Fuzzy-logic-based Relevant Color Extractor

J. Daniel Subias[1*], Ramon Fernandez-Gualda[2], Samuel Morillas[3] & Juan Luis Nieves[2]

[1]Universidad de Zaragoza, I3A, Spain
[2]University Institute of Pure and Applied Mathematics, Universitat Politècnica de València, Spain
[3]Universidad de Granada, Spain
*dsubias@unizar.es

Figure 1: *A fuzzy-logic-driven palette-extraction framework identifies the "relevant colors" present in an artwork. From a single input image of a painting (top), the method generates high-quality reconstructed versions that maintain strong perceptual fidelity to the original by emulating key aspects of human color perception (bottom). The colored circles appearing inside the gray boxes in the lower-right corner represent the palettes of "relevant colors" extracted using our method.*

Abstract

Color palette extraction is fundamental in art analysis, digital curation, and graphic design. This paper presents a fuzzy-logic-driven method for automatically deriving non-fixed color palettes from paintings. Drawing inspiration from the concept of "relevant colors" (the chromatic elements most salient to human observers), the approach employs fuzzy logic to characterize chromatic diversity using quantized color representations. This framework is particularly valuable in digital art preservation, restoration, and collection management, where accurately identifying the essential color gamut of an artwork is critical. The algorithm models perceptual color relevance by integrating luminance, chroma, and local color redundancy as input variables. A fuzzy inference system subsequently assigns relevance scores using linguistic rules, enabling adaptable palette generation tailored to each painting. Evaluated across a large dataset of artworks, the proposed method surpasses state-of-the-art techniques in image reconstruction tasks grounded in palette extraction. Furthermore, when assessed against human annotations from the Prado Museum dataset, the resulting palettes exhibit strong alignment with observer preferences and faithfully reflect the perceptual color structure of the paintings.

CCS Concepts
• ***Computing methodologies*** → *Image compression; Image processing;* • ***Applied computing*** → *Media arts;*

1. Introduction

Research in color perception aims to elucidate how the human visual system interprets the spatial and chromatic structure of the light entering the eye. The retina exhibits a remarkable capacity to discriminate a vast range of colors through its three classes of cone photoreceptors: L-, M-, and S-cones; which provide the initial stage of trichromatic encoding. These signals are subsequently transformed and integrated by post-receptoral mechanisms and higher-level visual pathways, ultimately giving rise to the rich perceptual experience of color. While the exact number of discernible colors (NDC) that normal trichromats can perceive is a matter of ongoing debate and research, estimates suggest that it falls within the range of millions number of discernible colors at an early stage of vision [PA98, LPN08, TCT13]. This extensive gamut of colors is fundamental to the art of painting, allowing artists to evoke emotions, create depth, and bring their visions to life. Understanding the breadth of discernible colors enhances our comprehension of visual perception and lays the foundation for exploring how this extensive gamut is distilled into a manageable subset of "relevant colors". These "relevant colors", rather than the full spectrum of discernible colors, are crucial for the development of tools and technologies, such as color palette extractors, which aid artists in selecting and utilizing a representative palette that captures the essence of an image while maintaining its perceptual impact. In parallel, color naming and categorization have been extensively studied in recent years [WG18] due to their direct applications in computational color and object classification using machine-learning-related approaches [SRSS22, FFM*26].

Despite the vast number of discernible colors, an observer cannot differentiate or retain such a large quantity of discernible colors when viewing a natural or artificial scene. The initial number of millions of possible signals from photoreceptors at the retinal level is drastically reduced when the information leaves the retina, resulting in only about one million ganglion cells that finally send signals to the visual cortex. In this bottleneck context, the concept of "relevant color" has been introduced to characterize the colors that an observer would highlight in an image, regardless of whether they are colorimetrically discernible or not [NGRCR20]. Relevant colors provide a simplified way to quantify the information that a painting contains, reducing the spatial and spectral complexity, obtaining a smaller discrete set that still captures the majority of the significant colorimetric information conveyed by the image. The image processing method that computationally reduces the number of colors to a set of "relevant colors" is based on color quantization. This process can be summarized into two phases [CPD23]: the design of the color palette and the pixel mapping. The first phase focuses on creating a color palette representing a small subset of the input colors, and the second phase assigns each pixel in the input image to one of the palette colors. Color quantization algorithms can be classified into two categories: image-independent algorithms and image-dependent algorithms. Image-independent algorithms create a universal color palette that is not influenced by the input image, while image-dependent algorithms design a custom color palette based on the input image's color distribution. Modern applications of color quantization include image dehazing, image compression, color-to-gray-scale conversion, image watermarking, image segmentation, saliency detection, skin detection, and posterization. These color quantization algorithms aim to obtain a set of representative colors from the image, retaining essential visual information [Cel23].

Different computer vision algorithms have been used to extract colors, with clustering methods such as k-means and fuzzy logic being common techniques [CSG21, CFL*15, HDI*23, AIP20]. On the other hand, color plays a pivotal role in art creation [GF07]. Numerous studies have established a direct link between color perception and emotion [WO18, GFL16, KE04]. For instance, primary hues (e.g., red or yellow) elicit stronger positive emotions compared to intermediate hues (e.g., yellow-red or blue-green hues) and achromatic colors (e.g., black and white) [LMS24]. Research by Dael et al. [DPM*16] suggests that intrinsic attributes of color, such as saturation, hue, and brightness, significantly influence affective processing. Each painter possesses a distinctive color gamut, and certain colors are often linked to specific painting styles. However, while color is a crucial element, it is not sufficient by itself to automatically categorize a painting style [WFC*09, KB-VdWF14].

For colorimetric characterization, spatial parameters such as fractal dimension, power spectrum, entropy and complexity have been studied in the computational analysis of paintings. However, Mureika [Mur05] asserted in a study that the best system for examining the representation of color in paintings is the CIELAB color space [cie18]. CIELAB is designed to be nearly perceptually uniform, meaning that the Euclidean distances in the space correlate well with visual color differences, making it particularly well-suited for accurately capturing and analyzing the colors used in artworks.

There are also the so-called "color-naming algorithms", which aim to discretize color characterization in an image based on the concept of "basic color themes" [ST20]. This approach has a notable difference compared to our concept of "relevant colors". The key difference lies not only in the absolute number of colors but also from the method used to determine the number of relevant colors (NRC). These colors are selected from a fixed palette of basic color names, independent of the specific image being analyzed.

This study proposes a new algorithm to estimate a reliable color palette in paintings based on the novel notion of "relevant colors", those most likely to be prioritized by a human observer [NGRCR20] (Figure 1). For this, we will use fuzzy logic because it allows for implementing human knowledge expressed in linguistic rules into a computational method in a quite simple way. In summary, we present the following contributions:

- A set of linguistic rules that model how humans determine a color as relevant.
- A novel fuzzy logic system that outperforms convex-hull-based methods and achieves comparable performance to the algorithm of Nieves et al. [NGRCR20].
- A metric to measure the similarity between the colors selected by observers and color palettes estimated by color palette extraction algorithms.

Our code is available at https://graphics.unizar.es/projects/fuzzy_palette_extractor_2026/

2. Related Work

There are few studies that deal with the colorimetric characterization of "relevant colors", even though color composition is a critical aspect of visual applications in art, graphic design, and visualization. Perhaps, the so-called "color theme" is the most similar concept, which best represents an image's color distribution, is extracted by algorithms (such as convex hulls or k-means clustering) that limit the number of colors, typically ranging between three and seven palette colors [CSG21, CFL*15, LH13, WLX19, TEG18, CNS19]. Recently, Weingerl et al. [WHJ20] propose a machine-learning learning framework to extract the most salient color theme on natural images by using LASSO regression (Least Absolute Shrinkage and Selection Operator), suggesting that the most influential factors on the perception are associated with color coverage, color lightness and chroma, and diversity of colors.

