Why the Circadian Protection Factor Is Not a Useful Circadian Display Metric

At the Eyesafe 3.0 test condition, CPF reduces to an affine rescaling of the melanopic Daylight Efficacy Ratio: CPF = 93 – 24·(mel-DER). Every measured commercial display, every measured Korrus and partner backlight, and every realistically modeled architecture clears the CPF35 pass mark by 26 to 40 points; the threshold binds only on a theoretical cyan-primary construct nobody ships, and even when it discriminates it scores a static, dose-, duration-, and time-of-day-blind quantity that does not represent circadian impact at the eye.

Ben Harrison, Chief Innovation Officer & Chief Scientist · June 2026

Key Insights

Closed form. At Eyesafe's 200 cd/m² test condition, CPF = 93 – 24·(mel-DER), reproduced to r^2 = 1.000 across 36,250 sampled designs

Non-binding for real displays. 40 measured consumer displays sit at CPF_raw 64–74; 11 measured and canonical architectures span CPF_raw 61–80; no real or realistic design comes within 26 points of the CPF35 fail line

Binds only on a theoretical construct. A deliberately built cyan primary at the melanopic peak (488–490 nm) can fail CPF35, but only by collapsing the gamut onto NTSC 72%; no shipping panel approaches it

Wrong quantity for circadian impact. mel-DER carries no dose, no field fraction, no distance, no duration, and no time of day; a static reward for monotonic melanopic reduction is sign-wrong in the daytime

Summary

The Eyesafe Display Requirements 3.0 specification, issued in February 2025, introduces a Circadian Protection Factor (CPF) as an optional component of its display certification and pairs the CPF35 pass mark with the marketable label "circadian-protective." This paper makes the case that CPF, as defined and as deployed, is not a useful circadian metric for the displays consumers actually own. Two claims, each independent of the other, support that conclusion.

First, the CPF pass/fail does not discriminate among real products. At Eyesafe's 200 cd/m² test condition, CPF reduces to a clean affine function of the melanopic Daylight Efficacy Ratio: CPF = 93 − 24·(mel-DER). Across 40 commercial displays measured with a NIST-traceable spectroradiometer, 11 measured and canonical display architectures, and 36,000 synthetic Gaussian-RGB designs sampled inside standard monitor-color parameter ranges, *no* color-compliant display approaches the CPF35 fail line. Every real or realistically modeled display sits 26 to 40 CPF points above the pass mark. The threshold *can* bind, but only on a theoretical construct: a deliberately built Gaussian design with the "blue" primary placed onto the melanopic peak at 488–490 nm and squeezed through the loosest gamut criterion the standard offers (NTSC ≥ 72%). That construct is a saturated cyan with a collapsed gamut, sitting well above the blue-primary range of every measured display in the evidence base, including the fluxometer-measured consumer displays whose blue primaries cluster in 446–462 nm (with one historical CCFL exception at 436 nm). Failing CPF35 is theoretically possible. It requires deliberate engineering of a non-viewable cyan primary. No commercial display is anywhere near it.

Second, the scored quantity is the wrong quantity for circadian impact. mel-DER is scale-invariant by construction: it judges spectral shape against a fixed 200-nit benchtop condition only. It carries no information about absolute melanopic Equivalent Daylight Illuminance (mEDI) at the eye, field fraction of the display in the viewer's visual field, viewing distance, exposure duration, time of day, or observer adaptation state. Circadian impact is a function of all of those. A static "lower melanopic is better" score is also mis-signed against the daytime endpoint: a higher CPF can correspond to a *worse* daytime entrainment signal. And because the metric is a spectral-shape ratio at fixed luminance, two displays at identical white point and chromaticity can differ in mel-DER by 2.5 to 3× at the same nominal "color," which means CPF treats colorimetrically identical displays as biologically very different. Color is not a circadian proxy. CPF, at the certified-output level, is.

The Eyesafe definition rests on a legitimate input. The CIE S 026:2018 alpha-opic toolbox and the melanopic Daylight Efficacy Ratio are sound, well-supported instruments for describing the spectral content of a light source's effect on the melanopic photoreceptor (CIE, 2018; Lucas et al., 2014). The fault is not in mel-DER. It is in what CPF wraps around mel-DER: a static, condition-free 200-nit normalization that drops the dose-and-time information; an affine rescaling that pretends to be an independent score; and a threshold calibrated to a colorimetric edge case that no shipping display crosses. The credential conveys no information that distinguishes the displays a buyer can actually buy.

This paper does not propose a replacement specification. The constructive direction (an absolute, time-resolved mEDI-at-the-eye target measured under representative use) is the subject of a companion Korrus standard paper. The job here is the diagnostic.

