Sensor systems

Rethinking the Concept of Pixel Intensity Contrast From a Machine Learning Perspective

Sanush Abeysekera, Melanie Po-Leen Ooi, Ye-Chow Kuang, Shah Faisal, Yaminn Thawdar, Geoffrey Holmes, Dale Fletcher, and Peter Reutemann

[1] *School of Engineering, The University of Waikato, Hamilton 3216, New Zealand*
[2] *School of Engineering, Technology, Sunway University, Selangor 47500, Malaysia*
[3] *School of Computing, Mathematical Sciences, The University of Waikato, Hamilton 3216, New Zealand*
[*] *Member, IEEE*
[**] *Senior Member, IEEE*

Manuscript received 15 September 2025; revised 20 October 2025; accepted 27 October 2025. Date of publication 30 October 2025; date of current version 17 November 2025.

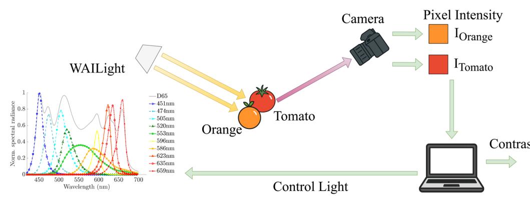

Abstract—Image contrast is a critical factor for machine vision tasks. A promising approach for enhancing contrast involves the use of algorithmically optimized, spectrally tunable illumination. However, the very definition of "contrast" is often rooted in principles of human perception, which may not be optimal for a machine observer. For an algorithm, contrast is an objective, task-driven metric that can be mathematically defined. To investigate the impact of this definition, we first use eigenvalue-based optimization algorithms to compute optimal illumination spectra. We then systematically evaluate these spectra using four distinct, physically realizable contrast formulations. Our analysis reveals that the performance of a given optimization algorithm is entirely dependent on the subsequent choice of evaluation metric. An illumination spectrum considered optimal under one metric can be significantly suboptimal when measured by another. This demonstrates that the choice of contrast metric is not a passive measurement, but an active design parameter with tangible physical consequences. From a machine learning perspective, the choice of this "loss function" should be codesigned with the physical hardware and the ultimate downstream task to achieve true system-level optimization.

Index Terms—Sensor systems, cameras, image processing, machine learning, machine vision, optimization algorithms, pixel contrast.

I. INTRODUCTION

The field of optical imaging historically treated hardware and software as separate, sequential stages. However, advancements in artificial intelligence and electronics allow us to consider an integrated framework that jointly optimizes both optical systems (the physical layer) and computational algorithms (the digital layer) in a field called "computational imaging." In this approach, physical optics are engineered to encode information into the light field, which is then captured by a sensor and decoded by a dedicated algorithm. Our work on optimizing illumination is a direct application of this principle, where a tunable light source acts as a programmable component of the "optical encoder," actively shaping the information captured by the sensor.

Contrast is a fundamental feature for the vast majority of machine vision algorithms. Historically, illumination design has relied heavily on manual techniques and heuristics, such as bright-field and dark-field setups, backlighting, and the strategic use of color filters, to improve image quality and feature visibility. Our previous work [1] introduced a framework to move beyond these heuristics by algorithmically computing the optimal illumination spectrum based on the measured spectral properties of the objects and the imaging sensor.

This letter addresses a more fundamental question that arises from that work: how should "contrast" be defined for a machine? For human observers, contrast is a well-studied perceptual phenomenon, and many image quality assessment metrics are designed to model the Human Visual System [2]. However, in machine vision, the ultimate "observer" is not the camera sensor itself, but the downstream algorithm chosen for high-level tasks, such as classification, segmentation, or object detection. This algorithm's performance is governed by a mathematical loss function. Therefore, for a machine, the most effective definition of contrast is one that maximizes separability in a way that directly benefits this loss function. This reframing of a performance metric as a task-specific, tunable objective is an emerging theme in sensor design, with recent works using machine learning to optimize sensor models [3] and even tailoring loss functions to improve quantification tasks [4]. We therefore posit that the concept of "contrast" is thus a definable, task-specific component of the system. This letter explores the consequences of this freedom by asking: How does the choice of this engineered definition impact the physical solution to the optimization problem?

