The ability of generative AI to express emotions through abstract images: a preliminary study design

Fabrizio Serrao[1*], Alberto Gallace[1], Marcello Gallucci[1], and Alessandro Gabbiadini[1]

[1] *Department of Psychology, University of Milano Bicocca, Piazza dell'Ateneo Nuovo 1, 20126, Milan, Italy.*

Abstract

The advent of Artificial Intelligence (AI) has transformed various aspects of human life, including art. Generative AI technologies can produce complex graphic artworks based solely on textual descriptions and beyond literal prompts. Here, we present the experimental protocol of an ongoing study focusing on generative AI's ability to express emotions through non-figurative elements. A body of literature suggests that such perceptual characteristics can be inherently expressive, regardless of the content of the artwork. Some authors tested this hypothesis by asking artists to express emotion through abstract artworks and evaluated the expressiveness of those artworks. Although they found that specific emotions were consistently associated with specific artworks, these conclusions can hardly be generalized. Generative AI can theoretically circumvent this issue and produce a larger and more unbiased set of emotion-laden abstract stimuli. In our study, participants will view six sets of AI-generated abstract images that are intended to express six target emotions (joy, sadness, fear, anger, disgust, surprise). We will assess whether the images effectively convey these emotions. Additionally, if the AI succeeds, we will investigate the impact of knowing that the "artist" is not a human on the expressiveness of those images. Indeed, the awareness that an image is generated by an AI was shown to influence the aesthetic judgment of it.

Keywords

Art evaluation, neuroaesthetics, artificial intelligence, emotions.

1. Introduction

In recent years, Artificial Intelligence (AI) has revolutionized many aspects of human life, including artistic production [1]. The ability of AI to generate artistic content, once considered a distinctive feature of humans, becomes increasingly refined. Thanks to models trained on millions of illustrations and images, generative AIs -like Midjourney, Dall-E, and Stable Diffusion- have become capable of producing visual content potentially indistinguishable from those by visual artists. However, the real challenge for AI is not merely recreating the perceptual features of artworks but rather capturing the emotional expression that often characterizes artistic works.

Humankind has a long history of expressing emotions through art, both through figurative subjects and abstract elements. While emotions in figurative works can be more obviously identified because of the represented subjects, abstract artworks can also convey emotions through the use of specific colors, lines, shapes, and other specific features [2].

This raises an important question: can AI capture and convey emotions through abstract elements as effectively as human artists? While it was shown that AI models can predict human emotional

Italian Workshop on Artificial Intelligence for Human Machine Interaction (AIxHMI 2024), November 26, 2024, Bolzano, Italy
*Corresponding author.

f.serrao1@campus.unimib.it (F. Serrao); alberto.gallace1@unimib.it (A. Gallace); marcello.gallucci@unimib.it (M. Gallucci); alessandro.gabbiadini@unimib.it (A. Gabbiadini)

0009-0007-7867-4899 (F. Serrao); 0000-0002-4561-8926 (A. Gallace); 0000-0003-3546-0093 (M. Gallucci); 0000-0002-7593-8007 (A. Gabbiadini)

reactions to artworks by inferring the emotional connotation of the represented subjects and scenes [3], the ability to evoke emotions through abstract elements remains a complex challenge.

In addition to that, the awareness that a given content was generated by an AI might produce emotional reactions *per se*. Understanding such reactions is crucial for the future development of this technology in the field of artistic production and visual communication.

Psychological studies have already begun to explore how people perceive and interact with AI-generated artworks [4, 5, 6]. Knowing that a work was created by AI can influence aesthetic judgment [4, 6], but the emotional impact of this awareness remains unclear. As far as we know, only one study has addressed this issue [5], yet they used neutral and thus poorly moving stimuli. Here, we describe the design of two studies aimed at filling this gap.

2. Study 1

This experimental paradigm is reminiscent of the one of Takahashi [7] and Damiano et al. [8]. These authors asked artists to represent a series of emotions through abstract drawings -i.e., by drawing only lines and geometric shapes- and subsequently asked other participants to judge the emotions conveyed by such drawings. However, in the present paradigm, there is a fundamental difference: the images will be produced not by human artists but by a generative AI.

It is important to note that asking human artists to express emotions has limitations: while the aforementioned studies [7, 8] found that a set of drawings did express a given emotion and shared certain features, it can only be concluded that those specific drawers effectively expressed emotion in such a way. Because an AI is trained on billions of images by a multiplicity of artists, it would enable us to draw more general conclusions.

2.1. Methods

The visual stimuli will be generated using the web platform Midjourney (www.midjourney.com). The algorithm will be instructed to generate emotionally stirring images characterized exclusively by abstract elements. The following prompt will be used:

/imagine: the emotion of [name of the emotion] abstract--no pictures, objects, symbols, people, person, faces, humans, animals, expressions, figurative, landscapes, mouth, eyes, street, sun, buildings.

