Automatic Image Transformation for Inducing Affect

Afsheen Rafaqat Ali
afsheen.ali@itu.edu.pk

Mohsen Ali
mohsen.ali@itu.edu.pk

Intelligent Machines Lab
Information Technology University of
Punjab (ITU)
Lahore, Pakistan

Abstract

Current image transformation and recoloring algorithms try to introduce artistic effects in the photographed images, based on user input of target image(s) or selection of pre-designed filters. These manipulations, although intended to enhance the impact of an image on the viewer, do not include the option of image transformation by specifying the affect information. In this paper we present an automatic image-transformation method that transforms the source image such that it can induce an emotional affect on the viewer, as desired by the user. Our proposed novel image emotion transfer algorithm does not require a user-specified target image. The proposed algorithm uses features extracted from top layers of deep convolutional neural network and the user-specified emotion distribution to select multiple target images from an image database for color transformation, such that the resultant image has desired emotional impact. Our method can handle more diverse set of photographs than the previous methods. We conducted a detailed user study showing the effectiveness of our proposed method. A discussion and reasoning of failure cases has also been provided, indicating inherent limitation of color-transfer based methods in the use of emotion assignment.

1 Introduction

Over the last decade we have seen a dramatic increase in both the creation of multimedia content and its mass sharing. Social networks and content sharing websites make it easy to share photographs, consume the ones posted by others and to solicit the input from the viewers. Photographers use image enhancement techniques to either overcome the limitation of their sensors or to induce an artistic effect. Many existing applications allow easy manipulation of an image to garner more positive feedback from the viewers or enhance the context for sharing of personal experience. Most of the currently available tools restrict themselves to a limited range of low-level filters (including smoothing and sharpening filters) and manipulations (e.g. giving vintage look to image, converting to black & white or enhancing colors of landscape image). These manipulations are restricted to low level changes like adjusting hue, saturation and brightness, with user given minimum control over these parameters. Even the recently popular selfie or face filters introduced in commercial image sharing applications are limited to few designed alternations that could be applied to facial features.

One of the important factors to judge the value of any multimedia content is it's emotional impact on the viewer. Images that evoke negative emotions tend to be more memorable [9]

Figure 1: Given source image and target emotion distribution, our system recolors the image by automatically selecting target images (from Emotion6 dataset [10]) that are semantically similar to source image and their emotion distribution matches provided target emotion distribution. **Results are best visible in color print.**

and online advertisement videos inducing surprise and joyous emotions have been shown to retain more viewers’ concentration and attention [10]. Unfortunately, most of the popular photo and video editing tools do not allow users to transform an image by indicating the emotional impact on it’s viewers, as desired by the photographer. Fig. (1) shows few of the possible examples of such transformations.

In this paper, we present a framework for image transformation such that the transformed source image has the affect, as desired by the user. Unlike previous methods [4, 10, 11], in our algorithm user has to input only the source image and seven dimensional discrete probability distribution representing ratio of seven emotions (anger, disgust, fear, joy, sadness, surprise and neutral). We use features extracted from top layers of deep convolutional networks (that encapsulate both in-image-context and content information) to select target images that have similar content and spatial context as the source image. In addition to that, we also make sure that emotion distribution of selected target images matches with the desired emotion distribution. To perform emotion transformation, low-level color features are extracted from these selected images and weighted combination of these features is then applied to the provided source image. Fig. (1) shows transformed images generated by our algorithm alongwith the automatically selected target images. The results of the performed user study shows that transformed images are more close to the desired emotion distributions (shown below target images) as compared to the corresponding source images.

Learning to generate a desired affect, either to change an image or to enhance it, will empower a photographer to manipulate the image according to the message he/she wants to convey or to align it with the content. For example, same photograph of scenic landscape can be manipulated to have gloomy feeling or joyous feeling of spring. Similarly, our method could be used to edit the multimedia content for the VR and AR environments.

2 Related Work

An important factor to judge the value of any multimedia content is the emotional impact it evokes. Music, lighting/illumination, camera pose and color etc. have been used by movie directors to enhance the desired mood of the scene and its emotional impact on the viewers. For example, dark and gloomy colors are used to induce depression or fear, warm colors are used for positive emotions, fast moving camera is used to capture energy in the scene and tempo of the music is used to grab the attention of the viewer(s). Photographers on the other hand don't have all these tools and are restricted to manipulation of colors and textures only once the photograph has been taken.

