

StyleCLIPDraw: Coupling Content and Style in Text-to-Drawing Translation

Peter Schaldenbrand¹, Zhixuan Liu², Jean Oh^{1*}

¹The Robotics Institute, Carnegie Mellon University

²School of Data Science, The Chinese University of Hong Kong
pschalde@andrew.cmu.edu, zhixuanliu@cuhk.edu.cn, jeanoh@cmu.edu

Abstract

Generating images that fit a given text description using machine learning has improved greatly with the release of technologies such as the CLIP image-text encoder model; however, current methods lack artistic control of the style of image to be generated. We present an approach for generating styled drawings for a given text description where a user can specify a desired drawing style using a sample image. Inspired by a theory in art that style and content are generally inseparable during the creative process, we propose a coupled approach, known here as StyleCLIPDraw, whereby the drawing is generated by optimizing for style and content simultaneously throughout the process as opposed to applying style transfer after creating content in a sequence. Based on human evaluation, the styles of images generated by StyleCLIPDraw are strongly preferred to those by the sequential approach. Although the quality of content generation degrades for certain styles, overall considering both content and style, StyleCLIPDraw is found far more preferred, indicating the importance of style, look, and feel of machine generated images to people as well as indicating that style is coupled in the drawing process itself. Our code, a demonstration, and style evaluation data are publicly available.

1 Introduction

Creating visual content is a challenging task that requires skill, time, and resources. Painting lessons can be expensive and prohibitively time consuming for many people to participate in. While visual content creation skills may be a privilege, there are common skills that the majority of people can use for sharing their ideas, such as verbal communication. For instance, people may find it difficult to draw a complex scene, but they are likely able to describe it using language. In this context, our research studies how Artificial Intelligence (AI) can act as a tool to overcome this barrier and empower people to appreciate the joy of creative practices that they otherwise would be incapable of. In this paper, we specifically

*Corresponding Author

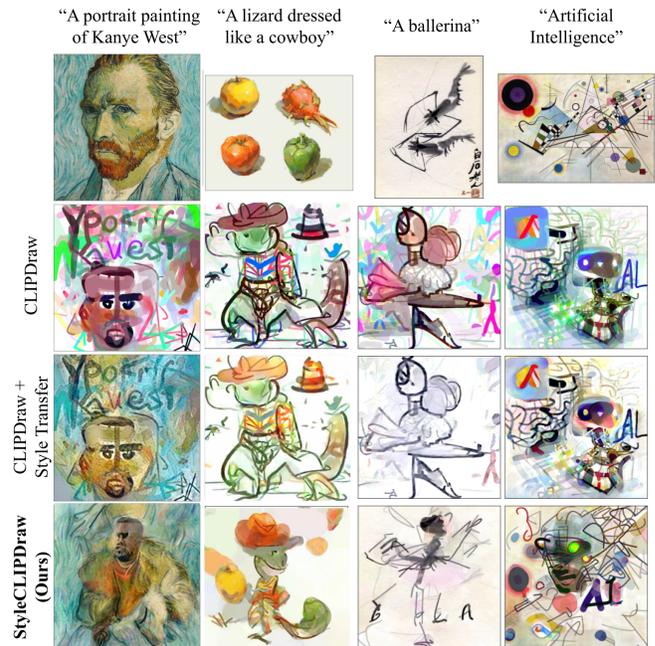


Figure 1: Comparison between our (StyleCLIPDraw) method and the baseline (CLIPDraw + Style Transfer). The top row shows the language input followed by the style image input.

focus on generating drawings from text description.

Language is a common form of communication to describe visual information at some level of abstraction. Recent technological advances such as OpenAI’s Contrastive Language-Image Pre-training (CLIP) [Radford *et al.*, 2021] have enabled high quality language-vision translation methods such as Dall-E [Ramesh *et al.*, 2021], GLIDE [Nichol *et al.*, 2021], and CLIPDraw [Frans *et al.*, 2021]. Using these technologies, people can write short, language descriptions to create an image in collaboration with an AI system; however, as shown in the drawings produced by CLIPDraw, e.g., the second row of Figure 1, all have a similar visual style. The language description can be used to alter the style of the drawing but only to a small extent, as shown in Figure 3, since it is not trivial to describe a drawing style in words.

