

Residual Attention Network vs Real Attention on Aesthetic Assessment

Author

Mandal, R, Becken, S, Connolly, RM, Stantic, B

Published

2021

Conference Title

Communications in Computer and Information Science

Version

Accepted Manuscript (AM)

DOI

[10.1007/978-981-16-1685-3_26](https://doi.org/10.1007/978-981-16-1685-3_26)

Rights statement

© Springer Nature Singapore Pte Ltd. 2021. This is the author-manuscript version of this paper. Reproduced in accordance with the copyright policy of the publisher. The original publication is available at www.springerlink.com

Downloaded from

<http://hdl.handle.net/10072/404292>

Griffith Research Online

<https://research-repository.griffith.edu.au>

Residual Attention Network vs Real Attention on Aesthetic Assessment

Ranju Mandal¹, Susanne Becken², Rod M. Connolly³, and Bela Stantic¹

¹ School of Information and Communication Technology, Griffith University,
Brisbane, Australia

² Australian Rivers Institute, Griffith University, Brisbane, Australia

³ Griffith Institute for Tourism, Brisbane, Australia

{r.mandal,s.becken,r.connolly,b.stantic}@griffith.edu.au

Abstract. Photo aesthetics assessment is a challenging problem. Deep Convolutional Neural Network (CNN)-based algorithms have achieved promising results for aesthetics assessment in recent times. Lately, few efficient and effective attention-based CNN architectures are proposed that improve learning efficiency by adaptively adjusts the weight of each patch during the training process. In this paper, we investigate how real human attention affects instead of CNN-based synthetic attention network architecture in image aesthetic assessment. A dataset consists of a large number of images along with eye-tracking information has been developed using an eye-tracking device⁴ power by sensor technology for our research, and it will be the first study of its kind in image aesthetic assessment. We adopted a Residual Attention Network and ResNet architectures which achieve state-of-the-art performance image recognition tasks on benchmark datasets. We report our findings on photo aesthetics assessment with two sets of datasets consist of original images and images with masked attention patches, which demonstrates higher accuracy when compared to the state-of-the-art methods.

Keywords: Photo Aesthetic Assessment · Image Aesthetic Evaluation · Great Barrier Reef · Aesthetic Scoring · Deep Learning

1 Introduction

Image quality assessment and predict photo aesthetic values have been a challenging problem in image processing and computer vision, as aesthetic assessment is subjective (i.e. influenced by individual's feelings, tastes, or opinions) in nature. A significant number of existing photo aesthetics assessment methods are available ([7, 13, 16, 20, 21, 30, 33]) in the literature using extraction of visual features and then employ various machine learning algorithms to predict photo aesthetic values. Aesthetic assessment techniques aim to quantify semantic level characteristics associated with emotions and beauty in images, whereas technical quality assessment deals with measuring low-level degradations such as noise, blur, compression artifacts, etc.

⁴ <https://www.tobii.com/group/about/this-is-eye-tracking/>

Recently, deep learning methods have shown great success in various computer vision tasks, such as object recognition [28, 27, 17], object detection [9, 26], and image classification [32]. Deep learning methods, such as deep convolutional neural network and deep belief network, have also been applied to photo quality/aesthetics assessment and have shown good results. As most deep neural network architectures require fixed-size inputs, recent methods transform input images via cropping, scaling, and padding, and design dedicated deep network architectures, such as double-column or multi-column networks, to simultaneously take multiple transformed versions as input.

Based on the available image assessment techniques in the literature, full-reference and no-reference approaches are the two main categories of image quality assessment. While the availability of a reference image is assumed in the former (metrics such as PSNR, SSIM [35], etc.), typically blind (no-reference) approaches rely on a statistical model of distortions to predict image quality. A quality score is to predict that relates well with human perception is the main goal in both cases.

Broadly, the task involved to distinguish computationally the aesthetic attributes of an image [10] for a related assumption. The literature proposes several methods to solve such challenging classification and scoring problems. The earlier approaches can be categorised into two groups, based on visual feature types (hand-crafted features and deep features based on Convolutional Neural Network), and evaluation criteria, dataset characteristics and evaluation metrics (examples include: Precision-recall, Euclidean distance, ROC curve, and mean Average Precision). More specifically, the term "hand-crafted" features refer to properties derived employing various algorithms using the information present in an image. As an example, edges and corners are two simple features that can be extracted from images. A basic edge detector algorithm works by finding areas where the image intensity "suddenly" changes. For example, the shell of a turtle can be identified as an edge. Likewise, the so-called Histogram of Gradients (HoG) [5] is another type of handcrafted feature that can be applied in many different ways.

