DLE4FC: a Deep Learning Ensemble to Identify Fabric Colors

(Discussion Paper)

Alessia Amelio[1,†], Gianluca Bonifazi[2,†], Francesco Cauteruccio[2,†], Enrico Corradini[2,†], Michele Marchetti[2,†], Domenico Ursino[2,*,†] and Luca Virgili[2,†]

[1] *INGEO, University "G. D'Annunzio" of Chieti-Pescara*
[2] *DII, Polytechnic University of Marche*

Abstract

The study of colors has attracted Artificial Intelligence researchers for many years. Nevertheless, there are aspects of this issue that are still little analyzed. One of them is the investigation of fabric colors, which has some peculiarities, such as the necessity to handle textures, that are not found in other scenarios. In this paper, we present DLE4FC, a deep learning ensemble for identifying fabric colors. Specifically, we introduce the general basic model, which consists of a particular Convolutional Neural Network, define three versions of it and integrate them into an ensemble to get better results. Finally, we test our ensemble and compare it with other already known systems.

Keywords

Color Classification, Ensemble Learning, Identification of Fabric Colors, Classification of Fabric Colors, Convolutional Neural Networks

1. Introduction

The study of colors has attracted, and is attracting, researchers from several areas [1, 2]. One of the most investigated issues concerns the identification of colors. An application context most interested in this issue is textile industry. In fact, in this industry, color plays a key role in manufacturer-customer relationships, as well as in R&D and marketing activities [3].

Various methods have been proposed in the literature to recognize color components, patterns and shades, as well as the layout of color yarns from images. To achieve their goals, these methods use very different approaches, such as fuzzy C-Means [4] and X-means [5], rule-based techniques [6], Support Vector Machine [7], Random Forest [8], eXtreme Gradient Boosting (XGBoost) [9]. Other methods have been proposed to estimate textile whiteness [10]. Finally,

SEBD 2023: 31st Symposium on Advanced Database Systems, July 02–05, 2023, Galzignano Terme, Padua, Italy

*Corresponding author.

† These authors contributed equally.

a.amelio@unich.it (A. Amelio); g.bonifazi@univpm.it (G. Bonifazi); f.cauteruccio@univpm.it (F. Cauteruccio); e.corradini@univpm.it (E. Corradini); m.marchetti@pm.univpm.it (M. Marchetti); d.ursino@univpm.it (D. Ursino); luca.virgili@univpm.it (L. Virgili)

0000-0002-3568-636X (A. Amelio); 0000-0002-1947-8667 (G. Bonifazi); 0000-0001-8400-1083 (F. Cauteruccio); 0000-0002-1140-4209 (E. Corradini); 0000-0003-3692-3600 (M. Marchetti); 0000-0003-1360-8499 (D. Ursino); 0000-0003-1509-783X (L. Virgili)

others have been thought to estimate tolerance with respect to standard fabrics by matching the color of interest with a reference palette [11, 12].

In the past, deep learning has been used for color identification in various contexts, including vehicles and traffic lights [13, 14], flowers [15], stool medical images [16], facial images [17] etc. In all these cases they have shown better behavior than other machine learning approaches [13]. Nevertheless, these methods have been little used for color recognition in the context of textiles.

In this paper, we want to make a contribution in this setting and propose DLE4FC, a new deep-learning based approach for identifying colors in textiles. The architecture of DLE4FC is based on an ensemble of Convolutional Neural Networks (CNNs, for short) whose input belongs to the color difference domain. The latter is obtained by considering the difference between the input image and a set of reference color images. Using it allows for better capture of color variations, shades and patterns, and ultimately allows DLE4FC to learn features for color recognition in fabric images. In addition, adopting an ensemble strategy ensures greater robustness, and ultimately greater accuracy in the results obtained. DLE4FC focuses on a particular aspect of color recognition in the textile industry, namely color identification of fabric images. This aspect is important in this application context as it reduces inefficiency, wasted time and subjectivity. This ensures higher quality products capable of better meeting customer desires. In the paper, we also describe some of the experiments we conducted to evaluate the accuracy of DLE4FC compared to a variety of other color identification systems proposed in the past literature. The results obtained show that the peculiarities of DLE4FC enable it to achieve better results than pattern recognition and color identification frameworks proposed in the past.

