

Recognizing Image Style

Sergey Karayev¹

¹ University of California, Berkeley

Matthew Trentacoste²

² Adobe

Helen Han¹

Aseem Agarwala²

Trevor Darrell¹

Aaron Hertzmann²

Holger Winnemoeller²

Abstract

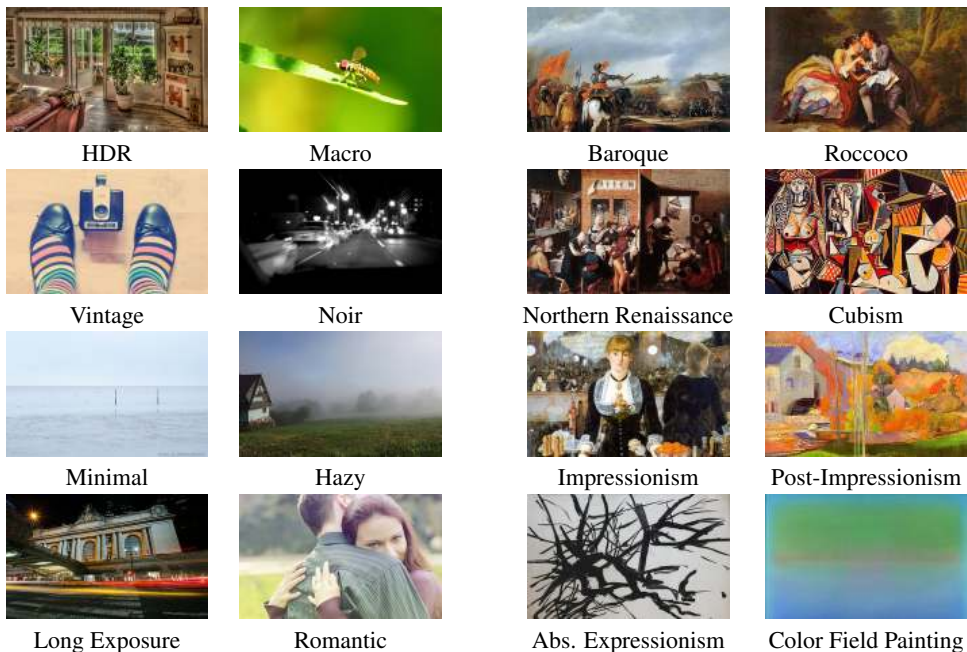
The style of an image plays a significant role in how it is viewed, but style has received little attention in computer vision research. We describe an approach to predicting style of images, and perform a thorough evaluation of different image features for these tasks. We find that features learned in a multi-layer network generally perform best – even when trained with object class (not style) labels. Our large-scale learning methods results in the best published performance on an existing dataset of aesthetic ratings and photographic style annotations. We present two novel datasets: 80K Flickr photographs annotated with 20 curated style labels, and 85K paintings annotated with 25 style/genre labels. Our approach shows excellent classification performance on both datasets. We use the learned classifiers to extend traditional tag-based image search to consider stylistic constraints, and demonstrate cross-dataset understanding of style.

1 Introduction

Deliberately-created images convey meaning, and *visual style* is often a significant component of image meaning. For example, a political candidate portrait made in the lush colors of a Renoir painting tells a different story than if it were in the harsh, dark tones of a horror movie. Distinct visual styles are apparent in art, cinematography, advertising, and have become extremely popular in amateur photography, with apps like Instagram leading the way. While understanding style is crucial to image understanding, very little research in computer vision has explored visual style.

Although is it very recognizable to human observers, visual style is a difficult concept to rigorously define. Most academic discussion of style has been in an art history context, but the distinctions between, say, Rococo versus pre-Raphaelite style are less relevant to modern photography and design. There has been some previous research in image style, but this has principally been limited to recognizing a few, well-defined optical properties, such as depth-of-field.

We define several different *types* of image style, and gather a new, large-scale dataset of photographs annotated with style labels. This dataset embodies several different aspects



Flickr Style: 80K images covering 20 styles.