2.1 Colorization and Color/Style Transfer

Image color transformation has been employed for variety of tasks, from colorization of the gray scale images to color de-rendering and image transfer. Where colorization of gray-scale images [14] and color de-rendering [14] use learned models for the task (requiring only the input image); most of the image transfer algorithms require both source and target images as input. Pouli and Reinhard [13] use the histogram matching algorithm to transfer the color and tone from the target image to the source image. [5] recolorizes the image such that style of the target image is transferred to the input image. Recently Gatys et al. [6] used deep convolutional network trained for object detection to transfer style of the target image to the source image. They claim that their algorithm understand how to separate the content and style because they use appropriate combination of high level and low-level features. However, their transformation is style transfer at the cost of realism.

2.2 Image Emotion Assignment

Affective computing and emotion assignment to the image has been a popular topic in the last decade in the multimedia and computer graphics community. For brevity, we describe few of the proposed techniques in this section. Works like [9, 10] have tried to represent emotion by a concept and assigned each concept a color pallet designed on the basis of color emotion literature. On the other hand, example based approaches require a target image, representing desired affect, to be provided [11]. First technique described in [9, 10] suffers from the problem of modeling human perception that restricts it to few concepts and [11] puts the burden on the user to provide an appropriate target image. Initially it looks an appropriate constraint to place on this under-constrained problem. However, finding an image which is both spatially similar to the input image and evokes emotion as user desires is difficult, ambiguity ridden and over-burdens the user.

Recent works have explored learning from data to understand the relationship between the images and different models of affect [9, 11]. Ali et al. [10] have studied how high level concepts present in the images are related to the affect they induce. [9, 11, 12] have used the datasets to learn and/or transfer the emotion to the input image. Xue et al. [12] model low-level color and tone features to represent different emotions of the movie clip clustered on the basis of genre and director of the film. Xu et al. [11] trained Gaussian Mixture Model (GMM) on the super-pixels of the images that have been clustered into emotion-scene subgroups. The input image is first matched to a scene-subgroup in the desired emotion and then each super-pixel is transformed by minimizing energy equation that tries to find mapping from learned GMM to the super-pixels of input image. The above method uses

low-level features and cluster their data into discrete classes of emotions. These methods are not suitable for our case where we have distribution of emotion and emotion-scene clustering is not feasible.

Kim et al. [10] instead of over-segmenting the input image use semantic segmentation. For each input-image-segment, semantically compatible segments are searched from the database while minimizing on position, scale, lightness of the segment and closeness of user supplied Valance Arousal (VA) score to VA score of the image to whom selected segment belongs to. Due to their reliance on the semantic segmentation they constraint themselves to the landscape images, they also rely on VA scores which are not easily interpret-able by the humans and each segment’s color is changed by just changing the mean value of the segment.

For our experimental settings, we use the dataset introduced by Peng et al. [11] where discrete emotion distribution is associated with each image. We use features extracted from the top layer of Deep CNN that capture both the content and spatial context information including the color and texture information at different locations of the image. We are not constrained by the semantic segmentation of an image and allow much more control of desired emotion by providing user ability to pick any discrete distribution of the emotions. Emotion distribution remains human interpret-able with the freedom to choose variety of combination of emotions.

3 Dataset

Emotion6 The results reported in this paper use a well known affective image dataset named Emotion6 [12] that consists of 1980 images. Each image in Emotion6 dataset has an emotion distribution associated with it that indicates the probability of six Ekman’s basic emotions [13] and a neutral one being evoked in its viewers. The Fig. (2) shows an example image from Emotion6 dataset with its corresponding emotion distribution. The higher the probability of an emotion, there are more chances for particular emotion being evoked in its viewers.

Figure 2: An example image from Emotion6 with its corresponding Emotion Distribution.

ArtPhoto Few selective images of Artphoto dataset [14] are also used in our experiments. In ArtPhoto dataset, images are carefully taken by artists to induce specific emotion in their viewers by keenly selecting interpretable aesthetic features related to image composition, colors and lighting effects, etc.

4 Methodology

We propose a novel method that unlike previous methods (e.g. [15]) requires only a source image and a target emotion distribution for affective image transformation. By eliminating the need for the suitable target image, we free the user from the burden of searching it in order to obtain the transformed image that could elicit desired emotion distribution. Semantic

Figure 3: Block diagram of proposed Affect assignment algorithm. The AlexNet and GoogleNet network visualizations are created using [10].

information of source image is taken in consideration while performing emotion transformation to ensure that color assigned to the objects present in the image lies in the color space in which those objects naturally exist. This is also crucial for the naturalness of the transformed image, as blue sky or blue sea might induce pleasant feelings in its viewers but blue tree would look unnatural and therefore may not elicit pleasantness in the observers.

