Inferring Individualized Color-Vision Distortions from fMRI Hue-Representation Geometry

Anonymous Authors[1]

Abstract

Many health-related data carry structured distortions that can act as phenotypes, yet interpreting such signatures and translating them into intervention remains challenging. Color vision deficiency (CVD), for example, distorts pairwise relations between neural hue representations while categorical recognition remains relatively preserved. Here we formulate fMRI-measured hue responses as structured neural representations and infer low-dimensional distortions that explain their geometry. We quantify this distortion using the Leave-One-Color-Out (LOCO) interpolation vulnerability profile, which measures how well each held-out hue can be interpolated from the remaining hue manifold. Applied to Ishihara-confirmed CVD participants across visual areas (V1–hV4), the framework recovers subject-specific distortions through both a one-parameter retinal cone-shift model and a two-parameter retinal–cortical model. Finally, we show that retinal and cortical models converge on detection performance but prescribe divergent corrections, yielding concrete predictions for planned behavioral validation. These findings show that detection-level agreement is insufficient for choosing an intervention and demonstrate that structured distortions in high-dimensional biological data can be resolved into low-dimensional, subject-specific models for personalized intervention.

1. Introduction

Many health-related data carry distorted structure that can act as a phenotype—systematic, repeatable patterns in how stimuli relate within a high-dimensional biological response space. Color vision deficiency (CVD), for example, produces distorted pairwise relations between neural hue representations even while categorical recognition remains relatively preserved. Other instances of such structured-phenotype data include geometric deformations of cardiac anatomy (Biffi et al., 2020) and directional shifts in gene-expression latent spaces under drug perturbation (Lotfollahi et al., 2019). Yet such phenotypes are challenging to interpret and, especially, to translate into individually-targeted intervention. When such structural distortion can be parameterized, it can be inverted for correction—a route absent from current generic correction approaches.

In a geometric framework, responses to stimuli can be summarized as a *representational geometry*, where pairwise distances capture what a neural population distinguishes (Kriegeskorte et al., 2008). In CVD, this geometry becomes the natural inference target: the pairwise relationships among hues can be systematically distorted even when categorical decoding remains preserved. Despite this perspective, structured distortions have rarely been used to derive individualized perceptual corrections; the closest analogue—a cochlear simulator that inverts a hearing-impairment deformation model (Irino, 2023)—targets simulation of sensory loss, not its correction.

CVD arises from altered spectral sensitivity of retinal cone photoreceptors, affecting approximately 8% of males (Machado et al., 2009), and produces structured distortions in cortical color representations that vary across individuals. Current assistive approaches, including commercial spectral-notch filters, apply a fixed, generic correction; recent evidence indicates that such filters alter color *appearance* but have minimal effect on color *discrimination* at threshold (Somers et al., 2024), suggesting that effective correction may require individually-tailored strategies.

While categorical color decoding can remain preserved—leave-one-run-out decoding shows no HC–CVD difference ($p = 0.668$)—continuous hue relationships may be systematically altered; classification accuracy is insensitive to this relational structure. CVD is also heterogeneous across individuals; even within a single diagnostic category, cortical compensation can vary substantially (Tregillus et al., 2021), motivating individual-level rather than group-level inference. In CVD, the pairwise geometry of stimulus-evoked

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Submitted to the ICML 2026 Workshop on Structured Data for Health. Do not distribute.

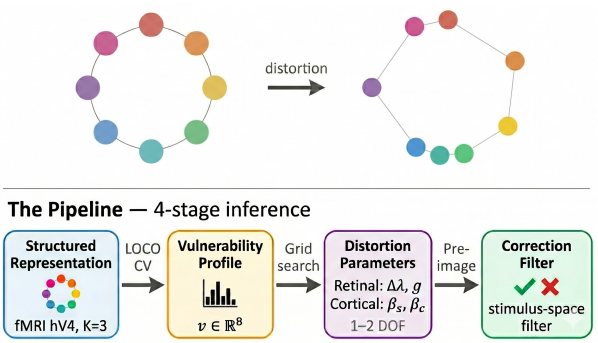

Figure 1. Structured distortion and the inference pipeline. **Top**: CVD distorts the continuous geometry of hue representations in hV4—some neighboring hues collapse while others expand—even though classification accuracy is preserved. **Bottom**: The inference pipeline quantifies the structured distortion via the LOCO vulnerability profile, fits parsimonious distortion parameters (1–2 DOF), and assesses correction feasibility via pre-image inversion.

fMRI responses is itself distorted in low-dimensional, parameterizable ways; we model this distortion as a mapping that can be fit and inverted.

