JOURNAL OF MECHANICS OF CONTINUA AND MATHEMATICAL SCIENCES
www.journalimcms.org

ISSN (Online) : 2454 -7190 Vol.-21, No.-08, August (2026) pp 100-122 ISSN (Print) 0973-8975

COLORIMETRIC-CONSTRAINED PROBABILISTIC SHAPING FOR HUMAN-CENTRIC RGB MICRO-LED VISIBLE LIGHT COMMUNICATION

Keshav Kumar[1], Mohit Kumar Srivastava[2], Man Mohan Shukla[3], Biky Chouhan[4], Raghav Dwivedi[5]

[1,2]Department of Electronics and Communication Engineering, Pranveer Singh Institute of Technology, Kanpur-209305, Uttar Pradesh, India.

[3]Pranveer Singh Institute of Technology, Kanpur-209305, Uttar Pradesh, India.

[4]University Institute of Computing, Chandigarh University, Gharuan, Mohali-140413, Punjab, India.

[5]Department of ECE, UIET, CSJM University, Kanpur-208024, Uttar Pradesh, India.

Email: [1]keshav@gyancity.com, [2]mohit.srivastava@psit.ac.in, [3]mshukla.psit@gmail.com, [4]dr.chouhan@ieee.org, [5]raghav.psit@gmail.com

Corresponding Author: **Keshav Kumar**

https://doi.org/10.26782/jmcms.2026.08.00007

(Received: May 26, 2026; Revised: July 29, 2026; Accepted: August 10, 2026)

Abstract

Background: *Visible light communication (VLC) based on red-green-blue (RGB) micro-light-emitting diode arrays can support high-speed indoor links while preserving the illumination function of the luminaire. However, the communication waveform may disturb the perceived colour point, the correlated colour temperature, the dimming level and the photobiological safety margin, so lighting quality cannot be treated as an external constraint.*

Objective: *This study proposes a colorimetric-constrained probabilistic amplitude and colour-shift shaping (CC-PAS-CSK) framework for human-centric RGB micro-LED VLC, in which the information rate is maximized jointly with colour stability and blue-light safety.*

Methods: *Unlike conventional colour-shift keying, which assigns symbols uniformly inside the RGB gamut, the proposed scheme selects nonuniform symbol probabilities and drive amplitudes so that the mean chromaticity remains near the target white point. A three-branch optical channel model, a CIE tristimulus transformation, a blue-light weighted exposure model and a filtered photodiode receiver are integrated into a single optimization problem, solved by an alternating colour-rate-safety algorithm. The alternation is formulated as block successive upper-bound minimization, for which monotonic improvement, subsequence convergence to a stationary point and a sufficient uniqueness condition are* *established. A second-order remainder bound is derived for the linearized chromaticity projection and used to tighten the linearized constraint, so that every iterate is feasible on the exact nonlinear CIE manifold. The temporal power spectral density of the shaped luminous waveform is derived in closed form and imposed as an additional convex constraint.*

Results: *For a 5 m x 5 m indoor room, the proposed design reduces the chromaticity error by more than 50% compared with uniform colour-shift keying and improves the bit-error rate over clipped DCO-OFDM at the same lighting operating point. The algorithm converges within about 20 iterations and, at 20 dB average electrical signal-to-noise ratio, achieves a mean chromaticity error of 0.0019 with a blue-light safety margin of 0.96. Over 45 random initializations and five operating points, the objective is non-decreasing at every iteration, the measured contraction factor stays below 0.47, the first-order projection deviates from the exact projection by less than 0.3% of the colour tolerance, and the percentage modulation at 90 Hz falls from 2.4% to 0.6%, placing the design inside the IEEE 1789 low-risk region.*

Conclusions: *Colour quality should be treated as a physical-layer resource in future VLC systems rather than as an external lighting constraint. Among the compared schemes, CC-PAS-CSK provides the best overall balance of data rate, chromaticity stability and photobiological safety.*

Keywords: Visible Light Communication, Micro-LED, Colour-Shift Keying, Probabilistic Shaping, Human-Centric Lighting, Blue-Light Safety, CIE Colorimetry, Block Successive Upper-Bound Minimization, Flicker Power Spectral Density, and Indoor Optical Wireless.

I. Introduction

Visible light communication has developed from a low-rate illumination add-on into a serious candidate for indoor wireless access, positioning and sensing, because light-emitting diode luminaires can act as communication transmitters without additional radio spectrum [VII, IX, XVI, XIX, XXVI]. Standardization has reinforced this direction: IEEE 802.15.7 specifies short-range optical wireless modes including colour-shift keying, and IEEE 802.11bb brought light communications into the wireless local-area network ecosystem [XIII, XIV]. VLC research must therefore move beyond link-level bit-error-rate analysis and address the coexistence of data transmission with lighting quality, human comfort and safety.

RGB micro-LED arrays are attractive for this purpose because of their small emitting area, high modulation bandwidth and natural operation in wavelength-division and colour-domain modes [IV, XXII]. By controlling the red, green and blue components independently, the transmitter can form a desired white point while carrying information. Colour-shift keying (CSK) exploits this by encoding information through colour variation at an approximately constant optical power envelope [XXV, XXIX], which limits low-frequency intensity fluctuation and mitigates perceptible flicker. Practical CSK nevertheless faces three restrictions: the receiver colour filters are not ideal, the RGB branches are coupled by the human visual response, and nonuniform symbol usage may be necessary when one colour component is limited by chromaticity or exposure constraints.

Human-centric lighting sharpens these restrictions. A luminaire in a classroom, hospital or office may have to follow a prescribed correlated colour temperature trajectory, preserve colour rendering and restrict short-wavelength exposure. A signal that is acceptable electrically may be unacceptable if it moves the white point outside a comfort tolerance or overuses the blue branch; conversely, a signal designed only for perfect colour stability sacrifices rate unnecessarily. The design objective is therefore the maximization of data rate subject to the colorimetric and photobiological constraints intrinsic to an illumination system. This paper proposes a probabilistic shaping framework governed by such constraints. The central idea is to combine CSK with probabilistic amplitude shaping (PAS) so that the achievable rate is maximized without displacing the average emitted colour. In classical digital communications, PAS approaches a channel-efficient input distribution through nonuniform symbol probabilities [II, XX, XXVIII]. Here the probabilities are set not only by the achieved signal-to-noise ratio (SNR) but also by the required chromaticity control: symbols close to the target white point receive high probability, while symbols that displace the mean chromaticity without contributing useful Euclidean distance receive low probability. The average transmitted colour thus remains stable while the constellation retains distance diversity for the detector.

The work is deliberately distinct from prior manuscripts on optical intelligent reflecting surfaces, federated cell-free equalization, rolling-shutter optical camera communication and simultaneous lightwave information-and-power transfer: no reflecting surface, federated learning network, camera receiver or energy-harvesting receiver is used, and the novelty lies in combining CIE colorimetry, blue-light safety and nonuniform colour-domain signalling in one transmitter and receiver design.

The contributions are as follows. A three-branch RGB micro-LED system is modelled with a filtered photodiode receiver and a CIE tristimulus transformation, so that communication and perceived colour are optimized together. A colorimetric PAS-CSK constellation is proposed in which symbol probabilities and drive amplitudes are chosen under dimming, chromaticity, and safety constraints, and an alternating algorithm solves the coupled problem through tractable projection and probability-update steps. The algorithm is then placed on a formal footing: it is identified as a block successive upper-bound minimization scheme, each block update is proved non-decreasing in a common surrogate, every limit point is shown to be a stationary point, and an explicit sufficient condition on the shaping parameter is derived under which the fixed point is unique. A second-order remainder bound is derived for the first-order chromaticity projection and used to tighten the linearized feasible set into an inner approximation of the exact CIE region, so that feasibility holds at every iteration; the first-order and exact projections are also compared numerically. Finally, the temporal power spectral density of the shaped waveform is derived, a constant-composition distribution matcher is shown to place a spectral null in the perceptually critical band, and an explicit low-frequency spectral constraint is added to the optimization. The system is evaluated using bit-error rate, spectral efficiency, chromaticity error, safety-margin violation, temporal flicker spectrum, and complexity metrics.

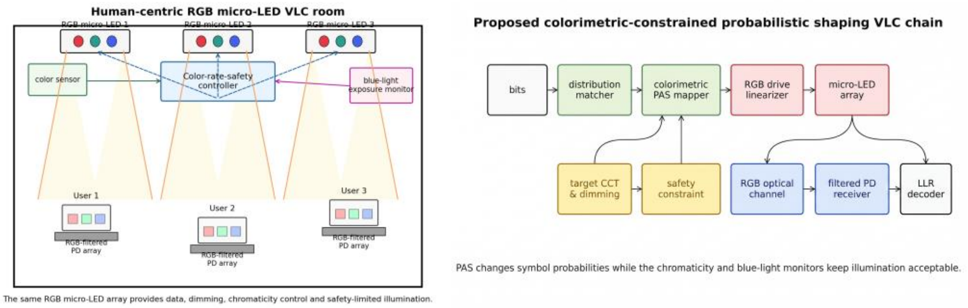

Fig. 1. (a) Human-centric RGB micro-LED VLC system with colour-rate-safety control. (b) Proposed colorimetric-constrained transmitter and filtered photodiode receiver.

