On the cardinality of color stimulus properties

Peter Morovič, Ján Morovič, HP Inc.

Abstract

When describing "how many colors" can be represented or reproduced by a given system or device, volume (or gamut volume) is often used. While it has a rightful place and is a valid way to describe a range of colors, it is an incomplete and indirect answer to the original "how many?" question. In this paper a first principles based approach is applied to this question, starting with delimiting the validity of any possible answer by being explicit about assumptions about domain (dimensionality, encoding bit-depth, coordinate system or color space. etc.). Furthermore, no real-world data exists without a reference to a measurement device, which in turn has to consider sources of noise: physical surfaces whose color or spectral measurements differ by less than the level of noise of that device, cannot be considered as distinct. These assumptions directly affect possible answers to the initial question of how many colors or color stimulus properties there are. Starting from reflectances, their quantization and measurement noise as well as colorimetry are analysed in a variety of color spaces and the concepts of cardinality and colorimetric errors are used. The result is not a single answer but both an exploration of the effects of the assumptions and a characterization of their dependence. For example, results show how quantization and measurement error can have a significantly larger impact than may be intuited, likewise, the same analysis in CIE LAB, CIE CAM16 and CAM16-UCS yields substantially different answers (and may be applicable in different contexts), which also highlights current limitations of these spaces and color difference metrics. Additionally, a cardinality analysis can also be performed in a biological domain of retinal responses, which bridges the physical and psychophysical domains. As will be shown, precise answers can be given under specific and explicit assumptions, but in a general context, the answer always has to be "it depends".

Introduction

Questions about how many somethings there are have occupied the minds of humans since prehistoric times. From counting lunar phases 42 000 years ago in South Africa's Lebombo mountains by carving notches into bones [1], via the elaborate accounting records of ancient Egyptians [2], to Peter Drucker's 20th century mantra that, "if you can't measure it, you can't manage it," questions about quantity have been constitutive of the human condition.

Looking at a narrow slice of this question, the present paper will ask how many "somethings" there are that relate to color. The obvious thing here is to jump straight to asking: "how many colors are there?" – and the paper will eventually arrive at it, but it may prove illuminating to work towards it instead.

As has been shown previously by Morovic *et al.* [3], the question of how many colors there are is elusive, both because of the myriad combinations of objects, viewing and lighting environments and observers and observer states, and because of the constraints of the ground truth psychophysical data on which colorimetry and color appearance models are based.

Such constraints notwithstanding, there are broader questions to ask that were not considered previously.

Before proceeding further, it is also worth spelling out the implicit qualifier that the question about how many colors there are contains, which is that these colors need to be distinct from each other. In other words, a set of colors contains only distinct colors if for each color in the set there is no other color that looks the same.

Starting from colors, or more specifically the color attributes of visual percepts (where percepts are the result of the visual perception of stimuli), it can be asked how best to quantify their number. The question can further be extended to the source of those percepts, to being about quantities of biological states in the human visual system and then quantities of stimuli, i.e., that which the human visual system perceives. It then becomes: "how many stimuli are there that result in color percepts?" or "how many biological states are there in the human visual pathway that result in color percepts?" This perspective also allows for an umbrella question to be formed that combines stimuli, biological states and percepts, since percepts can be thought of also as properties of biological states and stimuli in the sense that they correspond to them in function of an observer. The end result then becomes an enquiry into color stimulus properties – both direct and observer-mediated.

Furthermore, these questions ought to be asked not only in some abstract, continuous, theoretical domain, but in the context of physical systems that are prone to variation and noise and where quantities about them are established via measurement, which is itself subject to those same fluctuations and whose results are also expressed with finite precision.

Following some groundwork on cardinality, which will let us ask not only questions about "how many" somethings there are but also about what kind of infinity applies in continuous cases, the paper will proceed from stimulus, via biological state towards percept, and explore the impact both of domain and of physical constraints on the question of how many distinct color stimulus properties there are.

Cardinality

Cardinality is defined as the measure of the size of a set or the number of elements in a set. While it traces its roots to ancient Babylonian and Egyptian mathematics where systems of enumeration were used to count quantities of objects, it was Euclid who introduced the concept of correspondence, where the size of a set is expressed as the number of elements that can be put in correspondence with another set. More recently, Georg Cantor extended this theory to countably and uncountably infinite sets and formalized it into what today is commonly used in set theory to talk about cardinality [24].

For finite countable sets, cardinality simply refers to the number of elements in that set. For example, the set of A = {0, 1, 2, 3} is a countably finite set whose cardinality, denoted as |A|, is 4 as it has 4 elements. Another way to think about this is that this set has a cardinality of 4 because its members can be paired up (put in one-to-one correspondence) with the members of any other set that also has the same cardinality of

4. For sets that have infinite elements, cardinality is also expressed in terms of correspondence and Cantor defines some classes of infinities, such as $\aleph_0$ (aleph-0) representing the cardinality of all natural numbers $|\mathbb{N}|$, also referred to as countably infinite sets. Any other set for which a correspondence (a one-to-one mapping) between the natural numbers $\mathbb{N}$ and itself exists, also has cardinality $\aleph_0$ (such as the set of all integers, the set of all multiples of 42, the set of all rational numbers, etc.). Sets of numbers such as the real numbers $\mathbb{R}$ or irrational numbers $\mathbb{P}$ then have a higher cardinality since there is no one-to-one mapping to $\mathbb{N}$, this cardinality is $2^{\aleph_0}$ denoted $\aleph_1$.

These uncountably infinite sets furthermore break some of the intuitions from countably finite sets, so that the cardinality of the real interval [0, 1] is the same as that of the entire set of real numbers $\mathbb{R}$, or even its higher dimensional Euclidean spaces, by virtue of a one-to-one mapping from [0, 1] to $\mathbb{R}$ to $\mathbb{R}^n$ (see Cantors' diagonalization proof [24]).

