



RECOLOR: A MOBILE VISION ASSISTIVE TOOL FOR COLOR VISION DEFICIENCY USING ADAPTIVE DALTONIZATION AND IMAGE PROCESSING ALGORITHMS

Mark Risen B. Caparas

Project Manager

Lyceum of the Philippines University–Cavite

Cavite, Philippines

Email: caparasmarkrisen.mrc@gmail.com

Charles Justine B. Rull

Algorithm Engineer

Lyceum of the Philippines University–Cavite

Cavite, Philippines

Email: charlesrull18@gmail.com

Johann Nikkolai L. Cepeda

Front-End/Back End Engineer

Lyceum of the Philippines University–Cavite

Cavite, Philippines

Email: nixxd0617@gmail.com

Stephen L. Lacs

UI/UX Designer

Lyceum of the Philippines University–Cavite

Cavite, Philippines

Email: princepipen@gmail.com

Arcell R. Hadlocon, LPT, MIT

Thesis Adviser

Lyceum of the Philippines University–Cavite

Cavite, Philippines

Email: arcellhadlocon1984@gmail.com

Corresponding Author

Arcell R. Hadlocon, LPT, MIT

Thesis Adviser

Lyceum of the Philippines University–Cavite

Cavite, Philippines

Email: arcellhadlocon1984@gmail.com

Abstract

Color vision deficiency (CVD) affects roughly 1 in 12 males, yet mobile tools for people with CVD usually address a single function—screening, simulation, or enhancement—and the Philippine National Vision Screening Program does not currently cover CVD. To design, implement, and evaluate ReColor, an Android application that combines a digitized Ishihara screening module, adaptive color enhancement, CIELAB-based color identification, and GPU-accelerated CVD simulation. ReColor was developed through a Modified Agile–Spiral process (15



sprints). We deployed four deterministic pipelines: full-image daltonization by error redistribution, HSV hue rotation as a comparative baseline, CIELAB nearest-neighbor color naming, and a Skia (SkSL) shader for simulation; we trained a U-Net segmentation model as a non-deployed research asset. We evaluated algorithms against synthetic and standard ground truths (verified confusion pairs, X-Rite ColorChecker Classic 24, 30 W3C named colors, and a float64 CPU reference) through system testing and acceptability ratings from 5 participants with CVD, 30 general users, 10 IT experts, and 3 eye-care professionals. Daltonization produced larger discrimination gains than hue rotation for every CVD type (mean +6.91 vs +1.88 ΔE_{00} across 156 verified confusion pairs), although neither method reached SSIM > 0.90 for most non-neutral test images. The color identifier classified 54 of 54 reference colors correctly, and the GPU shader deviated from the CPU reference by a mean ΔE_{00} of 0.027. All executed system tests passed. Mean acceptability on a 4-point scale was 3.92 (clinical experts), 3.60 (general users), 3.38 (participants with CVD), and 3.33 (IT experts); participants with CVD rated the enhancement module lowest (2.73). The Philippine Eye Research Institute (PERI) approved the Ishihara module with revisions; its diagnostic accuracy has not yet been established. ReColor is a feasible integrated screening-and-assistance platform. Evidence of functional benefit for people with CVD is preliminary ($n = 5$) and rests on perception ratings rather than task performance. Next steps include a clinical pilot with normal-vision controls, behavioral testing of enhancement, and on-device latency benchmarking.

Keywords: *color vision deficiency; daltonization; CIELAB; Ishihara screening; mobile assistive technology; image processing; React Native*



Introduction

Color vision deficiency (CVD) impairs cone-mediated color discrimination. Its commonest form, congenital red–green deficiency, is X-linked and affects up to about 8% of males and 0.5% of females (Birch, 2012), with deutan defects predominating (Jeong et al., 2025). In the Philippines, a screening survey of male students in a public high school found a CVD prevalence of 5.17% (65 of 1,258), with 98% of them deutan (Cruz et al., 2010). People with CVD have difficulty reading color-coded maps and charts, interpreting signals, and choosing clothing, and several occupations—aviation and other transport roles, electrical work, and military service among them—set color-vision standards (Rodriguez-Carmona et al., 2012).

CVD is usually detected with pseudoisochromatic plates such as the Ishihara test, which is quick, inexpensive, and sensitive for red–green defects but does not grade severity (Birch, 1997); the anomaloscope remains the reference standard but is costly and seldom available locally (Garcia et al., 2022). Smartphone versions of the Ishihara test have been compared with the printed booklet (Khizer et al., 2022), which makes a phone a plausible screening platform provided display variability is controlled.

Assistive options include tinted optical filters, which are costly and lighting-dependent (Almutairi & Alharbi, 2022; Pattie et al., 2022), and software methods that (a) simulate dichromatic vision (Brettel et al., 1997; Viénot et al., 1999; Machado et al., 2009), (b) recolor images by shifting hues away from confusion axes, or (c) “daltonize” them by adding the information a dichromat cannot perceive back into channels that the person can use (Fidaner et al., 2005). Widely used apps each cover one of these functions: Color Blind Pal names colors and applies hue-shift filters, Daltonize Me Camera performs real-time daltonization, and the Chromatic Vision Simulator simulates CVD for designers and educators (Fiorentini, 2015; Lima, 2014; Asada, 2012). In the authors’ review of these three apps, none combined screening, enhancement, color identification, and education in one application; no systematic app review was reported.

