https://doi.org/10.2352/EI.2026.38.16.AVM-112

Mahalanobis Distance as a Robust Color (Signal) Separation Metric

Robin B. Jenkin; NVIDIA Corporation; Santa Clara, California, USA

Abstract

Quantifying the ability of an imaging system to distinguish between colors or signals is fundamental to evaluating camera performance, particularly for applications such as autonomous driving and surveillance. Traditional metrics such as ΔE_{2000} and $\Delta a^* b^*$ provide perceptually motivated color differences but are not designed to account for sensor noise, nor are they invariant to the linear signal processing stages common in imaging pipelines. This paper proposes the use of the Mahalanobis distance as a robust, noise-referred metric for color and signal separation. We demonstrate that the Mahalanobis distance is invariant to affine transformations—including white balance, color correction matrices, and linear color-space conversions—and therefore provides a stable figure of merit regardless of where in the linear pipeline the measurement is taken. We further examine practical considerations including the effects of sensor saturation, nonlinear transformations such as gamma and CIELAB conversion, spatial gradients, region-of-interest size, and target quality. Experimental results are presented for multiple color filter array configurations across a range of illumination levels, demonstrating the utility of the metric for both full signal separation (Yuv) and chrominance-only color separation (uv). The work is conducted within the context of the IEEE P2020 automotive image quality standard.

Introduction

The ability of a camera system to separate colors is a critical performance parameter for a wide range of applications. In automotive imaging, for example, reliably distinguishing a red traffic light from an amber one, or a pedestrian's clothing from a similarly colored background, directly impacts safety. In surveillance and machine vision, color separation determines whether objects of interest can be reliably detected and tracked under varying illumination conditions [7].

Traditionally, color differences have been quantified using perceptual metrics defined in the CIELAB color space, such as ΔE^*_{ab} [2] or its more sophisticated successor ΔE_{2000} [3]. While these metrics are well suited to predicting human perceptual judgments of color similarity, they present several shortcomings when applied to the evaluation of imaging system performance. First, they are not noise-referred: two color patches may have a large ΔE but still be indistinguishable in the presence of sensor noise. Second, the non-linear transformations required to convert from sensor RGB to CIELAB (including gamma correction and the cube-root nonlinearity) mean that the measured color difference depends on where in the image processing pipeline the computation is performed.

Mahalanobis distance [1] may be utilized as a metric that addresses some of these concerns. The Mahalanobis distance naturally incorporates the noise covariance structure of the image, providing a noise-referred separation measure in units of mutual standard deviations. Moreover, it is invariant to affine transformations, meaning that common linear image processing operations such as white balance gain adjustment, color correction matrix (CCM) application, and linear color-space conversion (e.g., RGB to Yuv) do not alter the measured distance. This invariance greatly simplifies system-level comparisons, since the metric need only be evaluated once at any convenient point in the linear portion of the pipeline. The IEEE P2020 working group was established to create standardized metrics for automotive camera image quality [4]. Within this framework, color performance has been decomposed into three tasks: color separation (the ability to distinguish between two color signals), color fidelity (how accurately a color is reproduced), and display color standards for human-in-the-loop applications such as backup cameras. This paper addresses the first of these tasks—color separation—and proposes the Mahalanobis distance as a principled metric for its measurement. The remainder of this paper is organized as follows. Existing color difference metrics and their limitations are reviewed after which the Mahalanobis distance is defined and its affine invariance demonstrated algebraically. Dimensionality reduction is discussed enabling the metric to be applied to either full signal separation or chrominance-only color separation. Practical considerations for measurement are examined, including saturation, non-linear transforms, spatial gradients, region-of-interest sizing, and target quality. Finally experimental results from simulations across multiple CFA configurations and illumination levels are discussed.

Background and Motivation

Perceptual Color Difference Metrics

The CIE 1976 ΔE^*_{ab} metric computes the Euclidean distance in CIELAB space:

$$\Delta E^*_{ab} = \sqrt{(\Delta L^*)^2 + (\Delta a^*)^2 + (\Delta b^*)^2} \tag{1}$$

where L^*, a^*, and b^* are the lightness and chrominance coordinates, respectively [2]. The ΔE_{2000} formula introduces weighting functions, a rotation term for the blue region, and parametric correction factors to improve perceptual uniformity [3, 5].

While ΔE_{2000} correlates well with human perception, it was not designed to characterize the discrimination capability of an electronic imaging sensor. A sensor with high noise may be unable to distinguish two patches whose ΔE_{2000} is well above the just-noticeable-difference threshold of approximately 1.0.

Noise-Referred Euclidean Distance

An intuitive first improvement to creating a noise-aware metric is to normalize the color difference by the pooled standard deviation, yielding a signal-to-noise ratio (SNR) style metric. For two patches with standard deviations $\sigma_{L,1}, \sigma_{L,2}$ in each $L^*a^*b^*$ channel, pooled standard deviations can be computed as:

$$\sigma_L = \sqrt{\frac{\sigma_{L,1}^2 + \sigma_{L,2}^2}{2}},\ \sigma_a = \sqrt{\frac{\sigma_{a,1}^2 + \sigma_{a,2}^2}{2}},\ \sigma_b = \sqrt{\frac{\sigma_{b,1}^2 + \sigma_{b,2}^2}{2}} \tag{2}$$

and combined in quadrature:

$$\sigma_{Lab} = \sqrt{\sigma_L^2 + \sigma_a^2 + \sigma_b^2} \tag{3}$$

The noise-referred ΔE_{ab}^* is then:

$$\text{Noise Referred } \Delta E_{ab}^* = \frac{\Delta E_{ab}^*}{\sigma_{Lab}} \tag{4}$$

This metric behaves correctly with respect to noise: as noise increases, the metric decreases; as noise decreases, the metric increases.