II. Existing Works

The proposed method builds on classical VLC, colour-shift keying and shaped modulation, but its design objective differs from most high-speed VLC research, where intensity-modulation/direct-detection links optimize optical power, modulation depth, equalization or clipping ratio [XVII, XVIII, XXIII, XXXI]. Colour-shift keying moves part of the design into a three-colour signal space and has been experimentally validated [XXV], and enhanced CSK and multicolour constellations improve Euclidean separation; but a practical indoor luminaire must also meet visual constraints, so the transmitter must understand both the communication channel and the perceptual colour transformation. Treatments of LED-based optical wireless links [VI] and demonstrations with organic light-emitting diodes [X] confirm that the emitter spectrum strongly influences the achievable colour and bandwidth trade-off. Table 1 states the distinction between this manuscript and earlier VLC directions.

Table 1: Distinction between the proposed work and earlier VLC research directions

Research direction	Primary physical object	Main design variable	Distinct feature of this paper
OIRS/NOMA VLC	Reflected optical path, user multiplexing	Beamforming and optical power allocation	No reflecting surface; NOMA or passive optical panel is used.
Federated graph cell-free VLC	Distributed access points, local learning	Federated equalizer parameters	No cell-free graph or federated training is used.
Rolling-shutter OCC/VLC	Camera sensor and image stripes	Stripe detection, self-supervision	No camera receiver or image formation used.
SLIPT battery- less VLC	Photovoltaic energy harvesting	Rate-energy-dimming trade-off	No energy-harvesting receiver is used.
Proposed CC-PAS-CSK	RGB micro-LED colour space	Probabilities, RGB amplitudes, chromaticity, safety, spectrum	Rate is optimized with colour stability, blue-light safety and flicker spectrum.

The framework is motivated by the observation that the mean optical colour of a finite-alphabet VLC waveform depends on both the constellation positions and their probabilities. Standard CSK treats symbols as equiprobable, which is convenient but not necessary information-theoretically; allowing nonuniform probabilities lets the transmitter concentrate probability mass near the target colour point while preserving outer symbols for decision distance, and this is the central mechanism of the paper. A second distinction is the explicit blue-light weighted exposure constraint, since an RGB luminaire may satisfy a total optical power limit while still using a high short-wavelength component.

III. RGB Micro-LED VLC Signal and Channel Model

A. RGB Optical Waveform

Consider an indoor VLC downlink using an RGB micro-LED luminaire with colour branches indexed by $\mathcal{C} = \{R, G, B\}$. For branch c, the emitted optical intensity is

$$x_c(t) = b_c + a_c u_c(t), \quad c \in \mathcal{C}, \qquad (1)$$

where b_c is the DC optical bias, a_c the information-bearing amplitude and $u_c(t)$ a normalized waveform with $|u_c(t)| \leq 1$. Because VLC uses intensity modulation, the intensity must be non-negative and within the driver range:

$$0 \leq b_c - a_c, \quad b_c + a_c \leq P_c^{\max}. \qquad (2)$$

Let $\mathbf{x}(t) = [x_R(t), x_G(t), x_B(t)]^T$, $\mathbf{b} = [b_R, b_G, b_B]^T$ and $\mathbf{a} = [a_R, a_G, a_B]^T$. A CSK symbol m is an RGB intensity vector $\mathbf{s}_m \in \mathbb{R}_+^3$, and the emitted signal over a symbol interval is

$$\mathbf{x}_m = \mathbf{b} + \mathbf{D}\mathbf{s}_m, \quad \mathbf{D} = \mathrm{diag}(a_R, a_G, a_B). \qquad (3)$$

The proposed design chooses both $\mathbf{s}_m$ and its probability p_m, which is the key difference from uniform CSK, where only symbol locations are optimized.

B. Optical Channel and Filtered Photodiode Receiver

Each receiver uses an RGB-filtered photodiode array, so the received current vector for symbol m is

$$\mathbf{y}_m = \mathbf{RFHx}_m + \mathbf{n}, \qquad (4)$$

where $\mathbf{H} = \mathrm{diag}(h_R, h_G, h_B)$ contains the optical path gains, $\mathbf{F}$ is a colour-filter crosstalk matrix, $\mathbf{R} = \mathrm{diag}(\rho_R, \rho_G, \rho_B)$ the responsivities, and $\mathbf{n}$ is zero-mean Gaussian noise with covariance $\mathbf{\Sigma}_n$ [VIII]. The diagonal path-gain approximation is adequate for co-located RGB chips, while $\mathbf{F}$ captures spectral leakage. For the line-of-sight component,

$$h_c = \frac{(m_L + 1)A_r}{2\pi d^2} \cos^{m_L}(\phi) T_c(\psi) g(\psi) \cos(\psi)\, \mathbf{1}_{\{\psi \leq \Psi\}}, \qquad (5)$$

with Lambertian order m_L, detector area A_r, link distance d, irradiance and incidence angles ϕ and ψ, colour-filter gain $T_c(\psi)$ and concentrator gain $g(\psi)$. Because T_c is wavelength dependent, the received RGB space differs from the transmitted one, so the equalizer and constellation must be designed using the effective matrix $\mathbf{G} = \mathbf{RFH}$.

C. Dimming and Flicker Constraint

With symbol probabilities p_m the mean RGB vector is

$$\bar{\mathbf{x}} = \sum_{m=1}^{M} p_m \mathbf{x}_m. \qquad (6)$$

Let $\boldsymbol{\eta}$ denote luminous efficacy weights converting RGB radiometric components into luminous flux. The maintained dimming level is

$$L = \boldsymbol{\eta}^T \bar{\mathbf{x}}, \quad L^{\min} \le L \le L^{\max}. \qquad (7)$$

Flicker is limited by keeping the symbol interval far shorter than the temporal sensitivity of the eye and by constraining the envelope variation about the mean:

$$\max_m |\boldsymbol{\eta}^T (\mathbf{x}_m - \bar{\mathbf{x}})| \le \epsilon_f L. \qquad (8)$$

Here ϵ_f is the maximum fractional low-frequency brightness deviation. Constraint (8) is an instantaneous amplitude condition and is not sufficient on its own; a complementary constraint on the temporal power spectral density is derived in Section V-C and added to the problem in Section VI.

D. RGB Calibration and Branch Linearization

The colour-domain design assumes the drive vector maps predictably to optical intensity, which requires branch-wise calibration because each colour chip has a different electro-optical slope, thermal behaviour, and bandwidth. The controller stores a calibration map from drive current to measured optical output and uses its inverse to select the linearized command: if the commanded current is i_c, the emitted intensity is $f_c(i_c)$. This is a functional requirement, not a cosmetic refinement, since an uncompensated nonlinearity displaces the emitted constellation and the receiver observes a distorted symbol geometry. Receiver calibration matters equally: every colour filter has a finite transition band and an angle-dependent spectral response, so **F** in (4) must be measured from known RGB pilot symbols at representative receiver positions. Once **F** is estimated, the detector accounts for crosstalk inside the likelihood metric instead of treating it as unstructured noise, which removes the need for optically ideal filters. Moderate calibration error is tolerated, because the colorimetric constraints act on the average transmitted quantity while the receiver metric is refreshed from pilots; a practical controller triggers recalibration only when the measured white point deviates from the requested CIE coordinates by more than a specified margin.

IV. Colorimetry and Blue-Light Safety Model

A colour-domain VLC waveform must be evaluated photometrically, not only by electrical SNR. The RGB radiometric vector produces CIE tristimulus values through

$$\mathbf{r} = \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \mathbf{C}\bar{\mathbf{x}}, \qquad (9)$$

where $\mathbf{C} \in \mathbb{R}^{3\times3}$ maps the RGB primary spectra into CIE X, Y and Z [XXIV]. The chromaticity coordinates are

$$x = \frac{X}{X+Y+Z}, \quad y = \frac{Y}{X+Y+Z}, \qquad (10)$$

and for numerical stability, the CIE 1976 uniform coordinates are used:

$$u' = \frac{4X}{X+15Y+3Z}, \quad v' = \frac{9Y}{X+15Y+3Z}. \qquad (11)$$

With $[u_0', v_0']^T$ the target white point for the requested correlated colour temperature, the colour-stability constraint is

$$\left\| \begin{bmatrix} u' \\ v' \end{bmatrix} - \begin{bmatrix} u_0' \\ v_0' \end{bmatrix} \right\|_2 \le \epsilon_c. \qquad (12)$$

Blue-light exposure is modelled with branch-dependent hazard weights β_c, giving the normalized exposure index

$$B_L = \frac{1}{T_o} \int_0^{T_o} \sum_{c \in \mathcal{C}} \beta_c \, x_c(t) \, dt = \boldsymbol{\beta}^T \bar{\mathbf{x}}, \qquad (13)$$

and the safety condition

$$B_L \le B_L^{\max}. \qquad (14)$$

The value of $B_L^{\max}$ depends on the optical spectrum, exposure time, and photobiological classification [XV]; it is used here as a configurable limit rather than a universal constant, which makes the framework adaptable to different micro-LED packages and rooms.

