Mining Probabilistic Color Palettes for Summarizing Color Use in Artwork Collections

Ying Cao
City University of Hong Kong

Antoni B. Chan
City University of Hong Kong

Rynson W.H. Lau
City University of Hong Kong

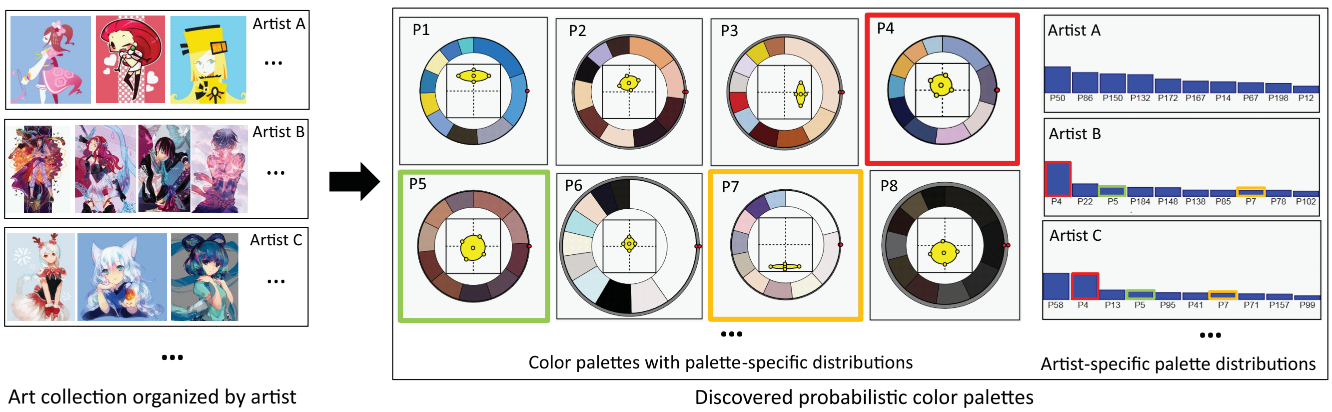

Figure 1: Given a collection of artworks organized by artist, we discover a set of color palettes and their associated probability distributions (collectively referred to as *probabilistic color palettes*), to summarize how colors are used within the collection. Each color palette is associated with three *palette-specific distributions*: color distribution, position distribution and size distribution. The color distribution is a categorical distribution, visualized as a pie chart where the arc length of each sector is proportional to its color frequency. The position distribution is a 2D Gaussian distribution, visualized as an oriented ellipse over a square area (normalized canvas space). It represents the spatial range that the palette is most likely applied to. The size distribution is a 1D Gaussian distribution, whose mean and standard deviation are reflected by the radius of the color pie chart and the size of the gray region outside the color pie chart. It represents the most likely sizes of regions using the palette. Each artist is associated with an artist-specific palette distribution, describing which color palettes that the artist prefers to use. The artist-specific palette distributions can capture both commonality and discrepancy among the artists. For example, both artists B and C prefer to use P4, P5 and P7 (outlined in colors), while artists A and B have no high-probability palettes in common. The artwork images for artist A, B and C are by Cindysuke, sakonma and Lyiet on DeviantArt, respectively.

ABSTRACT

Artists and designers often use examples to find inspirational ideas for using colors. While growing public art repositories provide more examples to choose from, understanding the color use in such large artwork collections can be challenging. In this paper, we present a novel technique for summarizing the color use in large artwork collections. Our technique is based on a novel representation, *probabilistic color palettes*, which can intuitively summarize the contextual and stylistic use of colors in a collection of artworks. Unlike traditional color palettes that only encapsulate *what* colors are used using a compact set of representative colors, probabilistic color palettes encode the knowledge of *how* the colors are used in terms of frequencies, positions, and sizes, using an intuitive set of probability distributions. Given a collection of artworks organized by artist, we learn the probabilistic color palettes using a probabilistic colorization model, which describes the colorization process in a probabilistic framework and considers the impact of both spatial and semantic factors upon the colorization process. The learned probabilistic color palettes allows users to quickly understand the color use within the collection. We present results on a large collection of artworks by different artists, and evaluate the effectiveness of our probabilistic color palettes in a user study.

SA '17 Symposium on Visualization, November 27-30, 2017, Bangkok, Thailand

ACM ISBN 978-1-4503-5411-0/17/11…$15.00
https://doi.org/10.1145/3139295.3139296

CCS CONCEPTS

• **Human-centered computing → Visualization techniques**;
• **Computing methodologies → Machine learning**;

KEYWORDS

Visual analytics, Color palettes, Probabilistic modeling

ACM Reference format:
Ying Cao, Antoni B. Chan, and Rynson W.H. Lau. 2017. Mining Probabilistic Color Palettes for Summarizing Color Use in Artwork Collections. In *Proceedings of SA '17 Symposium on Visualization, Bangkok, Thailand, November 27-30, 2017,* 8 pages.
https://doi.org/10.1145/3139295.3139296

1 INTRODUCTION

Color is one of the most fundamental ingredients in art and graphics. However, creating attractive colorings can be difficult for both professional artists and amateurs. Artists often use color palette books or online color palette galleries (e.g., Adobe Kuler [Kuler 2015]) to explore the large space of coloring before adding colors to an artwork. However, color palettes do not encode the information regarding how the colors are used *in situ* (e.g., how frequently and where the colors are used). Hence, artists and designers often work with existing reference examples that show how the colors and contents are combined, to find coloring inspiration [Busche 2015; Meier et al. 2004].

On-line communities, e.g., DeviantArt [DeviantART 2015] and Designflavr [Designflavr 2015], provide artists with a place to share their artworks, resulting in large-scale public art galleries comprising millions of artworks. While the art galleries serve as an excellent source for finding inspirational ideas of selecting and using colors [Chapman 2009], exploring the color use in such large galleries can be a non-trivial task. The art galleries are mostly organized by some semantic labels, e.g., category and artist. While this can make it easy to quickly find a corpus of artworks for a particular class, there exists no support for users to systematically explore the use of colors within the class. The users cannot easily understand what and how colors are used without looking at all the artworks in a class that may contain tens of thousands of examples. Thus, it is desirable to develop computational methods and systems to facilitate the exploration of color use in large artwork collections.

