



UNIVERSITÀ POLITECNICA DELLE MARCHE
Repository ISTITUZIONALE

Defining a deep neural network ensemble for identifying fabric colors

This is the peer reviewed version of the following article:

Original

Defining a deep neural network ensemble for identifying fabric colors / Amelio, A.; Bonifazi, G.; Corradini, E.; Di Saverio, S.; Marchetti, M.; Ursino, D.; Virgili, L.. - In: APPLIED SOFT COMPUTING. - ISSN 1568-4946. - 130:(2022). [10.1016/j.asoc.2022.109687]

Availability:

This version is available at: 11566/306601.7 since: 2024-05-08T15:00:02Z

Publisher:

Published

DOI:10.1016/j.asoc.2022.109687

Terms of use:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. The use of copyrighted works requires the consent of the rights' holder (author or publisher). Works made available under a Creative Commons license or a Publisher's custom-made license can be used according to the terms and conditions contained therein. See editor's website for further information and terms and conditions.

This item was downloaded from IRIS Università Politecnica delle Marche (<https://iris.univpm.it>). When citing, please refer to the published version.

(Article begins on next page)

Defining a deep neural network ensemble for identifying fabric colors

Alessia Amelio¹, Gianluca Bonifazi², Enrico Corradini², Simone Di Saverio³, Michele Marchetti², Domenico Ursino^{2*}, and Luca Virgili²

¹ INGEO, University “G. D’Annunzio” of Chieti-Pescara

² DII, Polytechnic University of Marche

³ Imola Informatica

* Contact Author

alessia.amelio@unich.it; g.bonifazi@univpm.it; e.corradini@pm.univpm.it;
sdisaverio@imolainformatica.it; m.marchetti@pm.univpm.it; d.ursino@univpm.it;
luca.virgili@univpm.it

Abstract

Colors characterize each object around us. For this reason, the study of colors has played a key role in Artificial Intelligence (think, for instance, of image classification, object recognition and segmentation). However, there are some topics about colors still little explored. One of them concerns fabric colors. This is a particular topic since fabrics have some characteristics, such as specific textures, that are not found in other contexts. In this paper, we want to propose a new Convolutional Neural Network (CNN) based model for identifying fabric colors. After introducing this model, we consider three different versions of it and create an ensemble of the corresponding CNNs to get better results. Finally, through a series of experiments, we show that our ensemble is able to improve the state-of-the-art on the identification of fabric colors.

Keywords: Color Classification; Ensemble Learning; Identification of Fabric Colors; Classification of Fabric Colors; Convolutional Neural Networks

1 Introduction

Visual information is one of the most important mechanisms by which people communicate (think, for instance, of documents, images, videos, movies, etc.). Color is an essential component of visual communication. It is also a relevant topic in several research fields, for example in physics to model light and reflectance, in physiology to model perception, in biology to model the visual system, in art to model beauty. As a result, the study of colors has fascinated, and continues to fascinate, researchers from a wide variety of fields [24, 35].

One of the general problems with colors concerns the identification of the color of objects and, more generally, of any entity. This issue is critical because it allows, for example, the characterization of the content of objects, the understanding of emotional and psychological states, and the discovery of recurrent patterns and anomalous conditions.

One area in which color recognition is essential is textile industry. In this scenario, color is one of the most important information to be transmitted between the manufacturer and the customer.

Furthermore, in the same industry, it is one of the most important characteristics that can be leveraged in marketing and R&D, as well as one of the most critical features that designers employ to succeed. The attraction that the color of a fabric can exert in a customer’s final decision has also been studied scientifically (see, for instance, [31]).

In recent years, several authors have investigated the issue of color identification in several fields including the textile industry. This has led to various methods for recognizing color components, patterns and shades, and for automatically detecting the layout of color yarns from images. These methods use several approaches, such as fuzzy C-Means [48, 21] and X-means [9], rule-based techniques [1], Support Vector Machine [32, 50], Random Forest [6] and deep neural networks [28], to reach their objectives. Other works have focused on textile whiteness estimation [39]. Others have investigated color prediction and recognition using deep neural networks [17, 52, 19, 44, 2] and eXtreme Gradient Boosting (XGBoost) [53]. Finally, several works on color recognition match the color of interest with a reference palette to detect the tolerance with respect to fabric standards [42, 46].

In the past, deep learning based approaches have been used for color identification in different contexts, including color recognition of vehicles and traffic lights [52, 19], flowers [2], stool medical images [44], etc. In all of these cases, they performed better than other machine learning approaches [52]. Despite this, the use of deep learning for color recognition in the context of textile is still poorly investigated. Moreover, the proposed approaches have some limitations in extracting colors within recurrent patterns, shades and layout variations, which are very important features for fabric analysis.

In this paper, we want to make a contribution in this setting by proposing a new deep learning based approach for fabric color identification. To overcome the limitations described above, we propose a new architecture based on an ensemble of Convolutional Neural Networks (hereafter, CNNs), whose input is in the color difference domain. This is obtained by considering the difference between the original input image and a set of reference color images. Representing the input in the color difference domain allows color variations, shades and patterns to be easily captured. This can make easier for our network to learn features for color recognition in fabric images. Furthermore, adopting an ensemble strategy provides greater robustness than a single CNN, which, in turn, results in greater accuracy in returned results. We focus on one particular aspect of color recognition in the textile industry, namely color identification of fabric images. This issue is interesting to address because it can help reduce inefficiency, wasted time and subjectivity, which can ultimately lead to higher quality products that better meet customer needs. We performed experiments to evaluate our approach on a well-known dataset of fabric images and compared our results with other state-of-the-art deep neural networks for color identification. The results obtained show that a CNN ensemble in the color difference domain not only outperforms single CNNs but also obtains better results than those returned by well-known deep learning architectures widely used for pattern recognition and color identification.

Summarizing, the main contributions of the proposed methodology, compared to the related ones already existing in the literature, are:

- a technique for encoding input images in the color difference domain;
- the use of ensemble learning for color recognition in the context of textiles.

The outline of this paper is as follows: Section 2 reviews the related literature, with a particular focus on the analysis of colors in the textile scenario. Section 3 presents our approach and, in particular, its two phases involving mapping the color space in the color difference space and defining an ensemble of CNNs. Section 4 illustrates the experiments we performed to test and evaluate our approach.

Finally, in Section 5, we draw our conclusions and look at possible future developments of our research efforts.

2 Related Literature

Color identification is a challenging image processing task in the context of Computer Vision. In fact, color represents a fundamental feature in the study of digital images and is the starting point for many object recognition and tracking techniques [26, 40]. For this reason, in the past literature, many color recognition techniques have been proposed. They can be divided into three distinct threads, which adopt handcrafted features, machine learning and deep learning, respectively, to reach their objective.

Handcrafted feature based approaches are the most traditional ones because they are easy to implement. They generally have a high execution speed against a lower accuracy, due to a poor generalization capability. One of the most critical features of digital color images is the RGB value of pixels [47, 1, 25, 33]. In this context, the authors of [1] present a rule-based approach to recognize different shades of basic colors of fabric images. They extract the RGB color features to obtain the mean and standard deviation of red, green and blue shades. Then, thanks to these features, they define the rules for color classification considering ten shades for each color. The main drawback of this approach is in the definition of rules, because it is difficult to find effective rules capable of capturing the different color shades. Sometimes, the recognition of colors by means of single features may not provide satisfactory results. For this reason, the authors of [7] propose an approach that combines multiple features. In particular, it combines transformed color histogram, hue histogram, normalized RGB histogram and color moment as feature patches to recover information from each feature.

As for machine learning based approaches, several studies have adopted the Gaussian Mixture Model (GMM, for short) to obtain an unsupervised classifier of color regions [18, 45, 23]. In particular, the authors of [18] present an approach for finding color regions in digital images. It consists of two steps. In the first one, it estimates GMM parameters by applying Expectation Maximization on a reference image that includes the regions of interest. For this purpose, it considers two chrominance features, i.e., Cb and Cr channels in the YCbCr color space model. The second step searches and segments the color regions in the input images using GMM parameters extracted from the reference images. The authors of [53] use the XGBoost classification model to classify the color of solid wood flooring. First they analyze all features; then, they remove those with low variance in order to keep only the most essential information and improve classification speed and accuracy. The two most widely used machine learning algorithms for color classification and clustering are Support Vector Machine and K-Means [42, 32, 41, 11, 3]. In [41], the authors present an approach that uses K-Means to perform color matching in textiles, whose ultimate goal is measuring quality and increasing efficiency. They show that K-Means provides better results than the standard Euclidean distance method for color detection. Recent studies use an evolution of K-Means, called fuzzy C-Means. In this case, data points can belong to more than one cluster with a given probability. Thanks to this choice, this approach obtains comparatively better results for overlapped datasets [48, 47, 21].

Deep learning based approaches have been used recently in several domains, such as health care and face and object recognition [44, 12, 19]. Many of them focus on vehicle color recognition [52, 51, 34, 14]. To this end, they train deep learning models capable of recognizing vehicle colors in photos taken in a road environment, despite the presence of adverse conditions, such as the existence of shadows or reflections. In [34], the authors propose a method for vehicle color recognition using Convolutional

Neural Networks. They show that this method is robust and accurate in handling the variation of lighting and environment, and outperforms approaches using only traditional handcrafted feature descriptors.