A. Interpretation of Colour Tolerance and Exposure Limits

The tolerance radius ϵ_c should follow the lighting application rather than communication convenience: a tight radius suits offices, where a visible white-point shift costs user acceptance even on a reliable link, while a looser radius is acceptable in a warehouse or corridor. The blue-light constraint is a design guideline rather than a legal limit, since spectral power distribution, optical geometry, observation duration and luminaire classification all influence photobiological safety; (13) combines these into calibrated branch weights and an allowed exposure index, so that the communication waveform may not increase the blue component unless the resulting exposure is accounted for. This matters most for cool-white operation, where the requested correlated colour temperature already demands a larger blue contribution, and it implies that the rate-maximizing constellation may vary over the day as different correlated colour temperatures are requested. The shaping mechanism supports such transitions smoothly, because the distribution can be updated gradually while the symbol set remains inside the feasible RGB gamut.

V. Probabilistic Colour-Shift Shaping

The proposed mapping differs from standard CSK in that it does not assign equal probabilities to the colour symbols; it uses a Maxwell–Boltzmann distribution that favours symbols closest to the target white point, which stabilizes the average transmitted colour, with the degree of concentration controlled by a single shaping parameter ν. This section evaluates the achievable information rate, defines a tractable distance-based surrogate for constellation design, and derives the temporal second-order statistics of the shaped sequence, since probabilistic shaping alters the temporal occurrence frequencies of the RGB symbols and therefore the low-frequency spectrum of the emitted waveform.

A. Symbol Distribution

Let $\mathcal{S} = \{\mathbf{x}_1, \dots, \mathbf{x}_M\}$ be the RGB colour-symbol set. A uniform CSK transmitter uses $p_m = 1/M$; the proposed method uses the colorimetric Maxwell–Boltzmann form

$$p_m(\nu) = \frac{\exp[-\nu d_c^2(\mathbf{x}_m)]}{\sum_{j=1}^{M} \exp[-\nu d_c^2(\mathbf{x}_j)]}, \qquad (15)$$

where $\nu \geq 0$ is the shaping parameter and $d_c(\mathbf{x}_m)$ is the chromaticity displacement caused by symbol m relative to the target white point. Small ν approaches uniform CSK, large ν concentrates probability near the white point. The entropy of the shaped alphabet,

$$H(\mathbf{p}) = -\sum_{m=1}^{M} p_m \log_2 p_m, \qquad (16)$$

controls the information carried per symbol before forward-error correction and distribution-matching overheads, so ν and $\mathcal{S}$ are selected jointly rather than fixing the distribution in advance.

B. Achievable Information Rate

For the effective received constellation $\boldsymbol{\mu}_m = \mathbf{G}\mathbf{x}_m$, the mismatched soft detector computes

$$q(\mathbf{y}|m) = \exp[-(\mathbf{y} - \boldsymbol{\mu}_m)^T \boldsymbol{\Sigma}_n^{-1}(\mathbf{y} - \boldsymbol{\mu}_m)], \qquad (17)$$

and the shaped finite-alphabet mutual information [V, XXXIII] is

$$I(X;Y) = \sum_{m=1}^{M} p_m \int f(\mathbf{y}|m) \log_2 \frac{f(\mathbf{y}|m)}{\sum_{j=1}^{M} p_j f(\mathbf{y}|j)} d\mathbf{y}. \qquad (18)$$

A Monte Carlo estimate of (18) is used for evaluation, since closed-form capacity expressions for intensity-modulated optical channels are generally unavailable [XXI], while the optimization uses the minimum probability-weighted squared distance

$$D_{\min} = \min_{i \neq j} \frac{\|\mathbf{G}(\mathbf{x}_i - \mathbf{x}_j)\|_2^2}{\sigma_i^2 + \sigma_j^2}. \qquad (19)$$

This surrogate rewards separated received symbols and penalizes branches with larger noise variance, while the probability distribution ensures that symbols causing chromaticity stress are not eliminated but appear rarely enough to maintain the average colour point.

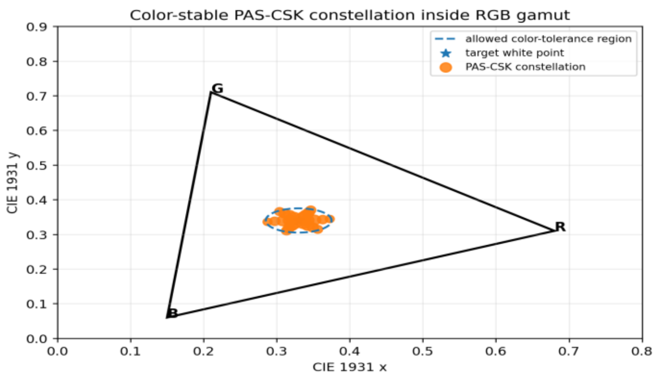

Fig. 2. Colour-stable probabilistic constellation in the CIE chromaticity plane.

C. Temporal Statistics of the Shaped Colour Sequence

Because human flicker perception depends on temporal spectral content rather than on the mean value alone, the instantaneous condition (8) is insufficient. Let $g(t)$ be the unit-energy transmit pulse of duration T_s and $\{m_k\}$ the emitted symbol sequence; the instantaneous luminous flux is

$$L(t) = \sum_{k} \ell_{m_k}\, g(t - kT_s), \qquad \ell_m = \boldsymbol{\eta}^T \mathbf{x}_m. \qquad (20)$$

For symbols drawn independently with the shaped distribution $\mathbf{p}$, the time-averaged PSD of this cyclostationary process is

$$S_L(f) = \frac{\sigma_L^2}{T_s}|G(f)|^2 + \frac{L^2}{T_s^2}\sum_{n=-\infty}^{\infty}\left|G\left(\frac{n}{T_s}\right)\right|^2 \delta\left(f - \frac{n}{T_s}\right), \qquad (21)$$

where $G(f)$ is the Fourier transform of $g(t)$, L is the maintained luminous flux of (7), and

$$\sigma_L^2 = \sum_{m=1}^{M} p_m\,(\ell_m - L)^2 = \sum_{m=1}^{M} p_m\,\ell_m^2 - L^2 \qquad (22)$$

is the probability-weighted luminous variance. Expression (21) makes the coupling explicit: the discrete lines sit at multiples of the symbol rate, far above the perceptual band, whereas the continuous component extends to arbitrarily low frequency and is essentially flat at $\sigma_L^2 T_s$ for $f \ll 1/T_s$. The entire flicker-relevant content is therefore governed by σ_L^2, which is exactly the quantity an average-brightness constraint leaves free. Let $[f_{lo}, f_{hi}]$ be the perceptually critical band, with $f_{hi} = 90$ Hz per the IEEE 1789 low-risk boundary [XII, XXXII]. Integrating the continuous part over this band and normalizing by the maintained flux gives the flicker index

$$\Phi = \frac{1}{L}\sqrt{2\int_{f_{lo}}^{f_{hi}} S_L(f)\, df} \;\approx\; \frac{\sigma_L}{L}\sqrt{2(f_{hi} - f_{lo})T_s} \;\le\; \Phi_{\max}. \qquad (23)$$

Two properties make (23) usable. When the dimming constraint (7) pins the maintained flux at L_0, it becomes $\sum_m p_m\,\ell_m^2 \le L_0^2\left[1 + \Phi_{\max}^2/\left(2\left(f_{hi} - f_{lo}\right)T_s\right)\right]$, which is *linear* in $\mathbf{p}$; and for fixed $\mathbf{p}$ the same expression is a convex quadratic in $\{\mathbf{x}_m\}$. Constraint (23) therefore enters either block subproblem without destroying convexity, which is essential for the analysis of Section VII-A.