While the concept of cardinality formalizes the notion of "how many" of something there are in a set, another complementary measure to use is that of size or volume. For example, the volume of the [0, 1] hypercube in arbitrary dimensions is always 1 (the length of a segment of [0, 1] is 1, the area of a square of side 1 is 1, the volume of a cube of side 1 is 1, etc.), while its cardinality is the same as that of $\mathbb{R}$, namely $\aleph_1$. Instead, for countably finite sets this relationship breaks. E.g., for a quantization in the aforementioned hypercube (e.g., instead of continuous over $\mathbb{R}$, 8-bit over integers), the set of all vectors will be a finite number of elements that directly depends on the quantization, while still having a volume of 1.
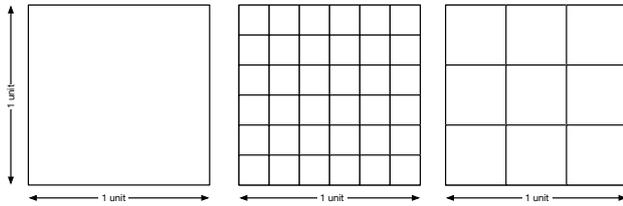
Figure 1. *A square of side one at three quantization levels – volume remains the same (=1) but the cardinality changes (from uncountably infinite to 49 to 16).*

Fig. 1 shows such an example where the three squares all have volume (area) 1.0, but the first square (assuming a representation in $\mathbb{R}^2$) has cardinality $\aleph_1$ (same as $\mathbb{R}$) while the second has cardinality 49 (there are 49 unique coordinates) and the last cardinality 16 (there are 16 unique coordinates).

Spectral domain

Moving on to the domain of color stimulus properties, these can be described in physical terms in the spectral domain by means of spectral reflectances or spectral power distributions, which give rise to color percepts. Without loss of generality, the spectral domain of Lambertian surfaces will be used here, assuming reflectances are constrained between 0% and 100% – no less than no light and no more than all light is reflected – (or the [0.0, 1.0] interval).

In theory, reflectances can be considered as arbitrary functions defined over a continuous sub-interval of $\mathbb{R}$ (e.g. some range of visible wavelengths) whose image, the values of the functions, also fall into a continuous sub-interval of $\mathbb{R}$ (e.g. [0, 1] for Lambertian surfaces). The cardinality of this

case of all possible continuous reflectances is $2^{\aleph_1}$ (often referred to as $\aleph_2$).

In practice, reflectances are represented in some higher-dimensional discrete space such as e.g. d=8, 16, 24 or 31 spectral samples on the 400nm to 700nm interval (40, 20, 12.5 or 10nm steps). This domain, also referred to as the Object Color Solid (OCS) [26], can therefore be considered a hypercube in n dimensions of side 1.0 and as mentioned earlier, has a volume of 1.0 and cardinality of $\aleph_1$, independent of n.

However, to answer the question of how many spectral reflectances there are, continuity is never the case in the real world, since all reflectances are the result of physical object properties measured using devices that operate at a particular spectral sampling and represent reflectances at a particular bit depth, even postponing the topic of variability and noise for a moment.

So, for a given dimension n, the volume of a hypercube (the set of possible reflectance stimuli) S, as mentioned before is defined as:

$$Vol(S) = r^n \tag{1}$$

And for $r = 1.0$ (Lambertian surfaces) Vol is equal to 1.0, and is independent of n (the dimensionality of the spectral samples). Introducing quantization, let smp be the number of unique samples that can be encoded per dimension, then the cardinality of unique stimuli at a given bit-depth can be written as:

$$|S| = smp^n \tag{2}$$

Where smp in the case of considering bit-depth can be expressed as 2^b, where b is the number of bits, such that for an 8-bit representation 256 samples per dimension can be represented and the cardinality of an 8-bit encoded 31 sampled spectral domain becomes:

$$|S| = 2^{b^n} = 2^{8^{31}} = 4.52e74 \tag{3}$$

While this seems an unimaginably large number, it needs to be compared with the cardinality of the continuous case ($\aleph_1$) which is uncountably infinite. Furthermore, bit-depth (at a fixed sampling representation) is exponentially related to the number of unique reflectances that can be represented, as Tab. 1 shows (with highlighted 8, 10, 12 and 16bit representations), also showing the orders of magnitude jump from 8bits (e74) to 10 (e93), 12 (e111), 16 (e149) and beyond: note that there are 10^{75} times more unique reflectances in a 16-bit than in an 8-bit encoding.

Table 1: Unique reflectances as a function of bit-depth (in a 31D spectral sampling domain)

Bits	Samples / dim	# of unique reflectances
8	**256**	**4.52E+74**
9	512	9.71E+83
10	**1024**	**2.09E+93**
11	2048	4.48E+102
12	**4096**	**9.62E+111**
13	8192	2.07E+121
14	16384	4.44E+130
15	32768	9.53E+139
16	**65536**	**2.05E+149**
17	131072	4.39E+158

18	262144	9.43E+167
19	524288	2.03E+177
20	1048576	4.35E+186
21	2097152	9.34E+195
22	4194304	2.01E+205
23	8388608	4.31E+214
24	16777216	9.25E+223

These results already transmit a sense of scale however they are not intuitive when it comes to their effect on color measurement. So, to express the impact of bit-depth error, a random 100K reflectances were generated and quantization error evaluated for the same bit-depth as in Tab. 1 but in terms of ΔE2000. The approach here is a worst-case analysis where uniformly distributed random noise (with $\mu = 0$ and range of -0.5 to 0.5) scaled by the magnitude of quantization (e.g., $\frac{1}{2}^{16}$ in the case of 16-bits) is added to the reference reflectances and compared colorimetrically against the reference reflectance. Fig. 2 below shows these results and while their magnitude is relatively small, as expected, the errors are non-negligible (even without taking repeatability or noise into account yet). At an 8-bit representation the median quantization error can be over 0.1 ΔE2000 and can go as high as 0.35. At 16 bits the errors are significantly lower, with a median below 0.05 and a maximum of 0.15.
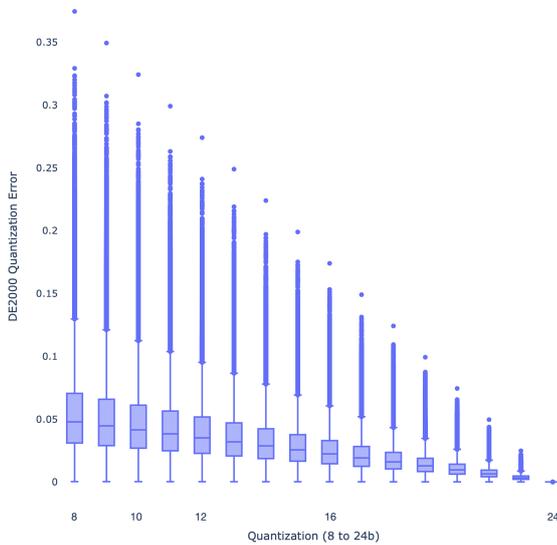
Figure 2. *Number of bits per spectral dimension (31) vs the ΔE2000 statistics of quantization error, over a set of 100K random reflectances.*

Sphere packing

While the analysis so far has dealt with continuous domains or with their per-dimension-quantised representations, incorporating the effect of measurement noise requires a different kind of approach.

