

Aesthetic Quality Classification of Photographs Based on Color Harmony

Masashi Nishiyama¹, Takahiro Okabe¹, Imari Sato², Yoichi Sato¹

¹Institute of Industrial Science, The University of Tokyo, Japan

²National Institute of Informatics, Japan

masashi1@iis.u-tokyo.ac.jp

Abstract

Aesthetic quality classification plays an important role in how people organize large photo collections. In particular, color harmony is a key factor in the various aspects that determine the perceived quality of a photo, and it should be taken into account to improve the performance of automatic aesthetic quality classification. However, the existing models of color harmony take only simple color patterns into consideration—e.g., patches consisting of a few colors—and thus cannot be used to assess photos with complicated color arrangements. In this work, we tackle the challenging problem of evaluating the color harmony of photos with a particular focus on aesthetic quality classification. A key point is that a photograph can be seen as a collection of local regions with color variations that are relatively simple. This led us to develop a method for assessing the aesthetic quality of a photo based on the photo’s color harmony. We term the method ‘bags-of-color-patterns.’ Results of experiments on a large photo collection with user-provided aesthetic quality scores show that our aesthetic quality classification method, which explicitly takes into account the color harmony of a photo, outperforms the existing methods. Results also show that the classification performance is improved by combining our color harmony feature with blur, edges, and saliency features that reflect the aesthetics of the photos.

1. Introduction

Many people these days have large photo collections, thanks to the widespread use of digital cameras and the Internet. To manipulate a large amount of photos, it would be useful to assess the aesthetic quality of a photo automatically, *i.e.*, whether the photo elicits a high or low level of affection in a majority of people ¹. The perceived quality of a photograph depends on various aspects, e.g., color, composition, lighting, and subjects. In this paper, we focus

¹In the same manner as existing methods of aesthetic quality classification [10, 5, 18, 14, 13], we do not consider artistic photos any differently than other photos in terms of quality.




Figure 1. Complexity of color harmony of photos. We wanted to bridge the gap between simple color patterns and actual colors appearing in photos.

on color harmony assessment because the colors of photos have a significant influence on their perceived quality.

The existing models [16, 8, 9, 15] proposed in the field of color science consider the color harmony of a simple color pattern like the ones shown in Fig. 1 (a). These models can be used to assess the colors of many man-made objects such as cars, clothes, and Web site content. Unfortunately, however, they perform poorly in assessing the color harmony of a photo like the one shown in Fig. 1 (b) because the color distribution of a photo is often significantly more complex in both color and spatial domains than those considered in the existing color harmony models. This is why these models cannot be used for assessing the color harmony of photos with complex color distributions.

The color harmony of photos has been largely ignored in the existing methods of aesthetic quality classification [10, 5, 18, 14, 13]. In those methods, a global color histogram computed from a whole image is used as one of the image features for evaluating the aesthetic quality of the image. Only recently has the research community started to look at the color harmony of photos, e.g., inferring affective words from photos of a limited category [22].

In this paper, we address the challenging problem of assessing the color harmony of photos. We propose a method for automatically evaluating it to enhance the performance of aesthetic quality classification. The key point we started from is that a photo can be seen as a collection of local regions with color variations that are relatively simple. Our preliminary experiments showed that the sum of color har-

mony scores computed from the local regions of a photograph is closely related to its aesthetic quality. On the basis of this observation, we developed a method for color harmony assessment of photos. In our method, a color harmony model is applied to each local region of an image to evaluate the distributions of relative values of hue, lightness, and chroma to the dominant color in the region, and then the result is integrated to describe the whole image in the *bag-of-features* framework. Experimental results demonstrate that our aesthetic quality classification method, which explicitly takes into account the color harmony of a photo outperforms the existing methods.

The rest of this paper is organized as follows. Section 2 briefly summarizes related work, and Section 3 describes the proposed method for assessing the color harmony of photos. Section 4 demonstrates its effectiveness through experiments and analyses, and Section 5 reports how the aesthetic quality classification can be improved by combining various features with our color feature. Our concluding remarks are given in Section 6.

2. Related work

Our present work is closely related to two topics: aesthetic quality classification and color harmony models. This section summarizes prior work in both of these areas.