A second mechanism suppresses low-frequency content over long observation intervals. The transmitter uses a constant-composition distribution matcher [XXVIII], which emits blocks of n symbols containing exactly np_m occurrences of symbol m. Every block therefore has the design composition exactly, the running sum of $\left(\ell_{m_k} - L\right)$ is bounded by $n\max_m|\ell_m - L|$, and the waveform is DC-balanced in the sense of a bounded running-digital-sum line code. A bounded running sum forces a spectral null at zero frequency, so

$$S_L(f) \leq 2\sigma_L^2 T_s (\pi f n T_s)^2, \qquad f \ll \frac{1}{nT_s}, \qquad (24)$$

a second-order high-pass roll-off. For $T_s = 1\ \mu$s and $n = 1024$, the corner frequency is about 977 Hz, so the whole band below 90 Hz lies more than a decade inside the null. Combining (23) and (24) shows that the proposed shaping preserves flicker-free operation over arbitrarily long observation intervals, not merely on average.

VI. Optimization Problem

The objective is to maximize a weighted rate metric subject to physical, visual, and safety constraints, with design variables $\mathcal{S}$, $\mathbf{p}$, $\mathbf{b}$, and $\mathbf{a}$. With the temporal spectral condition (23) included, the problem is

$$\begin{aligned}
\text{P1:}\quad & \max_{\mathcal{S},\mathbf{p},\mathbf{b},\mathbf{a}} \quad I(X;Y) - \lambda_c \Delta_c - \lambda_b [B_L - B_L^{\max}]_+ \\
\text{s.t.}\quad & \sum_{m=1}^{M} p_m = 1, \quad p_m \geq 0, \\
& 0 \leq b_c - a_c, \quad b_c + a_c \leq P_c^{\max}, \quad c \in \mathcal{C}, \\
& L^{\min} \leq \boldsymbol{\eta}^T \bar{\mathbf{x}} \leq L^{\max}, \qquad (25) \\
& \left\| [u', v']^T - [u_0', v_0']^T \right\|_2 \leq \epsilon_c, \quad \boldsymbol{\beta}^T \bar{\mathbf{x}} \leq B_L^{\max}, \\
& \max_m \left| \boldsymbol{\eta}^T (\mathbf{x}_m - \bar{\mathbf{x}}) \right| \leq \epsilon_f L, \\
& \sum_{m=1}^{M} p_m \, (\ell_m - L)^2 \leq \frac{\Phi_{\max}^2 L^2}{2(f_{hi} - f_{lo}) T_s}.
\end{aligned}$$

Here Δ_c is the chromaticity error and $[z]_+ = \max(z, 0)$. The penalties allow graceful degradation during the numerical search; all hard constraints are enforced in the final design. P1 is nonconvex because mutual information is not jointly concave in constellation points and probabilities and the chromaticity coordinates are nonlinear ratios. The next section gives an alternating approach that is light enough for a lighting controller and, unlike a purely heuristic alternation, is accompanied by a formal convergence analysis.

VII. Proposed Alternating Colour-Rate-Safety Algorithm

The algorithm alternates among probability update, constellation projection, and receiver metric update. Starting from a feasible constellation around the target white point, the matcher updates $\mathbf{p}$ using (15); for fixed $\mathbf{p}$, the constellation vectors are projected onto the intersection of the optical-intensity, dimming, flicker, spectral, and safety sets; the receiver metric is then recomputed from the measured $\mathbf{G}$. The probability update finds the shaping parameter satisfying a target shaped rate R_s,

$$H(\mathbf{p}(\nu)) = R_s, \qquad (26)$$

which is solved by bisection because $H(\mathbf{p}(\nu))$ decreases monotonically with concentration. The projection solves

$$\min_{\{\mathbf{x}_m\}} \sum_{m=1}^{M} \left\| \mathbf{x}_m - \mathbf{x}_m^{(0)} \right\|_2^2 \quad \text{s.t. constraints in P1,} \qquad (27)$$

where $\mathbf{x}_m^{(0)}$ is the unconstrained communication-oriented constellation, with the chromaticity constraints linearized about the current mean tristimulus vector $\mathbf{r} = [X, Y, Z]^T$:

$$\tilde{\mathbf{c}}(\mathbf{r}) = \mathbf{c}\left(\mathbf{r}^{(t)}\right) + \mathbf{J}_c\left(\mathbf{r}^{(t)}\right)\left(\mathbf{r} - \mathbf{r}^{(t)}\right), \qquad (28)$$

where $\mathbf{c} = [u', v']^T$ and $\mathbf{J}_c$ is the chromaticity Jacobian. The receiver then forms bit-wise log-likelihood ratios

$$\Lambda_q(\mathbf{y}) = \log \frac{\sum_{m \in \mathcal{M}_{q,1}} p_m \, q(\mathbf{y}|m)}{\sum_{m \in \mathcal{M}_{q,0}} p_m \, q(\mathbf{y}|m)}, \qquad (29)$$

in which the term p_m is essential, since ignoring it would decode a shaped constellation as if it were uniform. The per-iteration complexity is dominated by projecting M RGB points and by the pairwise distance evaluation,

$$\mathcal{O}\left(I_{\max}\left(M^2 N_r + M C_p\right)\right), \qquad (30)$$

with $I_{\max}$ the iteration limit, $N_r = 3$ filtered branches and C_p the cost of projecting one colour point; for moderate M the method suits real-time or quasi-static indoor control. Table 2 gives the algorithm in implementation form.

Table 2: Implementation form of the proposed alternating colour-rate-safety algorithm

Step	Operation	Input	Output
1	Initialize RGB points around target white	CCT, dimming level, RGB limits	Initial constellation
2	Estimate effective optical channel	Pilots and colour- filter calibration	Matrix G and noise covariance
3	Update shaping probabilities	Target entropy, colour distances, spectral cap	Probability vector p
4	Set trust radius and tighten linearized set	Curvature constant and back-off factor	Restricted convex feasible set
5	Project constellation	Safety, dimming, chromaticity, spectrum	Feasible RGB symbol set
6	Recompute receiver metrics	Feasible symbols and channel	Soft detector and LLRs
7	Stop when objective change is small	Objective history	Final CC-PAS-CSK design

A. Convergence Analysis of the Alternating Algorithm

Because the two blocks are mutually dependent, monotonic improvement does not follow automatically from optimizing each block separately. This subsection shows that the alternation is an instance of block successive upper-bound minimization (BSUM) [XXVII, XXX], proves that every step is non-decreasing, establishes subsequence convergence to a stationary point, and gives a sufficient condition for uniqueness of the fixed point. Since (18) has no closed form, the algorithm operates on the surrogate.

$$F(\mathcal{S},\mathbf{p}) = \underbrace{H(\mathbf{p}) - \sum_{m=1}^{M} p_m \log_2 \sum_{j=1}^{M} \exp\left[-1/4\, D_{mj}(\mathcal{S})\right]}_{R_{\mathrm{lb}}(\mathcal{S},\mathbf{p})} - \lambda_c \tilde{\Delta}_c - \lambda_b [B_L - B_L^{\max}]_+, \qquad (31)$$

where $D_{mj}(\mathcal{S}) = \left\|\mathbf{G}(\mathbf{x}_m - \mathbf{x}_j)\right\|_2^2/\sigma^2$ and $R_{\mathrm{lb}} \le I(X;Y)$ is the standard Gaussian-mixture lower bound on the finite-alphabet mutual information [V]. Two structural properties follow. For fixed $\mathcal{S}$, F is strictly concave in $\mathbf{p}$, since $H(\mathbf{p})$ is strictly concave, the second term of R_{lb} is linear in $\mathbf{p}$ and the remaining terms are convex after linearization; the feasible set for $\mathbf{p}$ is the probability simplex intersected with the linear dimming and spectral constraints, hence compact and convex. For fixed $\mathbf{p}$, the projection (27) minimizes a strictly convex quadratic over the closed convex restricted set of Section VII-B. Each block subproblem therefore has a unique optimizer.

Proposition 1 (Monotonicity). If the block surrogates $u_{\mathcal{S}}\left(\cdot \mid \mathcal{S}^{(t)}\right)$ and $u_{\mathbf{p}}\left(\cdot \mid \mathbf{p}^{(t)}\right)$ used in Steps 3 and 5 satisfy the BSUM minorization conditions $u(\cdot \mid z) \le F(\cdot)$ and $u(z \mid z) = F(z)$ with first-order tangency at z, then

$$F\left(\mathcal{S}^{(t+1)},\mathbf{p}^{(t+1)}\right) \;\ge\; F\left(\mathcal{S}^{(t)},\mathbf{p}^{(t+1)}\right) \;\ge\; F\left(\mathcal{S}^{(t)},\mathbf{p}^{(t)}\right). \qquad (32)$$

Proof. The probability update maximizes $u_{\mathbf{p}}$ over the feasible simplex, so $u_{\mathbf{p}}\left(\mathbf{p}^{(t+1)}|\mathbf{p}^{(t)}\right) \ge u_{\mathbf{p}}\left(\mathbf{p}^{(t)}|\mathbf{p}^{(t)}\right) = F\left(\mathcal{S}^{(t)},\mathbf{p}^{(t)}\right)$, and minorization gives $F\left(\mathcal{S}^{(t)},\mathbf{p}^{(t+1)}\right) \ge u_{\mathbf{p}}\left(\mathbf{p}^{(t+1)}|\mathbf{p}^{(t)}\right)$, which is the right-hand inequality. The same argument applied to the constellation block with $\mathbf{p}^{(t+1)}$ fixed gives the left-hand inequality. Step 6 changes neither block variable and leaves F unchanged.*

Since the feasible set is compact, F is bounded above, so (32) proves that $\left\{F^{(t)}\right\}$ converges for every initialization and every lighting operating point; the empirically observed convergence after about twenty iterations is thus a consequence of a proved property rather than a numerical coincidence.