Wikipaintings: 85K images for 25 art genres.

Figure 1: Typical images in different style categories of our datasets.

of visual style, including photographic techniques (“Macro,” “HDR”), composition styles (“Minimal,” “Geometric”), moods (“Serene,” “Melancholy”), genres (“Vintage,” “Romantic,” “Horror”), and types of scenes (“Hazy,” “Sunny”). These styles are not mutually exclusive, and represent different attributes of style. We also gather a large dataset of visual art (mostly paintings) annotated with art historical style labels, ranging from Renaissance to modern art. **Figure 1** shows some samples.

We test existing classification algorithms on these styles, evaluating several state-of-the-art image features. Most previous work in aesthetic style analysis has used hand-tuned features, such as color histograms. We find that deep convolutional neural network (CNN) features perform best for the task. This is surprising for several reasons: these features were trained on object class categories (ImageNet), and many styles appear to be primarily about color choices, yet the CNN features handily beat color histogram features. This leads to one conclusion of our work: mid-level features derived from object datasets are generic for style recognition, and superior to hand-tuned features.

We compare our predictors to human observers, using Amazon Mechanical Turk experiments, and find that our classifiers predict Group membership at essentially the same level of accuracy as Turkers. We also test on the AVA aesthetic prediction task [22], and show that using the “deep” object recognition features improves over the state-of-the-art results.

Applications and code. First, we demonstrate an example of using our method to search for images by style. This could be useful for applications such as product search, storytelling, and creating slide presentations. In the same vein, visual similarity search results could be

filtered by visual style, making possible queries such as “similar to this image, but more Film Noir.” Second, style tags may provide valuable mid-level features for other image understanding tasks. For example, there has increasing recent effort in understanding image meaning, aesthetics, interestingness, popularity, and emotion (for example, [10, 12, 14, 16]), and style is an important part of meaning. Finally, learned predictors could be a useful component in modifying the style of an image.

All data, trained predictors, and code (including results viewing interface) are available at <http://sergeykarayev.com/recognizing-image-style/>.

2 Related Work

Most research in computer vision addresses recognition and reconstruction, independent of image style. A few previous works have focused directly on image composition, particularly on the high-level attributes of beauty, interestingness, and memorability.

Most commonly, several previous authors have described methods to predict aesthetic quality of photographs. Datta et al. [4], designed visual features to represent concepts such as colorfulness, saturation, rule-of-thirds, and depth-of-field, and evaluated aesthetic rating predictions on photographs; The same approach was further applied to a small set of Impressionist paintings [18]. The feature space was expanded with more high-level descriptive features such as “presence of animals” and “opposing colors” by Dhar et al., who also attempted to predict Flickr’s proprietary “interestingness” measure, which is determined by social activity on the website [6]. Gygli et al. [10] gathered and predicted human evaluation of image interestingness, building on work by Isola et al. [12], who used various high-level features to predict human judgements of image memorability. In a similar task, Borth et al. [3] performed sentiment analysis on images using object classifiers trained on adjective-noun pairs.

Murray et al. [22] introduced the Aesthetic Visual Analysis (AVA) dataset, annotated with ratings by users of DPChallenge, a photographic skill competition website. The AVA dataset contains some photographic style labels (e.g., “Duotones,” “HDR”), derived from the titles and descriptions of the photographic challenges to which photos were submitted. Using images from this dataset, Marchesotti and Peronnin [20] gathered bi-grams from user comments on the website, and used a simple sparse feature selection method to find ones predictive of aesthetic rating. The attributes they found to be informative (e.g., “lovely photo,” “nice detail”) are not specific to image style.

Several previous authors have developed systems to classify classic painting styles, including [15, 25]. These works consider only a handful of styles (less than ten apiece), with styles that are visually very distinct, e.g., Pollock vs. Dalí. These datasets comprise less than 60 images per style, for both testing and training. Mensink [21] provides a larger dataset of artworks, but does not consider style classification.

