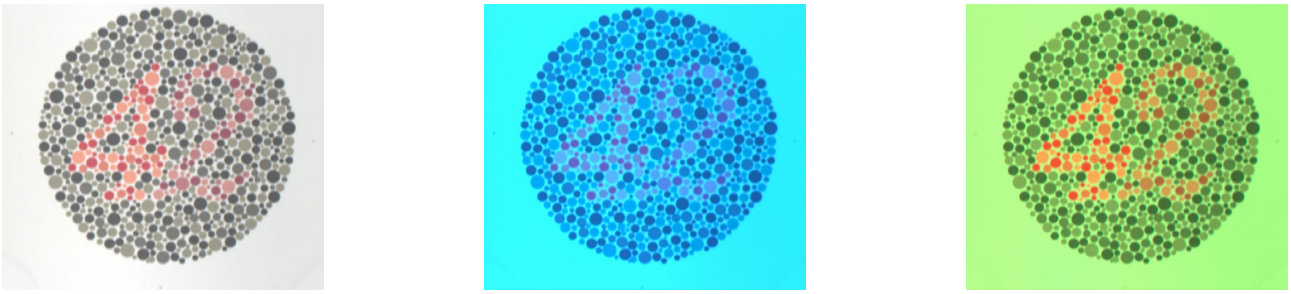
Color Science Laboratory · University of Eastern Finland · COSI 2025

Spectrally Optimal Light Sources

Designing LED Illumination to Hide or Reveal an Ishihara Numeral

Minh Khang Tran

Lab work September–October 2025 · report revised October 2026



(a) D65, $\Delta E^*_{ab} = 31.8$ **(b)** ΔE^*_{ab}-min light, 7.5 **(c)** ΔE^*_{ab}-max light, 57.7

Figure 1. Plate "42" of the Ishihara book, rendered from its measured reflectance under three illuminants. The two LED spectra were found by particle swarm optimisation to shrink or enlarge the colour difference between the numeral and the background, and were then reproduced on a tunable LED source.

Summary. Two dots on an Ishihara plate, one from the numeral and one from the background, were measured with a Specim V10E hyperspectral camera. Ten LED channels of a Ledmotive Spectra Tune Lab were measured with a Hamamatsu PMA-12 spectrometer. Particle swarm optimisation (PSO) searched the mix of the ten channels for six objectives: minimise or maximise CIELAB ΔE^*_{ab}, sRGB ΔRGB, and Michelson luminance contrast between the two dots. Against D65 ($\Delta E^*_{ab} = 31.8$), the optimised lights reach $\Delta E^*_{ab} = 7.5$ and 57.7. Each light needs at most three of the ten LEDs, and the ΔRGB and contrast objectives pick the same mix. All six lights were set on the Ledmotive, and their emitted spectra match the predictions with Pearson $r = 0.997$–0.999 and 3.5 % to 6.6 % relative error. The match is in the emitted spectrum; the perceived colour difference under the real lights was not measured.

1 Question

Colour appearance depends on the light as much as on the surface. Two surfaces whose reflectances differ only in some wavelengths look alike under a light with little power there, and look very different under a light concentrated there. An Ishihara plate is built on that idea: numeral and background dots are close in lightness and differ mainly in the long wavelengths, which a red–green colour-deficient observer weighs differently.

With a light source whose spectrum can be set channel by channel, this can be turned around: *which mix of the available LEDs makes the numeral as hard, or as easy, as possible to see?* The lab answered this for one plate and three measures of difference, and then checked that the hardware produces the spectra the optimiser asked for.

2 Setup and data

Plate. Plate "42" of the *Ishihara Color Deficiency Test* book was scanned with a Specim V10E push-broom camera (2144 pixels per line, 586 lines, 270 bands at about 2.5 nm from 323 to 1 008 nm). Reflectance was computed per band as $R = (S - D)/(W - D)$ from the sample, white-reference and dark-reference scans, and the 120 bands between 400 and 700 nm, the range the LED source covers, were kept. One pixel was taken from a pink dot of the numeral (R_1) and one from a grey dot of the background (R_2). Figure 3 shows the two spectra: they track each other up to about 570 nm, after which R_1 rises to nearly 1 while R_2 stays near 0.4.