The contribution of this article is demonstrating the physical consequences of the choice of "loss function." We use optimization objectives from [1] to compute and physically realize optimal spectra with our tunable LED illuminant. We then systematically evaluate these physical spectra using four distinct, physically realizable contrast metrics. This evaluation reveals that the choice of optimization objective directly alters the physical light spectrum. This change, in turn, dictates performance when measured by the four evaluation metrics. Therefore, the choice of the "loss function" is a key design parameter in end-to-end computational imaging systems.

Corresponding author: Sanush Abeysekera. (e-mail: sanush.abeysekera@waikato.ac.nz).
Associate Editor: Flavio Esposito.
Digital Object Identifier 10.1109/LSENS.2025.3627240

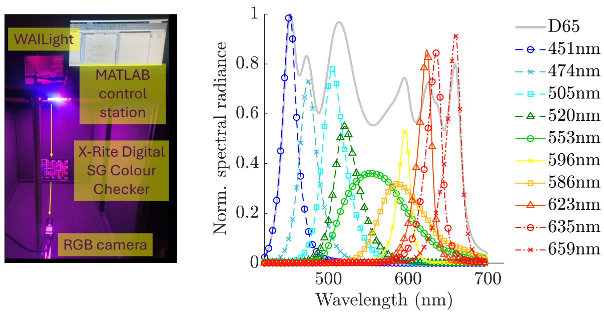

(a) (b)

Fig. 1. Imaging setup. (a) Tuneable luminaire (WAILight) illuminates the test sample (a CC SG), which is observed by the camera. (b) Individual LED channels of WAILight and a reproduction of the noon-daylight D65 spectrum.

II. METHODOLOGY

The experimental setup is shown in Fig. 1(a). The objective is to programmatically control a spectrally tunable luminaire (WAILight), to maximize image contrast between selected pairs of color patches from a ColorChecker Semi-Gloss (CC-SG) target, as observed by an 8-bit Basler ace acA1300-200uc red, green, blue (RGB) camera. The CC-SG is selected because it consists of flat, diffusely reflecting patches, ensuring that measured intensities are governed solely by the absorption spectra of the materials. This provides an ideal testbed, enabling the generation of 4851 unique pairwise color combinations. As illustrated in Fig. 1(b), WAILight comprises ten distinct, narrow-band LED channels, and the desired illumination spectrum is synthesized by modulating the brightness of each channel.

The pixel intensity $p_k \in \{p_r, p_g, p_b\}$ recorded by the kth camera channel for an object with reflectance $R \in \mathbb{R}^{1\times\lambda}$, measured at λ discrete wavelengths (400–700 nm at 5 nm resolution), can be expressed as

$$p_k = e\,R\,\omega_k\,L\beta$$

where $\omega_k \in \mathbb{R}^{\lambda\times\lambda}$ is a diagonal matrix representing the spectral sensitivity of the kth channel and e denotes the camera exposure, set to ensure a linear response [1], [5]. Here, $L \in \mathbb{R}^{\lambda\times 10}$ is the matrix of LED spectral outputs at maximum brightness [see Fig. 1(b)] and $\beta \in \mathbb{R}^{10\times 1}$ contains the modulation coefficients controlling the brightness of each LED channel. For an RGB camera, the pixel intensity is represented as a 3-D vector: $I = [p_r, p_g, p_b]$.

To enhance contrast between two objects R_1 and R_2, the task is to determine β that maximizes

$$D_k\beta = e(R_1 - R_2)\omega_k L\beta.$$

In our previous work [1], we introduced three eigenvalue-based algorithms: EIG1, EIGs, and EIG3. The simplest, EIG1, finds β that maximizes each camera channel independently

$$\max_{\beta} \frac{\beta^T D_k^T D_k \beta}{\beta^T\beta} = \max_{\beta} \frac{\beta^T H_k \beta}{\beta^T\beta}. \quad (1)$$

Equation (1) is solved analytically, with constraints $\beta \in [0, 1]^{10\times 1}$ as detailed in [1]. The channel with the highest contrast is selected as the optimal spectrum.

