

Photo Aesthetics Ranking Network with Attributes and Content Adaptation

Shu Kong¹(✉), Xiaohui Shen², Zhe Lin², Radomir Mech²,
and Charless Fowlkes¹

¹ UC Irvine, Irvine, USA

{skong2,fowlkes}@ics.uci.edu

² Adobe Research, San Jose, USA

{xshen,zlin,rmech}@adobe.com

Abstract. Real-world applications could benefit from the ability to automatically generate a fine-grained ranking of photo aesthetics. However, previous methods for image aesthetics analysis have primarily focused on the coarse, binary categorization of images into high- or low-aesthetic categories. In this work, we propose to learn a deep convolutional neural network to rank photo aesthetics in which the relative ranking of photo aesthetics are directly modeled in the loss function. Our model incorporates joint learning of meaningful photographic attributes and image content information which can help regularize the complicated photo aesthetics rating problem.

To train and analyze this model, we have assembled a new aesthetics and attributes database (AADB) which contains aesthetic scores and meaningful attributes assigned to each image by multiple human raters. Anonymized rater identities are recorded across images allowing us to exploit intra-rater consistency using a novel sampling strategy when computing the ranking loss of training image pairs. We show the proposed sampling strategy is very effective and robust in face of subjective judgement of image aesthetics by individuals with different aesthetic tastes. Experiments demonstrate that our unified model can generate aesthetic rankings that are more consistent with human ratings. To further validate our model, we show that by simply thresholding the estimated aesthetic scores, we are able to achieve state-of-the-art classification performance on the existing AVA dataset benchmark.

Keywords: Convolutional neural network · Image aesthetics rating · Rank loss · Attribute learning

1 Introduction

Automatically assessing image aesthetics is increasingly important for a variety of applications [1, 2], including personal photo album management, automatic photo editing, and image retrieval. While judging image aesthetics is a subjective task, it has been an area of active study in recent years and substantial progress has been made in identifying and quantifying those image features that are predictive of favorable aesthetic judgements by most individuals [1–5].



Fig. 1. Classification-based methods for aesthetic analysis can distinguish high- and low-quality images shown in the leftmost and rightmost columns, but fail to provide useful insights about borderline images displayed in the middle column. This observation motivates us to consider rating and ranking images w.r.t aesthetics rather than simply assigning binary labels. We observe that the contribution of particular photographic attributes to making an image aesthetically pleasing depends on the thematic content (shown in different rows), so we develop a model for rating that incorporates joint attributes and content. The attributes and ratings of aesthetics on a scale 1 to 5 are predicted by our model (displayed on top and right of each image, respectively).

Early works formulate aesthetic analysis as a classification or a regression problem of mapping images to aesthetic ratings provided by human raters [4–8]. Some approaches have focused on designing hand-crafted features that encapsulate standard photographic practice and rules of visual design, utilizing both low-level statistics (*e.g.* color histogram and wavelet analysis) and high-level cues based on traditional photographic rules (*e.g.* region composition and rule of thirds). Others have adopted generic image content features, which are originally designed for recognition (*e.g.* SIFT [9] and Fisher Vector [10, 11]), that have been found to outperform methods using rule-based features [12]. With the advance of deep Convolutional Neural Network (CNN) [13], recent works propose to train end-to-end models for image aesthetics classification [2, 3, 14], yielding state-of-the-art performance on a recently released Aesthetics Visual Analysis dataset (AVA) [15].

Despite notable recent progress towards computational image aesthetics classification (*e.g.* [1–3]), judging image aesthetics is still a subjective task, and it is difficult to learn a universal scoring mechanism for various kinds of images. For example, as demonstrated in Fig. 1, images with obviously visible high- or low-aesthetics are relatively easy to classify, but existing methods cannot generate reliable labels for borderline images. Therefore, instead of formulating image aesthetics analysis as an overall binary classification or regression problem, we argue that it is far more practical and useful to predict relative aesthetic rankings among images with similar visual content along with generating richer descriptions in terms of aesthetic attributes [16, 17].

To this end, we propose to train a model through a Siamese network [18] that takes a pair of images as input and directly predicts relative ranking of their aesthetics in addition to their overall aesthetic scores. Such a structure allows us to

deploy different sampling strategies of image pairs and leverage auxiliary side-information to regularize the training, including aesthetic attributes [1, 3, 7] and photo content [4, 15, 19]. For example, Fig. 1 demonstrates that photos with different contents convey different attributes to make them aesthetically pleasing. While such side information has been individually adopted to improve aesthetics classification [1, 3], it remains one open problem to systematically incorporate all the needed components in a single end-to-end framework with fine-grained aesthetics ranking. Our model and training procedure naturally incorporates both attributes and content information by sampling image pairs with similar content to learn the specific relations of attributes and aesthetics for different content sub-categories. As we show, this results in more comparable and consistent aesthetics estimation results.