In this work, we take a step towards this goal by proposing a novel technique for summarizing the use of colors within a collection of artworks. Our key idea is to formulate the colorization of the artworks as a probabilistic process, where artists sample a set of *hidden* color palettes to add colors to artworks. Given a collection of artworks organized by artist, we propose to describe the probabilistic colorization process of the artworks using a latent topic model where the hidden color palettes are treated as latent structure. As shown in Figure 1, by learning the model from the artwork collection, we can discover the hidden color palettes along with a set of probability distributions (collectively termed as probabilistic color palettes in our context), which characterize how artists use colors within the collection in an intuitive and compact form. It is worth noting that our probabilistic color palettes are meant to summarize the color use within a *collection* of artworks, rather than explaining the colors of each *single* artwork. Thus, our end goal is different from those of the existing color theme extraction methods that aim to extract a color palette for a single image.

To evaluate the effectiveness of our technique, we show results on a large collection of illustrations from a public art repository. We conduct a perceptual study to demonstrate that our probabilistic color palettes are significantly preferred by users against other baseline representations in summarizing the color use of art galleries.

2 RELATED WORKS

2.1 Color harmonization

Many theories and models have already been proposed in the psychology and computer graphics communities to understand and quantify the phenomenon of color harmony and preference [Cohen-Or et al. 2006; O'Donovan et al. 2011; Palmer and Schloss 2010; Whelan and Sutton 2004]. Unlike these works, our focus is to capture the color use in artwork collections, rather than modeling color compatibility. Thus, our probabilistic color palettes do not necessarily encompass harmonious colors, but rather encode the stylistic color combinations that are commonly used by artists.

2.2 Color palettes

Prior research has explored how to extract a color palette from a single image [Delon et al. 2007; Lin and Hanrahan 2013; Weeks and Hague 1997]. In contrast to these works on creating a color palette from a single image, we aim to discover a compact set of color palettes that collectively represent a set of artworks. In addition, our probabilistic color palettes contain not only a group of representative colors, but also probability distributions that summarize the contextual use of the palettes. Several works make use of color palettes as intuitive interfaces for color editing. For example, given a color palette, Wang *et al.* [2010] recolored a natural image, by matching its colors to the given color palette while still maintaining the realism of the image. Chang *et al.* [2015] developed an interactive tool to allow users to recolor an image by directly editing the colors of a color palette. In contrast, we use color palettes as an interface for summarizing and exploring the coloring of artwork collections. The system by [Meier et al. 2004] displays an artwork together with its color palette to provide information on how colors are used in context. However, users still have to go through every example to understand coloring patterns, which makes it practically infeasible for large artwork collections. In contrast, we only require users to browse through a few compact and informative representations, which makes the exploration far more efficient.

2.3 Analyzing big art data

The growing availability of large collections of digital artworks has embarked recent trends in analyzing large dataset of artworks for retrieval and classification tasks. Crowley *et al.* [2014] proposed to search objects from a large collection of paintings by learning an object classifier from natural images. Kim *et al.* [2014] performed statistical analysis on a large dataset of paintings based on low-level visual features to find artistic patterns in the paintings. Saleh *et al.* [2016] explored various visual features and metric learning methods for computing a similarity measure between paintings and used it for large-scale classification of painting's style and artist. Jahanian *et al.* [2017] jointly modeled colors and words from a corpus of designs using Latent Dirichlet Allocation (LDA) to discover the correlation between colors and semantics. We also analyze large collections of artworks but aim to learn the representation that well explains the color use in the collections, which has not been addressed by these prior works.

2.4 Exploration interfaces

As repositories of 2D images and 3D shapes continue to grow, there is an emerging interest in developing techniques for exploring such big visual data [Bao et al. 2013; Fish et al. 2014; Kim et al. 2013; Ovsjanikov et al. 2011; Ritchie et al. 2011; Talton et al. 2009; Zhu et al. 2014]. Our work is informed by a line of research on deriving intuitive representations for exploring interesting modes in the data. For example, Zhu *et al.* [2014] summarized big visual data using weighted averages of an image collection, and interactively edited the average images to explore visually important modes in the collection. Ovsjanikov *et al.* [2011] extracted a template deformable model from a collection of similar 3D shapes, which could be manipulated to explore the continuous variability of 3D shapes. We address the problem of exploring the color use in an artwork collection, which has not been studied by any prior research.

3 PROBABILISTIC MODELING OF COLORIZATION

Given an input collection of artworks organized by artist, our main goal is to derive a representation to summarize the stylistic patterns of using colors in the collection. We expect a desirable representation to satisfy two conditions. First, it should be *compact*, so that users can quickly understand the information presented in it. Second, it should be *informative* to encapsulate the information about how the colors are used by different artists. Note that each of the conditions alone is not sufficient for our purpose. For example, while traditional color palettes are in a simple and compact form that most artists and designers are familiar with, they do not encode the information about the contextual use of colors. On the other hand, artwork examples provide specific information about color use, but the overwhelming amount of information in a large number of examples makes it infeasible to explore them directly.

Our observation is that, to colorize an artwork, an artist often selects one or several color palettes, each containing a set of distinctive colors, and then assign the colors in the palettes to different regions of the artwork. In this process, the decisions that the artist makes, such as which color palettes should be selected, which colors from the selected palettes should be used more frequently, and where the colors should be used, essentially reflect how the artist uses colors. Unfortunately, those color palettes and their use are hidden in artists' expertise domain and cannot be directly observed from the completed artworks. Hence, our goal is to recover the hidden color palettes and their use from the input artwork collection. Our key idea is to assume artwork colorization as a probabilistic process, where we explicitly model the hidden color palettes and their use in a probabilistic form. Based on this assumption, we propose a probabilistic colorization model, which can be learned from the artwork collection, to produce a set of probabilistic color palettes that characterize the color use within the collection.