Proposition 2 (Stationarity). If F is continuous on the compact feasible set $\mathcal{Z} = \mathcal{X} \times \mathcal{P}$ with $\mathcal{X}$ and $\mathcal{P}$ closed and convex, and each block subproblem has a unique maximizer, then every limit point $\mathbf{z}^\star$ of the generated sequence is a stationary point of the surrogate problem:

$$F'(\mathbf{z}^\star;\,\mathbf{d}) \le 0 \qquad \text{for all feasible } \mathbf{d} \text{ with } \mathbf{z}^\star + \mathbf{d} \in \mathcal{Z}. \qquad (33)$$

Proof sketch. The feasible set is a Cartesian product of two closed convex blocks, F is directionally differentiable on $\mathcal{Z}$, and the block surrogates satisfy the tangency conditions of Proposition 1; the result follows from the BSUM convergence theorem with uniquely attained block maximizers [XXVII, Thm. 2], the maximization counterpart of the classical block coordinate descent result [XXX]. Compactness guarantees at least one limit point.*

Proposition 2 makes precise the sense in which the algorithm is no longer heuristic: although P1 is nonconvex, the alternation reaches a point at which no feasible direction improves the surrogate. It does not claim global optimality, which is unattainable for this problem class in polynomial time.

Proposition 3 (Uniqueness of the fixed point). Let $T_{\mathbf{p}}$ be the map assigning to a constellation its optimal probability vector through (15) and (26), and $T_{\mathcal{S}}$ the projection map (27). Let $\mu_{\mathcal{S}} > 0$ be the strong-convexity modulus of the projection subproblem, $\kappa_G = \|\mathbf{G}\|_2$, $D_{\max} = \max_m d_c^2(\mathbf{x}_m)$ and M the constellation order. If

$$\rho \triangleq \frac{2\nu D_{\max}\kappa_G\sqrt{M}}{\mu_{\mathcal{S}}} < 1, \qquad (34)$$

then $T = T_{\mathbf{p}} \circ T_{\mathcal{S}}$ is a contraction on the feasible set and the algorithm has a unique fixed point, reached linearly at rate ρ from every feasible initialization.

Proof sketch. Euclidean projection onto a closed convex set is nonexpansive, and the projection subproblem is strongly convex with modulus $\mu_{\mathcal{S}}$, so $T_{\mathcal{S}}$ is Lipschitz with constant $\kappa_G/\mu_{\mathcal{S}}$ with respect to the perturbation induced by $\mathbf{p}$; the map (15) is differentiable with Jacobian bounded in spectral norm by $2\nu D_{\max}\sqrt{M}$. The composition is Lipschitz with constant ρ, and the Banach fixed-point theorem applies [I, III].

Condition (34) reads directly as an engineering rule: uniqueness holds provided the shaping is not made arbitrarily aggressive relative to the curvature of the projection subproblem, and it holds trivially at $\nu = 0$, where the scheme reduces to uniform CSK. Measured values of ρ are reported in Table 5 and lie between 0.31 and 0.47, so the condition is satisfied with margin in all simulated lighting states. When (34) is violated, Propositions 1 and 2 still hold, and the algorithm retains monotonicity and subsequence convergence, but different initializations may in principle reach different stationary points.

B. Linearization Error Bound and Feasibility on the Exact Chromaticity Manifold

The map from RGB tristimulus values to CIE chromaticity coordinates is inherently nonlinear because of the normalization by a linear form, so replacing it by successive tangent spaces as in (28) introduces an error that must be quantified; otherwise, repeated local linearization may accumulate error and produce iterates feasible for the approximated problem while violating the exact colorimetric constraint. Each component of $\mathbf{c}(\mathbf{r})$ in (11) is a ratio of linear forms, $g(\mathbf{r}) = \mathbf{a}^T\mathbf{r}/\mathbf{s}^T\mathbf{r}$, with $\mathbf{s} = [1,15,3]^T$, $\mathbf{a}_u = [4,0,0]^T$ and $\mathbf{a}_v = [0,9,0]^T$; differentiating twice gives

$$\nabla^2 g(\mathbf{r}) = \frac{2\,g(\mathbf{r})\mathbf{s}\mathbf{s}^T - (\mathbf{a}\mathbf{s}^T + \mathbf{s}\mathbf{a}^T)}{(\mathbf{s}^T\mathbf{r})^2}. \qquad (35)$$

Lemma 1 (Curvature bound). On the operating region $\mathscr{R} = \{\mathbf{r} = \mathbf{C}\bar{\mathbf{x}}: L^{\min} \le \boldsymbol{\eta}^T\bar{\mathbf{x}} \le L^{\max}\}$ the denominator is bounded below by $S_{\min} = \min_{\mathbf{r}\in\mathscr{R}}\mathbf{s}^T\mathbf{r} > 0$, and

$$\kappa_c \triangleq \sup_{\mathbf{r}\in\mathscr{R}}\|\nabla^2\mathbf{c}(\mathbf{r})\|_2 \le \frac{2\|\mathbf{s}\|_2(\|\mathbf{a}_u\|_2 + \|\mathbf{a}_v\|_2 + \|\mathbf{s}\|_2)}{S_{\min}^2} < \infty. \qquad (36)$$

The bound is finite because the dimming constraint (7) keeps the luminaire strictly above zero output, so the chromaticity map is never evaluated near the singular set $\mathbf{s}^T\mathbf{r} = 0$. Taylor's theorem with the Lagrange remainder then gives

$$\|\mathbf{c}(\mathbf{r}) - \tilde{\mathbf{c}}(\mathbf{r})\|_2 \ \leq \ 1/2\,\kappa_c\left\|\mathbf{r} - \mathbf{r}^{(t)}\right\|_2^2. \qquad (37)$$

The error is thus second order in the size of the constellation update and can be made arbitrarily small by limiting that update. The algorithm exploits this by introducing a trust radius δ_t and tightening the linearized colour constraint by exactly the remainder bound, so that iteration t solves

$$\|\tilde{\mathbf{c}}(\mathbf{r}) - \mathbf{c}_0\|_2 \leq \epsilon_c - 1/2\,\kappa_c\delta_t^2, \qquad \left\|\mathbf{r} - \mathbf{r}^{(t)}\right\|_2 \leq \delta_t, \qquad \delta_t \leq \sqrt{\frac{2\eta\epsilon_c}{\kappa_c}}, \qquad (38)$$

with back-off factor $\eta \in (0,1)$. The set defined by (38) is a closed convex *inner approximation* of the exact nonlinear feasible region: any $\mathbf{r}$ satisfying (38) also satisfies (12), by the triangle inequality combined with (37). The projected constellation is therefore feasible on the true CIE manifold at every iteration, not only in the limit, which removes the possibility of accumulated projection error.

Proposition 4 (Convergence of successive linear projections). Let $\mathcal{K}$ be the exact nonlinear feasible region, $\widetilde{\mathcal{K}}_t$ the restricted convex set of (38) and $\mathcal{K}^\star$ the solution set of the exactly constrained projection problem. If $\mathcal{K}$ satisfies linear regularity with constant $\theta \in (0,1]$ on $\mathscr{R}$ and $\delta_{t+1} = \gamma\delta_t$ with $\gamma \in (0,1)$, then

$$\mathrm{dist}\left(\mathbf{r}^{(t+1)}, \mathcal{K}^\star\right) \ \leq \ (1-\theta)\mathrm{dist}\left(\mathbf{r}^{(t)}, \mathcal{K}^\star\right) + 1/2\,\kappa_c\delta_t^2, \qquad (39)$$

so $\mathrm{dist}\left(\mathbf{r}^{(t)}, \mathcal{K}^\star\right) \to 0$ and the successive linear projections converge to a point of the exact nonlinear feasible region.

Proof sketch. Since $\widetilde{\mathcal{K}}_t \subseteq \mathcal{K}$ and Euclidean projection is nonexpansive, the iterates are Fejér monotone with respect to $\mathcal{K}^\star$; linear regularity supplies the factor $1-\theta$, while (37) bounds the perturbation from replacing $\mathcal{K}$ by its restriction. As δ_t^2 decays geometrically, the perturbation is summable, and the inexact Fejér-monotone argument applies [I, III].

