate of the metric, then we should expect that our estimated eigenvectors do not deviate too far from the true eigenvectors. Specifically, we appeal to the Davis-Kahan theorem (Yu et al., 2015), which informally states that if the operator norm error between an estimated metric $\widehat{\boldsymbol{\Sigma}}_n$ and true metric $\boldsymbol{\Sigma}_n^\star$ is bounded by some threshold τ_n, then the angular deviation, which we denote α_n, between our estimated eigenvector $\widehat{\boldsymbol{u}}_2^{(n)}$ and the true eigenvector $\boldsymbol{u}_2^{(n)}$ is bounded as

$$\alpha_n \leq \frac{\tau_n}{|\widehat{\lambda}_1^{(n)} - \widehat{\lambda}_2^{(n)}|}. \tag{2}$$

Based on this result, we construct *error cones* of angle α_n, which we denote $\mathcal{C}_n(\alpha_n)$, for each of the N estimated metrics. The n-th cone has vertex at the reference point $\boldsymbol{z}_n$, is symmetric about the estimated eigenvector $\widehat{\boldsymbol{u}}_2^{(n)}$, and has boundaries that are an angle $\alpha_n/2$ rotated from the estimated eigenvector. We precisely define this error cone in Appendix A.

For each reference point and estimated metric, we compute the error cone, which defines a viable angular region around the estimated eigenvector in which the true eigenvector must lie in. This enables us to estimate the copunctal point, as we discuss next.

Step 3: Estimating the copunctal point With error cones for each reference point in hand, we now turn to estimating the copunctal point. Because each of the error cones contains the true eigenvector, and hence the true confusion line, the *intersection* of all of the error cones is non-empty and, crucially, contains the copunctal point. As a result, we solve the following feasibility program to obtain an estimate $\widehat{\boldsymbol{w}}$.

$$\widehat{\boldsymbol{w}} \in \mathcal{C}_n(\alpha_n) \quad \text{for } n = 1, \ldots, N \tag{3}$$

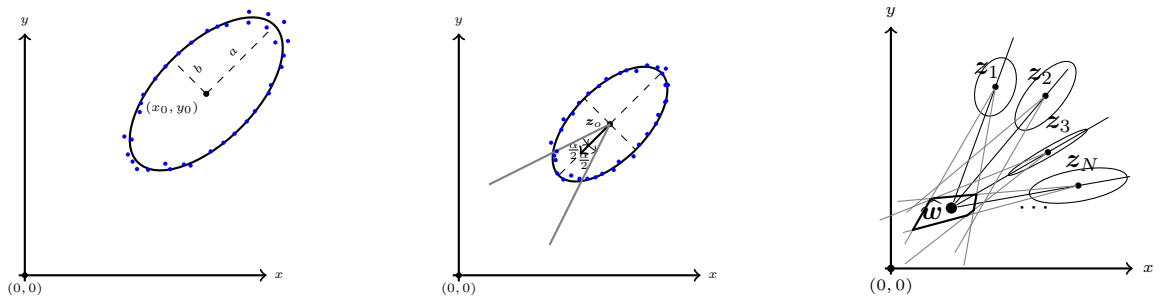

(a) (b) (c)

Figure 4. Three steps for copunctal point estimation: (a) Estimate metrics from PAQ responses; (b) Translate metric estimates to cones; (c) Estimate a copunctal point inside code intersection.

Since each error cone is defined by the two boundary lines originating from the reference point, the above optimization program can be rewritten as a linear feasibility program, and can be solved by off-the-shelf solvers. Specifically, if we write $\widehat{\boldsymbol{u}}_2^{(n)} = [\widehat{\boldsymbol{u}}_{2,1}^{(n)}, \widehat{\boldsymbol{u}}_{2,2}^{(n)}]^\top$, $\boldsymbol{w}^\star = [\boldsymbol{w}_1^\star, \boldsymbol{w}_2^\star]^\top$, and $\boldsymbol{z}_n = [\boldsymbol{z}_{n,1}, \boldsymbol{z}_{n,1}]^\top$, then the copunctal point $\boldsymbol{w}^\star$ must satisfy both of the following inequalities.