Color recognition approaches are also slowly gaining popularity in the textile and fashion industry, where they are used to accurately identify fabric colors [23, 28]. In [23], the authors focus on fashion product portraits and propose an approach that autonomously finds the garment region in an image and, then, retrieves information on colors and patterns of items. In [28], the authors propose a regression model for color retrieval in fashion garments. Their approach focuses mainly on objects belonging to a different category of garments. To this end, it modulates the importance of each pixel by adopting the so-called color name-attention technique, which selects only those pixels sufficiently close to the main color. This way of proceeding guarantees it a good performance. However, this decreases when the approach has to deal with fabrics without garments, because the color name-attention technique cannot be employed in that case.

In Table 1, we propose a summarization of the related approaches described in this section. It provides a schematic summary of their main characteristics.

<i>Reference work</i>	<i>Model adopted</i>	<i>Color recognition method</i>
Anami <i>et al.</i> [1]	Handcrafted features	Rule-based approach with RGB color features
Chen <i>et al.</i> [7]	Handcrafted features	Feature context, with features encoded as histograms
Kim [18]	Machine learning	Color region detection based on GMM
Zhuang <i>et al.</i> [53]	Machine learning	XGBoost classification model on selected features
Vitthal <i>et al.</i> [41]	Machine learning	K-Means color clustering algorithm
Su <i>et al.</i> [34]	Deep Learning	Convolutional Neural Networks trained for color recognition
Li <i>et al.</i> [23]	Deep Learning	Image segmentation, garment region selection, color coordinate matching
Rame <i>et al.</i> [28]	Deep Learning	Prediction of RGB values through color regression and attention modules

Table 1: Schematic summary of the main related approaches proposed in the past literature

3 Description of the proposed approach

We can refer to a color space as a specific organization of colors. To define it, we need to use a color model, which is an abstract mathematical model describing how colors can be represented as tuples of numbers. Examples of color models are RGB (Red, Green, Blue), HSV (Hue, Saturation, Value), CMYK (Cyan, Magenta, Yellow, Key). Currently, the most widely used color space is RGB, which describes the light emitted by a given color through a triplet of numbers representing the combination of three basic colors, namely red, green and blue. Although this model is sufficient to represent digital images, it does not allow the extraction of handcrafted features sufficient to guarantee high performance in the color classification task [47, 1, 7]. For this reason, our approach starts from the RGB model and then maps the images to another domain that represents color difference. This adds more information about the input images and provides additional perspectives that could be leveraged for color recognition.

3.1 Mapping the color space into the difference domain

To the best of our knowledge, the deep learning approaches for color identification proposed in the past literature do not consider textural patterns (which are heavily present in clothes images), since they can hamper the learning of the underlying neural network. This way of proceeding derives at least in part from the color representation used for images [47, 1, 7], which does not provide sufficient

information for color classification. To overcome this limitation, we introduce a new color space called difference space, which greatly simplifies the color recognition issue.

The reasoning underlying the difference space consists in computing the distances between the colors of an input image and a set of reference colors. First, we need to choose a set of p reference colors and determine the corresponding RGB encoding. Then, we compute the differences between the colors of the input image and the reference colors and obtain p images represented using the RGB encoding. After that, we concatenate all the p images thus obtained to create the corresponding input to the deep learning model. This input has the same width and height as the initial images, but it has a third deeper dimension, of size $q = p \times 3$. It is obtained by concatenating the p images, each represented by means of the 3 RGB channels.

Formally speaking, we define a function $\mathcal{F}: \mathbb{R}^{m \times n \times 3} \rightarrow \mathbb{R}^{m \times n \times q}$, where m and n are the width and height of the image, and q is the size of the difference domain. \mathcal{F} is the concatenation of the element-wise distance between the RGB channels of the input image and the corresponding ones of the reference color:

$$\mathcal{F} = \text{Concat}_{i=1}^p \{(R_{input} - R_{color_i}), (G_{input} - G_{color_i}), (B_{input} - B_{color_i})\} \quad (1)$$

Here, R_{input} , G_{input} and B_{input} are the corresponding Red, Green and Blue channels of the input image, while R_{color_i} , G_{color_i} , and B_{color_i} are the values of the RGB encoding of the i -th reference color. If a subtraction between the input color and the reference one returns less than 0, we set the corresponding result equal to 0. Actually, in principle, our model (viewed abstractly, without considering the semantics of the data it is processing) could also operate receiving negative values as input. However, we must take into account that these data represent values in the RGB color space that, based on the semantics underlying this space, cannot be negative. Therefore, in order to preserve the semantics of all the results returned, we have imposed that, where negative values are present, they are replaced with zeros. This way of proceeding presents the problem that the original images could not be reconstructed. However, we believe that this is a necessary price to pay for preserving the semantic clarity of results. We also considered redoing the computation leaving negative values as input to the model, without replacing them with zeros, to see how much the final results differed. At the end of this test we were able to verify that the difference in the final results was, in fact, negligible. Indeed, the changes in performance measures between the two cases were less than 0.5%.

For the purpose of this paper (i.e., identifying colors from fabric images), we use 12 reference colors, namely Orange, White, Blue, Cyan, Yellow, Magenta, Black, Red, Earth Brown, Green, Emerald Green, and Purple. The selection of colors to use is not straightforward because the potential colors would be thousands. To carry out this task, we conducted some interviews with managers of fabric companies. They provided us with two important insights. The first concerns the fact that the colors used in the context of fabric have a slight different RGB representation from the pure definition of colors. An example of this concerns the black color. The pure black color corresponds to the hexadecimal RGB code #000000. However, in the context of textiles, there are several shades of black, and these are a bit far from the pure black color, as well as a bit far from each other. For this reason, in our method, we used the hexadecimal color #1C1C1C, instead of the pure black color, to represent black. This is because, based on our interviews, in the fabric context, the black color corresponding to the hexadecimal code #1C1C1C is the most commonly used. The second insight concerns the number of colors to be adopted in our color recognition task. Our interviews revealed that a set of 12 colors is capable of representing most of the possible cases. In this way, we were able to reach a reasonable trade-off between the ability to represent reality, the complexity of model and

the accuracy that can be achieved. In any case, we point out that our model can be easily scaled up to include more colors. At present, we have not seen fit to proceed in that direction. However, we do not exclude doing so in the future, if we see the need for it. As a consequence, for each image to be classified, we create 12 images, each representing the distance of the image itself from one of the reference colors. This distance is expressed in the RGB model. To give an idea of our way of proceeding, in Figure 1 we show an example of the differences between an input image and the set of reference colors. Finally, we concatenate the 12 images thus obtained on the third axis of the RGB channels to build the input for the tasks of the next phase.




Figure 1: An input image and the differences between it and the set of reference colors

3.2 Ensemble learning for color identification

In the previous section, we have seen how, by switching to the color difference domain, given a fabric image, we obtain a concatenation of 12 images, each of which can be represented using the RGB model. Since this model has 3 channels, it follows that each image can be represented by 36 channels.

To process such data we introduce an extension of the traditional CNN that we call CNN^Δ . It consists of an input layer that receives images of any width and height through 36 channels. After this layer, there is a convolutional layer with 3×3 kernels and stride 1, followed by a max pooling layer of size 2×2 and stride 1. Then, there is a convolutional layer with the kernel size identical to the previous one followed by a pooling layer of size 2×2 and stride 1. After them, we have a dense layer with 128 units connected to a final dense layer with a softmax activation function. The output of this network is a vector of 12 elements each representing the classification probability of the corresponding reference color. The total number of parameters in CNN^Δ is 175,476. The schematic representation of CNN^Δ is shown in Figure 2. In this case, it is designed for an input image having a width and a height of 100 pixels.




Figure 2: Architecture of CNN^Δ designed for an input image having a width and a height of 100 pixels

To further improve the performance and generalization of our CNN^Δ model, we adopt an ensemble learning technique according to a way of proceeding well known in the literature [29]. To this end, we construct a deep learning model consisting of multiple models operating together to predict the class of an input image. In particular, our model consists of an ensemble of three different versions of CNN^Δ , each having a unique set of hyperparameters and training methods, as shown in Table 2.

Actually, it is possible to define an ensemble learning technique using different types of learning models. This is often done in the literature. In our case, we also tested this hypothesis by combining in various ways Convolutional Neural Networks, Residual Networks, Densely Connected Networks and Inception Networks, which represent the state-of-the-art in image classification. However, the various combinations of heterogeneous models we tried returned worse performances, in terms of Precision, Recall, F1-score, Accuracy and average epoch training time, than those returned by our proposed strategy. Due to space limitations, we do not report the various trials and the corresponding results in this paper. However, in Section 4.2.1, we will have a way to highlight how the individual CNNs performed better than the other individual network types in our reference context.