This way, we will generate 20 abstract images for each of the following basic emotions [9]: joy, sadness, fear, anger, disgust, and surprise (some examples are shown in Figure 1). Subsequently, the generated images will be validated by four independent judges who will exclude images containing any figurative element. We will retain only those images that all judges will deem to be abstract (Cohen's k =1). From these images, 30 will be randomly selected in total: 5 for joy, 5 for sadness, 5 for fear, 5 for anger, 5 for disgust, and 5 for surprise.

Figure 1: Sample of the images generated by the AI. The images in each row are intended to express four of the six target emotions: joy (top row), sadness (upper middle row), fear (lower middle row), and anger (bottom row).

Regarding the number of participants, a sample size of 30 or more is often considered sufficient to achieve a normal distribution [10]. Therefore, considering potential dropouts, 40 participants will be involved.

Using a snowball sampling technique, participants will be invited to participate in the study via social networks and instant messaging platforms. Each participant will evaluate the 30 images (presented at a resolution of 720x720 pixels, 72 dpi) in random order on Qualtrics platform (https://www.qualtrics.com). For each image, participants will be asked to evaluate the emotions it expresses without trying to identify recognizable objects or scenes. Specifically, participants will be asked to rate how much each image expresses each of the 6 emotions (joy, sadness, fear, anger, disgust, and surprise) on a Likert scale from 1 to 7. No information about the source of the images will be provided. We will then compare consistent ratings -i.e., when the rated emotion matches the emotion in the generative prompt- with inconsistent ratings.

In addition, we will analyze the features of the images to uncover correlations between visual features and emotions. Color features -e.g., mean saturation, mean lightness, mean hue, lightness entropy, and color entropy- and line features -e.g., orientation anisotropy and orientation entropy- will be extracted through the Aesthetic Toolbox [11].

First, we will investigate what features are shared by AI-generated images within each emotional category and compare them across categories. This will enable us to study the underlying generative criteria. Second, we will analyze the images that will be rated as more expressive (average score >4) for each emotion, whether or not they were meant to convey that emotion. This way, we want to elucidate what perceptual features make abstract images more likely to express a given emotion. In case all images in one category have a mean score lower than 4, we will consider the three images with the highest mean score.

2.2. Expected results

Midjourney and similar AI software have been trained on billions of images, including abstract artworks. Some of these works are emotionally stirring. Therefore, we expect Midjourney to be able to extract perceptual features that induce emotions and generate abstract compositions that can evoke such emotions. That is, we expect consistent ratings to be higher than inconsistent ratings for each emotion category. For example, we expect images created to express joy to be rated as more joyful than sad, fearful etc.

Even if the results reveal that Midjourney cannot evoke the desired emotion, it will still be interesting to analyze the similarities among the most emotionally stirring images for each emotion and identify unknown patterns.

Additionally, we expect the most emotionally engaging images (average score > 4) to show some of the perceptual features discussed in the literature [2, 7, 8]. This would support the hypothesis that such features are inherently moving and that they are interpreted as signals of the author's emotional state [2, 12]. This would help provide visual artists with a validated toolbox for generating emotional content.

3. Study 2

The aim of this paradigm is to verify whether the awareness that images are generated by AI moderates the emotional impact of such images. Indeed, the effect of contextual knowledge - including the artist's intentions and emotions- on artistic evaluation is well-established [13]. This seems to be particularly true in the case of abstract images [12].

3.1. Methods

For each emotional category, the images from the previous study that will have scored the highest (i.e., >4) in that emotion will be used. If all images in one emotional category will have a score lower than 4, we will drop that category. Yet, if this will apply to all categories, we will select the three images with the highest mean score.

The stimuli will be grouped into blocks by emotion. The order of the blocks will be randomized, as will the order of the stimuli within each block.

The necessary sample size for hypothesis testing was calculated using G*Power software, considering an effect size f=0.25, a power = .80, and 2 groups. The analysis suggests that an adequate sample size corresponds to 128 participants. Considering potential dropouts, 150 participants will be involved.

Participants will be invited to participate in the study on a voluntary basis and contacted using a snowball sampling technique.

They will be randomly assigned to one of the two experimental conditions (AI images vs. artist images). Although the images to be evaluated will be the same for both groups, the first group will be told the images are generated by artificial intelligence, while the second group will be told that they are produced by human artists.

Participants will rate how much each image expresses the emotion that it was prompted to express. They will view the images through Qualtrics platform as in the previous study and rate them on a Likert scale from 1 to 7. Note that, this time, participants will not express judgments regarding the other five emotions. For example, when viewing images created to express joy, participants will only rate how much those images express joy. The purpose is to test the influence of information priming on the images that have proven to be the most emotionally engaging.

This experimental design will be a *between-within* subject type. Specifically, the study will test which of the six emotions investigated is elicited most strongly (*within* factor) and whether there is a difference based on the purported source of the image (AI vs. Human, *between* factor).