In EIGs, instead of selecting one optimum channel, the three EIG1 illuminants are synthetically combined into one output (see Fig. 2). In contrast, EIG3 produces a single optimum β that simultaneously optimizes all three channels by replacing the Hessian in (1), with $H = \sum_{k=1}^{3} H_k = H_r + H_g + H_b$.

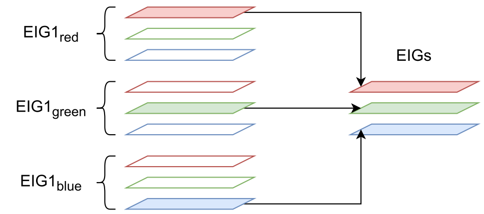

Fig. 2. EIGs algorithm, where pixel intensities obtained under each EIG1 illuminant are combined into a synthetic image.

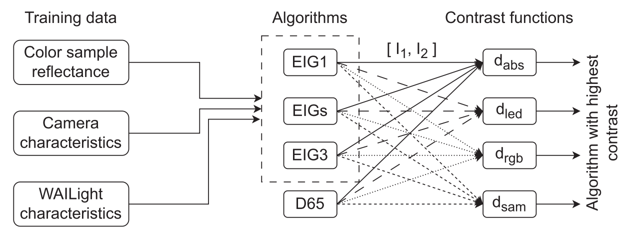

Fig. 3. Experimental workflow: Contrast between pixel intensities I_1 and I_2 is evaluated under EIG1, EIGs, EIG3, and D65 illuminants using four contrast functions (d_{abs}, d_{led}, d_{rgb}, and d_{sam}).

These algorithms were benchmarked against the D65 reference illuminant (noon daylight), as shown in Fig. 1(b). While that work concentrated on algorithmic strategies, it implicitly relied on a specific definition of contrast.

A key question, not addressed in prior studies, is how the choice of contrast metric influences the solution. For machine vision, "contrast" is not a fixed perceptual quantity but a design parameter that determines the optimal illumination.

A. Experimental Flow

The experimental flow for evaluating the contrast metrics is shown in Fig. 3. The optimization algorithms use the spectral reflectance, camera, and WAILight characteristics as input. As output, each algorithm, along with the D65 reference, generates a pair of pixel intensities I_1 and I_2 for the two target samples.

This letter compares four distinct contrast functions: d_{abs}, d_{led}, d_{rgb}, and d_{sam}. Each contrast function receives four sets of pixel intensity pairs: three optimized by the EIG algorithms and one under the D65 illumination. For each contrast function, the performance of the EIG algorithms and D65 is evaluated based on the contrast value it yields. This allows for a systematic assessment of which illumination strategy provides the "best" contrast as defined by each specific metric, thereby highlighting the impact of the chosen contrast function itself.

B. Contrast as a Malleable Objective Function

In this study, we investigate four distinct objective functions for contrast, each with a different physical and conceptual motivation. These metrics, used to evaluate the performance of different illumination optimization algorithms, are summarized in Table 1. While this does not cover all possible contrast functions, these metrics are particularly suitable for evaluating a binary contrast optimization problem.

Absolute contrast (d_{abs}) is the most direct method, maximizing the absolute Euclidean distance between the pixel intensity vectors. While simple, it is susceptible to favoring brighter overall illumination, which can lead to sensor saturation and may not be energy-efficient. The second metric, illumination normalized contrast (d_{led}), addresses this

Table 1. Summary of the Contrast Metrics Used in the Study

Metric	Formula	Physical interpretation	References
Absolute Contrast	$d_{abs} = \|\|I_1 - I_2\|\|_2$	Maximizes absolute contrast. Prioritizes overall intensity difference.	
Illumination Normalized Contrast	$d_{led} = \frac{\|\|I_1 - I_2\|\|_2}{\|\|l\|\|_1}$	Maximizes contrast per unit illumination energy $l = L\beta$. Emphasizes the spectral shape of illumination over its brightness.	[1], [6]
Michelson Contrast	$d_{rgb} = \frac{\|\|I_1 - I_2\|\|_2}{\|\|I_1 + I_2\|\|_2}$	Maximizes relative intensity difference. Invariant to uniform scaling of overall brightness.	[7], [8]
Spectral Angle Mapper	$d_{sam} = \cos^{-1}\left(\frac{I_1 \cdot I_2}{\|\|I_1\|\|\|\|I_2\|\|}\right)$	Quantifies the angular separation (rad) between two color vectors. Measures chromatic differences, insensitive to intensity variations.	[9]