However, a simple scalar normalization ignores the correlation between color channels. In a Bayer color filter array (CFA) sensor, for example, the red and blue channels are spatially subsampled and subsequently demosaiced, introducing inter-channel correlations that a scalar noise figure does not capture. Furthermore, the Euclidean distance—even when noise-referred—is not invariant to linear transformations: applying a 3×3 color correction matrix changes both the signal separation and the noise covariance in a way that alters the Euclidean distance as demonstrated in Figure 1. Two 2D-point clouds are shown before and after a random affine transformation ($\det(\boldsymbol{A}) = 1.0$). Figure 1(a) shows the point clouds before transformation separated by a distance of ten standard deviations. The Euclidean distance is calculated as 10, the noise-referred Euclidean distance as slightly lower but close to 10 (9.75) due to the randomness of the point clouds. Figure 1(b) shows the result of transforming the data using a 2×2 transform. Both the Euclidean and noise-referred Euclidean distances increase to 12.49 and 11.59, respectively. The rotation of the point clouds in Figure 1(a) is arranged such that the data aligns with the first right singular vector, or the direction of maximum stretch in the output Figure 1(b). In Figure 1(c) we rotate the point clouds to align with the second right singular vector, or the direction of minimum stretch in the output Figure 1(d). In Figure 1(d) we can see this results in the Euclidean and noise-referred Euclidean distances dropping to 8 and 7.45. This has significant implications; depending on the dot product of the difference vector of the transformed colors and the singular vectors of the transformation, the Euclidean and noise-referred Euclidean distances can increase beyond their original values. This would imply that information can arbitrarily be created by simply applying a linear transformation to data which is clearly nonsensical.

Mahalanobis Distance

Consider two groups of observations in p dimensions (e.g., $p = 3$ for an RGB sensor). Let group i ($i = 1, 2$) have sample mean vector $\bar{\boldsymbol{x}}_i \in \mathbb{R}^p$, sample covariance matrix $\boldsymbol{S}_i \in \mathbb{R}^{p \times p}$, and n_i observations. The *pooled covariance* is defined as

$$\boldsymbol{S}_{\text{pool}} = \frac{(n_1 - 1)\boldsymbol{S}_1 + (n_2 - 1)\boldsymbol{S}_2}{n_1 + n_2 - 2} \tag{5}$$

and the Mahalanobis distance between the two groups is

$$D_M = \sqrt{(\bar{\boldsymbol{x}}_1 - \bar{\boldsymbol{x}}_2)^T \boldsymbol{S}_{\text{pool}}^{-1} (\bar{\boldsymbol{x}}_1 - \bar{\boldsymbol{x}}_2)}. \tag{6}$$

This quantity gives the separation between the two group centroids in units of pooled standard deviations, taking into account the full covariance [1, 6]. Intuitively, the inverse covariance matrix *whitens* the space, removing inter-channel correlations before computing the distance. The result is a dimensionless number that can be interpreted as the number of mutual standard deviations separating the two distributions.

Affine Invariance

A key property of the Mahalanobis distance is its invariance under affine transformations. Let $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b}$ be an affine mapping with $\boldsymbol{A} \in \mathbb{R}^{p \times p}$ non-singular and $\boldsymbol{b} \in \mathbb{R}^p$. Under this transformation the group means transform as $\bar{\boldsymbol{y}}_i = \boldsymbol{A}\bar{\boldsymbol{x}}_i + \boldsymbol{b}$ and the covariance matrices transform as $\boldsymbol{S}_i^{(y)} = \boldsymbol{A}\boldsymbol{S}_i\boldsymbol{A}^T$. Substituting into Eq. (6):

$$\begin{aligned} D_M^{(y)} &= \sqrt{(\boldsymbol{A}\boldsymbol{d})^T (\boldsymbol{A}\boldsymbol{S}_{\text{pool}}\boldsymbol{A}^T)^{-1} (\boldsymbol{A}\boldsymbol{d})} \\ &= \sqrt{\boldsymbol{d}^T \boldsymbol{A}^T (\boldsymbol{A}^T)^{-1} \boldsymbol{S}_{\text{pool}}^{-1} \boldsymbol{A}^{-1} \boldsymbol{A}\,\boldsymbol{d}} \\ &= \sqrt{\boldsymbol{d}^T \boldsymbol{S}_{\text{pool}}^{-1} \boldsymbol{d}} = D_M^{(x)} \end{aligned} \tag{7}$$

where $\boldsymbol{d} = \bar{\boldsymbol{x}}_1 - \bar{\boldsymbol{x}}_2$. The translation vector $\boldsymbol{b}$ cancels in the difference of means.

This invariance has immediate practical significance for imaging pipelines. The following common operations are all affine transformations of the sensor signal and therefore do not affect the Mahalanobis distance:

- **Black level correction (translation):** Subtracting the sensor pedestal corresponds to $\boldsymbol{b} \neq \boldsymbol{0}, \boldsymbol{A} = \boldsymbol{I}$.
- **White balance (scaling):** Channel-wise gain adjustment is a diagonal $\boldsymbol{A}$ matrix.
- **Color correction matrix (shearing/rotation):** The 3×3 CCM is a general non-singular linear map.
- **Linear color space conversion:** Transformations from RGB to Yuv or similar linear spaces correspond to multiplication by a fixed matrix.

Figure 1(a), (b), (c) and (d) demonstrate the invariance. The Mahalanobis distance remains as 9.72 despite the transformation of the data and the dot product between the difference vector and the first and second singular vectors of the transformation.

Signal Separation vs. Color Separation

The Mahalanobis distance as defined in Eq. (6) operates on the full p-dimensional observation vector. For a three-channel sensor, this measures the total *signal separation*, encompassing both luminance and chrominance differences. In many applications, however, it is desirable to evaluate *color separation* independently of intensity.

