Modeling of the discrete RGB digital color space to the visible color spectrum

J. GRĄTKIEWICZ

jakub.gratkiewicz@wat.edu.pl

Military University of Technology, Faculty of Cybernetics
Kaliskiego St. 2, 00-908 Warsaw, Poland

The author asks a basic question: how does the world of color work, and how does a continuous physical spectrum become a perceived hue. This perspective is then contrasted with digital practice, where colors are encoded discretely as RGB triplets and reproduced by displays that emit artificial light with device-specific spectral characteristics. The article examines whether a conversion from digital colors to a physical, wavelength-based description is possible, and whether such a conversion is meaningful. By considering metamerism, inter-device variability, viewing conditions, and observer-dependent perception, the author argues that assigning an RGB value to a single "true spectrum" is inherently non-unique and only weakly justified from a perceptual standpoint. A further conclusion is practical but important: even the finite RGB space is beyond direct human comprehension. Not all differences between RGB codes are perceptually discriminable, and even when they are, discrimination does not imply that people can reliably assign unique names to every visible difference. In effect, the infinite spectral domain is reduced to a finite set of digital RGB codes and then reduced again to an even smaller set of linguistic color categories. Building on this, the paper proposes heuristics that are sufficient for precise everyday use of color terminology, while acknowledging their heuristic nature: (1) an author-defined RGB color cube with vector-based region assignment for named colors, (2) quantitative sampling with a nearest labeled neighbor rule, and (3) data-driven classification on the sampled datasets, including supervised models and clustering methods.

Keywords: computer graphics, colors, ML.

DOI: 10.5604/01.3001.0055.7728

1. Introduction

In conventional digital computers, any continuous quantity must ultimately be encoded in a discrete form. This follows from how hardware stores and processes information: numbers are represented in binary, and the available memory has a finite word length. With N bits, a register (or memory cell) can represent at most 2^N distinct patterns, so only a finite set of values can be stored exactly. Therefore, the real numbers-an infinite, uncountable set-cannot be represented without loss. In practice, real-valued quantities are approximated with a chosen precision using finite expansions such as fixed-point formats or floating-point types (e.g., IEEE 754 single and double precision). This is a deterministic kind of approximation: increasing the number of bits increases the representable range and improves resolution, but true continuity and the "full" infinity of R remain out of reach for any finite digital representation [1].

This limitation is not just a theoretical curiosity; it has direct consequences in fields where the underlying phenomena are continuous, including computer graphics and color processing. In the physical world, color stimuli arise from electromagnetic radiation, and visible light spans a continuous band of wavelengths roughly from 380 nm to 780 nm. More generally, light is described by a spectral power distribution (SPD)-a function that assigns energy to wavelengths across that range. A computer, however, cannot store a continuum of wavelengths or an arbitrary SPD; it must work with discrete codes. In everyday imaging, the most common such code is RGB, where a color is specified by three channel values that belong to a finite set (typically 8 bits per channel in consumer pipelines).

At this point, an important conceptual gap appears: an RGB triplet is not a wavelength, and it does not uniquely determine a physical spectrum. RGB values are tied to a particular encoding and device model (e.g., a specific RGB color space, display primaries, and viewing assumptions), and many different spectra can look identical to a human observer (metamerism). Nevertheless, it is often useful – especially for modeling, visualization, and educational purposes – to construct a deterministic mapping that assigns a representative spectral description to a given digital color. In this work, we treat the problem explicitly as a projection: from the discrete set of RGB codes onto a wavelength-based description that can be interpreted in terms of the visible spectrum.

The goal of the article is therefore not to "reconstruct the true spectrum" behind an RGB value, but to formalize a repeatable, well-defined procedure that converts a digital RGB input into an estimated dominant wavelength (or an equivalent spectral proxy) under stated assumptions. We first introduce the necessary theoretical background and the historical context of digital color representation. Next, we review selected literature on bridging digital color codes and physical color stimuli. We then propose a mathematical model for mapping RGB values onto wavelength-defined spectral colors and finally provide an example implementation: a function that estimates dominant wavelengths for RGB colors in a way that is deterministic and reproducible.

