

# Segmentation-Based Parametric Painting

Manuel Ladrón de Guevara  
Carnegie Mellon University  
rldg.manuel@gmail.com

Matt Fisher  
Adobe Research  
matfishe@adobe.com

Aaron Hertzmann  
Adobe Research  
hertzman@dgp.toronto.edu




Figure 1. Comparison of our method against previous work on parametric painting algorithms. Panels (a) through (c) demonstrate our algorithm’s capability to render scenes in various styles: abstract with simplified forms and abstract with a heavy emphasis on thick loose strokes, texture, and a lack of clear form (a), painterly impressionistic style with dynamic strokes and emphasis on effect (b), and realistic style capturing both micro and macro structures with high fidelity (c). These are juxtaposed against existing methods: Stylized Neural Painting (SNP) [40] (e), Paint Transformer (PT) [23] (f), and Learning to Paint (L2P) [15] (g), which exhibit limitations such as imprecise edge depiction and a lack of stylistic diversity. SNP and PT use rigid, blocky strokes that are unable to capture fine details, edges and shapes. L2P captures details better, although yields blurry results with visible seams at large resolutions. The input image is displayed in panel (d) for reference. This comparison underscores our method’s superior ability to balance detail retention with stylistic expression. Zoom-in is encouraged to see the nuances of each work.

## Abstract

We introduce a novel image-to-painting method that facilitates the creation of large-scale, high-fidelity paintings with human-like quality and stylistic variation. To process large images and gain control over the painting process, we introduce a segmentation-based painting process and a dynamic attention map approach inspired by human painting strategies, allowing optimization of brush strokes to proceed in batches over different image regions, thereby

*capturing both large-scale structure and fine details, while also allowing stylistic control over detail. Our optimized batch processing and patch-based loss framework enable efficient handling of large canvases, ensuring our painted outputs are both aesthetically compelling and functionally superior as compared to previous methods, as confirmed by rigorous evaluations. Code available at: [https://github.com/manuelladron/semantic\\_based\\_painting.git](https://github.com/manuelladron/semantic_based_painting.git)*

## 1. Introduction

Painting is the world’s oldest technology for creating visual imagery, and a continuing source of fascination and artistic expression. Representational painting is a complex process where an artist creates a picture of a realistic scene from a series of brush strokes. Artists use a variety of strategies to determine where and how to paint, and different strategies produce different painting styles. For example, a painter might work in a coarse-to-fine manner, beginning with general shapes and then refining details, or in a semantic-region-based manner, painting areas such as sky, mountains, and buildings sequentially; the latter strategy may produce a more segmented style than the former, with different styles for different objects.

These strategies, in part, are a function of human working memory and foveal vision: unlike optimizers, artists cannot simultaneously work on every part of a picture in parallel. They work piece-by-piece, focusing on one region or aspect of a painting at a time [3, 4, 7, 28, 35, 36], before moving on to the next.

The computational *stroke-based painting* problem is: given an input photograph, generate a painting to represent the photograph, parameterized by a set of brush strokes, which themselves are parameterized by their colors, control curves and thicknesses, with optional texturing [10, 13]. These methods allow styles to be defined directly in terms of brush stroke properties, with further additional user control possible on the strokes themselves, e.g., [27]. Current methods treat this as an optimization problem: using a limited set of brush strokes, approximate the input image as closely as possible [5, 12, 15, 23, 25, 31, 40]. By varying the level of approximation, a user should be able to select from more- or less- abstract styles, and selecting different placement strategies should produce qualitatively different styles. Moreover, robotic painting systems require explicitly-defined brush strokes, e.g., [1, 8, 21, 31], which are not available with pixel-based stylization techniques.

Despite recent progress, existing methods for this problem suffer from a limited ability to perform this optimization well for large images and limited user control. State-of-the-art methods either jointly optimize a large collection of strokes, or train a network to generate them. These approaches are limited by the problem of generating strokes uniformly for an input image, which becomes impractical as the input size increases. Hence, these methods do not generate high-fidelity results for even moderate-sized input images, thereby limiting the level of abstraction and control [15, 23, 34, 40]. Moreover, they lack mechanisms to mimic varied artistic strategies, such as applying different styles for different scene elements, and coarse-to-fine painting.