Moreover, as individuals have different aesthetics tastes, we argue that it is important to compare ratings assigned by an individual across multiple images in order to provide a more consistent training signal. To this end, we have collected and will publicly release a new dataset in which each image is associated with a detailed score distribution, meaningful attributes annotation and (anonymized) raters' identities. We refer to this dataset as the "Aesthetics with Attributes Database", or AADB for short. AADB not only contains a much more balanced distribution of professional and consumer photos and a more diverse range of photo qualities than available in the exiting AVA dataset, but also identifies ratings made by the same users across multiple images. This enables us to develop novel sampling strategies for training our model which focuses on relative rankings by individual raters. Interestingly, this rater-related information also enables us to compare the trained model to each individual's rating results by computing the ranking correlation over test images rated by that individual. Our experiments show the effectiveness of the proposed model in rating image aesthetics compared to human individuals. We also show that, by simply thresholding rated aesthetics scores, our model achieves state-of-the-art classification performance on the AVA dataset, even though we do not explicitly train or tune the model for the aesthetic classification task.

In summary, our main contributions are three-fold:

1. We release a new dataset containing not only score distributions, but also informative attributes and anonymized rater identities. These annotations enable us to study the use of individuals' aesthetics ratings for training our model and analyze how the trained model performs compared to individual human raters.
2. We propose a new CNN architecture that unifies aesthetics attributes and photo content for image aesthetics rating and achieves state-of-the-art performance on existing aesthetics classification benchmark.
3. We propose a novel sampling strategy that utilizes mixed within- and cross-rater image pairs for training models. We show this strategy, in combination with pair-wise ranking loss, substantially improves the performance *w.r.t.* the ranking correlation metric.

2 Related Work

CNN for Aesthetics Classification: In [2, 3, 14], CNN-based methods are proposed for classifying images into high- or low-aesthetic categories. The authors also show that using patches from the original high-resolution images largely improves the performance. In contrast, our approach formulates aesthetic prediction as a combined regression and ranking problem. Rather than using patches, our architecture warps the whole input image in order to minimize the overall network size and computational workload while retaining compositional elements in the image, *e.g.* rule of thirds, which are lost in patch-based approaches.

Attribute-Adaptive Models: Some recent works have explored the use of high-level describable attributes [1, 3, 7] for image aesthetics classification. In early work, these attributes were modeled using hand-crafted features [7]. This introduces some intrinsic problems, since (1) engineering features that capture high-level semantic attributes is a difficult task, and (2) the choice of describable attributes may ignore some aspects of the image which are relevant to the overall image aesthetics. For these reasons, Marchesotti *et al.* propose to automatically select a large number of useful attributes based on textual comments from raters [20] and model these attributes using generic features [12]. Despite good performance, many of the discovered textual attributes (*e.g.* *so_cute*, *those_eyes*, *so_close*, *very_busy*, *nice_try*) do not correspond to well defined visual characteristics which hinders their detectability and utility in applications. Perhaps the closest work to our approach is that of Lu *et al.*, who propose to learn several meaningful style attributes [3] in a CNN framework and use the hidden features to regularize aesthetics classification network training.

Content-Adaptive Models: To make use of image content information such as scene categories or choice of photographic subject, Luo *et al.* propose to segment regions and extract visual features based on the categorization of photo content [4]. Other work, such as [15, 19], has also demonstrated that image content is useful for aesthetics analysis. However, it has been assumed that the category labels are provided both during training and testing. To our knowledge, there is only one paper [21] that attempts to jointly predict content semantics and aesthetics labels. In [21], Murray *et al.* propose to rank images w.r.t aesthetics in a three-way classification problem (high-, medium- and low-aesthetics quality). However, their work has some limitations because (1) deciding the thresholds between nearby classes is non-trivial, and (2) the final classification model outputs a hard label which is less useful than a continuous rating.

Our work is thus unique in presenting a unified framework that is trained by jointly incorporating the photo content, the meaningful attributes and the aesthetics rating in a single CNN model. We train a category-level classification layer on top of our aesthetics rating network to generate soft weights of category labels, which are used to combine scores predicted by multiple content-adaptive branches. This allows category-specific subnets to complement each other in rating image aesthetics with shared visual content information while efficiently re-using front-end feature computations. While our primary focus is

Table 1. Comparison of the properties of current image aesthetics datasets. In addition to score distribution and meaningful style attributes, AADB also tracks raters’ identities across images which we exploit in training to improve aesthetic ranking models.

	AADB	AVA [15]	PN [5]	CUHKPQ [6,22]
Rater’s ID	Y	N	N	N
All real photo	Y	N	Y	Y
Attribute label	Y	Y	N	N
Score dist	Y	Y	Y	N



Fig. 2. Our AADB dataset consists of a wide variety of photographic imagery of real scenes collected from Flickr. This differs from AVA which contains significant numbers of professional images that have been highly manipulated, overlayed with advertising text, etc.

on aesthetic rating prediction, we believe that the content and attribute predictions (as displayed on the right side of images in Fig. 1) represented in hidden layers of our architecture could also be surfaced for use in other applications such as automatic image enhancement and image retrieval.

3 Aesthetics and Attributes Database

To collect a large and varied set of photographic images, we download images from the Flickr website¹ which carry a Creative Commons license and manually curate the data set to remove non-photographic images (*e.g.* cartoons, drawings, paintings, ads images, adult-content images, etc.). We have five different workers then independently annotate each image with an overall aesthetic score and a fixed set of eleven meaningful attributes using Amazon Mechanical Turk (AMT)². The AMT raters work on batches, each of which contains ten images. For each image, we average the ratings of five raters as the ground-truth aesthetic score. The number of images rated by a particular worker follows long tail distribution, as shown later in Fig. 6 in the experiment.