weakness by normalizing the pixel difference by the total energy of the illumination spectrum, $l = L\beta$, where $l \in \mathbb{R}^{1\times\lambda}$. The total energy $|l|$ is calculated as the L1-norm. This metric, derived from our original work [1] and also implied by Lee et al. [6], effectively maximizes the contrast achieved per unit of input light energy. It isolates the contribution of the illumination's spectral shape from its overall power, making it a measure of spectral efficiency.

The Michelson contrast (d_{rgb}) normalizes the Euclidean distance between two pixel intensity vectors by their sum. This long-standing metric, often used for "visibility" in vision science, is also conceptually similar to the contrast-to-noise ratio widely employed in medical imaging to quantify the separability of a region of interest from its background. Its utility lies in its invariance to uniform scaling of overall brightness, ensuring it effectively captures relative intensity differences.

Finally, complementing these intensity-based metrics, the spectral angle mapper (d_{sam}) quantifies contrast by measuring the angular separation between two color vectors. Unlike Euclidean distance metrics, d_{sam} is insensitive to the magnitude (brightness) of the pixel intensities. Instead, it exclusively measures their chromatic or spectral similarity, focusing on the "direction" of the color vector. This makes it highly effective for distinguishing colors based on their hue and saturation, independent of their absolute luminance, and is widely used in applications, such as hyperspectral imaging and material classification.

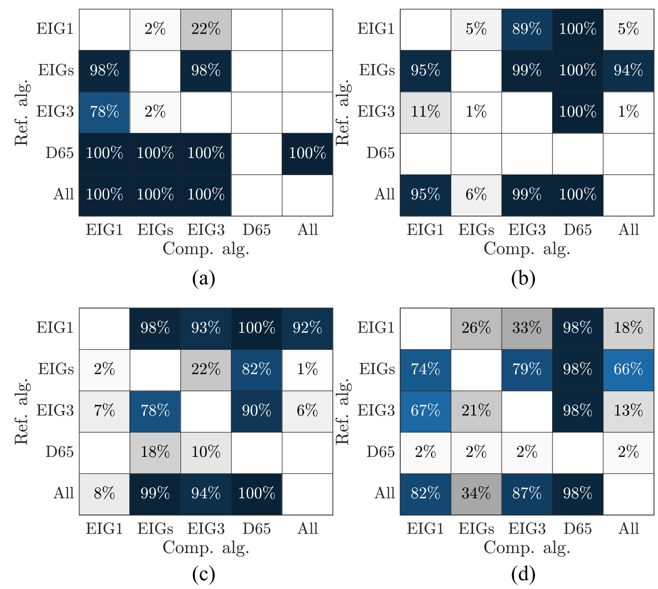

Fig. 4. Pairwise comparison of win percentages over 4851 tests. Each cell shows the percentage of cases where the row algorithm outperformed the column algorithm (ties are excluded) for a given metric. (a) d_{abs}, (b) d_{led}, (c) d_{rgb}, and (d) d_{sam}. The "All" row summarizes the win rate of an algorithm against all others.

III. EXPERIMENTAL ANALYSIS AND DISCUSSION

The CC SG generates 4851 unique pairwise combinations of color patches. These pairs were evaluated computationally, as physical acquisition is resource-intensive. This model's validity was established in [1], where its predictions matched 26 740 physical measurements (140 CC patches under 191 illuminants), with an $R^2 = 0.98$. EIG1, EIGs, and EIG3 provide the analytical solutions for the optimum illumination spectrum $l = L\beta$. We then calculate the pixel intensities generated as a result of l. As illustrated in Fig. 3, these pixel intensities are subsequently used as input for the contrast functions to compute their respective contrast values.

A. Metric-Dependent Algorithm Performance

For a comparative study, such as this, relative performance is more insightful than absolute contrast values. Therefore, we use pairwise comparison matrices, as shown in Fig. 4, to summarize the "win rate" of each algorithm under the four contrast definitions.