The block diagram of our emotion transfer algorithm is shown in Fig. (3) where a user is requested to input only a source image I_s and the desired emotion distribution P_d . Similar to [10], in our method color tone is adjusted using the algorithm proposed in [10] that relies on the color histogram of the target image to perform the transformation. However, we don't rely on any user input targeted image for the calculation of target color histogram.

Let, D be database constructed from image-emotion distribution pair. Currently, we are only using Emotion6 dataset as our database since it is the only dataset where each image I is associated with an emotion probability distribution $\mathbf{P} = \{p^1, p^2, \dots, p^7\}$ representing probability distribution of Ekman's six basic emotions and a neutral one.

Given an input image I_s and the desired emotion distribution P_d , we construct a histogram $\mathbf{h}$ as weighted summation of histograms of target images in the database D ; $\mathbf{h} = \sum_{i \in T} (w_i * \mathbf{h}_i)$

Figure 4: The number of images being transformed towards each emotion.

where $|T| = K$ and $T \subseteq D$. T is subset of K images such that $d = \sum_{I_i \in T} \text{distance}(P_d, P_i) + \text{distance}(I_i, I_s)$ is minimized. To minimize d we solve the problem in two steps. First we construct subset $D^* \subseteq D$ by selecting only those images $I_i \in D$ having emotion distribution similar to the input emotion distribution. Bhattacharyya coefficient [14] is used as similarity measure between emotion distributions

$$bc_i = \text{distance}(P_d, P_i) = \sum_{k=1}^7 \sqrt{P_d(k) \cdot P_i(k)} \quad (1)$$

To find the required target image we search our database D and instead of selecting just one target image, we select subset of images D^* that has emotion distribution similar to the desired emotion distribution, as described in eq. (2).

$$D^* = \{I_i : bc_i > \omega\}, \quad \text{where} \quad \omega = 1.5 * \text{mean}(bc_{i=\{1,2,\dots,N\}}) \quad (2)$$

In order to compute similarity between the images I_s and $I_i \in D^*$ we consider not only similar concepts that two images have in common but also the spatial and color combinations in which these concepts exist. A recent study in [15] shows that top layers of a Convolutional Neural Network (CNN) capture high level concepts along with discriminant low level features. Therefore, features extracted from top layers of pre-trained networks could be used as to compute semantic and context similarity us get target images that are semantically more closer to the source image. In order to select features that capture both the object and scene information, we use AlexNet [8] and GoogleNet networks [16], trained on ImageNet and Places2 datasets respectively. We used *fc7* and *pool5* layers of the respective networks to extract features. For each image, AlexNet and GoogleNet models give us 4096 and 1024 dimensional output vectors which are then normalized and concatenated together to get a single 5120 dimensional vector denoted as $\mathbf{f}$.

We apply K-nearest-neighborhood algorithm to select set of K target images $T = \{I_{t_k}\}_{k=1}^K$ from D^* that are semantically closer to the provided source image, using $\mathbf{f}$, eq.(3). We chose K to be 10 in our experiments because mean distance (representing semantic gap) between source and nearest neighborhood images increases drastically as we further increase K and implications of which is that our results deteriorate qualitatively.

$$\underset{T}{\operatorname{argmin}} \|\mathbf{f}_s - \mathbf{f}_j\| \quad \forall j \in D^* \quad (3)$$

where $\mathbf{f}_s$ represents the source image features. For each selected target image i in K , a histogram $\mathbf{h}_i$ is computed as in [13] and we take weighted summation using the Bhattacharyya coefficient bc_i of K selected images $h = \sum_{i=1}^K (bc_i * \mathbf{h}_i)$. Note that, bc_i measures similarity between target emotion distribution and the selected images. In this way the target histogram is more affected by the image whose distribution is more closer to the target emotion distribution. After computing target image histogram, we used color transfer algorithm in [13] to compute the final transformed image.

5 Experiments and Results

We apply our algorithm on images in Emotion6 dataset and few of the famous images. Fig. (1), Fig. (5), Fig. (6) and Fig. (7) demonstrate our emotion transfer results. Considering that emotions are subjective in nature, we conduct a user study to evaluate how good

Target Emotion	Source Image	Target Emotion Distribution	Selected Images from Emotion6 Dataset	Transformed Image
Anger
Disgust
Fear
Joy
Sadness

Figure 5: Examples of transferring evoked emotion distribution. Second column contains the source image. Third and fourth column indicates the Target Emotion Distribution provided by the user and the images selected by our algorithm respectively.

our transformed image represent the targeted emotion. The emotion with highest value in provided emotion distribution is chosen to be target emotion.