After consulting professional photographers, we selected eleven attributes that are closely related to image aesthetic judgements: *interesting_content*,

¹ www.flickr.com.

² www.mturk.com.

object_emphasis, *good_lighting*, *color_harmony*, *vivid_color*, *shallow_depth_of_field*, *motion_blur*, *rule_of_thirds*, *balancing_element*, *repetition*, and *symmetry*. These attributes span traditional photographic principals of color, lighting, focus and composition, and provide a natural vocabulary for use in applications, such as auto photo editing and image retrieval. The final AADB dataset contains 10,000 images in total, each of which have aesthetic quality ratings and attribute assignments provided by five different individual raters. Aggregating multiple raters allows us to assign a confidence score to each attribute, unlike, e.g., AVA where attributes are binary. Similar to previous rating datasets [15], we find that average ratings are well fit by a Gaussian distribution. For evaluation purposes, we randomly split the dataset into validation (500), testing (1,000) and training sets (the rest). The supplemental material provides additional details about dataset collection and statistics of the resulting data.

Table 1 provides a summary comparison of AADB to other related public databases for image aesthetics analysis. Except for our AADB and the existing AVA dataset, many existing datasets have two intrinsic problems (as discussed in [15]), (1) they do not provide full score distributions or style attribute annotation, and (2) images in these datasets are either biased or consist of examples which are particularly easy for binary aesthetics classification. Datasets such as CUHKPQ [6, 22] only provide binary labels (low or high aesthetics) which cannot easily be used for rating prediction. A key difference between our dataset and AVA is that many images in AVA are heavily edited or synthetic (see Fig. 2) while AADB contains a much more balanced distribution of professional and consumer photos. More importantly, AVA does not provide any way to identify ratings provided by the same individual for multiple images. We report results of experiments, showing that rater identity on training data provides useful side information for training improved aesthetic predictors.

Consistency Analysis of the Annotation: One concern is that the annotations provided by five AMT workers for each image may not be reliable given the subjective nature of the task. Therefore, we conduct consistency analysis on the annotations. Since the same five workers annotate a batch of ten images, we study the consistency at batch level. We use Spearman’s rank correlation ρ between pairs of workers to measure consistency within a batch and estimate p -values to evaluate statistical significance of the correlation relative to a null hypothesis of uncorrelated responses. We use the Benjamini-Hochberg procedure to control the false discovery rate (FDR) for multiple comparisons [23]. At an FDR level of 0.05, we find 98.45 % batches have significant agreement among raters. This shows that the annotations are reliable for scientific research. Further consistency analysis of the dataset can be found in the supplementary material.

4 Fusing Attributes and Content for Aesthetics Ranking

Inspired by [2, 24], we start by fine-tuning AlexNet [13] using regression loss to predict aesthetic ratings. We then fine-tune a Siamese network [18] which takes

image pairs as input and is trained with a joint Euclidean and ranking loss (Sect. 4.2). We then append attribute (Sect. 4.3) and content category classification layers (Sect. 4.4) and perform joint optimization.

4.1 Regression Network for Aesthetics Rating

The network used in our image aesthetics rating is fine-tuned from AlexNet [13] which is used for image classification. Since our initial model predicts a continuous aesthetic score other than category labels, we replace the softmax loss with the Euclidean loss given by $loss_{reg} = \frac{1}{2N} \sum_{i=1}^N \|\hat{y}_i - y_i\|_2^2$, where y_i is the average ground-truth rating for image- i , and \hat{y}_i is the estimated score by the CNN model. Throughout our work, we re-scale all the ground-truth ratings to be in the range of $[0, 1]$ when preparing the data. Consistent with observations in [2], we find that fine-tuning the pre-trained AlexNet [13] model performs better than that training the network from scratch.

4.2 Pairwise Training and Sampling Strategies

A model trained solely to minimize the Euclidean loss may still make mistakes in the relative rankings of images that have similar average aesthetic scores. However, more accurate fine-grained ranking of image aesthetics is quite important in applications (*e.g.* in automating photo album management [25]). Therefore, based on the Siamese network [18], we adopt a pairwise ranking loss to explicitly exploit relative rankings of image pairs available in the AADB data (see Fig. 3(a)). The ranking loss is given by:

$$loss_{rank} = \frac{1}{2N} \sum_{i,j} \max(0, \alpha - \delta(y_i \geq y_j)(\hat{y}_i - \hat{y}_j)) \quad (1)$$

where $\delta(y_i \geq y_j) = \begin{cases} 1, & \text{if } y_i \geq y_j \\ -1, & \text{if } y_i < y_j \end{cases}$, and α is a specified margin parameter.

By adjusting this margin and the sampling of image pairs, we can avoid the

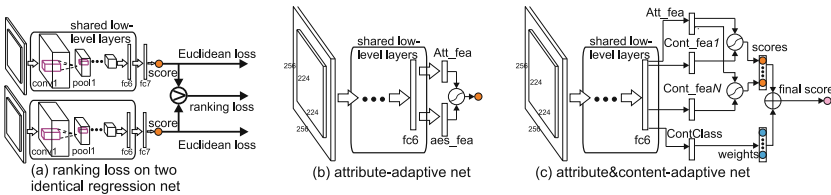


Fig. 3. Architectures for our different models. All models utilize the AlexNet front-end architecture which we augment by (a) replacing the top softmax layer with a regression net and adopting ranking loss in addition to Euclidean loss for training, (b) adding an attribute predictor branch which is then fused with the aesthetic branch to produce a final attribute-adapted rating and (c) incorporating image content scores that act as weights to gate the combination of predictions from multiple content-specific branches.

need to sample triplets as done in previous work on learning domain-specific similarity metrics [18, 26, 27]. Note that the regression alone focuses the capacity of the network on predicting the commonly occurring range of scores, while ranking penalizes mistakes for extreme scores more heavily.