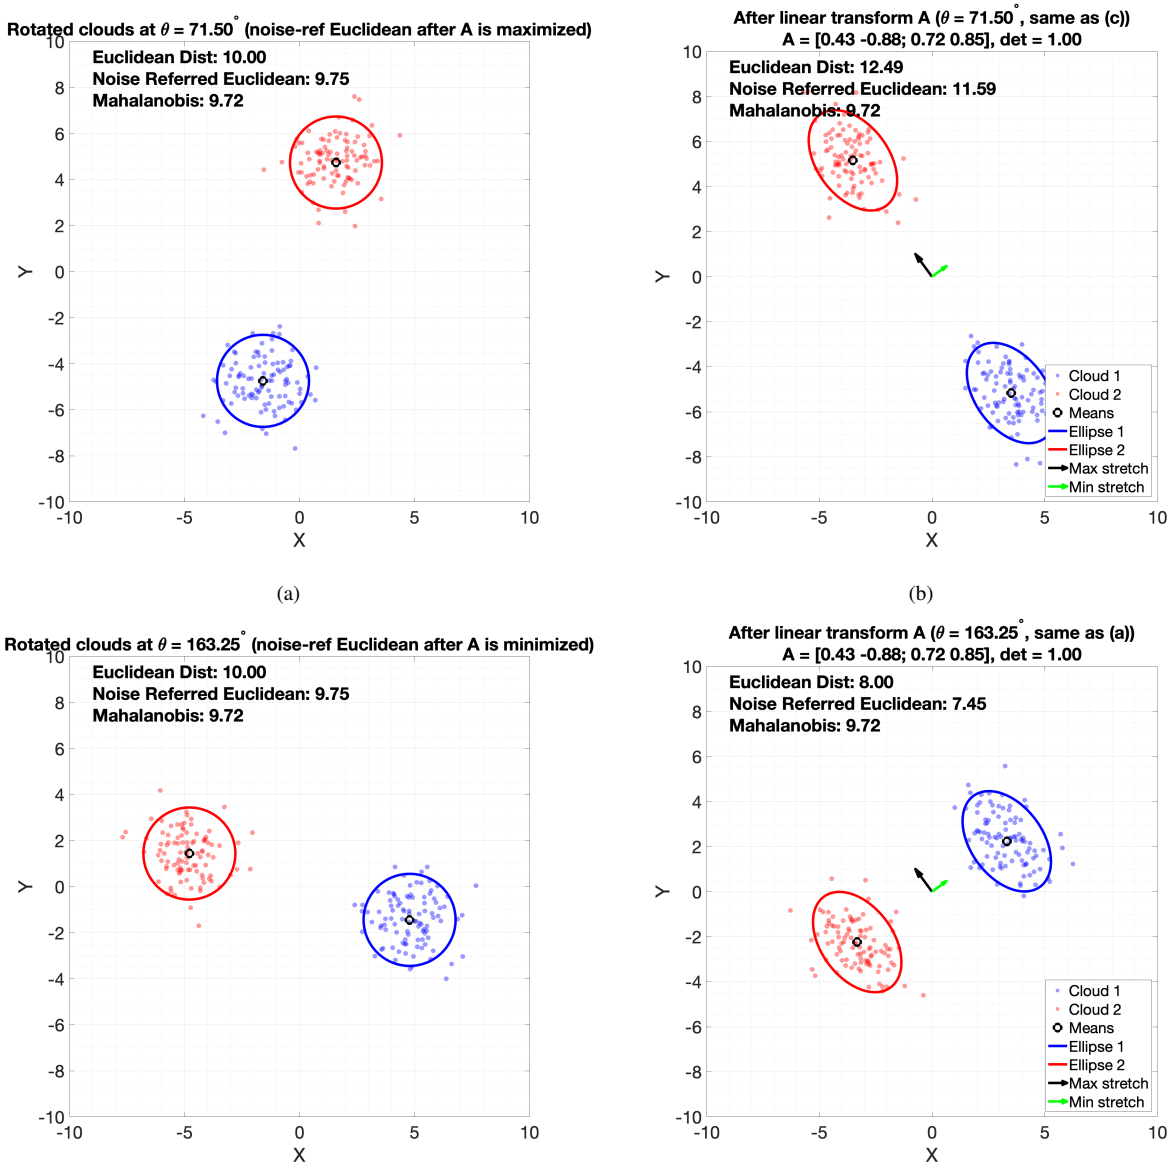

Figure 1: Demonstration of affine invariance. Two 2D-point clouds are shown before and after a random affine transformation ($\det(\boldsymbol{A}) = 1.0$). (a) shows the point clouds before transformation separated by a distance of ten standard deviations. The Euclidean distance is calculated as 10, the noise-referred Euclidean distance as slightly lower but close to 10 (9.75) due the randomness of the point clouds. Likewise Mahalanobis distance is slightly lower at 9.72. (b) shows the result of transforming the data using a $2x2$ transform. Both the Euclidean and noise-referred Euclidean distance increase to 12.49 and 11.59 respectively, yet the Mahalanobis distance remains the same, confirming the theoretical invariance of Eq. (7). The rotation of the point clouds in (a) is arranged such that the data aligns with the first right singular vector, or the direction of maximum stretch in the output (b). In (c) we rotate the point clouds to align with the second right singular vector, or the direction of minimum stretch in the output (d). In (d) we can see this results in the Euclidean and noise-referred Euclidean distances dropping to 8 and 7.45.

Because the Mahalanobis distance is invariant to linear transformations, one may freely convert from RGB to a luminance–chrominance space such as Yuv without altering the full signal separation metric. To obtain a color-only metric, one simply projects onto the chrominance subspace (the u–v plane) before computing the distance. This is equivalent to reducing the dimensionality from $p = 3$ to $p = 2$, discarding the luminance component Y. More generally, the metric can be applied in any dimensionality $p' < p$ by first projecting both groups onto the subspace of interest and then computing Eq. (6) with p'-dimensional mean vectors and $p' \times p'$ covariance matrices. For instance, projecting onto the Y axis alone yields a one-dimensional luminance separation metric. A key property when reducing dimensionality is that the Mahalanobis distance always decreases:

$$D_M^{(m)} \leq D_M^{(n)}, \quad m < n \tag{8}$$

This is consistent with information theory: removing or averaging channels discards information. For example, computing the Mahalanobis distance on the single Y channel yields a monochrome separation value that is consistently lower than the full RGB or YUV result. This property provides a useful sanity check.

Comparison with Other Distance Metrics

Several prior works have applied the Mahalanobis distance in color-related contexts. Imai et al. [10] demonstrated its use for perceptual color difference in complex images. Tsumura et al. [11] applied Mahalanobis-based color gamut mapping for endoscope image reproduction under different illuminants. Khongkraphan [12] applied the Mahalanobis distance for color edge detection. Al-Otum [13] proposed morphological operators for color image processing based on Mahalanobis distance measures. The present work differs in that it applies the metric specifically to CFA color separation measurement in the context of automotive imaging and safety-critical standards.