Earlier proposed techniques designed hand-crafted aesthetics features according to aesthetics perception of people and photography rules [2, 6, 7] and obtained encouraging results while handcrafted feature design for aesthetic assessment is a very challenging task. More robust feature extraction techniques were proposed later on to leverage more generic image features (e.g. Fisher Vector [18, 23, 24] and the bag of visual words [30]) for photo aesthetics evaluation. Generic feature-based representation of images is not ideal for image aesthetic assessment as those features are designed to represent natural images in general, and not specifically for aesthetics assessment.

In contrast, Convolutional Neural Network (CNN)-based features are learning from the training samples, and they do this by using dimensionality reduction and convolutional filters. Recent approaches to image aesthetic assessment mostly apply more complex and robust deep Convolutional Neural Networks (CNN) architectures ([37, 14, 4, 3]). Availability of large-scaled labeled and scored

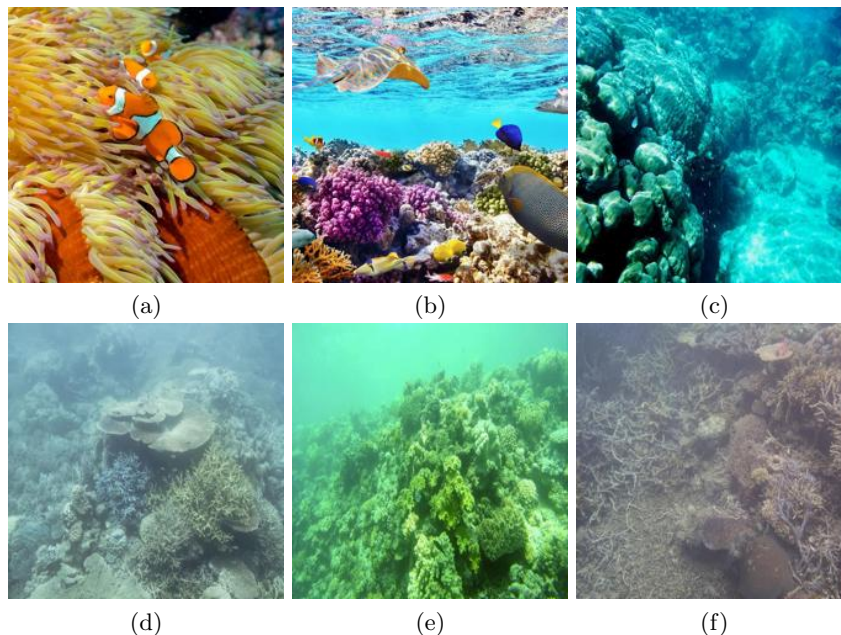


Fig. 1. Some example images from GBR dataset with score $\mu(\pm\sigma)$, where μ and σ represent Mean and Standard Deviation (SD) of score, respectively. (a) high aesthetics and low SD ($\mu = 9.583$, $SD = 0.64$), (b) high aesthetics and low SD ($\mu = 9.462$, $SD = 0.746$), (c) high aesthetics and high SD ($\mu = 5.916$, $SD = 3.773$), (d) low aesthetics and low SD ($\mu = 2.273$, $SD = 1.42$), (e) low aesthetics and low SD ($\mu = 3.0$, $SD = 1.0$), (f) low aesthetics and high SD ($\mu = 3.454$, $SD = 3.23$)

images from online repositories have enabled CNN-based methods to perform better than previously proposed non-CNN approaches [19, 25]. Moreover, having access to pre-trained models (e.g. ImageNet [27]) for network training initialization and fine-tuned the network on subject data of image aesthetic assessment have been proven more effective technique for typical deep CNN approach [34].