We carried out a random search procedure for tuning the values of the hyperparameters of the three CNN models. We chose this technique because, compared to other hyperparameter optimization techniques existing in the literature, it is intuitive to combine with early stopping methods, which we chose to use during the training phase for reducing the probability of overfitting. The combination of these two techniques allows us to increase the efficiency in narrowing down the search space. More specifically, we considered two possible kinds of optimizer, namely ADAM (Adaptive Moment Estimation) and SGD (Stochastic Gradient Descent). In addition, we set the range of the learning rate

η to $[0, 1]$. As for ADAM, we set the range of the parameter ϵ to $(0, 0.01]$. As for SGD, we set the default decay value to 0.01 and we adopted the Nesterov momentum. In the random search procedure, we set the number of trials to 5. In each trial, we picked up the suitable hyperparameters for each CNN $^\Delta$ and we randomly selected their values from the corresponding ranges. Then, we trained each CNN $^\Delta$ using the values of the selected hyperparameters. Finally, we computed the values of F1-score, Accuracy and average epoch training time of the ensemble by applying it to the test set. At the end of this procedure, we chose as the final values of the hyperparameters of the CNN $^\Delta$ models of the ensemble, the ones leading it to obtain the highest values of F1-score and Accuracy and the lowest values of average epoch training time.

In the first trial our choices were as follows:

- as for CNN $^\Delta$ 1, we selected the ADAM optimizer; furthermore, we set $\eta = 0.05$ and $\epsilon = 0.003$;
- as for CNN $^\Delta$ 2, we selected the ADAM optimizer; furthermore; we set $\eta = 0.01$ and $\epsilon = 0.005$;
- as for CNN $^\Delta$ 3, we selected the ADAM optimizer; furthermore, we set $\eta = 0.1$ and $\epsilon = 0.005$.

For this trial, we obtained an F1-score equal to 0.70 and an Accuracy equal to 0.70. Instead, the average epoch training time was 5.23 seconds.

In the second trial our choices were as follows:

- as for CNN $^\Delta$ 1, we selected the ADAM optimizer; moreover, we set $\eta = 0.01$ and $\epsilon = 0.005$;
- as for CNN $^\Delta$ 2, we selected the SGD optimizer, we set $\eta = 0.02$ and the decay value equal to 0.01; furthermore, we adopted the Nesterov momentum;
- as for CNN $^\Delta$ 3, we selected the SGD optimizer, and we set $\eta = 0.01$ and the decay value equal to 0.01; furthermore, we adopted the Nesterov momentum.

For this trial, we obtained an F1-score equal to 0.72 and an Accuracy equal to 0.73. Instead, the average epoch training time was 6.19 seconds.

In the third trial our choices were as follows:

- as for CNN $^\Delta$ 1, we selected the SGD optimizer, we set $\eta = 0.1$ and the decay value equal to 0.01; finally, we adopted the Nesterov momentum;
- as for CNN $^\Delta$ 2, we selected the ADAM optimizer; furthermore, we set $\eta = 0.04$ and $\epsilon = 10^{-6}$;
- as for CNN $^\Delta$ 3, we selected the ADAM optimizer; furthermore, we set $\eta = 0.02$ and $\epsilon = 0.01$;

For this trial, we obtained an F1-score equal to 0.72 and an Accuracy equal to 0.72. Instead, the average epoch training time was 5.01 seconds.

In the fourth trial our choices were as follows:

- as for CNN $^\Delta$ 1, we selected the ADAM optimizer; moreover, we set $\eta = 0.001$ and $\epsilon = 10^{-7}$;
- as for CNN $^\Delta$ 2, we selected the ADAM optimizer; moreover, we set $\eta = 0.0001$ and $\epsilon = 10^{-8}$;
- as for CNN $^\Delta$ 3, we selected the SGD optimizer, we set $\eta = 0.01$ and the decay value equal to 0.01; furthermore, we adopted the Nesterov momentum.

For this trial, we obtained an F1-score equal to 0.79 and an Accuracy equal to 0.80. Instead, the average epoch training time was 4.68 seconds.

In the fifth trial our choices were as follows:

- as for CNN^Δ_1 , we selected the SGD optimizer, we set $\eta = 0.0001$ and the decay value equal to 0.01; finally, we adopted the Nesterov momentum;
- as for CNN^Δ_2 , we selected the SGD optimizer, we set $\eta = 0.6$ and the decay value equal to 0.01; finally, we adopted the Nesterov momentum;
- as for CNN^Δ_3 , we selected the SGD optimizer, we set $\eta = 0.08$, the decay value equal to 0.01; finally, we adopted the Nesterov momentum.

For this trial, we obtained an F1-score equal to 0.69 and an Accuracy equal to 0.70. Instead, the average epoch training time was 6.68 seconds.

As can be seen from the above description, the hyperparameter values that guaranteed the best performance in terms of F1-score, Accuracy and average epoch training time of the ensemble are those corresponding to the fourth trial of the random search procedure. They are shown in Table 2. In the following, we call $\mathcal{E}_{\text{CNN}^\Delta}$ this ensemble.

Network	Optimizer	Learning rate η	Momentum	ϵ
CNN^Δ_1	ADAM	0.001	-	10^{-7}
CNN^Δ_2	ADAM	0.0001	-	10^{-8}
CNN^Δ_3	SGD	0.01	Nesterov	-

Table 2: Main characteristics of the three CNN^Δ models used in $\mathcal{E}_{\text{CNN}^\Delta}$.

$\mathcal{E}_{\text{CNN}^\Delta}$ works as follows. During the prediction step, each CNN^Δ provides a classification probability vector for the reference colors. In this way, three probability vectors are obtained. These are processed by a soft-voting function. This function first computes the average of the classification probabilities, then it identifies the highest values thus obtained and, finally, returns the corresponding class as final output. In Figure 3, we report a graphical representation of $\mathcal{E}_{\text{CNN}^\Delta}$.

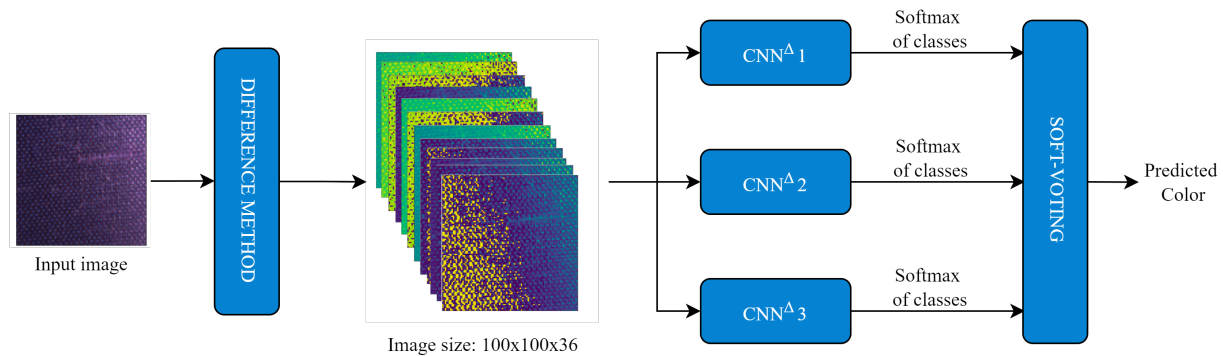


Figure 3: A graphical representation of $\mathcal{E}_{\text{CNN}^\Delta}$

4 Experiments

4.1 Testbed

We tested $\mathcal{E}_{CNN\Delta}$ on the Fabrics Dataset [16]. It consists of a benchmark collection of 2,000 different color images related to clothing and textiles. The format of the images is PNG and their size is 400×400 . They were acquired in clothing stores and labeled with the material composition provided by the manufacturer, as well as with information about the cloth origin (e.g., jacket, sweater, pants). Each image type was acquired using a photometric stereo sensor under 3 or 4 different illumination conditions, resulting in a total of 7,757 actual images. The dataset is not balanced because it was conceived to represent the real-world distribution of fabrics. Most of the images depict cotton and polyester clothes, while some of them regard silk and linen ones. The garments of many clothes are composed of two or more fabrics (blended fabrics). Figure 4 shows four sample images from the Fabrics Dataset.




Figure 4: Sample images from the Fabrics Dataset

Since the images in the Fabrics Dataset are not labeled with color, we had to manually specify the ground truth for each of them. For this purpose, we performed a procedure based on a visual inspection of them. Specifically, we showed each image to 11 different people and asked them to label it with one color among the 12 reference ones specified in Section 3.1. We assigned to an image the color most labeled by the 11 people to whom it was shown. We discarded all the images with which no color could be associated.

At the end of this procedure, we kept 4,591 images while discarding 3,166. The distribution of the ground truth colors relative to the images retained is shown in Figure 5.