3.1 Colorization as a Probabilistic Process

We assume that each artwork is composed of disjoint color regions, each of which is assigned one of K colors. Therefore, we quantize each artwork in terms of both spatial and color space. For spatial quantization, we perform a color-based segmentation to segment an artwork into a set of regions based on color similarity. For color quantization, we represent each segmented region using its median color, and quantize the region colors of all the artworks into $K = 300$ clusters, forming a visual dictionary of K representative colors. Finally, each region is assigned its closest representative color. Refer to Section 1 of the supplemental for more details.

After segmenting an artwork into color regions, we aim to assign each region a semantic label to represent its semantic meaning (e.g., face and hair). Such semantic information will be used to guide our method to discover more meaningful color palettes as discussed later. However, semantic annotation of regions in artworks can be quite challenging, mainly because the regions of the same semantic class often undergo great variations in appearance, shape and perspective across different artworks. To workaround this problem, we make use of a deep convolution neural network (CNN) in [Saito and Matsui 2015] to extract a high-level semantic feature vector from each region. The regions are then clustered into L semantic classes, and the cluster labels are used as the semantic labels. Since the exact meanings of the clusters are unknown a priori, we refer to the labels as *latent semantic labels*, which provides information about which regions belong to the same (but unknown) semantic class. After the semantic labeling, each color region is labeled as one of L latent semantic classes. Refer to Section 2 of the supplemental for more details.

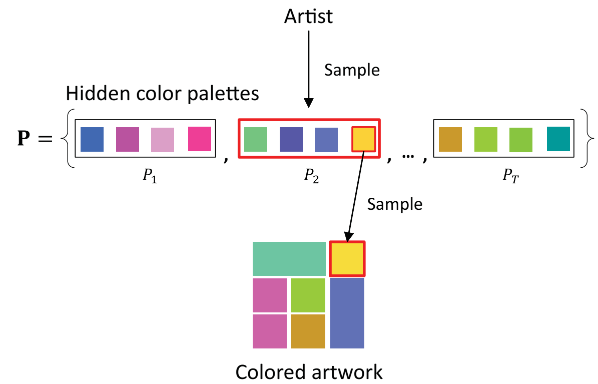

Figure 2: Illustrative example of probabilistic colorization.

After color quantization, each artwork can be described as a union of color regions, each taking one of K possible colors. In this way, colorizing an artwork can be cast as choosing a color among K possible colors for each region. Suppose that there exists a set of hidden color palettes $\mathbf{P} = \{P_1, \ldots, P_T\}$ that artists could select. We formulate the colorization of an artwork as a probabilistic process based on $\mathbf{P}$ as follows. For each region in an artwork, we first choose a color palette P_i from $\mathbf{P}$ based on a palette selection probability distribution. We then choose a color from P_i based on a color selection probability distribution, and assign the selected color to the region. This process is illustrated in Figure 2. Note that our formulation assume a single artwork can be associated with multiple color palettes. This differs from the traditional assumption that there is only one palette per artwork. The reason behind this design choice is that our goal is not to explain the coloring of a single artwork, but rather summarize the color use in an artwork collection. As compared with using one palette per artwork, this assumption enables us to capture more fine-grained, spatially-varying coloring patterns across the entire artwork collection. For example, for a collection of artworks composed of sky and ground, our representation can learn two separate palettes, one for sky and one for ground, while using a single palette would mix the sky and ground colors, resulting in a coarse and uninformative representation.

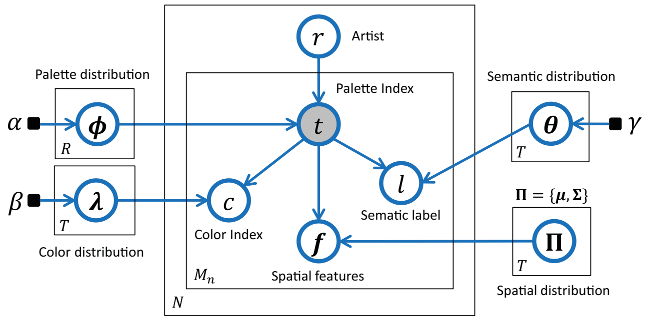

Figure 3: Graphical model representation of our colorization model.

3.2 Probabilistic Colorization Model

Based on the probabilistic colorization process in Section 3.1, we aim to develop a probabilistic colorization model that may characterize the generative process of coloring. The model will allow us to infer the hidden probabilistic color palettes (i.e., both the hidden color palettes and their associated probability distributions) from a collection of artworks, to capture the coloring patterns in the collection. Given the probabilistic colorization process, we notice that each artwork is, in essence, a mixture of hidden color palettes, and each of the palettes is a distribution of a fixed set of discrete colors. This motivates us to build our model upon the topic modeling framework. Here, we first briefly summarize the topical modeling algorithm, and then present our model in detail.

Topic modeling algorithms were originally developed to find patterns in large collections of documents [Blei 2012]. The topic models assume that a document is a random mixture over a set of hidden *topics* and that each topic is a distribution over a set of co-occurring *words*. The most basic topic model is Latent Dirichlet Allocation (LDA). LDA defines a generative process that a collection of documents arise from. More specifically, to generate a word in each document, we choose a topic according to a topic distribution, and then choose a word according to a word distribution for the chosen topic. In this way, LDA induces correlation between different words. For example, observing the word "rendering" implies a higher probability of the topic "computer graphics" and it is therefore more probable to see the word "geometry". Our modeling task can naturally fit into the topic modeling framework. In our context, each artwork (document) is a mixture of hidden color palettes (topics), each of which is a cluster of co-occurring colors (words) that are used together. In particular, our model is built upon an author-topic model [Steyvers et al. 2004], which assumes that each document contains one or more authors (i.e., artists in our context).