The Eyesafe CPF construction

Eyesafe Display Requirements 3.0 §5.7 defines the Circadian Protection Factor as a transformation of the display's melanopic Daylight Efficacy Luminance (MEDL) under a fixed test condition. The display under test is set to 200 cd/m² of luminous output and the full-screen white spectral power distribution is captured. MEDL is the melanopic-EDI luminance at that condition: it is the photopic luminance (200 cd/m²) multiplied by the display's mel-DER, where mel-DER is the standard CIE S 026 melanopic-to-photopic action-spectrum integral of the display white normalized so that CIE Standard Illuminant D65 returns exactly 1.000. CPF is then defined as

> *CPF_raw = −0.12 · MEDL + 93,*

with a rounding rule that floors the raw value to the nearest 10, with a special 35-bin for raw values in [35, 40), and a marketable pass threshold at 35. Certified scores can therefore land on any multiple of 10, plus the 35-bin; the public-facing scale around the pass mark is {35, 40, 50, 60, 70, 80}, while a failing design rounds to 30, 20, or below. Substituting MEDL = 200·(mel-DER) collapses this to a closed form:

> *CPF_raw = 93 − 24 · (mel-DER).*

This identity is exact. It is not an approximation that holds for most displays; it is what the formula resolves to at the 200-nit test condition, by elementary substitution. A simulation harness implementing the pipeline from first principles (Gaussian RGB primaries, white-balance to a target daylight chromaticity, integration against the CIE S 026 melanopic action spectrum and the photopic luminosity function V(λ), and Eyesafe rounding) verifies it across the entire design space: 36,250 realizable samples, mean residual zero, maximum residual below 1×10^{-9}, linear fit slope −24.000 and intercept 93.000 with r^2 = 1.000000. The D65 reference returns mel-DER = 1.000, MEDL = 200, CPF_raw = 69, certified CPF 60.

Three consequences follow from the closed form, and they are the entire diagnostic case.

First, CPF is not an independent scoring axis. It is a one-to-one, monotonic, affine rescaling of mel-DER. Anything CPF says about a display is something mel-DER already said. The number 93 sets where D65 maps to; the number −24 sets the slope. Neither carries information about the display beyond what is already in mel-DER. Calling the output by a different name (and dressing it in pass/fail thresholds and rounded ladders) does not make it a different measurement.

Second, the CPF35 pass mark corresponds to a single sharp mel-DER threshold. Solving the identity at CPF_raw = 35 gives mel-DER = (93−35)/24 = 2.4166…. Any display with mel-DER above 2.42 fails; any display below it passes. The cutoff at the rounded score 35 is therefore equivalent to a hard mel-DER cutoff at 2.42, which is roughly 2.4× the D65 reference. To fail CPF35, a display has to emit more than two and a half times the melanopic content of equal-luminance D65 daylight in its white SPD. That is a strong requirement.

Third, CPF carries no information about dose, time, or anything other than spectral shape. The construction starts by *fixing* luminance at 200 cd/m², so absolute melanopic illuminance is destroyed before mel-DER is computed; mel-DER is a ratio, so it is scale-invariant by definition; and nothing in the pipeline asks about field fraction, viewing distance, exposure duration, time of day, or observer adaptation state. None of those quantities is even in the construction's vocabulary. Whatever CPF measures, it does not measure those things.

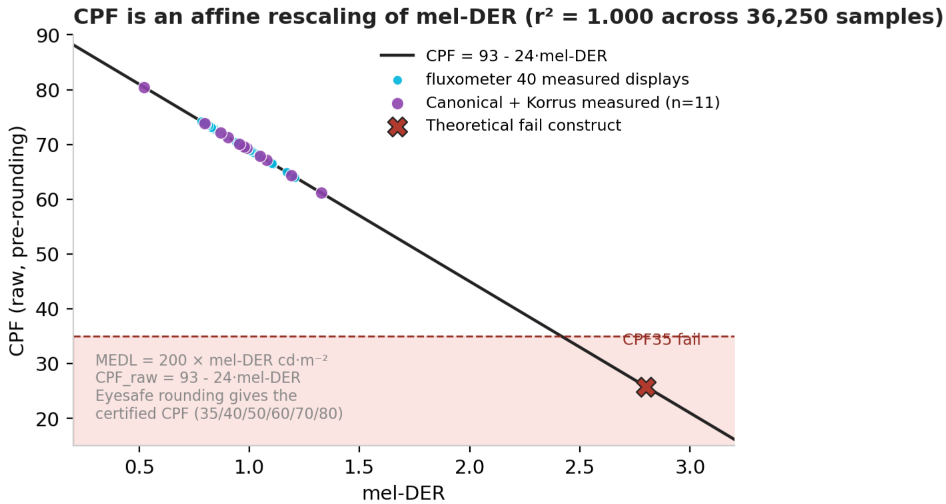

Figure 1. CPF is an affine rescaling of mel-DER. The solid line is CPF = 93 – 24·(mel-DER); every realizable sample lands exactly on it. The 40 fluxometer-measured commercial displays (cyan) and the 11 canonical and measured architectures (purple) cluster between mel-DER 0.78 and 1.33, well above CPF 60. The theoretical fail construct (×) sits at mel-DER 2.80, CPF 26, deep inside the shaded fail region; the next subsection shows what kind of "display" it actually is.