In order to anchor the scores output by the ranker to the same scale as user ratings, we utilize a joint loss function that includes both ranking and regression:

$$loss_{reg+rank} = loss_{reg} + \omega_r loss_{rank}, \quad (2)$$

where the parameter ω_r controls the relative importance of the ranking loss and is set based on validation data. The network structure is shown in Fig. 3(a).

Such a structure allows us to utilize different pair-sampling strategies to narrow the scope of learning and provide more consistent training. In our work, we investigate two strategies for selecting pairs of images used in computing the ranking loss. First, we can bias sampling towards pairs of images with a relatively large difference in their average aesthetic scores. For these pairs, the ground-truth rank order is likely to be stable (agreed upon by most raters). Second, as we have raters’ identities across images, we can sample image pairs that have been scored by the same individual. While different raters may have different aesthetics tastes which erode differences in the average aesthetic score, we expect a given individual should have more consistent aesthetic judgements across multiple images. We show the empirical effectiveness of these sampling strategies in Sect. 5.

4.3 Attribute-Adaptive Model

Previous work on aesthetic prediction has investigated the use of attribute labels as input features for aesthetics classification (*e.g.* [7]). Rather than independently training attribute classifiers, we propose to include additional activation layers in our ranking network that are trained to encode informative attributes. We accomplish this by including an additional term in the loss function that encourages the appropriate attribute activations. In practice, annotating attributes for each training image is expensive and time consuming. This approach has the advantage that it can be used even when only a subset of training data comes with attribute annotations. Our approach is inspired by [3] which also integrates attribute classifiers, but differs in that the attribute-related layer shares the same front-end feature extraction with the aesthetic score predictor (see Fig. 3(b)). The attribute prediction task can thus be viewed as a source of side-information or “deep supervision” [28] that serves to regularize the weights learned during training even though it is not part of the test-time prediction, though could be enabled when needed.

We add an attribute prediction branch on top of the second fully-connected layer in the aesthetics-rating network described previously. The attribute predic-

tions from this layer are concatenated with the base model to predict the final aesthetic score. When attribute annotations are available, we utilize a K -way softmax loss or Euclidean loss, denoted by $loss_{att}$, for the attribute activations and combine it with the rating and ranking losses

$$loss = loss_{reg} + \omega_r loss_{rank} + \omega_a loss_{att} \quad (3)$$

where ω_a controls relative importance of attribute fine-tuning. If we do not have enough data with attribute annotations, we can freeze the attribute layer and only fine-tune through the other half of the concatenation layer.

4.4 Content-Adaptive Model

The importance of particular photographic attributes depends strongly on image content [4]. For example, as demonstrated by Fig. 1, vivid color and rule of thirds are highly relevant in rating landscapes but not for closeup portraits. In [15, 19], contents at the category level are assumed to be given in both training and testing stages, and category-specific models are then trained or fine-tuned. Here we propose to incorporate the category information into our model for joint optimization and prediction, so that the model can also work on those images with unknown category labels.

We fine-tune the top two layers of AlexNet [13] with softmax loss to train a content-specific branch to predict category labels³ (as shown by ContClass layer in Fig. 3(c)). Rather than making a hard category selection, we use the softmax output as a weighting vector for combining the scores produced by the category specific branches, each of which is a concatenation of attribute feature and content-specific features (denoted by Att_fea and Cont_fea respectively in Fig. 3(c)). This allows for content categories to be non-exclusive (*e.g.* a photo of an individual in a nature scene can utilize attributes for either portrait and scenery photos). During training, When fine-tuning the whole net as in Fig. 3(c), we freeze the content-classification branch and fine-tune the rest network.

4.5 Implementation Details

We warp images to 256×256 and randomly crop out a 227×227 window to feed into the network. The initial learning rate is set at 0.0001 for all layers, and periodically annealed by 0.1. We set weight decay $1e - 5$ and momentum 0.9. We use Caffe toolbox [29] extended with our ranking loss for training all the models.

To train attribute-adaptive layers, we use softmax loss on AVA dataset which only has binary labels for attributes, and the Euclidean loss on the AADB dataset which has finer-level attribute scores. We notice that, on the AVA dataset, our attribute-adaptive branch yields 59.11 % AP and 58.73 % mAP for attribute

³ Even though category classification uses different features from those in aesthetics rating, we assume the low-level features can be shared across aesthetics and category levels.

prediction, which are comparable to the reported results of style-classification model fine-tuned from AlexNet [2]. When learning content-adaptive layers on the AVA dataset for classifying eight categories, we find the content branch yields 59% content classification accuracy on the testing set. If we fine-tune the whole AlexNet, we obtain 62% classification accuracy. Note that we are not pursuing the best classification performance on either attributes or categories. Rather, our aim is to train reasonable branches that perform well enough to help with image aesthetics rating.