Figure 3 shows a graphical representation of our model with all the variables and their dependencies. Our model assumes that there are T color palettes $\{\boldsymbol{\lambda}_t\}_{t=1}^{T}$, where $\boldsymbol{\lambda}_t \in \mathbb{R}^K$ is a distribution over a fixed set of K possible colors. All the color palettes are shared by R different artists, to generate the coloring of N different artworks. The n-th artwork contains M_n color regions and is created by an artist $r_n \in \{1, \ldots, R\}$. To assign a color $c_{m,n} \in \{1, \ldots, K\}$ to region $m \in \{1, \ldots, M_n\}$ in artwork n, artist r_n first selects a color palette $t_{m,n}$ based on his/her own preference. This color palette preference of artist r_n is modeled as an artist-specific distribution $\boldsymbol{\Phi}_{r_n} \in \mathbb{R}^T$ over the T color palettes. Color $c_{m,n}$ is then drawn from color palette $t_{m,n}$ based on color distribution $\boldsymbol{\lambda}_{t_{m,n}}$.

Since the human perception of colors is contingent upon not only the color frequency, but also the spatial attributes (e.g., position and size) of color regions [Meier et al. 2004], we introduce a feature vector $\boldsymbol{f}_{m,n} = (x, y, s)^T$ to describe the spatial features of region m in artwork n, where x and y are the spatial coordinate of the region's centroid and s is the area of the region, both normalized w.r.t. the whole artwork. We model $\boldsymbol{f}_{m,n}$ as being drawn from a Gaussian distribution parameterized by $\boldsymbol{\Pi}_{t_{m,n}} = \{\boldsymbol{\mu}_{t_{m,n}}, \Sigma_{t_{m,n}}\}$, where $\boldsymbol{\mu}_{t_{m,n}}$ and $\Sigma_{t_{m,n}}$ are the mean and variance. It should be noted that such spatial feature vector would result in a distribution of spatial attributes of color regions for each color palette. Hence, for each color palette t, our model will capture not only the frequencies of colors being used via $\boldsymbol{\lambda}_t$, but also the spatial statistics of the color regions where the palette is applied via $\boldsymbol{\mu}_{t_{m,n}}$ and $\Sigma_{t_{m,n}}$.

What colors can appear in a region is also closely related to the semantic meaning of the region. For example, skin colors have a higher probability of occurrence on a face region than on a sky region. To account for the semantics in our model, we introduce a semantic label variable $l_{m,n} \in \{1, \ldots, L\}$ for region m in artwork n, which is assumed to be drawn from a semantic distribution $\boldsymbol{\theta}_{t_{m,n}}$. $l_{m,n}$ can take on L possible values corresponding to L latent semantic classes obtained in Section 3.1. Similar to $\boldsymbol{f}_{m,n}$, $l_{m,n}$ can induce a distribution of semantic labels of color regions for each color palette. $\boldsymbol{f}_{m,n}$ and $l_{m,n}$ would encourage our model to find the palettes where colors frequently co-occur in both *spatially* and *semantically* consistent regions. In our model, colors $\mathbf{c} = \{c_{m,n}\}$, spatial attributes $\mathbf{f} = \{\boldsymbol{f}_{m,n}\}$ and semantic labels $\mathbf{l} = \{l_{m,n}\}$ of all regions are observed random variables, while palette assignments $\mathbf{t} = \{t_{m,n}\}$ of all regions are latent random variables.

Given our model, our probabilistic color palettes can be parameterized as follows. A color palette t has three palette-specific distributions, including color, position and size distributions. The color distribution for palette t is a categorical distribution parameterized by $\boldsymbol{\lambda}_t$. It represents how frequently colors are used in the palette. The position distribution for palette t is a bivariate Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_t^p, \Sigma_t^p)$, describing the spatial range where the palette is likely used. $\boldsymbol{\mu}_t^p, \Sigma_t^p$ are taken from the entries for x and y of $\boldsymbol{\mu}_t, \Sigma_t$. The size distribution for palette t is a Gaussian distribution $\mathcal{N}(\mu_t^s, \sigma_t^s)$, which specifies the likelihood of using the palette for regions in certain sizes. μ_t^s, σ_t^s are taken from the entries for s of $\boldsymbol{\mu}_t, \Sigma_t$. In addition, the artist-specific palette distribution of an artist r is a categorical distribution parameterized by $\boldsymbol{\Phi}_r$, which describes the tendency of the artist using all the palettes. Note that, since the exact meanings of latent semantic classes are unkown, the semantic distribution for palette t, parameterized by $\boldsymbol{\theta}_t$, is not explicitly used in our representation. However, it is useful to guide the learning of our model towards discovering a semantics-aware representation.

3.3 Discovering the Probabilistic Color Palettes

The probabilistic color palettes are discovered by learning our model from the input artwork collection. In particular, we estimate the color probabilities $\{\boldsymbol{\lambda}_t\}_{t=1}^{T}$, the palette probabilities $\{\boldsymbol{\Phi}_r\}_{r=1}^{R}$, the parameters of spatial feature distributions $\{\boldsymbol{\mu}_t, \Sigma_t\}_{t=1}^{T}$, and the semantic probabilities $\{\boldsymbol{\theta}_t\}_{t=1}^{T}$, from the training data $D = \{r_n, \{c_{m,n}, f_{m,n}, l_{m,n}\}_{m=1}^{M_n}\}_{n=1}^{N}$, where N is the number of artworks and M_n is the number of color regions in the n-th artwork.

To learn our model, we adopt a sampling-based method as in [Griffiths and Steyvers 2004]. In particular, we adopt Gibbs sampling to draw samples from posterior distribution $p(\mathbf{t}|D)$. After sufficient number of iterations (over all the variables in $\mathbf{t}$), we take the most recent samples $\mathbf{t}^*$ and use them to estimate the parameters. Refer to Section 3 of the supplemental for details.