$$\frac{\boldsymbol{w}_2^\star - \boldsymbol{z}_{n,2}}{\boldsymbol{w}_1^\star - \boldsymbol{z}_{n,1}} \geq \frac{-\sin(\frac{\alpha_n}{2})\widehat{\boldsymbol{u}}_{2,1}^{(n)} + \cos(\frac{\alpha_n}{2})\widehat{\boldsymbol{u}}_{2,2}^{(n)}}{\cos(\frac{\alpha_n}{2})\widehat{\boldsymbol{u}}_{2,1}^{(n)} + \sin(\frac{\alpha_n}{2})\widehat{\boldsymbol{u}}_{2,2}^{(n)}} \tag{4}$$

$$\frac{\boldsymbol{w}_2^\star - \boldsymbol{z}_{n,2}}{\boldsymbol{w}_1^\star - \boldsymbol{z}_{n,1}} \leq \frac{\sin(\frac{\alpha_n}{2})\widehat{\boldsymbol{u}}_{2,1}^{(n)} + \cos(\frac{\alpha_n}{2})\widehat{\boldsymbol{u}}_{2,2}^{(n)}}{\cos(\frac{\alpha_n}{2})\widehat{\boldsymbol{u}}_{2,1}^{(n)} - \sin(\frac{\alpha_n}{2})\widehat{\boldsymbol{u}}_{2,2}^{(n)}}. \tag{5}$$

To estimate a copunctal point with N error cones, we solve a linear program with $2N$ linear constraints. We refer to the optimization program with the above constraints as $\texttt{LP}(\mathcal{C}_1(\alpha_1), \ldots, \mathcal{C}_N(\alpha_N))$. In solving this linear program, we handle eigenvector orientation and consider the geometry of pairs of error cone intersections with some practical considerations in Appendix B and Appendix C, respectively.

3.2. Deterministic copunctal point estimation

Suppose we have estimated N metrics $\widehat{\boldsymbol{\Sigma}}_1, \ldots, \widehat{\boldsymbol{\Sigma}}_N$ corresponding to N distinct reference points. For all metrics, we assume that each of our estimated metrics satisfies the operator norm bound $\|\widehat{\boldsymbol{\Sigma}}_n - \boldsymbol{\Sigma}_n\|_{\text{op}} \leq \tau_n$, where for now, we assume exact knowledge of τ_n.

In our analysis, we consider pairs of intersecting cones. Recall from Appendix C that for a pair of error cones, α_n denotes the error cone angle and θ_n denotes the deviation angle for reference point n. We make the following assumption on each pair (i, j) of reference points.

Assumption 3.1. For all pairs of error cones, the quantities (α_i, θ_i) and (α_j, θ_j) satisfy

$$\theta_i + \frac{\alpha_i}{2} \leq \frac{\pi}{2} \quad \text{and} \quad \theta_j + \frac{\alpha_j}{2} \leq \frac{\pi}{2}.$$

Algorithm 1 Error cone-based copunctal point estimator

Input: Reference points $\{\boldsymbol{z}_n\}_{n=1}^N$, estimated metrics $\{\widehat{\boldsymbol{\Sigma}}_n\}_{n=1}^N$, error bounds $\{\tau_n\}_{n=1}^N$

1: $\texttt{ErrorCones} \leftarrow \{\}$.
2: **for** $n = 1, \ldots, N$ **do**
3: **Step 1:** For each $\widehat{\boldsymbol{\Sigma}}_n$, obtain its eigenvalues $\widehat{\lambda}_1^{(n)}, \widehat{\lambda}_2^{(n)}$ and eigenvectors $\widehat{\boldsymbol{u}}_1^{(n)}, \widehat{\boldsymbol{u}}_2^{(n)}$ via eigenvalue decomposition.
4: **Step 2:** Set the error cone angle for each reference point as the upper bound (2):

$$\alpha_n = \frac{\tau_n}{|\widehat{\lambda}_1^{(n)} - \widehat{\lambda}_2^{(n)}|} \tag{6}$$

5: Construct error cone $\mathcal{C}_n(\alpha_n)$ from α_n and $\boldsymbol{z}_n$.
6: $\texttt{ErrorCones} \leftarrow \mathcal{C}_n(\alpha_n)$
7: **end for**
8: **Step 3:** Transform error cones into linear constraints (See equation (4)) and solve a linear program

$$\widehat{\boldsymbol{w}} \leftarrow \texttt{LP}(\mathcal{C}_1(\alpha_1), \ldots, \mathcal{C}_N(\alpha_N)). \tag{7}$$

Output: Estimated copunctal point $\widehat{\boldsymbol{w}}$

With this assumption, we are ready to present our first result, which is a bound on the estimation error of Algorithm 1.