Here we present a framework that quantifies this structured distortion through the interpolation vulnerability profile—how well each hue can be predicted from its neighbors—and fits low-dimensional simulator parameters that reproduce the observed profile. The fitted model can then be used to generate candidate corrections through pre-image search. The unit of inference is each individual's distortion model, not the population mean. Our contributions are:

1. We formulate distorted neural geometry as a structured-distortion inference problem and quantify it through the interpolation vulnerability profile, which captures subject-specific deficits in pairwise hue relations beyond classification accuracy.

2. We infer low-dimensional, interpretable distortion parameters (1–2 DOF) from high-dimensional neural representations, showing that different CVD individuals require different parsimonious models within a closed-testing escalation.

3. We show that the inferred distortions generate level-specific, falsifiable correction predictions: retinal and cortical models prescribe divergent corrections whose behavioral adjudication constitutes a concrete future test design.

2. Methods

2.1. Data Acquisition and Structured Representation

Ten participants (7 healthy controls [HC], 3 females, age 22.7 ± 2.5 years; 3 color vision deficient [CVD], all male: two deutan and one protan) were assessed using the 14-plate Ishihara test (Ishihara, 1917). All three CVD participants are Ishihara-confirmed; one deutan case (Sub-10) is null at both the LOCO and SRM representational levels used in this work and is therefore reserved as a future-work probe of alternative representational dimensions. Given the small CVD sample, all CVD results are interpreted as individual case demonstrations, not population-level effects (Crawford & Howell, 1998).

Stimuli consisted of eight colors equally spaced from 0° to 315° in CIE $L^*a^*b^*$ space ($L^* = 75$, chroma distance 40). Participants performed a modified Rapid Serial Visual Presentation (RSVP) task (Brouwer & Heeger, 2009) across six runs, viewing 1.5 s color patches. Functional images (Siemens 3T; TR = 1.5 s; 2 mm isotropic; 24 oblique slices) were preprocessed with fMRIPrep and projected to retinotopically-defined ROIs (V1, V2, hV4) via the Wang probabilistic atlas at 50% threshold (Wang et al., 2015). Response amplitudes were estimated using a two-stage FIR-then-GLM procedure (Brouwer & Heeger, 2009) and Procrustes-aligned across runs.

We modeled hue-selective neural populations with half-wave rectified and squared sinusoidal basis functions (Brouwer & Heeger, 2009): $r_k(\theta) = \max(0, \cos(\theta - \theta_k))^2$ for K channels. The weight matrix W mapping channels to voxel responses was estimated via ridge regression with GCV regularization (Golub et al., 1979). This encoding model transforms raw voxel data into a structured representation in which the relational geometry across hues is explicitly parameterized.

2.2. Structured Distortion Quantification: Interpolation Vulnerability

To capture distortions in relational geometry, we extract an interpolation vulnerability profile that reveals which hues are geometrically fragile in a participant's neural representation. In leave-one-color-out (LOCO) cross-validation, the encoding model W is trained on 7 of 8 hues and evaluated on the held-out hue; the 8 prediction errors form a vulnerability vector $\mathbf{v} \in \mathbb{R}^8$, capturing structural deficits invisible to categorical classification.

This profile is specific to continuous geometry: leave-one-*run*-out (LORO) classification shows no HC–CVD difference, while LOCO interpolation error is significantly elevated in CVD. Among the ROIs tested, the effect was most robust in hV4 ($p = 0.017$; K=3), consistent with hV4's established role in perceptual color representation (Brouwer

& Heeger, 2009; Bannert & Bartels, 2018). We therefore focus distortion inference on hV4, with V1 analyses reported as cross-ROI validation in Section A.