2. Digital colors

One of the most common models for describing color in computer science is the RGB model, where a color is specified by three components: red (R), green (G), and blue (B). Each component takes a discrete value from a fixed range, most commonly 0-255 (8 bits per channel). This yields a total of 256 x 256 x 256 = 16,777,216 possible combinations, often referred to as true color (24-bit color) [3]. Although this is a very dense sampling of the RGB space and is sufficient for the vast majority of typical imaging tasks, it remains a finite and discrete set, whereas physical light stimuli form a continuous domain. In practice, 24-bit RGB is widely treated as "visually continuous" for human observers under normal viewing conditions. Adjacent code values (e.g., a change of 1 out of 255 in a channel) often produce differences that are small, and visible artifacts such as banding are more strongly influenced by image content, tone mapping, and rendering choices (including dithering) than by the raw count of available codes [4]. For this reason, true color has become the de facto standard in consumer and general-purpose graphics pipelines [3]. Still, it is important to keep the conceptual distinction clear: a digital color code is a deterministic, device-defined description, not a direct description of a physical spectrum. Historically, when graphics memory and bandwidth were more limited, systems relied on much smaller palettes and lower bit depths. Early computer displays were often monochrome, followed by color modes constrained to a handful of simultaneously displayable colors. For example, IBM's CGA adapter (1981) provided a 16-color palette, but typical graphics modes addressed only a subset of that palette at a time (e.g., 2 bits per pixel – up to 4 colors on screen). The later EGA standard (1984) expanded the palette to 64 colors and enabled 16-color display modes (4 bits per pixel). With VGA (1987), 256-color modes became common, while the hardware DAC typically supported 18-bit color output (262,144 possible analog RGB levels). By the 1990s, true color graphics hardware became widespread, providing approximately 16 million directly addressable colors without palette tables.

A simplified summary of common historical color depths is shown below:

- **Monochrome**: 1 bit/2 colors,
- **CGA**: 2 bits/4 colors,
- **EGA**: 4 bits/16 colors,
- **VGA**: 8 bits/256 colors,
- **High Color, XGA**: 16 bits/65,536 colors,
- **True Color, SVGA**: 24 bits/16,777,216 colors,
- **True Color + Alpha**: 32 bits/16,777,216 colors + alpha channel [3], [4].

True colors in a computer are often stored in hexadecimal notation, as so-called hex triples. A hex triple uses six hexadecimal digits, RRGGBB, to encode the red, green, and blue channel values as bytes (00-FF in hexadecimal corresponds to 0-255 in decimal).

Fig. 1. Color wheel
Source: https://www.deviantart.com/geckobot/art/Colour-Wheel-RGB-666094738]

For example, pure red is written as #FF0000 (maximum red, no green, no blue), while gold with components (R = 255, G = 215, B = 0) can be written as #FFD700. This notation is widely used in computer graphics, especially in web technologies (e.g., HTML and CSS). In practice, it is interpreted within a specific RGB color space – most commonly sRGB – which makes the encoding unambiguous once the color space is specified.

As display and imaging technology evolved, higher precision formats became common in professional workflows, such as 30-bit color (10 bits per channel) or even 36-bit variants. These increase the number of representable RGB codes to over a billion and enable smoother gradients, reduced banding, and, in appropriate pipelines, support for wider dynamic range (e.g., HDR) and more subtle tone reproduction. The fundamental principle, however, remains unchanged: digital colors form a deterministic, discrete set of (R, G, B) combinations, whereas physical light stimuli are continuous functions of wavelength (or mixtures of such continuous spectra). In other words, digital color spaces provide a practical approximation of human-visible color experience, not a direct representation of the underlying spectral reality.

3. Visible color spectrum

Light is electromagnetic radiation that is emitted or reflected by objects. The portion visible to humans is only a small slice of the electromagnetic spectrum and spans roughly from about 380 nm to 780 nm in wavelength [6]. Within this interval, wavelength is closely linked to hue: shorter wavelengths (about 380–450 nm) are typically perceived as violet and blue, mid-range wavelengths (about 500–570 nm) as green and yellow, and longer wavelengths (about 600–750 nm) as orange and red [7]. At the photon level, shorter wavelengths correspond to higher photon energy, while longer wavelengths correspond to lower photon energy.

$$E = \frac{hc}{l} \tag{1}$$

where:
h is Planck's constant,
c is the speed of light,
and λ is the wavelength.
When photons reach the retina, they can activate one of three cone receptor classes: S (short), M (medium), and L (long), which are most sensitive – in broad terms – to what we call blue, green, and red light, respectively [8]. The brain interprets color based on the relative stimulation of these three cone types. Importantly, most real-world light sources do not emit a single wavelength; instead, they produce a spectrum, meaning a distribution of intensities across many wavelengths.

In an idealized sense, "simple" (spectral) colors are associated with light whose energy is concentrated in a very narrow wavelength band – for example, a laser centered around 532 nm is perceived as green. Such sources can be treated as nearly single-wavelength stimuli, close to a delta-like peak in the spectral distribution. For orientation, spectral hues are often roughly grouped into wavelength bands: violet around 380–450 nm, blue around 450–495 nm, green around 495–570 nm, yellow around 570–590 nm, orange around 590–620 nm, and red around 620–750 nm:

- **Violet**: 380–450 nm,
- **Blue**: 450–495 nm,
- **Green**: 495–570 nm,
- **Yellow**: 570–590 nm,
- **Orange**: 590–620 nm,
- **Red**: 620–750 nm.