Light source. The Ledmotive Spectra Tune Lab has ten channels. Each was driven alone and measured with the PMA-12 spectrometer (85 ms exposure, mean of 10 spectra), then resampled to the 120 reflectance bands (Figure 4). The peaks lie at 430, 447, 466, 474, 504, 524, 546, 596, 637 and 660 nm; channel 7 is a broad phosphor-type emitter, the others are narrow.

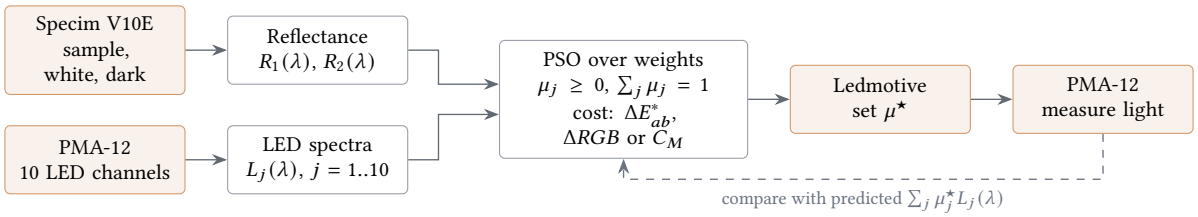


Figure 2. Pipeline. Orange boxes are lab instruments; white boxes are computation. Only the last comparison runs on the real light source.

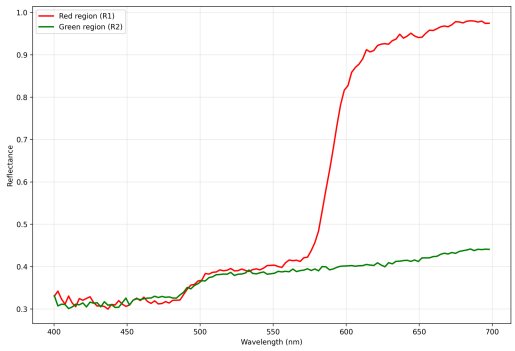


Figure 3. Reflectance of the numeral dot R_1 (red) and the background dot R_2, plotted by the 2025 run (the legend there calls R_2 "green"; it is the grey background). The raw cube has since been lost, so this plot is the only remaining record of the two spectra.

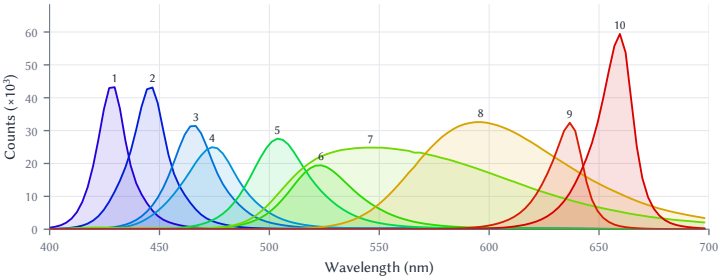


Figure 4. The ten LED channels as measured by the PMA-12, in raw counts on the reflectance grid. Numbers mark the channel index.

3 Method

Colour of a dot under a light. For a light $L(\lambda)$ and a reflectance $R(\lambda)$ the CIE 1931 tristimulus values are

$$X = k \sum_{\lambda} R(\lambda)\, L(\lambda)\, \bar{x}(\lambda), \qquad k = \frac{100}{\sum_{\lambda} L(\lambda)\, \bar{y}(\lambda)}, \tag{1}$$

and likewise for Y and Z, so a perfect white has $Y = 100$ under any light. Three costs follow:

- ΔE^*_{ab}: Euclidean distance in CIELAB, with the D65 white point for every light (Section 5).
- ΔRGB: Euclidean distance between gamma-encoded sRGB values in $[0, 1]$.
- Michelson contrast $C_M = |g_1 - g_2|/(g_1 + g_2)$, with grey level $g = 0.2125R + 0.7154G + 0.0721B$.

A light that pushes either dot outside the sRGB gamut is treated as infeasible, which keeps the optimiser to colours a display could show.