Theorem 3.2. *Suppose Assumption 3.1 holds for all pairs of cones. There exists a universal positive constant c such that the following is true. Define $C_{ij} := \max\{(3 \vee \frac{6}{\tan(\theta_i \vee \theta_j)}) \cdot (1 + \tan^2(\theta_i \vee \theta_j)), \frac{1}{\sin(\theta_i \vee \theta_j)}\}$ for each pair (i, j). If $\|\widehat{\Sigma}_n - \Sigma_n^*\|_{op} \leq \tau_n$ for each $n \in [N]$, then*

$$\|\widehat{\boldsymbol{w}} - \boldsymbol{w}^\star\|_2 \leq c \min_{i,j \in [N]} C_{ij} \cdot \|\boldsymbol{z}_i - \boldsymbol{z}_j\|_2 \cdot \left(\frac{2\pi\tau_i}{|\widehat{\lambda}_1^{(i)} - \widehat{\lambda}_2^{(i)}|} \vee \frac{2\pi\tau_j}{|\widehat{\lambda}_1^{(j)} - \widehat{\lambda}_2^{(j)}|} \right) \tag{8}$$

Theorem 3.2 is a *deterministic* result and stated broadly to

accommodate estimates of metrics under various estimation procedures. In the next section, we specialize this result to the case of metric estimation with a specific type of query. A few remarks are warranted about this result.

The denominator contains the absolute difference of the two eigenvalues of the estimated metric, which we denote as the eigenvalue gap. This term arises naturally as a measure of the inherent difficulty in estimating the true metric's eigenvectors. If $\widehat{\lambda}_1^{(i)}$ and $\widehat{\lambda}_2^{(i)}$ are close in value, then the level-set lines of the true metric are likely to be circular, making it difficulty to discern the directions of the true eigenvectors. Conversely, if the difference between $\widehat{\lambda}_1^{(j)}$ and $\widehat{\lambda}_2^{(j)}$ is large, then the level-set lines of the true metric are more elliptical, resulting in easier identification of the major and minor axes. The effect of the eigenvalue gap is explored in simulation in Section 5.1.

4. Copunctal point estimation from perceptual adjustment queries

In this section, we consider the specific problem of estimating the copunctal point from perceptual adjustment queries (PAQs) (Xu et al., 2024). We first provide a brief overview of the PAQ mechanism in Section 4.1 and then present estimation bounds in Section 4.2.

4.1. Review of the perceptual adjustment query (PAQ)

A key component of estimating the copunctal point is how users are queried for perception information. Queries and responses $(\gamma_m^{(n)}, \mathcal{A}_m^{(n)})$ are defined broadly in Section 2.3 to accommodate a wide range of potential querying mechanisms. For example, paired comparisons between a reference color $\boldsymbol{z}_n$ and another color $\boldsymbol{a}_m^{(n)}$, then the query set is the single vector $\mathcal{A}_m^{(n)} = \{\boldsymbol{a}_m^{(n)}\}$ and the response $\gamma_m^{(n)}$ is a binary response that indicates if the user can distinguish the two colors. Ordinal responses like triplets (Mason et al., 2017), have been analyzed in the broader context of metric learning. The continuous nature of color space and the need to identify precise transition regions make perceptual adjustment queries (PAQs) (Xu et al., 2024) particularly amenable to our task. Here, we provide an overview of the PAQ mechanism.

A PAQ consists of a reference item and a continuous path of items that start from the reference and vary gradually. The user is asked to select the first item along this path that is dissimilar (or inversely, similar) to the reference. Operationally, PAQs can be implemented via a slider where a user is asked to adjust a slider that gradually changes a color to compatre against the reference color. The stopping point from the user is recorded as the response collected by the UI; see in Figure 2.