Given this collection of images, we applied an undersampling procedure on the most numerous classes and an oversampling procedure on the least numerous ones to achieve a balance of classes. At the end of this procedure, we obtained a new dataset consisting of 1,200 images, in particular 100 images per class. Since images were acquired under different lighting conditions, the color identification task can be made more difficult by the presence of differently illuminated points. For example, this can happen due to shadows and flash glare from reflective fabrics. To address this issue, we first performed a pre-processing step on the 1,200 images for smoothing the differences in illumination caused by shade and flash glare. To address the shadow problem we used RIS-GAN [49], which is a Generative Adversarial Network for shadow removal. It consists of four generators in the encoder-decoder structure and three discriminators. The four generators are used to construct negative residual images, intermediate shadow-removal images, inverse illumination maps and refined shadow-removal images. To solve the flash glare problem, we adopted the approach described in [30]. It first derives priors, such as gradient sparsity, asymmetry, distinct color tint and piecewise constancy of the reflection

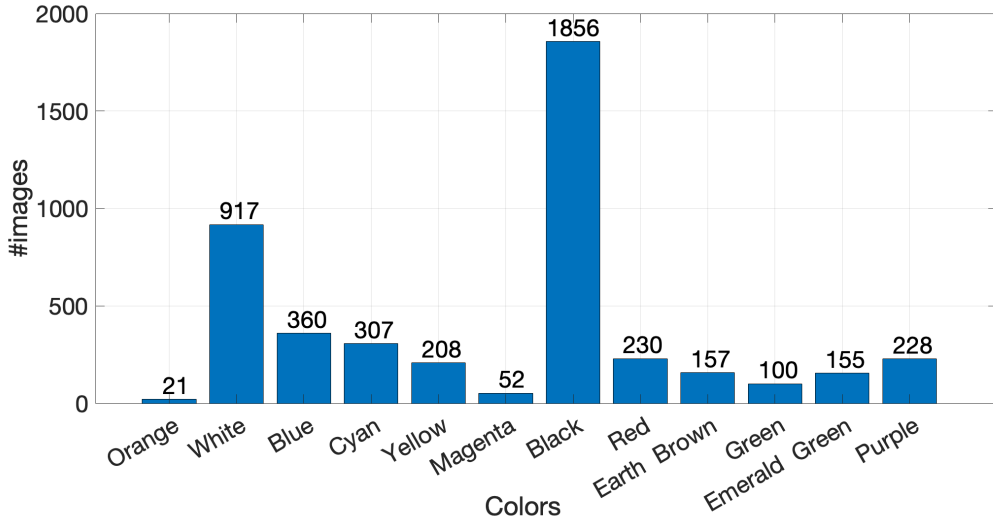


Figure 5: Distribution of the ground truth colors of the images kept in the dataset

layer. Then, it embeds these priors in an optimization framework, which finally produces the reflection-free image. In the following, we call \mathcal{D} the final dataset thus obtained.

We compared $\mathcal{E}_{CNN\Delta}$ with different deep learning state-of-the-art architectures. In particular, we considered: (i) the deep residual network (i.e., ResNet) [13], (ii) the inception network (i.e., Inception-ResNet) [36], and (iii) the dense convolutional network (i.e., DenseNet) [15]. We also considered: (iv) the Color Deep network [51], which is a reference network in the color recognition literature. Specifically, we selected two types of ResNet, i.e., ResNet50v2 and ResNet152v2, two types of DenseNet, i.e., DenseNet169v2 and DenseNet201v2, and InceptionResNetV2. We chose them because they revealed to achieve the best experimental results. Furthermore, to verify the effectiveness of adopting an ensemble strategy, we compared $\mathcal{E}_{CNN\Delta}$ with a single CNN $^{\Delta}$. For all the networks involved in the comparison task, the input consisted of 36 channels representing the difference between the image to be classified and the 12 reference colors (see Section 3.1). Since Color Deep could not simultaneously take in input the 36 channels corresponding to the 12 differences, we gave to this network the 12 images corresponding to the differences in cascade, and then aggregated the corresponding results.

To avoid overfitting in the classification of colors, we applied a data augmentation procedure in the images of \mathcal{D} . This activity consists of two tasks, namely: (i) random flip augmentation, which flips the image horizontally and vertically, and (ii) random zoom augmentation, which randomly increases the image size by adding new pixels around or interpolating the values of pixels.

We provided the original images of \mathcal{D} in input to the various models to be tested. The identification of colors in them is complex because of the presence of textural patterns.

We also extracted only the color components from the original fabric images and generated images with the only dominant color. We call \mathcal{D}^{-} the corresponding dataset. By the term “dominant color” we mean the most diffuse color in the image area. To extract it from each image we initially exploited the Modified Median Cut Quantization (MMCQ) algorithm [5] with the goal of clustering the color space of the image. Then, we selected as dominant color the one corresponding to the center of the largest cluster thus identified.

At the end of all these procedures, we had two datasets at disposal, namely: (i) \mathcal{D} , storing the images with the original colors, and (ii) \mathcal{D}^{-} , storing the same images as \mathcal{D} but pre-processed to contain

only the dominant color.

For the Color Deep network, we also employed a transfer learning procedure. Specifically, we pre-trained the network with the Car Dataset [22], which contains 16,185 images of colored vehicles. After the pre-training task, we fine-tuned the Color Deep network with the fabric images for identifying the corresponding color. To label car images with the ground truth color, we chose the color list proposed in [27] and used the approach for extracting the dominant color described above, followed by an approach for performing color search. The latter considered a palette of 865 reference colors in the CIELAB color space. Each color of the palette had associated a name and a hexadecimal code. For a given color image, we computed the Euclidean distance between the corresponding color and each color of the palette. We chose the color corresponding to the minimum distance as the one of the image. To evaluate these deep learning models, we divided each original dataset (i.e., \mathcal{D} and \mathcal{D}^-) into two partitions comprising 80% and 20% of the dataset, respectively. We randomly chose the items to be included in each partition. Then we applied 5-fold cross validation to the first partition. More specifically, we divided it into 5 folds. After that, we used each fold in turn as a test set and the remaining folds as training sets. Finally, we computed the values of the performance measures by averaging the corresponding values obtained through the test sets in the different iterations. In this way, for each model, we computed the average values of Precision, Recall, F1-score and Accuracy. After that, we used the second partition (consisting of data never exploited to train the model) for performing a validity check of the model. Specifically, we calculated the values of Precision, Recall, F1-score and Accuracy on this new partition and checked the maximum difference between them and the corresponding values obtained from the first partition. After performing this task, we obtained that the maximum difference in the values of Precision (resp., Recall, F1-score, Accuracy) between the first and second partitions was 0.02 (resp., 0.04, 0.03, 0.015). These values are extremely low and allow us to say that the Precision, Recall, F1-score and Accuracy values obtained by means of 5-fold cross validation were reliable.

We trained each CNN^Δ of $\mathcal{E}_{\text{CNN}\Delta}$ from scratch and independently of the others. Since there are very few trained deep learning models in the color recognition literature, and none of them was suitable for our difference space, we randomly initialized the weights of the three CNN^Δ models.

To limit overfitting, we employed the early stopping technique. For this purpose, we monitored (cross-)validation loss and stopped the training task when there was no change in this metric for three consecutive iterations. To determine the optimal values of the hyperparameters for the three CNN^Δ models, we used a random search procedure [4]. The latter, combined with early stopping strategies, was able to increase the efficiency for narrowing down the search space [37]. Furthermore, in accordance with the results found in [20], we considered low values of the ϵ parameter and adopted the ADAM optimizer for CNNs. In Table 3, we report the parameter setting for the different learning models adopted in our experiments. The last three rows of this table refer to the three CNN^Δ models that compose $\mathcal{E}_{\text{CNN}\Delta}$ and whose parameters were already partially defined in Table 2.

The interested reader can find the fabric dataset, along with the source code to train and evaluate the deep neural networks of our ensemble and those we will use for comparisons (see Section 4.2.1) at the GitHub address <https://github.com/lucav48/fabric-color-classification>.

Finally, we performed our experiments by means of the free version of Google Colab, which provided a GPU NVIDIA Tesla K80, 12 GB RAM, and two Intel Xeon CPUs 2.30GHz. The programming environment consisted of Python 3.7, Tensorflow 2.6, OpenCV 4.2 and Scikit-learn 1.0. In particular, we used the virtual machine mentioned above to test all learning models. Furthermore, we did not adopt any time limit as an exit criterion for any learning model. In fact, we strived to have each learning model operate under conditions allowing it to guarantee the best performance in terms of

<i>Deep neural network</i>	<i>Optimizer</i>	<i>Learning rate η</i>	<i>Decay</i>	<i>Momentum</i>	ϵ
Color Deep	SGD	0.001	10^{-7}	0.9	-
ResNet50v2, ResNet152v2	SGD	0.1	-	0.001	-
DenseNet169v2, DenseNet201v2	SGD	0.1	-	0.001	-
InceptionResNetV2	SGD	0.01	-	0.001	-
CNN $^{\Delta}$	ADAM	0.001	-	-	10^{-7}
CNN $^{\Delta}$ 1 of $\mathcal{E}_{CNN^{\Delta}}$	ADAM	0.001	-	-	10^{-7}
CNN $^{\Delta}$ 2 of $\mathcal{E}_{CNN^{\Delta}}$	ADAM	0.0001	-	-	10^{-8}
CNN $^{\Delta}$ 3 of $\mathcal{E}_{CNN^{\Delta}}$	SGD	0.01	-	Nesterov	-

Table 3: Parameter setting of the deep neural networks employed in our experiments

Precision, Recall, F1-score and Accuracy. In this way, we aimed to ensure fair conditions for all learning models under evaluation.

4.2 Results

Table 4 shows the values of Precision, Recall and F1-score for each class of colors obtained by applying $\mathcal{E}_{CNN^{\Delta}}$ on the test set of \mathcal{D} . Instead, in Table 5, we report the values of these three parameters averaged over all classes, along with the Accuracy value. Finally, Figure 6 shows the corresponding confusion matrix. The values in Table 5 and Figure 6 also refer to the test set of \mathcal{D} .