Search space. A candidate light is a convex mix of the ten channels,

$$L_\mu(\lambda) = \sum_{j=1}^{10} \mu_j L_j(\lambda), \qquad \mu_j \geq 0, \quad \sum_j \mu_j = 1. \tag{2}$$

Only the relative mix matters, because k removes the overall level.

Optimiser. PSO moves a swarm of candidate weight vectors, each pulled towards its own best position and the swarm's best. After every step the weights are clipped at zero and rescaled to sum to one. The run used 200 particles, 200 iterations, inertia $c_0 = 0.95$ and attraction weights $c_1 = c_2 = 5$, once per objective (six runs).

4 Results

4.1 The optimised lights

Every objective moved well away from D65 (Table 1 and Figure 5). The largest relative gains are in the maximisation direction. The minimisations are bounded by the gamut constraint and by the dots' shared reflectance below 570 nm: no mix of these LEDs makes the two dots identical, and luminance contrast falls by only a fifth.

Table 1. Cost under D65 and under the optimised light. Change is relative to D65.

	D65	Minimised	Change (%)	Maximised	Change (%)
ΔE^*_{ab}	31.8197	7.5120	−76.4	57.7057	+81.4
ΔRGB	0.2991	0.1988	−33.5	0.7450	+149.1
Contrast C_M	0.0494	0.0398	−19.4	0.1124	+127.6

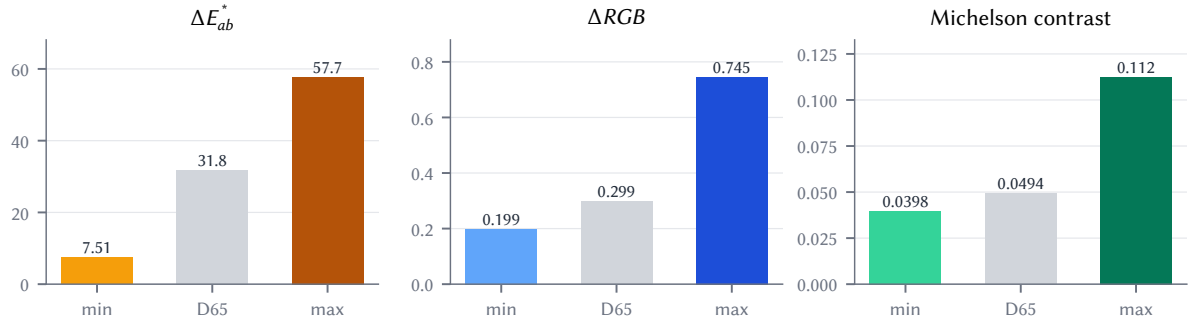


Figure 5. The same numbers as Table 1: each cost under the minimising light, D65, and the maximising light.

Figure 6 shows how the optimiser got there.

1. **Sparse solutions.** Each light uses at most three of the ten channels. Maximising lights use the 504 nm cyan-green and 660 nm deep-red channels; minimising lights use 430 nm violet plus the 524 nm and 546 nm greens.
2. **The mechanism is the long-wavelength gap.** To reveal the numeral the light puts power above 600 nm, where R_1 and R_2 differ most. To hide it the light removes that power and lights the plate where the two reflectances coincide.
3. ΔRGB **and contrast agree.** Both maximising runs returned the same weights (0.616 / 0.384), and the minimising runs differ by less than 0.001 per channel. ΔE^*_{ab} prefers a more even cyan/red split (0.49 / 0.51) when maximising.

Figure 7 renders the whole plate under each light. The renderings are not white-balanced, so the background takes the colour of the light; what matters is the numeral against its surround.

4.2 Reproduction on the LED source

Each set of weights was entered into the Ledmotive and the emitted light measured with the PMA-12. The measured spectrum was resampled to the 120-band grid and scaled by least squares onto the prediction, since the two measurements were taken at different levels. Figure 8 and Table 2 compare them.

Table 2. Agreement between predicted and measured emission. Relative RMSE is $\|p - s\,m\|/\|p\|$ with s the least-squares scale. Peak wavelengths agree exactly in all six cases.