Policy context matters in the Philippines. Republic Act 11358 established the National Vision Screening Program for kindergarten learners, implemented by the Department of Education and the Department of Health with PERI as the technical center (PERI, 2019). Current protocols focus on refractive error and amblyopia and do not include CVD.

Two design questions motivated this study. First, learned segmentation models such as U-Net (Ronneberger et al., 2015) can localize color classes but add model size and latency on mid-range phones, whereas deterministic methods are fast and predictable. Second, most enhancement studies judge output by its similarity to a normal-vision reference (ΔE , SSIM), a target that cannot be reached when a cone type is missing; a task-relevant question is whether previously confusable colors become distinguishable.

We therefore designed, implemented, and evaluated ReColor with five objectives:

1. to implement an adaptive color enhancement and identification pipeline comprising (a) daltonization by error redistribution, (b) CIELAB ΔE color naming against a reference database, and (c) GPU-accelerated CVD simulation;
2. to add multimodal (haptic and auditory) feedback for non-visual support;
3. to include career-awareness content on professions that require normal color vision;
4. to evaluate technical performance with CIEDE2000 ΔE , structural similarity (SSIM), ground-truth accuracy, and processing latency; and
5. to assess acceptance and software quality among users with CVD, general users, IT experts, and eye-care professionals, using Android core app quality criteria and WCAG 2.2.

Methods

Study design and development process

This was a mixed-methods design-and-development study conducted from August 2025 to May 2026. ReColor was built using a Modified Agile–Spiral process with weekly sprints (15 sprints across four phases: planning and foundation, core feature development, integration and testing, and evaluation and finalization). PERI served as the clinical partner and specified the Ishihara plate timing, scoring thresholds, and test modes.

System architecture and scope

ReColor is a React Native application that uses react-native-vision-camera for capture, @shopify/react-native-skia for GPU shaders, and Firebase (Authentication and Firestore, with offline persistence) with AsyncStorage for local settings. A five-layer architecture separates presentation, navigation, algorithmic processing, data persistence, and deployment; 22 screens are grouped into authentication, core features, screening, data collection, and education. Ishihara results are written twice: to a private, user-linked history and to a separate anonymized research collection, in keeping with the Data Privacy Act of 2012 (RA 10173). The interface never relies on color alone: states, alerts, and navigation carry text labels, icons, or haptic cues, and a dismissible banner states that the Ishihara module is not a clinical diagnosis. Design followed the Android core app quality guidelines (Android Developers, 2023) and WCAG 2.2 (W3C, 2023).

All three camera-based features—enhancement, identification, and simulation—operate on captured still images rather than live video; images are normalized to a width of 1040 px for enhancement and simulation and 640 px for identification. The design scope is protanomaly and deuteranomaly, the most prevalent subtypes locally; tritan modes are implemented but were not a design target. The target platform is a mid-range Android phone with an 8-MP autofocus rear camera and a GPU supporting OpenGL ES 3.0 or Vulkan, used at roughly 300–500 lux; accuracy may degrade below 50 lux. Figure 1 summarizes the deployed pipelines.

Figure 1. Deployed ReColor pipelines. Camera and gallery images feed simulation, daltonization, hue rotation, and color identification; Ishihara plates are static images scored separately. The U-Net model is not part of the deployed pipeline.




Image-processing pipelines

Gamma handling

Pixels are converted from sRGB to linear light with a $\gamma = 2.2$ power law in the shader and enhancement pipelines; the CIELAB conversion uses the exact piecewise sRGB transfer function.

CVD simulation

The simulation projects linear RGB onto the confusion plane of the missing cone type with a 3×3 matrix M_t (Table 1), followed by clamping and re-encoding. For camera and freeze-frame images, the matrix is applied by an SkSL runtime shader in half precision; a 1040×780 image is rendered in about 4 ms, against about 400 ms for an equivalent JavaScript pixel loop. Gallery images use a double-precision CPU loop with the same matrices.

Table 1. Simulation matrices M_t (linear RGB)

CVD type	Output row	R coefficient	G coefficient	B coefficient
Protan	R	0.152286	1.052583	−0.204868
	G	0.114503	0.786281	0.099216
	B	−0.003882	−0.048116	1.051998
Deutan	R	0.367322	0.860646	−0.227968
	G	0.280085	0.672501	0.047413
	B	−0.011820	0.042940	0.968881
Tritan	R	1.255528	−0.076749	−0.178779
	G	−0.078411	0.930809	0.147602
	B	0.004733	0.691367	0.303900

Daltonization

For each pixel in linear RGB, the algorithm (i) simulates the deficiency, $RGB_{sim} = M_t \cdot RGB_{lin}$; (ii) computes the error, $E = RGB_{lin} - RGB_{sim}$, which represents information the user cannot access; and (iii) redistributes it, $RGB_{out} = RGB_{lin} + S_t \cdot E$, scaled by the user-set enhancement intensity, then clamps and re-encodes. The redistribution matrices S_t (Table 2), following Fidaner et al. (2005), zero the row of the missing channel and route error into the remaining channels with weights of 0.7 (protan, tritan) or 0.6 (deutan). No segmentation mask is used in the deployed system; because perceivable pixels produce near-zero error, the method “self-gates”, leaving them essentially unchanged.

