The Emotional GAN: Priming Adversarial Generation of Art with Emotion

David Alvarez-Melis
CSAIL, MIT
dalvmel@mit.edu

Judith Amores
MIT Media Lab
amores@mit.edu

Abstract

We propose a neural network method for turning emotion into art. Our approach relies on a class-conditioned generative adversarial network trained on a dataset of modern artworks labeled with emotions. We generate this dataset through a large-scale user study of art perception with human subjects. Preliminary results show our framework generates images which, apart from aesthetically appealing, exhibit various features associated with the emotions they are conditioned on.

1 Introduction and Background

The influence of art in emotional state has been long studied in psychology and neuroscience (e.g., [11] and reference therein). For example, Neuroesthetics studies the neural bases behind aesthetic experiences such as artwork creation or contemplation [1]. More generally, the field of color psychology has studied how hue influences human behavior [14]. It is known that people react differently to color stimuli based on their previous experiences and biological components [5]. Despite this subjectiveness, there are certain color associations that are common across individuals. The association of the color black and death is innate, which can lead to aggressiveness on people exposed to it [4]. Valdez and Mehrabian [13] measured the emotional effect of color by means of continuous *pleasantness* and *arousal* scales, concluding that blue-green-purple hues are more often perceived as positive, whereas yellow and green-yellow are the least positive. On the other axis, green-yellow, blue-green, and green are the most arousing, while purple-blue and yellow-red are the opposite.

From the artificial intelligence perspective, automatic generation of art has been a long-standing objective [2]. Recent advances in deep generative models have been successfully applied to the artwork domain. For example, Elgammal et al. [3] propose a method to generate art by learning about styles and deviating from style norms. Departing from previous work on art generation, here we focus specifically on the interaction between emotions and aesthetics. Specifically, we aim to invert the `stimuli`→`emotion` mapping, i.e., given an emotion we seek to generate an artwork that conveys it. For this, we use a Generative Adversarial Network (GAN) with Auxiliary Classifier (AC-GAN)[9], an extension of GAN that conditions generation on class labels. In this architecture, the generator is fed a class label in addition to the noise variables, and is penalized through an additional loss term from a classifier that attempts to predict the class label only from the generated image. Training this method requires a large dataset with class labels. Thus, we construct a large dataset of modern art by collecting images from various museum collections, conducting a large-scale user study to obtain human annotations for a subset of the dataset, and labeling the rest of the collection with a purpose-built emotion classification network. Our preliminary results show that images generated by our method exhibit features that are commonly associated with the emotions they are conditioned on, such as aggressive strokes for the negative emotions and colorful, rounded shapes for the positive ones, confirming findings of previous work in the psychology literature.

31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA.

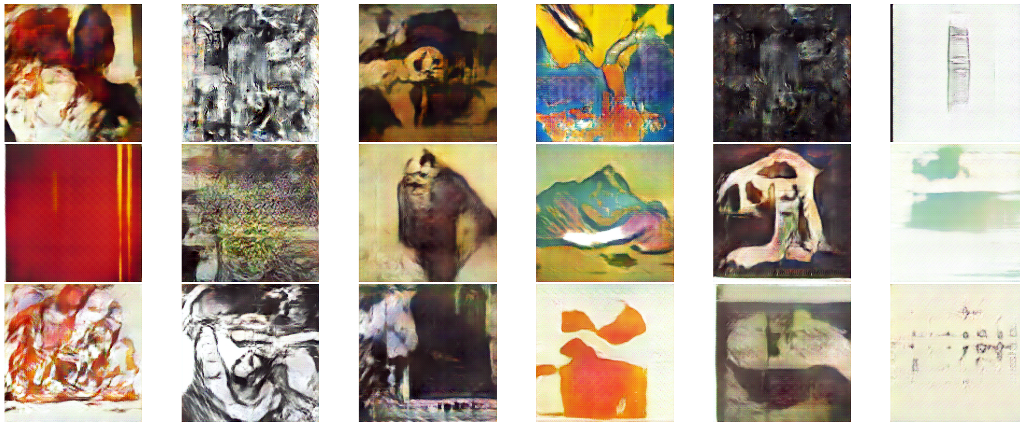

(a) Anger (b) Anxiety (c) Fear (d) Joy (e) Sadness (f) Neutral

Figure 1: Artworks generated by the Emotional GAN conditioning on various emotion categories.