A way to make the context explicit, in which a cardinality question is asked here, is the following. Given a volume in a space (e.g., a reflectance space of certain dimensionality) and a distinguishability threshold (e.g., a distance below which differences cannot be considered meaningful), what is the cardinality of the largest set of samples that can be distributed across that volume such that the difference between any pair

of samples is greater than or equal to the distinguishability threshold.

More formally, let S be a point in volume V, T be a distinguishability threshold and $d(S_1, S_2)$ be the minimum distance between any two points S_1 and S_2. Then, the cardinality m of the largest set A of points S such that the distance between each S and all the other Ss is greater than or equal to T can be written as:

$$m = \max\{|A| : \forall S_1, S_2 \in A, S_1 \neq S_2 \rightarrow d(S_1, S_2) \geq T\} \quad (4)$$

Finding the set A that maximizes m then corresponds to the n-sphere-packing problem in n dimensions [4]. In 2D an n-sphere is a circle and finding the m-maximizing set A is the same as finding the set of circles of radius T/2 that fills an area without overlapping, where all circle centers are inside that area and where the cardinality of A is maximal.

In 2D the answer to the m-maximizing arrangement of circles is having their centers on a grid of regular hexagons of side T (Fig. 3), whose optimality was proved by Lagrange in 1773 [5].
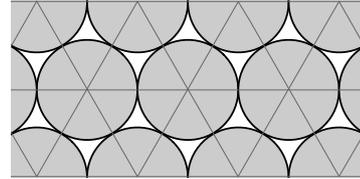
Figure 3. *Optimal circle packing, showing regular hexagonal grid formed by their centers as gray lines.*

A characteristic feature of n-sphere packings is their packing density, which refers to the proportion of a space enclosed by a set of packing n-spheres. In 2D, for circles, the highest packing density is $\pi/\sqrt{12}$, which means that an optimal circle packing covers 90.69% of the plane.

For the purposes of the analysis in this paper, it is useful to derive the number of packing n-spheres per unit volume from packing density, which can be obtained as follows:

$$C_U = pd/B \quad (5)$$

where C_U is the number of packing n-spheres per unit volume, pd is packing density and B is the volume of the packing n-balls (an n-ball being the space enclosed by an (n-1)-sphere, e.g., a circle, being a 1-sphere, encloses a 2-ball with a certain volume), computed as follows [6]:

$$B = \frac{\pi^{\frac{n}{2}}}{\Gamma\left(\frac{n}{2}+1\right)} r^n \quad (6)$$

where r is the n-ball's radius, n is the number of dimensions and Γ is Euler's gamma function. This is the case because packing density is the proportion of a unit space that is occupied by a certain packing and dividing it by packing n-ball volume gives the number of n-spheres that it takes to occupy that unit space proportion.

The number of packing spheres C at a given packing density within a certain volume V then is:

$$C = V \times C_U \quad (7)$$

During the following analysis, it would be desirable to compute the cardinality of packing sphere sets in the typical reflectance spaces that are 16 or 31 dimensional (for 20 and 10 nm samplings of the 400-700 nm range) and in 3 dimensions for colorimetry and color appearance. However, since optimal

packing densities for 16 and 31 dimensions are unknown, 8 and 24 dimensions will be used instead as a way to bracket the 16D case. Tab. 2 therefore shows the optimal packing densities in 3 [7], 8 [8] and 24 [9] dimensions alongside n-ball volume and the n-balls per unit volume factor obtained from them that will be used later.

Table 2: Optimal packing densities and n-ball volumes

Dimen–sions (n)	Optimal packing density (pd)	n-ball volume (B)	n-balls / unit volume (C_U)
3	$\frac{\pi}{3\sqrt{2}} = 74.05\%$	$\frac{4\pi}{3}r^3$	$\frac{1}{4\sqrt{2}r^3}$
8	$\frac{\pi^4}{384} = 25.37\%$	$\frac{\pi^4}{4!}r^8$	$\frac{4!}{384r^8}$
24	$\frac{\pi^{12}}{12!} = 0.19\%$	$\frac{\pi^{12}}{12!}r^{24}$	$\frac{1}{r^{24}}$

Measurement noise

Any physical measurement device has some level of noise due to environmental changes, current fluctuations, sensor variability, etc. [28,29] As a result, measuring a surface (in the same location under nominally the same conditions) multiple times, results in measured reflectances that vary numerically. If a device measures the same surface as different reflectances, consequently, different physical surfaces whose measurements vary to the same degree as this measurement noise cannot be distinguished and cannot be considered as different.

Hence, at a given level of noise, it is not possible to distinguish whether two measurements correspond to the same surface or to different surfaces. While quantization could be analyzed by counting the number of unique stimuli that can be represented at a given bit-depth, when it comes to measurement noise, the sphere-packing approach described above needs to be employed to determine cardinality.

The level of spectral repeatability noise is not often provided by manufacturers; however, to have a sense of its magnitude, it can be determined experimentally from multiple measurements of the same surfaces.

Here, a Minolta FD-9 spectrophotometer was used, taking 4 measurements of the same chart (under the same conditions and measured consecutively without moving the chart or changing any of the measurement conditions). Computing the pairwise differences among the 4 measurements and taking the per-wavelength maximum and minimum differences of all measured surfaces delimits a noise envelope interval of -1.9 to +2.9% in the spectral domain, as shown in Fig. 4. What is noteworthy is that these errors are an order of magnitude larger than those of, e.g., 8-bit quantization (1/256 = 0.39%).