Hue rotation

Hue rotation, the comparative baseline, converts pixels to HSV, skips near-gray pixels ($S < 0.15$), and rotates hues inside a CVD-specific confusion band by a shift that is maximal at the band center and tapers linearly to zero at its edges (weight = $1 - \text{distance}/\text{range}$, with circular hue distance). Parameters are in Table 2.

Table 2. Enhancement parameters by CVD type

CVD type	Redistribution matrix S_t (rows R, G, B; columns E_R, E_G, E_B)	Hue-rotation center	Band width	half-	Max shift
Protan	[0 0 0; 0.7 1 0; 0.7 0 1]	0° (red)	$\pm 60^\circ$		+40°
Deutan	[1 0.6 0; 0 0 0; 0 0.6 1]	0° (red)	$\pm 60^\circ$		+40°
Tritan	[1 0 0.7; 0 1 0.7; 0 0 0]	240° (blue)	$\pm 60^\circ$		-30°

Color identification

A tapped pixel is converted sRGB $\rightarrow$ XYZ $\rightarrow$ CIELAB and matched to its nearest entry in a reference database of named shades, each assigned to one of ten color families (Red, Orange, Yellow, Green, Cyan, Blue, Violet, Pink, Brown, Neutral) following the ISCC–NBS system (Kelly & Judd, 1955). The distance is a weighted CIE76-type ΔE with the L^* term down-weighted by 0.5, so that dark and light variants of a hue receive the same name; pixels with chroma $C^* < 12$ are matched only to Neutral to prevent noise from producing spurious color names. Match confidence is derived from the CIEDE2000 distance to the nearest entry (100% at $\Delta E = 0$, 0% at $\Delta E \geq 50$). The result is announced by text-to-speech and a vibration.

Ishihara screening module

The module digitizes the Ishihara plates in a Quick Mode (14 plates, 10 scored) and a Comprehensive Mode. Plates are scored with category weights (Table 3), shown for 3 s each in Fisher–Yates shuffled order after a demonstration plate, and require a minimum screen brightness. Severity is set by weighted percentage ($\geq 80\%$ normal; 65–79% indeterminate; 45–64% mild; 25–44% moderate; $< 25\%$ severe), and Protan/Deutan classification uses the diagnostic plates. Results carry a referral prompt and a non-diagnostic disclaimer.

Table 3. Ishihara plate categories and scoring weights (Comprehensive Mode)

Category	Plates	Weight	Purpose
Demonstration	1	0	Verification; excluded from scoring
Screening	2–9, 26–38	1	Detect red–green CVD
Vanishing	10–17	1	Visible to normal vision only; failure to read suggests CVD



Category	Plates	Weight	Purpose
Hidden	18–21	1	Visible to CVD only
Diagnostic	22–25	2	Protan vs deutan differentiation

U-Net research model

A U-Net with three encoder–decoder levels (32/64/128 feature maps) and a 10-class softmax output was trained for per-pixel color-class segmentation on a synthetic 10-class color-patch dataset with augmentation, then converted to TensorFlow Lite with INT8 quantization (about 8 MB). It is bundled as a research asset but not invoked at runtime.

Evaluation

Enhancement fidelity

We applied both enhancement algorithms to six synthetic images (red–green gradient, blue–yellow gradient, ColorChecker-24 chart, natural scene, saturated primaries, neutral grays) for each of the three CVD types. CIEDE2000 ΔE (ΔE_{00}) and SSIM (Wang et al., 2004) were computed between the CVD-perceived enhanced image and the original normal-vision image, with targets of $\Delta E_{00} < 2.0$ and $SSIM > 0.90$. Because normal-vision equivalence is physiologically unattainable, these values were treated as relative comparisons.

Discrimination gain

The primary effectiveness metric used pairs of colors that are distinguishable under normal vision ($\Delta E_{00} > 8$ in Set A; > 15 in Set B) but confused after simulation of the target deficiency ($\Delta E_{00} < 5$, a threshold the thesis attributes to Sharma & Bala (2002)). Each pair was enhanced, re-simulated with the same matrix, and its discrimination gain computed as the post-enhancement minus pre-enhancement ΔE_{00} ; a gain above 2.0 was counted as meaningful. Set A comprised 156 pairs (60 protan, 60 deutan, 36 tritan) and Set B 54 pairs (20, 20, 14).

Color identifier, simulation fidelity, and system tests

The identifier was tested on the 24 patches of the X-Rite ColorChecker Classic (published CIELAB values) and 30 W3C CSS Level 3 named colors, 54 samples in total across the ten families, reporting overall and per-class accuracy, a confusion matrix, and confidence. Simulation fidelity was assessed in two ways: (i) the float16 GPU shader against a float64 CPU implementation with identical matrices (six 720×480 images $\times$ three CVD types = 18 conditions; pass if mean $\Delta E_{00} < 1.0$) and (ii) app output against a reference computation using the piecewise sRGB transfer function for 12 test colors $\times$ three CVD types. System testing comprised consolidated module tests, end-to-end functionality tests, and Android core tests, run on five devices spanning entry-level to upper mid-range chipsets; stability was also recorded on the devices of one general user and one participant with CVD.