3.4 Visualizing the Probabilistic Color Palettes

To help users intuitively understand the color use within the input artwork collection, we visualize the learned probabilistic color palettes as follows. For each color palette, we visualize its probability distributions in an unified form, as shown in Figure 4(b). In particular, we visualize the color distribution as a pie chart, where each sector is a color and its arc length represents its relative frequency. For visual clarity, we only display the 10 most frequently used colors. The position distribution is visualized as a 2D oriented ellipse on a square canvas located at the center of the color pie chart. The center of the ellipse is the mean of the distribution, and the major and minor axes are the scaled eigenvectors of covariance matrix of the distribution. For the size distribution, we use the radius r of the color pie chart to represent the mean of the distribution, i.e., $r = r_{\min} + \rho * \mu^s$, where $r_{\min}$ is the minimum radius of the pie chart, μ^s is the normalized mean of the distribution and ρ is a constant coefficient. We use a gray margin region outside the color pie chart to reflect the variance of the distribution. The size of the margin region is proportional to the standard deviation of the distribution. Moreover, for each artist, we visualize its artist-specific palette distribution as a bar chart, where the height of a bar represents the probability of the corresponding color palette being used by the artist. Figure 4 shows a visualization example of the art gallery of an artist. From the two examples and the color palettes A and B, we can observe that this artist tends to use the colors in palette A for moderately sized regions around the center of the artworks (i.e., foreground), while the white and gray colors (most frequently used colors in palette B) are more likely used for the larger regions near the top part of the artworks (i.e., background). In addition, from the artist-specific palette distribution (Figure 4(c)), we can observe that this artist uses palette A more often than palette B.

4 RESULTS AND EVALUATION

We have tested our method on a PC with a 3GHz i7 CPU and 12 GB RAM. Our critical parameters are set as follows. We set hyperparameters $\alpha = 0.005$, $\beta = 0.005$ and $\gamma = 0.005$ to encourage the sparsity of color distributions and palette distributions, to avoid getting uninformative uniform distributions. To determine the number of colors in visual vocabulary K and the number of color palettes T, we define a range of possible values and perform cross-validation to choose the optimal values that maximize the marginal likelihood of the model, resulting in $K = 300$ and $T = 200$. When training our model, we run the Gibbs sampler for 500 iterations. We have found that more iterations do not give a significant difference in the estimated parameters.

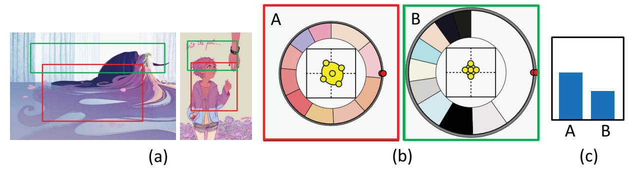

Figure 4: Visualization of probabilistic color palettes. (a) Artwork examples from the gallery of an artist (by celiere on DeviantArt). (b) High-probability color palettes with palette-specific distributions. (c) Artist-specific palette distribution. The rectangles on the examples in (a) indicate the regions where the color palettes in (b) are most likely used. The color of a rectangle indicates the correspondence between the region and its palette.

To evaluate our method, we use a test dataset containing 8,000 artworks from 30 different artists, which are crawled from an online public art gallery, DeviantArt. In particular, we randomly identify 30 artists, and for each artist, we download the artworks that are uploaded during the most recent 3 years from their gallery, in order to ensure that the coloring style of each artist is moderately consistent across all their artworks. The gallery of each artist in our dataset contains from 30 to 400 artworks.

4.1 Results

Figure 5 shows some examples of learned probabilistic color palettes organized by artist. By viewing the most frequently used color palettes and a few artwork examples, we can quickly understand how colors are used by an artist. For example, by observing all the color palettes in the first row of Figure 5, we find that this artist prefers to use gray, white and black colors in his artworks. In addition, by inspecting the color palettes along with just a few artwork examples, we find more interesting coloring patterns. In particular, this artist's artworks mostly depict one or several subjects captured in long or medium shots. Looking at the color palettes whose spatial distributions are around the centers of the square canvas, such as P67, P43, we find that gray and black colors are mainly used for body regions (almost around the center of an artwork). Meanwhile, this artist prefers to use some blue and pink colors to decorate some small regions, as reflected in P86. Furthermore, the size of the color pie chart in P44 is rather large, suggesting that the artist tends to paint large background regions using the colors in P44 (i.e., mainly black and white). Finally, by observing the color palettes whose spatial distributions are around the upper parts of the square canvas (i.e., P160), we know what colors are often used for face or hair regions.

Our probabilistic color palettes can also automatically capture composition patterns present in a gallery of artworks by an artist. In particular, if a spatial arrangement of several semantic regions with similar coloring frequently re-occurs across the artworks of an artist, we can find it easily, by observing the position distributions of the color palettes. For example, in the second row of Figure 5, we can find an apparent top-middle-bottom composition style, by looking at P54 (top), P186, P95, P100 (middle), and P146 (bottom). A few artwork examples by the artist show that this artist often paints