The outline of this paper is as follows: Section 2 provides a technical presentation of DLE4FC. Section 3 presents the tests performed to evaluate its accuracy while also comparing the latter with the one of other systems proposed in the past. Finally, Section 4 draws conclusions and outlines some possible future developments.

2. Description of DLE4FC

In this section, we present the technical details of DLE4FC. In particular, in Subsection 2.1, we introduce the color difference domain underlying DLE4FC. In Subsection 2.2, we describe in detail the architecture and behavior of our framework.

2.1. The color difference domain underlying DLE4FC

A color space is a specific color organization based on a color model. Currently, the most widely used color model is RGB (Red, Green, Blue). Although it is sufficient for representing digital images, it does not allow the extraction of handcrafted features, which play an important role in the color classification task [18, 6, 19].

To the best of our knowledge, deep learning approaches for color identification proposed in the past do not consider textural patterns (which are frequent in fabric images). In fact, the latter can hamper the learning of the underlying neural network. This limitation stems at least in part from the color representation used for the images, which does not provide sufficient information for color classification [18, 6, 19]. To overcome this problem, DLE4FC introduces a new color space called color difference space. This space is based on calculating the distances between the colors of an input image and a set of reference colors. To define it, we first choose a set of p reference colors and determine the corresponding RGB encoding. Then, we calculate the differences between the input image colors and the reference ones and obtain p images represented using RGB encoding. Then we concatenate all the p images thus obtained to create the input to DLE4FC. The images of this input have the same width and height as the initial image, but now there is a third dimension of size $p \times 3$, obtained by concatenating the p images each represented by means of the 3 RGB channels.

To construct our difference space, we used 12 reference colors, namely Orange, White, Blue, Cyan, Yellow, Magenta, Black, Red, Earth Brown, Green, Emerald Green, and Purple. The selection of these colors was not straightforward since there would be thousands of potential colors. To make this selection we consulted some managers from fabric companies, who provided us with two important insights. The first is that colors used in fabrics have a slightly different RGB representation from the pure color definition. For example, the pure black color corresponds to the hexadecimal code #000000. In contrast, in the context of textiles, there are several shades of black, the most widely used of which is that corresponding to the hexadecimal RGB code #1C1C1C. The second insight concerns the number p of colors to be adopted. The experts interviewed told us that 12 colors are capable of handling most of the possible cases and represent a very good tradeoff between the ability to represent reality, the complexity of the framework and the accuracy of its results. However, we would like to emphasize that our model can be easily scaled to include more colors, should that be necessary in this or other application contexts.

As a consequence of our way of proceeding, for each image we create 12 images, each representing its distance from one of the reference colors, expressed in the RGB model. To give an idea of this, in Figure 1 we show an example of the differences between an input image and the 12 reference colors. Finally, we concatenate the resulting 12 images to give them as input to DLE4FC.

2.2. Architecture and behavior of DLE4FC

In the previous section, we have seen that, given a fabric image, when we switch to the color difference domain, we get 12 images, each represented in the RGB model. Since this model has 3 channels, it follows that each fabric image can be represented by 36 channels. To process such data we construct an extension of the traditional CNN which we call CNN$^{\Delta}$. It has an input layer that receives images of any width and height through 36 channels. After this layer, there is a convolutional layer with kernels of size 3×3 and stride 1, followed by a max pooling layer of size 2×2 and stride 1. After that there is a convolutional layer with the same kernel size as the previous one and, then, another pooling layer of size 2×2 and stride 1. The latter is followed by a dense layer with 128 units connected to a final dense layer with a softmax activation function. The output of CNN$^{\Delta}$ is a vector of 12 elements each representing the classification probability of the corresponding reference color. In total, CNN$^{\Delta}$ has 175,476 parameters.

Actually, as indicated by its name, DLE4FC consists of an ensemble of three CNN$^{\Delta}$ models. This choice allows our system to improve its performance and generalization ability [20]. To tune the values of the hyperparameters of the three CNN$^{\Delta}$ models composing DLE4FC, we