The results clearly demonstrate that the choice of contrast metric directly determines which optimization strategy performs best. As shown in Fig. 4(a), the D65 illuminant overwhelmingly outperforms the other algorithms when evaluated by absolute contrast (d_{abs}). Conversely, D65 is the worst performing illuminant under all other metrics. The dominant algorithm shifts depending on the objective: EIGs excels under illumination-normalized (d_{led}) and spectral angle (d_{sam}) metrics, while EIG1 is most effective for the Michelson contrast (d_{rgb}). This provides strong quantitative evidence that the "best" optimization strategy is conditional on the mathematical definition of contrast.

B. Physical Manifestation of the Objective Function

The numerical contrast differences result from the different physical spectra each algorithm produces. Thus, a changed objective function leads to a measurable change in the physical light. This demonstrates a core computational imaging principle: computation directly controls physical reality.

We can explore the dominance of D65 under d_{abs} by examining the relationship between brightness and measured intensity. As shown in Fig. 5(a) and (b), imaging a white patch with increasing LED brightness results in a linear increase in pixel intensity. This indicates that any metric not normalized for brightness, such as d_{abs}, will inherently favor the most powerful illumination, regardless of its spectral shape. Because the D65 illuminant is broadband, it contains more total power than the spectrally optimized light, and therefore consistently yields the highest absolute contrast. For example, the EIG1 algorithm, which outperforms D65 under the d_{led}, d_{rgb}, and d_{sam} metrics, does so while requiring on average 71% less illumination power ($|l|$) than D65.

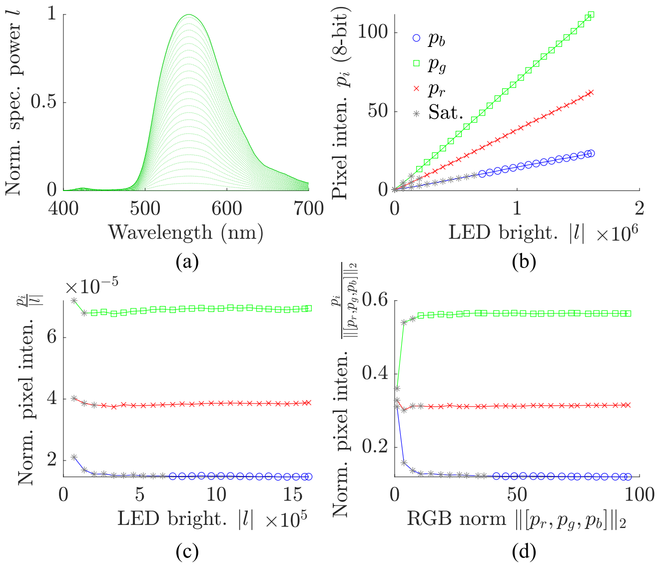

Fig. 5. Effect of illumination power on pixel intensity. (a) Normalized spectral power of an LED l at different brightness levels. (b) Pixel intensity increases linearly with $|l|$ within the sensor's dynamic range; values beyond this range (10–230; 8 bit) are discarded [1]. (c) Normalizing intensity by illumination power and (d) by the RGB norm removes dependence on $|l|$.

Conversely, normalization, as shown in Fig. 5(c) and (d), removes the effect of overall brightness and provides a more meaningful measure of contrast based on spectral efficiency. This also emphasizes that energy efficiency is not a secondary consideration, but is inherently embedded in the choice of contrast metric.

The term "accurate," however, depends on the application's goal. If the objective is to achieve the highest possible signal difference, then d_{abs} is a sufficient metric, implying a physical design that prioritizes high total power and thermal management. However, if the goal is to optimize the contrast-to-power ratio or selectively enhance color separability, then normalized metrics are more suitable. These task-centric metrics would drive a physical design focused on spectral precision and diversity—such as including more, carefully selected narrow-band channels—rather than on sheer brightness. This choice, therefore, directly influences the hardware requirements of the illumination system.