2 System Description

Raw data collection. We created the emotion-labeled modern art dataset as follows. Starting from the full MoMA [8] and Wikiart [15] collections, we filtered the data by removing works that were unlikely to carry emotion, such as maps, photographies or architecture plans. For WikiArt, we selected only those paintings that were produced from the 1860s onwards. After eliminating very large and broken images, we were left with about 100K paintings in the dataset.

Labeling the Data. We randomly split the dataset in two subsets. The first one (13,000 paintings from MoMA) was labeled with emotions in a large-scale user study, while the remaining 20,419 images were labeled automatically using a CNN classifier trained on the former. The user study was conducted using the HIVEAI platform[1]. For each artwork image, the user had to select the closest feeling that evoked on them (*anger, anxiety, disgust, excitement, fear, joy, calmness, sadness, lust* or *neutral*). The quality of the responses was ensured by mixing gold-labeled images with the tasks, and keeping only annotators with accuracy on these tasks above a certain threshold. We filtered out images for which inter-annotator agreement fell below a 66 % accuracy and removed those that were returned as inconclusive. This step yielded a total of 2,948 labeled images. We then trained a CNN emotion classifier on these gold-labeled artwork-emotion pairs, and used it to the label the remaining artworks. We discarded all examples where the emotion was predicted with less than 70% confidence by the classifier. This resulted in a final labeled dataset with 33,418 examples.

The model. We used the AC-GAN as proposed in [9] with default parameters, generating images at 128x128 pixel resolution. To enhance the quality of the generated artwork, we use a Super-Resolution GAN [7] to scale the images to a final 512x512 pixel resolution.

3 Results

We train the network with the Adam optimizer [6] with initial learning rate $l = 0.0002$. We observed that early in the training there are few subtle differences between images generated conditioned on different emotions, which get steadily amplified as training continues. Figure 1 shows a set of representative artworks generated by our method for a subset of emotions after 100 epochs. Interestingly, our method generates a variety of styles with some high-level emotional features that agree with previous psychology literature: e.g., bright colors for arousal, such as red for `anger` and dark colors for unpleasant and low arousal, such as `sadness and fear`. Joy is an emotional state with a medium-high arousal and positive valence [10], agreeing with the literature, our generated artwork has a mix of green-yellow and blue-green colors. In addition, we observed that many examples in the `joy` category resemble natural landscapes (also demonstrated to have positive psychophysiological effects [12]). More generated artworks can be found at `emotion2art.media.mit.edu`.

[1] thehive.ai

References

[1] A. Chatterjee. Neuroaesthetics: A coming of age story. *Journal of Cognitive Neuroscience*, 23 (1):53–62, 2011. doi: 10.1162/jocn.2010.21457. URL `https://doi.org/10.1162/jocn.2010.21457`. PMID: 20175677.

[2] S. DiPaola and L. Gabora. Incorporating characteristics of human creativity into an evolutionary art algorithm. *Genetic Programming and Evolvable Machines*, 10(2):97–110, 2009.

[3] A. Elgammal, B. Liu, M. Elhoseiny, and M. Mazzone. Can: Creative adversarial networks, generating" art" by learning about styles and deviating from style norms. *arXiv preprint arXiv:1706.07068*, 2017.

[4] M. G. Frank and T. Gilovich. The dark side of self-and social perception: black uniforms and aggression in professional sports. *Journal of personality and social psychology*, 54(1):74, 1988.

[5] A. C. Hurlbert and Y. Ling. Biological components of sex differences in color preference. *Current Biology*, 17(16):R623–R625, 2007.

[6] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[7] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint arXiv:1609.04802*, 2016.

[8] MoMa. The museum of modern art (moma) collection data, 2017.

[9] A. Odena, C. Olah, and J. Shlens. Conditional image synthesis with auxiliary classifier GANs. In *ICML*, pages 2642–2651, 2017.

[10] J. A. Russell. Evidence of convergent validity on the dimensions of affect. *Journal of personality and social psychology*, 36(10):1152, 1978.

[11] P. J. Silvia. Emotional responses to art: From collation and arousal to cognition and emotion. *Review of general psychology*, 9(4):342, 2005.

[12] R. S. Ulrich. Natural versus urban scenes – Some psychophysiological effects. *Environ Behav*, 13(5):523–556, 1981.

[13] P. Valdez and A. Mehrabian. Effects of color on emotions. *Journal of experimental psychology: General*, 123(4):394, 1994.