5 Experimental Results

To validate our model for rating image aesthetics, we first compare against several baselines including the intermediate models presented in Sect. 4, then analyze the dependence of model performance on the model parameters and structure, and finally compare performance of our model with human annotation in rating image aesthetics.

5.1 Benchmark Datasets

AADB dataset contains 10,000 images in total, with detailed aesthetics and attribute ratings, and anonymized raters' identity for specific images. We split the dataset into training (8,500), validation (500) and testing (1,000) sets. Since our dataset does not include ground-truth image content tags, we use clustering to find semantic content groups prior to training content adaptive models. Specifically, we represent each image using the **fc7** features, normalize the feature vector to be unit Euclidean length, and use unsupervised k -means for clustering. In our experimental comparison, we cluster training images into $k = 10$ content groups, and transform the distances between a testing image and the centroids into prediction weights using a softmax. The value of k was chosen using validation data (see Sect. 5.3). Figure 4 shows samples from four of these clusters, from which we observe consistencies within each cluster and distinctions across clusters.

AVA dataset contains approximately 250,000 images, each of which has about 200 aesthetic ratings ranging on a one-to-ten scale. For fair comparison, we follow the experimental practices and train/test split used in literature [2, 3, 15] which results in about 230,000 training and 20,000 test images. When fine-tuning AlexNet for binary aesthetics classification, we divide the training set into two categories (low- and high-aesthetic category), with a score threshold of 5 as used in [2, 3, 15]. We use the subset of images which contain style attributes and content tags for training and testing the attribute-adaptive and content-adaptive branches.

5.2 Performance Evaluation

To evaluate the aesthetic scores predicted by our model, we report the ranking correlation measured by Spearman’s ρ between the estimated aesthetics scores and the ground-truth scores in the test set [30]. Let r_i indicate the rank of the i th item when we sort the list by scores $\{y_i\}$ and \hat{r}_i indicate the rank when ordered by $\{\hat{y}_i\}$. We can compute the disagreement in the two rankings of a particular element i as $d_i = r_i - \hat{r}_i$. The Spearman’s ρ rank correlation statistic is calculated as $\rho = 1 - \frac{6 \sum d_i^2}{N^3 - N}$, where N is the total number of images ranked. This correlation coefficient lies in the range of $[-1, 1]$, with larger values corresponding to higher correlation in the rankings. The ranking correlation is particularly useful since it is invariant to monotonic transformations of the aesthetic score predictions and hence avoids the need to precisely calibrate output scores against human ratings. For purposes of comparing to existing classification accuracy results reported on the AVA dataset, we simply threshold the estimated scores $[\hat{y}_i > \tau]$ to produce a binary prediction where the threshold τ is determined on the validation set.

5.3 Results

For comparison, we also train a model for binary aesthetics classification by fine-tuning AlexNet (AlexNet_FT_Conf). This has previously been shown to be a strong baseline for aesthetic classification [2]. We use the softmax confidence score corresponding of the high-aesthetics class as the predicted aesthetic rating. As described in Sect. 4, we consider variants of our architecture including the regression network alone (Reg), along with the addition of the pairwise ranking loss (Reg+Rank), attribute-constraint branches (Reg+Rank+Att) and *content-adaptive branches* (Reg+Rank+Cont). We also evaluate different pair-sampling strategies including within- and cross-rater sampling.

Model Architecture and Loss Functions: Tables 2 and 3 list the performance on AADB and AVA datasets, respectively. From these tables, we notice several interesting observations. First, AlexNet_FT_Conf model yields good ranking results measured by ρ . This indicates that the confidence score in softmax



Fig. 4. Example images from four content clusters found in the training set. These clusters capture thematic categories of image content present in AADB without requiring additional manual labeling of training data.

Table 2. Performance comparison of different models on AADB dataset.

Methods	ρ
AlexNet_FT_Conf	0.5923
Reg	0.6239
Reg + Rank (cross-rater)	0.6308
Reg + Rank (within-rater)	0.6450
Reg + Rank (within- & cross-)	0.6515
Reg + Rank + Att	0.6656
Reg + Rank + Cont	0.6737
Reg+Rank+Att+Cont	0.6782

Table 3. Performance comparison of different models on AVA dataset.

Methods	ρ	ACC (%)
Murray <i>et al.</i> [15]	-	68.00
SPP [31]	-	72.85
AlexNet_FT_Conf	0.4807	71.52
DCNN [3]	-	73.25
RDCNN [3]	-	74.46
RDCNN_semantic [19]	-	75.42
DMA [2]	-	74.46
DMA_AlexNet_FT [2]	-	75.41
Reg	0.4995	72.04
Reg + Rank	0.5126	71.50
Reg + Att	0.5331	75.32
Reg + Rank + Att	0.5445	75.48
Reg + Rank + Cont	0.5412	73.37
Reg+Rank+Att+Cont	0.5581	77.33