<i>Color class</i>	<i>Precision</i>	<i>Recall</i>	<i>F1-score</i>
Orange	0.90	0.95	0.92
White	0.91	1.00	0.95
Blue	0.79	0.95	0.86
Cyan	0.67	0.70	0.68
Yellow	0.78	0.70	0.74
Magenta	0.95	0.95	0.95
Black	0.95	1.00	0.97
Red	0.81	0.85	0.83
Earth Brown	0.72	0.65	0.68
Green	0.76	0.65	0.70
Emerald Green	0.72	0.90	0.80
Purple	0.58	0.35	0.44

Table 4: Precision, Recall and F1-score for each class of colors obtained by $\mathcal{E}_{CNN^{\Delta}}$ when applied on the test set of \mathcal{D}

<i>Average Precision</i>	<i>Average Recall</i>	<i>Average F1-score</i>	<i>Accuracy</i>
0.80	0.80	0.80	0.80

Table 5: Average Precision, Average Recall, Average F1-score and Accuracy obtained by $\mathcal{E}_{CNN^{\Delta}}$ when applied on the test set of \mathcal{D}

From the analysis of Tables 4 and 5 and Figure 6, we can see that $\mathcal{E}_{CNN^{\Delta}}$ obtains very satisfactory results. As far as the single colors are concerned, Orange, White, Magenta and Black (bold marked in Table 4) reach the highest values of the performance measures. In fact, their values fall in the range $[0.90, 1]$ and the number of misclassifications involving them in the confusion matrix is very low. Instead, $\mathcal{E}_{CNN^{\Delta}}$ has difficulty in classifying Purple. This can be seen in both Table 4 and the confusion




Figure 6: Confusion matrix obtained by $\mathcal{E}_{CNN\Delta}$ when applied on the test set of \mathcal{D}

matrix of Figure 6. Such a behavior may be due to the texture of some fabric images having lighter purple nuances that make the identification of Purple more difficult. In these cases, the corresponding images were classified with colors, like Blue and Cyan, closely related to Purple in the color palette.

To show the advantage of using an ensemble strategy, in Table 6 we report the results obtained from the three $CNN\Delta$ models separately on the same dataset considered to evaluate $\mathcal{E}_{CNN\Delta}$ previously.

In this case, we can see that the first network (bold marked) obtains the best results, with performance values in the range $[0.76, 0.78]$. Comparing these values with the ones of Table 5, we can see that the adoption of the ensemble strategy produces an improvement of 2.56% in Precision, 5.26% in Recall, 3.89% in F1-score and 5.26% in Accuracy.

4.2.1 Comparisons

Tables 7 and 8 show the results on color identification for our approach and the ones we chose for comparison (see Section 4.1). Specifically, Table 7 refers to the test set of the original fabric images (i.e., images of the test set of \mathcal{D}), while Table 8 refers to the test set of the fabric images with the dominant color (i.e., images of the test set of \mathcal{D}^-). Furthermore, for Color Deep we considered both

$CNNs^\Delta$	Precision	Recall	F1-score	Accuracy
Network 1	0.78	0.76	0.77	0.76
Network 2	0.75	0.75	0.75	0.75
Network 3	0.74	0.73	0.73	0.73

Table 6: Average Precision, Average Recall, Average F1-score and Accuracy obtained by the three $CNNs^\Delta$ composing \mathcal{E}_{CNN^Δ} when working individually on the test set of \mathcal{D}

the average and the best result.

Deep neural network	Precision	Recall	F1-score	Accuracy
\mathcal{E}_{CNN^Δ}	0.80	0.80	0.80	0.80
CNN^Δ	0.78	0.76	0.77	0.76
ResNet50v2	0.74	0.65	0.69	0.71
ResNet152v2	0.67	0.65	0.66	0.65
InceptionResNetV2	0.73	0.69	0.71	0.69
DenseNet169v2	0.72	0.69	0.70	0.69
DenseNet201v2	0.69	0.66	0.67	0.66
Color Deep (avg.)	0.73	0.72	0.72	0.72
Color Deep (max.)	0.78	0.77	0.77	0.77

Table 7: Results of \mathcal{E}_{CNN^Δ} and the other models chosen for comparison when the images of the test set of \mathcal{D} are provided in input

Deep neural network	Precision	Recall	F1-score	Accuracy
\mathcal{E}_{CNN^Δ}	0.79	0.78	0.78	0.80
CNN^Δ	0.70	0.66	0.68	0.67
ResNet50v2	0.56	0.55	0.55	0.56
ResNet152v2	0.59	0.58	0.58	0.59
InceptionResNetV2	0.60	0.59	0.59	0.59
DenseNet169v2	0.59	0.55	0.57	0.56
DenseNet201v2	0.60	0.59	0.59	0.57
Color Deep (avg.)	0.51	0.62	0.56	0.64
Color Deep (max.)	0.66	0.66	0.66	0.67

Table 8: Results of \mathcal{E}_{CNN^Δ} and the other models chosen for comparison when the images of the test set of \mathcal{D}^- are provided in input

Table 7 shows that, for the images of the test set of \mathcal{D} , we can observe that \mathcal{E}_{CNN^Δ} overcomes the other deep learning architectures in all performance measures. Color Deep is able to obtain good performances in its second configuration. However, these are 2-3% lower than those obtained by \mathcal{E}_{CNN^Δ} . Another result emerging from Table 7 is that the ensemble strategy obtains better results than the usage of a single CNN^Δ . In fact, the value of Precision (resp., Recall, F1-score, Accuracy) obtained by \mathcal{E}_{CNN^Δ} is higher of 2.56% (resp., 5.26%, 3.89%, 5.26%) than the corresponding one obtained by a single CNN^Δ .

The examination of Table 8, and its comparison with Table 7, shows that providing input images with dominant color to the models under examination does not lead to improved results. On the contrary, for all the models except \mathcal{E}_{CNN^Δ} , a marked decrease is observed. The best results after the one of \mathcal{E}_{CNN^Δ} are achieved by the single CNN^Δ , whose Precision (resp., Recall, F1-score and Accuracy) value is 11.39% (resp., 15.38%, 12.82% and 16.25%) lower than that obtained by \mathcal{E}_{CNN^Δ} .

The other models perform even worse and, among them, again the best one is Color Deep in its second version. Once again we can observe that the ensemble strategy obtains better results than the use of a single CNN Δ .

Contrary to what we might have expected, using the images with the dominant color obtains worse results than the ones achieved when employing the original fabric images. This can be due to the flattening of some main color nuances, which increases the difficulty to discriminate the true color by the deep learning models into examination.

At this point, we found it necessary to conduct a formal analysis of the results obtained for verifying whether the differences in the performance of the various approaches were statistically significant or not. From Tables 7 and 8, we can see that we have four different performance measures (i.e., Precision, Recall, F1-score and Accuracy) to evaluate the various classifiers. With reference to the current formal analysis, each performance measure provides a dataset of performance values; specifically, there is one value for each classifier in each dataset. The values in the datasets are all real numbers between 0 and 1, regardless of which performance measure they refer to. The higher the value, the better the performance. Since our goal was comparing more classifiers across multiple datasets, we thought it appropriate to carry out our formal analysis by applying the approach reported in [8]. In this paper, the author proposes a set of simple, safe and robust non-parametric tests for statistical comparison of classifiers. In particular, in presence of more than two classifiers and multiple datasets, he recommends using the Friedman test [10], followed by suitable post-hoc tests.

Therefore, we initially used the Friedman test as an omnibus test on the results of Tables 7 and 8. Recall that the omnibus test allows us to test whether, among a set of models, there is at least one that performs statistically significantly different from the others (this is the case if the corresponding p-value is less than 0.05). Applying the Friedman test on the values of Table 7 (resp., Table 8), we obtained a p-value equal to $2.4 \cdot 10^{-4}$ (resp., $5.3 \cdot 10^{-4}$), which is much smaller than 0.05. As a consequence, we can conclude that, in both cases, there is at least one classifier whose performance is statistically significantly different from that of another classifier.

At this point, we proceeded by applying a post-hoc test to find which classifier actually differed. In our case, we decided to use the Tukey test [38] as post-hoc test. The results obtained by applying this test on the data in Table 7 (resp., Table 8) are shown in Figure 7 (resp., Figure 8). In this figure, there is a line for each pair of classifiers. On the x-axis, the mean differences in the performance value are reported. More specifically, given a pair of classifiers, a negative (resp., null, positive) value of the x-axis denotes that the mean difference in the performance values between the first and the second classifier is negative (resp., null, positive). Actually, for each pair of classifiers, not a single value but a confidence interval is reported; the confidence value adopted is 0.95. It follows that, when the confidence interval intersects $x = 0$, it is not possible to conclude that there is a significant difference between the two corresponding classifiers; in this case, in the figure, with the corresponding confidence interval is colored red. In all the other cases, the confidence interval is colored green.