Participants and instruments

Forty-eight participants were recruited by purposive and convenience sampling through PERI's clinical network, social media, and LPU–Cavite (Table 4). PERI advised a minimum of five participants with CVD for a pilot.



Testing was distributed in person and online, at participants' own locations under normal indoor lighting, with a minimum screen brightness of 80% for the Ishihara module.

Table 4. *Study participants and instruments*

Group	n	Role	Instrument
Participants with CVD	5	Referred through PERI; tested the Ishihara and enhancement modules	End-user survey (32 items, analyzed separately)
General users (non-CVD)	30	Usability, accessibility, awareness	End-user survey (32 items; 29 scored, since the three enhancement items apply to CVD users only)
Eye-care professionals	3	Clinical validation of screening, simulation, enhancement, disclaimers	Clinical expert validation (28 items); PERI Ishihara validation form
IT experts	10	Technical quality, performance, privacy, accessibility	IT expert checklist (55 scored items)

Statistical and qualitative analysis

Survey data were summarized as item and dimension means and standard deviations. Cronbach's α was computed for the end-user survey (reliability threshold ≥ 0.70) and is also reported for the IT checklist and the CVD sub-sample; it was not computed for the three clinical experts, and it is unstable in very small samples (Nunnally, 1978). Open-ended responses were analyzed by thematic analysis (Braun & Clarke, 2006). No inferential tests were run: the planned pre- versus post-enhancement Ishihara comparison (paired t-test or Wilcoxon signed-rank test) was not carried out.

Ethical considerations

Participants gave informed consent in English or Filipino. Records were pseudonymized with participant codes; the code-to-name key was stored separately in encrypted form, and data handling followed RA 10173. Screening was framed as preliminary; positive results triggered referral information rather than a diagnosis. Ethical clearance was obtained through the institutional review process of the College of Information Technology and Computer Science, LPU-Cavite.



Results

System implementation

The development process produced a working Android application with four deployed pipelines (simulation, daltonization, hue rotation, and color identification), the Ishihara module, an education module with career-awareness content, an enhancement-intensity slider, before/after comparison by long-press, a gallery mode that applies simulation to stored photographs, and haptic and text-to-speech feedback in the color identifier. Two engineering problems were resolved during development: a camera hardware-abstraction race condition on rapid screen transitions, addressed with a 700–900 ms delayed activation and a listener that deactivates the camera before unmount, and unhandled runtime failures in camera-heavy scenarios, addressed with a three-tier error-handling scheme (rolling file log, global handlers, and error boundaries on every camera screen).

We chose the deployed enhancement method after comparing it with the U-Net. The U-Net reached 95.8% pixel accuracy and 81.1% mean IoU on its validation set. Self-gating daltonization needed 0.4 s per image against 2.8 s for the segmentation-gated approach, removed the 8 MB model dependency, and removed the risk of classifier errors.

Color enhancement

Fidelity to the normal-vision original

Table 5 summarizes fidelity across the six test images. Mean ΔE_{00} was 15.72 for daltonization and 15.31 for hue rotation; differences between the algorithms were at most 1.1 ΔE_{00} within any CVD type and changed direction across types. Hue rotation had the higher mean SSIM for protan and tritan, and daltonization for deutan. The neutral-gray control gave $\Delta E_{00} = 0$ and SSIM = 1.000 for every CVD type and algorithm, confirming that achromatic pixels pass through unchanged. Distances from the original were largest (roughly 19–28 ΔE_{00}) for the red–green gradient, the saturated primaries, and, for tritan, the blue–yellow gradient, where enhancement necessarily makes large chromatic changes. The $\Delta E_{00} < 2.0$ target was not met in any non-control condition. The SSIM > 0.90 target was met in 10 of 36 image $\times$ CVD type $\times$ algorithm combinations; six of these were the neutral-gray control, and four were the ColorChecker chart under deutan and tritan (both algorithms).

Table 5. *Fidelity of enhanced images to the normal-vision original*

CVD type	ΔE_{00} Daltonization	ΔE_{00} Hue rotation	SSIM Daltonization	SSIM Hue rotation
Protan	15.27	15.04	0.763	0.824
Deutan	14.51	14.63	0.787	0.771
Tritan	17.38	16.27	0.742	0.768
Mean	15.72	15.31	0.764	0.788

Discrimination gain

On the task-relevant metric, daltonization outperformed hue rotation for every CVD type in both confusion-pair sets (Tables 6 and 7; Figure 2). In Set A, daltonization's mean gain was +6.91 ΔE_{00} (unweighted mean of the three type means; +7.50 weighted by pairs) against +1.88 (+2.09 weighted) for hue rotation. In Set B, daltonization's gains were +22.06, +12.72, and +3.36 ΔE_{00} for protan, deutan, and tritan against +8.39, +4.65, and +1.19 for hue rotation, and hue rotation's median gain was about zero in every type, meaning that it changed most pairs by less than the perceptual threshold. Daltonization's gain for the tritan condition was small in both sets, and hue rotation's tritan gain stayed below the 2.0 threshold in both.