Inspired by these observations, we propose a new approach to generating stroke-based paintings from photographs. We first structure the entire painting process in

layers by isolating different semantic parts of the canvas. Then dynamic attention maps guide the painting process to refine areas that need more detail, and differentiable strokes are optimized sequentially within each region. This sequential optimization process allows us to generate strokes over a large image canvas efficiently, in contrast to methods that produce pictures as a single joint optimization. It also approximates human painting strategies that operate sequentially on different image regions, rather than jointly optimizing all strokes at once. Several objective function parameters are provided to control the style, including limiting the number of strokes and level of detail for each semantic region, providing a continuum from realism to abstraction.

Quantitative and qualitative evaluations show that our method achieves a greater range of high-fidelity reconstructions to abstract imagery, while also producing high-quality results at large image sizes.

## 2. Related Work

Early stroke-based-rendering (SBR) approaches varied from procedural, rule-based systems without optimization [10, 22], including layering strokes in a coarse-to-fine fashion. Procedural methods effectively produce captivating paintings, but their rigidity in stroke placement is a notable limitation. In principle, optimization-based SBR overcomes this constraint, offering greater adaptability in stroke positioning and more accurately encoding artist choices [12]. Such methods have evolved from heuristic-based optimization to more advanced techniques such as EM-like packing algorithms for arranging non-overlapping stroke primitives [13, 29], with applications extending to 3D model stylization [24, 32]. Recent SBR methods employ deep networks trained to simulate a painting agent through reconstruction losses, involving pixel and perceptual losses [17, 38], facilitated by differentiable rendering engines [20]. However, the difficulty of optimizing large collections of strokes over an entire image remains challenging.

Reinforcement Learning (RL) has become prominent in SBR for simulating decision-making in painting agents [5, 15, 16, 25, 30, 34, 37], with diverse applications ranging from detailed reconstructions to robotic painting. SPIRAL [5] and its enhanced version SPIRAL++ [25] are noteworthy for their adversarial training algorithms, despite producing less defined images. Nonetheless, both methods struggle with interpretability, style control, and accurate input reconstructions. Some RL methods focus on accurate depiction [15, 34]. Huang *et al.*’s method [15] needs thousands of tiny strokes to reconstruct an image, although becomes blurry for high resolution images, and it is stylistically limited.

Deep learning methods like Paint Transformer [23] utilize a CNN-Transformer architecture to predict stroke sequences, showing promise in generalization despite pattern

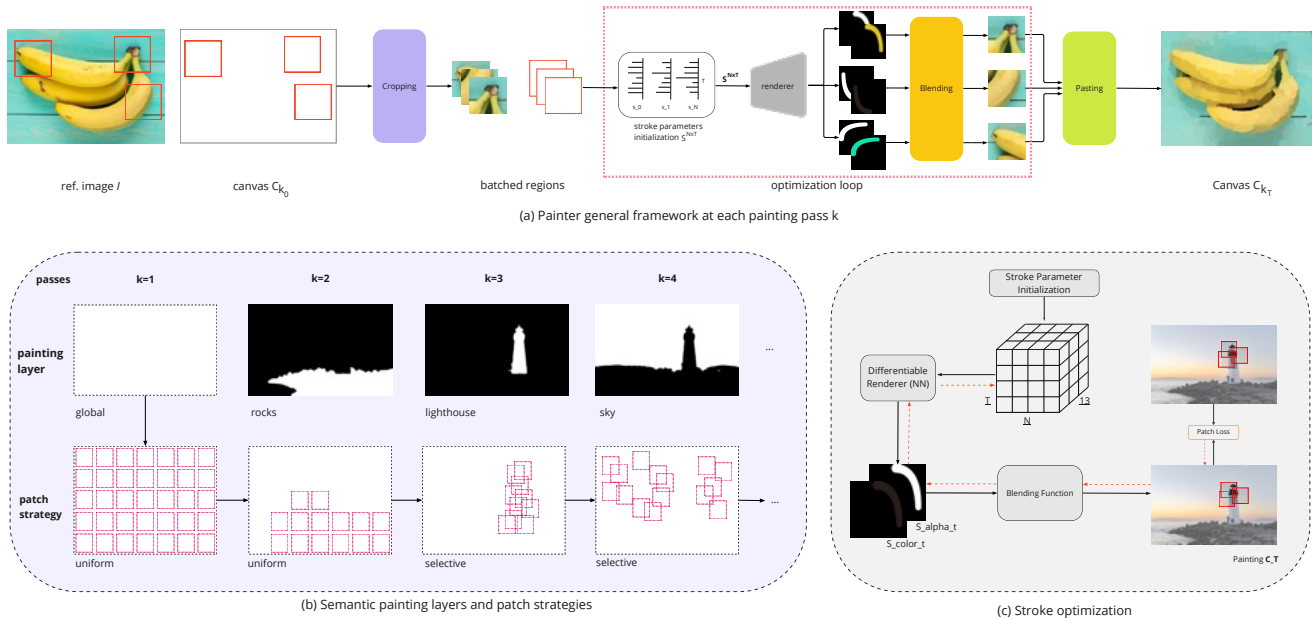


Figure 2. Overall Painting Framework. Our framework models the painting process as a set of sequential painting passes. (a) Pipeline of each painting pass  $k$ , which is composed by four parts: attention patches, cropping, optimization, and pasting. (b) Semantic layers and patch strategies. (c) Stroke optimization module with patch-loss.