Early color quantization algorithms attempted to design a universal palette without regard to the input image, through independent uniform scalar quantization to each color component independently [Gen90], which was initially implemented using bit-cutting. However, uniform quantization is costly to implement and less effective than non-uniform quantization. The latter is equivalent to the k-means algorithm in one dimension and produces less distortion than a uniform scalar quantizer. Regardless of this, a universal palette often yields inferior results compared to a custom palette. Consequently, researchers turned to image-dependent algorithms in the early 1980s. Among image-dependent algorithms are hierarchical algorithms that recursively find nested clusters in a top-down (divisive) or bottom-up (agglomerative) manner. Studies like Celebi et al. [CWH15] and Ueda et al. [UKSU17] propose divisive algorithms. The former splits the cluster with the greatest sum of squared errors along the color axis with the greatest variance at the mean point, and the second splits the cluster with the greatest product of its dominant eigenvalue and size. On the side of agglomerative algorithms, the three most prominent methods are the octree [GP88], the pairwise nearest neighbor algorithm [Equ89] and the Balanced Iterative Reducing and Clustering using Hierarchies (Birch) [ZRL97] algorithms. However, hierarchical algorithms present disadvantages: their splits are irreversible, leading to a cumulative effect of erroneous decisions; and many of them do not optimize an objective function. Therefore, some of these algorithms were later used to initialize partitional algorithms (k-means or fuzzy logic algorithms).

Computer vision algorithms have attempted to extract the colors that describe an image. The most commonly approaches have been based on clustering techniques such as k-means and fuzzy logic. Charlin and Cifuentes [CC21] use a k-means clustering algorithm to select color palettes, limited to three colors, which were manually configured beforehand. They aimed to study the relationship between the price and color of paintings, along with other spatial parameters such as dominant colors, color harmony, and color emotions. On the other hand, Chang et al. [CFL*15] propose a deterministic variant of the k-means algorithm for color palette extraction, dividing RGB colors into a fixed number of three-dimensional bins so that the computation is constant and invariant to image resolution. Muratbekoya and Shamoi [MS24] employed fuzzy sets to classify emotions in paintings. First, they introduce the fuzzy color representation model; then, at the fuzzification stage, they process the paintings tagged with emotions, extracting fuzzy dominant colors linked to specific emotions, revealing specific associations between emotions and colors (gratitude with green, brown and orange; brown with anger, orange with shame, yellow with happiness and gray with fear). The color palette obtained by the algorithm processing each dataset image contains 15 fuzzy colors maximum. Generally, these algorithms are simple to implement but require some manual intervention to pre-estimate the number of clusters or colors from the start. Other online color palette extractor tools, such as Colormind [Col19], use classifiers trained on a hand-picked list of examples, resulting in the disadvantage of a fixed and limited number of colors.

Another color quantization method involves histogram-storage, repeating colors multiple times despite being a simple and fast implementation. Each bin of such a histogram stores the frequency of the three color components, and with this data structure, it only requires accessing each bin once to analyze the data. In the case of 3D representations, it may require an unnecessary amount of memory if the colors occupy only a small part of the color space, as is often the case with natural images [Cel23].

Additionally, public image datasets have been published, such as the one by Celebi and Pérez-Delgado [CPD23], which consists of 100 high-quality 24-bit color images aimed at evaluating color quantization algorithms. Classification algorithms, known as "color-naming algorithms", attempt to map pixels to color names as similar as human would, using machine learning algorithms such as Support Vector Machine, Bayes or k-Nearest Neighbor; or deep neural networks with color spaces as input features: RGB, HSV, Lab, or CYMK. This helps to ensure the relationship between content and colors by constraining color names. However, the range of color names they offer is between 11 and 15 for all linguistic color categories [ST20, XVCH*24].

In an earlier work, Nieves et al. [NGRCR20] developed an algorithm for the characterization of the chromatic range of a set of paintings. The algorithm works by converting the RGB values of a painting into the perceptually uniform CIELAB color space, then dividing this space into cubes of 20 units in each dimension to identify clusters of significant colors. A threshold is applied to determine relevance, ensuring only those clusters representing a notable percentage of pixels are considered. Additional criteria account for high chroma and brightness, refining the selection of perceptually relevant colors. It obtained an average NRC of 18. This surpassed the average number of basic color names found in other categorical color studies. Continuing on this path without requiring predefined color categories, this study aims to develop a more reliable algorithm based on fuzzy logic.

3. Methodology

Fuzzy logic was introduced by Zadeh [Zad65] to mimic human thinking and the decision process through reasoning, when dealing with uncertain data. Fuzzy logic provides a theoretical framework to derive meaningful conclusions even from imprecise or incomplete knowledge. This remarkable versatility has led to its widespread adoption across various scientific and technolog-

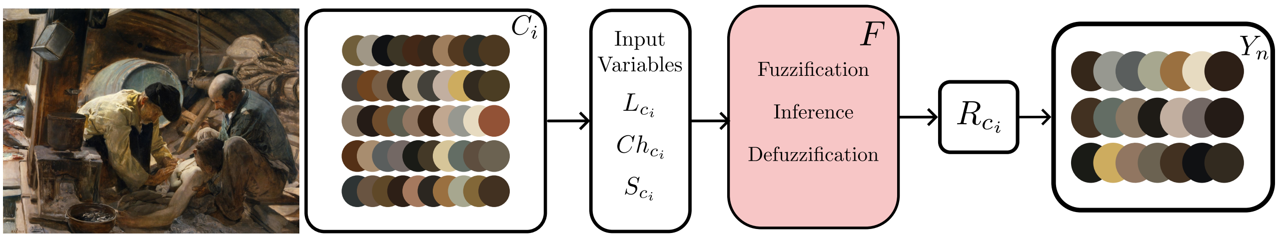

Figure 2: *High-level overview of the proposed system. The fuzzy logic system F receives, for each color in the extracted subset C_i, three input variables: luminance L_{c_i}, chroma Ch_{c_i} and number of similar colors S_{c_i}. Through the fuzzification, inference, and defuzzification, the system computes the set of relevance R_{c_i} for every color. The most relevant top-k colors are then selected to form the output palette Y_n. A final classification step, based on Fisher's linear discriminant [Fis36], separates "relevant" from "non-relevant" colors.*

ical fields, such as image processing tasks [SWK07, SDWN*07, VDVNVdW*03].

In this section, we introduce our fuzzy logic system to extract a color palette from a given image of a painting. In Section 3.1 we introduce our goal and provide an overview of our proposed fuzzy logic system. Then, Section 3.2 describes the initialization of our fuzzy logic system to compute the values of the input variables from the colors that compose the input image. Section 3.3 describes the fuzzification process to transform the input variables into fuzzy linguistic terms by applying the membership functions of our fuzzy logic system. Then, in Section 3.4, we detail the inference process of our fuzzy system, based on the set of rules of our system. Finally, Section 3.5 describes the defuzzification process to obtain a color palette, based on the fuzzy relevance values measured by our fuzzy system for the colors composing the input image.