Aesthetic quality classification:

Some researchers have proposed aesthetic quality classification methods [10, 5, 14, 13] and the use of the aesthetic quality for applications such as image cropping and recomposition [18, 12, 1]. In contrast to the conventional image quality measures [24, 25], aesthetic image quality measures are advantageous because they are more closely correlated with our impressions of images.

The existing aesthetic quality classification methods distinguish between high- and low-quality photos by using a classifier trained from a large collection of image samples, typically collected from the Internet, with manually provided aesthetic quality scores. For better classification accuracy, it is essential to choose good image features. In the previous methods, each photo is described with empirically chosen features such as the rule of thirds, color histogram, and the size and position of a salient region in the image. Here, we argue that simple color features like color histograms are insufficient, and it is essential to consider color harmony to evaluate aesthetic quality properly.

Color harmony models:

There have been two major models proposed in the field of color science for evaluating color harmony: the Moon-Spencer model [16] and the Matsuda model [15].

The Moon-Spencer model handles a simple color pattern consisting of two colors, as shown in Fig. 1 (a). This model, which is based on psychological experiments, evaluates the relationship between a color pattern and the affection it elic-




Figure 2. The Moon-Spencer model for color harmony assessment. The model assesses the relationship between two colors by comparing a chosen color value with a target color value. When the relative value between the two does not appear in the region of ‘Ambiguity,’ the model tells us that the two colors are in harmony.

its when the pattern is shown as a stimulus. The model computes the relative value between two colors in the Munsell color system. As Fig. 2 illustrates, when the relative values of hue, chroma, and lightness do not appear in ‘Ambiguity,’ the two colors are considered to be in harmony. This model presents three types of color harmony: ‘Contrast,’ which is a target color that is significantly different from a chosen color, ‘Similarity,’ which is a resembling color, and ‘Identity,’ which is the same color.

The Matsuda model was developed for designing clothes based on simple patterns with a few colors. Matsuda presented nine harmonic templates that define ranges where colors are in harmony on the hue circle. This model also uses relative hue values such as the Moon-Spencer model, and it has been used in certain computer graphics and pattern recognition applications. For example, Cohen-Or *et al.* [2] proposed a method for color transfer while maintaining color harmony by fitting the model to a hue histogram counted in an object region. Moorthy *et al.* [17] exploited the model to assess the color harmony of videos by comparing the model with hue histograms counted from the whole image.

Unfortunately, we cannot apply these methods of color harmony to the problem of aesthetic quality classification of images because the color distribution of images is significantly more complex than a simple combination of a few color patches. As described later and shown in Fig. 9, using the Moon-Spencer or the Matsuda models for a whole photo is not suitable for color harmony assessment.

3. Assessing color harmony of photos

3.1. Overview

As shown in Fig. 1, the distribution of colors in a photo is more complicated than simple color patterns in both the color and spatial domains. Therefore, there is a large gap between the previously proposed color harmony models for simple color patterns and that would be effective for a photo.

To overcome this difficulty, we exploit the fact that, when




Figure 3. Assumptions on assessing the color harmony of photos. We observed that (a) a local region is regarded as a simple color pattern and assume that (b) a photo consists of a collection of simple color patterns. We developed a method for assessing the color harmony of a photo by using a collection of local regions.




Figure 4. Histogram features representing color harmony. Our method prepares representative local regions as prior knowledge for assessing the color harmony of photos and counts the frequency of appearances of local regions that are similar to the representative regions in a photo. The histogram features tell us the difference in the color harmony between photos.

we observe a local region of a photo, the distribution of colors within the region is relatively simple, as illustrated in Fig. 3 (a). We can consider each region to have a simple color pattern and evaluate the color harmony of the region by using the color harmony models for simple color patterns. Because a photo is a collection of local regions, the color harmony of the whole photo can be computed from the color harmony scores of all the local regions.

Our preliminary experiments demonstrate that we can represent a photo as a collection of local regions with color patterns and use the collection to assess the whole photo. As we describe in greater detail in Section 4, the sum of the color harmony scores computed from local regions of a photo by using the Moon-Spencer model is positively correlated with the aesthetic quality of the photo. In addition, a photo with high (low) aesthetic quality often contains a large number of local color patterns with high (low) color harmony scores.