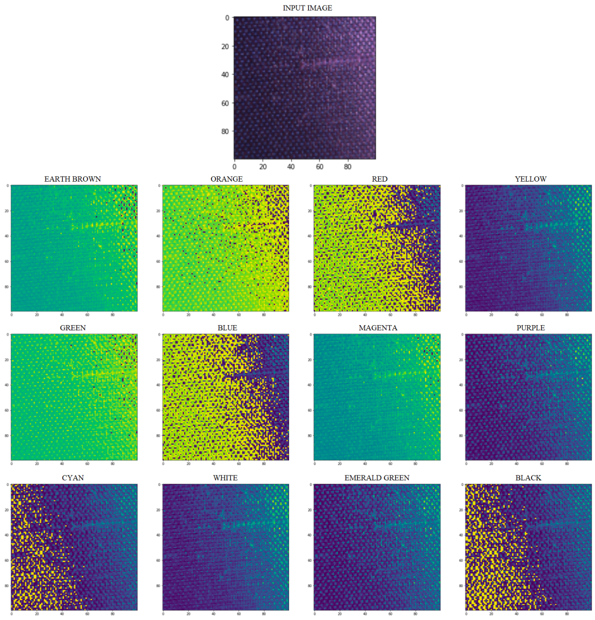

Figure 1: An input image and the differences between it and the set of reference colors

performed a random search procedure. We chose this procedure because, compared with other existing hyperparameter optimization techniques, it can be easily combined with early stopping methods, which we adopted during the training phase for reducing the overfitting probability. The combination of these two techniques allows us to narrow the search space more efficiently.

For space limitations, we cannot describe in detail the whole implementation of this technique. The interested reader can refer to [21]. Here, we only say that, at the end of this procedure, the values of the hyperparameters that provided the best performance in terms of F1-score, Accuracy and Average epoch training time of the ensemble are those shown in Table 1.

Network	*Optimizer*	*Learning rate η*	*Momentum*	ϵ
CNN$^{\Delta}$ 1	ADAM	0.001	-	10^{-7}
CNN$^{\Delta}$ 2	ADAM	0.0001	-	10^{-8}
CNN$^{\Delta}$ 3	SGD	0.01	Nesterov	-

Table 1
Main characteristics of the three CNN$^{\Delta}$ models used in DLE4FC

Having discussed the structure of DLE4FC, we can now describe its behavior. It is shown in Figure 2. As can be seen from this figure, first the input image is transformed into 12 corresponding images in the color difference domain. Then, these images are given as input to the three CNN$^{\Delta}$ models, each of which returns a classification probability vector for the reference colors. In this way, three probability vectors are obtained. These are processed by a soft-voting function that first averages the classification probabilities, then identifies the highest values thus obtained and finally returns the corresponding class as output.

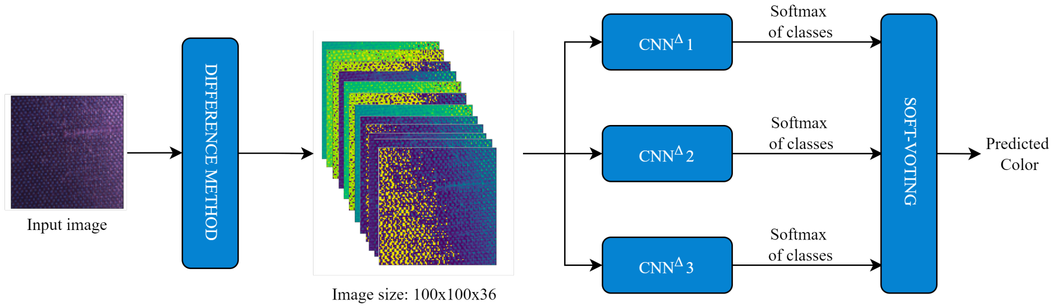

Figure 2: Representation of the behavior of DLEFC

3. Experiments

We tested DLE4FC on Fabric Dataset [22]. This contains 2000 color images concerning clothing and textiles. Each image was acquired using a photometric stereo sensor under 3 or 4 different illumination conditions. Therefore, the total number of actual images is 7,757. The images are in PNG format and their size is 400×400. They are not labeled with colors so we had to manually specify the ground truth for each of them. To achieve this goal, we showed each image to 11 different people and asked each of them to label it with one of the 12 reference colors specified in Section 2.1. Where different people assigned different colors to the same image, we assigned it the color specified by the majority of people. Finally, we discarded all images to which no color could be associated. At the end of this procedure we retained 4,591 images and discarded 3,166 ones. Afterwards, we performed an undersampling task on the most numerous classes, and an oversampling one on the least numerous classes, in order to balance them. At the end of these activities, we obtained a new dataset of 1,200 images, with 100 images per class. We call $\mathcal{D}$ this new dataset.