What the empirics show

A static metric calibrated to a particular spectral threshold only does work if real displays vary across the threshold. They do not. Three independent populations were evaluated through the same pipeline.

40 measured consumer displays (fluxometer.com)

The fluxometer.com archive (f.lux Software LLC, CC BY 4.0) publishes full visible-band spectroradiometer measurements (PR-655, NIST-traceable) for 40 consumer displays at their default 6500 K white, full brightness, captured under the same protocol. The set spans roughly fifteen years of consumer hardware: phones (iPhone 3GS through iPhone X, multiple Galaxy and Nexus generations), tablets (iPad 1 through iPad Pro), laptops (MacBook Pro 2009, MacBook Pro 16-inch mini-LED 2021, Thinkpad and Asus models), and standalone monitors (Thunderbolt 27, Surface Pro 3). Architectures include WLED-LCD, KSF-phosphor LCD, IPS, AMOLED, and mini-LED.

Run through the harness: mel-DER spans 0.78 to 1.21, CPF_raw spans 64.0 to 74.3, certified CPF is 60 or 70, and zero displays fail CPF35. The closest-to-binding measured display is the iPhone 3GS at mel-DER 1.21, CPF_raw 64. It sits 1.21 mel-DER units (or 29 CPF points) below the fail line. Across the 40 fluxometer displays, 39 blue-primary peaks fall in 446–462 nm; one outlier from a 2008 CCFL-backlit iMac sits at 436 nm. None is anywhere near the 488–490 nm wavelength range at which the threshold binds.

Canonical architectures and measured Korrus and partner backlights

Eleven additional displays were evaluated: six canonical commercial architectures constructed from published peak and FWHM disclosures for the standard architecture families (WLED-LCD, QD-LCD in Samsung QLED style, OLED RGB stripe, QD-OLED in Samsung S95B style, WOLED+RGBW in LG C-series style, KSF + β-SiAlON LCD; published peak and FWHM values for canonical consumer architectures are reported in the OLED vision-science measurement work of Abu Haila et al., 2025), and five measured backlights: two Korrus reference backlights, two partner sample backlights, and a third Korrus backlight included specifically as a stress case for the upper end of product-side melanopic content. The intent of including that stress-case backlight is diagnostic: across the populations sampled, it produces the highest melanopic content among the Korrus measurements, and is included here to mark that envelope. Spectral parameters of the Korrus measurements are held internal; only the resulting mel-DER and CPF readouts are reported.

The highest mel-DER among these 11 is 1.33 (a partner sample backlight, CPF_raw 61). The Korrus stress-case backlight reaches mel-DER 1.19, CPF_raw 64. The lowest mel-DER is 0.52, CPF_raw 80, from one of the Korrus reference backlights. The 11-architecture range spans CPF 61 to CPF 80, a factor of 2.6 in mel-DER, and never within 26 CPF points of the fail line. The Korrus stress-case backlight is still 1.23 mel-DER units short of failing CPF35. Read against the dose-and-time argument developed below, a backlight's higher daytime melanopic content is the desirable property for daytime alignment; CPF treats it as a demerit. The mis-signaling, not the backlight, is the metric's problem.

36,000 synthetic Gaussian-RGB designs in the standard's monitor-color window

To probe the full feasible design space rather than just shipped hardware, 36,250 Gaussian-RGB primary sets were swept across the standard monitor-color parameter ranges (R peak 600–640 nm, G peak 510–540 nm, B peak 440–470 nm; FWHM 15–60 nm per primary; target CCT 5,500–7,000 K), white-balanced to the daylight locus at each target CCT, and evaluated through the same identity. 34,337 samples (94.7%) satisfy the standard's color-compliance criterion (CCT in range *and* gamut coverage above one of {sRGB ≥ 95%, DCI-P3 ≥ 90%, Adobe RGB ≥ 90%, NTSC ≥ 72%}). Of those, compliant CPF_raw ranges from 49.1 to 79.8 (median 68.1, mean 67.8), compliant mel-DER ranges from 0.57 to 1.83, and zero compliant samples fail CPF35. The maximum compliant mel-DER (1.83) sits 0.59 below the fail line at 2.42, leaving 24% of the fail-side margin unused even at the worst spectral edge of the modeled monitor-color space.