As mentioned, the optimal solution of the sphere-packing problem is only known for certain dimensions, hence the analysis below will use 8 and 24 dimensions as the spectral representation. While 8 is unusually low for spectral data, 24 is already reasonable as some spectrophotometers measure 16 spectral samples (400nm to 700nm with 20nm intervals).

Using noise values from ½^16 (a 16-bit representation without other noise) to 0.029 (maximum spectral noise from the Minolta FD-9 test), the number of n-balls – i.e. uniquely distinguishable reflectances in the spectral hypercube's volume of 1.0 – is shown in Fig. 5 both in 8D (blue dots) and 24D (green dots).
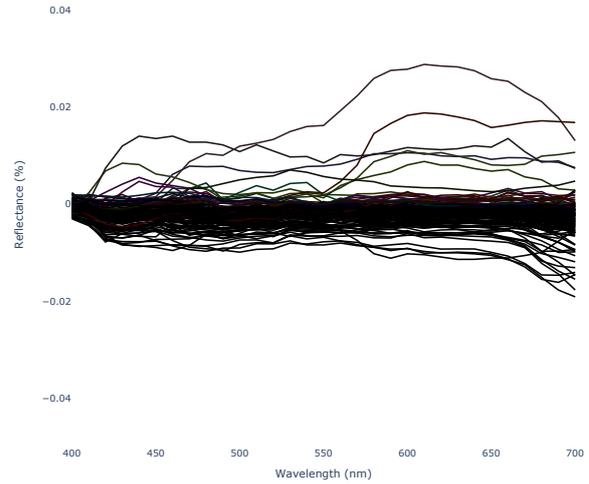
Figure 4. Spectral noise based on 4 repeated measurements of the same physical chart using a Minolta FD-9 spectrophotometer shown as the maximum spectral difference for each sample of the four measurements.
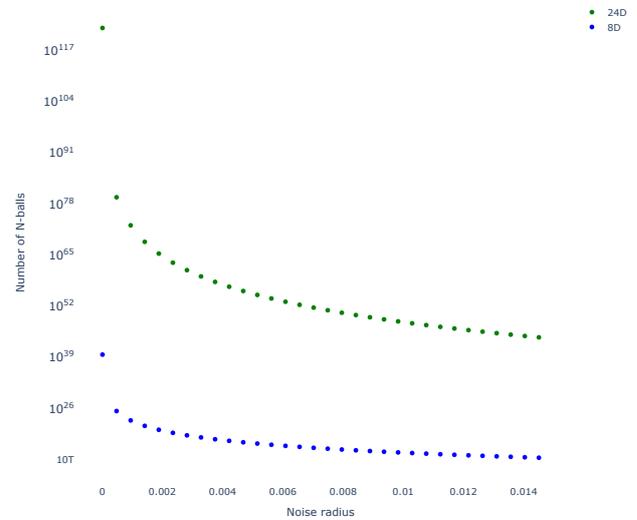
Figure 5. Spectral noise based on 4 repeated measurements of the same physical chart using a Minolta FD-9 spectrophotometer and corresponding cardinality.

Computing cardinality in terms of n-balls differs from the earlier, simpler analysis of quantization where, implicitly, hypercube (n-cube) subsampling was used instead of n-balls. In considering noise, using n-balls is more representative since for any point in reflectance space, a hyper-sphere encloses all possible deviations from the reference value in all directions. A hypercube (of the same size as the radius of an n-sphere) introduces non-uniformity as is evident also in 3D – e.g., a point at the center of a cube varies in its distance to the cube surface, while a point at the center of a sphere is equidistant to all points on the sphere.

The values in Fig. 5 show the relative impact of noise on the number of n-balls in the reflectance hypercube (shown at log-scale). As can be seen this relationship too is exponential and even apparently small levels of noise have a dramatic effect. As before, these values are not immediately relatable to perceptual magnitudes. Fig. 6 therefore shows a similar analysis

to Fig. 2, in terms of color differences due to measurement noise. In this case, the 53,489-member SOCS set [25] of measured reflectances is used, represented at 31 spectral samples and evaluated using the same levels of noise as in Figure 5, from $\frac{1}{2}^{16}$ to 0.029.

In each case a noise-scaled uniform random error (as before, $\mu = 0$ and range of -0.5 to 0.5) was added to the reference reflectance set and its $\Delta E2000$ color difference evaluated vs the original ("noise free") reference. As can be seen here too, even apparently low levels of spectral noise (less than 3%), can result in substantial colorimetric errors with a median of 0.19 and a maximum of 6.5 $\Delta E2000$.
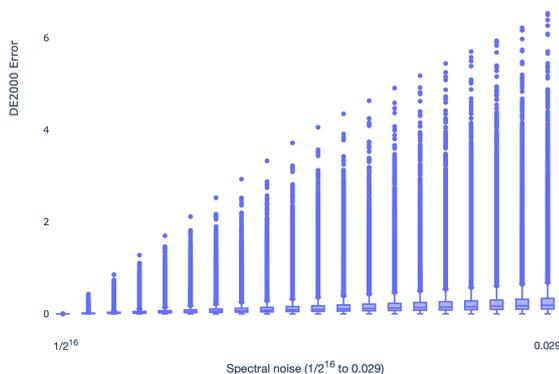
Figure 6. *Spectral noise interpreted in $\Delta E2000$ color difference terms shown as box-plots for each level of noise.*
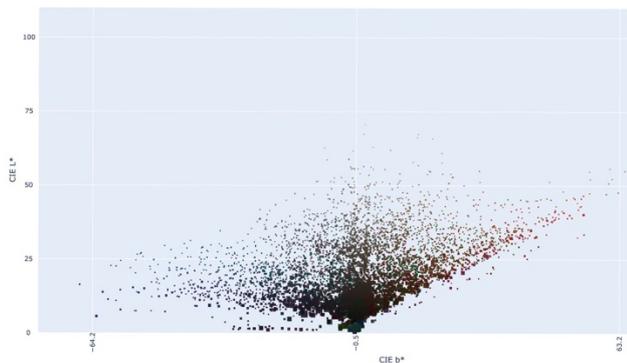
Figure 7. *Spectral noise interpreted in $\Delta E2000$ color difference terms shown as CIE LAB plot (L*b* projection) to illustrate the distribution of largest errors (above 1.0 $\Delta E2000$) – the values correspond to the errors analysed for a 0.029 per wavelength noise level (i.e. largest error).*

The distributions shown in Fig. 6 clearly show a long tail (outliers, beyond the inter-quartile range indicated by the box and even outside the 'whiskers' that correspond to 1.5 times the inter-quartile range) and Fig. 7 plots the color distribution of the largest errors in CIE LAB space. As can be seen, these outliers (where the size of the square markers indicates the magnitude of the error – scaled for visualization purposes) lie at the bottom of the gamut, meaning they correspond to the darker samples, where noise can have the biggest impact.