3.1. Goal and Overview

Our goal is to generate a color palette Y_n given a colored painting. The output color palette Y_n is composed of N "relevant colors", so that $n \in \{0,...,N-1\}$ that would stand out for an observer when just glancing at the input painting. To achieve this goal, we propose a fuzzy logic system F that measures the fuzzy relevance R_{c_i} for each color in a subset of 50 colors C_i of the input painting, where $i \in \{0,...,49\}$. The colors with the highest relevance make up the final color palette Y. For this purpose, we automatically classify the colors into "non-relevant colors" (rejected) and "relevant colors" (retrieved) by using the Fisher's linear discriminant [Fis36], which can be applied recursively. Our fuzzy logic system F takes as input variables: the luminance L_{c_i}, chroma Ch_{c_i}, and number of similar colors S_{c_i} of each color in C_i. A high-level overview of our fuzzy logic system F is shown in Figure 2.

3.2. Fuzzy System Initialization

To obtain the initial set of colors C_i we compute k-means in the CIELAB color space, to extract a set colors from the pixel colors of the input painting, similar to the works of Nieves et al. [NGRCR20], Chang et al. [CFL*15] and Chao et al. [CSG21]. Thus, the color set C_i is comprised of the centroids given by k-means. We have set the number of centroids to 50, taking into account that no observer exceeds the 41 NRC selected for a painting from the Prado Museum collection [Mus24] (Section 4.1), thereby providing a higher confidence interval. For each color in C_i we compute the input variables, i.e., the luminance L_{c_i}, chroma Ch_{c_i} and number of similar colors S_{c_i}. The luminance channel L^* in the CIELAB color space corresponds with the input variable luminance L_{c_i}. On the other hand, the input variable chroma Ch_{c_i} is defined as $\sqrt{a^{*2}+b^{*2}}$, where a^* and b^* are the green-red and blue-yellow components in the CIELAB color space respectively. Finally, the number of similar colors S_{c_i} corresponds with the percentage of pixels of the input painting associated with each centroid in C_i. We normalize the luminance L_{c_i} and chroma Ch_{c_i} to be in the range $[0,1]$, while the number of similar colors S_{c_i} is in the range $[0,100]$. The set of fuzzy rules of our fuzzy logic system F is defined on linguistic terms considering linguistic variables of both inputs (i.e., luminance, chroma and number of similar colors) and outputs (i.e., relevance). A fuzzy rule is a simple IF-THEN rule with an antecedent and a consequent [Men95]. Considering the IF–THEN structure, our fuzzy logic system F determines the relevance R_{c_i} of each color through expert evaluation rules based on the outputs of the membership functions. For our fuzzy logic system F, we define the following rules:

- Low relevance (R_{c_i}):
  1. **IF** few S_{c_i} **AND** low L_{c_i} **AND** low Ch_{c_i} **THEN** low R_{c_i}
  2. **IF** few S_{c_i} **AND** low L_{c_i} **AND** very low Ch_{c_i} **THEN** low R_{c_i}
  3. **IF** few S_{c_i} **AND** very low L_{c_i} **AND** very low Ch_{c_i} **THEN** low R_{c_i}
  4. **IF** few S_{c_i} **AND** very low L_{c_i} **AND** low Ch_{c_i} **THEN** low R_{c_i}
  5. **IF** Rule 1 **OR** Rule 2 **OR** Rule 3 **OR** Rule 4 **THEN** low R_{c_i}
- Medium relevance (R_{c_i}):
  1. **IF** few S_{c_i} **AND** medium L_{c_i} **AND** (medium Ch_{c_i} **OR** high Ch_{c_i}) **THEN** medium R_{c_i}
  2. **IF** low Ch_{c_i} **AND** many S_{c_i} **THEN** medium R_{c_i}
  3. **IF** Rule 1 **OR** Rule 2 **THEN** medium R_{c_i}
- High relevance (R_{c_i}):
  1. **IF** many S_{c_i} **AND** low L_{c_i} **THEN** high R_{c_i}
  2. **IF** many S_{c_i} **AND** very low L_{c_i} **THEN** high R_{c_i}
  3. **IF** many S_{c_i} **AND** medium L_{c_i} **THEN** high R_{c_i}
  4. **IF** (many S_{c_i} **OR** few S_{c_i}) **AND** high L_{c_i} **THEN** high R_{c_i}
  5. **IF** (many S_{c_i} **OR** few S_{c_i}) **AND** very high L_{c_i} **THEN** high R_{c_i}
  6. **IF** (many S_{c_i} **OR** few S_{c_i}) **AND** high Ch_{c_i} **THEN** high R_{c_i}

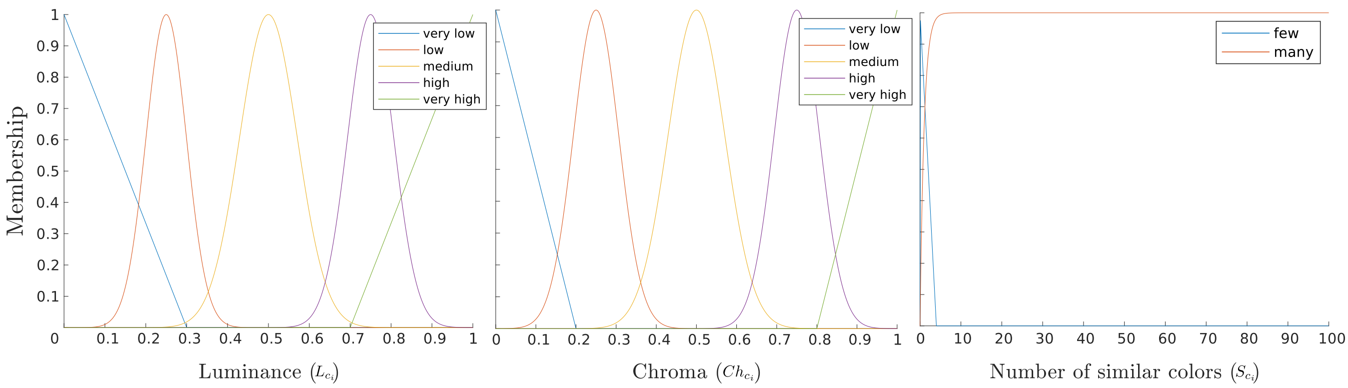

Figure 3: *Fuzzy membership functions for the different levels of the input variables. The membership functions output the degree of belonging, in the range $[0,1]$, of each input variable to each of the levels. The luminance L_{c_i} and chroma Ch_{c_i} are in the range $[0,1]$ and are fuzzified into five levels: very low, low, medium, high and very high (left and center). The number of similar colors S_{c_i} is in the range $[0,100]$ and fuzzified into two levels: few and many (right).*

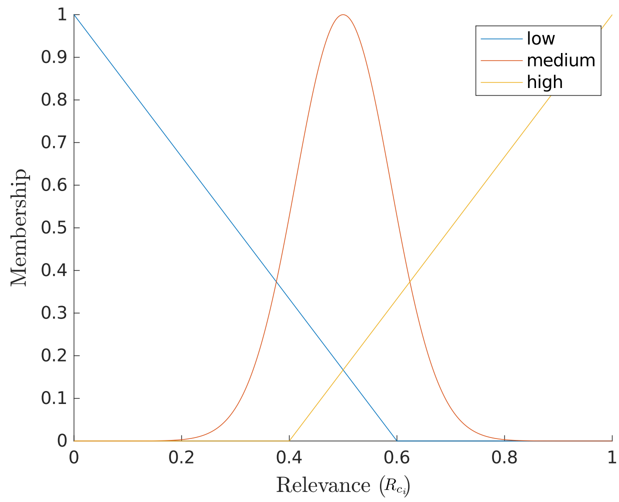

Figure 4: *Fuzzy membership functions defining the levels of color relevance R_{c_i}. Each membership function returns a degree of belonging in the range $[0,1]$ for the corresponding relevance level. The relevance value R_{c_i}, also bounded in $[0,1]$, is fuzzified into three linguistic categories: low, medium, and high.*

7. **IF** (many S_{c_i} **OR** few S_{c_i}) **AND** very high Ch_{c_i} **THEN** high R_{c_i}
8. **IF** Rule 1 **OR** Rule 2 **OR** Rule 3 **OR** Rule 4 **OR** Rule 5 **OR** Rule 6 **THEN** high R_{c_i}

Where the conjunction (AND) and disjunction (OR) operations are applied to compute the certainty of the whole antecedent rule. For that purpose, an appropriate t-norm is used to represent the AND operation and an s-norm for the OR operation. We use the classical product as t-norm for the AND operation, and the $max(\cdot,\cdot)$ s-norm for the OR operation.