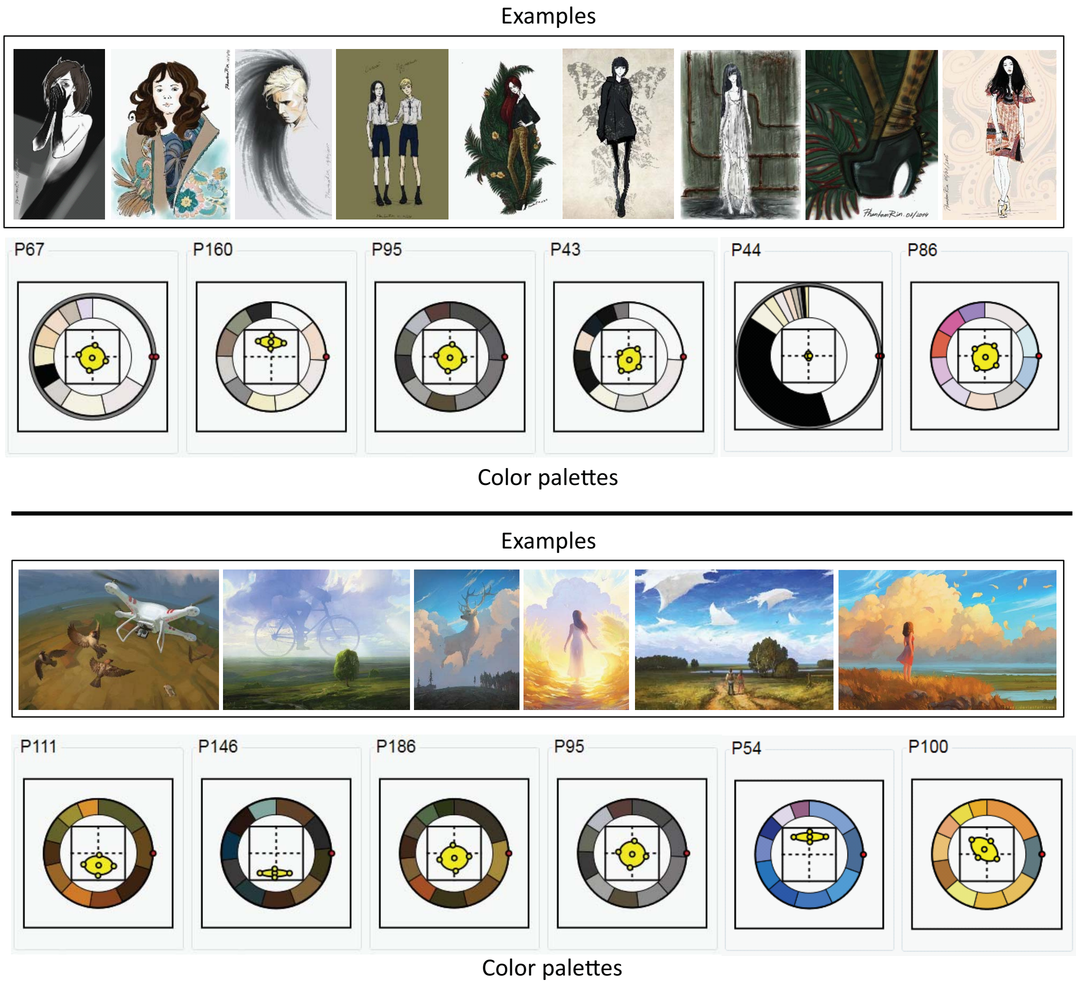

Figure 5: Learned probabilistic color palettes. Each row shows some examples from the gallery of an artist and the 6 most frequently used color palettes. The artwork images at the first and second rows are by PhantomRin and RHADS on DeviantArt, respectively.

outdoor scenes with three important semantic regions including sky (top), people (middle) and ground (bottom). Hence, we can easily figure out what colors are usually used for the three regions from the respective color palettes. Note that some color palettes may have similar position distributions around the center of an artwork, but have different color distributions (e.g., P186, P95, P100 in the second row of Figure 5). Since artists tend to place objects of interest around the center of the artwork, these color palettes encode rich coloring variations around the positions where objects are located. Refer to Section 4 of the supplemental for more examples.

The artist-specific palette distributions can be used to understand artists' coloring styles. In particular, we encode each artist-specific palette distribution as a T-dimensional feature vector, and use MDS to embed them into a 2D space, where the distance between two artists reflects their coloring style similarity. By visualizing the embedding space as shown in Figure 6, we can easily understand how artists are related to each other in coloring styles. From Figure 6, we can see that artists B and C have similar coloring styles, while artists E, F and G have similar coloring styles. Artists A and D are far away from other artists, implying that their coloring styles are more distinctive from the other artists.

4.2 Evaluation of the Probabilistic Color Palettes

We have conducted a user study to evaluate the effectiveness of our probabilistic color palettes in summarizing the color use within artwork collections, as compared to two baseline representations: 1) global color histogram and 2) spatial-binning color histogram. The study was performed on 5 test galleries from 5 different artists that were randomly chosen from our test dataset. For each gallery, we generated three types of representations and visualized them in the same way. All the representations were based on the color-based segmentation and quantization in Section 3.1, and the 300 representative colors computed from our entire test dataset. For our representation (Ours), we visualized the top 6 probabilistic color

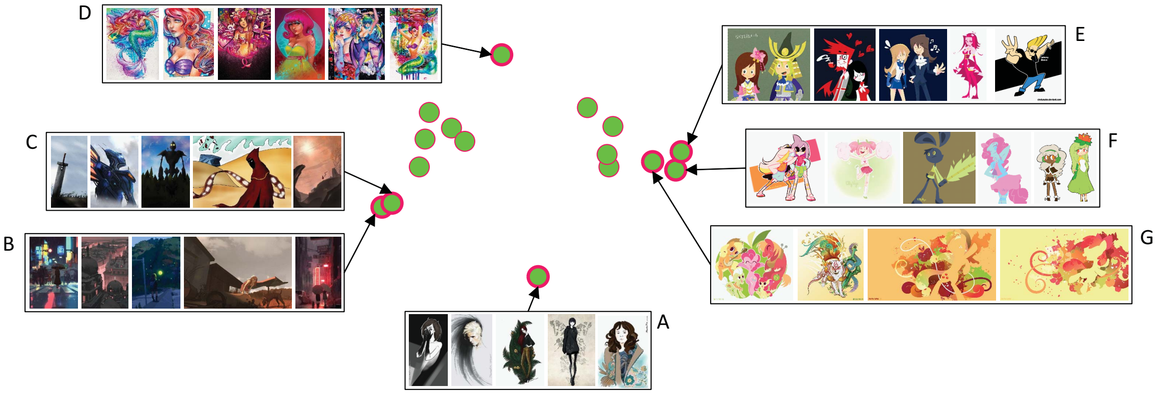

Figure 6: 2D projection of the artist-specific palette distributions. Each point represents an artist, and the distance between two points indicates the coloring style similarity between the corresponding artists. Some artworks from several artists are shown nearby the corresponding points highlighted by red circles. The artwork images for A, B, C, D, E, F, G are by PhantomRin, snatti89, TacoSauceNinja, rianbowart, Cindysuke, GrumpyBuneary, SambaNeko on DeviantArt, respectively.