[14] T. Whitfield and T. Whiltshire. Color psychology: a critical review. *Genetic, social, and general psychology monographs*, 1990.

[15] WikiArt. Wikiart, 2017.

A User Study

A.1 Instructions

The instructions page shown to the labelers is shown in Figure 2.

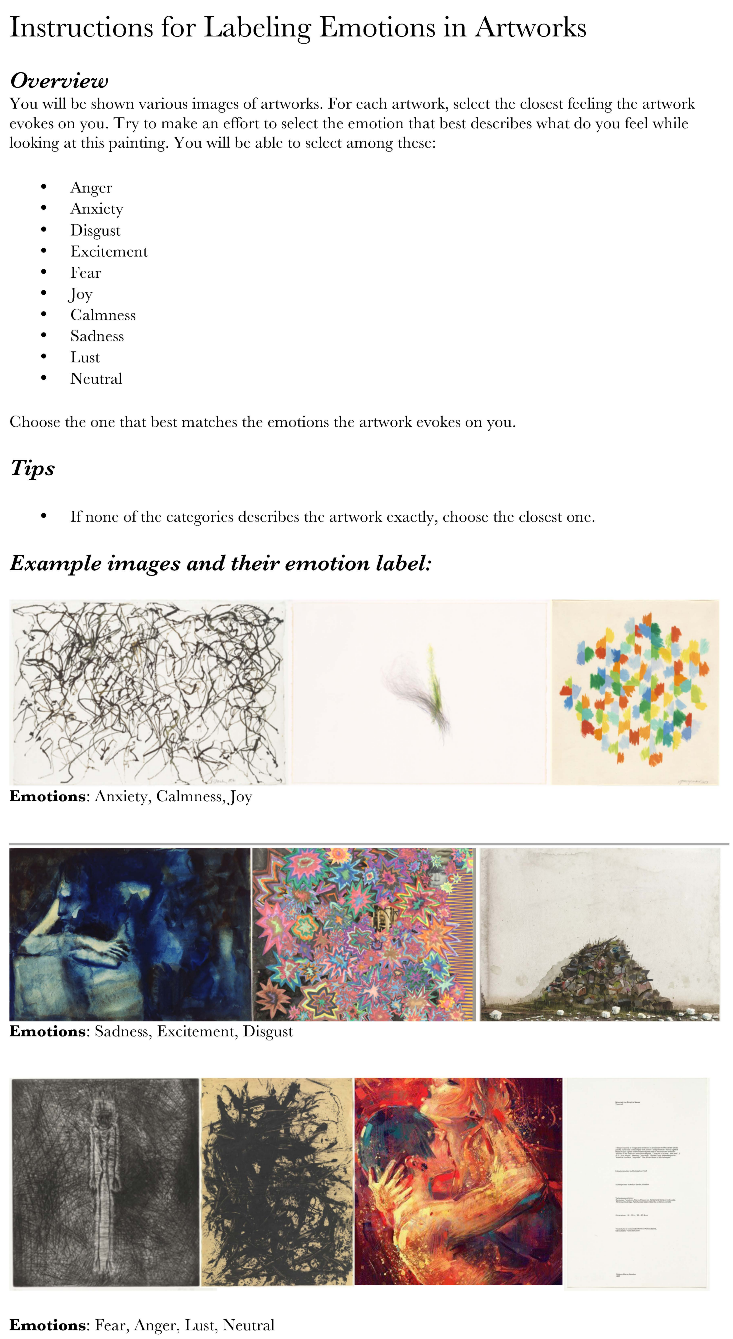

Instructions for Labeling Emotions in Artworks

Overview

You will be shown various images of artworks. For each artwork, select the closest feeling the artwork evokes on you. Try to make an effort to select the emotion that best describes what do you feel while looking at this painting. You will be able to select among these:

- Anger
- Anxiety
- Disgust
- Excitement
- Fear
- Joy
- Calmness
- Sadness
- Lust
- Neutral

Choose the one that best matches the emotions the artwork evokes on you.

Tips

- If none of the categories describes the artwork exactly, choose the closest one.

Example images and their emotion label:

Emotions: Anxiety, Calmness, Joy

Emotions: Sadness, Excitement, Disgust

Emotions: Fear, Anger, Lust, Neutral

Figure 2: Instructions for artwork labeling task