Table 6. *Discrimination gain, Set A*

CVD type	Pairs	Daltonization mean gain	Daltonization % > 2.0	Hue rotation mean gain	Hue rotation % > 2.0
Protan	60	+13.62	98.3%	+3.29	61.7%
Deutan	60	+4.03	86.7%	+1.88	55.0%
Tritan	36	+3.08	83.3%	+0.46	50.0%
Mean of type means	156	+6.91		+1.88	
Pair-weighted mean	156	+7.50		+2.09	

Table 7. *Discrimination gain, Set B*

CVD type	Pairs	Daltonization mean / median	Daltonization % > 2.0	Hue rotation mean / median	Hue rotation % > 2.0
Protan	20	+22.06 / 17.64	95.0%	+8.39 / 0.00	45.0%
Deutan	20	+12.72 / 8.12	90.0%	+4.65 / 0.00	25.0%
Tritan	14	+3.36 / 3.00	64.3%	+1.19 / 0.10	35.7%

Figure 2. Mean discrimination gain (CIEDE2000 ΔE) by algorithm and CVD type for Set A (left) and Set B (right). The dashed line marks the 2.0 perceptual-significance threshold.



Color identification

The identifier classified all 54 reference samples correctly (24 ColorChecker patches and 30 CSS colors; Table 8), with a perfectly diagonal confusion matrix, including the low-chroma Neutral class and adjacent pairs such as red–orange, blue–cyan, and pink–violet. Mean confidence across the 24 ColorChecker patches was about 93% (range 83–99%); the lowest confidence, 83%, was for the blue-flower patch at the violet–blue boundary, which was still correctly named Violet. The thesis notes that these results reflect controlled reference samples and that camera noise, lighting, and sensor variation may lower accuracy in real use.

Table 8. Color identifier accuracy by family

Family	Correct/total	Confidence range	Family	Correct/total	Confidence range
Blue	6 / 6	91–98%	Orange	3 / 3	90–96%
Brown	4 / 4	89–96%	Pink	5 / 5	88–97%
Cyan	5 / 5	92–98%	Red	4 / 4	91–99%
Green	6 / 6	87–95%	Violet	5 / 5	83–94%
Neutral	12 / 12	85–99%	Yellow	4 / 4	89–97%
Total	54 / 54 (100%)	mean 93%			

CVD simulation fidelity

The float16 GPU shader matched the float64 CPU reference closely: overall mean ΔE_{00} was 0.027 (maximum 1.199, in a noisy natural-scene region under protan simulation), all 18 conditions passed the < 1.0 criterion, and saturated primaries and neutral grays were bit-exact. Against a reference computation that uses the piecewise sRGB transfer function, pure primaries matched exactly ($\Delta E_{00} = 0$), whereas naturalistic colors deviated by up to 5.60 ΔE_{00} under protan simulation (forest green; maximum channel error 22), 1.85 under deutan, and 1.82 under tritan. The thesis attributes these deviations to the $\gamma = 2.2$ approximation, which diverges from the piecewise function at low luminance.

Table 9. CVD simulation fidelity

Test	Conditions	Result
GPU float16 vs CPU float64 (same matrices)	6 images $\times$ 3 CVD types = 18	Mean ΔE_{00} 0.027; max 1.199; 18/18 below 1.0
App output vs piecewise-sRGB reference	12 colors $\times$ 3 CVD types = 36	Max ΔE_{00} : protan 5.60; deutan 1.85; tritan 1.82; pure primaries 0.000

System testing and cross-device stability

All executed system test cases passed—with no critical failures. Across the seven devices used by testers and participants, operation was stable on six; camera lag was observed on the entry-level phone of the participant with CVD, which the thesis attributes to the capture-and-process pipeline (Table 10). No crashes were recorded. Processing latency was not benchmarked systematically across device tiers; the timings available are those given in Sections 2.3.2 and 3.1.

Table 10. Cross-device stability by participant device

ID	Participant type	Device	Chipset	Tier	Result
P1	System tester	Huawei Nova 5T	Kirin 980	Upper mid-range	Stable
P2	System tester	Samsung Galaxy A54 5G	Exynos 1380	Upper mid-range	Stable
P3	System tester	Nubia Neo 3 5G	Unisoc T8300	Mid-range	Stable
P4	System tester	Oppo A5 Pro 5G	Dimensity 700	Entry-level	Stable
P5	System tester	Samsung Galaxy A15 5G	Dimensity 700	Entry-level	Stable

ID	Participant type	Device	Chipset	Tier	Result
P6	General user	Xiaomi Redmi 10C	Snapdragon 680	Entry-level	Stable
P7	Participant with CVD	Samsung Galaxy A22 5G	Dimensity 700	Entry-level	Camera lag observed

Ishihara screening module

PERI's review (signed 23 April 2026) verified the answer keys for both modes against standard Ishihara references and rated them Highly Acceptable; confirmed the scoring thresholds and the Protan/Deutan classification logic; and approved the minimum-brightness control, the exclusion of the demonstration plate from scoring, and the preliminary-screening disclaimer. Two PERI recommendations were implemented before evaluation: a fixed 3-s presentation with automatic advance, to discourage analytic strategies, and explicit "professional screening/diagnosis" wording on the results screen. PERI suggested considering blue-light filter detection; the enforced brightness minimum was retained instead. The outcome was "Approved with Revisions", with all items Acceptable or Highly Acceptable. The three eye-care professionals gave the module's implementation a mean of 4.00 (SD = 0.00; Table 12).