3 Data Sources

Building an effective model of photographic style requires annotated training data. To our knowledge, there is only one existing dataset annotated with visual style, and only a narrow range of photographic styles is represented [22]. We would like to study a broader range of styles, including different *types* of styles ranging from genres, compositional styles, and

moods. Moreover, large datasets are desirable in order to obtain effective results, and so we would like to obtain data from online communities, such as Flickr.

Flickr Style. Although Flickr users often provide free-form tags for their uploaded images, the tags tend to be quite unreliable. Instead, we turn to Flickr groups, which are community-curated collections of visual concepts. For example, the Flickr Group “Geometry Beauty” is described, in part, as “Circles, triangles, rectangles, symmetric objects, repeated patterns”, and contains over 167K images at time of writing; the “Film Noir Mood” group is described as “Not just black and white photography, but a dark, gritty, moody feel...” and comprises over 7K images.

At the outset, we decided on a set of 20 visual styles, further categorized into types:

- **Optical techniques:** Macro, Bokeh, Depth-of-Field, Long Exposure, HDR
- **Atmosphere:** Hazy, Sunny
- **Mood:** Serene, Melancholy, Ethereal
- **Composition styles:** Minimal, Geometric, Detailed, Texture
- **Color:** Pastel, Bright
- **Genre:** Noir, Vintage, Romantic, Horror

For each of these stylistic concepts, we found at least one dedicated Flickr Group with clearly defined membership rules. From these groups, we collected 4,000 positive examples for each label, for a total of 80,000 images. Example images are shown in [Figure 1a](#). The exact Flickr groups used are given in the Supplementary Materials.

The derived labels are considered clean in the positive examples, but may be noisy in the negative examples, in the same way as the ImageNet dataset [5]. That is, a picture labeled as *Sunny* is indeed *Sunny*, but it may also be *Romantic*, for which it is not labeled. We consider this an unfortunate but acceptable reality of working with a large-scale dataset. Following ImageNet, we still treat the absence of a label as indication that the image is a negative example for that label. Mechanical Turk experiments described in [section 6.1](#) serve to allay our concerns.

Wikipaintings. We also provide a new dataset for classifying painting style. To our knowledge, no previous large-scale dataset exists for this task – although very recently a large dataset of artwork did appear for other tasks [21]. We collect a dataset of 100,000 high-art images – mostly paintings – labeled with artist, style, genre, date, and free-form tag information by a community of experts on the [Wikipaintings.org](#) website.

Analyzing style of non-photorealistic media is an interesting problem, as much of our present understanding of visual style arises out of thousands of years of developments in fine art, marked by distinct historical styles. Our dataset presents significant stylistic diversity, primarily spanning Renaissance styles to modern art movements (Supplementary Materials provides further breakdowns). We select 25 styles with more than 1,000 examples, for a total of 85,000 images. Example images are shown in [Figure 1b](#).

4 Learning algorithm

We learn to classify novel images according to their style, using the labels assembled in the previous section. Because the datasets we deal with are quite large and some of the features

are high-dimensional, we consider only linear classifiers, relying on sophisticated features to provide robustness.

We use an open-source implementation of Stochastic Gradient Descent with adaptive subgradient [1]. The learning process optimizes the function

$$\min_w \lambda_1 \|w\|_1 + \frac{\lambda_2}{2} \|w\|_2^2 + \sum_i \ell(x_i, y_i, w)$$

We set the L_1 and L_2 regularization parameters and the form of the loss function by validation on a held-out set. For the loss $\ell(x, y, w)$, we consider the hinge ($\max(0, 1 - y \cdot w^T x)$) and logistic ($\log(1 + \exp(-y \cdot w^T x))$) functions. We set the initial learning rate to 0.5, and use adaptive subgradient optimization [8]. Our setup is of multi-class classification; we use the One vs. All reduction to binary classifiers.