By contrast, composite colors arise when the spectrum contains many wavelengths with different relative intensities. White light is a classic example: it contains a broad range of visible wavelengths with approximately balanced power. Another instructive case is magenta, which does not correspond to any single wavelength; it is perceived when red and blue components reach the eye with relatively little green, and the visual system interprets this pattern as a distinct hue even though there is no "magenta wavelength" in the spectrum [9]. In this sense, a simple color is dominated by a narrow wavelength band, while a complex color reflects a mixture of wavelengths.

Color perception depends on the shape of the spectrum – that is, how radiation power is distributed as a function of wavelength. Two physically different spectra can produce the same perceived color if they stimulate the S, M and L cones in the same way; this effect is known as metamerism [10]. For practical discussion, spectra are sometimes grouped into narrowband (nearly monochromatic) sources such as lasers, sources with one dominant peak but noticeable spread such as many LEDs, and broad spectra such as sunlight or incandescent lamps, which cover a large part of the visible range:

- **monochromatic**: with a very narrow range (e.g., laser, monochromatic rays);

- **coherent but broad**: with one dominant maximum, but also with a wider range (e.g., LED light);
- **broad incoherent**: e.g., sunlight or incandescent light, which covers a large part of the spectrum.

We typically perceive a "simple" (spectral) color when the spectrum contains a clearly dominant component – that is, when power is concentrated in a narrow wavelength band and one region of the spectrum outweighs the rest. In such cases, cone responses are driven mainly by that narrowband stimulus; for example, light concentrated around about 625 nm tends to produce a strong response in the L cones and is commonly perceived as red. When no single band dominates and the spectrum stimulates S, M and L cones in a more balanced or multi-peaked way, the percept is no longer well described by a single wavelength: it may appear as a composite hue (produced by multiple spectral components) or, if cone responses are close to equal, as an achromatic sensation such as white, gray, or black.

In 1853, Hermann Grassmann proposed a set of empirical regularities describing how mixtures of lights are perceived. These principles – commonly referred to as Grassmann's laws – became a cornerstone of colorimetry, because they support an approximately linear, algebraic treatment of additive color mixing under photopic (cone-mediated) vision. In the literature, these laws are sometimes summarized and numbered in slightly different, but largely equivalent, ways. In this article, I use the term "Grassmann's laws" in a pragmatic, modeling-oriented sense: as a compact label for the empirical properties that are most relevant to the later parts of this paper – continuity of mixtures, proportionality, additivity, and metamerism.

First law *– trichromatic characterization (operational form). A perceived color can be specified, for a standard observer under fixed viewing conditions, by three independent quantities* (*tristimulus description*).

I interpret this as a practical statement that, for photopic vision, the visual system effectively reduces a full spectrum to three cone responses. As a result, different spectra can share the same three-component description and still look identical, which is the starting point for metamerism.

Second law – *proportionality* (*scaling*). *If the intensity of a light stimulus is scaled by a factor, its perceived color remains the same while its brightness scales accordingly* (*within the normal photopic regime*).

In practice, this means that changing the power of a single light source moves us along the same "direction" in color space: hue does not change, only the magnitude of the stimulus. This property is what makes it meaningful to treat intensity changes as scalar multiplication in an algebraic model.

Third law *– additivity of mixtures. The appearance of a mixture changes continuously with the proportions of its components, and the perceived result depends on the component contributions in an approximately additive way.*

This is the key justification for linear models of additive mixing: if we increase the contribution of one component while keeping others fixed, the perceived mixture shifts accordingly. Under this approximation, mixing behaves like vector addition, which later connects naturally to tristimulus coordinates and CIE-based spaces.

Fourth law *– metamerism invariance under addition. Physically different spectra can appear identical* (*metamers*). *Moreover, if two stimuli are metamers for a given observer, adding the same third stimulus to both preserves that match* (*and, in the same idealized framework, applying an equivalent filtering operation does not break the match*).

I treat this as a perceptual "equivalence relation" defined by cone responses: if two spectra produce the same S, M and L stimulation, then subsequent equal additive modifications keep them equivalent from the observer's perspective. This emphasizes that perceived color is determined by receptor responses rather than by the underlying spectral composition itself.

These ideas provided the conceptual basis for later experimental work by Helmholtz and Maxwell and, ultimately, for the establishment of the CIE 1931 standard observer and the corresponding CIE color spaces. In modern colorimetry, the practical consequence is that we can represent a light stimulus by three matching coefficients: the amounts of three fixed primaries required to match that stimulus for a standard observer. Under the additivity assumption, mixing two stimuli corresponds (approximately) to adding these coefficients component-wise.