There are several limitations to the above analysis. Firstly, the measurement noise is considered to be the same at each wavelength, which need not be the case and even in the limited experimental data shown in Fig. 4, it appears to vary and result in higher levels of noise at the longer wavelengths vs the shorter ones. On the other hand, the errors shown in Fig. 6 also

don't constitute a worst-case scenario but a random sample. Further analysis with a wavelength-dependent model and more extensive noise modelling can be done in the future.

Biological domain

The stimuli whose cardinalities were considered so far are the source of the human visual system's response, triggering signals that travel from the retina to the brain's visual cortex. The question of how many "somethings" there are can therefore be also asked about the cardinality of biological states that are possible in a human observer and that then result a variety of color experiences.

A first attempt at quantifying them can be made by considering that the visual system has three color channels, in terms of cone responses and then cone-opponent signals that travel from the retina's ganglion cells along the optic nerve. Since the cone-opponent signals are the result of additive and substrative operations on the output of short- (S), middle- (M) and long-wavelength (L) sensitive cones, a cardinality analysis can be applied either directly in cone space or in cone-opponent space, yielding the same results.

Studies of the visual system's adaptation behavior [10–12] show approximately 1000 distinct signal levels for each of the ON and OFF ganglion cell types, which corresponds to a combined 2000 signal states for each of the three cone-opponent channel types, leading to a total of $2000^3 = 8 \times 10^9$ signal states leaving retinal ganglion cells.

While this does express the number of possible, distinct signal combinations at 2000 levels in a three-channel system, it would be an overestimate of biological states leaving retinal ganglion cells, due to the overlapping spectral responsivities of the three cone types. Taking the reflectances of the OCS, computing LMS cone responses [27] and from there cone-opponent values in L-M, (L+M)-S and L+M+S channels yields a convex volume that can then be intersected with a 8×10^9 grid spanning the full opponent space. The result are 2.09×10^9 possible, distinct cone-opponent signals (Fig. 8).
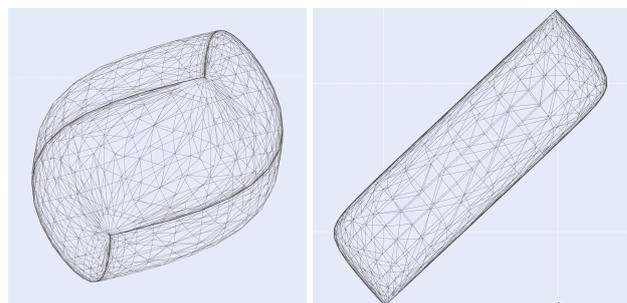
Figure 8. *(L+M)-S versus L-M (left) and L+M+S versus (L+M)-S (right) projections of the Object Color Solid convex hull under D65 in cone-opponent space.*

Colorimetric and color appearance domain

The question of cardinality can also be asked in a colorimetric domain in terms of sets of samples taken from the OCS gamut. Unlike in the spectral case, where sampling from a unit n-cube (i.e., n-dimensional hypercube) could be handled directly, sampling from a color gamut requires additional considerations when done in color spaces that are non-linear transformations of reflectance and CIE XYZ.

Since transforming the OCS into CIE LAB or CIE CAM16 does not preserve convexity, its gamut boundary will be determined using alpha shapes [13] with an ISO–recommended radius of 40 [14]. Because the result is a triangulated surface delimiting the OCS gamut, the question of what the largest number of spheres is that it can contain is distinct from how many spheres can pack a certain volume in 3D.

Addressing that question directly would involve computationally optimizing the placement of spheres inside the gamut boundary, which is an expensive optimization problem [15] and which yields lower packing densities (of around 65%) than those of regular lattices that in 3D can reach the theoretical optimum of $\pi/(3\sqrt{2})$=74%.

The approach taken here will be to first characterize how the positioning of a regular lattice versus a given boundary in 3D determines the number of enclosed lattice vertices and to then use the resulting statistics when approximating the maximal cardinality of enclosed color sets. This is done by taking a hexagonal close packing (HCP) lattice [4] (Fig. 2) that optimally packs in 3D and counting how many of its vertices are enclosed per unit volume in boundaries of increasing volume (Fig. 9).
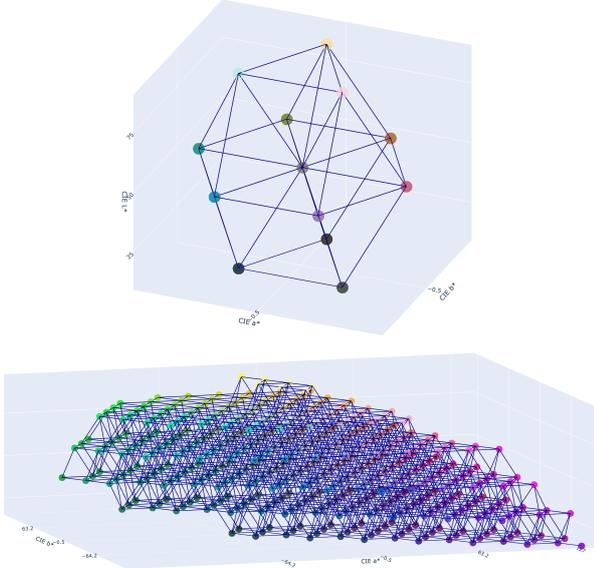
Figure 9. (top) A HCP cell and (bottom) the OCS under D65 in CIE LAB filled with a coarse HCP grid.
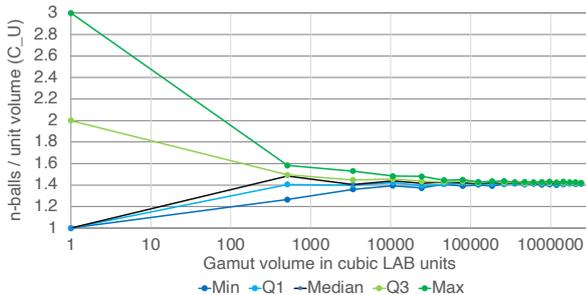
Figure 10. HCP n-balls per unit volume range as a function of gamut volume.