Our approach is based on the assumption that adding more control to visual content creation systems is a way to enable users of the systems to feel that they have artistic ownership



Figure 2: The top row shows the style image used to generate the following images. The second row used the text prompt “Albert Einstein dancing” and the last row, “A sheep wearing a top hat.”

and control of the output. It has been reported that humans, in order to feel independent and take ownership, prefer having control to being served by a fully autonomous system [Selvaggio *et al.*, 2021; Bretan *et al.*, 2016]. If a machine performs a task for a human user without the human’s input, the user may be satisfied with the product but feel that they did not participate enough in the process to feel that they own the artistic aspects of the output as in the case for professional animators [Bateman, 2021]. In this context, we propose an approach where a user can specify a style using an example image in addition to text descriptions.

The task of machine style transfer has long been studied [Gatys *et al.*, 2016; Kolkin *et al.*, 2019] where a style is applied to a given content image. This leads to an intuitive decoupled approach using existing technologies in succession, e.g., CLIPDraw generates the drawing from text, then a style transfer algorithm styles the drawing using a given style image. While this method does combine style and content, the combination is decoupled.

In art, style (which is often referred to as “form”) and content are generally linked. Style and content could be applied agnostic to each other, but generally one informs the other. [Robertson, 1967] explains that the same content represented with different styles creates distinct artworks with varying messages or meanings, and that only a small, abstract resemblance of the content exists between each of the different styled artworks; therefore, the content, meaning, and message of an artwork is contingent on the style choices for that work. Following this theory, our approach, StyleCLIPDraw, generates drawings using language and style as input with the prior knowledge that style and content are coupled in the drawing process.

This paper includes the following contributions: 1) we propose a coupled text-and-style-to-image generation approach; 2) we release our full human evaluation results and 352 AI-generated drawings publicly as the first dataset of its kind for evaluating the style of images thoroughly¹; 3) we open

¹<https://www.kaggle.com/pittsburghskeet/>



A turtle wearing a top hat A watercolor painting of a turtle wearing a top hat A line drawing of a turtle wearing a top hat

Figure 3: Controlling the style of CLIPDraw generated drawings by altering the text description.

source StyleCLIPDraw with a public Google Colab notebook that allows people to use it easily for free² ; 4) An online demonstration that does not require programming knowledge is available³.

2 Related Work

2.1 Text-to-Image Synthesis Models

The most general and successful methods for text-to-image synthesis recently either involve training a model in a supervised fashion to perform the task or using a pre-trained image generation model with CLIP [Radford *et al.*, 2021] to guide generation. Examples of the former include Dall-E [Ramesh *et al.*, 2021] and GLIDE [Nichol *et al.*, 2021] which were trained using roughly 400 million image-text pairs. Training these models on such a huge dataset took an incredible amount of resources. Neither Dall-E nor GLIDE have been released publicly in full.

2.2 CLIP-Guided Text-to-Image Synthesis

A more computationally efficient method of text-to-image synthesis is creating one text-to-image translation at a time