On the basis of these observations, we developed a method for assessing the color harmony of photos. Our method assumes that a photo is described by a collection of simple color patterns (Fig. 3 (b)) and classifies the aesthetic quality of the photo on the basis of the frequency of appearance, *i.e.*, the histogram of those color patterns (Fig. 4). Specifically, local regions are sampled from a photo, and

each region is described by a feature based on the color harmony models for simple color patterns. These features are then quantized and the photo is represented by the histogram of the quantized features. Finally, the aesthetic quality of the photo is determined by a classifier trained with labeled photos.

Our method has similarities with a technique called *bags-of-features* [4, 20] for generic object recognition and image retrieval. One of the difficulties in these research areas is that the appearance of an object drastically varies depending on imaging conditions such as camera viewpoint, object pose, and illumination. To overcome this difficulty, the *bags-of-features* technique represents an image as a set of local features insensitive to imaging conditions. Thus, in this sense, our method for assessing the color harmony of photos can be termed *bags-of-color-patterns*. The descriptor of a local region computed from the color harmony models is a counterpart of the descriptor of local features such as SIFT for generic object recognition.

3.2. Bags-of-color-patterns

The proposed method (1) samples local regions of a photo, (2) describes each local region by features based on color harmony models for simple color patterns, (3) quantizes these features, and (4) represents the photo as a histogram of quantized features. Finally, it uses a support vector machine (SVM) classifier trained by using sample photos with aesthetic qualities labeled by various individuals. The steps are described below.

3.2.1 Sampling local regions

Our method uses a grid-sampling technique to extract a set of local regions. The operators used in existing methods for object recognition [4, 20], *e.g.*, the Difference of Gaussians algorithm and the Harris-Laplace detector, are not suitable for color harmony assessment because not only the colors around the edges and corners but also those in uniform regions affect the perceived aesthetic quality of a photograph.

The grid-sampling technique extracts local regions of a fixed size from equally spaced positions. We empirically determine the size of the regions and the sampling density to determine the number of regions cropped from a single photo. We experimentally confirmed that aesthetic quality classification is not necessarily sensitive to the parameters for grid-sampling (see Section 4.4).

The Moon-Spencer model (Fig. 2) tells us that a simple color pattern is in harmony when colors in the pattern are ‘Identity,’ ‘Similarity,’ or ‘Contrast.’ This means that uniform regions are usually in harmony but regions around edges and corners are in harmony only when the colors within a region are similar or contrasting. In order to incorporate the reason why the color of a local region is in




Figure 5. Sampling local regions. We extract local regions from a whole photo using a grid-sampling technique, and distinguish the regions that contain color boundaries. The pixel value in (a) represents a color boundary computed from the photo shown in Fig. 1 (b). We make two sets of local regions with/without color boundaries in (b).

harmony into our method, we distinguish uniform regions from regions around edges and corners and treat them separately. We divide a whole photo into segments by using mean shift segmentation [3], and detect color boundaries by using discriminant analysis [21] (Fig. 5 (a)). We then split a set of local regions into those with/without color boundaries (Fig. 5 (b)).

3.2.2 Describing local regions

The Moon-Spencer model shows that the color harmony of simple color patterns such as one with two colors can be described by the difference between two colors: the hue, chroma, and lightness values relative to the chosen color. Inspired by the Moon-Spencer model, we find the dominant color of a local region and describe the local region by using the hue, chroma, and lightness values relative to the dominant color. It is worth noting that the existing methods for aesthetic quality classification use absolute values, *i.e.*, the standard pixel values in RGB channels, and not the relative values.