can provide information about relative rankings. Second, the regression net outperforms the AlexNet_FT_Conf model, and ranking loss further improves the ranking performance on both datasets. This shows the effectiveness of our ranking loss which considers relative aesthetics ranking of image pairs in training the model. More specifically, we can see from Table 2 that, by sampling image pairs according to the the averaged ground-truth scores, *i.e.* cross-rater sampling only, Reg + Rank (cross-rater) achieves the ranking coefficient $\rho = 0.6308$; whereas if only sampling image pairs within each raters, we have $\rho = 0.6450$ by Reg + Rank (within-rater). This demonstrates the effectiveness of sampling image pairs within the same raters, and validates our idea that the same individual has consistent aesthetics ratings. When using both strategies to sample image pairs, the performance is even better by Reg + Rank (within- & cross-), leading to $\rho = 0.6515$. This is possibly due to richer information contained in more training pairs. By comparing the results in Table 3 between “Reg” (0.4995) and “Reg+Rank” (0.5126), and between “Reg+Att” (0.5331) and “Reg+Rank+Att” (0.5445), we clearly observe that the ranking loss improves the ranking correlation. In this case, we can only exploit the cross-rater sampling strategy since rater’s identities are not available in AVA for the stronger within-rater sampling approach. We note that for values of ρ near 0.5 computed over 20000 test images on AVA dataset, differences in rank correlation of 0.01 are highly statistically significant. These results clearly show that the ranking loss helps enforce overall ranking consistency.

To show that improved performance is due to the side information (*e.g.* attributes) other than a wider architecture, we first train an ensemble of eight rating networks (Reg) and average the results, leading to a $\rho = 0.5336$ (*c.f.* Reg + Rank + Att which yields $\rho = 0.5445$). Second, we try directly training

Table 4. Ranking performance ρ vs. rank loss weighting ω_r in Eq. 2.

ω_r	0.0	0.1	1	2
AADB	0.6382	0.6442	0.6515	0.6276
AVA	0.4995	0.5126	0.4988	0.4672

Table 5. Ranking performance (ρ) of “Reg + Rank” with different numbers of sampled image pairs on AADB dataset.

#ImgPairs	2 million	5 million
cross-rater	0.6346	0.6286
within-rater	0.6450	0.6448
within- & cross-rater	0.6487	0.6515

the model with a single Euclidean loss using a wider intermediate layer with eight times more parameters. In this case we observed severe overfitting. This suggests for now that the side-supervision is necessary to effectively train such an architecture.

Third, when comparing Reg + Rank with Reg + Rank + Att, and Reg + Rank with Reg + Rank + Cont, we can see that both attributes and content further improve ranking performance. While image content is not annotated on the AADB dataset, our content-adaptive model based on unsupervised K -means clustering still outperforms the model trained without content information. The performance benefit of adding attributes is substantially larger for AVA than AADB. We expect this is due to (1) differences in the definitions of attributes between the two datasets, and (2) the within-rater sampling for AADB, which already provides a significant boost making further improvement using attributes more difficult. The model trained with ranking loss, attribute-constraint and content-adaptive branches naturally performs the best among all models. It is worth noting that, although we focus on aesthetics ranking during training, we also achieve the state-of-the-art binary classification accuracy in AVA. This further validates our emphasis on relative ranking, showing that learning to rank photo aesthetics can naturally lead to good classification performance.

Model Hyperparameters: In training our content-adaptive model on the AADB dataset which lacks supervised content labels, the choice of cluster number is an important parameter. Figure 5 plots the ρ on validation data as a function of the number of clusters K for the Reg + Cont model (without ranking loss). We can see the finer clustering improves performance as each content specific model can adapt to a sub-category of images. However, because the total dataset is fixed, performance eventually drops as the amount of training data available for tuning each individual content-adaptive branch decreases. We thus fixed $K = 10$ for training our unified network on AADB.

The relative weightings of the loss terms (specified by ω_r in Eq. 2) is another important parameter. Table 4 shows the ranking correlation test performance on both datasets w.r.t. different choices of ω_r . We observe that larger ω_r is favored in AADB than that in AVA, possibly due to the contribution from the within-rater image pair sampling strategy. We set ω_a (in Eq. 3) to 0.1 for jointly fine-tuning attribute regression and aesthetic rating. For the rank loss, we used validation performance to set the margin α to 0.15 and 0.02 on AVA and AADB respectively.

Number of Sampled Image Pairs: Is it possible that better performance can be obtained through more sampled pairs instead of leveraging rater’s information? To test this, we sample 2 and 5 million image pairs given the fixed training images on the AADB dataset, and report in Table 5 the performance of model “Reg + Rank” using different sampling strategies, *i.e.* within-rater only, cross-rater only and within-&cross-rater sampling. It should be noted the training image set remains the same, we just sample more pairs from them. We can see that adding more training pairs yields little differences in the final results, and even declines slightly when using higher cross-rater sampling rates. These results clearly emphasize the effectiveness of our proposed sampling strategy which (perhaps surprisingly) yields much bigger gains than simply increasing the number of training pairs by 2.5x.

Classification Benchmark Performance: Our model achieves state-of-the-art classification performance on the AVA dataset simply by thresholding the estimated score (Table 3). It is worth noting that our model uses only the whole warped down-sampled images for both training and testing, without using any high-resolution patches from original images. Considering the fact that the fine-grained information conveyed by high-resolution image patches is especially useful for image quality assessment and aesthetics analysis [2,3,14], it is quite promising to see our model performing so well. The best reported results [2] for models that use low resolution warped images for aesthetics classification are based on Spatial Pyramid Pooling Networks (SPP) [31] and achieves an accuracy of 72.85 %. Compared to SPP, our model achieves 77.33 %, a gain of 4.48 %, even though our model is not tuned for classification. Previous work [2,3,14] has shown that leveraging the high-resolution patches could lead to additional 5 % potential accuracy improvement. We expect a further accuracy boost would be possible by applying this strategy with our model.