Figure 5: *Subset of paintings used to evaluate our fuzzy logic system. The images show four paintings (the first two belong to Painting-91 dataset and the last two to the Prado dataset) present in the evaluation dataset. From left to right: "Rainy Night on Bridlington Promenade", "Marilyn", "Vase of Flowers (sixteenth century)" and "Mancorbo Canal in Picos de Europa".*

3.3. Fuzzification

In the fuzzification process, we transform the input variables (i.e., luminance, chroma and number of similar colors) into linguistic terms through the application of membership functions. In our fuzzy logic system F, we fuzzified the luminance L_{c_i} and chroma Ch_{c_i} into five levels: very low, low, medium, high and very high. Concerning the number of similar colors S_{c_i}, we fuzzified this input variable into two levels: few and many. The membership functions output the degree of membership of each input variable to each of its levels, in the range $[0,1]$. As we can see in Figure 3 (left and center), for the luminance L_{c_i} and chroma Ch_{c_i}, we adopt the popular triangular fuzzy membership functions for the extreme levels; and exponential membership functions for the intermediate levels. Concerning the number of similar colors S_{c_i}, we use the Pareto distribution as a membership function for the many level and a triangular fuzzy membership function for the few level, as we can see in Figure 3 (right).

3.4. Inference

After the fuzzification process, by applying a set of rules, we transform the outputs of the membership functions into the output rele-

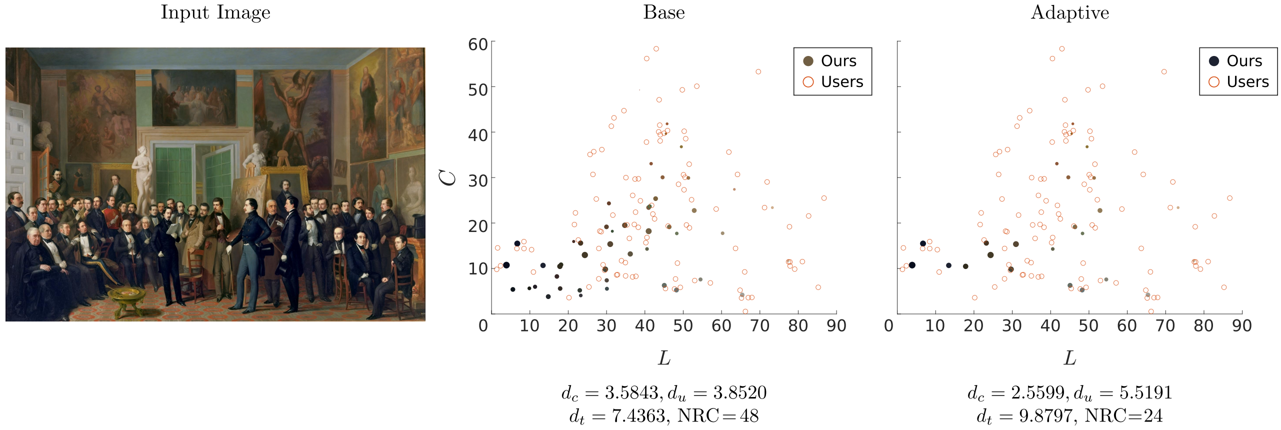

Figure 6: *Distribution of our color palette Y_n (color dots) and the color set U_m selected by the observers (orange dots) for the input painting "Contemporary Poets" in the L–C (luminance–chroma) plane. The size of the color dots denotes the number of similar colors S_{c_i} of each color, a larger size implies a larger number of similar colors. We can see how our color palettes follow a similar distribution regarding to the set of colors selected by the observers.*

vance R_{c_i}. Linguistically, our fuzzy logic system F considers three levels of relevance: low, medium and high. As we can see in Figure 4, for the low and high levels we use the popular triangular fuzzy membership functions, while we use an exponential fuzzy membership function for the medium level.

The inference process is the mathematical operation used to determine the relevance R_{c_i} of a color, given the degree of membership of each input variable to each of its levels. The certainty of the consequent rule is assigned to the certainty of the antecedent computed in the previous step. This certainty represents the degree to which the relevance R_{c_i} can be considered to be at the level corresponding to the rule. The overall result is a fuzzy certainty value that measures the relevance R_{c_i} at each of the three defined levels (i.e., low, medium, high).

3.5. Defuzzification

Finally, in the defuzzification process, the relevance measures for each level (i.e., low, medium and high) must be defuzzified to obtain a final output relevance R_{c_i} for each color in C_i. We perform the defuzzification according to the membership functions (Figure 4) of the output relevance R_{c_i}. Different methods can be used for defuzzification. In our fuzzy logic system F, the Center of Gravity method is adopted [Ker99]. To obtain a numerical value of relevance R_{c_i} from the qualitative evaluation given by the fuzzy rules, an area is formed by the membership function where each function is bounded by the corresponding degree of certainty of the output level (i.e., low, medium, high) that it represents. The center of gravity of this area is then computed and its projection on the relevance R_{c_i} axis provides the quantitative level of R_{c_i} of the input color. Finally, to obtain the colors with the highest relevance, we use the Fisher's linear discriminant [Fis36] recursively to divide the set of colors C_i into "non-relevant colors" and "relevant colors". The latter are the colors that make up the final color palette Y_n.

4. Results

We evaluate the fuzzy-logic framework F on an extensive collection of paintings and benchmark its performance in image-reconstruction tasks against color sets identified as relevant by human observers. Section 4.1 describes the evaluation dataset and the quantitative metrics used to assess the method. Section 4.2 analyzes the effect of applying Fisher's linear discriminant [Fis36] recursively in the reconstruction pipeline. Section 4.3 demonstrates the robustness of the approach through comparisons with observer-selected palettes. Finally, Section 4.4 presents both qualitative and quantitative comparisons with state-of-the-art palette-extraction and reconstruction methods.

4.1. Evaluation Dataset and Metrics

To assess the performance of the proposed method in image-reconstruction tasks, we first use the Painting-91 dataset [KBVdWF14], which contains 4,266 paintings produced by 91 artists. During preprocessing, a small number of images were discarded because their RGB distributions were nearly degenerate (e.g., near-grayscale or strongly clustered), which led to systematic failures in the convex-hull simplification and mesh-reduction stages. After removing these problematic cases, 4,243 paintings remained for the experiments. To evaluate the degree of alignment between the estimated palettes and human color perception, we also rely on the annotations collected by Nieves et al. [NOGRR21] for 20 paintings from the Prado Museum collection [Mus24]. In this dataset, six observers with normal color vision (aged 30–60; three women and three men) were asked to freely select a non-fixed set of NRC for each painting. Two participants were naïve and unaware of the study's purpose. Each painting is therefore associated with a set of observer-selected colors U_m, for $m \in \{0,\ldots,M-1\}$, where M denotes the total number of colors chosen across observers. For our study, we excluded one observer's annotation for the painting number 20 (*"The Triumph of Death"*), due to the anomalous NRC selection

(68 colors), given that the average across users and paintings is 15.3250 ± 9.0304. For clarity, we refer to this annotated collection as the Prado dataset throughout the paper; see Figure 5 where representative examples from both datasets are shown.