Specifically, we describe local regions based on the Moon-Spencer model as follows. First, we convert pixel values within a local region from RGB color space to the Munsell color system². Then, we find the dominant color of the region based on the hue values. We compute the hue values relative to the dominant color by subtracting the hue value of the dominant color from that of each pixel and obtain the histogram of the relative hue values as illustrated on the top side of Fig. 6 (b). Motivated by the analogy with the difference of hue values in the Moon-Spencer model (shown in Fig. 2 (a)), we use this histogram to describe a local region. In addition, we compare the average values of chroma and lightness on the pixels with the dominant color and subtract these average values from the chroma and

²We used a table published by Berns *et al.* of the Rochester Institute of Technology (<http://www.cis.rit.edu/mcsl/online/munsell.php>). The table provides values moving from the Munsell color system to CIE xyY color space. We calculated values not in the table by using a linear interpolation technique.




Figure 6. Local descriptor for assessing color harmony in a local region. To compute relative values, we determine the dominant color in a local region, such as red in (a). Next, we plot the two histograms of (b) by using relative colors with respect to the dominant color.




Figure 7. Quantizing local descriptors. We compute local descriptors of (a) from local regions with/without color boundaries to generate codebooks. We plot the two histograms of (b) by using visual words in the codebooks. We concatenate the histograms to make a feature representing the whole photo.

lightness of each pixel. Thus, we obtain a 2D histogram of the relative chroma and lightness values (illustrated on the bottom side of Fig. 6 (b)). This histogram is analogous to the differences of chroma and lightness in Fig. 2 (b). Finally, we concatenate the histogram of relative hue values and the 2D histogram of relative chroma and lightness into a single vector and use it as the descriptor of a local region.

3.2.3 Quantizing local descriptors

As is often the case with *bags-of-features* for generic object recognition and image retrieval, we also quantize local descriptors by using visual words in codebooks. First, we obtain a large number of local descriptors from training samples, as shown in Fig. 7 (a). Then, we generate two codebooks for local descriptors with/without color boundaries by using the k-means clustering method.

3.2.4 Representing a whole photo

We use these codebooks to compute the frequency of appearance, *i.e.*, the histograms of quantized local descriptors with/without color boundaries, as illustrated in Fig. 7 (b). We concatenate the histograms for local regions with/without color boundaries into a single vector. Unfortunately, however, the vector cannot represent the distribution of colors in spatial domain because the descriptors of local regions have no information about their spatial locations.

Accordingly, in order to incorporate spatial information into the descriptor for a whole photo, we divide a photo into rectangular segments and obtain a vector (histogram) from each segment. Finally, we concatenate those vectors into a single vector and use it for to represent the whole photo.




Figure 8. The histogram in (a) shows the aesthetic score given to each photo. The one in (b) shows the number of photos evaluated by each individual. The horizontal axis represents the number of individuals who gave aesthetic scores to the photos. Both histograms are sampled from the category ‘Landscape.’

4. Experimental analysis on assessing color harmony of photos

We evaluated the effectiveness of our method on the datasets described below in Section 4.1. Here, we report the accuracy of color harmony assessment (Section 4.2) and show examples of visual words (Section 4.3). We also evaluate the sensitivity of local region parameters in Section 4.4.

4.1. Datasets

We tested our method using a photo collection available on the Internet (DPChallenge [6]). On the Web site, various individuals have given aesthetic scores to various photos. We collected 124,664 photos in 14 categories (Abstract, Animals, Cityscape, Floral, Landscape, Macro, Nature, Portraiture, Rural, Seascapes, Sky, Still-life, Urban, and Water) and removed all the sepia-tone and black-and-white photos. Let us first provide the details of this photo collection. Figure 8 (a) shows how many photos received aesthetic scores, and Fig. 8 (b) shows the number of individuals who provided scores for ‘Landscape.’ About 200 people gave a score to more than 4,000 photos. In our experiments, the top and bottom 10% of scores were deemed the high- and low-quality photos. The other categories have the same tendency as the one shown in Fig. 8. We trained our system using half the photo collection and tested our system using the other half.

It should be noted that the aesthetic scores on DPChallenge are given to photo based on various aspects including, among others, color harmony, composition, subject, blur, and contrast. However, most of the top 10% (bottom 10%) photos show high (low) color harmony, and it seems reasonable to assume that the top 10% (bottom 10%) of the distribution of the scores do not contains examples of low (high) color harmony.