Clearly, we are interested in the confidence intervals colored green. In addition, we are interested in checking whether there is a classifier that always performs better than the other, i.e., one whose confidence interval is always to the right of the graph, when it turns out to be the first element of the pair, and always to the left, when it turns out to be the second element of the pair. As far as Figure 7 is concerned, this never happens. The classifier that comes closest to this condition is $\mathcal{E}_{CNN\Delta}$, which always performs statistically significantly better than the other classifiers except for CNN Δ and Color Deep (max.), in which case it has a comparable performance (and, in fact, the corresponding confidence intervals intersect $x = 0$). Instead, in Figure 8, there is a classifier that performs statistically significantly better than all the other ones, and this classifier is $\mathcal{E}_{CNN\Delta}$.

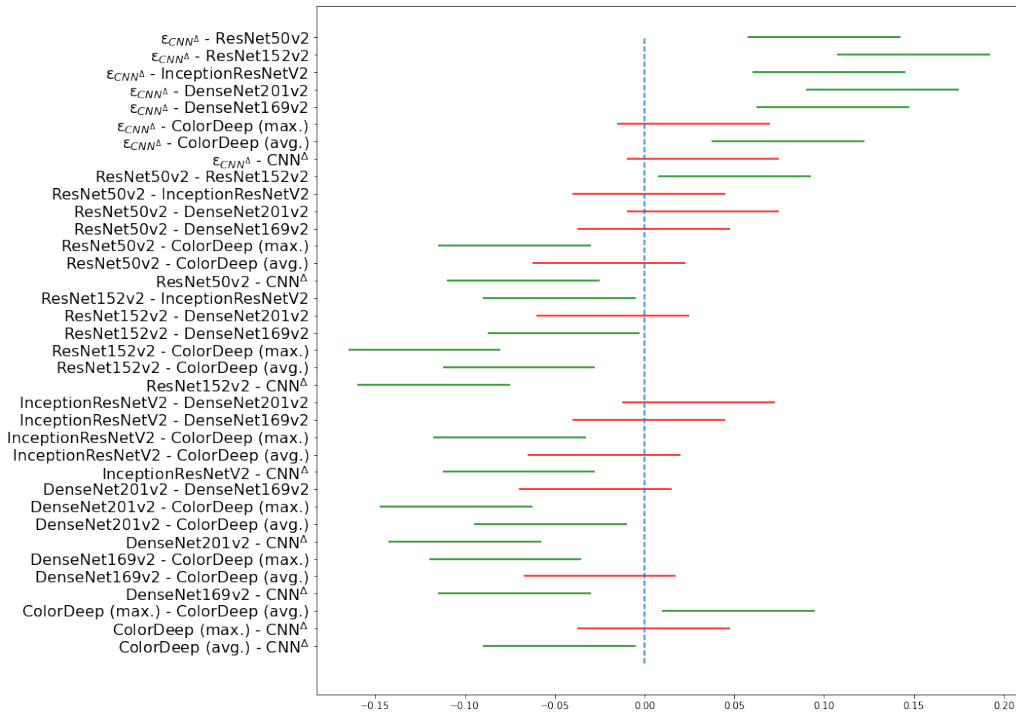


Figure 7: Results of the Tukey test on the data in Table 7




Figure 8: Results of the Tukey test on the data in Table 8




Figure 9: Values of Precision obtained by applying \mathcal{E}_{CNN^Δ} , a single CNN^Δ and Color Deep

Having made a formal analysis of the results obtained, we can now proceed to deepen the analyses reported in Tables 7 and 8. Figure 9 (resp., 10, 11) show the values of Precision (resp., Recall, F1-score) obtained by applying \mathcal{E}_{CNN^Δ} , a single CNN^Δ and Color Deep on the original fabric images. From the examination of this figure we can observe that the Precision of \mathcal{E}_{CNN^Δ} is particularly good for Orange, Magenta, Black, White and Red. Conversely, Color Deep achieves a high Precision for Red, Earth Brown, Green and White. Both models reach a high Precision for White and Red, while both of them obtain a low Precision for Cyan and partly for Yellow. As could be expected from the results reported in Tables 7 and 8, \mathcal{E}_{CNN^Δ} always has a higher Precision than a single CNN^Δ .

Instead, Figure 10 shows that the Recall of \mathcal{E}_{CNN^Δ} is very high for Orange, White, Blue, Magenta, Black, and partially for Red, Yellow and Emerald Green. In contrast, the Recall of this model is low for Purple. Regarding Color Deep, we can see that the Recall is very high for Orange, White, Blue, Magenta, Red and Earth Brown and high for Green and Purple. Color Deep reaches the lowest values of Recall for Cyan and Emerald Green. Orange, White, Blue, Magenta, Black and Red are colors with a high Recall for both models. Once again, we can observe that \mathcal{E}_{CNN^Δ} always reaches a better Recall than a single CNN^Δ .

Finally, examining Figure 11, we deduce that the F1-score of \mathcal{E}_{CNN^Δ} is very high for Orange, White, Blue, Magenta, Black and Red and is high for Yellow and Emerald Green. Instead, it is low for Purple. Color Deep has high values of F1-score for Orange, White, Blue, Red, Brown and Green while it has low values for Cyan. We also observe that Orange, White, Blue and Red show high values of F1-score for both approaches. Interestingly, there is no color that shows low values of F1-score for both \mathcal{E}_{CNN^Δ} and Color Deep. This result could be usefully exploited in future developments, as we will see below. Finally, the values of F1-score are always better for \mathcal{E}_{CNN^Δ} than for a single CNN^Δ .

As a final experiment, we computed the average epoch training time required for each deep learning model into evaluation. The results obtained are shown in Figure 12.

From the analysis of this figure, we can observe that CNN^Δ requires the lowest average epoch training time, while ResNet152v2 requires the highest one. A key factor affecting the epoch training

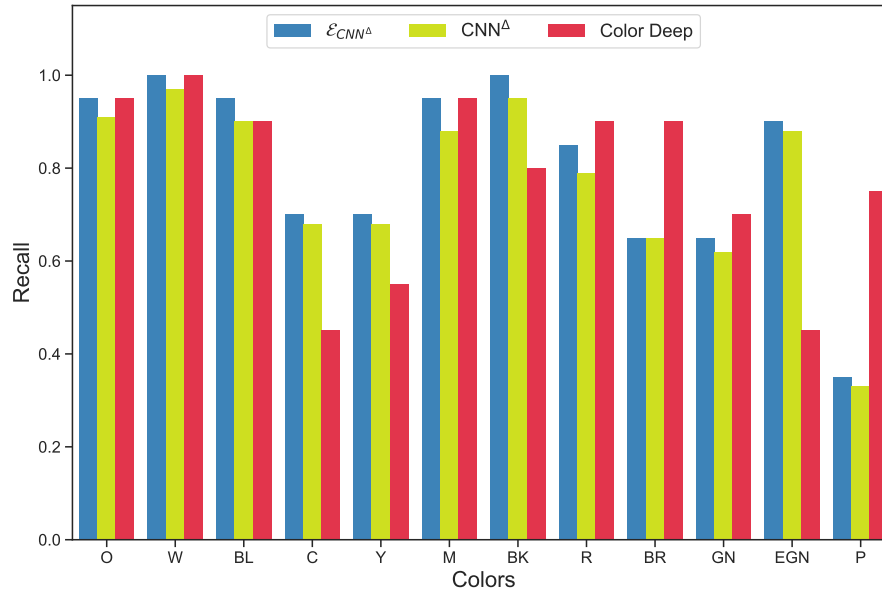


Figure 10: Values of Recall obtained by applying \mathcal{E}_{CNN^Δ} , a single CNN^Δ and Color Deep




Figure 11: Values of F1-score obtained by applying \mathcal{E}_{CNN^Δ} , a single CNN^Δ and Color Deep

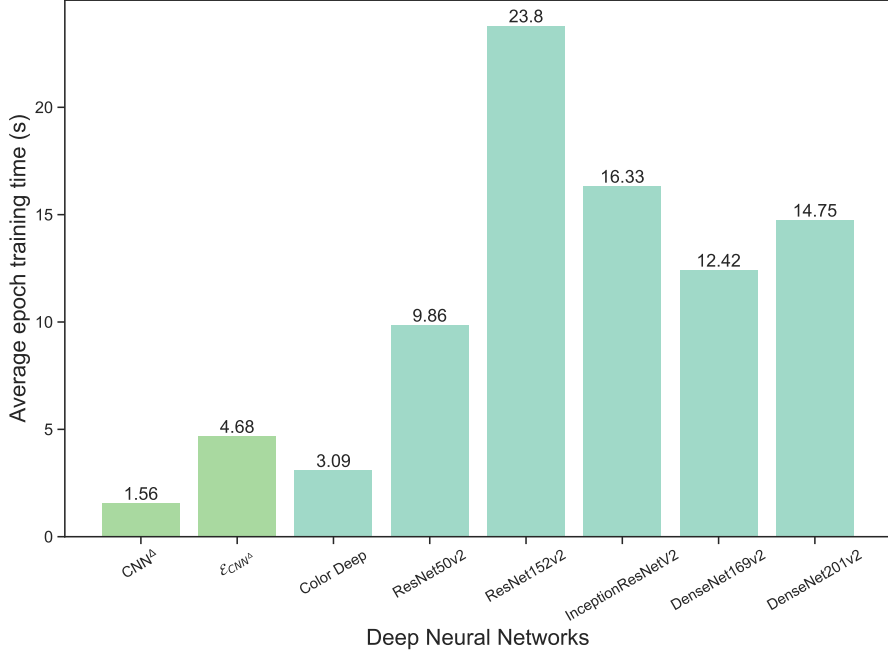


Figure 12: Average epoch training time for each deep learning model

time is the number of layers in these models. In fact, ResNet50v2, ResNet152v2, InceptionResNetv2, DenseNet169v2 and DenseNet201v2 have more convolutional layers than CNN^Δ , \mathcal{E}_{CNN^Δ} and Color Deep, and, furthermore, these layers are also denser. Moreover, note that \mathcal{E}_{CNN^Δ} , which is a combination of three different CNN^Δ models, has a competitive training time compared with Color Deep. More specifically, the average epoch training time of CNN^Δ is 50.49% less than that of Color Deep. On the other hand, the average epoch training time of \mathcal{E}_{CNN^Δ} is 51.49% higher than that of Color Deep. This increase can be explained by considering that \mathcal{E}_{CNN^Δ} is an ensemble of three different deep learning models. This fact, while making the average epoch training time of \mathcal{E}_{CNN^Δ} slightly longer than the one of Color Deep, also enables it to achieve a higher performance in color classification.