2 Methodology

We propose two novel Deep CNN architecture for image aesthetics assessment adapted from the recently published state-of-the-art image classification model ([11, 34]). Both the models used in our experiments have been well tested as image classifier for a large number of classes. In experiments, our aim for predictions with higher correlation with human ratings, instead of classifying images to low/high score or mean score regression, the distribution of ratings are predicted as a histogram [8]. The squared EMD (Earth Mover’s Distance) loss-based assessment was proposed by Talebi et al. [8], which shows a performance boost inaccurate prediction of the mean score. All network models for aesthetic assessment

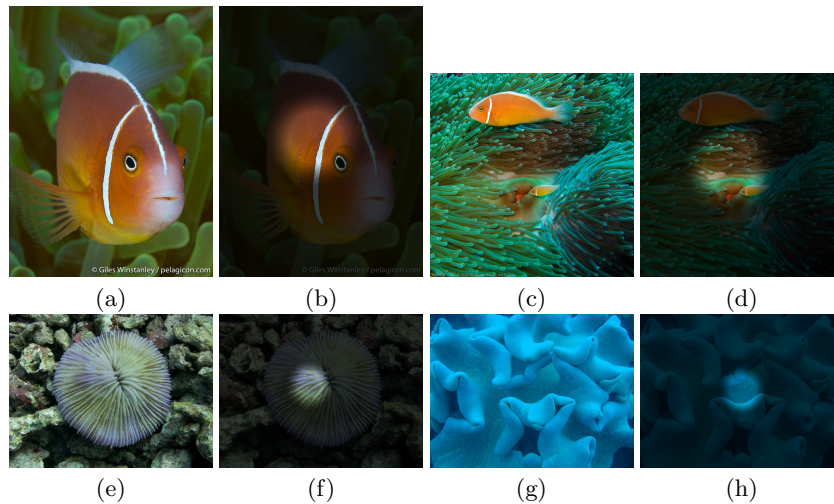


Fig. 2. Four example images from GBR dataset along with same processed images after studying eye movements. (a), (c), (e), (g) are original images. (b), (d),(f), (h) are masked images of (a), (c), (e), (g) respectively, after masked using attention patch information obtained from the Tobii eye-tracking device.

are based on image classifier architectures. Two different architectures (with and without attention mechanism) state-of-the-art networks are explored for the proposed applications. Networks used in our experiments were first trained and then fine-tuned using the large-scale aesthetics assessment AVA dataset [19]. The AVA dataset has 250,000 images, which is very useful for training such a large deep neural network model. The complete architecture of this project consists of different sub-modules, and each of these sub-modules consists of building blocks, such as pooling, filters, activation functions, and so forth. The following sections provide more information on the sub-modules. A more detailed description of the architecture and different modules is provided below.

ResNet: Theoretically, neural networks should get better results as added more layers. A deeper network can learn anything a shallower version of itself can, plus possibly some more parameters. The intuition behind adding more layers to a deep neural network was that more layers progressively learn more complex features. The first, second, third, layers learn features such as edges, shapes, objects, respectively, and so on. He et al. [11] empirically presented that there is a maximum threshold for depth with the traditional CNN model. As more layers are added, the network gets better results until at some point; then as continue add extra layers, the accuracy starts to drop. The reason behind failures of the very deep CNN was mostly related to optimization function, network weights initialization, or the well-known vanishing/exploding gradient problem. Vanishing gradients are especially blamed, however, He et al. [11] argue that the use of Batch Normalization ensures that the gradients have healthy norms. In contrast, deeper networks are harder to optimize due to add more difficulty in

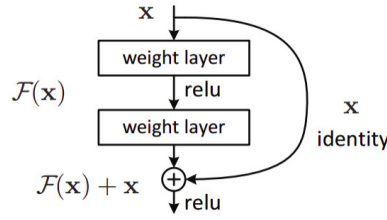


Fig. 3. Residual learning: a building block.

the process of training; it becomes harder for the optimization to find the right parameters.

The problem of training very deep networks has been attenuated with the introduction of a new neural network layer -The Residual Block 3. Residual Networks attempt to solve this issue by adding the so-called skip connections. A skip connection is depicted in Fig 3. If, for a given dataset, there are no more things a network can learn by adding more layers to it, then it can just learn the identity mapping for those additional layers. In this way, it preserves the information in the previous layers and can not do worse than shallower ones. So, the most important contribution to ResNet architecture is the ‘Skip Connection’ identity mapping.