drawings-with-style-evaluation-styleclipdraw

²<https://github.com/pschaldenbrand/StyleCLIPDraw>

³<https://replicate.com/pschaldenbrand/style-clip-draw>

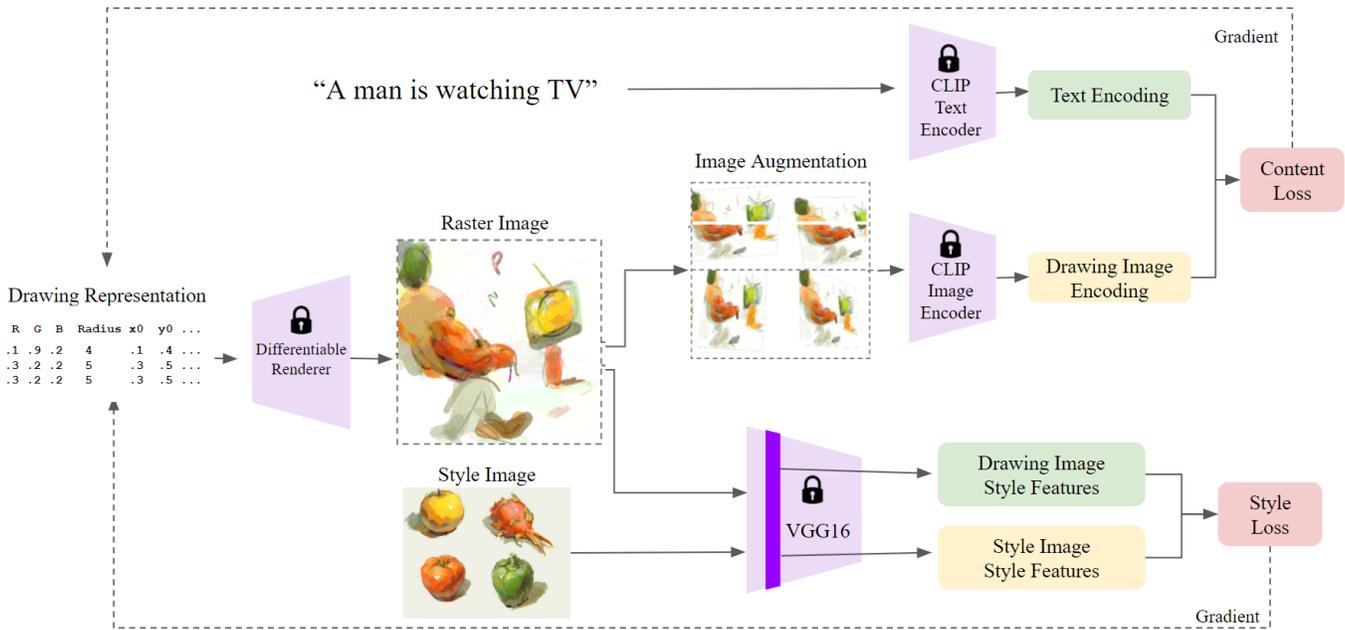


Figure 4: StyleCLIPDraw optimizes a drawing representation by computing two losses: one for content using the text description and the other for style using a style image. The drawing representation is rasterized, then style and content features are extracted using CLIP and VGG16 models respectively. Style and content features are compared using distance functions to compute a loss.

and utilizing pre-trained models. Any pretrained image generation model can be used as the image generator. A latent vector used as input to the generator model is randomly initialized and acts as the parameter space to be optimized. CLIP encodes the image generated from the latent space as well as the text description given by a user. The latent space is optimized such that the cosine distance between the generated image’s encoding and the text encoding is minimized. This method is very popular in the technological art community now and has been implemented many times with different pre-trained models [Galatolo. *et al.*, 2021; Smith and Colton, 2021].

2.3 Text-to-Drawing Synthesis

Previously mentioned methods generate raster images directly. In this work, we focus on drawings which are represented by strokes that can be rendered into raster images. Prior work with drawings often focuses on decomposing images into brush strokes, referred to as stroke-based rendering (SBR). More recent methods of SBR train deep reinforcement learning or transformer models to generate the brush strokes in a probabilistic fashion [Liu *et al.*, 2021; Schaldenbrand and Oh, 2020; Huang *et al.*, 2019].

CLIPDraw is a method that combines the action space of SBR and the methodology of CLIP-guided image generation. Rather than optimizing a model to perform every text-to-drawing synthesis task imaginable, each translation is optimized separately. So instead of modifying the parameters of a neural network model, the brush strokes themselves are modified to optimize the objective.

2.4 Drawing and Style Evaluation

In related work to altering the style of images such as style transfer [Gatys *et al.*, 2016; Kolkin *et al.*, 2019], style is not properly defined or it is defined only as relating to the color and texture of an image leaving out important concepts such as shape and space. Artists usually use art features, such as color, line, texture, and shape [Thorson, 2020], to describe artworks. In the computer vision area, research shows that utilizing these features to describe images could positively affect image retrieval [Alsmadi, 2020] and image classification [Banerji *et al.*, 2013]. These features can be used as strong representatives of the style of an image, therefore some are used as evaluation methods for artworks [Kim *et al.*, 2022; Kim, 2010]. In this work, we define the style of a drawing as a combination of the 7 elements of art drawn from [Thorson, 2020] and [Beck, 2014] which are each described in Table 1. Note that Form is not relevant to drawings as drawings are two dimensional images.

3 Method

Our approach is depicted in Figure 4. Two losses are computed to optimize for both content and style in a coupled manner at each iteration. We build on the drawing representation and the content loss methodology of CLIPDraw [Frans *et al.*, 2021] for which we include brief descriptions for self-containment.