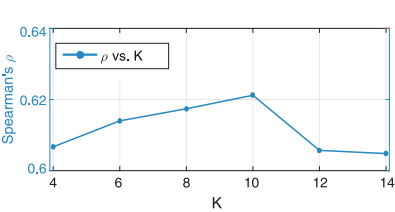


Fig. 5. Dependence of model performance by varying the number of content clusters. We select $K = 10$ clusters in our experiments on AADB.

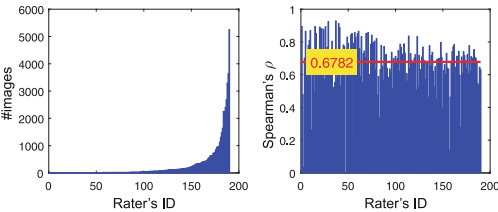


Fig. 6. Panels show (left) the number of images labeled by each worker, and the performance of each individual rater w.r.t Spearman’s ρ (Right). Red line shows our model’s performance. (Color figure online)

Table 6. Human perf. on the AADB dataset.

#Images	#Workers	ρ
>0	190	0.6738
>100	65	0.7013
>200	42	0.7112
Our best	–	0.6782

Table 7. Cross dataset train/test evaluation.

Spearman’s ρ		Test	
		AADB	AVA
Train	AADB	0.6782	0.1566
	AVA	0.3191	0.5154

5.4 Further Comparison with Human Rating Consistency

We have shown that our model achieves a high level of agreement with average aesthetic ratings and outperforms many existing models. The raters’ identities and ratings for the images in our AADB dataset enable us to further analyze agreement between our model each individual as well as intra-rater consistency. While human raters produce rankings which are similar with high statistical significance, as evaluated in Sect. 3, there is variance in the numerical ratings between them.

To this end, we calculate ranking correlation ρ between each individual’s ratings and the ground-truth average score. When comparing an individual to the ground-truth, we do not exclude that individual’s rating from the ground-truth average for the sake of comparable evaluations across all raters. Figure 6 shows the number of images each rater has rated and their corresponding performance with respect to other raters. Interestingly, we find that the hard workers tend to provide more consistent ratings. In Table 6, we summarize the individuals’ performance by choosing a subset raters based on the number of images they have rated. This clearly indicates that the different human raters annotate the images consistently, and when labeling more images, raters contribute more stable rankings of the aesthetic scores.

Interestingly, from Table 6, we can see that our model actually performs above the level of human consistency (as measured by ρ) averaged across all workers. However, when concentrating on the “power raters” who annotate more images, we still see a gap between machine and human level performance in terms of rank correlation ρ .

5.5 Cross-Dataset Evaluation

As discussed in Sect. 3, AVA contains professional images downloaded from a community based rating website; while our AADB contains a much more balanced distribution of consumer photos and professional photos rated by AMT workers, so has better generalizability to wide range of real-world photos.

To quantify the differences between these datasets, we evaluate whether models trained on one dataset perform well on the other. Table 7 provides a comparison of the cross-dataset performance. Interestingly, we find the models trained on either dataset have very limited “transferability”. We conjecture there are two

reasons. First, different groups of raters have different aesthetics tastes. This can be verified that, when looking at the DPChallenge website where images and ratings in the AVA dataset were taken from. DPChallenge provides a breakdown of scores which shows notable differences between the average scores among commenters, participants and non-participants. Second, the two datasets contain photos with different distributions of visual characteristics. For example, many AVA photos are professionally photographed or heavily edited; while AADB contains many daily photos from casual users. This observation motivates the need for further exploration into mechanisms for learning aesthetic scoring that is adapted to the tastes of specific user groups or photo collections [32].

6 Conclusion

We have proposed a CNN-based method that unifies photo style attributes and content information to rate image aesthetics. In training this architecture, we leverage individual aesthetic rankings which are provided by a novel dataset that includes aesthetic and attribute scores of multiple images by individual users. We have shown that our model is also effective on existing classification benchmarks for aesthetic judgement. Despite not using high-resolution image patches, the model achieves state-of-the-art classification performance on the AVA benchmark by simple thresholding. Comparison to individual raters suggests that our model performs as well as the “average” mechanical turk worker but still lags behind more consistent workers who label large batches of images. All these observations motivate us to study individual-oriented mechanism for aesthetics rating, as an example of subjective problem study.