We additionally quantify pairwise structural distortion via representational distance matrices (RDMs). The 8×8 crossnobis distance matrix is computed in voxel space for each participant, and the difference $\Delta\text{RDM} = \text{RDM}_{\text{CVD}} - \overline{\text{RDM}}_{\text{HC}}$ measures how a CVD individual's pairwise geometry deviates from the HC mean. While $\mathbf{v}$ summarizes per-color vulnerability (8-dim), ΔRDM captures pairwise distortion ($\binom{8}{2} = 28$-dim)—used as a fitting regularizer (Equation (3)) and cross-criterion validation target (Section 3).

2.3. Inverse Inference of Low-Dimensional Distortion Fields

We infer low-dimensional distortion parameters from the observed profile $\mathbf{v}$ by testing two model families (Figure 1, bottom; notation in Section A.1). Both parameterize how a stimulus hue angle θ maps to a distorted angle θ', but at different processing levels: Family 1 targets retinal cone-shift; Family 2 targets cortical opponent-axis reorganization. Within Family 1, a 1-DOF model is tested first and escalated to 2-DOF only if non-significant (closed testing).

Machado cone shift (1 DOF). A retinal model (Machado et al., 2009): the anomalous cone's spectral peak is shifted by $\Delta\lambda \in [0, 20]$ nm, predicting hue distortion through altered L–M cone opponency.

Retinal + cortical gain (R+C; 2 DOF). Extends the cone shift with a cortical opponent-channel gain g:

$$o' = o_0 + (1+g) \cdot \underbrace{(o_{\Delta\lambda} - o_0)}_{\text{retinal shift}} \tag{1}$$

where o_0 and $o_{\Delta\lambda}$ are opponent-channel responses before and after the cone shift $\Delta\lambda$. Setting $g{=}0$ recovers pure Machado; $g{=}{-}1$ is exact cortical compensation; $g{<}{-}1$ is overcompensation—motivated by evidence for progressive cortical amplification in anomalous trichromats (Tregillus et al., 2021).

Cortical opponent-axis model (2 DOF). The second family models distortion at the cortical level as angular dilation along two opponent axes:

$$\theta'(c) = \theta(c) + \underbrace{\beta_s \cos(\theta(c){-}90^\circ)}_{\text{S-cone axis}} + \underbrace{\beta_c \cos(\theta(c){-}\theta_{\text{conf}})}_{\text{confusion axis}} \tag{2}$$

where β_s captures S-cone axis expansion, motivated by behavioral evidence for hue-angle reorganization in CVD (Emery et al., 2021), and $\theta_{\text{conf}} = 16^\circ$ (protan) or 150° (deutan) is the type-specific confusion axis.

Table 1. Inferred distortion parameters (hV4 LOCO, cone-based family). Within the cone-based family, the 1-DOF Machado model is tested first; the 2-DOF R+C is tested only if Machado is non-significant (closed testing). Sub-10 is null at this representational level (LOCO and SRM).

Subject	Model	DOF	ρ	p_{perm}
Sub-08 (deutan)	R+C	2	0.857	**0.005****
Sub-09 (protan)	Machado	1	0.762	**0.018***
Sub-10 (deutan)	—	—	−0.048	0.559

2.4. Specificity and Validation

We select optimal parameters by forward-simulating each model's distortion through the HC encoding model and comparing the resulting profile to the observed $\mathbf{v}$. The composite fitting loss balances per-color functional accuracy with pairwise structural coherence:

$$L = \underbrace{L_{\text{vuln}} + 0.5\, L_{\text{rank}}}_{\text{per-color (8-dim)}} + \underbrace{0.2\, L_{\text{rdm}}}_{\text{pairwise (28-dim)}} + 0.1\, L_{\text{smooth}} \tag{3}$$

where L_{vuln} is MSE between simulated and observed profiles, $L_{\text{rank}} = 1 - \rho$ penalizes rank disagreement, L_{rdm} encourages pairwise distance coherence, and L_{smooth} regularizes adjacent-color shifts. Effective DOF are 1–2, well below data dimensionality of 8; parameters are optimized over grids. Statistical significance is assessed by exact permutation of the 8-color labels ($8! = 40{,}320$) on the Spearman ρ between the best-fit simulated and observed vulnerability vectors. A valid model must reach significance ($p < 0.05$) on the per-color vulnerability profile; HC-based false-positive calibration is reported separately in Section A.2.