Such queries were originally proposed and theoretically analyzed in a metric learning context (Xu et al., 2024). In the metric learning setting, PAQ responses can be formalized as follows. Suppose we want to estimate an unknown metric $\boldsymbol{\Sigma}^\star$. For a reference color $\boldsymbol{z}$, the m-th query has query set $\mathcal{A}_m = \{\boldsymbol{a}_m\}$, where $\boldsymbol{a}$ is the direction along which colors change. That is, the user is presented with a path of items of the form $\{\boldsymbol{z} + \ell \boldsymbol{a}_m : \ell \in [0, \infty)\}$. The user then selects the first item $\boldsymbol{z} + \gamma_m \boldsymbol{a}_m$ that is perceivably different from $\boldsymbol{z}$, resulting in the response γ_m. This response γ_m can be viewed as a *scaling* of the query vector $\boldsymbol{a}_m$. Under the confusion region model, the response item $\boldsymbol{z} + \gamma_m \boldsymbol{a}_m$ should be a squared $\boldsymbol{\Sigma}^\star$-Mahalanobis distance y away from the reference color. For a noiseless PAQ response, we write this relationship as the quadratic measurement

$$y = \|(\boldsymbol{z} + \gamma \boldsymbol{a}) - \boldsymbol{z}\|_{\boldsymbol{\Sigma}^\star}^2 = \gamma^2 \boldsymbol{a}^\top \boldsymbol{\Sigma}^\star \boldsymbol{a}. \tag{9}$$

Since we do not expect human responses to be perfect, we adopt the noise model of (Xu et al., 2024), which models noise in the perception space. That is, for any direction $\boldsymbol{a}$, the user responds with an item that is a noisy distance $y + \eta$ away from the reference. This can be expressed as a noisy quadratic measurement

$$y + \eta = \gamma^2 \boldsymbol{a}^\top \boldsymbol{\Sigma}^\star \boldsymbol{a}. \tag{10}$$

Here, we assume that the noise η is an i.i.d. copy of a zero-mean random variable that is bounded, i.e., $\eta \leq \eta^\uparrow < y$ almost surely.

4.2. Metric and copunctal point estimation bounds for PAQs

Now we specialize the previous result to the scenario where we use PAQs to perform estimation on the N metrics. In order to do so, we first show a high-probability bound on the operator norm error when estimating with PAQs. In particular, we choose to apply an unregularized least squares estimator to recover $\widehat{\boldsymbol{\Sigma}}$. This procedure, as noted in previous work (Xu et al., 2024), is an inconsistent estimator. However, as we show in the following result, least squares produces a consistent estimate of a *scaled* version of the true matrix $\widehat{\boldsymbol{\Sigma}}$. Crucially, this preserves the directions of the eigenvectors, meaning the cone-based estimator (1) can still be utilized. Our PAQ estimation bound is presented below.

Theorem 4.1. *Let $c > 0$ be a universal constant. Let $\sigma = \frac{2\eta^\uparrow}{(1-\eta^\uparrow)(1+\eta^\uparrow)}$. Consider any $\delta \in (0, 1)$. Suppose the number of measurements satisfies*

$$M \geq cd \left\{ \log\left(\frac{M}{\delta}\right) + \log^2\left(\frac{M}{\delta}\right) \right\} \log\left(\frac{d}{\delta}\right). \tag{11}$$

Then with probability greater than $1-\delta$,

$$\|\widehat{\boldsymbol{\Sigma}} - \mathbb{E}[\frac{1}{1+\eta}]\boldsymbol{\Sigma}^\star\|_{op} \leq c\sigma \|\boldsymbol{\Sigma}^\star\|_F \sqrt{\frac{d^2\ \left(1+\texttt{polylog}(\frac{d}{\delta})\right)}{M}} \tag{12}$$

The σ quantity acts as a proxy for the noise variance of the bounded measurement noise η. Crucially, σ appears without any additive factors, meaning if the noise is zero, then we achieve perfect recovery. This is consistent with the fact that a set of $M \geq \frac{d(d+1)}{2}$ noiseless PAQ responses results in a solvable linear system, allowing for perfect recovery of the metric. With this result stated, we arrive at an error bound on the copunctal point estimation with PAQs.