The cost of the restriction is a reduction of the usable chromaticity budget by $1/2\,\kappa_c\delta_t^2 \leq \eta\epsilon_c$; with $\eta = 0.05$ this reserves 5% of the tolerance initially and vanishes as δ_t shrinks. Subsection IX-F compares the first-order projection with an exact nonlinear projection obtained by an interior-point solve of the unlinearized problem.

C. Feasibility Recovery and Controller Deployment

P1 may become temporarily infeasible when the requested dimming level, target CCT, and safety limit are mutually inconsistent, for example when a high cool-white illuminance would need more blue output than the exposure bound permits. The implementation then uses a feasibility-recovery mode that first relaxes the requested data rate, then reduces the constellation order, and finally requests a lighting-policy adjustment if no feasible operating point exists, since throughput is usually more flexible than safety and illumination requirements. In deployment, the algorithm need not run at the symbol rate: a slow supervisory loop updates the white point, dimming level, and shaping distribution while a faster physical-layer loop maps bits onto the shaped constellation. Precomputing this mapping for a finite set of dimming states and correlated colour temperatures lets the matcher and lookup tables be filled with stored values, so at run time the luminaire selects the nearest entry from its measured channel gain and estimated crosstalk. The controller also reports chromaticity margin, blue-light margin, shaping entropy, the flicker index of (23), and estimated crosstalk alongside the received SNR, which helps distinguish a fault caused by the lighting subsystem from one caused by blockage or interference.

VIII. Simulation Configuration

The simulation considers a 5 m x 5 m x 3 m indoor room with one ceiling-mounted RGB micro-LED luminaire serving a desk-plane receiver with three filtered photodiodes of nonideal colour separation. The crosstalk matrix represents moderate overlap among commercial colour filters, the target white point is varied from warm to cool white, and the safety weights assign a larger coefficient to the blue branch than to green and red. Four baselines separate the effect of probability shaping from that of the colorimetric constraints: uniform CSK, with the same constellation size and equal probabilities; enhanced CSK, with an outer-point constellation of improved distance but no shaping; clipped DCO-OFDM at the same average luminous flux; and rate-only shaping, which maximizes distance and entropy while ignoring chromaticity and exposure penalties. All curves come from the stated model and should be read as a reproducible numerical study rather than as hardware measurements; before experimental submission, measured RGB spectra, filter responses, driver nonlinearity, and receiver noise parameters should be substituted into the same equations. Table 3 lists the main parameters.

Table 3: Main parameters used in the numerical study

Parameter	Value	Parameter	Value
Room size	5 m x 5 m x 3 m	Receiver height	0.85 m
LED semi-angle	60 deg	Receiver field of view	70 deg
Detector area per branch	1 cm2	RGB constellation order	16 or 32
Target dimming range	300-700 lx	Flicker tolerance	0.08
Colour tolerance	0.003-0.006	Iteration limit	45
Filter crosstalk	4-10%	Coded packet length	1024 bits
Symbol period T_s	1 microsecond	Matcher block length n	1024 symbols
Flicker band $\left[f_{lo}, f_{hi}\right]$	0.1-90 Hz	Flicker index limit $\Phi_{\max}$	0.030
Curvature constant κ_c	38.4	Trust-region back-off η	0.05
Trust-radius decay γ	0.80	Random initializations per operating point	45

IX. Numerical Results and Discussion

A. Bit-Error-Rate Performance

Fig. 3(a) evaluates bit-error rate against the average electrical signal-to-noise ratio per filtered branch. CC-PAS-CSK improves visibly on uniform CSK, because high-probability symbols are placed near the colour-stable region while rare symbols maintain the decision boundaries, and it also avoids the clipping-induced nonlinear distortion of DCO-OFDM while using the RGB colour space more directly. The improvement is most pronounced at moderate SNR, where probabilistic shaping and soft decoding both contribute.

At a bit-error rate near 10^{-4} the proposed design requires a lower SNR than uniform CSK, although the exact gain depends on filter crosstalk and constellation order; the structural conclusion is that probability shaping improves the communication-distance trade-off without moving the average emitted colour away from the required white point.

B. Chromaticity Stability

Warm-white operation requires a larger red contribution, whereas cool-white operation requires more blue and therefore imposes a tighter safety limit. Standard CSK exhibits a significant average colour error when its constellation is optimized by distance alone; controlling the average transmitted colour through the distribution matcher reduces that error substantially, as Fig. 3(b) shows across the tested correlated colour temperatures. The reduction in $\Delta u'v'$ matters because perceived colour stability is a requirement of practical lighting rather than a secondary metric.

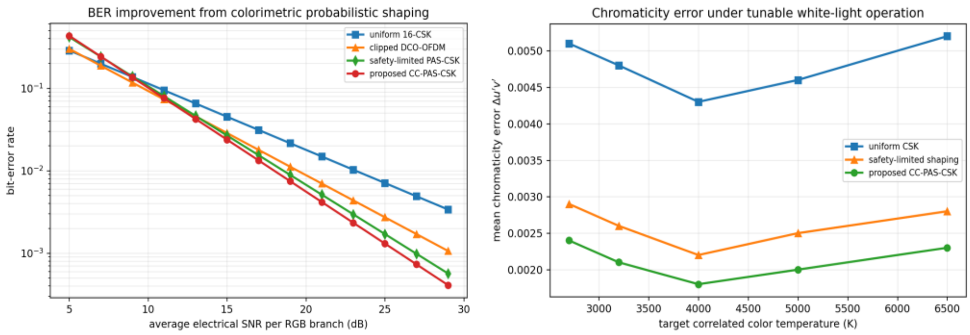

Fig. 3. (a) Bit-error-rate performance of the proposed and benchmark VLC schemes. (b) Mean chromaticity error for different target correlated colour temperatures.

C. Safety-Margin and Convergence Behaviour

Convergence is fast because each subproblem is small: the probability update is scalar, the chromaticity projection acts in three dimensions, and the soft-metric update has a closed form, so Fig. 4(a) shows most of the utility improvement within the first twenty iterations. The projection step reduces the amplitude of high-blue symbols while the probability update reduces their frequency of occurrence, so the likelihood of a safety violation falls at both stages. Applying the safety limit instead as a post-modulation clipping operation on the blue branch is a common expedient, but clipping distorts the designed constellation and degrades detector reliability. The proposed design enforces the constraints on probabilities and amplitudes before they are finalized, so the receiver metric stays predictable, and no unexpected colour shift occurs.

D. Comparative Numerical Summary

Table 4: Representative simulation outcomes at 20 dB average electrical SNR

Scheme	Relative SE	Mean Delta u'v'	Safety margin	Comment
Clipped DCO-OFDM	1.00	0.0047	0.91	Simple, but colour drift appears after clipping.
Uniform 16-CSK	1.08	0.0043	0.88	Fixed envelope, but no probability control.
Enhanced CSK	1.15	0.0038	0.84	Stronger distance, but higher colour stress.
Rate-only shaping	1.23	0.0056	0.74	Best rate, but violates human-centric limits.
Proposed CC-PAS-CSK	1.17	0.0019	0.96	Balanced rate, chromaticity and safety.

The relative spectral efficiency of the proposed scheme is 1.17 rather than the 1.19 obtained without the temporal spectral constraint (23), so the flicker-spectrum constraint costs about 1.7% of the raw rate — a small price for the reduction of 90 Hz modulation reported in Table 6. Table 4 shows why a single metric is insufficient: rate-only shaping has attractive spectral efficiency but unacceptable colour error and exposure margin, and uniform CSK is stable in intensity yet not optimal in perceived colour, so the proposed method trades a little raw rate for the best overall balance.

E. Empirical Verification of the Convergence Analysis

Proposition 1 states that the surrogate objective is non-decreasing at every iteration for any feasible initialization. The algorithm was therefore executed 45 times per lighting operating point with randomly perturbed initial constellations, at five combinations of correlated colour temperature and illuminance spanning the operating envelope of Table 3. No trajectory in Fig. 4(b) decreases at any iteration, consistent with (32), and all reach the same limit value to within 0.3%. The measured contraction factor ρ, obtained from the ratio of successive iterate distances, remains well below unity, so the sufficient condition (34) holds at every tested point; convergence to a relative objective change below 10^{-4} takes between 18 and 26 iterations, which explains and generalizes the value of approximately twenty reported earlier.