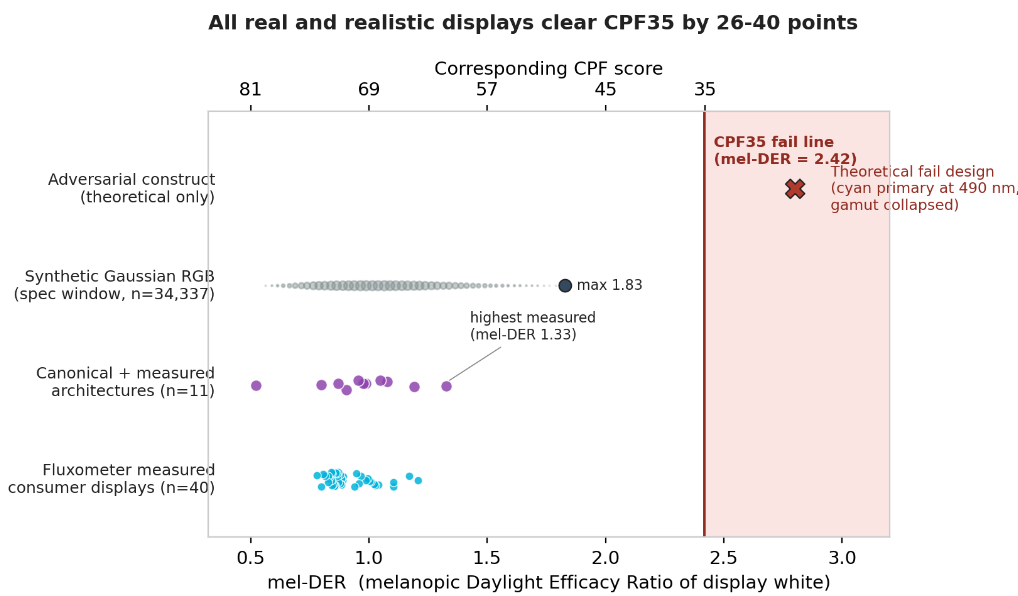

Figure 2. All four populations on the same mel-DER (and CPF) axis. From bottom: 40 fluxometer measured consumer displays cluster between mel-DER 0.78 and 1.21 (CPF 64–74); 11 canonical and Korrus-measured architectures span 0.52 to 1.33 (CPF 61–80); 34,337 synthetic Gaussian-RGB designs in the standard's monitor-color window reach up to 1.83 (CPF 49); the adversarial construct sits alone past 2.42, deep inside the CPF35 fail region. No overlap between the realistic distribution and the binding regime.

A single sentence the empirics support

Across 51 measured or canonical real-world displays and 34,337 color-compliant synthetic designs sampled inside standard monitor-color parameter ranges, zero approach the CPF35 fail line. Every real or realistic design clears the threshold by 26 to 40 CPF points. The credential is a uniform pass with no power to distinguish among the displays a buyer can purchase.

The threshold is not vacuous: where it binds, and why that does not save it

An obvious objection deserves to be conceded fully and addressed precisely. CPF35 *can* be failed. The naive universal claim ("no color-compliant display can ever fail CPF35") is false, and a careful reviewer would reject it immediately. The right empirical claim is the more restrained one above: no real or realistically modeled display approaches the fail line. To find a failing design at all, the harness has to step outside the modeled monitor-color window and into a regime that no commercial product occupies.

The binding regime, located by stress-testing the spec, is narrow and specific. It requires the blue primary's peak wavelength to sit at 488–490 nm, right on the peak of the CIE S 026 melanopic action spectrum (the action curve's peak is near 490 nm). It requires a narrow FWHM, on the order of 10 nm, so the entire emitter sits inside the 440–495 nm melanopic band. The display white must be balanced to the warm end of the compliance window (≈ 7,000 K) so the green and red weights do not pull the integral back down. And it must clear color-compliance through *one* gate only: NTSC 72%, the loosest of the four. A specific instance of this construction, found in the sweep, has μ_B = 490 nm, FWHM_B = 10 nm, μ_G = 520 nm, μ_R = 650 nm, CCT 7,000 K. It produces mel-DER 2.80, CPF_raw 26 (rounded 20). Its sRGB coverage is 65.7% (fails the 95% threshold). Its DCI-P3 coverage is 65.4% (fails 90%). Its Adobe RGB coverage is 71.4% (fails 90%). Its NTSC coverage is 72.7%, just past the 72% gate.

What this "display" actually is, in colorimetric terms, is a saturated cyan-primary architecture with a collapsed gamut. A Gaussian "blue" with peak at 490 nm sits well above the highest measured real-world blue primary in any tested population, and 26 to 42 nm above the fluxometer-measured consumer displays (446 to 462 nm range, with one historical CCFL outlier at 436 nm). Its chromaticity is approximately (0.04, 0.30), which is unambiguously cyan, not blue, on a CIE 1931 plot. The display triangle that results from substituting it for the standard blue corner is severely skewed: it stretches up the spectral locus toward the green region, covers nowhere near sRGB blue, and only clears NTSC 1953's wide reference triangle by virtue of NTSC 1953's particular geometry. Such a display has no usable blue. It cannot render the sky, deep blue clothing, the standard playback blue of any video standard, or the blue half of essentially every sRGB or P3 image asset on the open internet. It is a thing a competent display engineer can describe and a synthesis harness can construct, not a thing a panel maker can ship.