To quantitatively evaluate our color palettes, we rely on four widely used image quality metrics to measure the similarity between the input painting and the reconstructed image using our color palette Y_n: Mean Squared Error (MSE), Mean Absolute Error (MAE), Pixel-to-Signal Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM). On the other hand, we also measure the similarity d_t between our color palette Y_n returned by our fuzzy logic system F and the color set U_m selected by the observers on our evaluation dataset. To achieve this, we compute the average distance of the min distance for each color in Y_n to each color in U_m defined as:

$$d_c = \frac{1}{N} \sum_{n=0}^{n=N-1} \arg\min_{m} \Delta E_{00}^{*}(Y_n, U_m). \tag{1}$$

Where ΔE_{00}^{*} is the CIE Delta E_{00} Color Difference [SWD05] and N is the total NRC in Y_n. This distance d_c measures the precision of the fuzzy logic system, decreasing in the cases where the system correctly identifies a relevant color close to a color chosen by the observers (true positive instance), and increasing when a relevant color is not close to a color selected by a observer (false positive instance). On the other hand, we also compute the average distance of the min distance for each color in U_m to each color in Y_n as follows:

$$d_u = \frac{1}{M} \sum_{m=0}^{m=M-1} \arg\min_{n} \Delta E_{00}^{*}(U_m, Y_n). \tag{2}$$

Where M is the total number of colors in U_m. This distance d_u evaluates the recall of the system, the distance decreases when a relevant color is close to a set of colors selected by the observers (relevant instance), while it increases when the relevant color is far away of colors chosen by the observers (irrelevant instance). Therefore, the similarity d_t between a color palette and the colors chosen by the observers is defined as:

$$d_t = d_c + d_u. \tag{3}$$

4.2. Quantization Reconstruction

We assess the performance of different versions of our fuzzy logic system to extract consistent color palettes: our base version applying the Fisher's linear discriminant [Fis36] once to classify the colors into "non-relevant colors" and "relevant colors" (Base), applying Fisher's linear discriminant twice (Recursive) for further reduction of the NRC in Y_n and applying the recursion (with a maximum limit of five iterations) only when the NRC in the color palette Y_n is higher than 30 (Adaptive), to avoid a drastic color reduction. Once the relevant colors are determined, each pixel in the image is assigned to the closest relevant color based on the Euclidean distance in the CIELAB color space. In Table 1, we show the average image quality metrics (i.e., MSE, MAE, PSNR, and SSIM) on the evaluation dataset (both the Painting-91 and Prado datasets) for the different versions of our method (i.e., Base, Recursive, and Adaptive). The average NRC results for the Base/Recursive/Adaptive versions are 38.3438/20.9397/19.5019, respectively. We see that the Adaptive and Recursive versions achieve a similar performance, and both have lower performance than the base version, since the latter returns a high NRC.

Table 1: *Average MSE, MAE, PSNR, and SSIM reconstructing the input images on the evaluation dataset. We can see how our fuzzy logic system achieves better performance when applying the Fisher's linear discriminant [Fis36] once (Base) because it returns a greater NRC in comparison to applying Fisher's linear discriminant twice (Recursive) or only when the NRC is lower than 30 (Adaptive).*

Version	MSE ↓	MAE ↓	PSNR (dB) ↑	SSIM ↑
Base	**0.0015**	**0.0773**	**29.3951**	**0.9339**
	± 0.0017	± 0.0311	± 3.1867	± 0.0400
Recursive	0.0060	0.1388	24.4757	0.8613
	± 0.0110	± 0.0889	± 4.2171	± 0.0927
Adaptive	0.0058	0.1379	24.3442	0.8650
	± 0.0119	± 0.0881	± 3.7812	± 0.0813

4.3. Consistency of our Color Palettes

We evaluate the robustness of our fuzzy logic system F by comparing the distribution of the color palette Y_n with the set of colors U_m chosen by the observers on the Prado dataset. The average similarity d_t results on the Prado dataset for the Base/Recursive/Adaptive versions are 7.6141/10.8690/9.1898, respectively. Given that the Base and Adaptive versions are the ones that present results most aligned with the observers, hereinafter we focus on these versions of our fuzzy logic system. Figure 6 shows an example of how the color distribution of our color palette follows a similar distribution concerning the color set U_m selected by the observers within the space defined by the L–C (luminance–chroma) plane. As expected, the color palette Y_n returned by the base version of our method (Base) fits better with the distribution of colors selected by the observers due to the greater NRC. Therefore, when applying the Adaptive versions, the recall (i.e., distance d_u) increases since the system removes "relevant colors" that are close to the colors chosen by the observers, leading to an overall increase in the similarity d_t.

4.4. Comparison with Other Methods

We compare our results against the method of Nieves et al. [NGRCR20] which is also based on color palette extraction from a color subset of the painting using k-means and the selection of "relevant colors" is guided by colorimetric parameters and color statistics. We also compare our methods against methods based on convex hull color extraction, specifically, we show results of the methods of Tan et al. 16 [TLG16] (with a maximum NRC of 10) and Tan et al. 18 [TEG18] (with a variable NRC) whose color

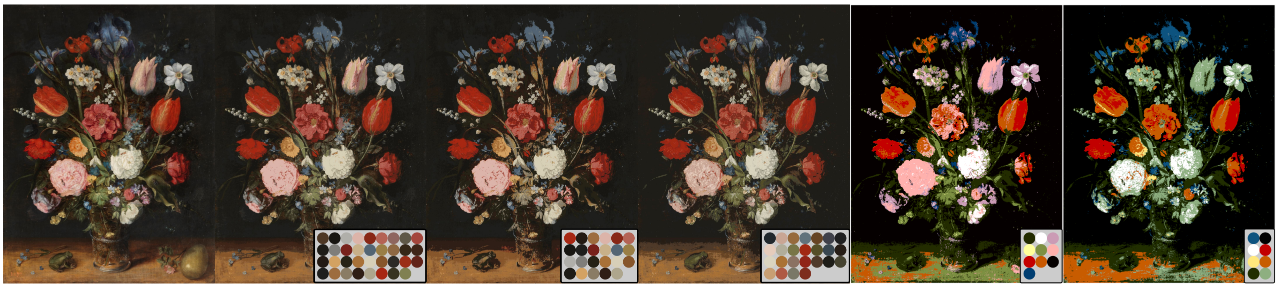

	Input Image	Base	Adaptive	Nieves et al.	Tan et al. 16	Tan et al. 18
MSE/MAE		**0.0013/0.0728**	0.0027/0.1026	0.0017/0.0894	0.0191/0.3625	0.0208/0.3817
PSNR/SSIM		**28.9119/0.9040**	25.6882/0.8506	27.7688/0.8748	17.1827/0.3341	16.8216/0.2360
NRC		35	23	25	10	8

Figure 7: *Quantizations representations of the input painting "Vase of Flowers" using the color palette Y_n extracted by our fuzzy logic system F both the base and adaptive versions, and the methods of Nieves et al. [NGRCR20], Tan et al. 16 [TLG16] and Tan et al.18 [TEG18]. We can see how the base version of our methods outperforms the other methods, while the adaptive version achieves a similar performance as the method of Nieves et al. and outperforms convex-hull-based methods. The colored circles appearing inside the gray boxes in the lower-right corner represent the palettes of "relevant colors" extracted using the evaluated methods.*

palette extraction is based on the image's RGB-space geometry, where the vertices of the convex hull formed by the pixel values in RGB-space correspond to the different "relevant colors". We provide both results reconstructing the input image using the color palettes extracted by the different methods, and evaluating the consistency regarding the colors selected by the observers.