3. Results

3.1. Subject-Specific Distortion Inference

Table 1 and Figure 2 present the central result: the framework predicts each participant's observed LOCO vulnerability profile significantly above the exact $8! = 40{,}320$ label-permutation null for two of the three Ishihara-confirmed CVD participants, with distinct low-dimensional distortions per subject; the third (Sub-10, deutan) is correctly null at both the LOCO and SRM representational levels.

Sub-09 (moderate protanomaly). The 1-DOF Machado model with $\Delta\lambda = 13.5$ nm predicts Sub-09's profile at $\rho = 0.762$ ($p_{\text{perm}} = 0.018$). The fitted $\Delta\lambda$ falls within the established range for moderate–severe protanomaly (Machado et al., 2009). Adding a cortical gain parameter does not improve fit ($g = 0.0$ at optimum), indicating that a purely retinal model is sufficient.

Sub-08 (mild deuteranomaly). The 1-DOF Machado model is only trend-level ($p_{\text{perm}} = 0.058$); following closed

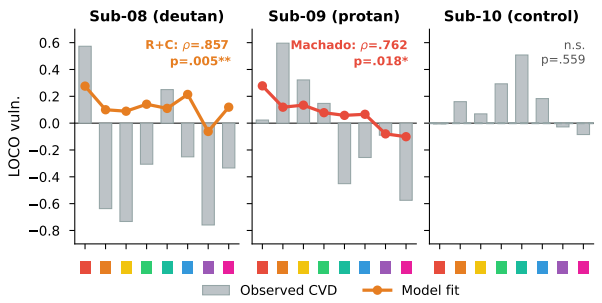

Figure 2. Subject-specific LOCO vulnerability profiles (hV4). Gray bars: observed CVD profile; colored line: best-model prediction. **Sub-08** (left): R+C (ρ=0.857, p_{perm}=0.005). **Sub-09** (middle): Machado (ρ=0.762, p_{perm}=0.018). **Sub-10** (right): flat profile (ρ=−0.048, p_{perm}=0.559, correct null). **Both Sub-08 and Sub-09 achieve statistically significant prediction of the observed CVD distortion structure against the exact** 8! **label-permutation null, while Sub-10 is correctly null.** See Figure 1 for pipeline overview.

testing, we escalate to the 2-DOF R+C model, which predicts Sub-08's profile at $\rho = 0.857$ ($p_{\text{perm}} = 0.005$) with $\Delta\lambda = 2.0$ nm and $g = +2.25$. The cortical gain may reflect amplification of a weak retinal signal, though it could also absorb unmodeled variance (see Section A).

Sub-10 (deutan). All models yield $\Delta\lambda \approx 0$ ($\rho = -0.048$, $p_{\text{perm}} = 0.559$, correct null). Sub-10 is also null at the SRM criterion, bounding the sensitivity of the representational levels used here.

3.2. Convergent Detection, Divergent Correction

The cortical-level 2-component model (Equation (2)) also reaches significance in hV4 for both fit-recovered participants: sub-08 ($\beta_s = 38°$, $\beta_c = -14°$; $\rho = 0.881$, $p = 0.004$) and sub-09 ($\beta_s = 6°$, $\beta_c = -22°$; $\rho = 0.690$, $p = 0.035$); sub-10 is trend-level under this model ($p = 0.058$) but remains null at the primary LOCO criterion. The two families *converge on detection*—both identify the same two individuals as distorted—but *diverge on correction*: for sub-08, the cone-based and cortical models prescribe correction directions with cosine $= -0.18$ (sign agreement 3/8), despite comparable fit quality ($\rho = 0.857$ vs. 0.881). This divergence is a falsifiable prediction: the retinal model prescribes spectral pre-distortion, while the cortical model prescribes opponent-axis rescaling; a behavioral 2AFC task would simultaneously validate the correction and localize the dominant processing level.

Correction feasibility itself is severity- and level-dependent: at the retinal level, sub-08's mild distortion admits exact pre-image inversion for all 8 colors, while sub-09's moderate arc compression ($360° \rightarrow \sim 96°$) makes exact inversion structurally impossible for 4/8 colors. The cortical-level model, by contrast, preserves bijectivity by construction and admits exact solutions for both participants (details in Section A).