Acceptability and expert evaluation

General users and participants with CVD

The 30 general users (19 male, 11 female; mean age 22.6 years, SD = 5.3, range 18–45) rated ReColor Highly Acceptable overall (grand mean 3.60, SD = 0.49; Cronbach's α = 0.9348 over 29 scored items). Multimodal feedback (3.48) was the only dimension below the Highly Acceptable threshold. The five participants with CVD gave a grand mean of 3.38 (SD = 0.81), Acceptable overall, with the simulation module rated highest (4.00) and the color-enhancement dimension lowest (2.73, SD = 0.88; Table 11). The α for the CVD sub-sample was 0.4875; given $n = 5$, we did not use it as a reliability estimate, and these results are exploratory.

Table 11. *Acceptability by dimension: general users and participants with CVD (4-point scale)*

Dimension (items)	General users (n = 30) M (SD)	CVD users (n = 5) M (SD)
Engagement (7)	3.58 (0.49)	3.29 (0.99)
Functionality (4)	3.58 (0.50)	3.35 (0.75)
Aesthetics (4)	3.67 (0.47)	3.35 (0.67)
Information quality (5)	3.67 (0.47)	3.64 (0.70)
Color enhancement (3)	Not applicable	2.73 (0.88)

Dimension (items)	General users (n = 30) M (SD)	CVD users (n = 5) M (SD)
CVD simulation (1)	3.70 (0.47)	4.00 (0.00)
Color identification (1)	3.60 (0.50)	3.40 (0.55)
Ishihara screening (1)	3.53 (0.51)	3.60 (0.55)
Multimodal feedback (2)	3.48 (0.54)	3.30 (0.95)
Career awareness (2)	3.53 (0.54)	3.70 (0.48)
Overall satisfaction (2)	3.65 (0.48)	3.40 (0.84)
Grand mean	3.60 (0.49) [29 items]	3.38 (0.81) [32 items]

IT experts and clinical experts

The ten IT experts rated the 55-item checklist Acceptable overall (grand mean 3.33, SD = 0.71; $\alpha = 0.9794$). Multimodal feedback and model performance (identifier and simulation) were Highly Acceptable (both 3.57). The lowest dimension was functionality (3.18), and the lowest items were those on camera rendering and processing latency (Q32 = 2.80; Q33 = 2.90). Overall, 46.7% of ratings were Highly Acceptable, 40.4% Acceptable, 12.2% Fairly Acceptable, and 0.7% Unacceptable. The three eye-care professionals rated all four clinical dimensions Highly Acceptable (grand mean 3.92, SD = 0.28); their ratings are descriptive only (n = 3, no α ; Table 12). All eight WCAG 2.2 success criteria selected for evaluation (including non-text content, use of color, contrast, text resizing, label in name, and target size) passed the internal audit.

Table 12. Expert evaluation by dimension (4-point scale)

Evaluator group and dimension (items)	M	SD	Interpretation
IT experts (n = 10)			
Visual experience: navigation, UI, accessibility (18)	3.32	0.67	Acceptable
Functionality (5)	3.18	0.77	Acceptable
Performance and stability (9)	3.27	0.78	Acceptable
Privacy and security (11)	3.29	0.78	Acceptable
Daltonization / color enhancement (4)	3.42	0.71	Acceptable

Evaluator group and dimension (items)	M	SD	Interpretation
Multimodal feedback (3)	3.57	0.57	Highly Acceptable
Model performance: identifier, simulation (3)	3.57	0.50	Highly Acceptable
Cross-device compatibility (2)	3.45	0.60	Acceptable
Grand mean (55)	3.33	0.71	Acceptable
Clinical experts (n = 3)			
Ishihara implementation (8)	4.00	0.00	Highly Acceptable
CVD simulation accuracy (6)	3.83	0.38	Highly Acceptable
Daltonization effectiveness (8)	3.92	0.28	Highly Acceptable
Medical disclaimers and referral (6)	3.89	0.32	Highly Acceptable
Grand mean (28)	3.92	0.28	Highly Acceptable

Qualitative findings

Thematic analysis of open-ended responses from general users and participants with CVD produced six themes (Table 13): improved color distinction, ease of use, awareness and empathy, real-world usefulness, customization and control, and performance limitations. Clinical experts' comments converged on suitability for preliminary screening rather than diagnosis, support for clinical workflow, educational value, and general acceptance, with three cautions: add instructions for users who cannot see the plate content, the risk of the app being used as a diagnostic tool, and a need for larger-scale validation. Findings across the clinical, technical, and user streams were consistent on two points—the tool's educational and screening value and its processing delay.