Instead of considering a single value per gamut volume, Fig. 10 shows a range, which is the result of taking the HCP grid and moving it relative to a gamut boundary in all three

dimensions in four steps over a one unit range, resulting in a jiggling that corresponds to the span of one packing sphere. As can be seen, at very low gamut volumes (below around 10K cubic LAB units) there is considerable variation as a result of the relative positioning of the grid versus the gamut. However, at gamut volumes above around 250K the full range of values is within 1-2% of the minimum C_U, with interquartile ranges below 1%. For the purposes of the analysis here, the *n*-balls per unit volume count of HCP in 3D will therefore be used to obtain sample counts from volume.

OCS cardinality in CIE LAB

As with reflectances, here too the effect of measurement noise can be taken into account when quantifying the number of distinguishable color coordinates that can be distributed across the OCS' volume of 2,267,996 cubic CIE LAB units (Fig. 11).
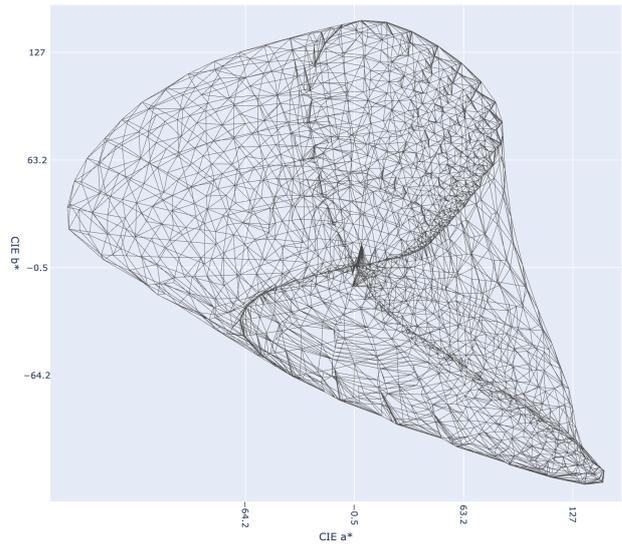
Figure 11. The radius=40 alpha hull of the Object Color Solid under D65 and for the 2° standard observer in CIE LAB.

Table 3: Cardinality of OCS-packing n-ball sets for different degrees of measurement repeatability

Measurement error (95th percentile, 0°/45°)	Packing n-ball radius	Packing n-ball set cardinality
0	0.5	3,207,431
0.2 (intra-instrument)	0.7	1,168,888
1.6 (inter-instrument)	2.1	43,292
2.2 (inter-vendor)	2.7	20,369

The cardinalities of OCS-packing color sets in Tab. 3 vary dramatically, from it being possible to pack the 2M volume with over 3M unit spheres that correspond to the implied distinguishability of CIE LAB, via a drop to around 1M just as a result of a small amount of measurement repeatability and further decreasing to tens of thousands when inter-instrumental differences typical of 0°/45° spectrophotometers are considered at a 95% level. In other words, given typical differences between spectrophotometers of different vendors, only around 20K colors from the OCS can reliably be known to be different on the basis of such measurements.

OCS cardinality in CIE CAM16 and CAM16-UCS

Since CIE LAB coordinates are not the best predictor of color, the analysis performed there can also be applied using a color appearance model that makes more accurate predictions about appearance attributes like lightness, chroma and hue. Doing so in CIE CAM16 [16] results in a volume of 2,013,465 cubic CIE CAM16 units (Fig. 12), which corresponds to 2,847,470 HCP-packed unit-distance-spaced points.
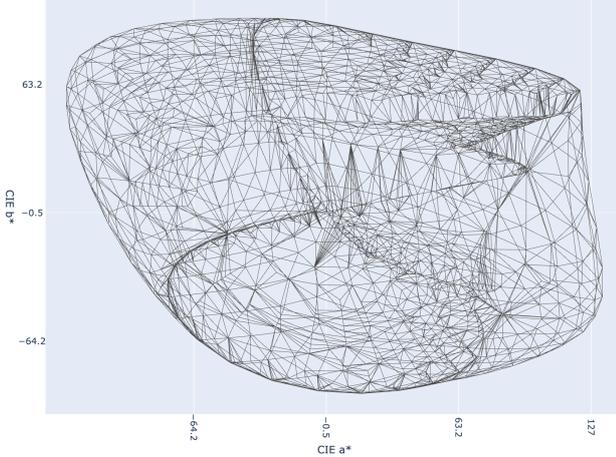
Figure 12. The radius=40 alpha hull of the Object Color Solid under D65 and for the 2° standard observer in CIE CAM16.

As was to be expected, the representations of the OCS in these two color spaces give rise to different cardinalities of the set of all color space coordinates that are a unit distance away from each other and that pack its color gamut boundary. In spite of their differences, the two spaces result in comparable cardinalities of 2.3M for CIE LAB and 2.0M for CIE CAM16. They are essentially answers to the same question though: in a space designed to represent individual colors, how many unit-distance-spaced coordinates fit in the boundary enclosing them?

These color spaces, however, have an important limitation, which is that they do not accurately represent how many just-noticeably different colors can be placed between two colors that are far apart. That is, while for a pair of colors that are two units apart, the prediction that only a single color can be placed between them that is distinguishable from both is reliable, but it does not follow that a pair of colors that is 20 units apart can accommodate 19 intermediate, pairwise-just-noticeably-different colors. Instead, only a smaller number of such intermediate colors would be possible. This "diminishing return" nature of color spaces like CIE LAB and CIE CAM16 corresponds to their being non-Riemannian, due to their additivity failure along a path connecting two colors [17, 18]:

$$||A-C|| > ||A-B|| + ||B-C|| \qquad (8)$$

where B is a color on the shortest path between colors A and C and $||x||$ is the L2 norm of vector x. This characteristic of color differences has also been long understood in the context of color difference equation research, where separate solutions have been developed for predicting small versus large differences.