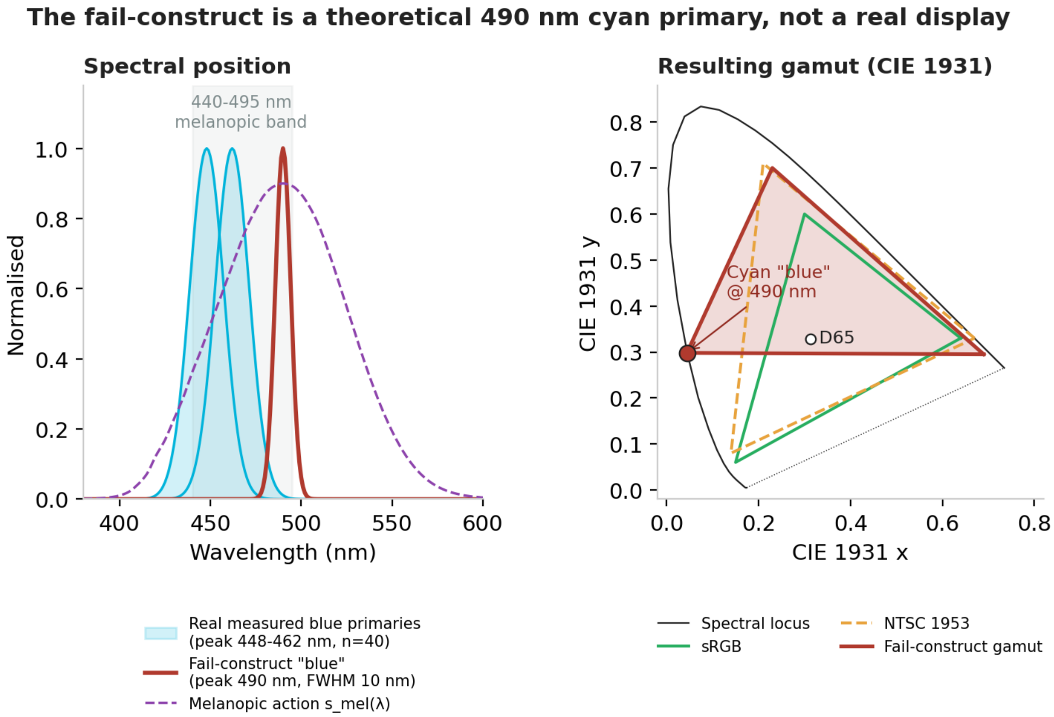

Figure 3. The fail construct, side by side with real displays. *Left:* the 40 measured consumer-display blue primaries (cyan envelope) sit between 448 and 462 nm; the fail-construct primary (red) sits at 490 nm, right at the peak of the CIE S 026 melanopic action curve (dashed). The wavelength gap is 25 to 40 nm. *Right:* the fail-construct's gamut triangle in CIE 1931. The "blue" corner has migrated up the spectral locus into the cyan region; the triangle no longer covers sRGB or DCI-P3, only clearing NTSC 1953's 72% gate by a thin margin.

The CPF35 threshold catches this design. The standard's gamut criteria do not. The threshold is doing real standards-engineering work: it forecloses a class of spectrally adversarial Gaussian-RGB displays that the gamut and CCT requirements would otherwise admit. That function is real and the brief here is not to deny it.

It does not save the metric, for three reasons. First, the class of designs being excluded is empty. No commercial display approaches the cyan-primary regime today, and the design move required to do so (a roughly 30 percentage point gamut concession against sRGB and DCI-P3, with coverage collapsing to about 65%) is on the order of a structural product compromise no current positioning supports. The threshold is catching a class of designs that does not exist outside the harness. Second, even if standards-driven incentives moved primary placements over time (the historical precedent for spectrum-driven design shifts under certification pressure is real), the move toward a 488–490 nm blue primary still requires that ~30-point gamut concession, which is qualitatively larger than the marginal shifts certifications have produced elsewhere. The critique therefore stands today and is not obviously time-limited; if it ever does become time-limited, it will be because the credential has succeeded in pulling commercial designs toward an undesirable colorimetric corner, which is its own problem. Third, even granting the exclusion is real, it does not make CPF a useful *circadian* metric for the displays that do exist. A credential whose entire discriminative power lives on a fictional cyan-primary architecture cannot distinguish the products a buyer can buy. Failing CPF is theoretically possible. It takes deliberate effort. No normal display is anywhere near it. That is the honest reading.

The relative-ranking objection, and why the score does not survive it

A buyer might concede the pass/fail point but use the underlying score as a relative ranker: two displays, both pass, one returns CPF_raw 70 and the other CPF_raw 60, choose the 70 all else equal. This is the most likely surviving defense of the credential. The score does not survive the move, for the same reasons the threshold does not, only sharper. Within the all-pass band, CPF is monotone in mel-DER by the identity CPF = 93 − 24·(mel-DER), so choosing the higher CPF means choosing the display with lower mel-DER. Lower mel-DER at fixed luminance is a weaker daytime entrainment signal, which is the wrong direction during the day; and lower mel-DER at the 200 cd/m² benchtop condition says nothing about what the display delivers at the eye in actual use, because the same SPD shape can be deployed at any luminance and any field fraction. Buyer-relative ranking on CPF, in other words, ranks displays by the wrong axis (static spectral shape) with the wrong sign for the daytime endpoint (lower-is-better) at a misrepresentative scale (200 cd/m² benchtop). Using the score as a finer ranker means giving the sign-wrong direction more granular weight, not less.