$$(R,G,B) = (R_1 + R_2, G_1 + G_2, B_1 + B_2) \qquad (2)$$

4. Generating the visible spectrum using digital technology

The accuracy of digital color reproduction in the real world depends strongly on the display device. Different technologies generate light with different spectral power distributions, so the same RGB triplet can produce slightly different physical stimuli – and therefore slightly different appearance – across screens. Classic CRT (cathode ray tube) monitors excite phosphors that typically emit relatively broad spectra for each primary, whereas LCD monitors with LED backlighting rely on a backlight spectrum shaped by filters (or color-conversion layers), which often yields narrower spectral bands and may include secondary components outside the main band [11]. Comparative studies report that, under typical conditions, LCDs can cover a larger effective color gamut than CRTs. Sharma notes that this advantage is largely driven by higher achievable luminance in the LCD prototypes (brighter backlighting), while also pointing out a practical trade-off: LCDs can be more sensitive to viewing angle, with noticeable color shifts when observed off-axis [12]. Industry-oriented comparisons sometimes summarize gamut capability using the NTSC percentage convention. In such tables, OLED panels are often reported to exceed 100% NTSC, while many mainstream LCD configurations cluster closer to roughly three-quarters of NTSC coverage, and emerging micro-LED systems are sometimes reported above 100% as well [13]. Interpreted cautiously, this suggests that OLED and related approaches (for example, quantum-dot enhanced displays) can deliver more saturated, "vivid" colors than older or lower--gamut screens, although the exact numbers depend on measurement methodology and the specific panel implementation.

Devices also differ in effective color depth. As analog devices, CRTs can be viewed as capable of continuous tonal variation, but in practice their electronics, noise, and phosphor behavior limit stability and repeatability. Modern flat panels are digitally driven and most commonly operate at 8 bits per channel (about 16 million addressable RGB combinations), while 10-bit and 12-bit modes are available in higher-end workflows, including reference displays used in color-critical production. Higher bit depth reduces visible banding in smooth gradients and improves the reproduction of subtle detail in shadows and highlights. Perceived color on a display is also shaped by viewing conditions, especially screen brightness and ambient illumination. High ambient brightness can introduce reflections and glare that reduce perceived contrast and visibility [14]. More generally, ambient lighting influences both the physical tristimulus values reaching the eye and the observer's color adaptation state, which can shift the perceived appearance of the same on-screen content [15]. For consistent evaluation, calibration guidance typically recommends controlled, moderate ambient illumination and a screen white luminance chosen to match the viewing environment; a commonly cited target range in office-like conditions is on the order of ~80–120 cd/m². Finally, perceived saturation and brightness are not fully separable in human vision: due to the Helmholtz–Kohlrausch effect, highly saturated colors can appear brighter (and subjectively more intense) than an achromatic stimulus of the same measured luminance [17].

Although screen resolution primarily affects spatial detail, it can also influence color perception indirectly. Each pixel is composed of tiny R, G, and B subpixels arranged next to one another. If pixel density is low – or if the viewer sits very close – these subpixels may become visible, appearing as a colored grid, fringing along edges, or a so-called screen-door effect. At sufficiently high pixel densities, the visual system integrates the subpixel emissions into a single perceived color per pixel. Under typical viewing conditions, modern displays with roughly a few hundred PPI (often around 200–400 PPI or higher) make individual subpixels effectively indistinguishable, producing smooth images with visually uniform colors. To further improve perceived sharpness and reduce color artifacts near the visibility threshold of subpixel structure, manufacturers may use alternative subpixel layouts (e.g., Pentile in some OLED panels) and rendering techniques such as subpixel rendering, which intentionally leverages the subpixel geometry.

In summary, converting a digital color (an RGB code) into an "actual" physical color requires accounting for the display's emission characteristics (the spectra of its phosphors, LEDs, or color-conversion layers), its achievable gamut, its effective bit depth, and the viewing environment. The color-science and color--management literature describes, how devices are characterized and calibrated (e.g., with ICC profiles and colorimetric transforms) so that a given digital code corresponds as closely as possible to a target visual stimulus [11]. However, even with careful calibration, gamut limitations remain fundamental: a standard monitor cannot reproduce every physically realizable color stimulus. For example, a highly saturated laser green may lie outside the display's achievable chromaticity region, and colors outside the display's gamut cannot be produced by mixing its primaries. For most everyday use the resulting differences are negligible, but they can matter in color-critical workflows and in edge cases involving extremely saturated or unusual spectra.

5. Mathematical model of RGB mapping onto the visible spectrum

To formalize the relationship between a digital color code and the corresponding physical light stimulus, we introduce a simple mathematical model of a display. Let λ denote wavelength (in nanometers) within the visible range. We assume that the screen has three primary emitting components associated with its subpixels, characterized by dominant wavelengths:

- $\boldsymbol{\lambda_R}$: the dominant wavelength emitted by the red subpixel (e.g., approx. 610–630 nm for typical red LEDs);
- $\boldsymbol{\lambda_G}$: the dominant wavelength for the green subpixel (e.g., approx. 530–550 nm);
- $\boldsymbol{\lambda_B}$: the dominant wavelength for the blue subpixel (e.g., approx. 450–470 nm).