palettes of each test gallery based on the palette distribution of the gallery. For the global histogram (Global), we constructed a single color histogram from all the artworks of each gallery. It was then visualized as a single probabilistic color palette, where the position and size distributions were obtained from the spatial properties of the regions of all the artworks in the gallery. For the spatial-binning histogram (Spatial), we partitioned each artwork space into $K \times K$ sub-regions of equal size, and built a single color histogram for each sub-region in the same way as the global histogram representation. The spatial histogram representation was finally visualized as K^2 probabilistic color palettes, each corresponding to a sub-region. We experimented with two different spatial-binnings, $K = 2$ (Spatial 2×2) and $K = 3$ (Spatial 3×3).

Our user study involves 30 participants, 15 of whom have more than 2 years of experience in drawing color illustrations or choosing colors for various graphic designs (expert user), while the others have less or no practice in drawing or coloring (average user). For each of our test galleries, the participants were asked to browse the entire gallery, as well as the 4 representations which were displayed in random order. They were then asked to rate the 4 representations, with respect to how well each representation could summarize the use of colors in the shown art gallery, on a scale of 1 to 5, with 1 being the worst and 5 being the best. The same rating could be given to different representations if they were regarded as similar. We ended up with 150 comparisons (30 participants × 5 test gallaries per participant). Figure 7 shows the average ratings of the four representations. Our representation is rated significantly higher than the others. In addition, the average rating from the expert users on our representation (3.93 ± 0.88) is slightly higher than that from the average users (3.71± 0.6). However, the difference between the groups of participants is not statistically significant (independent t-test, $p = 0.2$). This suggests that our representation is preferred by both professional artists and non-experts consistently. In Figure 8, we show the fraction that our representation is rated higher than each of the other representations, and the fraction that each representation is rated as the best or the worst in all the

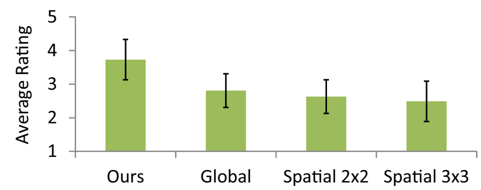

Figure 7: Average ratings of 4 different representations. These ratings have statistically significant differences (one-way repeated measures ANOVA test, $p < 0.001$). Ours is rated higher than others (paired t-test, $p < 0.001$).

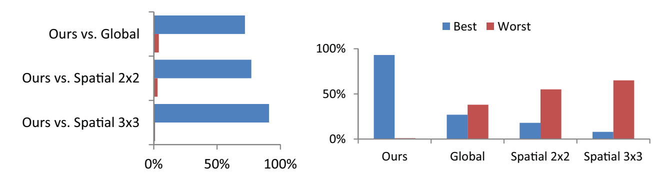

Figure 8: Left: fraction of our representation being rated higher (blue) or lower (red) than each of the other representations. Right: fraction of each representation being rated as the best (blue) and the worst (red).

comparisons. Both plots further confirm the significant advantage of our representation over other alternatives.

Figure 9 shows an example comparison used in our study. The global histogram can only summarize the color use coarsely, and is not able to capture the important coloring modes that vary spatially. For example, the artist uses a set of colors for the ground regions while another set of colors for the sky regions. The spatial-binning histogram can take spatial information into account. However, oblivious to composition patterns in the artworks, it may break some regions that have consistent color combination across the artworks due to image space partitioning. Hence, when building the histogram for a sub-region, it may group the color regions with

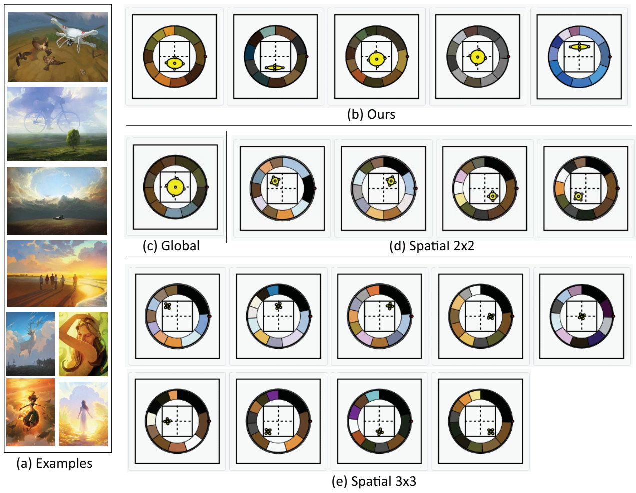

Figure 9: Comparison of 4 different representations for summarizing the color use in an artwork gallery. Some examples from the gallery (by RHADS on DeviantArt) are shown in (a).

inhomogeneous color combinations together, resulting in rather uniform color distributions as shown in Figure 9. In addition, it produces the color distributions that are dominated by unimportant but frequently occurring colors (e.g., black and white), failing to capture any interesting coloring modes. In contrast, our model can discover the color combinations that frequently co-occur in similar spatial configurations (i.e., sizes and positions), enabling our representation to capture some important spatially-varying coloring modes with more informative distributions. Refer to Section 5 of the supplemental for more comparisons.

5 CONCLUSION

In this paper, we take a step towards summarizing, visualizing and exploring the color use within large repositories of artworks. We formulate colorization as a probabilistic process. This allows us to discover probabilistic color palettes that can intuitively summarize coloring within an artwork collection. We hope that our initial solution would inspire others to further explore this exciting topic.