The Bhattacharyya distance [8] measures divergence between two distributions but its output is not in units of standard deviations. The squared Mahalanobis distance is proportional to the Bhattacharyya distance when the two classes share equal covariance matrices. The Kullback–Leibler divergence [9] is asymmetric and does not produce output with a direct probabilistic interpretation tied to separation. The Mahalanobis distance, by expressing separation in mutual standard deviations, provides the most direct pathway from measurement to probability of classification.

Practical Considerations

While the Mahalanobis distance offers attractive theoretical properties, several practical factors must be considered when applying it to real imaging data.

Sensor Saturation

Saturated pixels violate the assumption of multivariate normality that underlies the covariance estimate. When one or more color channels clip at the sensor's full-well capacity or a channel saturates after transformation, the observed distribution is truncated and the sample covariance matrix underestimates the true variance in those channels. This can artificially inflate the computed Mahalanobis distance. To mitigate this effect, pixels whose values are 0 or $2^b - 1$ should be identified and excluded from the analysis.

Non-Linear Transformations

The affine invariance property holds strictly only for linear (affine) operations. Non-linear transformations such as gamma correction ($y = x^\gamma$) and the conversion from XYZ to CIELAB (which involves a cube-root nonlinearity) *do* alter the Mahalanobis distance.

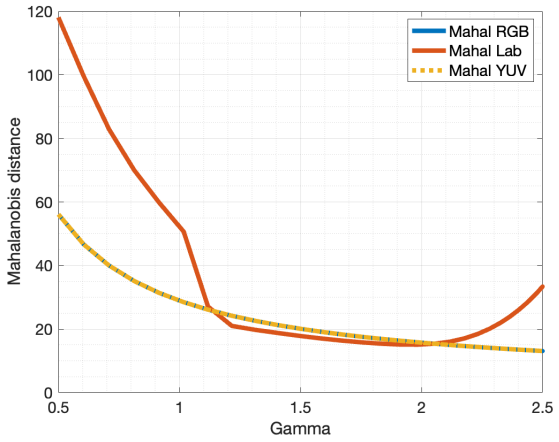

Figure 2: Mahalanobis distance as a function of gamma for three color representations (RGB, Lab, and Yuv) at high SNR (100:1). The RGB and Yuv curves overlap exactly, confirming affine invariance. The Lab curve diverges significantly, particularly at low gamma values, due to the cube-root nonlinearity in the CIELAB conversion.

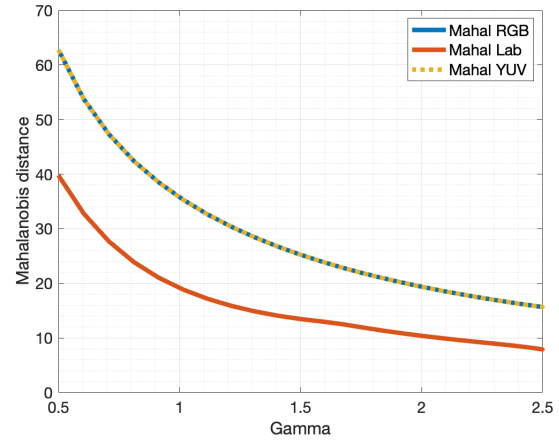

Figure 3: Mahalanobis distance versus gamma at a lower SNR (10:1) condition. Again the RGB and Yuv results are identical while Lab deviates, though the absolute differences are smaller than in the high-SNR case of Fig. 2.

Figure 2 illustrates this effect. When gamma is applied before computing the Mahalanobis distance, the RGB and Yuv results remain identical (since the Yuv conversion is linear), but the Lab result diverges markedly. At $\gamma = 1$ (the linear case), all three representations converge as expected. Figure 3 illustrates the change at lower SNR (10:1). Again the Lab result diverges

from the RGB or Yuv but is somewhat masked by the lower SNR and hence Mahalanobis distance to start with.

Anecdotally, while conversion to Lab is non-linear and changes the absolute Mahalanobis distance values, the relative ordering of color-pair separations appears to be largely conserved. A rigorous characterization of this approximate ordering preservation is left for future work.

Spatial Gradients

Illumination non-uniformities across a color patch introduce a spatial gradient in the measured signal. Because the Mahalanobis distance is computed from the sample covariance of the pixel values within a region of interest (ROI), any systematic spatial variation inflates the estimated variance and consequently deflates the computed distance.

At high SNR, spatial gradients may dominate the variance estimate, leading to a measured distance that characterizes the uniformity of the illumination rather than the sensor's intrinsic color discrimination. At low SNR, the photon shot noise may dominate, partially masking the gradient effect, but the bias remains. For accurate measurements, illumination should be as uniform as possible; alternatively, one may fit and subtract a low-order polynomial surface from each ROI before computing the covariance.

Region-of-Interest Size

The sample covariance matrix requires a sufficient number of observations for a stable estimate. For a p-dimensional measurement, the covariance matrix has $p(p+1)/2$ unique parameters; a minimum ROI size is needed to estimate these reliably.

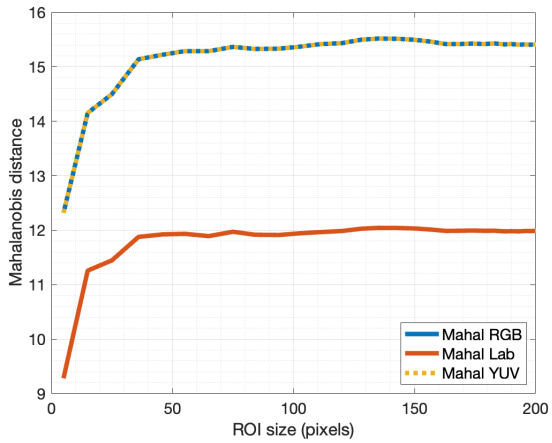

Figure 4: Mahalanobis distance as a function of ROI size (in pixels per side) for RGB, Lab, and Yuv. The metric stabilizes above approximately 30×30 pixels for both the linearly related RGB/Yuv pair and for Lab. Smaller ROIs yield unreliable estimates due to insufficient samples for robust covariance estimation.