repetition in results. RNN-based models have also been explored for distinct line drawing and sketching [9, 19, 26, 39], with our method adopting similar attention mechanisms [26].

Additionally, recent works have leveraged direct differentiable optimization for stroke placement [20, 40], with style variations often induced by style-transfer techniques in pixel space [6, 17].

### 3. Method

Our goal is to efficiently optimize a collection of stroke parameters that define a large-scale painting from a given image, capturing micro and macro structures with high fidelity, and providing control for stylistic variation. Our system takes as input an image  $X$ , and style parameters, encapsulated by three style presets: “realistic,” “painterly,” and “abstraction.” Then, our objective is to optimize a parameter vector  $\bar{Y}$  of strokes, which, when rendered on a canvas  $C$ , will approximate the given image  $X$  as a stylized painting. Our scalable solution accommodates varying sizes and aspect ratios through batch-patch optimization, cropping the input and canvas into a set of  $128 \times 128$  patches, which are reassembled after the optimization process to match the original image scale.

Inspired by human painting techniques, and to provide more control over the painting process, we first use a segmentation network that decomposes the painting into semantic areas, and then employ dynamic attention maps to

determine where to paint. An overview of our framework is shown in Figure 2. For information on the parametrization of styles and implementation details, refer to supplemental.

#### 3.1. Layered Painting and Patch-Based Optimization

The set of paint strokes is separated into  $K$  layers, one for each input semantic region. Applying separate painting styles or details for different regions allows stylistic control over separate scene elements. For each painting layer, the canvas is divided into a set of  $128 \times 128$  patches, where all stroke parameters in each patch are batch-optimized; different layers may have different numbers of patches, as a function of the semantic layer’s area.

The optimization is performed in a series of  $P$  passes, designed to progress from general, broader strokes to intricate, finer details, akin to the approach of a traditional artist [11]. For each painting pass  $p$ , our method applies an optimization phase to each of the  $K$  semantic layers, following a coarse-to-fine procedure.

Let  $P$  denote the total number of painting passes, and let’s denote the number of semantic layers as  $K$ . We define  $k_p$  as the layer index being refined in pass  $p$ , where  $k_p \in \{0, 1, 2, \dots, K - 1\}$ . Overall, our goal is to minimize a reconstruction loss  $L_{rec}$  summed over the entire image. We specify this loss later in the paper. At a given pass, the goal is to optimize just the strokes that appear in some layer  $k$ . The objective for each layer  $k$  during optimization is to optimize a subset of the loss just for the patches and strokes

in this layer:  $\arg \min_{\bar{y}_k} \sum_{i=1}^{N_k} \mathcal{L}_{rec}(X_{k,i}, C_{k,i}(\bar{y}_k))$ , where  $X_{k,i}$  is the  $i$ -th patch of the semantic region  $k$  of the input image, and  $C_{k,i}(\bar{y}_k)$  denotes the  $i$ -th patch of the canvas at layer  $k$ , parameterized by the stroke parameters  $\bar{y}_k$ . The number of layers  $K$  adapts to each painting based on the unique semantic content of the input image, and it is determined automatically by the segmentation network. In cases where the input image lacks semantic distinction, or if semantic assistance is disabled, a single layer approach is adopted ( $K = 1$ ), involving the entire canvas. We use the DETR model [2], which consists of a CNN (ResNet) backbone followed by an encoder-decoder Transformer, trained on the COCO Panoptic Segmentation task.

Starting with the input image  $X$ , we derive all segmentation layers and compute each binary mask, which is then used to initialize the vector of stroke parameters  $\bar{Y}_{k_p}$ . Upon the completion of the optimization steps, strokes that fall outside the designated segmented area are filtered out based on a threshold  $\gamma$ .