Light	Pearson r	Rel. RMSE (%)	Peak (nm)
ΔE^*_{ab} min	0.999	3.8	429.6
ΔE^*_{ab} max	0.997	6.6	659.5
ΔRGB min	0.999	4.1	429.6
ΔRGB max	0.999	3.5	659.5
Contrast min	0.999	4.1	429.6
Contrast max	0.999	4.1	659.5

The shapes match closely. The residual is systematic: the measured secondary peak (520 nm for the minimising lights, 660 nm for ΔE^*_{ab} max) comes out slightly stronger than predicted. A likely cause is that channel output is not exactly linear in the drive setting, so mixes do not add up exactly as single-channel measurements suggest.

5 Limitations

- **The validation covers the light, not the appearance.** It shows that the Ledmotive emits the requested spectra. The colour difference on the real plate under those lights was not measured (for example with a second hyperspectral scan or a calibrated camera), and the photos in Figure 9 are white-balanced by the phone.
- **No chromatic adaptation.** CIELAB uses the D65 white for every light. That describes an observer adapted to D65

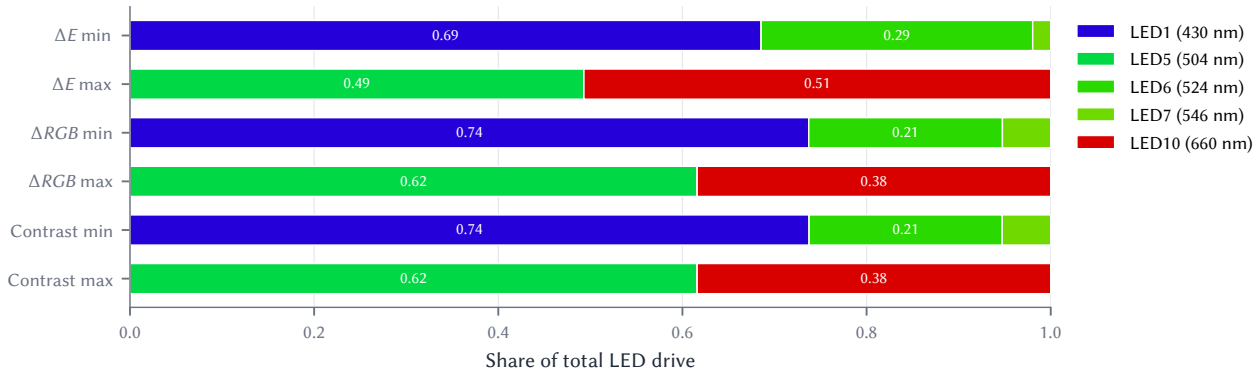


Figure 6. LED mix chosen for each objective. Unused channels are omitted.