Even where CPF discriminates, the quantity is wrong

Set aside the pass/fail question. Suppose the score itself (a number on a continuous scale from 0 to 93) is read as a comparative measure of “circadian protection.” It is not, because the underlying mel-DER is the wrong quantity for circadian impact at the eye. The argument has three threads.

CPF carries no dose information

The Eyesafe construction fixes luminance at 200 cd/m² before mel-DER is computed. mel-DER is then a ratio, scale-invariant by definition. The certified output of the CPF measurement is therefore identical for a 200 cd/m² display and a 20 cd/m² display with the same SPD shape; identical for a 6-inch phone and a 65-inch TV with the same SPD shape; and identical for a display held at arm’s length and one mounted across the room. Absolute melanopic Equivalent Daylight Illuminance at the eye, the quantity that drives the melanopic response (Brown et al., 2022; Lucas et al., 2014), depends on the source’s luminance *and* the angular size it presents at the eye *and* the viewer’s distance and gaze direction. CPF carries none of those.

Real-world melanopic dose at the eye from a display varies by more than an order of magnitude across plausible use conditions even for a single device. A 14-inch laptop at 200 cd/m² at 60 cm of viewing distance and a 65-inch TV at 200 cd/m² at 3 m of viewing distance deliver very different melanopic illuminances at the eye despite returning identical CPF scores. The certified CPF rating cannot tell those situations apart.

CPF carries no time-of-day information, and the sign is wrong in the daytime

The melanopic response is mediated by melanopsin-expressing intrinsically photosensitive retinal ganglion cells (ipRGCs), the photoreceptor class identified in the early 2000s through the work of Berson, Hattar, and colleagues (Hattar et al., 2003). Its behavioural and physiological effects depend strongly on the time of day at which the exposure occurs. Higher melanopic illuminance during the biological day is broadly beneficial: it consolidates entrainment, supports daytime alertness and mood, and reduces evening sleep latency. Higher melanopic illuminance during the biological evening and night is detrimental: it suppresses melatonin, delays the body’s evening melatonin onset, and degrades sleep depth and continuity. The published consensus recommendations for healthy adults make this time-of-day split explicit, with separate daytime, pre-bed, and night targets resolved by hour rather than by spectrum alone (Brown et al., 2022). A circadian metric must respect this sign reversal, because a quantity that is good in one window is bad in the other.

CPF rewards monotonic reduction of mel-DER at all times. A higher CPF is, by construction, always "better." A display that uniformly cuts its melanopic output across the entire day would maximize CPF and minimize evening melatonin suppression. It would also degrade daytime entrainment, because the daytime signal is exactly what entrainment requires. The metric is sign-wrong in the daytime by construction. This is not a tunable parameter to be added later; it is a structural property of taking a static, lower-is-better number and labeling it "circadian protection."

CPF treats colorimetrically identical displays as biologically different, and vice versa

At a fixed white point, mel-DER varies widely across architectures because the underlying spectra are very different beneath the same chromaticity. Among the 11 canonical and measured architectures at 6,508 K, mel-DER spans 0.52 to 1.33, a 2.6× range. Among synthetic compliant designs balanced to a 6,500 K daylight white, mel-DER spans 0.57 to 1.83, a 3.2× range. Among the fluxometer measurements at CCT close to 6,500 K, mel-DER spans 0.78 to 1.03, even within the modest spectral diversity of shipped hardware. Same white point. Same color appearance. Up to three times the melanopic content. Conversely, two displays at very different CCTs (e.g., a warm 5,500 K display and a cool 6,500 K display) can have identical mel-DER if their underlying SPDs happen to integrate the same way against the melanopic action curve.