We compared DLE4FC with several deep learning architectures proposed in the literature, i.e. ResNet [23], InceptionResNet [24], DenseNet [25], as well as Color Deep network [26], which is a reference network in the color recognition literature. Specifically, we selected two types of ResNet, namely ResNet50v2 and ResNet152v2, two types of DenseNet, namely DenseNet169v2 and DenseNet201v2, as well as InceptionResNetv2. We chose these networks because they have been shown to achieve the best experimental results. In addition, to test whether it was worthwhile to adopt an ensemble, we compared DLE4FC with a single CNN$^{\Delta}$. To all these networks we gave as input the 36 channels representing a single RGB image in the color difference domain. Since Color Deep could not receive 36 channels as input simultaneously, we gave it the 12 images corresponding to the differences in cascade, and then aggregated the corresponding results. To avoid overfitting in color classification, we applied a data augmentation procedure to the images of $\mathcal{D}$. This procedure involves two steps, namely: *(i)* a random flip augmentation, which flips the image horizontally and vertically, and *(ii)* a random zoom augmentation, which randomly increases the image size by adding new pixels around or interpolating the values of pixels.

At this point, for each original fabric image, we identified the dominant color, i.e., the most prevalent color in the image area. For this purpose, we used the Modified Median Cut Quantization (MMCQ) algorithm [27]. In this way, we clustered the color space of the image and selected as the dominant color the one corresponding to the center of the largest cluster thus identified. Following this, we generated a new image containing only that color. We call $\mathcal{D}^-$ the corresponding dataset. At the end of all these activities, we had two datasets available for experiments, namely: *(i)* $\mathcal{D}$, which stores the original images, and *(ii)* $\mathcal{D}^-$, which stores the same images but pre-processed to contain only the dominant color.

We provided the images of $\mathcal{D}$ and $\mathcal{D}^-$ as input to the various models to be tested. For each image, we calculated the Euclidean distance between the corresponding color and each color in the palette of 12 colors that we had previously identified. We chose the color corresponding to the minimum distance as the image color.

To perform our tests, we divided each of the two datasets $\mathcal{D}$ and $\mathcal{D}^-$ into two partitions comprising 80% and 20% of their images, respectively. We randomly chose the images to be included in each partition. Finally, we applied the 5-fold cross-validation technique to the first partition.

We trained each CNN$^\Delta$ of DLE4FC from scratch and independently of the others. In the color recognition literature there are few trained learning models, and none of them were suitable for our color difference space. Therefore, we had to randomly initialize the weights of the three CNN$^\Delta$ models. To limit overfitting, we used the early stopping technique. Specifically, we stopped the training task when there was no change in the (cross-)validation loss for three consecutive iterations. In Table 2, we report the parameter setting for the various learning models to be compared. The last three rows of this table refer to the three CNN$^\Delta$ models composing DLE4FC. The hyperparameters of these models have been already shown in Table 1; we report them again in Table 2 for convenience.

Deep neural network	*Optimizer*	*Learning rate* η	*Decay*	*Momentum*	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 DLE4FC	ADAM	0.001	-	-	10^{-7}
CNN$^\Delta$ 2 of DLE4FC	ADAM	0.0001	-	-	10^{-8}
CNN$^\Delta$ 3 of DLE4FC	SGD	0.01	-	Nesterov	-

Table 2
Parameter setting of the deep neural networks employed in our experiments

Table 3 (resp., 4) shows the values of Precision, Recall, F1-score and Accuracy obtained by the various learning models when they receive as input the images of $\mathcal{D}$ (resp., $\mathcal{D}^-$). For Color Deep we considered both the average and the best results.

From the analysis of Table 3, we can observe that, as far as the dataset $\mathcal{D}$ is concerned, the performance of DLE4FC is better than that of the other deep learning models for all metrics considered. Color Deep obtains very good performance values in its second configuration, although they are slightly lower than those of DLE4FC. Table 3 also reveals that the ensemble

Deep neural network	*Precision*	*Recall*	*F1-score*	*Accuracy*
DLE4FC	**0.80**	**0.80**	**0.80**	**0.80**
CNN$^{\Delta}$	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 3
Results of DLE4FC 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*
DLE4FC	**0.79**	**0.78**	**0.78**	**0.80**
CNN$^{\Delta}$	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 4
Results of DLE4FC and the other models chosen for comparison when the images of the test set of $\mathcal{D}^{-}$ are provided in input

strategy achieves better results than using a single CNN$^{\Delta}$. In fact, the value of Precision (resp., Recall, F1-score, Accuracy) obtained by DLE4FC is 2.56% (resp., 5.26%, 3.89%, 5.26%) higher than the corresponding value obtained by a single CNN$^{\Delta}$.