A consequence then is that wrapping a boundary around accurate predictions of the colors of the OCS and then determining the cardinality of the unit-distance-spaced coordinates that pack it will result in an overprediction of how many just-noticeably different colors the OCS accommodates.

To better quantify that cardinality, a space is needed in which small, local distances accurately predict just-noticeable differences, instead of a space that accurately predicts the color appearance attributes of stimuli. One way to do so would be to take a space like CIELAB and then distort it to make it uniform in terms of a good small-distance ΔE metric like CIE ΔE2000 [19], which could be done using Urban et al.'s approach [20]. However, a recent, large-scale study has shown that an even better predictor of small color differences than CIE ΔE2000 [21] is the CAM16 Uniform Color Space (UCS) [22].

Representing the OCS in CAM16-UCS results in a volume of 502,736 cubic CAM16-UCS units (Fig. 13), which can be HCP-packed with 710,976 HCP points at unit distance from each other.
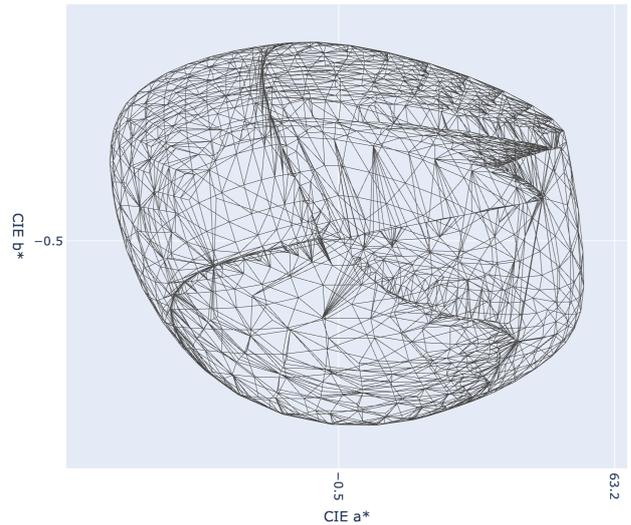
Figure 13. The radius=40 alpha hull of the Object Color Solid under D65 and for the 2° standard observer in CAM16-UCS.

In other words, the volume delimited by OCS colors would fit 2,847,470 HCP-packed unit-distance-spaced points in CIE CAM16, which is optimized for color appearance prediction. However, only 710,976 HCP-packed unit-distance-spaced points fit into the OCS in CAM16-UCS, which is optimized for small color difference prediction, and which is a better basis for answering the question of how many distinct colors there are within the OCS. The difference between the two is also an indication of the severity of non-additivity of color differences across the full OCS color gamut.

Instead of using unit diameter spheres that represent color difference perception between two spatially clearly separated and uniform color patches on a mid-gray background, it is also possible to use other thresholds for color difference in different contexts. E.g., when viewing continuous transitions of color, the perceptibility threshold may be somewhere in the region of 0.70 to 1.01 ΔE UCS units [23]. Using the lower value and packing the OCS gamut with 0.70 diameter spheres then gives 2,072,816 as the number of colors needed to construct continuous-appearing color transitions – around 4 times as many as can be perceived as different when viewed as uniform color patches.

Conclusions and next steps

In this paper an analysis of the cardinality of color stimuli and percepts has been presented taking a first-principles approach. Starting with the spectral domain, the effect of quantization and measurement noise was discussed and cardinality ("how many objects are there in a set") was introduced. Due to the nature of noise – whether spectral or colorimetric – a sphere-packing approach was presented and applied to both spectral and colorimetric data.

Considering quantization alone – applicable to all real data – a significant effect on the number of unique reflectances, as well as the colorimetric error due to quantization could be seen. Moving from an 8-bit encoding to 16-bit encoding results in an exponential increase in the number of unique reflectances that can be represented (by a factor of 10^{75}) which in turn results in going from an approximate potential median colorimetric error of 0.1 ΔE2000 (maximum of 0.35) to a median below 0.05 ΔE2000 (maximum of 0.15).

Turning to measurement noise and applying the HCP sphere-packing analysis to determine cardinality in the spectral domain (over a fixed volume of 1.0) showed that even low levels of measurement noise can have a significant effect on cardinality, e.g. without measurement noise (at a 16-bit precision and in 24 dimensions) 10^{122} unique reflectances can be considered, while at the maximum considered noise of 0.029 this number drops to 10^{44} (a loss by a factor of 10^{78}), corresponding to an increase in median colorimetric error from essentially 0.0 (0.00016) to 0.30 ΔE2000. Instead, in a biological domain, the responses to all spectra have a cardinality of only 2×10^{9} due to their much reduced dimensionality and the relative coarseness of retinal signal encoding.

Finally, applying the cardinality analysis via HCP sphere-packing in the CIE LAB, CIE CAM16 and CAM16-UCS color spaces, results show different cardinalities of color stimuli, from 2.3M for CIE LAB and 2.0M for CIE CAM16 to 0.7M for CAM16-UCS, assuming the same unit sphere diameter and for the viewing conditions and geometry under which canonical color difference equations are developed where color patches are viewed with a clear separation against a gray background. These results highlight some of the shortcomings of these color spaces, including their non-additivity. Context also greatly impacts color stimulus cardinality and it was shown how the perceptibility threshold for continuity of color transitions of 0.7 ΔE UCS yields 2.1M unique colors.

Acknowledgements

References

[1] F. D'Errico, L. Backwell, P. Villa, I. Degano, J. J. Lucejko, M. K. Bamford, T. F. G. Higham, M. P. Colombini, P. B. Beaumont, "Early evidence of San material culture represented by organic artifacts from Border Cave, South Africa". PNAS. 109 (33): 13214–13219. (2012). Doi:10.1073/pnas.1204213109

[2] E. Stevelinck, "Accounting in ancient times". Accounting Historians Journal, 12 (1): 1. (1985).

[3] J. Morovic, V. Cheung, P. Morovic, "Why We Don't Know How Many Colors There Are", *Proc. IS&T CGIV 2012 6th European Conf. on Colour in Graphics, Imaging, and Vision*, 49 – 53. (2012). https://doi.org/10.2352/issn

[4] J. H. Conway, N. J. H. Sloane, *Sphere Packings, Lattices and Groups*, Springer, New York, 3rd ed. (1998).