Quantization Comparisons We evaluate reconstruction quality using the color palette Y_n extracted by our fuzzy logic system F (both Base and Adaptive versions), and the color palettes returned by the methods of Nieves et al. [NGRCR20], Tan et al. 16 [TLG16] and Tan et al. 18 [TEG18], on both the Painting-91 and Prado datasets. Table 2 shows that reconstructions using the Base version of our method present a significant quality regarding the method of Nieves et al., while the Adaptive version achieves a slightly lower performance. Moreover, both versions of our method (i.e., the Base and the Adaptive) outperform methods based on convex hulls, which focus on computing image boundary colors and do not take into account colorimetric parameters and color statistics. Figure 7 shows the quality-assessment metrics used to compare the performance of the different methods for one of the images. We can see that our fuzzy logic system F and the method of Nieves et al. achieve more accurate color palettes and reconstruction quality in contrast with convex-hull-based methods. However, while the Adaptive version preserves the primary colors, it exhibits a loss of high-frequency detail compared to reconstructions produced by the Base version and the method of Nieves et al.

Consistency with the Observers Finally, we compute the similarity between our color palettes (using both the Base and Adaptive versions) against the colors chosen by the observers, and compare the results with the methods of Nieves et al. [NGRCR20], Tan et al. 16 [TLG16] and Tan et al. 18 [TEG18]. In Table 3, we show the average similarity d_t (on the Prado dataset) between the color palettes returned by the different methods and the colors chosen by the observers (first row). Our base fuzzy logic system outperforms

Table 2: *Average MSE, MAE, PSNR, and SSIM metrics reconstructing the input images on the evaluation dataset. We can see how the base version of our fuzzy logic system outperforms the methods proposed by: Nieves et al. [NGRCR20], Tan et al. 16 [TLG16] and Tan et al.18 [TEG18]; while the adaptive version of our fuzzy logic system achieves a similar performance to the method of Nieves et al.*

Method	MSE ↓	MAE ↓	PSNR (dB) ↑	SSIM ↑
Base	**0.0015** ± 0.0017	**0.0773** ± 0.0311	**29.3951** ± 3.1867	**0.9339** ± 0.0400
Adaptive	0.0058 ± 0.0119	0.1379 ± 0.0881	24.3442 ± 3.7812	0.8650 ± 0.0813
Nieves et al.	0.0024 ± 0.0014	0.1067 ± 0.0261	26.7072 ± 2.0831	0.8950 ± 0.0431
Tan et al. 16	0.0203 ± 0.0122	0.1126 ± 0.0410	17.7482 ± 2.8205	0.6025 ± 0.1476
Tan et al. 18	0.0439 ± 0.0268	0.1693 ± 0.0655	14.4745 ± 3.0076	0.4384 ± 0.1849

the other methods since it returns a greater NRC (second row). As shown in Figure 8 (in which we can see how the color palettes, extracted by our fuzzy logic system and the method of Nieves et al., follow a similar distribution close to the distribution of colors selected by the observers. On the other hand, the convex-hull-based methods extract color palettes that are far from following the distribution of the colors selected by the observers, since these algorithms are not designed to mimic human perception.), and compute the color boundary in the RGB color space without modeling human perception. The Adaptive version of our method achieves a similar performance to the method of Nieves et al., while outperforming convex-hull-based methods. To further investigate potential differences, we performed Shapiro-Wilk tests, which confirmed that data for both the Adaptive version ($W = 0.9871, p = 0.9914$) and the method of Nieves et al. ($W = 0.9651, p = 0.6494$) follow

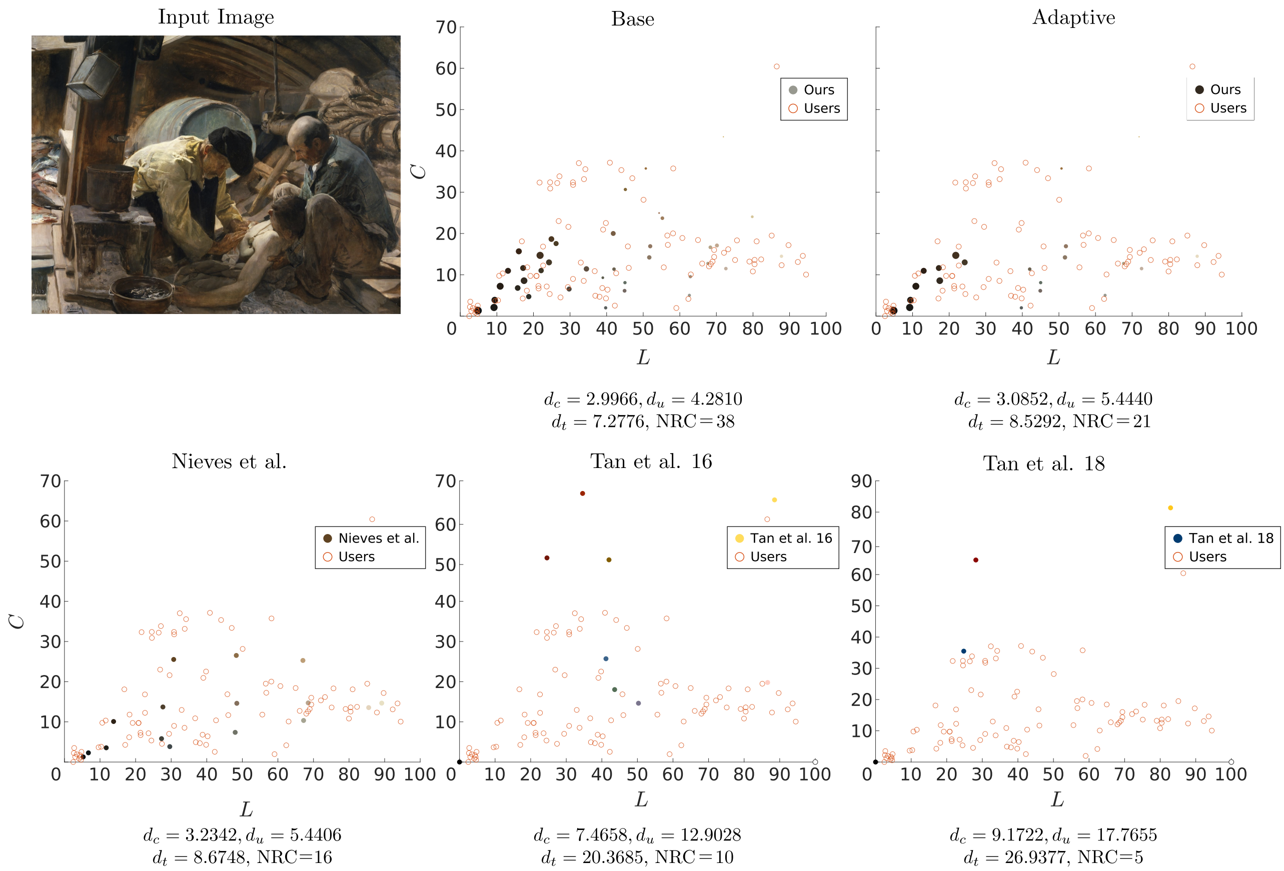

Figure 8: *Comparison of the distribution in the L–C (luminance–chroma) plane of the colors selected by the observers (color dots) and the relevant colors (orange dots) extracted by: our fuzzy logic system F (both the base and adaptive versions), the method of Nieves et al. [NGRCR20], the method of Tan et al. 16 [TLG16] and the method of Tan et al.18 [TEG18]; on the input painting "And they say that fish is expensive!".The size of the color dots denotes the number of similar colors S of each color when using our fuzzy logic system, a larger size implies a larger number of similar colors.*