Analyzing Table 4 and comparing it with Table 3, we can see that providing input images with only the dominant color to the models under consideration does not lead to better results. On the contrary, all models except DLE4FC show a marked decrease in performance values. The best results after those of DLE4FC are obtained by using only one CNN$^{\Delta}$. In this case, the Precision (resp., Recall, F1-score, Accuracy) is 11.39% (resp., 15.38%, 12.82%, 16.25%) worse than that obtained by DLE4FC. The other models show even lower performance results. Among them, again, the best is Color Deep in its second version. Even with the dataset $\mathcal{D}^{-}$ we can observe that the ensemble strategy gets better results than using only one CNN$^{\Delta}$. As previously pointed out, contrary to what might have been expected, using the images with only the dominant color returns worse results than using the original fabric images. This may be due to the flattening of some major color nuances, which increases the difficulty of recognizing the true color by the deep learning models under consideration.

4. Conclusion

In this paper, we have proposed DLE4FC, a new approach to identify colors in fabrics. This is a challenging issue since fabric images contain textural patterns that can make color classification difficult. DLE4FC is based on a new color space, called color difference space, which allows providing deep learning classifiers with much more information than the one that could be provided by the RGB color space. Because of this idea and the use of the ensemble strategy, DLE4FC manages to obtain equal or better performance than the main related approaches proposed in the past literature. DLE4FC is based on a particular CNN model, called CNN^{Δ}; specifically, it is an ensemble of three CNN^{Δ} models. In this paper, we also illustrated some experiments performed on the Fabric dataset through which we compared DLE4FC with other existing approaches.

In the future, we plan to enhance DLE4FC along several directions. First, we plan to increase the number of reference colors used during classification. Also, we plan to allow DLE4FC to assign multiple colors to a single fabric image. In addition, if we had access to a dataset much larger than Fabric, we could consider enriching DLE4FC with a vision transformer-based methodology. Indeed, the latter architecture is very promising but it needs the application of thousands of training images. Finally, the color difference space used by DLE4FC is currently built on the RGB color space. We could think of building this space on other color spaces, such as HSL (Hue, Saturation, Lightness), HSV (Hue, Saturation, Value) or CMYK (Cyan, Magenta, Yellow and Key Black). Indeed, we believe that they could provide additional information for DLE4FC to use in making the classification of fabric images.

References

[1] B. Liu, J. Gan, B. Wen, Y. LiuFu, W. Gao, An automatic coloring method for ethnic costume sketches based on generative adversarial networks, Applied Soft Computing 98 (2021) 106786. Elsevier.

[2] W. Sun, J. Hao, Y. Zang, X. Lai, A solution-driven multilevel approach for graph coloring, Applied Soft Computing 104 (2021) 107174. Elsevier.

[3] A. Shams-Nateri, E. Hasanlou, Computer vision techniques for measuring and demonstrating color of textile, in: W. Wong (Ed.), Applications of Computer Vision in Fashion and Textiles, The Textile Institute Book Series, Woodhead Publishing, 2018, pp. 189–220. doi:https://doi.org/10.1016/B978-0-08-101217-8.00008-7.

[4] C. Kou, C. Shih, C. Kao, J. Lee, Color and Pattern Analysis of Printed Fabric by an Unsupervised Clustering Method, Textile Research Journal 75 (2005) 9–12. doi:10.1177/004051750507500103.

[5] Z. Di, Z. Luoqing, 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 (2013) 1673 – 1689.

[6] B. Anami, M. 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 (2015) 389–400.

[7] M. Zhang, 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, Washington DC, USA, 2012, pp. 218–221.

[8] D. Borza, T. Ileni, 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), Springer, Auckland, New Zealand, 2018, pp. 438–449.

[9] Z. Zhuang, Y. Liu, F. Ding, Z. Wang, Online Color Classification System of Solid Wood Flooring Based on Characteristic Features, Sensors 21 (2021) 336.