Corollary 4.2. *Suppose the conditions of Theorem 3.2 and Theorem 4.1 hold. Consider any* $\delta \in (0,1)$. *Suppose the number of measurements for each of the* N *reference points satisfies*

$$M_i \geq cd\ \left\{\log\left(\frac{M}{\delta}\right) + \log^2\left(\frac{M}{\delta}\right)\right\}\log\left(\frac{d}{\delta}\right). \tag{13}$$

Then with probability greater than $1-N\delta$,

$$\|\widehat{\boldsymbol{w}} - \boldsymbol{w}^\star\|_2 \leq cd \min_{i,j\in[N]} C_{ij}\sigma \cdot \|\boldsymbol{z}_i - \boldsymbol{z}_j\|_2 \cdot \left(\frac{2\pi\tau_i}{|\widehat{\lambda}_1^{(i)} - \widehat{\lambda}_2^{(i)}|\sqrt{M_i}} \vee \frac{2\pi\tau_j}{|\widehat{\lambda}_1^{(j)} - \widehat{\lambda}_2^{(j)}|\sqrt{M_j}}\right) \cdot \sqrt{1+\texttt{polylog}(\frac{d}{\delta})}.$$

The proof of Corollary 4.2 is a straightforward application of Theorem 3.2 and Theorem 4.1. Once again, the estimation bound scales linearly with the noise variance. As such, with noiseless PAQ measurements, the exact recovery of the metrics also results in exact recovery of the copunctal point. This is consistent with our estimator: If we are able to recover all of the metrics perfectly, then the copunctal point must be precisely the point at the intersection of all of the major axes.

5. Experiments

In this section, we present our experimental results on both simulated and real-world data. In Section 5.1, we investigate the effects of various problem parameters and compare the performance of our estimator with PAQs against other popular query types in a numerical simulation study. We then present copunctal point estimation results on data collected from user a study in Section 5.2.

5.1. Numerical simulations

For all results, we report the average estimation error $\|\widehat{\boldsymbol{w}} - \boldsymbol{w}^\star\|_2$ and standard error of the mean over 20 trials. We first generate a ground truth copunctal point uniformly on $[0, 0.15]^2$, then select N reference color at random on $[0,1]^2 \backslash [0, 0.15]^2$. For each reference color, we generate a ground truth metric such that the second eigenvector aligns exactly with the line that connects the reference color and the true copunctal point. For each simulation, we pick a fixed difference $|\lambda_1 - \lambda_2|$ between the eigenvalues of all metrics. We denote this quantity as the *eigenvalue gap* $\lambda_{\texttt{gap}}$. We generate the largest eigenvalue $\lambda_1^{(n)}$ of the metric uniformly on $[\lambda_{\texttt{gap}}, 2\lambda_{\texttt{gap}}]$ and set the smallest eigenvalue to be $\lambda_2^{(n)} = \lambda_1^{(n)} - \lambda_{\texttt{gap}}$. For all query types, we assume the metric estimation operator norm error $\tau_n = c\sqrt{d^2/M}$ and cross-validate the constant. For all experiments, we utilize `cvxpy` (Diamond & Boyd, 2016; Agrawal et al., 2019) for both metric estimation and copunctal point estimation.

In the following experiments, we compare the performance of noiseless triplets, noiseless paired comparisons, and PAQs at various noise levels. To generate noiseless triplets, we randomly generate two colors $\boldsymbol{x}_1 = \boldsymbol{z} + \boldsymbol{a}_1$ and $\boldsymbol{x}_2 = \boldsymbol{z} + \boldsymbol{a}_2$, where $\boldsymbol{a}_i$ are i.i.d standard normal vectors. We select the color $\boldsymbol{x}_i$ that is closer to the reference point under the ground truth metric. Noiseless paired comparisons are generated in a similar manner: we generate an item $\boldsymbol{x} = \boldsymbol{z} + \boldsymbol{a}$, where $\boldsymbol{a}$ is a standard normal vector, and report if the distance between $\boldsymbol{x}$ and the reference color $\boldsymbol{z}$ is greater or less than the threshold y under the ground truth metric. For noiseless PAQs, we assume that the user responds with a color exactly a squared distance y away from the reference color. With noisy PAQs, we adopt the noise model (10), where the noise η is an i.i.d. copy of a zero-mean random variable that is bounded, i.e., $\eta \leq \eta^\uparrow < y$ almost surely. Following (Xu et al., 2024), we sample noise in PAQ responses to be uniform on the interval $[-\eta^\uparrow, \eta^\uparrow]$, with $\eta^\uparrow < y$. For low, medium, and high noise we set $\eta^\uparrow$ to be $0.1y, 0.5y$, and $0.75y$, respectively.