References

1. Marchesotti, L., Murray, N., Perronnin, F.: Discovering beautiful attributes for aesthetic image analysis. *Int. J. Comput. Vis.* **2**, 1–21 (2014)
2. Lu, X., Lin, Z., Shen, X., Mech, R., Wang, J.Z.: Deep multi-patch aggregation network for image style, aesthetics, and quality estimation. In: ICCV (2015)
3. Lu, X., Lin, Z., Jin, H., Yang, J., Wang, J.Z.: Rapid: rating pictorial aesthetics using deep learning. In: Proceedings of the ACM International Conference on Multimedia, pp. 457–466. ACM (2014)
4. Luo, W., Wang, X., Tang, X.: Content-based photo quality assessment. In: 2011 IEEE International Conference on Computer Vision (ICCV), pp. 2206–2213. IEEE (2011)
5. Datta, R., Joshi, D., Li, J., Wang, J.Z.: Studying aesthetics in photographic images using a computational approach. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) ECCV 2006. LNCS, vol. 3953, pp. 288–301. Springer, Heidelberg (2006). doi:[10.1007/11744078_23](https://doi.org/10.1007/11744078_23)
6. Ke, Y., Tang, X., Jing, F.: The design of high-level features for photo quality assessment. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, vol. 1, pp. 419–426. IEEE (2006)

7. Dhar, S., Ordonez, V., Berg, T.L.: High level describable attributes for predicting aesthetics and interestingness. In: 2011 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1657–1664. IEEE (2011)
8. Nishiyama, M., Okabe, T., Sato, I., Sato, Y.: Aesthetic quality classification of photographs based on color harmony. In: 2011 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 33–40. IEEE (2011)
9. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vis.* **60**(2), 91–110 (2004)
10. Perronnin, F., Sánchez, J., Mensink, T.: Improving the fisher kernel for large-scale image classification. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) *ECCV 2010*. LNCS, vol. 6314, pp. 143–156. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-15561-1_11](https://doi.org/10.1007/978-3-642-15561-1_11)
11. Perronnin, F., Dance, C.: Fisher kernels on visual vocabularies for image categorization. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2007*, pp. 1–8. IEEE (2007)
12. Marchesotti, L., Perronnin, F., Larlus, D., Csúrká, G.: Assessing the aesthetic quality of photographs using generic image descriptors. In: *ICCV*, pp. 1784–1791. IEEE (2011)
13. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems*, pp. 1097–1105 (2012)
14. Kang, L., Ye, P., Li, Y., Doermann, D.: Convolutional neural networks for no-reference image quality assessment. In: 2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1733–1740. IEEE (2014)
15. Murray, N., Marchesotti, L., Perronnin, F.: Ava: a large-scale database for aesthetic visual analysis. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2408–2415. IEEE (2012)
16. Geng, B., Yang, L., Xu, C., Hua, X.S., Li, S.: The role of attractiveness in web image search. In: *Proceedings of the 19th ACM International Conference on Multimedia*, pp. 63–72. ACM (2011)
17. San Pedro, J., Yeh, T., Oliver, N.: Leveraging user comments for aesthetic aware image search reranking. In: *Proceedings of the 21st International Conference on World Wide Web*, pp. 439–448. ACM (2012)
18. Chopra, S., Hadsell, R., LeCun, Y.: Learning a similarity metric discriminatively, with application to face verification. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2005*, vol. 1, pp. 539–546. IEEE (2005)
19. Lu, X., Lin, Z., Jin, H., Yang, J., Wang, J.: Rating pictorial aesthetics using deep learning. *IEEE Trans. Multimedia* **17**(11), 1 (2015)
20. Marchesotti, L., Perronnin, F., Meylan, F.: Learning beautiful (and ugly) attributes. In: *BMVC* (2013)
21. Murray, N., Marchesotti, L., Perronnin, F., Meylan, F.: Learning to rank images using semantic and aesthetic labels. In: *BMVC*, pp. 1–10 (2012)
22. Luo, Y., Tang, X.: Photo and video quality evaluation: focusing on the subject. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) *ECCV 2008*. LNCS, vol. 5304, pp. 386–399. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-88690-7_29](https://doi.org/10.1007/978-3-540-88690-7_29)
23. Benjamini, Y., Yekutieli, D.: The control of the false discovery rate in multiple testing under dependency. *Ann. Stat.* **29**, 1165–1188 (2001)
24. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition (2014). arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556)

25. Cui, J., Wen, F., Xiao, R., Tian, Y., Tang, X.: Easyalbum: an interactive photo annotation system based on face clustering and re-ranking. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 367–376. ACM (2007)
26. Wang, J., Song, Y., Leung, T., Rosenberg, C., Wang, J., Philbin, J., Chen, B., Wu, Y.: Learning fine-grained image similarity with deep ranking. In: 2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1386–1393. IEEE (2014)
27. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: a unified embedding for face recognition and clustering. In: CVPR, pp. 815–823 (2015)
28. Lee, C., Xie, S., Gallagher, P.W., Zhang, Z., Tu, Z.: Deeply-supervised nets. In: AISTATS (2015)
29. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: convolutional architecture for fast feature embedding (2014). arXiv preprint [arXiv:1408.5093](https://arxiv.org/abs/1408.5093)
30. Myers, J.L., Well, A., Lorch, R.F.: Research Design and Statistical Analysis. Routledge, New York (2010)
31. He, K., Zhang, X., Ren, S., Sun, J.: Spatial pyramid pooling in deep convolutional networks for visual recognition. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014, Part III. LNCS, vol. 8691, pp. 346–361. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10578-9_23](https://doi.org/10.1007/978-3-319-10578-9_23)
32. Caicedo, J.C., Kapoor, A., Kang, S.B.: Collaborative personalization of image enhancement. In: 2011 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 249–256. IEEE (2011)