Table 13. *Themes and illustrative excerpts*

Theme	Illustrative excerpt	Source
Improved color distinction	"The colors are easier to tell apart now."	General user
Ease of use	"Navigation was smooth and intuitive."	General user
Awareness and empathy	"Now I understand what people with CVD see."	General user



Theme	Illustrative excerpt	Source
Real-world usefulness	"I can use this for daily tasks."	General user
Customization and control	"The slider helps me adjust to what feels comfortable."	Participant with CVD
Performance limitations	"Processing takes a bit of time."	General user
Clinical: misuse risk	"Might be abused as diagnostic intervention."	Clinical expert
Clinical: need for guidance	"Add instruction if user cannot see anything..."	Clinical expert



Discussion

Principal findings

ReColor combines four functions that existing apps offer separately—screening, enhancement, identification, and simulation—in one Android application built from deterministic image-processing pipelines. The technical evidence is the study's strongest point: daltonization separated confusable colors more reliably than hue rotation in every CVD type, the GPU shader reproduced the CPU reference with negligible error, and the identifier named all 54 reference colors correctly. The weakest evidence concerns the outcome the app exists to deliver, namely better color discrimination in daily life for people with CVD, which rests on five participants' perception ratings.

Enhancement: which metric, and what it does and does not show

Faithfulness metrics (ΔE_{00} and SSIM against the normal-vision original) showed small and inconsistent differences between algorithms. That is expected: for a person missing a cone class, similarity to normal vision is not an attainable target, so the $\Delta E_{00} < 2.0$ goal was never a realistic criterion, and the $SSIM > 0.90$ goal was met only in a minority of conditions because effective enhancement must change pixel values. Discrimination gain measures the intended effect directly—whether colors that were confused become separable—and it favored daltonization consistently, with hue rotation's median gain near zero. Two cautions apply. First, this metric was not among those listed in the study's original evaluation framework (Table 12 of the thesis), and its magnitude depends on which pairs are chosen: deutan gains for daltonization differed threefold between the two sets. The ordering of the algorithms is therefore better supported than the absolute gains. Second, all gains are computed against a simulated observer.

That observer is a full dichromat, but the target users are anomalous trichromats (protanomaly and deuteranomaly), who retain partial L/M discrimination. Enhancement tuned to a complete deficiency can over-correct for milder cases, and gains measured against a dichromat model probably overstate the benefit for them. The intensity slider lets users scale the correction, and severity-parameterized models such as that of Machado et al. (2009) would allow gain to be evaluated at realistic severities. Most importantly, simulation-based gains do not show that people with CVD perform better with enhanced images. Behavioral evaluations of daltonization, such as visual-search accuracy and response time (Simon-Liedtke & Farup, 2016), are the appropriate next step. The pre- versus post-enhancement Ishihara protocol built into ReColor was not run, and the five participants with CVD rated the enhancement dimension lowest (2.73) even though their comments described easier color distinction. With $n = 5$, we treat those two signals as unresolved rather than contradictory.

Deterministic versus learned enhancement

Deploying self-gating daltonization instead of a segmentation-gated pipeline traded semantic targeting for speed and predictability: 0.4 s versus 2.8 s per image, no 8 MB model, and no classifier error path. Self-gating identifies pixels that a dichromat would confuse, not regions or objects, so it cannot decide that only some colored regions deserve correction. The U-Net's 95.8% pixel accuracy and 81.1% mean IoU were obtained on a synthetic color-patch dataset; they show that the architecture can learn the ten-class mapping, not that it would generalize to natural scenes, and the timing comparison lacks device details. As presented, the model is best regarded as a research asset for later region-specific enhancement.

Color identification

Perfect accuracy on 54 reference colors shows that the CIELAB nearest-neighbor design, with L^* down-weighting and a chroma gate, is coherent on standardized inputs. It should not be read as field accuracy. A convincing test would use a locked database and camera-captured images of physical swatches with independently measured



color, under varied lighting, including the near-gray boundary where the thesis reports the weakest performance. Ten coarse families are also a limited vocabulary for everyday tasks such as choosing clothing.

Simulation fidelity

The GPU shader introduces no meaningful precision loss (mean ΔE_{00} 0.027). The larger difference from the piecewise-sRGB reference, up to 5.6 ΔE_{00} for dark, saturated greens under protan simulation, comes from the simplified $\gamma = 2.2$ transfer function. Using the exact piecewise function in the shader is a low-cost fix, and the simulation is presented as an educational aid, where a small error matters less than it would for enhancement.

Screening and the National Vision Screening Program

PERI's review establishes that the digitized test follows standard Ishihara content, scoring, and classification logic; it does not establish diagnostic accuracy, which needs a sample that includes normal-vision controls so that specificity can be estimated. Display variability is a specific risk for on-screen color tests (Linhares et al., 2016), and the module's mitigations—minimum brightness, timed plates, and a non-diagnostic disclaimer—address it only in part. The experts' warning that the app might be misused as a diagnostic tool supports the strong disclaimers and referral prompts already implemented. The thesis maps ReColor's features to seven objectives of RA 11358 (early detection, standardized procedures, referral, research data, public awareness, accessibility, and functional assistance); six are supported by evaluation evidence, while the anonymized research data layer awaits integration with PERI's systems. The thesis recommends a clinical pilot with at least 15 participants with CVD and 10 without.

Acceptability

General users, who cannot judge enhancement, found the app highly acceptable, and the CVD simulation module received the highest single-item scores from both users and participants with CVD, which supports its educational role. IT experts were more critical, particularly of camera rendering and latency, consistent with the freeze-frame design and with camera lag on an entry-level phone. Multimodal feedback was rated lowest by general users; because they may have little need for non-visual cues, this dimension needs evaluation in a larger group of people with CVD (the five here also rated it Acceptable, 3.30).