3.2. Expected results

A significant main effect of the experimental manipulation (AI images vs. artist images) is hypothesized. Understanding the artist's intentions and emotions was shown to play a crucial role in the evaluation of art [2, 12, 13]. This applies especially to abstract images, which can be interpreted as emotional cues and thus promote empathy for the artists [2]. Therefore, we expect those aware that the images are generated by artificial intelligence to be less emotionally engaged [5]. Nevertheless, while generative AI lacks an inner life, it is still possible that participants will endow it with human characteristics, as shown by Paiva and colleagues [14]. This might nullify the difference between the two conditions. In either case, this second paradigm would produce useful results for those who intend to use AI-generated images to achieve an emotional impact.

4. Conclusion

The ability of generative AI to express emotions through abstract compositions would carry profound implications for the fields of art, psychology, and human-computer interaction. Considering that artistic expression has historically been rooted in human experience, this would invite us to reflect on artistic agency and creativity. It would force us to question the essence of emotional connection and the role of the artist. Ultimately, it would open new possibilities for artistic expression and innovative applications in fields like advertising, therapy, and education.

References

[1] C. Then, E. J. Soewandi, M. F. Danial, S. Achmad and R. Sutoyo, "The impact of artificial intelligence on art -a systematic literature review," in *2023 IEEE 9th Information Technology International Seminar (ITIS)*, 2023.

[2] F. Serrao, A. Chirico, A. Gabbiadini, A. Gallace and A. Gaggioli, "Enjoying art: an evolutionary perspective on the esthetic experience from emotion elicitors," *Frontiers in Psychology,* vol. 15, 2024.

[3] S. Chilappagari, R. S. Sidhu, P. Gayen, R. Mullick and P. Patnaik, "MACHINE, OH! MACHINE, WHAT IS THE EMOTION IN THIS PAINTING? ASSESSING EMOTIONS IN ABSTRACT ARTS THROUGH MACHINE LEARNING TECHNIQUES," in *Proceeding of the 3rd International Conference on Arts and Humanities*, 2016. D. Harel, First-Order Dynamic Logic, volume 68 of Lecture Notes in Computer Science, Springer-Verlag, New York, NY, 1979. doi:10.1007/3-540-09237-4.

[4] L. Bellaiche, R. Shahi, M. H. Turpin, A. Ragnhildstveit, S. Sprockett, N. Barr and P. Seli, "Humans versus AI: whether and why we prefer human-created compared to AI-created artwork," *Cognitive Research: Principles and Implications,* vol. 8, no. 1, p. 42, 2023.

[5] T. R. Demmer, C. Kühnapfel, J. Fingerhut and M. Pelowski, "Does an emotional connection to art really require a human artist? Emotion and intentionality responses to AI-versus human-created art and impact on aesthetic experience," *Computers in Human Behavior,* vol. 148, 2023.

[6] C. Di Dio, M. Ardizzi, S. V. Schieppati, D. Massaro, G. Gilli, V. Gallese and A. Marchetti, "Art made by artificial intelligence: The effect of authorship on aesthetic judgments," *Psychology of Aesthetics, Creativity, and the Arts,* 2023.

[7] S. Takahashi, "Aesthetic properties of pictorial perception," *Psychological review,* vol. 102, no. 4, p. 671, 1995.

[8] C. Damiano, P. Gayen, M. Rezanejad, A. Banerjee, G. Banik and P. Patnaik, "Anger is red, sadness is blue: emotion depictions in abstract visual art by artists and non-artists," *Journal of Vision,* vol. 23, no. 1, 2023.

[9] P. Ekman, "Are there basic emotions?," *Psychological Review,* vol. 99, no. 3, p. 550–553, 1992.

[10] A. Ghasemi and S. Zahediasl, "Normality tests for statistical analysis: a guide for non-statisticians," *International journal of endocrinology and metabolism,* vol. 10, no. 2, p. 486, 2012.

[11] C. Redies, R. Bartho, L. Koßmann, B. Spehar, R. Hübner, J. Wagemans and G. U. Hayn-Leichsenring, "A toolbox for calculating objective image properties in aesthetics research," 2024.
[12] D. Freedberg and V. Gallese, "Motion, emotion and empathy in esthetic experience," *Trends in Cognitive Sciences,* vol. 11, p. 197–203, 2007.
[13] M. Pelowski, P. S. Markey, M. Forster, G. Gerger and H. Leder, "Move me, astonish me… delight my eyes and brain: the Vienna integrated model of top-down and bottom-up processes in art perception (VIMAP) and corresponding affective, evaluative, and neurophysiological correlates," *Physics of Life Reviews,* vol. 21, p. 80–125, 2017.
[14] A. Paiva, I. Leite, H. Boukricha and I. Wachsmuth, "Empathy in virtual agents and robots: A survey," *ACM Transactions on Interactive Intelligent Systems (TiiS),* vol. 7, no. 3, pp. 1-40, 2017.