Experiment 1 In our first experiment, we select top 100 neutral images, based on the ground truth, from Emotion6 dataset [14] as source images. For each image, we use target emotion to be the one that selects target images with most similar probability distribution as the desired distribution. This is performed by computing Bhattacharyya coefficient (see equation 1) between desired distribution and emotion distribution of selected target images by our algorithm. The target emotion distribution is then obtained by setting the target emotion to 1 in the source image’s emotion distribution and then normalizing the distribution such that it sums to 1. Out of the total 100 images, the number of images transformed for

Figure 6: Emotion transformation on some images from ArtPhoto dataset. First row: change of affect class. Second row: enhancing the arousal value of induced emotion.

each emotion are shown in Fig. (4). A few source-transformed pairs of this user study are shown in Fig. (5). To evaluate that the emotion distribution of transformed image has higher value for targeted emotion than the neutral, we presented these 100 source-transformed image pairs in random order to different subjects and ask them to tag the one image that induce more target emotion in them. User study was conducted through an online web portal that collects subject’s responses to the shown images. The link was shared online, and undergraduate and graduate students of computer science department were requested to take part in the study. The demographics of the user study, therefore, consists of both male and female subjects, of age ranging from 17 to 25. We got a total of 1671 responses, where on average each image was tagged almost by 16 times. The statistics generated by this user study show that 65.0% of the times our emotion transfer algorithm successfully transformed neutral images towards the target emotion.

Fig. (4) shows that some of our neutral images failed to transform to joy emotion, this is because the images chosen contain the neural objects and simple color transformation cannot make these images elicit joyous feeling. An example of such transformation is given in first row and second column of figure 9. These results clearly demonstrate that high level concepts present in the images constraint how much the emotion of the image can be modified.

Experiment 2 In the second experiment we transformed emotion of some images from ArtPhoto dataset. The results in Fig. (6) show that we are able to successfully perform emotion transformation on these images. In addition to these images, we have also tried to transform induced emotion of few popular photographs. The results of these images are shown in Fig. (7). For all images in Fig. (6) and Fig. (7) majority of subjects voted that transformed images are more close to the target emotion as compared to the source images.

5.1 Which layer of CNN to use?

Our decision to choose fc7 layer of the AlexNet was based on multiple factors, including the size of the of the output from the layer and how much information is captured by that layer. The output layer quantize all the information to the object probability, this would have reduced similarity measure to just counting similar objects in the images. We wanted to measure similarity on the basis of spatial, context and content information. In order to identify

Figure 7: Emotion transformation on some popular photographs. In the first row, we have shown to increase anger with bit of sadness and fear, whereas, in the second image we have tried to increase the element of fear with a bit of sadness.

Figure 8: The left column shows transformation results using fc6 features of Alexnet network while right column shows results with fc7 features. The transformation direction is Sad to Joy in first row and Fear to Joy in second row.

the top layer of AlexNet network which has features that are more appropriate for selecting target images and performing emotion transformation, we perform emotion transformation using features from two top layers, named fc6 and fc7, separately. The Fig. (8) shows the results of this experiment on two images. These results clearly depict that fc7 features select more appropriate target images and hence generate transformation results that are more close to the target emotion and are more appealing.

6 Failure Cases and Question of High Level Concepts

As visible from the results in Fig. (9), one cannot transfer any image to incite all the emotions. One of the reason is as explained by Ali et al. in [14], the correlation between the high level concepts in the images and their corresponding elicited emotions, which means that content of the image defines and restricts the spectrum of the emotions that could be

elicited from the image. Since we are only transforming low level features (color features) we cannot escape that spectrum. For example, no amount of automatic color and texture manipulation can transform the image of a sad girl to one arousing joyous feelings, without HLC dependent manipulations. Similarly, the images containing neutral concepts (objects having no direct emotion label attach to them) can not be transformed to any other emotion as represented in the second image of first row of Fig. (9).

Figure 9: Failure examples of Emotion transformation

7 Conclusion

We are witnessing exponential increase in both creation of multimedia, especially photographs, and its sharing on the social network. Digital libraries with millions of images are available for people to share, comment and search on. However, affective analysis of the images, affect base image retrieval or image manipulation to induce emotion has gained much less traction in vision community. In this paper, we present a method that allows a user to manipulate an image just by providing the desired emotion distribution. We perform search in the existing database by minimizing over distance between input distribution and emotion distribution of image in database, and content features of the respective images. Histogram for color transformation is constructed from these images selected from database and [13] is used to transform input image. Since, we use features captured from top layers of CNN trained on object detection and scene identification, we are able to find the images which are similar in content and spatial structure as input image. This allows us to avoid semantic segmentation or patch matching which like previous methods could have restricted us to only few types of images. Use of emotion distribution makes the manipulation much more interpretable. We performed a detailed user-study that showed that transformed images generated through our method are better representative of target emotion than the original input image. The failure case and their reasoning is provided, highlighting the limitation of color-transformation based emotion transfer methods.