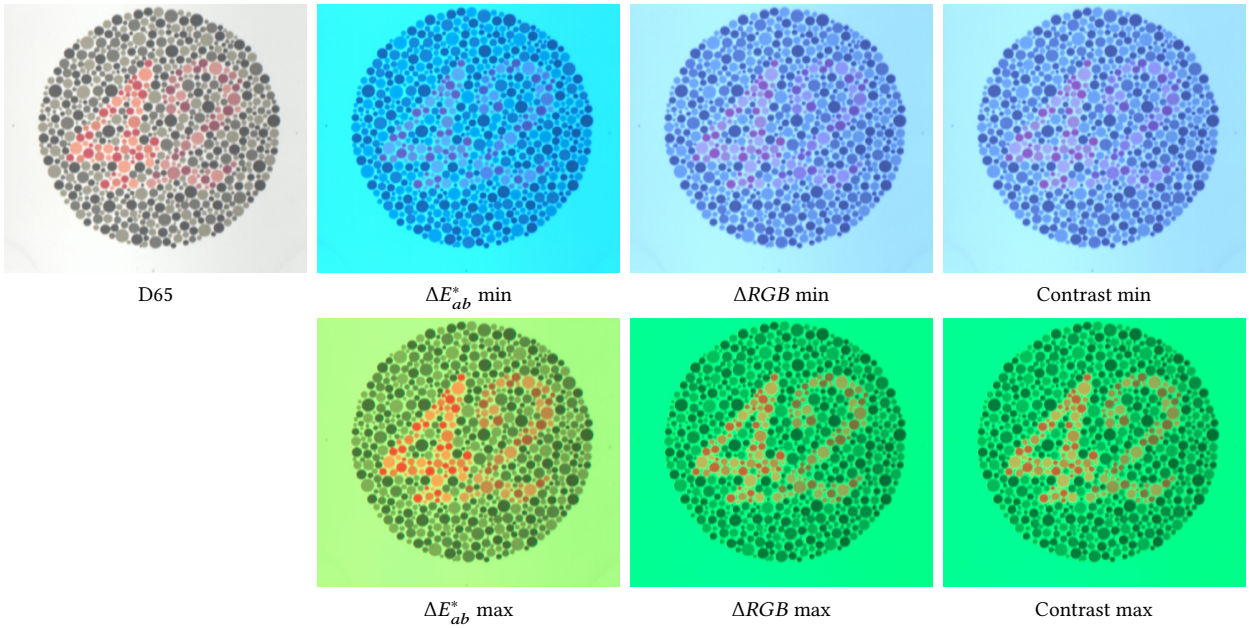


Figure 7. The plate rendered from its reflectance cube under each light (sRGB, no white balance). Under the minimising lights the "42" nearly disappears into the background; under the maximising lights it stands out in orange.

who suddenly sees the test light; an observer adapted to the test light would see a different, probably smaller, ΔE^*_{ab}. Using the test light's own white, or a von Kries step, would be the more faithful choice.

- **One pixel per region.** R_1 and R_2 are single pixels, so sensor noise is part of the input (visible below 450 nm in Figure 3). Averaging a patch inside each dot would be more robust. The reference scans were also averaged over all pixels rather than per detector column.
- **One run per objective.** PSO is stochastic and was run once per objective with no fixed seed. Repeated runs would show whether the reported optima are global.
- **No observer test.** The plate targets colour-deficient vision, but no observer, normal or colour-deficient, was tested under the lights.

6 Conclusion

With ten measured LED channels and two measured reflectances, PSO found illuminants that cut the CIELAB difference between an Ishihara numeral and its background from 31.8 to 7.5 or raised it to 57.7. All six lights were set on the Ledmotive Spectra Tune Lab, and the emitted spectra matched the predictions to $r \geq 0.997$. The solutions are simple (two or three LEDs) and easy to interpret: power above 600 nm reveals the numeral, and removing it hides it. The next steps would be to measure the plate under the real lights and to compute colour differences with chromatic adaptation.

Code and data. Code, LED spectra, optimised weights and the measured light spectra are available at `https://github.com/tmkhang1999/Spectrally-Optimal-Light-Sources`. The comparison in Figure 8 and Table 2