5 Image Features

In order to classify styles, we must choose appropriate image features. We hypothesize that image style may be related to many different features, including low-level statistics [19], color choices, composition, and content. Hence, we test features that embody these different elements, including features from the object recognition literature. We evaluate single-feature performance, as well as second-stage fusion of multiple features.

L*a*b color histogram. Many of the Flickr styles exhibit strong dependence on color. For example, *Noir* images are nearly all black-and-white, while most *Horror* images are very dark, and *Vintage* images use old photographic colors. We use a standard color histogram feature, computed on the whole image. The 784-dimensional joint histogram in CIELAB color space has 4, 14, and 14 bins in the L*, a*, and b* channels, following Palermo et al. [24], who showed this to be the best performing single feature for determining the date of historical color images.

GIST. The classic gist descriptor [23] is known to perform well for scene classification and retrieval of images visually similar at a low-resolution scale, and thus can represent image composition to some extent. We use the INRIA LEAR implementation, resizing images to 256 by 256 pixels and extracting a 960-dimensional color GIST feature.

Graph-based visual saliency. We also model composition with a visual attention feature [11]. The feature is fast to compute and has been shown to predict human fixations in natural images basically as well as an individual human (humans are far better in aggregate, however). The 1024-dimensional feature is computed from images resized to 256 by 256 pixels.

Meta-class binary features. Image content can be predictive of individual styles, e.g., *Macro* images include many images of insects and flowers. The mc-bit feature [2] is a 15,000-dimensional bit vector feature learned as a non-linear combination of classifiers trained using existing features (e.g., SIFT, GIST, Self-Similarity) on thousands of random ImageNet synsets, including internal ILSVRC2010 nodes. In essence, MC-bit is a hand-crafted “deep” architecture, stacking classifiers and pooling operations on top of lower-level features.

Table 1: Mean APs on three datasets for the considered single-channel features and their second-stage combination. As some features were clearly worse than others on the AVA Style dataset, only the better features were evaluated on larger datasets.

	Fusion x Content	DeCAF ₆	MC-bit	L*a*b* Hist	GIST	Saliency	random
AVA Style	0.581	0.579	0.539	0.288	0.220	0.152	0.132
Flickr	0.368	0.336	0.328	-	-	-	0.052
Wikipaintings	0.473	0.356	0.441	-	-	-	0.043

Deep convolutional net. Current state-of-the-art results on ImageNet, the largest image classification challenge, have come from a deep convolutional network trained in a fully-supervised manner [17]. We use the Caffe [13] open-source implementation of the ImageNet-winning eight-layer convolutional network, trained on over a million images annotated with 1,000 ImageNet classes. We investigate using features from two different levels of the network, referred to as DeCAF₅ and DeCAF₆ (following [7]). The features are 8,000- and 4,000-dimensional and are computed from images center-cropped and resized to 256 by 256 pixels.

Content classifiers. Following Dhar et al. [6], who use high-level classifiers as features for their aesthetic rating prediction task, we evaluate using object classifier confidences as features. Specifically, we train classifiers for all 20 classes of the PASCAL VOC [9] using the DeCAF₆ feature. The resulting classifiers are quite reliable, obtaining 0.7 mean AP on the VOC 2012.

We aggregate the data to train four classifiers for “animals”, “vehicles”, “indoor objects” and “people”. These aggregate classes are presumed to discriminate between vastly different types of images – types for which different style signals may apply. For example, a *Romantic* scene with people may be largely about the composition of the scene, whereas, *Romantic* scenes with vehicles may be largely described by color.

To enable our classifiers to learn content-dependent style, we can take the outer product of a feature channel with the four aggregate content classifiers.

6 Experiments

6.1 Flickr Style

We learn and predict style labels on the 80,000 images labeled with 20 different visual styles of our new Flickr Style dataset, using 20% of the data for testing, and another 20% for parameter-tuning validation.