Figure 4 shows the measured Mahalanobis distance as a function of ROI size. For ROIs smaller than approximately 30×30 pixels, the estimate is unstable and biased. Above this threshold, the metric converges to a stable value. For the linear representations (RGB and Yuv), the curves coincide as expected. We recommend a minimum ROI size of 30×30 pixels for three-channel measurements.

Target Quality

Physical color targets used for measurement are subject to degradation over time, including fading, scratching, and contamination. Such degradation increases the intra-patch variance and reduces the measured Mahalanobis distance, potentially leading to pessimistic assessments of system performance.

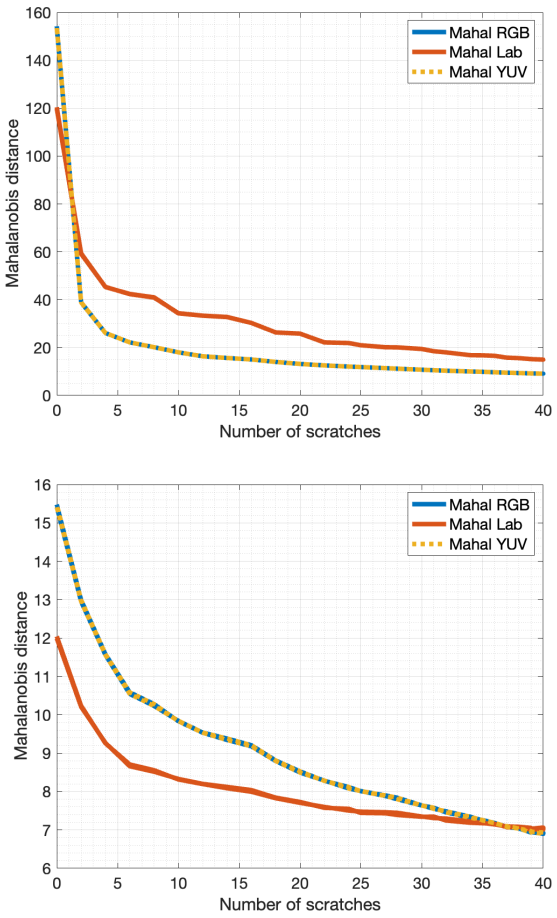

Figure 5: Effect of target degradation (simulated scratches) on the Mahalanobis distance at linear SNR of 100:1 (top) and 10:1 (bottom). At low SNR the impact is modest since sensor noise already dominates the variance. At high SNR, target imperfections significantly reduce the measured separation, underscoring the need for high-quality targets or transmissive test setups.

Figure 5 shows simulated target degradation at two SNR levels. At low linear SNR (10:1), sensor noise dominates the covariance and target scratches have a relatively modest effect. At high linear SNR (100:1), the same scratches dramatically reduce the Mahalanobis distance. For demanding measurements, programmable light sources or transmissive targets are recommended to eliminate the contribution of target surface imperfections.

Experimental Results

To demonstrate the practical utility of the Mahalanobis distance, we present measurements across multiple color filter array

(CFA) configurations and a range of illumination levels. Measurements were conducted using a radiometric sensor simulation based on an f/1.6 2.1 μm pixel pitch camera with an RGB CFA as detailed by Jenkin in [14] and [15].

Color patch pairs from a standard color target were imaged at illumination levels spanning approximately four decades (from 10^{-1} to 10^{3} lux), Figure 6. Six CFA configurations were evaluated: the standard Bayer RGB pattern, RCB, RYCy (red–yellow–cyan), RCG (red–clear–green), MYCy (magenta–yellow–cyan), and MCCy (magenta–clear–cyan). For each patch pair and illumination level, the Mahalanobis distance was computed in three representations: full-signal Yuv, chrominance-only uv, and raw RGB.

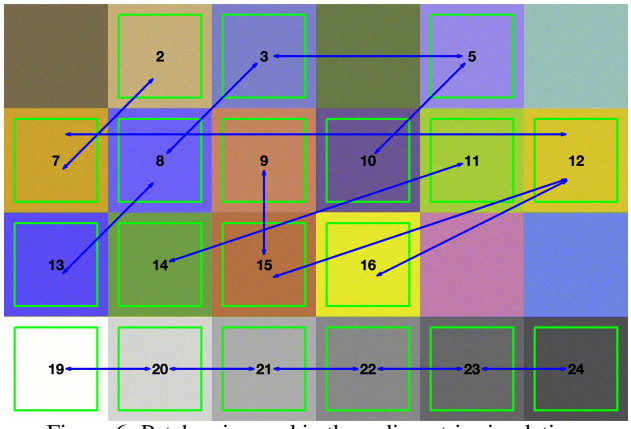

Figure 6: Patch pairs used in the radiometric simulation.

Mahalanobis Distance vs. Light Level

Figure 7 shows the Mahalanobis distance as a function of illumination level for a standard Bayer RGB CFA. The distance increases monotonically with light level, consistent with the expected improvement in SNR as photon flux increases. The relationship is approximately linear on a log–log scale, reflecting the square-root dependence of shot-noise-limited SNR on signal level. Different color patch pairs exhibit different absolute separations, but their relative ordering is generally preserved across the illumination range. When measuring the chrominance (uv) difference of the monochromatic pairs, notice that this drops to zero as would be expected for achromatic signals.

Comparison Across CFA Configurations

Figure 8 presents the full-signal (Yuv) Mahalanobis distance for all six CFA configurations. The standard Bayer RGB provides a familiar baseline. The alternative CFA designs—RCB, RYCy, RCG, MYCy, and MCCy—show distinct separation profiles. Certain configurations provide improved separation for specific color pairs, particularly at low light levels where noise dominates.

Figure 9 shows the corresponding chrominance-only (uv) Mahalanobis distance. By removing the luminance component, the metric isolates the color-discrimination ability of each CFA. The relative performance of the CFA configurations changes when only chrominance is considered, demonstrating that a CFA which provides superior total signal separation does not necessarily provide the best color separation. This distinction is important for applications where color information is the primary discriminant (e.g., traffic light classification) versus those where total signal contrast matters (e.g., object detection). Notice also that the patch pair ordering for each individual CFA changes. These simulations clearly demonstrate that while RCB provides good sensitivity, its color separation is poor compared to RGB. The MYCy CFA appears to be a reasonable trade of some sensitivity gain for a modest reduction in color separation.