Each subpixel emits a spectrum with nonzero bandwidth: it is not a single spectral line but a band whose shape depends on the emitter and any optical filters or conversion layers. To keep the model tractable, we assume each primary emits in a relatively narrow band around its dominant wavelength λ_R, λ_G, λ_B. We represent these emissions by spectral power distribution functions $P_R(\lambda), P_G(\lambda), P_B(\lambda)$, which describe the relative radiant power distribution of the display primaries as a function of wavelength. In an idealized limit, these primary spectra can be approximated by Dirac delta-like distributions concentrated at λ_R, λ_G, λ_B. This approximation is not physically exact, but it provides a convenient baseline model for deriving a deterministic mapping from RGB values to a simplified wavelength-domain description.

A digital color in the RGB model is specified by three channel values (R, G, B), typically in the range 0–255. We normalize each channel by dividing by the maximum value (255) to obtain fractional intensities:

$$x = \frac{X}{255} \tag{3}$$

where X denotes one of the RGB channel values. This yields three relative coefficients $r, g, b \in [0,1]$. Under the display model introduced above, the spectral power distribution emitted by a pixel with code (R, G, B) can be approximated as:

$$S(\lambda) = I[rP_R(\lambda) + gP_G(\lambda) + bP_B(\lambda)] \tag{4}$$

where I is an overall brightness multiplier related to pixel luminance. In this linear model, the emitted spectrum is the sum of the three primary spectra scaled by r, g and b. This assumes additivity and ignores nonlinear effects such as saturation, channel crosstalk, and other device-dependent interactions, which are typically small in a well-characterized display operating within its normal range.

With this simplified RGB-to-spectrum model, we can ask when a stimulus is likely to be perceived as close to a spectral (narrowband) color, when it is perceived as a composite hue, and when it appears achromatic. However, once metamerism is considered, it becomes clear that Eq. (4) does not define a unique "spectrum of a color" in the perceptual sense. What matters for appearance is the resulting cone (or tristimulus) responses, and many different spectra – including spectra from different devices – can produce essentially the same perceived color. Digital color codes are therefore better understood as recipes for stimulating the visual system rather than as direct descriptions of physical wavelengths.

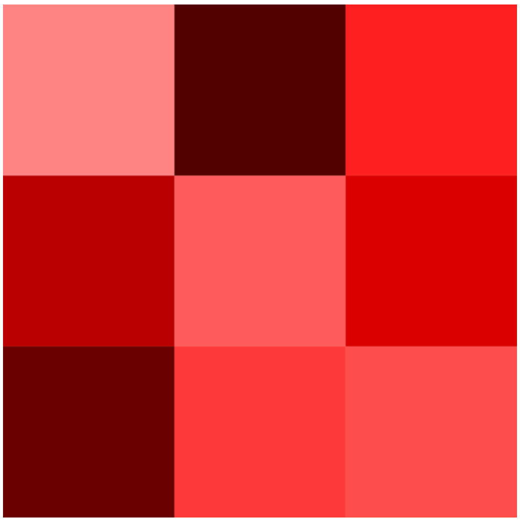

Fig. 2. Red hue
Source: https://pl.wikipedia.org/wiki/Barwa_czerwona#/media/Plik:Color_icon_red.svg

Changing the RGB components changes the cone excitations and thus the perceived color, consistent (as in the operational, trichromatic interpretation discussed earlier). At the same time, human color naming is categorical: even if two stimuli differ measurably in hue, observers may still label both with the same basic color term, especially when differences are small or context dependent. Using the monitor spectrum model in Eq. (4), we can introduce a simple dominance criterion for when an RGB-coded stimulus may be treated as approximately spectral in the sense of being driven mainly by one primary component. In practice, this requires one channel to be much larger than the other two:

$$x_1 \gg x_2 \ \wedge \ x_1 \gg x_3 \tag{5}$$

where x_1 is the dominant channel and x_2, x_3 are the suppressed channels. For example, for a red-dominated stimulus this condition becomes:

$$r \gg g \ \wedge \ r \gg b \tag{6}$$

This should be understood as a heuristic convenience rather than a statement about the true physiology of vision or the actual spectrum of the emitted light. As a further simplifying heuristic (not a one-to-one physiological mapping), we treat RGB channel dominance as a proxy for cone-weighted stimulation. We denote these approximate, cone-like responses as L, M and S, loosely corresponding to the red-, green-, and blue-associated components, respectively. Under this operational interpretation, a "red" percept corresponds to a stimulus, where the L-like response is present while the other components are negligible relative to it:

$$\frac{M}{L} \approx 0 \ \ \wedge \ \frac{S}{L} \approx 0 : L > 0 \tag{7}$$

6. Discussion and conclusions: why reconstructing a spectrum from RGB is not perceptually meaningful

In the previous sections, a simplified emission-based description of a display pixel was introduced (Eq. (4)), where a color code (r, g, b) is modeled as a weighted combination of the spectral components of the three primaries. Such a model is meaningful as a **device description** (e.g., for a specific LCD/OLED panel) and can be useful for analyzing gamut limitations, inter-display differences, or reproduction accuracy. However, two distinct objectives must be separated:

- a **physical** description of emission (the spectral power distribution $S(\lambda)$ produced by a device);
- a **perceptual** description of color (how the stimulus is encoded and interpreted by the visual system).