There are several performance metrics we consider. Average Precision evaluation (as reported in Table 1 and in detailed tables in the Supplementary Materials) is computed on a random class-balanced subset of the test data (each class has equal prevalence). We compute confusion matrices on the same data. Per-class accuracies are computed on subsets of the data balanced by the binary label, such that chance performance is 50%. We follow these decisions in all following experiments.

The best single-channel feature is DeCAF₆ with 0.336 mean AP; feature fusion obtains 0.368 mean AP. Per-class APs range from 0.17 [Depth of Field] to 0.62 [Macro]. Per-class accuracies range from 68% [Romantic, Depth of Field] to 85% [Sunny, Noir, Macro]. The

average per-class accuracy is 78%. We show the most confident style classifications on the test set of Flickr Style in Figure 3.

Upon inspection of the confusion matrices, we saw points of understandable confusion: Depth of Field vs. Macro, Romantic vs. Pastel, Vintage vs. Melancholy. There are also surprising sources of mistakes: Macro vs. Bright/Energetic, for example. To explain this particular confusion, we observed that lots of Macro photos contain bright flowers, insects, or birds, often against vibrant greenery. Here, at least, the content of the image dominates its style label.

To explore further content-style correlations, we plot the outputs of PASCAL object class classifiers (one of our features) on the Flickr dataset in Figure 2. We can observe that some styles have strong correlations to content (e.g., “Hazy” occurs with “vehicle”, “HDR” doesn’t occur with “cat”).

We hypothesize that style is content-dependent: a Romantic portrait may have different low-level properties than a Romantic sunset. We form a new feature as an outer product of our content classifier features with the second-stage late fusion features (“Fusion \times Content” in all results figures). These features gave the best results, thus supporting the hypothesis.

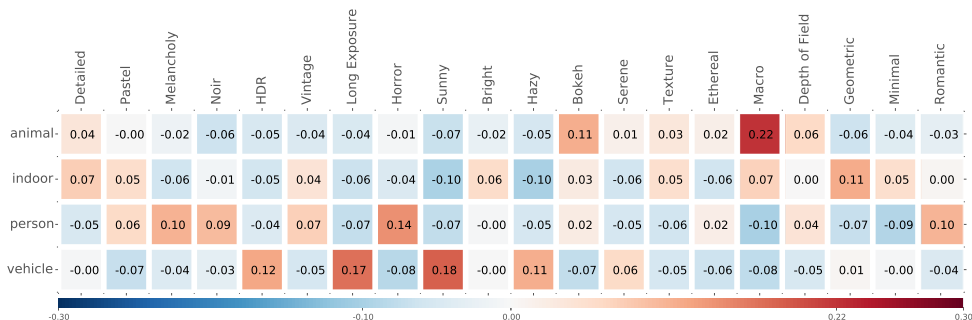


Figure 2: Correlation of PASCAL content classifier predictions (rows) against ground truth Flickr Style labels (columns). We see, for instance, that the Macro style is highly correlated with presence of animals, and that Long Exposure and Sunny style photographs often feature vehicles.

Mechanical Turk Evaluation. In order to provide a human baseline for evaluation, we performed a Mechanical Turk study. For each style, Turkers were shown positive and negative examples for each Flickr Group, and then they evaluated whether each image in the test set was part of the given style. We treat the Flickr group memberships as ground truth as before, and then evaluate Turkers’ ability to accurately determine group membership. Measures were taken to remove spam workers; see the Supplemental Material details on the experimental setup. For efficiency, one quarter of the test set was used, and two redundant styles (Bokeh and Detailed) were removed. Each test image was evaluated by 3 Turkers, and the majority vote taken as the human result for this image.

In total, Turkers achieved 75% mean accuracy (ranging from 61% [Romantic] to 92% [Macro]) across styles, in comparison to 78% mean accuracy (ranging from 68% [Depth of Field] to 87% [Macro]) of our best method. Our algorithm did significantly worse than Turkers on Macro and Horror, and significantly better on Vintage, Romantic, Pastel, Detailed, HDR, and Long Exposure styles.