Our approach has several limitations. First, our approach assumes that the coloring style is consistent across the artworks within an artist's gallery. Hence, for an art gallery with mixed coloring styles or from an artist without any coloring preference, our probabilistic color palettes might not capture any meaningful and interpretable coloring patterns. Second, we assume that the coloring style of an artist does not change significantly over time, which results in a static representation. While this is often the case over a short term, it is not unusual to find that some artists evolve their coloring styles over a long period of time. One solution is to exploit dynamic topic modeling [Blei and Lafferty 2006] to model how the probabilistic color palettes evolve over time. Third, while our model has considered region-level semantics, the context of the whole artwork (e.g., happy and sad) could also affect color use [Cousins 2014], but are not handled by this work. Thus, it would be interesting to incorporate such context information into our model.

ACKNOWLEDGMENTS

We would like to thank all the reviewers for their valuable comments on the paper. This work was supported by a GRF grant from the Research Grants Council of Hong Kong (CityU 11200314).

REFERENCES

F. Bao, D. Yan, N. Mitra, and P. Wonka. 2013. Generating and exploring good building layouts. In *ACM SIGGRAPH'13*.
D. Blei. 2012. Probabilistic topic models. *Commun. ACM* 55, 4 (2012).
D. Blei and J. Lafferty. 2006. Dynamic topic models. In *ICML'06*.
L. Busche. 2015. https://designschool.canva.com/blog/color-tips/. (2015).
H. Chang, O. Fried, Y. Liu, S. DiVerdi, and A. Finkelstein. 2015. Palette-based Photo Recoloring. In *ACM SIGGRAPH'15*.
C. Chapman. 2009. http://mashable.com/2009/03/16/design-inspiration/. (2009).
D. Cohen-Or, O. Sorkine, R. Gal, T. Leyvand, and Y.Q. Xu. 2006. Color harmonization. In *ACM SIGGRAPH'06*.
Carrie Cousins. 2014. http://designshack.net/articles/inspiration/the-science-behind-color-and-emotion/. (2014).
E. J. Crowley and A. Zisserman. 2014. In Search of Art. In *Workshop on Computer Vision for Art Analysis, ECCV*.
J. Delon, A. Desolneux, J. Lisani, and A. Petro. 2007. Automatic color palette. *Inverse Problems and Imaging* 1, 2 (2007).
Designflavr. 2015. http://illustratd.com/. (2015).
DeviantART. 2015. http://www.deviantart.com/. (2015).
N. Fish, M. Averkiou, O. Kaick, O. Sorkine-Hornung, D. Cohen-Or, and N. Mitra. 2014. Meta-representation of Shape Families. In *ACM SIGGRAPH'14*.
T. Griffiths and M. Steyvers. 2004. Finding Scientific Topics. *Proceedings of the National Academy of Sciences* 101 (2004).
A. Jahanian, S. Keshvari, S. V. N. Vishwanathan, and J. Allebach. 2017. Colors – Messengers of Concepts: Visual Design Mining for Learning Color Semantics. *ACM Transactions on Computer-Human Interaction* 24, 2 (2017).
D. Kim, S. W. Son, and H. Jeong. 2014. Large-Scale Quantitative Analysis of Painting Arts. *Scientific Reports* 4, 7370 (2014).
V. Kim, W. Li, N. Mitra, S. Chaudhuri, S. DiVerdi, and T. Funkhouser. 2013. Learning Part-based Templates from Large Collections of 3D Shapes. In *ACM SIGGRAPH'13*.
Adobe Kuler. 2015. https://color.adobe.com/. (2015).
S. Lin and P. Hanrahan. 2013. Modeling how people extract color themes from images. In *ACM SIGCHI'13*.
B.J. Meier, A.M. Spalter, and D.B. Karelitz. 2004. Interactive color palette tools. *IEEE CG&A* 24, 3 (2004).
P. O'Donovan, A. Agarwala, and A. Hertzmann. 2011. Color Compatibility From Large Datasets. In *ACM SIGGRAPH'11*.
M. Ovsjanikov, W. Li, L. Guibas, and N. Mitra. 2011. Exploration of continuous variability in collections of 3D shapes. In *ACM SIGGRAPH'11*.
S.E. Palmer and K.B. Schloss. 2010. An ecological valence theory of human color preference. *Proceedings of the National Academy of Sciences* 107, 19 (2010).
D. Ritchie, A. Kejriwal, and S. Klemmer. 2011. d.tour: style-based exploration of design example galleries. In *UIST'11*.
M. Saito and Y. Matsui. 2015. Illustration2Vec: A Semantic Vector Representation of Illustrations. In *SIGGRAPH Asia Technical Briefs*.
B. Saleh and A. Elgammal. 2016. Large-scale Classification of Fine-Art Paintings: Learning The Right Metric on The Right Feature. *International Journal for Digital Art History* 2 (2016).
M. Steyvers, P. Smyth, M. Rosen-Zvi, and T. Griffiths. 2004. Probabilistic Author-Topic Models for Information Discovery. In *ACM SIGKDD'04*.
J.O. Talton, D. Gibson, L. Yang, P. Hanrahan, and V. Koltun. 2009. Exploratory modeling with collaborative design spaces. In *ACM SIGGRAPH Asia'09*.
B. Wang, Y. Yu, T. Wong, C. Chen, and Y. Xu. 2010. Data-Driven Image Color Theme Enhancement. In *ACM SIGGRAPH'10*.
A. Weeks and G. Hague. 1997. Color segmentation in the hsi color space using the k-means algorithm. In *Proc. SPIE*.
B.M. Whelan and T. Sutton. 2004. *The Complete Color Harmony: Expert Color Information for Professional Color Results.* Rockport Publishers.
J. Zhu, Y. Lee, and A. Efros. 2014. AverageExplorer: interactive exploration and alignment of visual data collections. In *ACM SIGGRAPH'14*.