In any case, we can observe that there is a certain complementarity between the performances of \mathcal{E}_{CNN^Δ} and those of Color Deep. In fact, one of the two models often returns the best performances, in terms of Precision, Recall and F1-score, exactly for those colors where the other proves somewhat weaker. Also, the average epoch training time of Color Deep is not excessively higher than that of CNN^Δ and is somewhat lower than that of \mathcal{E}_{CNN^Δ} . This suggests us that, in the future, it is possible to define new ensembles involving both CNN^Δ and Color Deep and verify if they can achieve even better results than the ones reached by \mathcal{E}_{CNN^Δ} in this paper.

4.3 Discussion

As shown in Section 4.2, \mathcal{E}_{CNN^Δ} achieves very good results. These performances are reached thanks to two expedients, namely the difference method and ensemble learning. Working with fabrics is not easy because they often present textures, which can make color recognition very difficult. However, the two expedients mentioned above allow us to increase the amount of information provided in input to \mathcal{E}_{CNN^Δ} , without the risk of overfitting it.

The important role played by the difference method and ensemble learning in achieving high

performance is evident when comparing $\mathcal{E}_{CNN\Delta}$ with other deep learning architectures. In fact, we have seen that the difference method allows a single CNN^Δ and $\mathcal{E}_{CNN\Delta}$ to perform better than the other models. The usage of ensemble learning in $\mathcal{E}_{CNN\Delta}$ allows it to perform better than a single CNN^Δ . This shows that both the expedients used in the definition of $\mathcal{E}_{CNN\Delta}$ help to achieve the goals we had set for ourselves.

Our approach has a high potential in several scenarios. For example, it could be integrated into a recommender system that matches and suggests clothes to create a pleasing outfit. In turn, this system could be integrated into an app that could suggest what to wear based on the clothes in the closet. In addition, it could be used to help color blind people to recognize the correct colors of clothes.

However, besides these merits, our approach has also some limitations, which represent the subject of future investigations. First, in this paper, we considered a small number of colors (i.e., 12) in our experiments. These are the most frequent ones present in fabric images, but they are a fraction of the colors that can be found in this kind of image and that we could use in future experiments and applications. It is clear that increasing the number of reference colors, and therefore the number of classes, makes the classification task much more complex. Already a classification with 12 possible classes is not very common in the machine learning context. A further increase in the number of classes is possible, but it requires much more effort and is subject to many more errors. It must be taken into account that, continuing in this direction, the efforts needed could increase very much and the performance might decrease. So, at some point, a cost-benefit analysis must be done and the suitable trade-off must be found.

Going into more detail in the examination of results, we have seen that $\mathcal{E}_{CNN\Delta}$ has unsatisfactory performances on some colors (e.g., Purple and Cyan). This contributes to lower its overall performance but, more importantly, makes it not fully adequate in classifying fabrics whose color could fall into one of those classes in which it showed low performances.

5 Conclusion

In this paper, we have proposed a new approach for identifying fabric colors. We have seen that such a task is not trivial, because fabric images contain textural patterns that can make color classification hard. To address this issue, we have introduced a new color space, called difference space. Then, we have seen that this way of proceeding allows us to overcome the limitations of the previous representations and makes color recognition much more effective. Indeed, such an expedient allows us to provide much more information as input to the deep learning system.

Based on these results, we have designed a Convolutional Neural Network model, called CNN^Δ , to operate in this new color space. To further improve the performance, we have used ensemble learning and we have combined three versions of CNN^Δ models to build an ensemble model called $\mathcal{E}_{CNN\Delta}$. Then, we have carried out some experiments on the Fabrics Dataset [16] that showed how the adoption of the two expedients (difference method and ensemble strategy) allows $\mathcal{E}_{CNN\Delta}$ to obtain very satisfactory results, better than those reached by other learning models already proposed in the literature.

This paper should not be considered an ending point but, rather, a starting point for further research efforts. First, as we said above, we plan to increase the number of reference colors used in the classification task. Then, we think to extend our model in such a way that it can assign more colors to a single fabric image. In fact, there are fabrics with several colors and, therefore, it would be important to have an approach capable of handling this situation. These two factors could make the

problem of color identification in fabric images much more complex and could provide many interesting insights for researchers who want to address this issue in the future. Also, as we have seen in Section 4.2.1, there is some complementarity between CNN^Δ and Color Deep. In fact, one of the two models often returns the best performances exactly for those colors where the other proves somewhat weaker. This suggests us to verify the possibility to build an ensemble that includes CNN^Δ and Color Deep in order to exploit the strengths of both of them and reduce their weaknesses. Furthermore, if we had access to a much larger database than the Fabrics Dataset, it would be interesting to evaluate the application of a Vision Transformers based methodology [17, 43] for color classification of fabric images. In fact, we consider this architecture very promising, but it needs thousands of images for its application. It is our intention to try it in the future. Finally, as a further future development, we plan to evaluate the usage of different color spaces in our model. In fact, we currently use the RGB color space. However, we believe that the inclusion of other color spaces, such as HSL (Hue, Saturation, Lightness), HSV (Hue, Saturation, Value) or CMYK (Cyan, Magenta, Yellow and Key Black) could lead to an improvement in the accuracy of our color identification approach.

References

- [1] B.S. Anami and M.C. Elemmi. A rule based approach for classification of shades of basic colors of fabric images. *International Journal of Signal Processing, Image Processing and Pattern Recognition*, 8(2):389–400, 2015.
- [2] D.H. Apriyanti, L.J. Spreeuwens, P.J. Lucas, and R.N. Veldhuis. Automated color detection in orchids using color labels and deep learning. *PLoS ONE*, 16(10):e0259036, oct 2021.
- [3] S. Basar, A. Adnan, N.H. Khan, and S. Haider. Color image segmentation using k-means classification on rgb histogram. *Recent Advances In Telecommunications, Informatics and Educational Technologies*, pages 257–262, 2014.
- [4] J. Bergstra and Y. Bengio. Random Search for Hyper-Parameter Optimization. *Journal of Machine Learning Research*, 13(2):281–305, 2012. JMLR.org.
- [5] D.S. Bloomberg. Color quantization using modified median cut. In *Leptonica*, pages 1–10. 2008.
- [6] D. Borza, T. Ileni, and A. Darabant. A Deep Learning Approach to Hair Segmentation and Color Extraction from Facial Images. In *Proc. of the International Conference on Advanced Concepts for Intelligent Vision Systems (ACIVS’18)*, pages 438–449, Auckland, New Zealand, 2018. Springer.
- [7] P. Chen, X. Bai, and W. Liu. Vehicle color recognition on urban road by feature context. *IEEE Transactions on Intelligent Transportation Systems*, 15(5):2340–2346, 2014. IEEE.
- [8] J. Demšar. Statistical comparisons of classifiers over multiple data sets. *The Journal of Machine Learning Research*, 7:1–30, 2006. JMLR.
- [9] Z. Di, Z. Luoqing, and S. Jun. A novel feedback error-correcting algorithm for automatic recognition of the color and weave pattern of yarn-dyed fabrics. *Textile Research Journal*, 83:1673 – 1689, 2013.
- [10] M. Friedman. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the American Statistical Association*, 32(200):675–701, 1937. Taylor & Francis.
- [11] A.K. Gupta and D.J. Bora. A novel color image segmentation approach based on k-means clustering with proper determination of the number of clusters and suitable distance metric. *International Journal of Computer Science & Engineering Technology*, 7(09):395–409, 2016.
- [12] M. Hajjarbabi and A. Agah. Human skin color detection using neural networks. *Journal of Intelligent Systems*, 24(4):425–436, 2015. De Gruyter.
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *Proc. of the International Conference on Computer Vision and Pattern Recognition (CVPR’16)*, pages 770–778, Las Vegas, NV, USA, 2016.
- [14] C. Hu, X. Bai, L. Qi, P. Chen, G. Xue, and L. Mei. Vehicle color recognition with spatial pyramid deep learning. *IEEE Transactions on Intelligent Transportation Systems*, 16(5):2925–2934, 2015. IEEE.