Some of this variance may be due to subtle difference from the Turk tasks that we provided, as compared to the definitions of the Flickr groups, but may also due to the Flickr groups’ incorporating images that do not quite fit the common definition of the given style. For example, there may be a mismatch between different notions of “romantic” and “vintage,” and how inclusively these terms are defined.

We additionally used the Turker opinion as ground truth for our method’s predictions. In switching from the default Flickr to the MTurk ground truth, our method’s accuracy hardly changed from 78% to 77%. However, we saw that the accuracy of our Vintage, Detailed, Long Exposure, Minimal, HDR, and Sunny style classifiers significantly decreased, indicating machine-human disagreement on those styles. Detailed tables are provided in Supplemental Results.

6.2 Wikipaintings

With the same setup and features as in the Flickr experiments, we evaluate 85,000 images labeled with 25 different art styles. The results are given in [Table 1](#) and in Supplemental Materials. The best single-channel feature is MC-bit with 0.441 mean AP; feature fusion obtains 0.473 mean AP. Per-class accuracies range from 72% [Symbolism, Expressionism, Art Nouveau] to 94% [Ukiyo-e, Minimalism, Color Field Painting].

6.3 AVA Style

AVA [22] is a dataset of 250K images from `dpchallenge.net`. We evaluate classification of aesthetic rating and of 14 different photographic style labels on the 14,000 images of the AVA dataset that have such labels. For the style labels, the publishers of the dataset provide a train/test split, where training images have only one label, but test images may have more than one label [22]. For style classification, the best single feature is the DeCAF₆ convolution network feature, obtaining 0.579 mean AP. Feature fusion improves the result to 0.581 mean AP; both results beat the previous state-of-the-art of 0.538 mean AP [22].

In all metrics, the DeCAF and MC-bit features significantly outperformed more low-level features on this dataset. Accordingly, we do not evaluate the low-level features on the larger Flickr and Wikipaintings datasets.

6.4 Application: Style-Based Image Search

Style classifiers learned on our datasets can be used toward novel goals. For example, sources of stock photography or design inspiration may be better navigated with a vocabulary of style. Currently, companies expend labor to manually annotate stock photography with such labels. With our approach, any image collection can be searchable and rankable by style.

To demonstrate, we apply our Flickr-learned style classifiers to a new dataset of 80K images gathered on Pinterest (also available with our code release); some results are shown in [Figure 5](#). Interestingly, styles learned from photographs can be used to order paintings, and styles learned from paintings can be used to order photographs, as illustrated in [Figure 4](#).

6.5 Discussion

We have made significant progress in defining the problem of understanding photographic style. We provide a novel dataset that exhibits several types of styles not previously consid-

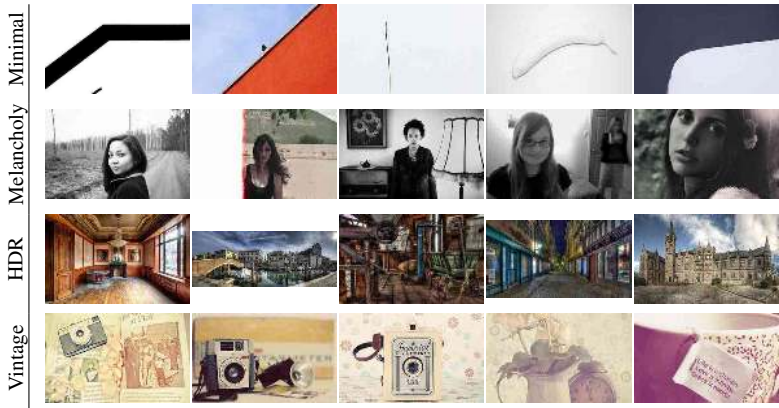


Figure 3: Top five most-confident positive predictions on the Flickr Style test set, for a few different styles. See Figures 1-3 of the Supplemental Material for more results.