This identity mapping does not have any parameters and is just there to add the output from the previous layer to the layer ahead. However, sometimes x and $F(x)$ will not have the same dimension. Recall that a convolution operation typically shrinks the spatial resolution of an image, e.g. a 3×3 convolution on a 32×32 image results in a 30×30 image. The identity mapping is multiplied by a linear projection W to expand the channels of shortcuts to match the residual. This allows for the input x and $F(x)$ to be combined as input to the next layer. A block with a skip connection as in the image above is called a residual block, and a Residual Neural Network (ResNet) is just a concatenation of such blocks.

$$y = F(x, W_i) + W_s x$$

The above equation shows when $F(x)$ and x have a different dimensionality such as 32×32 and 30×30 . This W_s term can be implemented with 1×1 convolutions, this introduces additional parameters to the model. The Skip Connections between layers add the outputs from previous layers to the outputs of stacked layers. This results in the ability to train much deeper networks than what was previously possible. He et al. [11] proposed their network with 100 and 1,000 layers and tested on bench mark datasets such as CIFAR-10, ImageNet dataset with 152 layers and achieved state-of-the-art performance.

Residual Attention Network: Residual Attention Network (Fig 4), a convolutional neural network that incorporates both attention mechanism and residual units which can incorporate with state-of-art feed forward network architecture in an end-to-end training fashion. Residual Attention Network is con-

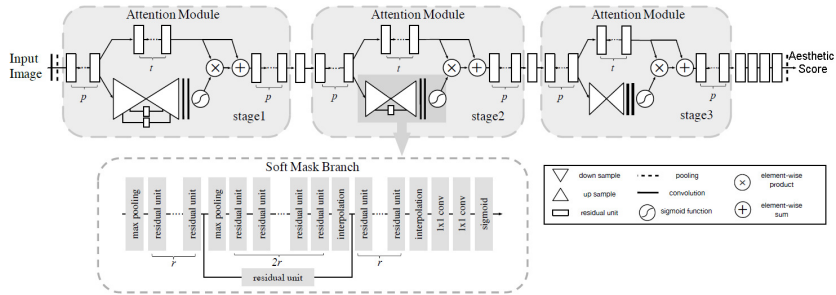


Fig. 4. An architecture the Residual Attention Network Architecture [34]. The output layer is modified to produce image aesthetic score instead class label.

structured by stacking multiple Attention Modules which generate attention-aware features. Residual unit is a basic component that utilizes skip-connections to jump over few layers with nonlinearities and batch normalizations which is the prominent feature is the attention module.

Each Attention Module (see Fig. 4) is divided into two branches: mask branch and trunk branch. The trunk branch performs feature processing with Residual Units and can be adapted to any state-of-the-art network structures. Mask Branch uses bottom-up top-down structure softly weight output features with the goal of improving trunk branch features. The Bottom-Up step collects global information of the whole image by downsampling (i.e. max pooling) the image. The Top-Down step combines global information with original feature maps by upsampling (i.e. interpolation) to keep the output size the same as the input feature map. Inside each Attention Module, bottom-up top-down feedforward structure is used to unfold the feedforward and feedback attention process into a single feedforward process. The attention-aware features from different modules change adaptively as layers going deeper. Importantly, an attention residual learning to train very deep Residual Attention Networks which can be easily scaled up to hundreds of layers. The experiment also demonstrates that Residual Attention Network is robust against noisy labels.

In Residual Attention Network , a pre-activation Residual Unit [12], ResNeXt [36] and Inception [31] are used as basic unit to construct Attention Module. Given trunk branch output $T(x)$ with input x , the mask branch uses bottom-up top-down structure [15, 22, 1] to learn same size mask $M(x)$ that softly weight output features $T(x)$. The bottom-up top-down structure mimics the fast feedforward and feedback attention process. The output mask is used as control gates for neurons of trunk branch similar to Highway Network [29]. The output of Attention Module H is:

$$H_i, c(x) = M_i, c(x) \times T_i, c(x)$$

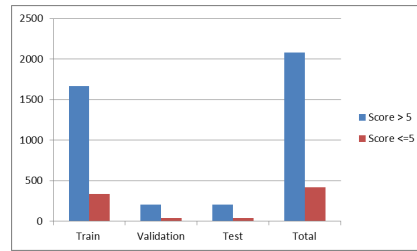


Fig. 5. Score distribution of the GBR dataset developed in-house, containing a total of 5417 images.