- [15] G. Huang, Z. Liu, L. Van Der Maaten, and K.Q. Weinberger. Densely Connected Convolutional Networks. In *Proc. of the International Conference on Computer Vision and Pattern Recognition (CVPR'17)*, pages 2261–2269, Honolulu, HI, USA, 2017. IEEE Computer Society.
- [16] C. Kampouris, S. Zafeiriou, A. Ghosh, and S. Malassiotis. Fine-Grained Material Classification Using Microgeometry and Reflectance. In *Proc. of the European Conference on Computer Vision (ECCV'16)*, pages 778–792, Amsterdam, The Netherlands, 2016. Springer.
- [17] S. Khan, M. Naseer, M. Hayat, S.W. Zamir, F.S. Khan, and M. Shah. Transformers in vision: A survey. *ACM Computer Surveys*, 2021. ACM.
- [18] D.K. Kim. Color Detection Using Gaussian Mixture Model. *Journal of Theoretical and Applied Information Technology*, 95(17):4313–4320, 2017.
- [19] H.K. Kim, J.H. Park, and H.Y. Jung. An Efficient Color Space for Deep-Learning Based Traffic Light Recognition. *Journal of Advanced Transportation*, page 2365414, 2018.
- [20] D.P. Kingma and J.L. Ba. Adam. A Method for Stochastic Optimization. In *Proc. of the International Conference on Learning Representations (ICLR'15)*, San Diego, CA, USA, 2015.
- [21] C.F.J. Kou, C.Y. Shih, C.Y. Kao, and J.Y. Lee. Color and Pattern Analysis of Printed Fabric by an Unsupervised Clustering Method. *Textile Research Journal*, 75(1):9–12, 2005.
- [22] J. Krause, M. Stark, J. Deng, and L. Fei-Fei. 3D Object Representations for Fine-Grained Categorization. In *Proc. of the International Conference on Computer Vision Workshops (ICCV'13)*, pages 554–561, Sydney, Australia, 2013.
- [23] Z. Li, G. Golwala, S. Sundaram, and J. Allebach. Use of Color Information in the Analysis of Fashion Photographs. *Electronic Imaging*, 2018(10):446–1, 2018. Society for Imaging Science and Technology.
- [24] B. Liu, J. Gan, B. Wen, Y. LiuFu, and W. Gao. An automatic coloring method for ethnic costume sketches based on generative adversarial networks. *Applied Soft Computing*, 98:106786, 2021. Elsevier.
- [25] G.H. Liu and J.Y. Yang. Content-based image retrieval using color difference histogram. *Pattern Recognition*, 46(1):188–198, 2013. Elsevier.
- [26] A.G. Oskouei, M. Hashemzadeh, B. Asheghi, and M.A. Balafar. CGFFCM: CLuster-weight and group-local feature-weight learning in fuzzy C-means clustering algorithm for color image segmentation. *Applied Soft Computing*, 113:108005, 2021. Elsevier.
- [27] R.F. Rachmadi and I. Purnama. Vehicle Color Recognition using Convolutional Neural Network. *arXiv preprint*, abs/1510.07391, 2015.
- [28] A. Rame, A. Douillard, and C. Ollion. CORE: Color Regression for Multiple Colors Fashion Garments. *ArXiv preprint*, abs/2010.02849, 2020.
- [29] O. Sagi and L. Rokach. Ensemble learning: A survey. *WIREs Data Mining and Knowledge Discovery*, 8(4):e1249, 2018.
- [30] T. Sandhan and J. Y. Choi. Anti-Glare: Tightly Constrained Optimization for Eyeglass Reflection Removal. In *Proc. of the International Conference on Computer Vision and Pattern Recognition (CVPR'17)*, pages 1675–1684, Honolulu, HI, USA, 2017.
- [31] A. Shams-Nateri and E. Hasanlou. Computer vision techniques for measuring and demonstrating color of textile. In W.K. Wong, editor, *Applications of Computer Vision in Fashion and Textiles*, The Textile Institute Book Series, pages 189–220. Woodhead Publishing, 2018.
- [32] J. Shen, H. Ma, W. Chen, and X. Zhou. A novel analysis of color component for top dyed melange yarn with support vector machine. *Color Research & Application*, 41(6):636–641, 2016.
- [33] S.M. Singh and K. Hemachandran. Content-based image retrieval using color moment and Gabor texture feature. *International Journal of Computer Science Issues*, 9(5):299, 2012.
- [34] B. Su, J. Shao, J. Zhou, X. Zhang, and L. Mei. Vehicle color recognition in the surveillance with deep convolutional neural networks. In *Proc. of the Joint International Mechanical, Electronic and Information Technology Conference (JIMET 2015)*, Chongqing, China, 2015.
- [35] W. Sun, J.K. Hao, Y. Zang, and X. Lai. A solution-driven multilevel approach for graph coloring. *Applied Soft Computing*, 104:107174, 2021. Elsevier.

- [36] C. Szegedy, L. Wei, J. Yangqing, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proc. of the International Conference on Computer Vision and Pattern Recognition (CVPR'15)*, pages 1–9, Boston, MA, USA, 2015.
- [37] Y. Tong and Z. Hong. Hyper-parameter optimization: A review of algorithms and applications. *CoRR*, abs/2003.05689, 2020. arXiv.
- [38] J.W. Tukey. Comparing individual means in the analysis of variance. *Biometrics*, pages 99–114, 1949. JSTOR.
- [39] T. Vafeiadis, N. Kolokas, N. Dimitriou, A. Zacharaki, M. Yildirim, H. G. Selvi, D. Ioannidis, and D. Tzovaras. A comparison of 2DCNN network architectures and boosting techniques for regression-based textile whiteness estimation. *Simulation Modelling Practice and Theory*, 114:102400, 2022.
- [40] M. Veluchamy and B. Subramani. Fuzzy dissimilarity color histogram equalization for contrast enhancement and color correction. *Applied Soft Computing*, 89:106077, 2020. Elsevier.
- [41] P.S. Vitthal, S. Balasubramanian, and P.S. Mane. Color Analysis and Classification Based on Machine Learning Technique Using RGB Camera Industrial Practice and Experience Paper. In *Proc. of the International Conference on Computational Intelligence and Computing Research (ICIC'17)*, pages 1–5, Coimbatore, India, 2017.
- [42] L. Xie, Y. Shen, and X. Chen. Study on Automatic Recognition of Fabric Color and Matching to Standard Color Chip by Computer Vision and Image Analysis Technology. *Journal of Fiber Bioengineering and Informatics*, 9(1):29–37, 2016.
- [43] Y. Xu, H. Wei, M. Lin, Y. Deng, K. Sheng, M. Zhang, F. Tang, W. Dong, F. Huang, and C. Xu. Transformers in computational visual media: A survey. *Computational Visual Media*, 8(1):33–62, 2022. Springer.
- [44] Z. Yang, L. Leng, and B.G. Kim. StoolNet for Color Classification of Stool Medical Images. *Electronics*, 8(12):1464, 2019.
- [45] H. Ye, L. Zheng, and P. Liu. Color detection and segmentation of the scene based on Gaussian mixture model clustering. In *Proc. of the International Conference on Electronics Information and Emergency Communication (ICEIEC'17)*, pages 503–506, Shenzhen, China, 2017.
- [46] Z. Zehra and M.N. Bashir. Color Fastness Grading System for Textile Industry Using CIEL*a*b Color Space. In *Proc. of the International Conference on Big Data and Smart City (ICBDSC'19)*, pages 1–5, Yokohama, Japan, 2019.
- [47] J. Zhang, R. Pan, E. Gao, and D. Zhu. Automatic detection of layout of color yarns of yarn-dyed fabric. Part 1: Single-system-mélange color fabrics. *Color Research & Application*, 40(6):626–636, 2015.
- [48] J. Zhang, R. Pan, W. Gao, and D. Zhu. Automatic recognition of the color effect of yarn-dyed fabric by the smallest repeat unit recognition algorithm. *Textile Research Journal*, 85:432 – 446, 2015.
- [49] L. Zhang, C. Long, X. Zhang, and C. Xiao. RIS-GAN: Explore Residual and Illumination with Generative Adversarial Networks for Shadow Removal. In *Proc. of the International Conference on Artificial Intelligence (AAAI'20)*, volume 34, pages 12829–12836, New York, NY, USA, 2020.
- [50] M. Zhang and Y. Jia. Color Recognition Model Based on Multiple SVMs for Bobbin Sorting Machine. In *Proc. of the International Symposium on Computational Intelligence and Design (ISCID'12)*, volume 2, pages 218–221, Washington DC, USA, 2012.
- [51] M. Zhang, P. Wang, and X. Zhang. Vehicle Color Recognition Using Deep Convolutional Neural Networks. In *Proc. of the International Conference on Artificial Intelligence and Computer Science (AICS'19)*, pages 236–238, Wuhan, China, 2019. Association for Computing Machinery.
- [52] Q. Zhang, L. Zhuo, J. Li, J. Zhang, H. Zhang, and X. Li. Vehicle color recognition using Multiple-Layer Feature Representations of lightweight convolutional neural network. *Signal Processing*, 147:146–153, 2018.
- [53] Z. Zhuang, Y. Liu, F. Ding, and Z. Wang. Online Color Classification System of Solid Wood Flooring Based on Characteristic Features. *Sensors*, 21(2):336, 2021.