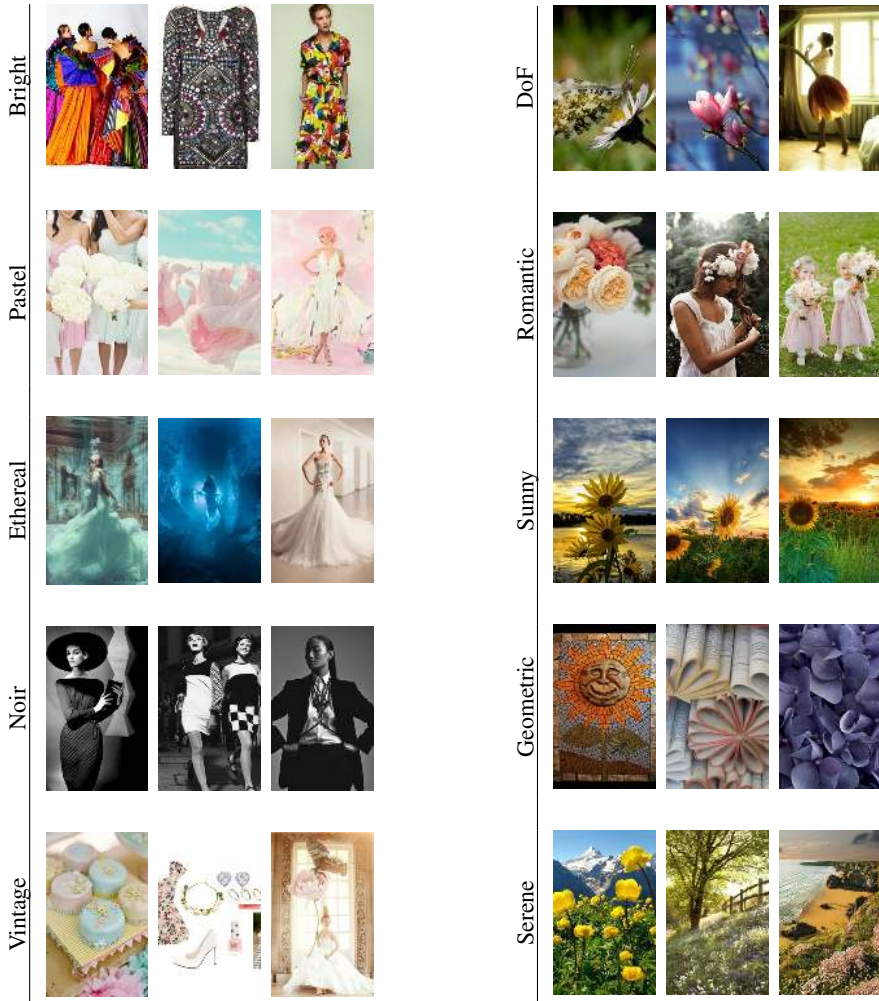


Figure 4: Cross-dataset style. On the left are shown top scorers from the Wikipaintings set, for styles learned on the Flickr set. On the right, Flickr photographs are accordingly sorted by Painting style. (Figure best viewed in color.)

ered in the literature, and we demonstrate state-of-the-art results in prediction of both style and aesthetic quality. These results are comparable to human performance. We also show that style is highly content-dependent.

Style plays a significant role in much of the manmade imagery we experience daily, and there is considering need for future work to further answer the question “What is style?”

One of the most interesting outcomes of this work is the success of features trained for object detection for both aesthetic and style classification. We propose several possible hypotheses to explain these results. Perhaps the network layers that we use as features are extremely good as general visual features for image representation in general. Another explanation is that object recognition depends on object appearance, e.g., distinguishing red from white wine, or different kinds of terriers, and that the model learns to repurpose these features for image style. Understanding and improving on these results is fertile ground for future work.



Query: "dress".

Query: "flower".

Figure 5: Example of filtering image search results by style. Our Flickr Style classifiers are applied to images found on Pinterest. The images are searched by the text contents of their captions, then filtered by the response of the style classifiers. Here we show three out of top five results for different query/style combinations.