4. Discussion

From structured distortion to intervention design. The LOCO vulnerability profile compresses high-dimensional voxel data into an 8-element geometry-sensitive summary. Combined with a composite loss that decomposes geometric fidelity into per-color and pairwise components (Equation (3)), this enables distortion inference that is both interpretable and invertible into candidate corrections. Future work could replace grid search with amortized inference (e.g., conditional normalizing flows) for real-time distortion estimation. This approach may generalize to other conditions where sensory representations are distorted rather than absent.

Limitations. We have *not* shown that corrections improve perceptual discrimination—behavioral validation is the framework's primary next step. Three CVD participants are insufficient for population-level claims; this work is a proof-of-concept, and the fMRI requirement limits scalability. Rank-based grid optimization on 8 data points yields elevated HC false-positive rates (15/21 model–subject pairs significant; Section A.2); the pipeline targets individual profiling of *known* CVD cases, not screening. The fit-recovered CVD cases are distinguished from HC fits by physiologically interpretable parameters and cross-criterion convergence (Table 4).

5. Conclusion

We present a framework that quantifies structured distortions in fMRI hue representations and fits parsimonious low-dimensional parameters, recovering subject-specific distortions that converge on detection but diverge on correction—generating falsifiable behavioral predictions as the immediate next step.

Impact Statement

The broader implication of this work is methodological: many health conditions leave structured rather than absent signatures in high-dimensional measurements. Mapping such signatures onto low-dimensional, invertible geometric features could turn descriptive phenotypes into actionable, individualized intervention targets.

References

Bannert, M. M. and Bartels, A. Human V4 activity patterns predict behavioral performance in imagery of object color.

Journal of Neuroscience, 38(15):3657–3668, 2018.

Biffi, C., Oktay, O., Tarroni, G., Bai, W., De Marvao, A., Doumou, G., Rajchl, M., Bedair, R., Prasad, S., Cook, S., O'Regan, D., and Rueckert, D. Explainable anatomical shape analysis through deep hierarchical generative models. *IEEE Transactions on Medical Imaging*, 39(6): 2088–2099, 2020. doi: 10.1109/TMI.2020.2964499.

Brouwer, G. J. and Heeger, D. J. Decoding and reconstructing color from responses in human visual cortex. *Journal of Neuroscience*, 29(44):13992–14003, 2009.

Crawford, J. R. and Howell, D. C. Comparing an individual's test score against norms derived from small samples. *The Clinical Neuropsychologist*, 12(4):482–486, 1998.

Emery, K. J., Volbrecht, V. J., Peterzell, D. H., and Webster, M. A. Variations in normal color vision. VII. relationships between color naming and hue scaling. *Vision Research*, 183:1–12, 2021.

Golub, G. H., Heath, M., and Wahba, G. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21(2):215–223, 1979.

Irino, T. Hearing impairment simulator based on auditory excitation pattern playback: WHIS. *IEEE Access*, 11:78419–78430, 2023. doi: 10.1109/ACCESS.2023.3298673.

Ishihara, S. *Tests for Colour-Blindness*. Handaya, Tokyo, 1917.

Kriegeskorte, N., Mur, M., and Bandettini, P. Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2:4, 2008. doi: 10.3389/neuro.06.004.2008.

Lotfollahi, M., Wolf, F. A., and Theis, F. J. scGen predicts single-cell perturbation responses. *Nature Methods*, 16: 715–721, 2019. doi: 10.1038/s41592-019-0494-8.

Machado, G. M., Oliveira, M. M., and Fernandes, L. A. F. A physiologically-based model for simulation of color vision deficiency. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1291–1298, 2009.

Somers, L. P., Franklin, A., and Bosten, J. M. Empirical tests of the effectiveness of EnChroma multi-notch filters for enhancing color vision in deuteranomaly. *Vision Research*, 218:108390, 2024. doi: 10.1016/j.visres.2024.108390.

Tregillus, K. E. M., Isherwood, Z. J., Vanston, J. E., Engel, S. A., and MacLeod, D. I. A. Color compensation in anomalous trichromats assessed with fMRI. *Current Biology*, 31(5):936–942, 2021.