Table 3: *Average distance d_t and NRC on the Prado dataset by the different our fuzzy logic system F (both base and adaptive), the method proposed by Nieves et al. [NGRCR20], and the convex-hull-based methods of Tan et al. 16 [TLG16] and Tan et al.18 [TEG18]. Our base fuzzy logic system outperforms the other methods due to the greater NRC, while the adaptive version achieves similar performance to the method proposed by Nieves et al. and significantly improves the performance of convex-hull-based methods.*

Base	Adaptive	Nieves et al.	Tan et al. 16	Tan et al. 18
7.6515	9.2272	9.0586	20.4624	26.2867
$\pm$ 1.3729	$\pm$ 1.9436	$\pm$ 0.8785	$\pm$ 2.0323	$\pm$ 4.3038
37.1	23.5	18.8	10.0	6.0

a normal distribution for the distance d_t. Consequently, a t-test was conducted to compare the two approaches. The results indicated no statistically significant difference ($p = 0.6671$), suggesting that both methods achieve similar performance.

5. Discussion and Conclusions

We have presented a fuzzy logic system for color palette extraction from paintings, focusing on colorimetric parameters such as luminance, chroma, and prevalence of colors to determine their relevance. Compared to traditional methods, such as convex-hull-based techniques or those solely reliant on RGB values, our system improves the image reconstruction quality and generates a more representative color palette. Although the Adaptive and Recursive versions performed similarly in image reconstruction, the Adaptive method was chosen over the Recursive version for comparisons with other methods due to its ability to extract a balanced NRC, while the color palette better aligns with the colors chosen by observers. Our Adaptive version achieves reconstruction quality and observer consistency comparable to the approach of Nieves et al. [NGRCR20]. The adaptive strategy applies a second iteration of Fisher's linear discriminant only when the initial number of colors exceeds a threshold, ensuring that the method adapts to varying levels of color richness across paintings. Through the integration of k-means clustering in the CIELAB color space and a rule-based fuzzy inference mechanism, we achieved significant alignment with human perception of relevant colors in paintings. The evaluation of our method on a painting dataset demonstrates its effectiveness in reconstructing images and its consistency with observer-selected palettes, as evidenced by low d_t, in particular the Base version. Another key contribution of this work is the introduction of a new evaluation metric that measures the similarity between the extracted relevant color palettes. This metric, derived from the CIE Delta E_{00} color difference [SWD05], accounts for both precision and recall, providing a comprehensive measure of perceptual relevance.

However, our method is not free of limitations. Despite its strong performance, the approach presents certain limitations. While the fuzzy-logic framework effectively identifies "relevant colors" with low chroma and luminance, its ability to recover highly saturated or very bright colors is reduced,in which case, the method proposed by Nieves et al. [NGRCR20] might be more suitable. Future work could explore adjustments to the rule base (particularly within the adaptive variant) to better encompass the chromatic gamut of both the images and the observer-selected color sets. Additionally, the palette extraction process relies on an initial subset of colors obtained using k-means clustering. A promising direction for improvement would involve replacing the standard squared Euclidean distance with a perceptually meaningful color metric for centroid computation, enabling a more faithful representation of the image's color distribution, or using a weighted mean by splitting the pixel colors into three-dimensional bins to compute the weights, while initializing the means deterministically, following the work of Chang et al. [CFL*15]. Another interesting option could be to use a combination of relevance and average distance to other "relevant colors" as a variable to calculate Fisher's linear discriminant. This would leave us with "relevant colors" that are different from other "relevant colors", i.e., with a wider gamut. Finally, our work could be extended to assess whether our color palettes are also consistent with other types of visual stimuli (e.g., natural images) and with a broader group of users, by collecting a test dataset annotated by a large number of observers, to test whether this yields the same conclusions.

6. Acknowledgments

This work has been supported by grant PID2022-141539NB-I00, funded by MICIU/AEI/10.13039/501100011033 and by ERDF, EU, and by the Government of Aragon's Departamento de Ciencia, Universidad y Sociedad del Conocimiento through the Reference Research Group "Graphics and Imaging Lab". Samuel Morillas acknowledges the support of Generalitat Valenciana under grant IMaLeVICS CIAICO-2022-051 and Spanish Agencia Estatal de Investigación under grants PID2022-140189OB-C21 and PID2023-152301OB-I00. J. Daniel Subias was supported by the CUS/702/2022 predoctoral grant. Juan Luis Nieves also ackowledges the Erasmus+ master Computational Colur and Spectral Imaging for supporting this work at the University of Granada. We would like to thank all the members of the Graphics and Imaging Laboratory who helped proofread the text

References

[AIP20] AZIMZADEH IRANI A., POURGHOLI R.: Segmentation assisted object distinction for direct volume rendering. *Journal of AI and Data Mining 8*, 1 (2020), 67–82. 2

[CC21] CHARLIN V., CIFUENTES A.: A general framework to study the price-color relationship in paintings with an application to mark rothko rectangular series. *Color Research & Application 46*, 1 (2021), 168–182. 3

[Cel23] CELEBI M. E.: Forty years of color quantization: a modern, algorithmic survey. *Artificial Intelligence Review 56*, 12 (2023), 13953–14034. 2, 3

[CFL*15] CHANG H., FRIED O., LIU Y., DIVERDI S., FINKELSTEIN A.: Palette-based photo recoloring. *Proc. ACM SIGGRAPH 34*, 4 (July 2015). 2, 3, 4, 10

[cie18] *CIE. Colorimetry, 4th Edition. Report No. CIE Publication 015:2018,CIE Central Bureau, Vienna* (2018). 2

[CNS19] CIOCCA G., NAPOLETANO P., SCHETTINI R.: Evaluation of automatic image color theme extraction methods. In *International Workshop on Computational Color Imaging* (2019), Springer, pp. 165–179. 3

[Col19] COLORMIND: Available online. http://colormind.io, 2019. accessed on 2 February 2024. 3

[CPD23] CELEBI M. E., PÉREZ-DELGADO M.-L.: cq100: a high-quality image dataset for color quantization research. *Journal of Electronic Imaging 32*, 3 (2023), 033019–033019. 2, 3

[CSG21] CHAO C.-K. T., SINGH K., GINGOLD Y.: Posterchild: Blend-aware artistic posterization. In *Computer Graphics Forum* (2021), vol. 40, pp. 87–99. 2, 3, 4

[CWH15] CELEBI M. E., WEN Q., HWANG S.: An effective real-time color quantization method based on divisive hierarchical clustering. *Journal of Real-Time Image Processing 10* (2015), 329–344. 3

[DPM*16] DAEL N., PERSEGUERS M.-N., MARCHAND C., ANTONIETTI J.-P., MOHR C.: Put on that colour, it fits your emotion: Colour appropriateness as a function of expressed emotion. *Quarterly Journal of Experimental Psychology 69*, 8 (2016), 1619–1630. 2

[Equ89] EQUITZ W. H.: A new vector quantization clustering algorithm. *IEEE Transactions on Acoustics, Speech, and Signal Processing 37*, 10 (1989), 1568–1575. 3

[FFM*26] FATHEMA S., FATHEMA S., MOUNIKA S. B., LAKSHMI R. N. S., SUNAIN S. H., RAJA RAJESWARI C.: Colour analysis and classification based on deep learning technique. In *ICT for Intelligent Systems* (2026), Springer Nature Singapore, pp. 401–410. 2