3 Results and Discussions

Dataset: For experimental purposes, both publicly available datasets and dataset developed in-house were used. For the dataset specific to the Great Barrier Reef (i.e. the GBR dataset), we used 5,417 underwater GBR images, which were downloaded from the Flickr social media platform. These images were sorted based on the content and then rated by participants in an online survey for their aesthetic beauty. At least 10 survey participants provided an aesthetic score for each image and the mean score was calculated. Most of the images (i.e. 80%) served as training material to enable the proposed Neural Network model to learn key feature parameters. The remaining 20% of the data were used during the test and validation phases. The validation dataset helps to understand the system performance in terms of accuracy during the training phase, whereas the test set is normally used once the training phase is completed and ready for deployment. To better understand the distribution of ‘beautiful’ and ‘ugly’ pictures in the dataset, Figure 5 presents the number of images with scores above or equal/below 5. More than 2,080 images were scored as highly aesthetic ($score > 5$) and only 420 ($score \leq 5$) images were scored as having low aesthetics. Figure 5 also shows how many images of high and low scores were used in each experimental stage.

The GBR dataset of 5,471 images is comparatively small for training a multi-layered deep Convolutional Neural Network. It was, therefore, necessary to complement the GBR data with a large-size, publicly available dataset (AVA, see Murry et al., 2012). This helped to train the system and allowed us to use the in-house GBR dataset for fine-tuning the algorithm. The detailed dataset description of the AVA [19] is described below.

- AVA1: We adopted the score of 0.5 (mean aesthetic score ranges between 0 and 1) as the threshold value to divide the dataset into high aesthetic value and low aesthetic value classes. By doing this, we obtained 74,056 images in the low aesthetic value class and 181,447 images in the high aesthetic value class. 229,954 and 25,549 (approximately 10% images) were used for testing system performance.

- AVA2: In a different experimental setup, and to increase the gap between images with high aesthetic and low aesthetic value, all images were sorted based on their mean scores. Then, the top 10% of images were considered as highly aesthetic and the bottom 10% images were classed as low aesthetic. Thus, 51,100 images (approximately, 20% of the full dataset) then formed the AVA dataset that was used for training.

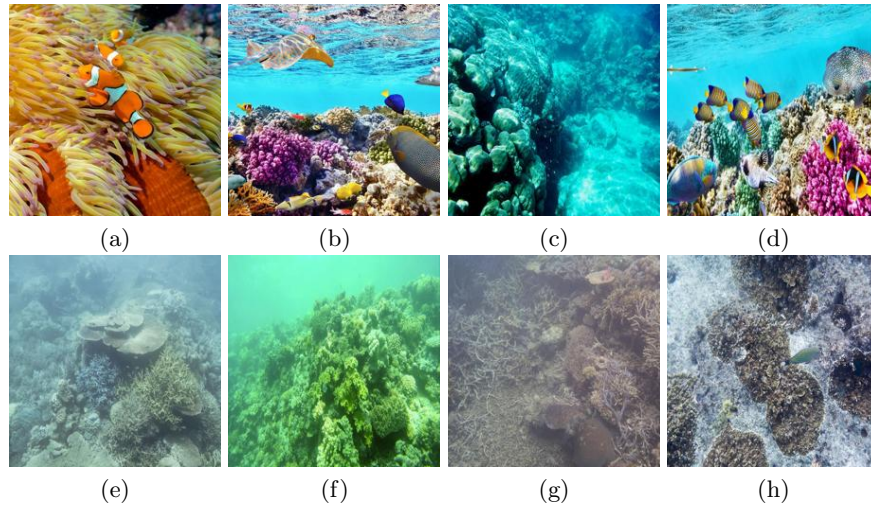


Fig. 6. Ranking some sampled images from GBR dataset with high and low aesthetic values using our proposed aesthetic assessment model using NIMA on Xception. Predicted (and ground truth) scores are shown together for every images (a) 9.583 (9.43), (b) 9.462 (9.22), (c) 5.916 (5.09), (d) 5.916 (5.60), (e) 2.273 (2.55), (f) 3.0 (3.5), (g) 3.454 (3.23), (h) 4.1 (4.22)

4 Conclusion

The objective of the work was to study the effectiveness of real human attention obtained using an eye-tracking device on deep Convolutional Neural Network architecture for automatic image aesthetic assessment and predicting an aesthetic score for images. The significance of aesthetic evaluation system was studied in details from the literature. A state-of-the-art deep Residual Attention Network architecture and ResNet were adapted to our need for the modeling of our GBR aesthetic assessment task. A wide range of experiments was conducted using a comprehensive analysis of performance is presented. Findings on photo aesthetics assessment of original images and images with masked attention patches demonstrates the higher accuracy of proposed method when compared to the state-of-the-art methods.