Wang, L., Mruczek, R. E. B., Arcaro, M. J., Bhatt, S., and Kastner, S. Probabilistic maps of visual topography in human cortex. *Cerebral Cortex*, 25(10):3911–3931, 2015.

A. Supporting Analyses

A.1. Notation

Table 2. Key notation.

Symbol	Description
$\theta(c)$, $\theta'(c)$	Stimulus / distorted hue angle for color c
$\mathbf{v} \in \mathbb{R}^8$	LOCO vulnerability profile (per-color interp. error)
W	Encoding weight matrix (K channels $\rightarrow$ voxels)
$\Delta\lambda$	Retinal cone-shift magnitude (nm)
o_0, $o_{\Delta\lambda}$	Opponent response: healthy / after cone shift
g	Cortical gain (0: retinal only; -1: full compensation)
β_s, β_c	S-cone / confusion-axis angular dilation ($^\circ$)
ΔRDM	Pairwise distance deviation from HC mean
ρ	Spearman corr. (simulated vs. observed $\mathbf{v}$)

A.2. HC Specificity Calibration

To empirically calibrate the false-positive rate of the grid-search fitting procedure, we applied the full hV4 LOCO pipeline to all 7 HC participants, treating each as if they were a CVD candidate (best-family selected per model).

Results. Table 3 reports the number of HC participants reaching $p < 0.05$ for each model. The observed false-positive rate substantially exceeds the nominal 5%: 15/21 model–subject tests are significant, compared to the expected 1.05 under the null. The 2-component model is worst (7/7), Machado is least inflated (3/7).

Table 3. HC false-positive rates (hV4 LOCO, $p < 0.05$).

Model	DOF	HC significant	FPR
Machado	1	3/7	43%
R+C	2	5/7	71%
2-Component	2	7/7	100%
Total		15/21	71%

Mechanism. The elevated FPR arises from two reinforcing factors:

1. **Grid-search flexibility on 8 data points.** With 1–2 DOF optimized over a fine grid (41–1,326 candidate parameter configurations) and rank-based loss ($L_{\text{rank}} = 1 - \rho$), most individuals—healthy or not—will have some configuration that achieves a nominally significant rank correlation with 8 elements. This is a known property of permutation tests under flexible model search.

2. **Regression to the mean.** HC individuals with low baseline LOCO performance (i.e., already poorly interpolated before any model fitting) have the most room for apparent improvement. We observe a strong negative correlation ($r = -0.894$) between HC baseline ρ and the improvement $\Delta\rho$ achieved by fitting, confirming that the signal captured in HC fits is predominantly regression to the mean rather than color-specific structure.

Why this does not invalidate the fit-recovered CVD results. The pipeline is designed for *individual profiling* of known CVD cases, not population screening; we do not claim specificity against healthy color vision from the present case series. The two fit-recovered CVD cases are distinguished from HC fits by three properties absent in HC: (i) inferred parameters fall within established physiological ranges (Sub-09 $\Delta\lambda = 13.5$ nm matches moderate protanomaly; Sub-08 $\Delta\lambda = 2.0$ nm matches mild deuteranomaly); (ii) the 2-component model shows cross-criterion convergence (LOCO and ΔRDM both significant for both fit-recovered participants; no HC achieves this); and (iii) the 2-component model's axes—S-cone and type-specific confusion—are grounded in known CVD physiology (Emery et al., 2021), constraining the fit to physiologically interpretable dimensions rather than arbitrary directions. In sum, HC fits reflect opportunistic overfitting to 8 data points,