[10] T. Vafeiadis, N. Kolokas, N. Dimitriou, A. Zacharaki, M. Yildirim, H. G. Selvi, D. Ioannidis, D. Tzovaras, A

comparison of 2DCNN network architectures and boosting techniques for regression-based textile whiteness estimation, Simulation Modelling Practice and Theory 114 (2022) 102400.
[11] L. Xie, Y. Shen, 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 (2016) 29–37.
[12] Z. Zehra, M. 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), Yokohama, Japan, 2019, pp. 1–5. doi:10.1109/ICBDSC.2019.8645600.
[13] Q. Zhang, L. Zhuo, J. Li, J. Zhang, H. Zhang, X. Li, Vehicle color recognition using Multiple-Layer Feature Representations of lightweight convolutional neural network, Signal Processing 147 (2018) 146–153. doi:https://doi.org/10.1016/j.sigpro.2018.01.021.
[14] H. Kim, J. Park, H. Jung, An Efficient Color Space for Deep-Learning Based Traffic Light Recognition, Journal of Advanced Transportation (2018) 2365414.
[15] D. Apriyanti, L. Spreeuwers, P. Lucas, R. Veldhuis, Automated color detection in orchids using color labels and deep learning, PLoS ONE 16 (2021) e0259036. doi:10.1371/journal.pone.0259036.
[16] Z. Yang, L. Leng, B. Kim, StoolNet for Color Classification of Stool Medical Images, Electronics 8 (2019) 1464. doi:10.3390/electronics8121464.
[17] A. Rame, A. Douillard, C. Ollion, CORE: Color Regression for Multiple Colors Fashion Garments, ArXiv preprint abs/2010.02849 (2020).
[18] J. Zhang, R. Pan, E. Gao, D. Zhu, Automatic detection of layout of color yarns of yarn-dyed fabric. Part 1: Single-system-mélange color fabrics, Color Research & Application 40 (2015) 626–636. doi:10.1002/col.21927.
[19] P. Chen, X. Bai, W. Liu, Vehicle color recognition on urban road by feature context, IEEE Transactions on Intelligent Transportation Systems 15 (2014) 2340–2346. IEEE.
[20] O. Sagi, L. Rokach, Ensemble learning: A survey, WIREs Data Mining and Knowledge Discovery 8 (2018) e1249.
[21] A. Amelio, G. Bonifazi, E. Corradini, S. Di Saverio, M. Marchetti, D. Ursino, L. Virgili, Defining a deep neural network ensemble for identifying fabric colors, Applied Soft Computing 130 (2022) 109687. Elsevier.
[22] C. Kampouris, S. Zafeiriou, A. Ghosh, S. Malassiotis, Fine-Grained Material Classification Using Micro-geometry and Reflectance, in: Proc. of the European Conference on Computer Vision (ECCV'16), Springer, Amsterdam, The Netherlands, 2016, pp. 778–792.
[23] K. He, X. Zhang, S. Ren, J. Sun, Deep Residual Learning for Image Recognition, in: Proc. of the International Conference on Computer Vision and Pattern Recognition (CVPR'16), Las Vegas, NV, USA, 2016, pp. 770–778. doi:10.1109/CVPR.2016.90.
[24] C. Szegedy, L. Wei, J. Yangqing, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, A. Rabinovich, Going deeper with convolutions, in: Proc. of the International Conference on Computer Vision and Pattern Recognition (CVPR'15), Boston, MA, USA, 2015, pp. 1–9. doi:10.1109/CVPR.2015.7298594.
[25] G. Huang, Z. Liu, L. V. D. Maaten, K. Weinberger, Densely Connected Convolutional Networks, in: Proc. of the International Conference on Computer Vision and Pattern Recognition (CVPR'17), IEEE Computer Society, Honolulu, HI, USA, 2017, pp. 2261–2269.
[26] M. Zhang, P. Wang, X. Zhang, Vehicle Color Recognition Using Deep Convolutional Neural Networks, in: Proc. of the International Conference on Artificial Intelligence and Computer Science (AICS'19), Wuhan, China, 2019, pp. 236––238. doi:10.1145/3349341.3349408, association for Computing Machinery.
[27] D. Bloomberg, Color quantization using modified median cut, in: Leptonica, 2008, pp. 1–10.