C. Discussion on a Machine Learning Perspective on Contrast

The brightness vector β that generates the optimal illumination spectrum is analogous to the learnable "weights" of a model, and the chosen contrast metric is its "loss function." Our results provide a physical demonstration of a principle well known in machine learning: the choice of loss function leads to different optimal models. Here, we demonstrate that the choice of a mathematical objective leads to different optimal physical systems.

This perspective raises new questions: If the contrast metric is a design choice, then could a more complex, multiobjective metric be learned? Our EIG3 algorithm provides a simple proof-of-concept for this approach. By combining the objectives for all three color channels ($H = H_r + H_g + H_b$), it seeks a single, balanced illumination spectrum rather than a separate, optimal one for each. Instead of a human preselecting from candidates, such as those in Table 1, could we develop an end-to-end differentiable system where the objective function itself is a learnable hyperparameter? For example, a neural network could be designed where the final layers perform a specific task (e.g., classification), while the initial layers learn not only the optimal illumination, but also the ideal coefficients (α_i) for a composite loss function, such as a weighted sum of the metrics we tested ($d_{\text{total}} = \alpha_1 \, f_1(d_{\text{abs}}) + \alpha_2 \, f_2(d_{\text{led}}) + \alpha_3 \, f_3(d_{\text{rgb}}) + \alpha_4 \, f_4(d_{\text{sam}})$), which best serves the end-task. Such a system would represent a truly adaptive sensor, capable of learning not only how to see, but what it means to see well for a given purpose.

IV. CONCLUSION

We demonstrated that the optimal illumination spectrum is a physical property of an imaging system, directly linked to the mathematical objective function used to derive it. This refutes the notion of a single, universally "best" illumination and establishes the contrast metric as an active, tunable design parameter.

Limitations of this study include the focus on static scenes and predefined, rather than learned, metrics. Future work will explore extending this framework to dynamic environments and developing task-driven, adaptive contrast metrics that can be optimized end-to-end with downstream machine learning models, advancing the design of intelligent, self-configuring sensor systems.

ACKNOWLEDGMENT

The authors would like to thank New Zealand Royal Society Te Apārangi for the Rutherford Discovery Fellowship to Ooi.

REFERENCES

[1] S. K. Abeysekera, M. P.-L. Ooi, Y.-C. Kuang, and V. Kalavally, "Spectral optimization of illumination for contrast enhancement in machine vision," *Opt. Lasers Eng.*, vol. 195, 2025, Art. no. 109257.

[2] W. Zhou, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.

[3] A. J. Mohler, M. McGeehan, K. G. Ong, and M. Hahn, "Evaluation of a multiaxial optical-based shear sensor using a multilayer perceptron artificial neural network model," *IEEE Sens. Lett.*, vol. 9, no. 5, pp. 1–4, May 2025.

[4] T. Matsumoto, L. Du, F. Rodino, Y. Thoma, C. Premachandra, and S. Carrara, "Optimized quantification of multiple drug concentrations by weightedMSE with machine learning on electrochemical sensor," *IEEE Sens. Lett.*, vol. 8, no. 10, pp. 1–4, Oct. 2024.

[5] S. K. Abeysekera, Y.-C. Kuang, and M. P.-L. Ooi, "Comprehensive modeling of camera spectral and color behavior," in *Proc. 2025 IEEE Int. Instrum. Meas. Technol. Conf.*, 2025, pp. 1–6.

[6] M.-H. Lee, D.-K. Seo, B.-K. Seo, and J.-I. Park, "Optimal illumination for discriminating objects with different spectra," *Opt. Lett.*, vol. 34, no. 17, pp. 2664–2666, Sep. 2009.

[7] E. Peli, "Contrast in complex images," *J. Opt. Soc. Am A*, vol. 7, no. 10, pp. 2032–2040, Oct. 1990.

[8] A. Erickson, G. Bruder, and G. F. Welch, "Adapting Michelson contrast for use with optical see-through displays," in *Proc. 2022 IEEE Int. Symp. Mixed Augmented Reality Adjunct*, 2022, pp. 409–410.

[9] F. Kruse et al., "The spectral image processing system (SIPS)—interactive visualization and analysis of imaging spectrometer data," *Remote Sens. Environ.*, vol. 44, no. 2, pp. 145–163, Jun. 1993.