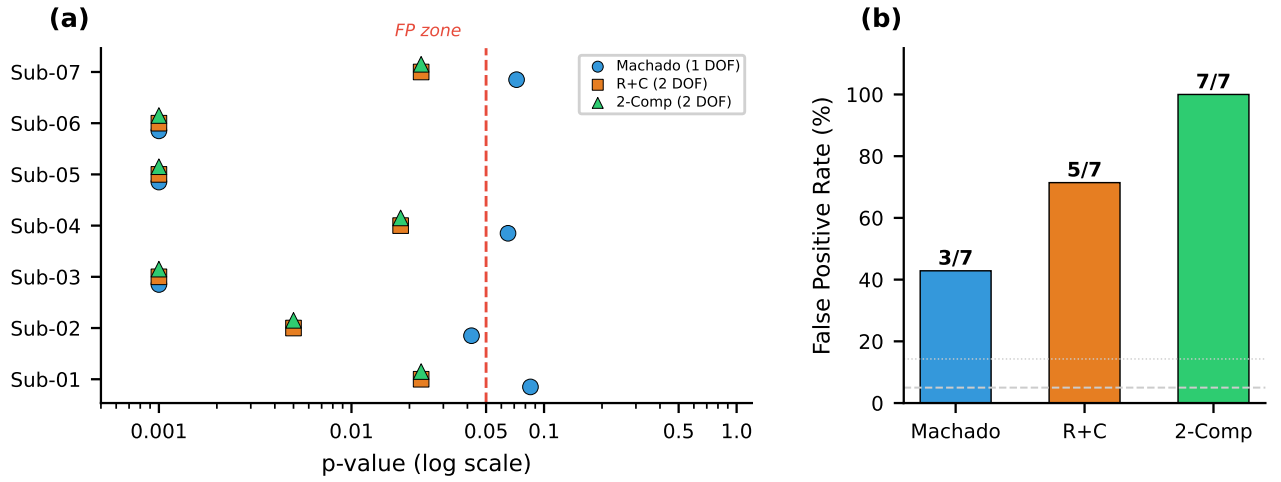

Figure 3. HC specificity calibration. **(a)** Per-subject p-values for each model (log scale). The red dashed line marks $\alpha = 0.05$; shaded region indicates the false-positive zone. **(b)** Aggregate false-positive rates. Dashed line: nominal $\alpha = 5\%$; dotted line: chance rate for n=7 (14.3%). Higher model complexity (DOF) yields higher false-positive rates.

whereas the fit-recovered CVD cases align with convergent external evidence. Sub-10 is Ishihara-confirmed CVD but is null at both the LOCO and SRM representational levels used here; probing this case at alternative representational dimensions is an explicit future-work direction.

A.3. Cross-ROI and Cross-Criterion Evidence

Beyond the primary hV4 LOCO inference, the 2-component model was validated across multiple ROIs and criteria. Sub-08 V1 LOCO reached $p = 0.001$ ($\beta_s = 50°$, $\beta_c = -14°$)—the strongest LOCO result in the pipeline. Sub-09 V1 ΔRDM reached $p = 0.007$ (crossnobis cosine). The 2-component model is the only model achieving significance under both LOCO and ΔRDM criteria for both CVD participants (Table 4).