[Fis36] FISHER R. A.: The use of multiple measurements in taxonomic problems. *Annals of Eugenics 7*, 2 (1936), 179–188. 4, 6, 7

[Gen90] GENTILE R. S.: Quantization of color images based on uniform color spaces. *Journal of Imaging Technology 16*, 1 (1990), 11–21. 3

[GF07] GRAHAM D. J., FIELD D. J.: Statistical regularities of art images and natural scenes: spectra, sparseness and nonlinearities. *Spatial vision 21* (2007). 2

[GFL16] GILBERT A. N., FRIDLUND A. J., LUCCHINA L. A.: The color of emotion: A metric for implicit color associations. *Food Quality and Preference 52* (2016), 203–210. 2

[GP88] GERVAUTZ M., PURGATHOFER W.: A simple method for color quantization: Octree quantization. In *New Trends in Computer Graphics: Proceedings of CG International'88* (1988), Springer, pp. 219–231. 3

[HDI*23] HEGEMANN L., DAYAMA N. R., IYER A., FARHADI E., MARCHENKO E., OULASVIRTA A.: Cocolor: Interactive exploration of color designs. In *Proc. International Conference on Intelligent User Interfaces* (2023), pp. 106–127. 2

[KBVdWF14] KHAN F. S., BEIGPOUR S., VAN DE WEIJER J., FELSBERG M.: Painting-91: a large scale database for computational painting categorization. *Machine Vision and Applications 25* (2014), 1385–1397. 2, 6

[KE04] KAYA N., EPPS H. H.: Relationship between color and emotion: A study of college students. *College Student Journal 38*, 3 (2004), 396–405. 2

[Ker99] KERRE, ETIENNE: *Fuzzy sets and approximate reasoning. Xian Jiaotong University Press. Xian, 254 p.* 1999. 6

[LH13] LIN S., HANRAHAN P.: Modeling how people extract color themes from images. In *Proc. SIGCHI Conference on Human Factors in Computing Systems* (2013), pp. 3101–3110. 3

[LMS24] LIN C., MOTTAGHI S., SHAMS L.: The effects of color and saturation on the enjoyment of real-life images. *Psychonomic Bulletin & Review 31*, 1 (2024), 361–372. 2

[LPN08] LINHARES J. M. M., PINTO P. D., NASCIMENTO S. M. C.: The number of discernible colors in natural scenes. *Journal of the Optical Society of America A 25*, 12 (2008), 2918–2924. 2

[Men95] MENDEL J.: Fuzzy logic systems for engineering: a tutorial. *Proc. IEEE 83*, 3 (1995), 345–377. 4

[MS24] MURATBEKOVA M., SHAMOI P.: Color-emotion associations in art: Fuzzy approach. *IEEE Access* (2024). 3

[Mur05] MUREIKA J. R.: Fractal dimensions in perceptual color space: a comparison study using jackson pollock's art. *Chaos: An Interdisciplinary Journal of Nonlinear Science 15*, 4 (2005). 2

[Mus24] MUSEUM P.: Painting collection, 2024. Available online: https://www.museodelprado.es/en/the-collection (only Spanish version). 4, 6

[NGRCR20] NIEVES J. L., GOMEZ-ROBLEDO L., CHEN Y.-J., ROMERO J.: Computing the relevant colors that describe the color palette of paintings. *Applied Optics 59*, 6 (Feb 2020), 1732–1740. 2, 3, 4, 7, 8, 9, 10

[NOGRR21] NIEVES J. L., OJEDA J., GÓMEZ-ROBLEDO L., ROMERO J.: Psychophysical determination of the relevant colours that describe the colour palette of paintings. *Journal of Imaging 7*, 4 (2021). 6

[PA98] POINTER M. R., ATTRIDGE G.: The number of discernible colours. *Color Research & Application 23*, 1 (1998), 52–54. 2

[SDWN*07] SCHULTE S., DE WITTE V., NACHTEGAEL M., MÉLANGE T., KERRE E. E.: A new fuzzy additive noise reduction method. In *Image Analysis and Recognition* (Berlin, Heidelberg, 07), Kamel M., Campilho A., (Eds.), Springer Berlin Heidelberg, pp. 12–23. 4

[SRSS22] SAXENA P., RAJESH E., SINGH P., SINGH A. S.: Colour detection in objects using nin implemented cnn. *Specialusis Ugdymas 1*, 43 (2022), 3989–4001. 2

[ST20] SAINUI J., TONGSAMRIT M.: Color naming for describing object in image using different classification algorithms and color spaces. In *Proc. International Conference on Computer and Automation Engineering* (2020), pp. 65–69. 2, 3

[SWD05] SHARMA G., WU W., DALAL E. N.: The ciede2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application 30*, 1 (2005), 21–30. 7, 10

[SWK07] SCHULTE S., WITTE V., KERRE E.: A fuzzy noise reduction method for color images. *IEEE Transactions on Image Processing 16* (06 2007), 1425–36. 4

[TCT13] THOMAS J.-B., COLANTONI P., TRÉMEAU A.: On the uniform sampling of cielab color space and the number of discernible colors. In *Proc. Computational Color Imaging* (2013), Springer, pp. 53–67. 2

[TEG18] TAN J., ECHEVARRIA J., GINGOLD Y.: Efficient palette-based decomposition and recoloring of images via rgbxy-space geometry. *ACM Transactions on Graphics (TOG) 37*, 6 (2018), 1–10. 3, 7, 8, 9

[TLG16] TAN J., LIEN J.-M., GINGOLD Y.: Decomposing images into layers via rgb-space geometry. *ACM Transactions on Graphics (TOG) 36*, 1 (Nov. 2016). 7, 8, 9

[UKSU17] UEDA Y., KOGA T., SUETAKE N., UCHINO E.: Color quantization method based on principal component analysis and linear discriminant analysis for palette-based image generation. *Optical Review 24* (2017), 741–756. 3

[VDVNVdW*03] VAN DE VILLE D., NACHTEGAEL M., VAN DER WEKEN D., KERRE E., PHILIPS W., LEMAHIEU I.: Noise reduction by fuzzy image filtering. *IEEE Transactions on Fuzzy Systems 11*, 4 (2003), 429–436. 4

[WFC*09] WALLRAVEN C., FLEMING R., CUNNINGHAM D., RIGAU J., FEIXAS M., SBERT M.: Categorizing art: Comparing humans and computers. *Computers & Graphics 33*, 4 (2009), 484–495. 2

[WG18] WITZEL C., GEGENFURTNER K. R.: Color perception: Objects, constancy, and categories. *Annual Review of Vision Science 4*, 1 (2018), 475–499. 2

[WHJ20] WEINGERL P., HLADNIK A., JAVORŠEK D.: Development of a machine learning model for extracting image prominent colors. *Color Research & Application 45*, 3 (2020), 409–426. 3

[WLX19] WANG Y., LIU Y., XU K.: An improved geometric approach for palette-based image decomposition and recoloring. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 11–22. 3

[WO18] WILMS L., OBERFELD D.: Color and emotion: effects of hue, saturation, and brightness. *Psychological Research 82*, 5 (2018), 896–914. 2

[XVCH*24] XUE D., VAZQUEZ-CORRAL J., HERRANZ L., ZHANG Y., BROWN M. S.: Palette-based color harmonization via color naming. *IEEE Signal Processing Letters* (2024). 3

[Zad65] ZADEH L.: Fuzzy sets. *Information and Control 8*, 3 (1965), 338–353. 3

[ZRL97] ZHANG T., RAMAKRISHNAN R., LIVNY M.: Birch: A new data clustering algorithm and its applications. *Data mining and knowledge discovery 1* (1997), 141–182. 3