Choice of query type. A crucial question in copunctal point estimation how to collect color perception feedback from users. Our first simulation explores this choice. We compare PAQs against two commonly used queries in metric learning: triplets ("Which one of these two colors is more similar to the reference color?") and paired comparisons ("Is this color distinguishable from the reference color?"). Specifically, we compare PAQs at four different noise levels against noiseless triplets and paired comparisons. In this experiment, we set $y = 5$, $N = 25$, and $\lambda_{\texttt{gap}} = 10$, and sweep the performance of various queries as the number of measurements increases. Our results, shown in Figure 5a, demonstrates that PAQs have marked advantages in copunctal point estimation when compared to other ordinal queries, with responses with relatively high noise levels outperforming *noiseless* triplet and paired comparison responses. This increase in performance highlights the improved expressive

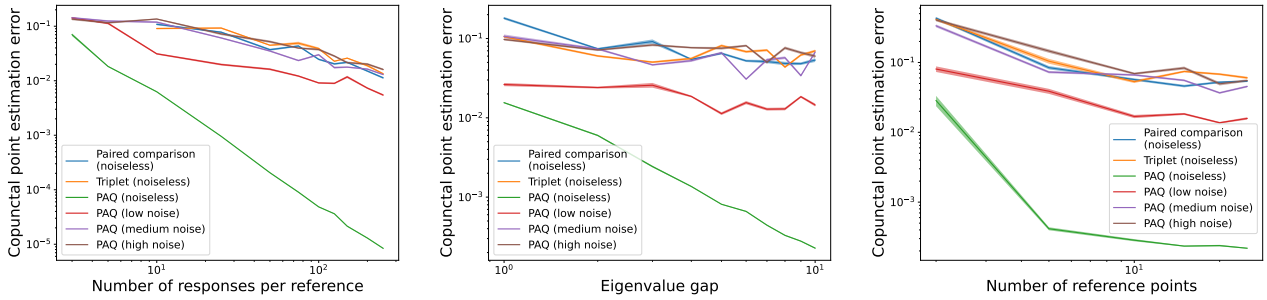

(a) (b) (c)

Figure 5: Numerical simulation results

power of each individual PAQ response.

Metric eigenvalue gap. Another key quantity that arises in our problem setting is the eigenvalue gap λ_{gap}, which serves as one measure of how difficult preserving the eigenvalue structure of the ground truth metric is in estimation. We show the effect of this quantity in Figure 5b, where we fix the number of responses to be 50 per reference for all query types and sweep the eigenvalue gap. Since we set the error cone angles to be the bound (2), the error cone angles are set inversely proportional to λ_{gap}. The smaller the gap, the larger the error cones, and the more imprecise the estimate of the copunctal point $\boldsymbol{w}^\star$ is. This effect is borne out to a modest degree in our experimental results, where estimation error with the same number of measurements tends to decrease as the eigenvalue gap grows.

Number of reference colors. We additionally test the effect of the number of reference colors for which we collect measurements and estimate metrics. We fix the number of measurements per reference to again be 50 and sweep the number of references. As seen in Figure 5c, as the number of references increases, the estimation error decreases to a certain point before flattening out, as expected. Because the copunctal point is solved for by a linear feasibility program and each reference color's error cone adds two additional constraints to the linear program, estimation error should decrease steadily when the number of references is small. However, as the number of references increases beyond a certain point, it is unlikely that any additional error cones will further reduce the area of the intersection, resulting in performance that levels off.