Discussion

The experimental results confirm the theoretical advantages of the Mahalanobis distance as a color and signal separation metric. Its affine invariance eliminates a significant source of measurement ambiguity in imaging system evaluation: the dependence on where in the processing pipeline the metric is computed. An engineer can evaluate the metric at the raw sensor output, after white balance, or after CCM application, and obtain the same result—provided these operations are linear.

The sensitivity to non-linear transformations should be viewed not as a limitation but as a feature of the metric's design philosophy. Non-linear operations genuinely alter the statistical separation of the data, and the Mahalanobis distance correctly reflects this change. The practical implication is that the metric should be computed in the linear domain of the pipeline for maximum interpretability and reproducibility. Studies on the effect of non-linear operations should be carried out.

The practical guidelines outlined provide a framework for reliable measurement. In summary:

- Exclude saturated pixels from the analysis.
- Compute the metric in the linear domain, before gamma or other non-linear operations.
- Use uniform illumination or correct for spatial gradients.
- Use ROIs of at least 30×30 pixels.
- Use high-quality targets, or preferably programmable illumination sources or transmissive targets, to minimize target-induced variance.

The ability to decompose the metric into signal separation (Yuv) and color separation (uv) provides system designers with a powerful tool for evaluating CFA design choices. As demonstrated in Figures 8 and 9, alternative CFA configurations may offer advantages for specific application requirements, and the Mahalanobis distance provides a principled basis for such comparisons. Further, separations expressed in standard deviations are easily converted to probabilities of separation via the use of z-scores.

Conclusion

We have presented the Mahalanobis distance as a robust, noise-referred metric for quantifying color and signal separation in imaging systems. The metric offers several key advantages over traditional approaches: it naturally incorporates the sensor's noise covariance structure; it is invariant to all affine transformations including white balance, color correction, and linear color-space conversion; and it provides an intuitive interpretation as the number of mutual standard deviations separating two color distributions.

Non-linear transformations such as gamma correction and conversion to CIELAB do affect the metric, but the metric should be computed in the linear domain for maximum utility. Anecdotally, while the conversion to Lab changes absolute values,

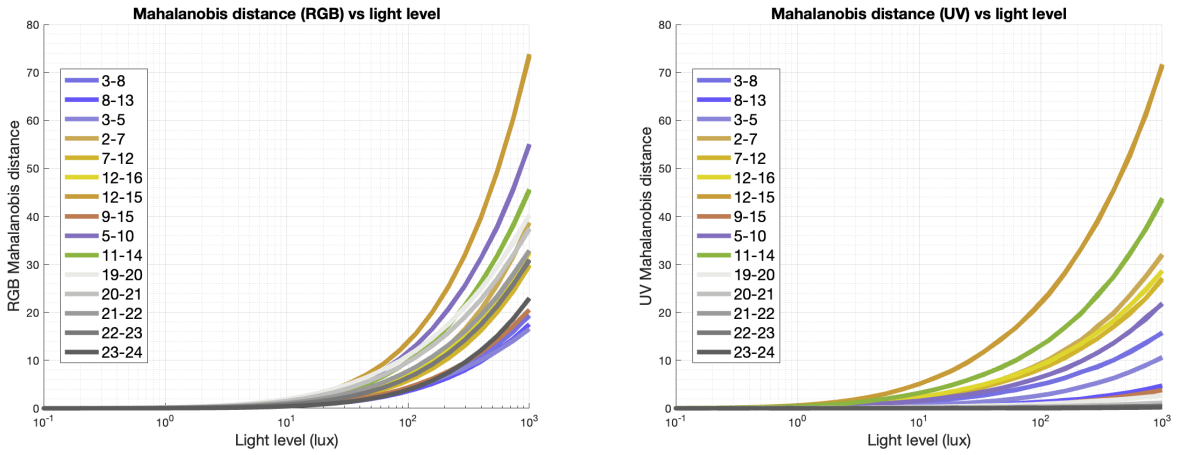

(a) (b)

Figure 7: Mahalanobis distance versus illumination level for a standard RGB CFA sensor. Left: RGB space. Right: uv (chrominance only). Each curve represents a different color patch pair from the target. As expected, separation increases monotonically with illumination due to the improving SNR. The uv metric isolates the chrominance component, showing different relative rankings among patch pairs compared to the full RGB metric. Notice that the chrominance (uv) separation of the monochromatic patches drops to zero even at significant light levels.