Limitations

- Only five participants with CVD, one clinical panel of three, and non-random sampling; results for CVD users are exploratory, and no inferential tests were run.
- Enhancement was judged against simulated observers using synthetic images and pairs, with a primary metric that was not in the original evaluation framework; no behavioral test of task performance was done.
- Simulation used complete-dichromacy matrices for a target population of anomalous trichromats; tritan modes were implemented but not a design target.
- Color-identification accuracy was measured on reference samples, possibly overlapping with those used to refine the database; no held-out camera test was reported.
- Ishihara diagnostic accuracy, specificity, and agreement statistics are not established; smartphone display variability remains a threat to screening validity.



- Processing is capture-and-process rather than live video; latency was not benchmarked across device tiers, and the U-Net timing comparison lacks device details.
- No head-to-head comparison with the commercial apps named in the Introduction was made.

Future work follows directly: a clinical pilot with normal-vision controls; behavioral evaluation of enhancement with people with CVD; severity-aware enhancement and simulation; a held-out, camera-based identifier test; exact sRGB handling in the shader; and on-device latency benchmarking, including native or JSI-based pixel processing to bring processing time under one second.

Conclusions

ReColor shows that screening, enhancement, color identification, and simulation can be delivered together as deterministic pipelines on mid-range Android phones. Daltonization gave the more reliable enhancement across CVD types, GPU simulation was faithful, and reference colors were identified accurately. Clinical and general-user ratings were favorable, while ratings from people with CVD were more moderate and low for enhancement. The claim that ReColor improves everyday color discrimination remains to be tested, as does the diagnostic accuracy of its Ishihara module. Both can be addressed with the clinical and behavioral studies outlined above.

Acknowledgments. The authors thank the Philippine Eye Research Institute for clinical validation of the Ishihara module and for referring participants, and all participants and expert evaluators.



References

- Android Developers. (2023). *Core app quality*. Google. <https://developer.android.com/docs/quality-guidelines/core-app-quality>
- Birch, J. (1997). Efficiency of the Ishihara test for identifying red–green colour deficiency. *Ophthalmic and Physiological Optics*, 17(5), 403–408. <https://doi.org/10.1111/j.1475-1313.1997.tb00072.x>
- Birch, J. (2012). Worldwide prevalence of red-green color deficiency. *Journal of the Optical Society of America A*, 29(3), 313–320. (The thesis list gives 2016; volume 29 corresponds to 2012.)
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. (Added; cited in the thesis but absent from its list.)
- Brettel, H., Viénot, F., & Mollon, J. D. (1997). Computerized simulation of color appearance for dichromats. *Journal of the Optical Society of America A*, 14(10), 2647–2655. (Added.)
- Cruz, E. M., Cerdana, H. G. S., Cabrera, A. M. B., Garcia, C. B., Santos-Morabe, E. T., & Nañagas, M. L. R. (2010). Prevalence of color-vision deficiency among male high-school students. *Philippine Journal of Ophthalmology*, 35(1), 20–24.
- Kelly, K. L., & Judd, D. B. (1955). The ISCC–NBS method of designating colors and a dictionary of color names (NBS Circular 553).
- Khizer, M. A., Ijaz, U., Khan, T. A., Khan, S., Liaqat, T., Jamal, A., Zahid, I., Shah, H. G., & Zahid, M. A. (2022). Smartphone color vision testing as an alternative to the conventional Ishihara booklet. *Cureus*. <https://doi.org/10.7759/cureus.30747> (Added.)
- Linhares, J. M. M., et al. (2016). Consistency of color representation in smartphones. *Journal of the Optical Society of America A*, 33(3), A300–A305.
- Machado, G. M., Oliveira, M. M., & Fernandes, L. A. (2009). A physiologically-based model for simulation of color vision deficiency. *IEEE Transactions on Visualization and Computer Graphics*, 15(6), 1291–1298. <https://doi.org/10.1109/TVCG.2009.113>
- Nunnally, J. C. (1978). *Psychometric theory* (2nd ed.). McGraw-Hill.
- Rodriguez-Carmona, M., O'Neill-Biba, M., & Barbur, J. L. (2012). Assessing the severity of color vision loss with implications for aviation and other occupational environments. *Aviation, Space, and Environmental Medicine*, 83(1).
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In N. Navab et al. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015* (pp. 234–241). Springer.
- Sharma, G., & Bala, R. (2002). *Digital color imaging handbook*. CRC Press.
- Simon-Liedtke, J. T., & Farup, I. (2016). Evaluating color vision deficiency daltonization methods using a behavioral visual-search method. *Journal of Visual Communication and Image Representation*, 35, 236–247. <https://doi.org/10.1016/j.jvcir.2015.12.014>



Viénot, F., Brettel, H., & Mollon, J. D. (1999). Digital video colourmaps for checking the legibility of displays by dichromats. *Color Research & Application*, 24(4), 243–252.

W3C Web Accessibility Initiative. (2023). *Web Content Accessibility Guidelines (WCAG) 2.2*. World Wide Web Consortium. <https://www.w3.org/TR/WCAG22/>