5.2. Real-world data collection and copunctal point estimation

We construct 80 questions by selecting four reference points in CIE xyY space: (0.25, 0.34, 0.5), (0.29, 0.40, 0.5), (0.37, 0.40, 0.5) and (0.35, 0.35, 0.5) from a region where the colors are relatively hard for people to differentiate in order to capture users' discrimination in a fine-grained manner. For each reference point, we equivalently space 20 sensing vectors over 360 degrees of angles. To construct the slider along each sensing vector's direction pointing towards a fixed reference point, we select a start point by hand-finding an appropriate length and sample 100 color points between the start point and the reference point. To prevent cheating, we randomly place the reference point on the slider and add a second path afterwards. We collect PAQ response data from 3 trichromats (color normal) users and 1 user with color-blindness, and adopt least square estimator to fit the chromaticity discrimination ellipses at four reference points. Figure 3 shows the four ellipses along with their major axes for each user. We make the following observations. For the three trichromats, their confusion lines still follow some directional patterns. For trichromats A and B, their confusion lines point diagonally in bottom left/top right; the confusion lines for trichromat C are more horizontal. The confusion lines also appear to have some converging behavior, and it appears plausible that they converge at some point (the copunctal point of the user). Moreover, the location of the copunctal point for trichromat C is different than the ones for trichromats A and B.

Comparing the visualization for the color-blind person D and the trichromats, the color-blind person has larger ellipses. The orientation of the confusion lines and the location of the copunctal point are quite different from all three trichromats. In this case, the color-blind person is known to have red-green color-blindness, and their copunctal point roughly aligns with the location for the protan type (with a location of (0.76, 0.24) derived theoretically from the literature (Smith & Pokorny, 2003). The person has one confusion line pointing in bottom left/bottom right. We observe that the ellipse is close to a circle for this particular reference color, and hypothesize that this orientation is due to the higher noise in estimating this major axis.

References

Agrawal, A., Verschueren, R., Diamond, S., and Boyd, S. A rewriting system for convex optimization problems,

2019.

Dain, S. J., Atchison, D. A., and Hovis, J. K. Limitations and precautions in the use of the farnsworth-munsell dichotomous d-15 test. *Optometry and Vision Science*, 96 (9):695–705, 2019.

Diamond, S. and Boyd, S. Cvxpy: A python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.

Hardy, L. H., Rand, G., and Rittler, M. C. Tests for the detection and analysis of color-blindness. i. the ishihara test: An evaluation. *J. Opt. Soc. Am.*, 35(4):268–275, Apr 1945.

Mason, B., Jain, L., and Nowak, R. Learning low-dimensional metrics. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Meyer, G. and Greenberg, D. Color-defective vision and computer graphics displays. *IEEE Computer Graphics and Applications*, 8(5):28–40, 1988.

Miquilini, L., Ratis, M. A. S., Lima, M. G., Bento-Torres, N. V. O., Lacerda, E. M. d. C. B., Cortes, M. I. T., Rodrigues, A. R., Silveira, L. C. d. L., and Souza, G. d. S. A proposed correction in the weighted method to score the ishihara test. *BMC research notes*, 12:1–6, 2019.

Moreira, H., Álvaro, L., Melnikova, A., and Lillo, J. *Colorimetry and dichromatic vision*. IntechOpen London, UK, 2018.

Smith, T. and Guild, J. The cie colorimetric standards and their use. *Transactions of the optical society*, 33(3):73, 1931.

Smith, V. C. and Pokorny, J. Color matching and color discrimination. *The science of color*, 2:103–148, 2003.

Tsekouras, G. E., Rigos, A., Chatzistamatis, S., Tsimikas, J., Kotis, K., Caridakis, G., and Anagnostopoulos, C.-N. A novel approach to image recoloring for color vision deficiency. *Sensors*, 21(8):2740, 2021a.

Tsekouras, G. E., Rigos, A., Chatzistamatis, S., Tsimikas, J., Kotis, K., Caridakis, G., and Anagnostopoulos, C.-N. A novel approach to image recoloring for color vision deficiency. *Sensors*, 21(8), 2021b.