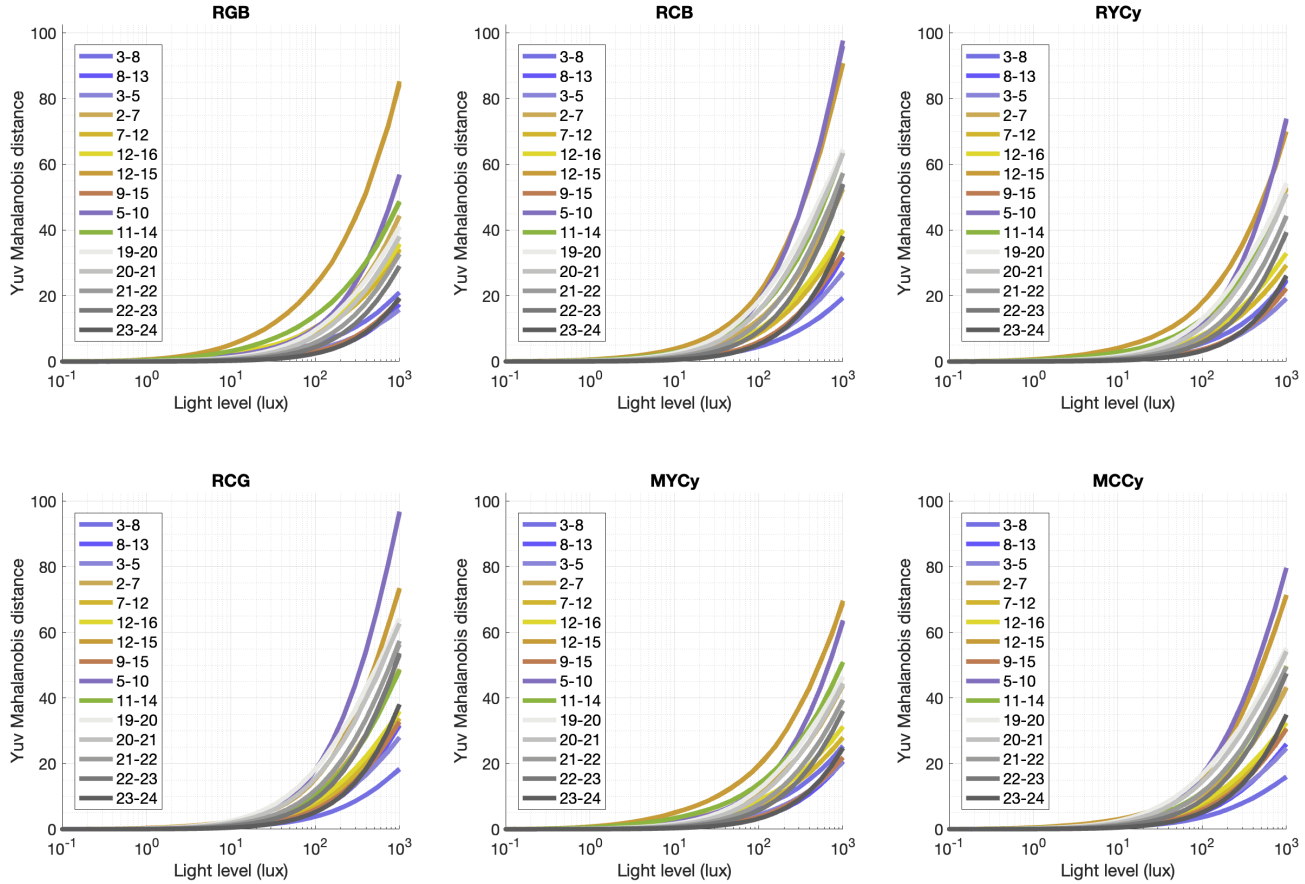

Figure 8: Yuv Mahalanobis distance (signal separation) versus light level for six CFA configurations: RGB, RCB, RYCy, RCG, MYCy, and MCCy. Each curve corresponds to a different color patch pair. The non-standard CFAs (particularly RYCy and MYCy) show different separation profiles, with some patch pairs exhibiting improved discrimination at low light levels compared to the standard Bayer RGB.

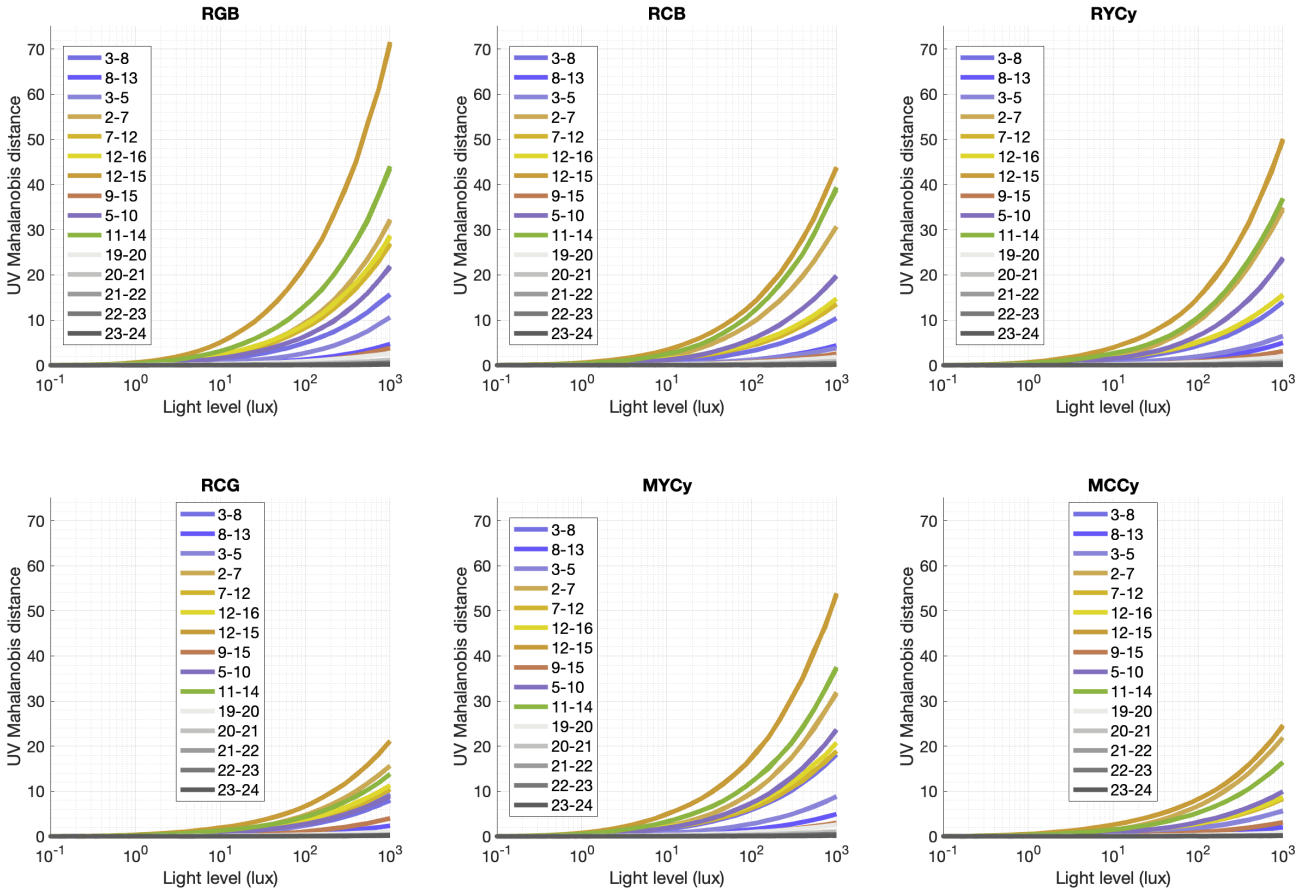

Figure 9: UV Mahalanobis distance (color separation only) versus light level for the same six CFA configurations. Compared to Fig. 8, the chrominance-only metric reveals different relative performance among the CFA designs. For example, MYCy shows improved color separation over RCB for certain patch pairs at intermediate illumination levels, highlighting the value of evaluating both signal and color separation independently.

relative relationships between color pairs appear to be largely conserved—a property that warrants further theoretical and experimental investigation.