3.1 Drawing Representation

Drawings are represented using a series of virtual brush strokes which each have parameters for a trajectory, color, and width as was defined in [Frans *et al.*, 2021]. The trajectories are described by cubic Bézier curves consisting of four

control points on a two dimensional coordinate plane. The color of the stroke is represented by four real numbers defining the amount of red, green, and blue in the color along with an alpha channel controlling the opacity of the stroke. The widths of the brush strokes are represented by the radius of the brush stroke given in pixel units.

The drawing representation cannot be utilized directly by models like CLIP or VGG16 which require raster images as input. To rasterize the drawing representation, we employ the DiffVG [Li *et al.*, 2020] package which renders vector graphics in a differentiable manner.

3.2 Drawing Augmentation

During optimization, minor flaws in the pre-trained model can be exploited, resulting in local minima that have high scores but poor results that are sometimes referred to as adversarial examples [Rebuffi *et al.*, 2021]. To stabilize optimization, the rasterized drawings are augmented with random crops and perspectives prior to being fed to CLIP as in [Frans *et al.*, 2021; Tian and Ha, 2021].

3.3 Content Loss

The content loss is computed following the methodology of [Frans *et al.*, 2021]. The input text description is the primary control for the content of the image. CLIP is used to connect the text description to the drawing representation. The drawing representation (x) is rasterized, augmented, and then encoded using CLIP’s image encoder (Eq. 1). The given text description is encoded by CLIP’s text encoder. These two encodings are compared using cosine similarity (S_C in Eq. 2) and averaged to form the loss function for the model which guides the drawings to resemble the text description. We multiply the cosine similarity by -1 because our optimization algorithms are designed to decrease the loss.

$$\begin{aligned} Drawing &= \text{DiffVG}(x) \\ Enc_{drawing} &= \text{CLIP}(\text{Augment}(Drawing)) \\ Enc_{text} &= \text{CLIP}(text) \end{aligned} \quad (1)$$

$$\mathcal{L}_{content} = -S_C(Enc_{drawing}, Enc_{text}) \quad (2)$$

3.4 Style Loss

The style of the generated drawing is specified through a given example style image. The style of the example image and generated drawing are made up of the texture, use of space, shapes of objects, lines, colors, and value of the colors in the images. To extract this style information from raster images, it is common to use pretrained neural networks [Gatys *et al.*, 2016; Kolkin *et al.*, 2019]. We extract style features based on the approach from STROTSS [Kolkin *et al.*, 2019]. Features are extracted from early layers of the VGG16 object detection model from both the drawing and the style image. Early layers of the model tend to extract low level image features such as texture and color [Chandrarathne *et al.*, 2020], so these features are known to correlate well with the style information that humans would agree on. We compare these features with Earth Mover’s Distance to compute a loss term for style which can be back propagated.

3.5 Optimization Algorithm

The drawing representation is randomly initialized to begin the optimization algorithm with a controllable number of brush strokes. As seen in Figure 4, the drawing is rasterized and used to compute losses for content and style. The objective of the optimization algorithm is to find the drawing representation (\hat{x} in Eq. 3) such that the weighted sum of the content ($\mathcal{L}_{content}$) and style (\mathcal{L}_{style}) losses are minimized. The weights, λ_1 and λ_2 , control how strong the text or style image should influence the appearance of the drawing. We report on the impact of these two weights in Section 5.1.

$$\hat{x} = \min_x [\lambda_1 \mathcal{L}_{content} + \lambda_2 \mathcal{L}_{style}] \quad (3)$$

In every iteration of the optimization algorithm, the losses are computed, the derivatives are calculated with backpropagation, and the drawing representation is altered to decrease the losses. The modification of the drawing representation is performed by the RMSProp algorithm. There is a separate instance of RMSProp for the stroke trajectories, stroke widths, and stroke colors since different learning rates (e.g., 0.3, 0.3, and 0.03 respectively) need to be used because of the varying magnitudes of the data.

The two loss terms are combined and optimized in one concerted backpropagation step as they are added together in Eq. 3. Alternatively, each loss can be optimized separately in an alternated fashion where content is optimized for some number of iterations followed by style optimized for several iterations. We call these approaches “concerted optimization” and “alternated optimization” respectively and discuss their efficiency in Section 5.2.

4 Experiments

Settings: All experiments were conducted using our publicly available notebooks on Google Colab.

4.1 Decoupled Baseline

Styled text-to-drawing synthesis can be conducted using existing technologies chained together in sequence. This formed the baseline in this work: The language input is converted into a drawing using CLIPDraw, then the drawing is styled with the style image using the STROTSS style transfer algorithm as a post-process.