Van Staden, D., Mahomed, F. N., Govender, S., Lengisi, L., Singh, B., and Aboobaker, O. Comparing the validity of an online ishihara colour vision test to the traditional ishihara handbook in a south african university population. *African Vision and Eye Health*, 77(1):1–4, 2018.

Xu, A., McRae, A., Wang, J., Davenport, M., and Panan-jady, A. Perceptual adjustment queries and an inverted measurement paradigm for low-rank metric learning. *Advances in Neural Information Processing Systems*, 36, 2024.

Ying, Y., Huang, K., and Campbell, C. Sparse metric learning via smooth optimization. *Advances in neural information processing systems*, 22, 2009.

Yu, Y., Wang, T., and Samworth, R. J. A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 2015.

A. Definition of error cones

Definition A.1 (Error cone at reference point)**.** Given a reference point $\boldsymbol{z}_n$ with the corresponding estimated metric $\widehat{\boldsymbol{\Sigma}}_n$ and its second eigenvector $\widehat{\boldsymbol{u}}_2^{(n)}$, we define a cone $\mathcal{C}_n$ with cone angle α_n as the set

$$\mathcal{C}_n(\alpha_n) \coloneqq \left\{ \boldsymbol{z}_n + t\boldsymbol{R}(\beta)\widehat{\boldsymbol{u}}_2^{(n)} : t \geq 0, -\frac{\alpha}{2} \leq \beta \leq \frac{\alpha}{2} \right\}, \tag{14}$$

where $\boldsymbol{R}(\beta) \coloneqq \begin{bmatrix} \cos(\beta) & -\sin(\beta) \\ \sin(\beta) & \cos(\beta) \end{bmatrix}$ is a rotation matrix of angle β.

B. Practical considerations: eigenvector orientation.

In our proposed approach, a key step in estimating the copunctal point is to ensure that the estimated eigenvectors all are oriented in the general direction as the copunctal point. Because $\widehat{\boldsymbol{u}}_2$ and $-\widehat{\boldsymbol{u}}_2$ are both valid eigenvectors of estimated an metric $\widehat{\boldsymbol{\Sigma}}$, we need to ensure the correct direction is chosen. One such approach to is to consider pairs of estimated eigenvectors. For a small number of pairs, choose the orientation of the eigenvectors such that they intersect, then orient the rest of the eigenvectors in the same direction. Specifically in the color context, we can also choose eigenvector directions to point toward general areas that copunctal points in which are known to be located.

C. The geometry of pairs of error cone intersections

To bound the error of the estimator (1), we must compute the diameter of the intersection of N error cones, which is impracticable for general N, but is feasible if pairs of cones are considered at a time. As such, we work with two error cones $\mathcal{C}_i(\alpha_i)$ and $\mathcal{C}_j(\alpha_j)$ associated with reference points $\boldsymbol{z}_i$ and $\boldsymbol{z}_j$, respectively.

We begin by considering a rotated geometry, where the line connecting reference points $\boldsymbol{z}_i$ and $\boldsymbol{z}_j$ is parallel to the x-axis. We next define the *deviation angle*, denoted θ_i, of the i-th cone as the acute angle between the line connecting $\boldsymbol{z}_i$ and $\boldsymbol{z}_j$ and the estimated eigenvector at the center of the i-th cone, $\widehat{\boldsymbol{u}}_2^{(i)}$. Note that the value of the deviation angle θ_i depends on the pair of cones chosen, because it is measured with respect to the connecting line segment between the two reference points. For a visualization with these products defined, see Figure 6.

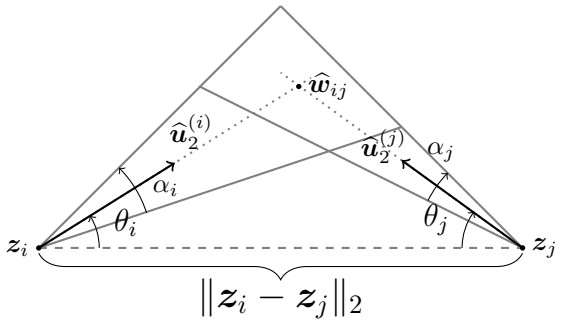

Figure 6: Geometry of the intersection of two error cones.