8 Acknowledgements

We thank Mr. Junaid Sarfraz for his assistance with the design and implementation of on-line web portal. The portal was used for conducting user study on perceptual analysis of transformed images.

References

- [1] Afsheen Rafaqat Ali, Usman Shahid, Mohsen Ali, and Jeffrey Ho. High-level concepts for affective understanding of images. In *IEEE Winter Conference on Applications of Computer Vision*, 2017.
- [2] Paul Ekman. An argument for basic emotions. *Cognition & Emotion*, 6(3-4):169–200, May 1992. doi: 10.1080/02699939208411068. URL <http://dx.doi.org/10.1080/02699939208411068>.
- [3] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [4] Li He, Hairong Qi, and Russell Zaretski. Image color transfer to evoke different emotions based on color combinations. *Signal, Image and Video Processing*, 9(8):1965–1973, 2015.
- [5] Hristina Hristova, Olivier Le Meur, Rémi Cozot, and Kadi Bouatouch. Style-aware robust color transfer. In *Proceedings of the workshop on Computational Aesthetics*, pages 67–77. Eurographics Association, 2015.
- [6] Aditya Khosla, Akhil S Raju, Antonio Torralba, and Aude Oliva. Understanding and predicting image memorability at a large scale. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2390–2398, 2015.
- [7] Hye-Rin Kim, Henry Kang, and In-Kwon Lee. Image recoloring with valence-arousal emotion model. In *Computer Graphics Forum*, volume 35, pages 209–216. Wiley Online Library, 2016.
- [8] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems 25*, pages 1106–1114, 2012. URL http://books.nips.cc/papers/files/nips25/NIPS2012_0534.pdf.
- [9] Jana Machajdik and Allan Hanbury. Affective image classification using features inspired by psychology and art theory. In *Proceedings of the international conference on Multimedia MM 10*, pages 83–92, 2010. ISBN 9781605589336. doi: 10.1145/1873951.1873965. URL <http://portal.acm.org/citation.cfm?doid=1873951.1873965>.
- [10] Naila Murray, Sandra Skaff, Luca Marchesotti, and Florent Perronnin. Towards automatic concept transfer. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Non-Photorealistic Animation and Rendering*, pages 167–176. ACM, 2011.
- [11] Netscope. Netscope-neural network visualizer, 2015. URL <http://ethereon.github.io/netscope/quickstart.html>. [Online; accessed 19-July-2017].
- [12] Kuan-Chuan Peng, Tsuhan Chen, Amir Sadovnik, and Andrew Gallagher. A Mixed Bag of Emotions: Model, Predict, and Transfer Emotion Distributions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 860–868. IEEE, jun 2015. doi: 10.1109/CVPR.2015.7298687. URL <http://ieeexplore>.

- ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7298687<http://research.google.com/pubs/AndrewGallagher.html>.
- [13] Tania Pouli and Erik Reinhard. Progressive histogram reshaping for creative color transfer and tone reproduction. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering*, pages 81–90. ACM, 2010.
 - [14] Muhammad Rushdi, Mohsen Ali, and Jeffrey Ho. Color de-rendering using coupled dictionary learning. In *ICIP*, pages 315–319, 2013.
 - [15] Thales Teixeira, Michel Wedel, and Rik Pieters. Emotion-induced engagement in internet video advertisements. *Journal of Marketing Research*, 49(2):144–159, 2012.
 - [16] Wikipedia. Bhattacharyya distance — wikipedia, the free encyclopedia, 2016. URL https://en.wikipedia.org/w/index.php?title=Bhattacharyya_distance&oldid=722480916. [Online; accessed 19-July-2017].
 - [17] Mengdi Xu, Bingbing Ni, Jinhui Tang, and Shuicheng Yan. Image re-emotionalizing. In *The Era of Interactive Media*, pages 3–14. Springer, 2013.
 - [18] Su Xue, Aseem Agarwala, Julie Dorsey, and Holly Rushmeier. Learning and applying color styles from feature films. In *Computer Graphics Forum*, volume 32, pages 255–264. Wiley Online Library, 2013.
 - [19] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *Computer vision—ECCV 2014*, pages 818–833. Springer, 2014.
 - [20] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, 2016.
 - [21] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning deep features for scene recognition using places database. In *Advances in Neural Information Processing Systems*, pages 487–495, 2014.