Table 1. Performance of the proposed method and few recently published architectures in predicting images from GBR dataset. Accuracy values are based on classification of photos to two classes high and low aesthetics(column 2). LCC (linear correlation coefficient) and are computed between predicted and ground truth mean scores (column 3) and standard deviation of scores (column 5). EMD measures closeness of the predicted and ground truth rating distributions.

Model	Accuracy	LCC-Mean	LCC-std.dev	EMD
NIMA(MobileNet)	80.36%	0.518	0.152	0.081
NIMA(VGG16)	80.60%	0.610	0.205	0.052
NIMA (Inception-Resnet)	81.51%	0.636	0.233	0.050
Residual Attention Network	81.74%	0.641	0.211	0.045

Acknowledgement

This research was partly funded by the National Environment Science Program (NESP) Tropical Water Quality Hub Project No 5.5.

References

1. Badrinarayanan, V., Handa, A., Cipolla, R.: Segnet: A deep convolutional encoder-decoder architecture for robust semantic pixel-wise labelling. CoRR **abs/1505.07293** (2015), <http://arxiv.org/abs/1505.07293>
2. Bhattacharya, S., Sukthankar, R., Shah, M.: A framework for photo-quality assessment and enhancement based on visual aesthetics. In: In Proc. ACM International Conference on Multimedia. pp. 271–280 (2010)
3. Bianco, S., Celona, L., Napoletano, P., Schettini, R.: On the use of deep learning for blind image quality assessment. In: arXiv preprint arXiv:1602.05531 (2016)
4. Bosse, S., Maniry, D., T.Wiegand, Samek, W.: A deep neural network for image quality assessment. In: In Proc. Computer Vision and Pattern Recognition. pp. 3773–3777 (2016)
5. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: In Proc. Computer Vision and Pattern Recognition (CVPR’05). vol. 1, pp. 886–893 (2005)
6. Datta, R., Joshi, D., Li, J., Wang, J.Z.: Studying aesthetics in photographic images using a computational approach. In: In Proc. European Conference on Computer Vision. pp. 288–301 (2006)
7. Dhar, S., Ordonez, V., Berg, T.: High level describable attributes for predicting aesthetics and interestingness. In: In Proc. Computer Vision and Pattern Recognition. pp. 1657–1664 (2011)
8. Esfandarani, H.T., Milanfar, P.: NIMA: neural image assessment. CoRR **abs/1709.05424** (2017), <http://arxiv.org/abs/1709.05424>
9. Girshick, R.B.: Fast r-cnn. In: ICCV (2015)