References

- [1] Alekh Agarwal, Olivier Chapelle, Miroslav Dudik, and John Langford. A Reliable Effective Terascale Linear Learning System. *Journal of Machine Learning Research*, 2012.
- [2] A. Bergamo and L. Torresani. Meta-class Features for Large-Scale Object Categorization on a Budget. In *CVPR*, 2012.
- [3] Damian Borth, Rongrong Ji, Tao Chen, and Thomas M Breuel. Large-scale Visual Sentiment Ontology and Detectors Using Adjective Noun Pairs. In *ACM MM*, 2013.
- [4] Ritendra Datta, Dhiraj Joshi, Jia Li, and James Z Wang. Studying Aesthetics in Photographic Images Using a Computational Approach. In *ECCV*, 2006.
- [5] Jia Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009.
- [6] Sagnik Dhar, Vicente Ordonez, and Tamara L Berg. High Level Describable Attributes for Predicting Aesthetics and Interestingness. In *CVPR*, 2011.
- [7] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, Trevor Darrell, Trevor Eecs, and Berkeley Edu. DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition. Technical report, 2013. arXiv:1310.1531v1.
- [8] John Duchi, Elad Hazan, and Yoram Singer. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *Journal of Machine Learning Research*, 2011.
- [9] M Everingham, L Van Gool, C K I Williams, J Winn, and A Zisserman. The PASCAL VOC Challenge Results, 2010.
- [10] Michael Gygli, Fabian Nater, and Luc Van Gool. The Interestingness of Images. In *ICCV*, 2013.
- [11] Jonathan Harel, Christof Koch, and Pietro Perona. Graph-Based Visual Saliency. In *NIPS*, 2006.
- [12] Phillip Isola, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. What Makes an Image Memorable? In *CVPR*, 2011.
- [13] Yangqing Jia. Caffe: an open source convolutional architecture for fast feature embedding. <http://caffe.berkeleyvision.org/>, 2013.
- [14] Jungseock Joo, Weixin Li, Francis Steen, and Song-Chun Zhu. Visual Persuasion: Inferring Communicative Intents of Images. In *CVPR*, 2014.
- [15] Daniel Keren. Painter Identification Using Local Features and Naive Bayes. In *ICPR*, 2002.
- [16] A. Khosla, A. Das Sarma, and R. Hamid. What Makes an Image Popular? In *WWW*, 2014.
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoff E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *NIPS*, 2012.
- [18] Congcong Li and Tsuhan Chen. Aesthetic Visual Quality Assessment of Paintings. *IEEE Journal of Selected Topics in Signal Processing*, 3(2):236–252, 2009.
- [19] Siwei Lyu, Daniel Rockmore, and Hany Farid. A Digital Technique for Art Authentication. *PNAS*, 101(49), 2004.
- [20] Luca Marchesotti and Florent Perronnin. Learning beautiful (and ugly) attributes. In *BMVC*, 2013.
- [21] Thomas Mensink and Jan van Gemert. The Rijksmuseum Challenge: Museum-Centered Visual Recognition. In *ICMR*, 2014.
- [22] Naila Murray, De Barcelona, Luca Marchesotti, and Florent Perronnin. AVA: A Large-Scale Database for Aesthetic Visual Analysis. In *CVPR*, 2012.
- [23] Aude Oliva and Antonio Torralba. Modeling the Shape of the Scene: A Holistic Representation of the Spatial Envelope. *IJCV*, 42(3):145–175, 2001.
- [24] Frank Palermo, James Hays, and Alexei A Efros. Dating Historical Color Images. In *ECCV*, 2012.
- [25] Lior Shamir, Tomasz Macura, Nikita Orlov, D. Mark Eckley, and Ilya G. Goldberg. Impressionism, Expressionism, Surrealism: Automated Recognition of Painters and Schools of Art. *ACM Trans. Applied Perc.*, 7(2), 2010.