A.4. Model-Dependent Correction Feasibility

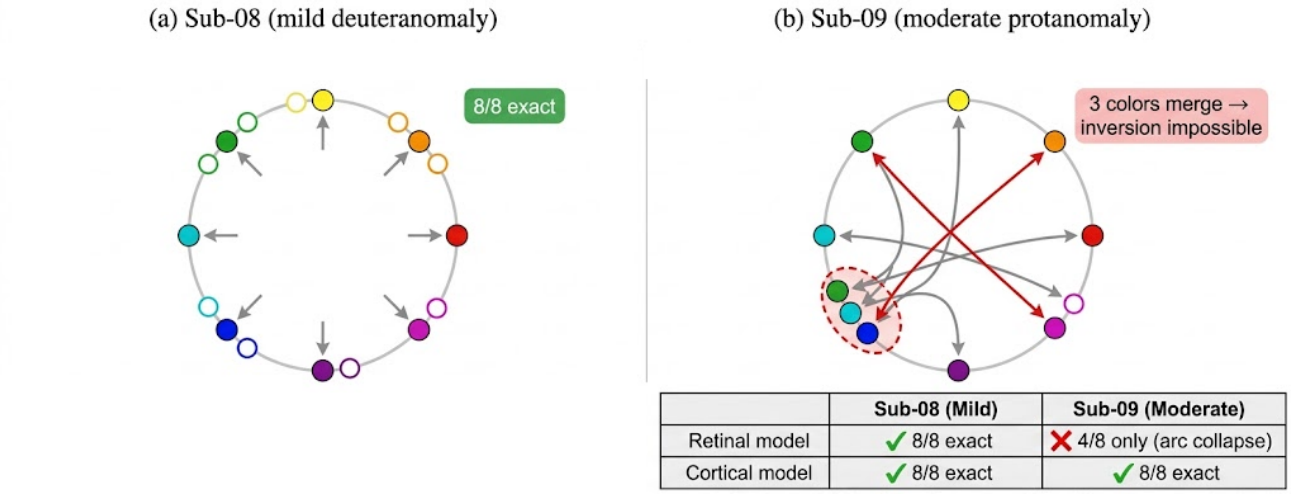

	Sub-08 (Mild)	Sub-09 (Moderate)
Retinal model	✓ 8/8 exact	✗ 4/8 only (arc collapse)
Cortical model	✓ 8/8 exact	✓ 8/8 exact

Figure 4. Correction geometry depends jointly on severity and modeling level. **(a)** Mild distortion (Sub-08): wide opponent arc preserved; all 8 colors admit exact pre-image inversion ($< 0.001°$ error). **(b)** Moderate distortion (Sub-09): arc compresses $360° \rightarrow \sim 96°$; three colors merge, making exact inversion impossible. The retinal model achieves exact correction only for mild distortion (4/8 for moderate), while the cortical 2-component model preserves bijectivity for both (bottom table).

Under the cortical-level 2-component model, both CVD participants admit exact pre-image solutions for all 8 colors (sub-08: mean $|\delta| = 46.3°$; sub-09: mean $|\delta| = 20.1°$; residual $< 0.001°$). This contrasts with the retinal Machado model, where sub-09's arc compression makes exact inversion impossible for 4/8 colors. The difference arises because angular dilation in opponent hue space preserves bijectivity by construction, whereas retinal spectral shifts can compress the opponent arc irreversibly.

For sub-08, the R+C and 2-component models prescribe substantially different correction directions (cosine between correction vectors = -0.18; sign agreement 3/8), despite comparable fit quality (R+C $\rho = 0.857$ vs. 2-component $\rho = 0.881$). Which correction actually improves perception requires behavioral validation.

A.5. Model Comparison: Full Table

Table 4. Full model comparison across hV4 LOCO, V1 LOCO, and V1 ΔRDM criteria. The 2-component model is the only model achieving significance under both LOCO and ΔRDM for both CVD participants.

Subject	Model	DOF	hV4 LOCO ρ	hV4 LOCO p	V1 LOCO p	V1 ΔRDM	Notes
Sub-08	Machado	1	0.619	0.058	0.033*	anti-corr.	Trending (hV4)
	R+C	2	**0.857**	**0.005****	0.047*	0.179	hV4 primary
	2-Comp	2	0.881	0.004**	**0.001*****	CI excl. 0	Cross-criterion
Sub-09	Machado	1	**0.762**	**0.018***	0.112	$+0.091$	hV4 primary
	R+C	2	0.762	0.018*	—	0.026*	g=0 (= Machado)
	2-Comp	2	0.690	0.035*	0.018*	**0.007*****	Cross-criterion
Sub-10	Cone-level	—	-0.048	0.559	—	—	Null at LOCO/SRM
	2-Comp	2	—	0.058†	—	—	Trend-level

† Trend-level 2-component fit while primary LOCO criterion is null; see Limitations discussion.

A.6. Correction Feasibility Details

Sub-08 (R+C pre-image). All 8 colors admit exact pre-image solutions (mean $|\delta| = 23.5°$, max $= 42.9°$, residual $< 0.001°$). The resulting 8-point lookup table is directly implementable on any digital display.

Sub-09 (Machado pre-image). At $\Delta\lambda = 13.5$ nm, the opponent-space arc compresses from 360° to ∼96°. Colors c4 (green), c5 (cyan), and c6 (blue) converge to a single perceived angle (∼282°). Exact pre-image: 4/8 colors. Separation-optimized fallback: min separation $1.03° \rightarrow 5.76°$ ($5.6\times$), mean separation $39.3° \rightarrow 24.3°$ (healthy: 70.7°).

Sub-09 (2-component pre-image). All 8 colors admit exact pre-image solutions (mean $|\delta| = 20.1°$, max $= 48.1°$, residual $< 0.001°$). The cortical-level angular dilation preserves bijectivity, unlike retinal spectral shift, enabling exact correction where the Machado model predicts irreversible collapse.