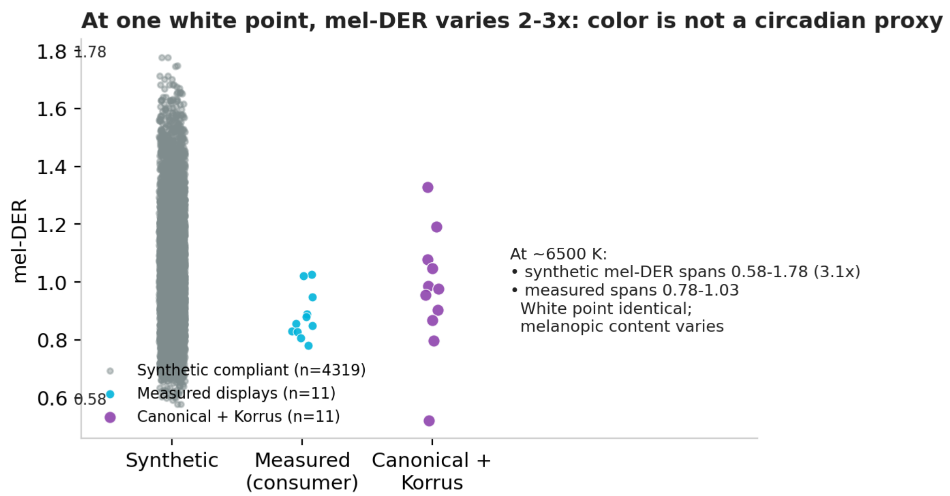

Figure 4. At a single white point (≈ 6,500 K), mel-DER (and therefore CPF) varies by a factor of 2.5 to 3 across the synthetic and architecture populations. Same chromaticity, same nominal color, very different melanopic content. Whatever CPF is measuring, it is not measuring "color" and color cannot be substituted for it; CCT and white point are not circadian proxies, and the converse holds equally.

This is the metameric spread result, and it has a corollary that Korrus has stated elsewhere and the data here corroborates: color and circadian content are independent dimensions. CCT is not a circadian proxy. "Warm dim" is not a circadian-relevant variable. The standard CCT and gamut criteria of a display tell a buyer nothing about the display's melanopic load. The CPF rating, viewed as a circadian metric, suffers from the same problem in reverse: it tells the buyer about static spectral shape at one luminance, which is not what a circadian outcome depends on. A static, single-number score for a system whose appropriate signal varies continuously across the day cannot, in principle, capture what the system requires.

A note on conflation

The Eyesafe 3.0 specification combines a Radiance Protection Factor (RPF), which addresses retinal hazard via the IEC 62471 blue-light-hazard action spectrum B(λ) peaking at 435–440 nm, with the CPF construction this paper examines, which addresses circadian effect via the CIE S 026 melanopic action spectrum peaking near 490 nm. These are different curves for different biological endpoints. CIE has explicitly stated that the term “blue light hazard” should not be used when referring to circadian rhythm disruption or sleep disturbance (CIE, 2019), and that this is a vocabulary distinction the field must hold. A specification that bundles both factors under a single display-certification scheme invites conflation in the downstream marketing, which the underlying standards (IEC 62471 for hazard; CIE S 026 for circadian) have worked to prevent. The blue-light-hazard sections of Eyesafe 3.0 are out of scope here and are the subject of a separate Korrus paper (KORRUS-WP-2026-001), which reaches a different conclusion about the regulatory basis for treating consumer displays as a blue-light hazard than the certification framework assumes; that conclusion is reported in that paper and is not relitigated here. The point relevant here is narrower: the bundling itself, irrespective of either section’s internal merits, drives a public conflation of hazard and circadian that the science distinguishes.

What a defensible circadian display specification requires

The constructive program is the subject of a separate paper. The diagnostic version, here, is one paragraph.

A defensible circadian display specification has to express the quantity that actually drives the circadian system: absolute melanopic Equivalent Daylight Illuminance at the eye, under representative use conditions, resolved by time of day. “At the eye” requires luminance, field fraction, and viewing distance to be specified or estimated for a representative use case rather than fixed at a benchtop value. “Under representative use” requires that the specified condition correspond to how the device is actually viewed. “Resolved by time of day” requires that the daytime endpoint and the evening endpoint be specified separately, because they have opposite valences, and a single number that minimizes both is impossible. The first condition (absolute mEDI at the eye) eliminates the scale-invariance error; the second (representative use) eliminates the 200-nit-benchtop error; the third (time-of-day resolution) eliminates the sign error. None of these conditions is exotic; they correspond to how the published consensus recommendations on healthy daytime and evening light exposure are framed (Brown et al., 2022). The full target language belongs in the standard paper.

What CPF captures vs. what circadian impact requires

Quantity	Captured by CPF
Spectral shape of display white	✓ yes
Absolute melanopic EDI at the eye	✗ no
Field fraction of display	✗ no
Viewing distance	✗ no
Exposure duration	✗ no
Time of day	✗ no
Observer adaptation state	✗ no

CPF is a static, scale-invariant ratio at one benchtop condition.
Circadian impact at the eye is a function of all seven rows.

Figure 5. What CPF captures versus what circadian impact requires. The metric represents the first row only. The remaining six are the levers that determine actual melanopic effect at the eye and that a defensible specification must address.

A note on standing and reproducibility

Korrus has a commercial interest in display standards and a companion paper proposes an alternative direction. The diagnostic case here stands or falls on the algebra and the measurements, not on the source. The algebra (CPF = 93 − 24·(mel-DER) at the 200 cd/m² test condition) is reproducible by inspection from the published Eyesafe formula. The synthetic populations are reproducible bit-for-bit from the deposited simulation harness with deterministic seed 42. The 40 consumer-display measurements are public-domain fluxometer.com spectra under CC BY 4.0; the 11 architecture data points include canonical architectures built from published peak and FWHM disclosures (Abu Haila et al., 2025) and measured Korrus and partner backlights, the latter of which can be substituted with any independently measured set without changing the conclusion. Nothing load-bearing in the critique depends on Korrus-internal spectral data. The headline finding (the wide separation between every realistic display and the failing construct) lives entirely in the published fluxometer set, the public synthetic sweep, and the Eyesafe formula.