4.2 Human Evaluation Survey

Main experiment: We hand crafted⁴ 22 text prompts and chose 8 different style images to generate 176 drawings for the StyleCLIPDraw and baseline methods. For each style image and prompt pair, we generated 4 drawings and selected the one with the best cosine similarity between CLIP encodings of the drawing and text description. We used the Amazon Mechanical Turk (AMT) platform to conduct the study. The 139 participants were first educated on what the different dimensions of art are and what they mean (Table 1). The participants were then shown the style image that was used to

⁴On a pilot study, prompts from the Flickr image caption dataset [Plummer *et al.*, 2015] were found to be too complex.

Element Description of Art	
Lines	Outlines capable of producing texture with their
Shapes	General geometry and makeup of objects within the drawings.
Space	Perspective and proportion between shapes and objects.
Texture	The way things look as if they might feel if touched.
Value	The Lightness or darkness of tones or colors.
Color	The hue, lightness, darkness, and intensity of color in a drawing.
Form	The three dimensional properties of the work (not used in the evaluation because it is less relevant to 2D images).

Table 1: The seven elements of art that make up the style of an artwork.

generate the drawings, and the two drawings: one from StyleCLIPDraw and the other from the baseline in a random order. There were 9 questions asked about each pair of drawings. The first 6 were for each style element in Table 1, i.e., “Which drawing is most similar to the Style Image in terms of [style element]?” The participants were also asked to pick which drawing overall fits the style of the style image best. Next, the text prompt used to generate the drawings was given, and the participants were asked to pick which drawing best fits the description, if neither fit at all, or if both fit it equally well. Lastly, the participants were asked which drawing (if any or both) fits the text *and* the style best.

Auxiliary questions: To compare the realism of the style images used in the study, we asked the AMT workers to rate the 8 style images used in our study in 5 categories: very simple, simple, neutral, detailed, and very detailed.

5 Results

We report on the findings from the analysis of our approach and then present the evaluation results.

5.1 Style and Content Weights

We qualitatively report on how the quality of the output drawings is impacted by the choice of the style and content weights in the proposed optimization method in Equation 3. We generated drawings using the same text prompt and style image and varied the style and content weights. From left to right in Figure 5, drawings were generated with decreasing style loss weight and increasing content loss weight (λ_2 and λ_1 respectively in Eq. 3). With exclusively style weight, the drawings capture the style well but lack resemblance to the text description, while with exclusively content weight the drawings represent the text well but lack style from the style image. Using these qualitative results we chose 1.0 for both λ_1 and λ_2 for the remainder of the paper for the balance between content and style strength.

5.2 Concerted versus Alternated Optimization

Based on an AMT study designed similar to the main survey described in Section 4.2 on 51 users, the concerted and al-

Comparison Dimension	Ours	Baseline	Neither	Both
Lines	75.4	24.6		
Shapes	76.6	23.4		
Space	82.1	17.9		
Texture	77.6	22.4		
Value	77.4	22.6		
Color	80.8	19.2		
Style Overall	84.9	15.1		
Content	24.3	52.1	9.1	14.4
Style & Content	48.1	23.3	19.8	8.9

Table 2: Percentage of preference of drawings in direct comparison between our StyleCLIPDraw approach and the baseline. More details in Section 4.2.

ternated optimization approaches are comparable in terms of content and style generation although a minor advantage was shown for alternated optimization in style generation; however, the concerted optimization method is drastically superior in terms of computational efficiency. Specifically, the runtime increase in the alternated optimization is ten fold longer than the concerted optimization and only two fold more preferred in terms of style. The concerted optimization approach takes about 1 minute with a P100 NVIDIA GPU on our Google Colab notebook which is promising news towards developing a real-time system.

5.3 Qualitative Results

In Figure 1, we compare the performance of StyleCLIPDraw with the original CLIPDraw [Frans *et al.*, 2021] and the decoupled baseline approach: CLIPDraw + Style Transfer (Section 4.1). Four prompts and style image pairs are used to generate drawings. On the content generation, all three approaches appear to fit the content of the given text description well. On the style generation, our coupled approach is able to generate the output images faithfully according to the given style input. By contrast, the decoupled baseline of applying style transfer as a postprocess to the CLIPDraw drawings alters the colors and textures of the images but lacks adaptation of other art elements such as shape, space, and lines.