The main conclusion of this article is that attempting to reconstruct a "proper" or "true" spectrum for a color specified by an RGB code is, in general, a poorly posed goal in the context of color perception.

In human photopic vision, spectral information is largely reduced to the responses of three cone classes (approximately summarized as L, M and S). This means that the full distribution $S(\lambda)$ is not preserved in perception as a high-dimensional object; instead, it is compressed into a small set of receptor responses (followed by opponent processing). As a result, metamerism occurs: two different spectra may produce the same (or practically indistinguishable) color appearance if they generate similar receptor excitations.

Consequently, any mapping of the form "RGB to unique spectrum" is inherently non-unique: infinitely many spectra can be perceptually equivalent. Even if Eq. (4) provides a plausible emission spectrum for a particular display, the resulting $S(\lambda)$ is not "the spectrum of the color" in a general sense; it is merely one of many spectra that can yield a similar perceptual outcome.

Engineering practice often relies on a standard observer (e.g., in the CIE framework), which is highly useful for standardization and comparison. However, real observers differ in visual physiology (e.g., ocular media properties, macular pigment density, and photopigment parameters). Consequently:

- stimuli that are metamers for one person may not be metamers for another;
- an accurate spectral "match" does not ensure identical appearance across the population.

Therefore, even a radiometrically perfect spectrum reconstruction for an RGB code does not solve perceptual agreement, because perception is observer dependent.

Color appearance depends not only on the stimulus spectrum, but also on viewing conditions: luminance level, adaptation state, surround illumination, background, contrast, scene context, and color constancy mechanisms. In practice, the "same RGB" may yield different appearances depending on:

- ambient brightness and illumination characteristics;
- the observer's adaptation level;
- he relationship between the stimulus and its surround.

Under this perspective, pursuing "spectrum reconstruction" as an end is often less relevant than properly specifying viewing conditions and using colorimetric/appearance approaches (e.g., CIE-based representations and appearance models).

A further argument against treating the spectrum as an absolute determinant of color appearance is the large variability across species. The number of photoreceptor types and their spectral sensitivities differ significantly among animals.

Consequently:

- "color" is not a property of the physical stimulus alone;
- it is relational: stimulus + sensory system + neural processing.

A spectrum that produces a certain color sensation in humans may be categorized very differently by another visual system (e.g., species with additional receptors or UV sensitivity).

A spectral model (such as Eq. (4)) is appropriate when the goal is device emission description, radiometric/energetic analysis, spectral simulations, or spectrum-dependent computations (e.g., light–material interactions, sensor response prediction, multispectral pipelines). However, for typical imaging and graphics applications where human color appearance is the target, a more appropriate objective is:

- converting to a device-independent colorimetric description (e.g., XYZ/xy for a defined observer and conditions),
- and, if necessary, incorporating appearance modeling and explicit viewing conditions.

In summary, reconstructing a "visible spectrum" from RGB is generally not important for human color perception because (i) metamerism makes spectra non-unique, (ii) perception varies between individuals, and (iii) it varies even more across species. Therefore, rather than searching for a single "true spectrum" behind an RGB code, the problem should be formulated in terms of colorimetry and perception, with an explicitly defined observer and viewing conditions.

7. Color classification heuristics

Based on the discussion so far, a single, unambiguous "true" color label cannot be derived from an RGB code in a universal sense. The same RGB triplet may correspond to different physical spectra on different devices, different spectra can produce the same percept (metamerism), and perceived appearance depends on viewing conditions and observer-specific physiology. Even if we restrict ourselves to a healthy human observer under photopic vision, color perception remains context dependent and variable across individuals.

A further practical limitation is that the discreteness of RGB does not translate into a comparably rich vocabulary. RGB systems contain millions of distinct codes, but this does not mean that the average observer can reliably discriminate, remember, and consistently name each small difference with a unique term. In everyday language, color naming is categorical: many distinguishable stimuli are still grouped under the same basic label (such as "red" or "green"), and different observers may draw category boundaries differently.

The first heuristic proposed here is an author-defined geometric construction: the RGB color cube. The cube represents the entire discrete RGB space as a 3D lattice of size 256 x 256 x 256, where the R, G and B channel values are treated as orthogonal axes, analogous to x, y and z coordinates. Each digital color becomes a point in this cube, and operations such as subtraction, distance measurement, and region membership can be defined directly in this coordinate system. Once the core idea of the color cube is established, naming becomes a problem of partitioning this discrete volume into labeled regions. For each name in the chosen label set, we define at least one reference point – a prototype RGB triplet that serves as a representative "starting color" for that category. The category itself is then modeled as a neighborhood around the prototype, described by an admissible displacement within the cube.