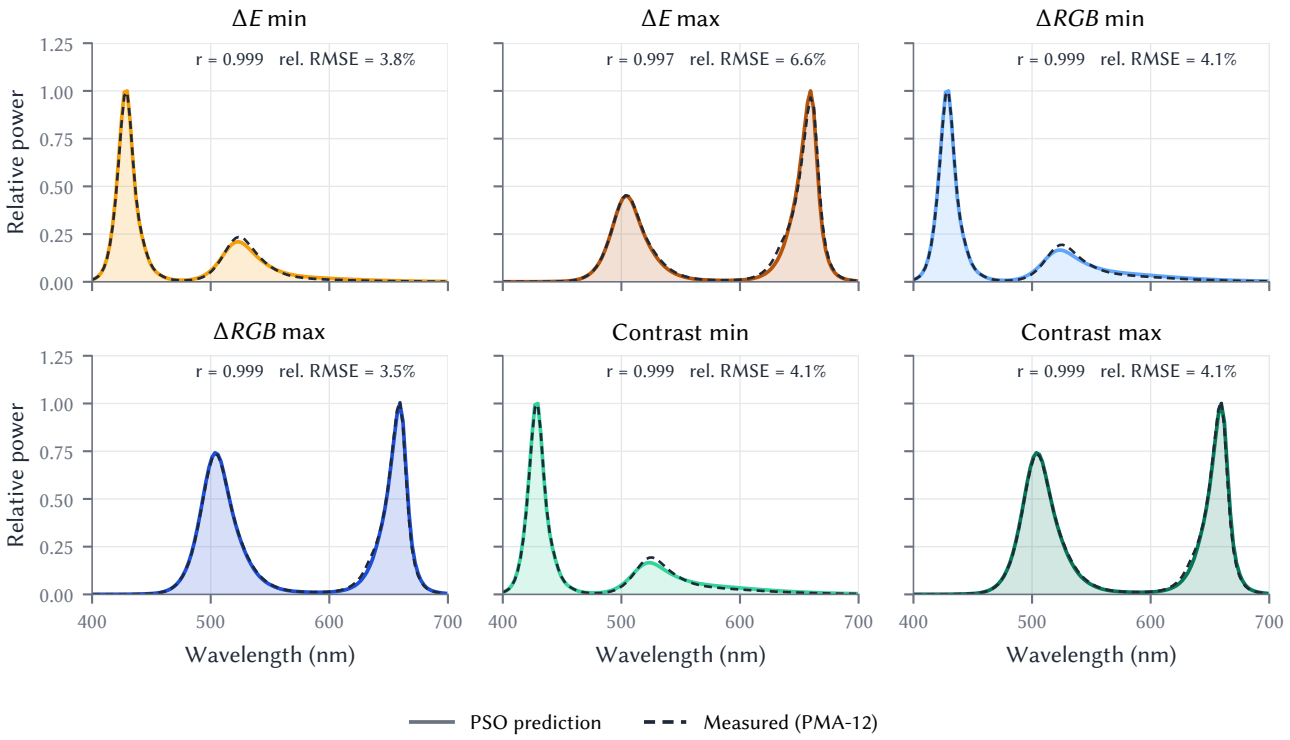
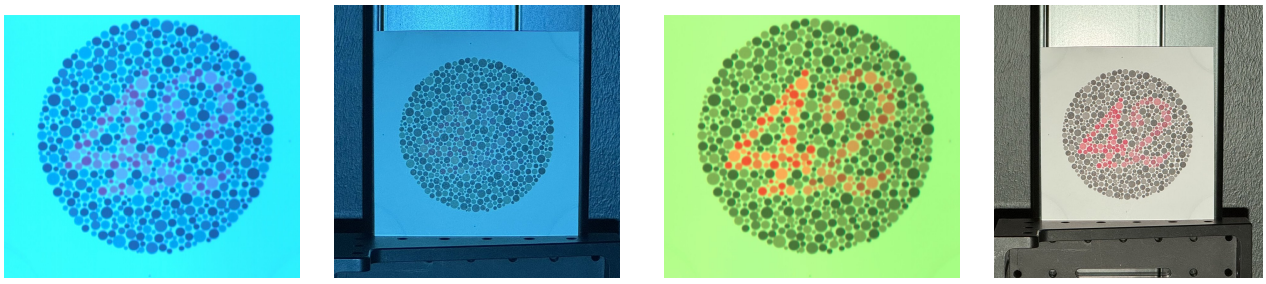


Figure 8. Predicted emission $\sum_j \mu_j^\star L_j(\lambda)$ (filled) against the light the Ledmotive produced (dashed), each normalised to the predicted peak.



(a) ΔE^*_{ab} min, rendered **(b)** ΔE^*_{ab} min, photo **(c)** ΔE^*_{ab} max, rendered **(d)** ΔE^*_{ab} max, photo

Figure 9. Renderings next to phone photographs of the plate under the reproduced lights. The camera's automatic white balance neutralises the light colour, so the photos show the effect only qualitatively: the numeral fades under the minimising light and stands out under the maximising one.

runs from those files alone; the optimisation needs the plate's reflectance cube.

References

[1] J. Kennedy and R. Eberhart. Particle swarm optimization. *Proc. IEEE International Conference on Neural Networks*, vol. 4, pp. 1942–1948, 1995.

[2] CIE 015:2018. *Colorimetry*, 4th edition. Commission Internationale de l'Éclairage, 2018.

[3] IEC 61966-2-1:1999. *Multimedia systems and equipment – Colour measurement and management – Part 2-1: Default RGB colour space – sRGB.*

[4] Y. Ohno. CIE fundamentals for color measurements. *Proc. IS&T NIP16 Conference*, pp. 540–545, 2000.