Additional examples of how StyleCLIPDraw responds to different prompts and style images are shown in Figure 2.

We note that, due to the lack of abstract style, the CLIPDraw drawings appear to capture the content of the text more clearly and literally. Certain styles of drawings are more constraining than others. For instance, the more painterly style images appear to give the approach more flexibility with the output compared to the highly abstract line drawing examples. Since StyleCLIPDraw applied the style far more than the baseline, it could be true that the style hurt the appearance of the content of the image more than the baseline. Using the auxiliary questions in our survey, we found that there is a strong, positive correlation (0.72 with a p-value of 0.04) between the divided opinion on realism of the style image and the appearance of content in the image, i.e., the more ambiguous the realism of the style image is, the less preferred the generated output’s content is.



Figure 5: From left to right: increasing the weight of the content loss and decreasing the weight of the style loss (λ_1 and λ_2 respectively in Eq. 3). The drawings were generated using the style image on the far left and the text prompt, “A person is walking down a city street.”

5.4 Quantitative Results

The results from the human evaluation study based on 139 participants are summarized in Table 2. As we speculated in the qualitative analysis, the decoupled method generated clearer contents that are easier for human users to recognize. On the style aspect, according to all style dimensions in Table 1, our StyleCLIPDraw approach was vastly preferred to the decoupled approach. This result provides strong evidence for our assumption that style and content generation is a coupled process and that the two objectives should be jointly optimized throughout the creation process.

Overall, when asked about both style *and* content, the participants vastly preferred our approach over the baseline, underscoring the importance of style, look, and feel of the generated images. This result also provides a promising signal for our assumption that an additional control over the output can potentially improve the user’s artistic ownership.

5.5 Limitations

Our approach relies on existing image-text models such as CLIP that are generally trained over a gigantic amount of non-public data. Based on our experiments, the complexity of text descriptions that existing models can handle is relatively simple as seen in the examples used in the paper.

As discussed in the results, using abstract style images seems to impede the content quality. Creating recognizable content using an abstract style generally requires more creativity. For a future direction, we consider including additional loss to also optimize for recognizability or gestalt.

Our current method is still not fast enough to be used in real-time applications such as communication assistance.

6 Conclusions

This paper addresses a question of how AI can be used to empower laypeople to engage in creative visual content generation. Building upon the assumption that adding more control by way of human input to an AI system can enable users to gain more fulfilling ownership and control over the output, we propose an approach for image generation where users can

have an example-based stylistic control as well as a text-based content control. Following the theory in art that the style of the drawing is intertwined with its content, we propose StyleCLIPDraw, a coupled approach where style and content are jointly optimized at every iteration from start to end.

Based on 139 random human evaluators, StyleCLIPDraw was preferred to the decoupled approach 84.9% of time in terms of various style elements, matching the underlying theory of our approach. Notably, despite the fact that it was less preferred in content solely, when considering both style *and* content, StyleCLIPDraw drawings were highly preferred. This result is an indication that style is as much as or even more important than the content of the image to people’s perceptions of generated images.

As part of contributions, we open source our approach to promote future research on this direction. Our online demonstration provides a tool for non-technical people to also benefit from our work. We also make the results from the study available publicly that include 352 digital drawings with comparisons across style as a whole, individual elements of style, and content. To our knowledge, this is the first public dataset with information about the quality of drawings with respect to the individual elements of art and can contribute to future automatic methods of evaluating the style of generated images. Future work on StyleCLIPDraw includes adding more artistic control in the form of compositional input and improving the runtime to near real-time.

Acknowledgements

We would like to thank Benjamin Stoler for meaningful discussions.

References

- [Alsmadi, 2020] Mutasem K Alsmadi. Content-based image retrieval using color, shape and texture descriptors and features. *Arabian Journal for Science and Engineering*, 45(4):3317–3330, 2020.