4.2. Color harmony assessment for photos

We tested the performance of our method in two aspects: (1) considering color harmony based on a ‘Whole’ photo or




Figure 9. Performance of assessing the color harmony of photos in the DPChallenge dataset.

dividing it into a set of ‘Local’ regions, and (2) evaluating harmony in color using its ‘Absolute’ or ‘Relative’ color values. A feature representing the color harmony of a photo is computed using existing models of color harmony [15, 16], absolute values in the colors of the photo, or its relative values. All of these methods are tested for ‘Whole’ and ‘Local’ cases. An SVM with a linear kernel was used for all methods except ‘Matsuda’ and ‘Moon-Spencer’, 2,304 local regions sized 32×32 were extracted from each photo. We compared the following methods.

Whole, Matsuda: A feature is designed from the harmonic templates of the Matsuda model [15]. We computed a hue histogram in CIE LCH color space from a whole photo and evaluated the similarity between the templates and the histogram using the technique described in [2, 17].

Whole, Moon-Spencer: A feature is extracted from a whole photo by using the Moon-Spencer model [16]. We computed the histograms of relative values of hue, chroma, and lightness and summed bins corresponding to the ‘Contrast,’ ‘Similarity,’ and ‘Identity’ in Fig. 2.

Local, Moon-Spencer: A feature is the sum of color harmony scores using the Moon-Spencer model from local regions of a photo.

Whole, Absolute (Chroma): A histogram of chroma values extracted from a whole photo in CIE LCH color space.

Whole, Absolute (Hue): A histogram of hue values extracted from a whole photo in CIE LCH color space.

Whole, Absolute (RGB): A histogram of pixel values extracted from a whole photo in RGB color space.




Figure 10. Examples of photos classified from their color as being of high quality (a) and low quality (b).

Local, Absolute (Chroma): A feature extracted from a collection of local regions. Local regions were described by using a chroma histogram. We used the proposed method (described in Section 3.2.3) to combine the descriptors.

Local, Absolute (Hue): A feature computed from a collection of local regions. Local region were described by using a hue histogram.

Local, Absolute (RGB): A feature computed from a collection of local regions. Local regions were described by using an RGB histogram.

Local, Relative: Our color harmony feature (described in Section 3.2). The feature was extracted from a collection of local regions. Local regions were described by using relative color values. Our local descriptor had 200 dimensions.

Figure 9 shows the classification performance as a recognition rate: the probability that the quality inferred using each feature matched the correct quality. The plot shows the average rate among the 14 categories. We can clearly see that the ‘Local’ feature using a set of local regions is superior to the ‘Whole’ feature extracted from the whole photo, as well as that it is effective for existing models of color harmony. Also, we see that the features of ‘Relative’ values are superior to the one of ‘Absolute’ value. Overall, our ‘Local, Relative’ method outperforms all other features and achieved about 66% accuracy in this difficult task.

Figure 10 shows the classification results by using the proposed method: (a) photos with high color harmony and (b) photos with low color harmony. A comparison showed that the photos in (a) have more pleasant color distributions and spatially better balanced.

4.3. Evaluation of visual words

To evaluate visual words used for computing the features representing a whole photo (described in Section 3.2.3), we compared the visual words for high harmony with ones for low harmony. We generated visual words using the proposed method, and assessed the color harmony score of each visual word by using the Moon-Spencer model. The




Figure 11. Examples of visual words for high harmony (a) and low harmony (b). The local regions correspond to individual visual words.

top 100 and bottom 100 visual words in the ‘Landscape’ category were used for this experiment.

We counted the frequency of appearance of local regions corresponding to the top 100 visual words of high harmony using the test samples with high- and low-quality labels of the DPChallenge dataset. 6,055 regions extracted as high-quality regions and 4,848 low-quality regions appeared. We also counted those corresponding to the bottom 100 visual words of low harmony. Results showed that 19,690 high-quality regions and 24,448 low-quality ones appeared. These results demonstrate that the visual words using our local descriptors have the capability to differentiate between high- and low-quality photos.

Figure 11 shows examples of visual words of high harmony (a) and low harmony (b) and examples of local regions corresponding to the visual words. Comparison of the




Figure 12. Performance of assessing the color harmony of larger local regions.