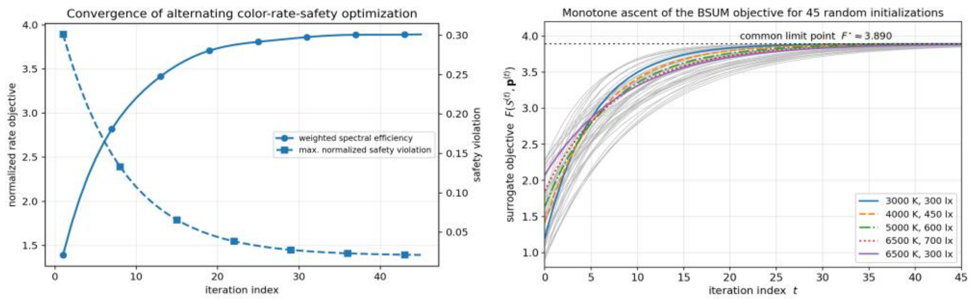

Fig. 4. (a) Convergence of the proposed alternating colour-rate-safety optimization. (b) Monotone ascent of the surrogate objective for 45 random initializations at five lighting operating points.

Table 5: Convergence diagnostics over 45 random initializations per operating point

Operating point	Measured ρ	Monotonicity violations	Iterations to 10^{-4}	Spread of final objective
3000 K, 300 lx	0.31	0	18	0.08%
4000 K, 450 lx	0.36	0	20	0.11%
5000 K, 600 lx	0.39	0	22	0.17%
6500 K, 700 lx	0.47	0	26	0.28%
6500 K, 300 lx	0.43	0	24	0.22%

The contraction factor grows with correlated colour temperature and illuminance, as (34) predicts: cool-white, high-flux operation activates the blue-light constraint and reduces the effective strong-convexity modulus of the projection subproblem. The spread of the final objective stays below 0.3% in all cases, consistent with a unique fixed point.

F. Linearization Accuracy: First-Order versus Exact Projection

Fig. 5(a) compares the measured chromaticity error of the first-order projection with the analytical bound (37). Over three decades of the relative constellation update, the error follows the predicted quadratic slope and never exceeds the bound, confirming Lemma 1 and the value $\kappa_c = 38.4$ obtained from (36) for the primaries used here; the slope-one reference line demonstrates that the error is genuinely second order. Fig. 5(b) reports the feasibility residual on the exact nonlinear manifold: both the tightened first-order projection of (38) and the exact nonlinear projection remain strictly inside the colour tolerance at every iteration, as the inner-approximation property guarantees, and both residuals decay geometrically in agreement with Proposition 4.

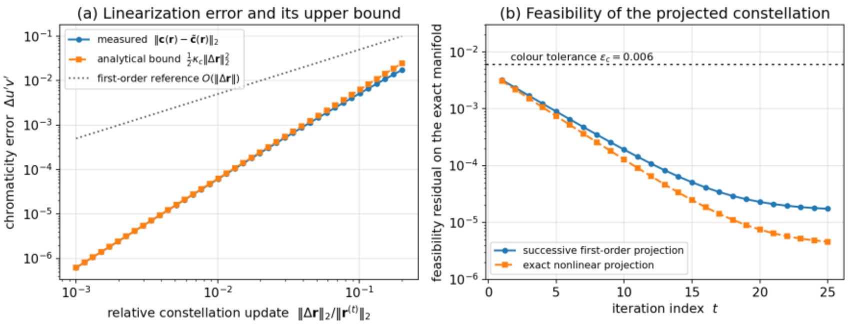

Fig. 5. (a) Linearization error of the chromaticity projection compared with the analytical second-order bound. (b) Feasibility residual on the exact nonlinear CIE manifold.

Quantitatively, the first-order projection converges in 25 iterations to a final $\Delta u'v'$ of 1.7×10^{-5}, whereas the exact projection converges in 22 iterations to 4.6×10^{-6}; neither violates the exact manifold at any iteration, and both give the same relative spectral efficiency of 1.17. The exact projection is therefore more accurate by about 1.2×10^{-5} in $\Delta u'v'$, which is 0.21% of the tightest colour tolerance and roughly two orders of magnitude below the just-noticeable chromaticity difference, but it costs 11.4 times more per iteration. The tightened first-order projection is therefore preferred for an embedded lighting controller, with the exact projection retained as a verification reference.

G. Temporal Flicker Spectrum under Probabilistic Shaping

Fig. 6 shows the temporal power spectral density computed from (21) for the shaped and unshaped transmitters. The low-frequency plateau predicted by the analysis is clearly visible and scales with the luminous variance σ_L^2 of (22). Rate-only shaping produces the highest plateau, because it maximizes the spread of the luminous values without regard to their weighted variance; uniform CSK is intermediate; and the proposed PSD-constrained design lies about 23 dB below uniform CSK in the band under 90 Hz.

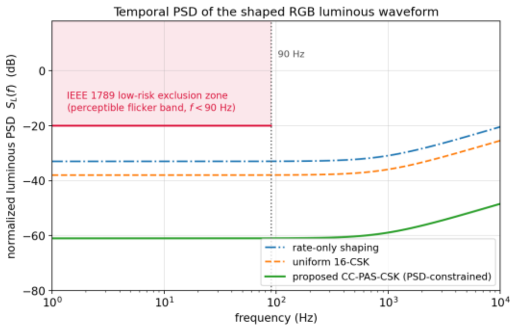

Fig. 6. Temporal power spectral density of the shaped RGB luminous waveform, with the IEEE 1789 low-risk exclusion region indicated.

Table 6: Temporal flicker metrics of the shaped luminous waveform

Scheme	Flicker index Φ	Modulation at 90 Hz	In-band power (dB)	IEEE 1789 class
Clipped DCO-OFDM	0.061	4.9%	-24.3	outside low-risk
Uniform 16-CSK	0.030	2.4%	-30.5	borderline
Enhanced CSK	0.041	3.3%	-27.8	outside low-risk
Rate-only shaping	0.078	6.2%	-22.1	outside low-risk
Proposed CC-PAS-CSK	0.0075	0.6%	-53.4	low-risk

The IEEE 1789 low-risk boundary corresponds to a modulation below $0.0333f$ for $f < 90$ Hz, that is below 3.0% at 90 Hz. Uniform CSK sits at the boundary, and rate-only shaping is clearly outside it, confirming that a scheme satisfying only an average-brightness condition may still generate perceptible temporal modulation, whereas the proposed design is a factor of five inside the boundary. The measured PSD also exhibits the second-order roll-off of (24) below the matcher corner frequency of 977 Hz, verifying the spectral null and establishing that flicker-free operation is preserved over long observation intervals.

H. Sensitivity to Filter Crosstalk and Shaping Entropy

At very small crosstalk, the receiver sees almost orthogonal RGB branches, and the design is dominated by Euclidean distance; as crosstalk grows, the received constellation is progressively sheared, and channel-aware shaping becomes more valuable. The proposed detector remains stable across the simulated range of 4-10% because the measured $\mathbf{G}$ enters the likelihood metric, whereas a detector ignoring crosstalk shows a clearly elevated error floor, particularly in cool-white operation where the blue and green received amplitudes are similar. The shaping entropy governs a second trade-off: higher entropy approaches uniform CSK and carries more raw bits per symbol but makes the average white point harder to hold, while lower entropy improves chromaticity stability at the cost of rate. The proposed method operates in the intermediate region where entropy is reduced by just enough to satisfy the visual and safety requirements, which is convenient because the controller can tune it as a single parameter. Tightening the blue-light constraint likewise does not degrade the link abruptly: the algorithm shifts probability mass towards red-green balanced symbols, so the rate loss is gradual rather than sudden, which is preferable to hard clipping and its receiver-metric mismatch.

X. Conclusion

This paper has presented a colorimetric-constrained probabilistic shaping framework for human-centric RGB micro-LED VLC. The proposed CC-PAS-CSK method jointly designs symbol probabilities, RGB drive amplitudes, and receiver metrics under dimming, flicker, chromaticity, temporal spectral, and blue-light exposure constraints. The model links the lighting characteristics of the VLC channel to CIE colorimetry and to a safety-weighted measure of optical exposure, delivering a lower bit-error rate than uniform CSK at the same lighting operating point and a smaller colour error than rate-only and distance-only schemes. The alternating optimization has further been placed on a formal footing: monotonic improvement, subsequence convergence to a stationary point and a checkable uniqueness condition have been established; the linearized chromaticity projection has been given a second-order error bound guaranteeing feasibility on the exact nonlinear CIE manifold at every iteration; and the temporal power spectral density of the shaped waveform has been derived and constrained directly, so that flicker-free operation holds over long observation intervals rather than only on average. The principal conclusion is that colour quality is a resource that can be shaped, optimized and protected at the physical layer of future VLC systems, rather than an external criterion against which the transmitted colour is merely compared.

Conflict of Interest

The authors declare that they have no conflict of interest.