- [Banerji *et al.*, 2013] Sugata Banerji, Atreyee Sinha, and Chengjun Liu. New image descriptors based on color, texture, shape, and wavelets for object and scene image classification. *Neurocomputing*, 117:173–185, 2013.
- [Bateman, 2021] Cole Bateman. Creating for creatives: A humanistic approach to designing ai tools targeted at professional animators. *Bachelor's thesis, Harvard College*, 2021.
- [Beck, 2014] Paula D Beck. *Fourth-Grade Students' Subjective Interactions with the Seven Elements of Art: An Exploratory Case Study Using Q-methodology*. Long Island University, CW Post Center, 2014.
- [Bretan *et al.*, 2016] Mason Bretan, Deepak Gopinath, Philip Mullins, and Gil Weinberg. A robotic prosthesis for an amputee drummer. *arXiv preprint arXiv:1612.04391*, 2016.
- [Chandrarathne *et al.*, 2020] Gayani Chandrarathne, Kokul Thanikasalam, and Amalka Piniidiyaarachchi. A comprehensive study on deep image classification with small datasets. In *Advances in Electronics Engineering*, pages 93–106. Springer, 2020.
- [Frans *et al.*, 2021] Kevin Frans, LB Soros, and Olaf Witkowski. Clipdraw: Exploring text-to-drawing synthesis through language-image encoders. *arXiv preprint arXiv:2106.14843*, 2021.
- [Galatolo *et al.*, 2021] Federico Galatolo., Mario Cimino., and Gigliola Vaglini. Generating images from caption and vice versa via clip-guided generative latent space search. *Proceedings of the International Conference on Image Processing and Vision Engineering*, 2021.
- [Gatys *et al.*, 2016] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016.
- [Huang *et al.*, 2019] Zhewei Huang, Wen Heng, and Shuchang Zhou. Learning to paint with model-based deep reinforcement learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8709–8718, 2019.
- [Kim *et al.*, 2022] Diana Kim, Ahmed Elgammal, and Marian Mazzone. Formal analysis of art: Proxy learning of visual concepts from style through language models. *arXiv preprint arXiv:2201.01819*, 2022.
- [Kim, 2010] Seong-in Kim. A computer system for the analysis of color-related elements in art therapy assessment: Computer_color-related elements art therapy evaluation system (c_creates). *The Arts in psychotherapy*, 37(5):378–386, 2010.
- [Kolkin *et al.*, 2019] Nicholas Kolkin, Jason Salavon, and Gregory Shakhnarovich. Style transfer by relaxed optimal transport and self-similarity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10051–10060, 2019.
- [Li *et al.*, 2020] Tzu-Mao Li, Michal Lukáč, Michaël Gharbi, and Jonathan Ragan-Kelley. Differentiable vector graphics rasterization for editing and learning. *ACM Transactions on Graphics (TOG)*, 39(6):1–15, 2020.
- [Liu *et al.*, 2021] Songhua Liu, Tianwei Lin, Dongliang He, Fu Li, Ruifeng Deng, Xin Li, Errui Ding, and Hao Wang. Paint transformer: Feed forward neural painting with stroke prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6598–6607, 2021.
- [Nichol *et al.*, 2021] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021.
- [Plummer *et al.*, 2015] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649, 2015.
- [Radford *et al.*, 2021] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. *CoRR*, abs/2103.00020, 2021.
- [Ramesh *et al.*, 2021] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. *arXiv preprint arXiv:2102.12092*, 2021.
- [Rebuffi *et al.*, 2021] Sylvestre-Alvise Rebuffi, Sven Gowal, Dan Andrei Calian, Florian Stimberg, Olivia Wiles, and Timothy A Mann. Data augmentation can improve robustness. *Advances in Neural Information Processing Systems*, 34, 2021.
- [Robertson, 1967] Duncan Robertson. The dichotomy of form and content. *College English*, 28(4):273–279, 1967.
- [Schaldenbrand and Oh, 2020] Peter Schaldenbrand and Jean Oh. Content masked loss: Human-like brush stroke planning in a reinforcement learning painting agent. *arXiv preprint arXiv:2012.10043*, 2020.
- [Selvaggio *et al.*, 2021] Mario Selvaggio, Marco Cognetti, Stefanos Nikolaidis, Serena Ivaldi, and Bruno Siciliano. Autonomy in physical human-robot interaction: A brief survey. *IEEE Robotics and Automation Letters*, 2021.
- [Smith and Colton, 2021] Amy Smith and Simon Colton. Clip-guided gan image generation: An artistic exploration. *Evo* 2021*, page 17, 2021.
- [Thorson, 2020] Mark Thorson. Formal elements of art. *Encounters With the Arts: Readings for ARTC150*, 2020.
- [Tian and Ha, 2021] Yingtao Tian and David Ha. Modern evolution strategies for creativity: Fitting concrete images and abstract concepts. *arXiv preprint arXiv:2109.08857*, 2021.