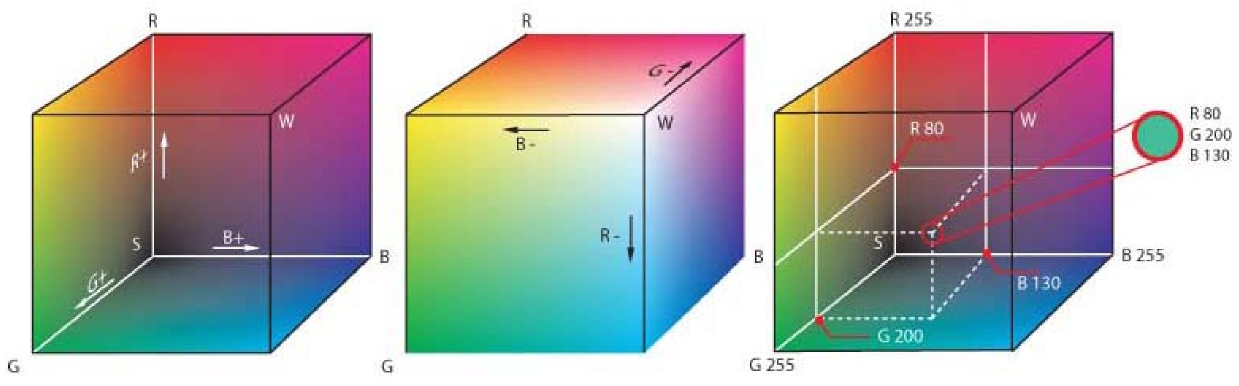

Fig. 3. RGB Cube
Source: https://upload.wikimedia.org/wikipedia/ commons/0/03/RGB_farbwuerfel.jpg

In the simplest version, this displacement can be expressed as a tolerance vector that determines, how far the category extends along each axis, which turns the category into a well-defined region in RGB space. Operationally, an RGB code belongs to a named category if its vector difference from the prototype lies within the category's allowed displacement region. If multiple categories overlap, the rule can be completed by a deterministic tiebreaker, such as choosing the nearest prototype under a chosen distance metric. The outcome is intentionally heuristic: it does not claim physiological uniqueness, but it provides a repeatable, explicit method for translating a large set of digital colors into a smaller set of human-usable names.

A second approach is to ground the mapping in quantitative user studies by sampling the RGB space. Since there are far too many RGB codes to evaluate each one experimentally, we select a subset of representative colors (for example, by using a regular grid step or stratified sampling across the cube) and collect human judgments for that subset. Each sampled RGB point is then associated with a name chosen by observers under controlled viewing conditions, producing an empirical dataset of (RGB –> label) pairs. Once such a dataset exists, a simple and practical heuristic is the "nearest labeled neighbor" rule. For any arbitrary RGB code, we assign the name of the closest sampled point that already has an experimental label. In geometric terms, the sampled points act as anchors inside the RGB cube, and every unsampled color inherits the label of the anchor it is closest to under a chosen distance metric. This produces a deterministic classifier that can be refined by increasing the sampling density in regions where observers show higher sensitivity or where naming boundaries are uncertain.

A third heuristic builds directly on the dataset from the second approach, but replaces hand-crafted neighborhood rules with data-driven classification models. Here, the experimentally labeled samples serve as training data, and the mapping from RGB to color names is learned by an algorithm. Depending on the goal, this can be formulated either as supervised classification (predicting one of the predefined names) or as unsupervised clustering (discovering groupings that are later assigned names). In the supervised setting, suitable models include k-nearest neighbors (k-NN) as a direct generalization of the nearest-neighbor heuristic, as well as decision trees and random forests (which produce interpretable, rule-like partitions of the cube), support vector machines (effective for high-dimensional boundaries, even though RGB is only 3D), and simple neural classifiers. In the unsupervised setting, clustering methods such as k-means can partition the RGB space into a chosen number of groups, while Gaussian mixture models can capture elliptical clusters and provide probabilistic assignments. Density-based methods such as DBSCAN are also useful when clusters are irregularly shaped or when "outlier" colors should be treated separately. After clustering, each cluster can be assigned a name by linking it to the majority label in the nearest set of human-annotated samples, or by selecting representative prototypes and labeling them manually.

8. Bibliography

[1] Maussion F., "Numbers and their representation in computers", Scientific Programming (SciPro WS2021), https://fabienmaussion.info/scipro_ws2021/week_04/02-Numbers.html (online: December 2025).

[2] Elsevier, "Color component – an overview", ScienceDirect Topics, https://www.sciencedirect.com/topics/computer-science/color-component (online: December 2025).