What this paper has established

At Eyesafe's 200 cd/m² test condition, CPF reduces in closed form to CPF = 93 − 24·(mel-DER), verified to $r^2 = 1.000$ across 36,250 sampled designs. CPF35 is equivalent to a hard mel-DER threshold at 2.42. Empirically, every one of 40 spectroradiometer-measured consumer displays clears that threshold with CPF_raw 64 to 74, every one of 11 measured and canonical commercial architectures clears it with CPF_raw 61 to 80, and every one of 34,337 color-compliant synthetic designs inside the standard monitor-color window clears it with CPF_raw 49 to 80. No real or realistic display approaches the fail line. The threshold *can* bind, but only on a deliberately constructed cyan primary at 488–490 nm with a collapsed gamut that only clears NTSC 72%, an architecture no commercial display occupies and no manufacturer would deliberately ship. The scored quantity, even where it varies, is the wrong quantity for circadian impact: it has no dose, no field fraction, no distance, no duration, no time of day, and no adaptation state, and a static lower-is-better number is sign-wrong against the daytime endpoint. The legitimate input (mel-DER from CIE S 026) is not the problem. The problem is what CPF wraps around it.

The defensible reading of the credential is therefore as follows. A CPF35 pass mark on a consumer display certifies that the display is not a deliberately constructed adversarial cyan-primary architecture. It certifies nothing about the display's circadian impact in use, nothing about its melanopic dose at the eye, and nothing about its appropriateness for the time of day at which it is being viewed. As a label intended to inform display buyers and to ground product marketing about "circadian protection," it is non-informative. The displays a consumer can buy all pass, by a wide margin. The displays that fail do not exist outside a simulation harness.

The constructive direction, an absolute, time-resolved mEDI-at-the-eye target measured under representative use conditions, is taken up in the companion Korrus standard paper.

References

Abu Haila, T., Kunst, K., Khanh, T. Q., & Wallis, T. S. A. (2025). Recent consumer OLED monitors can be suitable for vision science. *Journal of Vision* 25(2):11. https://doi.org/10.1167/jov.25.2.11 (Preprint: arXiv 2410.17019.)

Brown, T. M., Brainard, G. C., Cajochen, C., Czeisler, C. A., Hanifin, J. P., Lockley, S. W., Lucas, R. J., Münch, M., O'Hagan, J. B., Peirson, S. N., Price, L. L. A., Roenneberg, T., Schlangen, L. J. M., Skene, D. J., Spitschan, M., Vetter, C., Zee, P. C., & Wright, K. P. (2022). Recommendations for daytime, evening, and nighttime indoor light exposure to best support physiology, sleep, and wakefulness in healthy adults. *PLOS Biology* 20(3):e3001571. https://doi.org/10.1371/journal.pbio.3001571

CIE (International Commission on Illumination). (2018). *CIE S 026/E:2018, System for Metrology of Optical Radiation for ipRGC-Influenced Responses to Light.* CIE Central Bureau, Vienna.

CIE (International Commission on Illumination). (2019). *CIE Position Statement on the Blue Light Hazard (April 23, 2019).* CIE Central Bureau, Vienna. https://www.cie.co.at/publications/position-statement-blue-light-hazard-april-23-2019

Eyesafe. (2025). *Eyesafe Display Requirements 3.0.* Eyesafe Inc., released early 2025 (announced at CES 2025). CPF construction: §4 and §5.7.

f.lux Software LLC. (n.d.). *fluxometer.com: Spectroradiometer measurements of consumer displays.* Licensed CC BY 4.0. Available at https://fluxometer.com/rainbow/.

Hattar, S., Lucas, R. J., Mrosovsky, N., Thompson, S., Douglas, R. H., Hankins, M. W., Lem, J., Biel, M., Hofmann, F., Foster, R. G., & Yau, K. W. (2003). Melanopsin and rod-cone photoreceptive systems account for all major accessory visual functions in mice. *Nature* 424(6944):76–81. https://doi.org/10.1038/nature01761

Lucas, R. J., Peirson, S. N., Berson, D. M., Brown, T. M., Cooper, H. M., Czeisler, C. A., Figueiro, M. G., Gamlin, P. D., Lockley, S. W., O'Hagan, J. B., Price, L. L. A., Provencio, I., Skene, D. J., & Brainard, G. C. (2014). Measuring and using light in the melanopsin age. *Trends in Neurosciences* 37(1):1–9. https://doi.org/10.1016/j.tins.2013.10.004

Korrus internal measurements. Backlight spectra for three Korrus LCD backlights (two reference and one stress-case for elevated melanopic content) and two partner sample backlights (PR-655 spectroradiometer, Q1 2026). Spectral details are held internal; only mel-DER and CPF readouts are reported in this paper. Simulation pipeline and outputs are deposited in `cpf_sim/` of the project repository, with deterministic seed 42.