Figure 13. Performance of assessing the color harmony when the density of sampling is increased.

local regions in (a) with the ones in (b) showed that those in (b) elicited more affection. We believe that our visual words can effectively assess the color harmony of local regions.

4.4. Evaluation with the parameters of local regions

We evaluated the sensitivity of the parameters for the grid-sampling technique described in Section 3.2.1. One parameter is the size of one local region and the other is sampling density to determine the number of local regions extracted from a photo. Figure 12 shows the recognition performance while varying the size on a fixed density $48 \times 48 = 2,304$. The performance deteriorated when the size became too small and large or too large. Figure 13 shows the recognition performance while changing the density on a fixed size of 32×32 . The performance saturated as the sampling density increased.

5. Combining features for aesthetic quality classification

5.1. Algorithm for combining features

The aesthetic quality of photos in the DPChallenge should be evaluated considering the several aspects described in Section 4.1. In this section, we report how the performance of assessing the various aspects is improved as we take different features into consideration besides our color harmony feature.

Specifically, we consider edge, blur, and saliency as additional local features. We computed the feature of edges using DAISY [23], which is a histogram of gradients, the feature for blur using the frequency magnitude-based feature [19], which is a distribution of magnitude values in




Figure 14. Comparison of existing methods and our method of combining features for classifying aesthetic quality.

the frequency domain, and the feature for saliency, using a graph-based visual saliency [7] that gives an attention strength to each pixel. We generated a feature representing the whole photo by using the grid-sampling and the *bags-of-features* technique for each set of local features.

To classify the quality, we estimated the posterior probability in which each feature is matched with high quality labels and then computed the product of the probabilities of color harmony, edges, blur, and saliency by assuming that the features are independent. This assumption does not always hold because a saliency feature depends on color distribution. In our feature work, we plan to remove the redundant features. Each posterior probability was calculated by fitting the output of the SVM to a sigmoid function using a previously reported technique [11].

5.2. Evaluation of combining features

We compared the classification performance of our combination method described in Section 5.1, existing methods for classifying aesthetic quality [14, 18], and conventional methods for evaluating image quality [24, 25]. We briefly explain each method.

SSIM: The quality was determined by using the nearest neighbor technique to compare a given photo with training samples. This technique looks at the quality label of the training sample of the highest similarity. We computed a similarity using SSIM [24], which evaluates the variations of pixel values in equally divided blocks.

GSSIM & CS: The quality was determined by using the nearest neighbor technique and the similarity was computed using GSSIM and colorfulness [25], which evaluates the variations of edge strengths and RGB values.

Single subject: The quality was determined by using SVM and features extracted from a single subject and a background detected using Luo et al.'s method [14]. We used edge, color, and blur features.

Multiple subjects: The quality was determined by using SVM and features extracted from multiple subjects and a background detected using Nishiyama *et al.*'s method [18]. We used edge, color, and blur features.

Local regions: The quality was determined by using the combination method described in Section 5.1.

Figure 14 shows the classification performance using the average rate among the 14 categories in the DPChallenge. Our method (77.6%) outperformed the conventional methods for evaluating image quality (55.1%, 54.8%) and the methods for classifying aesthetic quality (67.4%, 71.7%). This significant gain in recognition performance demonstrates the effectiveness of considering a photo as a collection of its local regions and computing various descriptors from those regions.

6. Conclusion

We proposed a method for aesthetic quality classification based on the color harmony of photos. To tackle the difficulty of dealing with complex color distribution, we introduced the *bags-of-color-patterns* representation. The salient contributions of this work are the following.

- We observed that a photograph can be seen as a collection of local regions with color variations that are relatively simple. On the basis of this observation, the aesthetic quality of a whole image is then evaluated in the *bags-of-features* framework.
- We took blur, edges, and saliency features into consideration besides our color harmony feature and introduced a scheme that combine the features for aesthetic quality classification.

Our experimental results using a large dataset from the DPChallenge demonstrate that our method provides substantially more accurate aesthetic quality classifications than the existing methods. For future work, we plan to extend our method by taking into consideration scale variation when sampling local regions. We also intend to explore the real-world application of our method for various tasks, e.g., color transfer and automatic white balance.