Practical guidelines for measurement have been established, with particular emphasis on the need for adequate ROI size ($\geq 30 \times 30$ pixels for 3-dimensional data), uniform illumination, unsaturated data, and high-quality targets. Programmable light sources or transmissive targets are recommended for demanding applications.

Experimental results across six CFA configurations demonstrate that the Mahalanobis distance enables meaningful comparisons of color and signal separation performance, and that the choice between signal separation (Yuv) and color separation (uv) can reveal important differences in CFA design trade-offs.

Acknowledgments

The author thanks Dr. Brian Deegan for many useful conversations during the development of this work.

References

[1] P. C. Mahalanobis, "On the generalized distance in statistics," *Proceedings of the National Institute of Sciences of India*, vol. 2, no. 1, pp. 49–55, 1936.

[2] CIE, *Colorimetry*, 3rd ed., CIE Publication 15:2004, Vienna, Austria, 2004.

[3] G. Sharma, W. Wu, and E. N. Dalal, "The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations," *Color Research & Application*, vol. 30, no. 1, pp. 21–30, 2005.

[4] IEEE P2020 Working Group, White paper – IEEE P2020 automotive imaging, IEEE Standards Association, 2019.

[5] M. R. Luo, G. Cui, and B. Rigg, "The development of the CIE 2000 colour-difference formula: CIEDE2000," *Color Research & Application*, vol. 26, no. 5, pp. 340–350, 2001.

[6] G. J. McLachlan, *Discriminant Analysis and Statistical Pattern Recognition*, Wiley-Interscience, 2004.

[7] R. Lukac, *Computational Photography: Methods and Applications*, CRC Press, 2010.

[8] A. Bhattacharyya, On a measure of divergence between two statistical populations defined by their probability distributions, Bull. Calcutta Math. Soc., **35**, 99–109 (1943).

[9] S. Kullback and R. A. Leibler, On information and sufficiency, Ann. Math. Stat., **22**(1), 79–86 (1951).

[10] F. H. Imai, N. Tsumura, and Y. Miyake, Perceptual color difference metric for complex images based on Mahalanobis distance, J. Electron. Imaging, **10**(2), 385–393 (2001).

[11] N. Tsumura, F. H. Imai, T. Saito, H. Haneishi, and Y. Miyake, Color gamut mapping based on Mahalanobis distance for color reproduction of electronic endoscope image under different illuminant, Proc. IS&T/SID Color Imaging Conf., **5**, 158–161 (1997).

[12] K. Khongkraphan, An efficient color edge detection using the Mahalanobis distance, J. Inf. Process. Syst., **10**(4), 589–601 (2014).

[13] M. N. Al-Otum, Morphological operators for color image processing based on Mahalanobis distance measure, Opt. Eng., **42**(9), 2595–2606 (2003).

[14] R. B. Jenkin and P. J. Kane, Fundamental imaging system analysis for autonomous vehicles, Proc. IS&T Electronic Imaging: Autonomous Vehicles and Machines, IS&T, Springfield, VA, 2018.

[15] R. Jenkin, "The influence of CFA choice on automotive and other critical imaging systems," in *Proceedings of the International Colour Association (AIC) Conference 2021, 14th Congress*, Milan, Italy (Online), August 30–September 3, 2021. Invited paper.

Author Biography

Robin Jenkin received his BSc(Hons) Photographic and Electronic Imaging Science (1995) and a PhD (2001) in the field of image science from the University of Westminster. He also holds an M.Res Computer Vision and Image Processing from University College London (1996). Robin is a Fellow of The Royal Photographic Society, UK, and Executive Vice President of Society for Imaging Science and Technology. Robin is secretary of the IEEE P2020 Image Quality for Autonomous Vehicles working group. At NVIDIA Corporation, Robin is a distinguished engineer and models image quality for autonomous vehicle and other applications. He is a Visiting Professor at the University of Westminster within the Computer Vision and Imaging Technology Research Group and co-author of the 10th ed. "The Manual of Photography", Focal Press.

JOIN US AT THE NEXT EI!

electronic IMAGING

Imaging across applications . . . Where industry and academia meet!

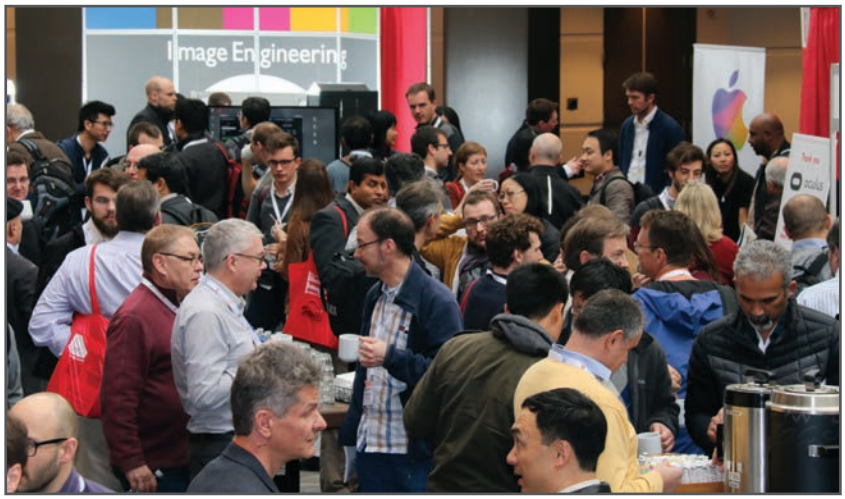

• SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •
• INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •

www.electronicimaging.org

ADVANCING IMAGING SCIENCE 75 1947 IS&T 2022 imaging.org