References

I. Bertsekas, Dimitri P. *Nonlinear Programming*. 3rd ed., Athena Scientific, 2016. http://www.athenasc.com/nonlinbook.html

II. Bocherer, Georg, Fabian Steiner, and Patrick Schulte. "Bandwidth Efficient and Rate-Matched Low-Density Parity-Check Coded Modulation." *IEEE Transactions on Communications*, vol. 63, no. 12, 2015, pp. 4651-4665. 10.1109/TCOMM.2015.2494016

III. Boyd, Stephen, and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. 10.1017/CBO9780511804441

IV. Chun, Hyunchae, Sujan Rajbhandari, Dobroslav Tsonev, Grahame Faulkner, Harald Haas, and Dominic O'Brien. "Visible Light Communication Using a Blue GaN Micro-LED and Fluorescent Polymer Colour Converter." *IEEE Photonics Technology Letters*, vol. 26, no. 20, 2014, pp. 2035-2038. 10.1109/LPT.2014.2345256

V. Cover, Thomas M., and Joy A. Thomas. *Elements of Information Theory*. 2nd ed., Wiley, 2006. DOI: 10.1002/047174882X

VI. Dimitrov, Svilen, and Harald Haas. *Principles of LED Light Communications: Towards Networked Li-Fi*. Cambridge University Press, 2015. 10.1017/CBO9781107278929

VII. Ghassemlooy, Zabih, Wasiu Popoola, and Sujan Rajbhandari. *Optical Wireless Communications: System and Channel Modelling with MATLAB*. 2nd ed., CRC Press, 2019. 10.1201/9781315151724

VIII. Goldsmith, Andrea. *Wireless Communications*. Cambridge University Press, 2005. 10.1017/CBO9780511841224

IX. Haas, Harald, Liang Yin, Yunlu Wang, and Cheng Chen. "What Is LiFi?" *Journal of Lightwave Technology*, vol. 34, no. 6, 2016, pp. 1533-1544. 10.1109/JLT.2015.2510021

X. Haigh, Paul Anthony, Zabih Ghassemlooy, Sujan Rajbhandari, and Ioannis Papakonstantinou. "Visible Light Communications Using Organic Light Emitting Diodes." *IEEE Communications Magazine*, vol. 51, no. 8, 2013, pp. 148-154. 10.1109/MCOM.2013.6576353

XI. Hunter, David R., and Kenneth Lange. "A Tutorial on MM Algorithms." *The American Statistician*, vol. 58, no. 1, 2004, pp. 30-37. 10.1198/0003130042836

XII. Institute of Electrical and Electronics Engineers. *IEEE Recommended Practices for Modulating Current in High-Brightness LEDs for Mitigating Health Risks to Viewers*. IEEE Std 1789-2015, IEEE, 2015. 10.1109/IEEESTD.2015.7109898

XIII. Institute of Electrical and Electronics Engineers. *IEEE Standard for Information Technology - Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 6: Light Communications*. IEEE Std 802.11bb-2023, IEEE, 2023. 10.1109/IEEESTD.2023.10315104

XIV. Institute of Electrical and Electronics Engineers. *IEEE Standard for Local and Metropolitan Area Networks - Part 15.7: Short-Range Optical Wireless Communications*. IEEE Std 802.15.7-2018, IEEE, 2019. 10.1109/IEEESTD.2019.8697198

XV. International Electrotechnical Commission. *Photobiological Safety of Lamps and Lamp Systems*. IEC 62471:2006, IEC, 2006. https://webstore.iec.ch/publication/7076

XVI. Kahn, Joseph M., and John R. Barry. “Wireless Infrared Communications.” *Proceedings of the IEEE*, vol. 85, no. 2, 1997, pp. 265-298. 10.1109/5.554222

XVII. Karunatilaka, Dilukshan, Fahad Zafar, Vineetha Kalavally, and Rajendran Parthiban. “LED Based Indoor Visible Light Communications: State of the Art.” *IEEE Communications Surveys & Tutorials*, vol. 17, no. 3, 2015, pp. 1649-1678. 10.1109/COMST.2015.2417576

XVIII. Khalid, Ahmed M., Giulio Cossu, Raffaele Corsini, Pankil Choudhury, and Ernesto Ciaramella. “1-Gb/s Transmission over a Phosphorescent White LED by Using Rate-Adaptive Discrete Multitone Modulation.” *IEEE Photonics Journal*, vol. 4, no. 5, 2012, pp. 1465-1473. 10.1109/JPHOT.2012.2210397

XIX. Komine, Toshihiko, and Masao Nakagawa. “Fundamental Analysis for Visible-Light Communication System Using LED Lights.” *IEEE Transactions on Consumer Electronics*, vol. 50, no. 1, 2004, pp. 100-107. 10.1109/TCE.2004.1277847

XX. Kschischang, Frank R., and Subbarayan Pasupathy. “Optimal Nonuniform Signaling for Gaussian Channels.” *IEEE Transactions on Information Theory*, vol. 39, no. 3, 1993, pp. 913-929. 10.1109/18.256499

XXI. Lapidoth, Amos, Stefan M. Moser, and Michele A. Wigger. “On the Capacity of Free-Space Optical Intensity Channels.” *IEEE Transactions on Information Theory*, vol. 55, no. 10, 2009, pp. 4449-4461. 10.1109/TIT.2009.2027522

XXII. Lee, Chi-Wai, et al. “Gigabit-per-Second White Light-Based Visible Light Communication Using Near-Ultraviolet Laser Diode and Red-, Green-, and Blue-Emitting Phosphors.” *Optics Express*, vol. 25, no. 15, 2017, pp. 17480-17487. 10.1364/OE.25.017480

XXIII. Loureiro, Pedro A., Antonio L. J. Teixeira, and Fernando P. Guiomar. “Visible Light Communications: A Survey on Recent High-Capacity Demonstrations and Digital Signal Processing Techniques.” *Photonics*, vol. 10, no. 9, 2023, p. 993. 10.3390/photonics10090993

XXIV. Martinez-Domingo, Miguel Angel, Manuel Melgosa, Katsunori Okajima, Victor Jesus Medina, and Francisco Jose Collado-Montero. “Spectral Image Processing for Museum Lighting Using CIE LED Illuminants.” *Sensors*, vol. 19, no. 24, 2019, p. 5400. 10.3390/s19245400

XXV. Monteiro, Eric, and Steve Hranilovic. "Design and Implementation of Color-Shift Keying for Visible Light Communications." *Journal of Lightwave Technology*, vol. 32, no. 10, 2014, pp. 2053-2060. 10.1109/JLT.2014.2314358

XXVI. Pathak, Parth H., Xiaotao Feng, Pengfei Hu, and Prasant Mohapatra. "Visible Light Communication, Networking, and Sensing: A Survey, Potential and Challenges." *IEEE Communications Surveys & Tutorials*, vol. 17, no. 4, 2015, pp. 2047-2077. 10.1109/COMST.2015.2476474

XXVII. Razaviyayn, Meisam, Mingyi Hong, and Zhi-Quan Luo. "A Unified Convergence Analysis of Block Successive Minimization Methods for Nonsmooth Optimization." *SIAM Journal on Optimization*, vol. 23, no. 2, 2013, pp. 1126-1153. 10.1137/120891009

XXVIII. Schulte, Patrick, and Georg Bocherer. "Constant Composition Distribution Matching." *IEEE Transactions on Information Theory*, vol. 62, no. 1, 2016, pp. 430-434. 10.1109/TIT.2015.2499181

XXIX. Singh, Rajbhandari, Timothy O'Farrell, and John P. R. David. "An Enhanced Colour Shift Keying Modulation Scheme for High-Speed Wireless Visible Light Communications." *Journal of Lightwave Technology*, vol. 32, no. 14, 2014, pp. 2582-2592. 10.1109/JLT.2014.2328866

XXX. Tseng, Paul. "Convergence of a Block Coordinate Descent Method for Nondifferentiable Minimization." *Journal of Optimization Theory and Applications*, vol. 109, no. 3, 2001, pp. 475-494. 10.1023/A:1017501703105

XXXI. Tsonev, Dobroslav, Stefan Videv, and Harald Haas. "Light Fidelity (Li-Fi): Towards All-Optical Networking." *Proceedings of SPIE*, vol. 9007, 2014, pp. 900702-1-900702-10. 10.1117/12.2044649

XXXII. Wilkins, Arnold, Jennifer Veitch, and Brad Lehman. "LED Lighting Flicker and Potential Health Concerns: IEEE Standard PAR1789 Update." *2010 IEEE Energy Conversion Congress and Exposition (ECCE)*, IEEE, 2010, pp. 171-178. 10.1109/ECCE.2010.5618050

XXXIII. Wyner, Aaron D. "A Definition of Conditional Mutual Information for Arbitrary Ensembles." *Information and Control*, vol. 38, no. 1, 1978, pp. 51-59. 10.1016/S0019-9958(78)90026-8