[3] Elsevier, "Truecolor image – an overview", ScienceDirect Topics, https://www.sciencedirect.com/topics/computer-science/truecolor-image (online: December 2025).

[4] Cambridge in Colour, "Bit Depth Tutorial", Cambridge in Colour, https://www.cambridgeincolour.com/tutorials/bit-depth.htm (online: December 2025).

[5] Tyson J., Carmack C., "How Computer Monitors Work: Color Depth", HowStuffWorks, 2000, https://computer.howstuffworks.com/monitor4.htm (online: December 2025).

[6] CIE, "Fundamentals of Color and Vision", International Commission on Illumination.

[7] Robertson A.R., "The CIE 1976 Color-Difference Formulae", *Color Res. Appl.*, vol. 2, no. 1 (1977).

[8] Thornton W.A., "Spectral Distribution of Typical Indoor Electric Light Sources", *J. Illum. Eng. Soc.* (1974).

[9] Hunt R.W.G., *The Reproduction of Colour*, Wiley, 2004.
[10] Fairchild M.D., *Color Appearance Models*, 3rd ed., Wiley, 2013.
[11] Clark R.N., "Color Part 2: Color Spaces and Color Perception", ClarkVision.com, May 23, 2019 (last updated Dec. 20, 2020), https://clarkvision.com/articles/color-spaces/ (online: December 2025).
[12] Sharma G., "LCDs Versus CRTs–Color--Calibration and Gamut Considerations," *Proc. IEEE*, vol. 90, no. 4, pp. 605–622, https://hajim.rochester.edu/ece/sites/gsharma/papers/lcdvscrtprocieee.pdf (online: December 2025).
[13] Wu Y., et al., "Full-Color Realization of Micro-LED Displays", *Nanomaterials*, vol. 10, no. 12, art. 2482 (2020), doi: 10.3390/nano10122482.
[14] Rodán A., et al., "Too Bright to Focus? Influence of Brightness Illusions and Ambient Light Levels on the Dynamics of Ocular Accommodation", *Vision* (Basel), vol. 9, no. 4, art. 81 (2025), doi: 10.3390/vision9040081.
[15] Baek Y., et al., "The Effect of Ambient Lighting on the Preferred Color Temperature of Television", *Proc. IS&T Int'l Symp. on Electronic Imaging: Color Imaging XXVI: Displaying, Processing, Hardcopy, and Applications*, 2021, doi: 10.2352/ISSN.2470-1173.2021.16.COLOR-327.
[16] Wikipedia contributors, "Grassmann's laws (color science)", Wikipedia, the free encyclopedia. Last edited Dec. 23, 2025, https://en.wikipedia.org/wiki/Grassmann%27s_laws_%28color_science%29 (online: December 2025).
[17] Fernandez-Maloigne C. (Ed.), *Advanced Color Image Processing and Analysis*, Springer, 2013, doi: 10.1007/978-1-4419--6190-7.

Modelowanie dyskretnej przestrzeni kolorów cyfrowych RGB do widzialnego spektrum barw

J. GRĄTKIEWICZ

Autor stawia pytanie, w jaki sposób funkcjonuje świat kolorów: jak ciągłe, fizyczne widmo promieniowania staje się subiektywnym wrażeniem barwy. Następnie konfrontuje ten opis z praktyka cyfrowa, w której kolory są kodowane dyskretnie jako RGB i odtwarzane przez ekrany emitujące sztuczne światło o urządzeniowo zależnych widmach. Artykuł analizuje, czy możliwa jest konwersja kolorów cyfrowych do przestrzeni fizycznej (widmowej) oraz czy taka konwersja ma realne znaczenie. Opierając się na metameriach, różnica między urządzeniami, w warunkach obserwacji i zmienności obserwatorów, autor dochodzi do wniosku, że przypisanie RGB do „jednego prawdziwego widma" jest w praktyce niejednoznaczne i słabo ugruntowane percepcyjnie. Kluczowym wnioskiem jest też to, że człowiek nie rozumie nawet skończonej przestrzeni RGB: nie wszystkie różnice są rozróżnialne, a nawet gdy są, nie przekłada się to na jednoznaczne i konsekwentne nazewnictwo. Tym samym problem nieskończonego spektrum zostaje uproszczony do skończonego zbioru kodów RGB, a następnie do jeszcze mniejszego zbioru nazw barw, którymi posługujemy się w języku. Na tej podstawie autor proponuje heurystyki klasyfikacji, które są „wystarczająco dobre" do praktycznego użycia terminologii barw: (1) autorska metoda kostki RGB i podziału przestrzeni na regiony nazw, (2) podejście empiryczne oparte na próbkowaniu i regule najbliższego oznaczonego sąsiada oraz (3) klasyfikacje bazujące na danych, w tym metody uczenia maszynowego i klasteryzacji.

Słowa kluczowe: grafika komputerowa, kolory a barwy, ML.