References

- [1] S. Bhattacharya, R. Sukthankar, and M. Shah. A framework for photo-quality assessment and enhancement based on visual aesthetics. *Proc. ACM Int'l Conf. Multimedia*, pages 271–280, 2010.
- [2] D. Cohen-Or, O. Sorkine, R. Gal, T. Leyvand, and Y. Q. Xu. Color harmonization. *ACM Trans. Graphics*, 25(3):624–630, 2006.
- [3] D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 24(5):603–619, 2002.
- [4] C. Dance, J. Willamowski, L. Fan, C. Bray, and G. Csurka. Visual categorization with bags of keypoints. *Proc. ECCV Int'l Workshop Statistical Learning in Computer Vision*, pages 1–22, 2004.
- [5] R. Datta, D. Joshi, J. Li, and J. Z. Wang. Studying aesthetics in photographic images using a computational approach. *Proc. European Conf. Computer Vision*, 3953:288–301, 2006.
- [6] DPChallenge. <http://www.dpchallenge.com>.
- [7] J. Harel, C. Koch, and P. Perona. Graph-based visual saliency. *Proc. Advances in Neural Information Processing Systems*, 19:545–552, 2007.
- [8] J. Itten. *The Art of Color*. Van Nostrand Reinhold Company, 1960.
- [9] D. B. Judd and G. Wyszecki. *Color in Business, Science, and Industry*. John Wiley & Sons, 1975.
- [10] Y. Ke, X. Tang, and F. Jing. The design of high-level features for photo quality assessment. *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 1:419–426, 2006.
- [11] H. T. Lin, C. J. Lin, and R. C. Weng. A note on platt's probabilistic outputs for support vector machines. *Technical report, National Taiwan University*, 2003.
- [12] L. Liu, R. Chen, L. Wolf, and D. Cohen-Or. Optimizing photo composition. *Computer Graphic Forum*, 29(2):469–478, 2010.
- [13] A. Loui, M. D. Wood, A. Scalise, and J. Birkelund. Multidimensional image value assessment and rating for automated albuming and retrieval. *Proc. IEEE Int'l Conf. Image Processing*, pages 97–100, 2008.
- [14] Y. Luo and X. Tang. Photo and video quality evaluation: Focusing on the subject. *Proc. European Conf. Computer Vision*, 5304:386–399, 2008.
- [15] Y. Matsuda. *Color design*. Asakura Shoten, 1995.
- [16] P. Moon and D. E. Spencer. Geometric formulation of classical color harmony. *J. Opt. Soc. Amer.*, 34:46–50, 1944.
- [17] A. K. Moorthy, P. Obrador, and N. Oliver. Towards computational models of the visual aesthetic appeal of consumer videos. *Proc. European Conf. Computer Vision*, 6315:1–14, 2010.
- [18] M. Nishiyama, T. Okabe, Y. Sato, and I. Sato. Sensation-based photo cropping. *Proc. ACM Int'l Conf. Multimedia*, pages 669–672, 2009.
- [19] M. Nishiyama, H. Takeshima, J. Shotton, T. Kozakaya, and O. Yamaguchi. Facial deblur inference to improve recognition of blurred faces. *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 1115–1122, 2009.
- [20] E. Nowak, F. Jurie, and B. Triggs. Sampling strategies for bag-of-features image classification. *Proc. European Conf. Computer Vision*, pages 490–503, 2006.
- [21] N. Otsu. A threshold selection method from gray-level histograms. *IEEE Trans. Systems, Man and Cybernetics*, 9(1):62–66, 1979.
- [22] Y. Shin and E. Y. Kim. Affective prediction in photographic images using probabilistic affective model. *Proc. ACM Int'l Conf. Image and Video Retrieval*, pages 390–397, 2010.
- [23] E. Tola, V. Lepetit, and P. Fua. Daisy: An efficient dense descriptor applied to wide-baseline. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 32(5):815–830, 2010.
- [24] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. Image Processing*, 13(4):600–612, 2004.
- [25] Y. Xiang, B. Zou, and H. Li. Selective color transfer with multi-source images. *Pattern Recognition Letters*, 30(7):682–689, 2009.