[5] Lagrange, J.-L. "Recherches d'arithmétique." *Nouv. Mém. Acad. Roy. Soc. Belles Lettres (Berlin)*, 265-312. (1773).

[6] G. Huber, "Gamma function derivation of n-sphere volumes," *Amer. Math. Monthly*, 89 (5): 301–302. (1982). Doi:10.2307/2321716

[7] C. F. Gauss, "Besprechung des Buchs von L. A. Seeber: Untersuchungen über die Eigenschaften der positiven ternären quadratischen Formen usw", *Göttingsche Gelehrte Anzeigen*. (1831).

[8] M. Viazovska, "The sphere packing problem in dimension 8", *Annals of Mathematics*, 185 (3): 991–1015. (2017). Doi:10.4007/annals.2017.185.3.7.

[9] H. Cohn, A. Kumar, S. D. Miller, D. Radchenko, M. Viazovska, "The sphere packing problem in dimension 24", Annals of Mathematics, 185 (3): 1017–1033. (2017). Doi:10.4007/annals.2017.185.3.8.

[10] A. T. Rider, G. B. Henning, A. Stockman, "Light adaptation controls visual sensitivity by adjusting the speed and gain of the response to light", PloS ONE, 14 (8), doi:10.1371/journal.pone.0220358. (2019).

[11] R. Shapley, C. Enroth-Cugell, "Visual Adaptation and Retinal Gain Controls." *Progress in Retinal Research*. 3:263–346, New York: Pergamon Press. (1984).

[12] S. G. Solomon, "Retinal ganglion cells and the magnocellular, parvocellular, and koniocellular subcortical visual pathways from the eye to the brain." *Handb Clin Neurol.*, 178:31-50. Doi: 10.1016/B978-0-12-821377-3.00018-0. (2021).

[13] H. Edelsbrunner, E. P. Mücke, "Three–Dimensional Alpha Shapes", *ACM Transactions on Graphics*, 13(1): 43–72, (1994).

[14] ISO/TS 18621-11:2021, Image quality evaluation methods for printed matter — Part 11: Colour gamut analysis. (2021).

[15] V. Baranau, D. Hlushkou, S. Khirevich, U. Tallarek, "Pore-size entropy of random hard-sphere packings", *Soft Matter*, 9(12): 3361-3372. (2013). http://dx.doi.org/10.1039/C3SM27374A

[16] CIE 248:2022, *The CIE 2016 Colour Appearance Model For Colour Management Systems: CIECAM16*, CIE, Vienna, doi: 10.25039/TR.248.2022. (2022).

[17] L. D. Griffin, D. Mylonas, Categorical colour geometry. *PloS One*, 14, e0216296. (2019).

[18] R. Bujack, E. Teti, J. Miller, E. Caffrey, T. L. Turton, The non-Riemannian nature of perceptual color space. *PNAS*, 119, e2119753119. (2022).

[19] ISO/CIE 11664-6:2022, *Colorimetry — Part 6: CIEDE2000 colour-difference formula*, CIE, Vienna. (2022).

[20] P. Urban, M. R. Rosen, R. S. Berns, D. Schleicher, "Embedding non-Euclidean color spaces into Euclidean color spaces with minimal isometric disagreement," *J. Opt. Soc. Am. A*, 24, 1516-1528. (2007).

[21] M. R. Luo, Q. Xu, M. Pointer, M. Melgosa, G. Cui, C. Li, K. Xiao, M. Huang, "A comprehensive test of colour-difference formulae and uniform colour spaces using available visual datasets," *Color Res Appl*, 48(3):267-282. (2023). Doi:10.1002/col.22844.

[22] M. R. Luo, G. Cui, C. Li, "Uniform colour spaces based on CIECAMO2 colour appearance model," *Color Res Appl*, 31(4):320-330. (2006).

[23] J. Morovic, P. Morovic, "The perceptibility of color differences in continuous transitions", *31st IS&T Color and Imaging Conference*, submitted for publication (2023).

[24] G. Cantor, "Ueber unendliche, lineare Punktmannichfaltigkeiten." *Math. Ann.* **15**:1–7 (1879). https://doi.org/10.1007/BF01444101

[25] Tajima J., Haneishi H., Ojima N., Tsukada M. (2002) Representative Data Selection for Standard Object Colour Spectra Database (SOCS), 10th Color Imaging Conference, 155-160.

[26] E. Schrödinger, "Theorie der Pigmente von größter Leuchtkraft", *Annalen der Physik* (Paris), 62:603-622, (1920).

[27] A. Stockman, L. T. Sharpe, "The spectral sensitivities of the middle- and long-wavelength-sensitive cones derived from measurements in observers of known genotype", *Vision Research*, 40(13):1711-1737, https://doi.org/10.1016/S0042-6989(00)00021-3. (2000).

[28] D. R. Wyble, D. C. Rich, "Evaluation of methods for verifying the performance of color-measuring instruments. Part I: Repeatability". *Color Res. Appl.*, 32: 166-175. https://doi.org/10.1002/col.20320. (2007).

[29] D. R. Wyble, D. C. Rich, "Evaluation of methods for verifying the performance of color-measuring instruments. Part II: Inter-instrument reproducibility." *Color Res. Appl.*, 32: 176-194. https://doi.org/10.1002/col.20308. (2007).

Author Biography

Ján Morovič received his Ph.D. in color science from the University of Derby (UK) in 1998, where he then worked as a lecturer. Since 2003 he has been at Hewlett-Packard in Barcelona as a senior color scientist and later master technologist. He has also served as the director of CIE Division 8 on Image Technology and Wiley and Sons have published his 'Color Gamut Mapping' book. He is the author of over 100 papers and has filed 180+ US patents (142 granted).

Peter Morovič received his Ph.D. in computer science from the University of East Anglia (UK) in 2002 and holds a B.Sc. in theoretical computer science from Comenius University (Slovakia). He has been a senior color and imaging scientist at HP Inc. since 2007, has published 65+ scientific articles and has 180+ US patents filed (140 granted) to date. His interests include color science, image processing, color vision, computational photography, computational geometry.