10. Haas, A., Guibert, M., Foerschner, A., Co, T., Calhoun, S., George, E., Hatay, M., Dinsdale, E., Sandin, S.A., Smith, J.E., Vermeij, M., Felts, B., Dustan, P., Salamon, P., Rohwer, F.: Can we measure beauty? computational evaluation of coral reef aesthetics. *PeerJ* 3:e1390 <https://doi.org/10.7717/peerj.1390> (2015)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016)
12. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. *CoRR* **abs/1603.05027** (2016), <http://arxiv.org/abs/1603.05027>
13. Jiang, W., Loui, A., Cerosaletti, C.: Automatic aesthetic value assessment in photographic images. In: In Proc. International Conference on Multimedia and Expo (ICME). pp. 920–925 (2010)
14. Kang, L., Ye, P., Li, Y., Doermann, D.: Convolutional neural networks for no-reference image quality assessment. In: In Proc. Computer Vision and Pattern Recognition. pp. 1733–1740 (2014)
15. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. *CoRR* **abs/1411.4038** (2014), <http://arxiv.org/abs/1411.4038>
16. Lu, X., Lin, Z., Jin, H., Yang, J., Wang, J.Z.: Rapid: Rating pictorial aesthetics using deep learning. In: In Proc. ACM International Conference on Multimedia. pp. 457–466 (2014)
17. Mandal, R., Connolly, R.M., Schlacher, T.A., Stantic, B.: Assessing fish abundance from underwater video using deep neural networks. In: 2018 International Joint Conference on Neural Networks, IJCNN 2018, Rio de Janeiro, Brazil, July 8–13, 2018. pp. 1–6. IEEE (2018), <https://doi.org/10.1109/IJCNN.2018.8489482>
18. Marchesotti, L., Perronnin, F., Larlus, D., Csurka, G.: Assessing the aesthetic quality of photographs using generic image descriptors. In: In Proc. International Conference on Computer Vision. pp. 1784–1791 (2011)
19. Murray, N., Marchesotti, L., Perronnin, F.: Ava: A large-scale database for aesthetic visual analysis. In: In Proc. Computer Vision and Pattern Recognition. pp. 2408–2415 (2012)
20. Nishiyama, M., Okabe, T., Sato, I., Sato, Y.: Aesthetic quality classification of photographs based on color harmony. In: In Proc. Computer Vision and Pattern Recognition. pp. 33–40 (2011)
21. Niu, Y., Liu, F.: What makes a professional video? a computational aesthetics approach. *IEEE Transactions on Circuits and Systems for Video Technology* **22(7)**, 1037–1049 (2012)
22. Noh, H., Hong, S., Han, B.: Learning deconvolution network for semantic segmentation. *CoRR* **abs/1505.04366** (2015), <http://arxiv.org/abs/1505.04366>
23. Perronnin, F., Dance, C.: Fisher kernels on visual vocabularies for image categorization. In: In Proc. Computer Vision and Pattern Recognition. pp. 1–8 (2007)
24. Perronnin, F., Sanchez, J., Mensink, T.: Improving the fisher kernel for large-scale image classification. In: In Proc. European Conference on Computer Vision. pp. 143–156 (2010)
25. Ponomarenko, N., Ieremeiev, O., Lukin, V., Jin, L., Egiazarian, K., Astola, J., Vozel, B., Chehdi, K., Carli, M., Battisti, F., Kuo, C.C.: A new color image database tid2013: Innovations and results. In: 15th International Conference on Advanced Concepts for Intelligent Vision Systems - Volume 8192. pp. 402–413 (2013)
26. Ren, S., He, K., Girshick, R.B., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. In: NIPS (2015)

27. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. *IJCV* **115**(3), 211–252 (2015). <https://doi.org/10.1007/s11263-015-0816-y>
28. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *CoRR* **abs/1409.1556** (2014)
29. Srivastava, R.K., Greff, K., Schmidhuber, J.: Training very deep networks. *CoRR* **abs/1507.06228** (2015), <http://arxiv.org/abs/1507.06228>
30. Su, H.H., Chen, T.W., Kao, C.C., Hsu, W.H., Chien., S.Y.: Scenic photo quality assessment with bag of aesthetics-preserving features. In: In Proc. ACM International Conference on Multimedia. pp. 1213–1216 (2011)
31. Szegedy, C., Ioffe, S., Vanhoucke, V.: Inception-v4, inception-resnet and the impact of residual connections on learning. *CoRR* **abs/1602.07261** (2016), <http://arxiv.org/abs/1602.07261>
32. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: *Computer Vision and Pattern Recognition (CVPR)* (2015), <http://arxiv.org/abs/1409.4842>
33. Tang, X., Luo, W., X.Wang: Content-based photo quality assessment. *IEEE Transactions on Multimedia* **15**(8), 1930–1943 (2013)
34. Wang, F., Jiang, M., Qian, C., Yang, S., Li, C., Zhang, H., Wang, X., Tang, X.: Residual attention network for image classification. In: In Proc. *Computer Vision and Pattern Recognition (CVPR)*. pp. 6450–6458 (2017)
35. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Multimedia* **13**(4), 600–612 (2004)
36. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: In Proc. *Computer Vision and Pattern Recognition (CVPR)*. pp. 5987–5995 (2017)
37. Xue, W., Zhang, L., Mou, X.: Learning without human scores for blind image quality assessment. In: In Proc. *Computer Vision and Pattern Recognition*. pp. 995–1002 (2013)