**Dynamic Attention Maps.** While painting, any human painter must continually decide where next to paint brush strokes, and different painters will make these choices differently. Our system represents this process by dynamic attention maps, with two distinct strategies: uniform, and selective (see Figure 2(b)). The former is based on a uniform distribution across the canvas, yielding patches in a consistently-spaced grid covering the entire image. The latter identifies the  $V$  patches with the highest error for attention, where  $V$  is a user-determined style parameter. This error is computed as  $L_1$  loss between the canvas and image patches. We found that just using  $L_1$  loss is enough to provide a good painting guidance, making the process more efficient in comparison to using perceptual losses.

The uniform approach ensures an unbiased and uniform distribution of detail throughout the painting, and is typically reserved for the foundational stages of a painting, generally producing higher realistic painting than the selective approach, which generates a more human-like, organic painterly style paintings.

**Stroke Initialization, Renderer, and Blending.** We use the stroke parameter representation and differentiable renderer,  $G$ , provided by [15] (for more details refer to supplemental). Each stroke is parameterized by a 13-dimensional tuple  $\{x_0, y_0, x_1, y_1, x_2, y_2, r_0, r_2, t_0, t_2, R, G, B\}$  that encodes the start, middle and end points of a quadratic Bézier curve, radii and transparency at the start and end points, and RGB color of the stroke. For each  $128 \times 128$  patch in the canvas, we distribute  $T$  stroke parameters evenly in a regular array.

Strokes are composited into a patch with soft blending:  $M \leftarrow M_{t-1} \odot (1 - G(s_j)_a) + G(s_j)_c$  where  $s_j$  are the pa-

rameters of the  $j$ -th stroke,  $M$  is a canvas patch, and  $G(s_j)_a$  and  $G(s_j)_c$  are the alpha channel and the stroke given by the differentiable renderer  $G$ , respectively. However, for semantic layers, the corresponding binary mask  $\alpha_k$  is used:  $M_k \leftarrow M_{k_{t-1}} \odot (1 - G(s_j)_a) + G(s_j)_c \odot \alpha_k$

### 3.2. Optimization and Loss Functions.

At a given pass  $p$ , for each layer  $k$  in such a pass, we optimize all stroke parameters from all  $N_k$  patches in batch. That is, we optimize a matrix  $S_k^{N \times T \times 13}$  altogether, where  $T$  is the stroke budget, and instead of computing the loss function over the entire canvas, we use a patch loss described earlier (see Figure 2 (c)). This allows us to perform these optimizations much more efficiently than optimizing over the entire canvas. Formally, given a set of canvas patches and the same set of reference image’s patches, our overall loss function is:

$$\mathcal{L}_{rec} = \frac{1}{N} \sum_{n=0}^N \alpha \mathcal{L}_{\text{pixel}_n} + \beta \mathcal{L}_{\text{perc}_n} \quad (1)$$

where  $N$  is the total number of patches. For our pixel loss, we use the  $L_1$  function,  $\mathcal{L}_{\text{pixel}} = |C - I|$ , where  $C$  is a canvas and  $I$  a reference image. For perceptual loss, we use the feature vectors from a pretrained VGG16 [38]. Specifically, let  $V_{ij} = \{V_{ij}^1, \dots, V_{ij}^k\}$  and  $W_{ij} = \{W_{ij}^1, \dots, W_{ij}^k\}$  be a set of  $k$  feature vectors extracted from image  $I$  and canvas  $C_T$ , respectively. We use cosine similarity as follows:

$$\mathcal{L}_{\text{perc}} = \cos \theta = \sum_k \sum_{i,j} \frac{V_i W_j}{\|V_i\| \|W_j\|} \quad (2)$$

where  $i, j$  index the spatial dimensions of the feature maps  $V$  and  $W$ , and  $K$  the extracted layers from VGG16 trained on ImageNet [33]. We use layers 1, 3, 6, 8, 11, 13, 15, 22 and 29, and optimize via gradient descent and use Adam optimizer [18] to find the set of strokes  $S_{k_p}$ .




Figure 3. Coarse-to-fine approach for a painterly portrait of a woman. (a) Initial coarse layer lays foundational strokes for a later refinement (b), guided by dynamic attention maps (insets).




Figure 4. Our approach offers enhanced precision over the canvas’s semantic regions, directing specific emphases in the artwork: (a) accentuates the foliage, capturing its essence, while (b) concentrates on highlighting the architectural details of the building.

## 4. Experiments

We now demonstrate stylistic effects achievable with our method, and compare with state-of-the-art methods.

The semantic control provided by our method allows a user to apply different styles to different regions, akin to classical artistic techniques where backgrounds might be abstracted for emphasis on detailed foregrounds and vice-versa, reminiscent of techniques observed in classical artworks [14], and specific subjects or objects may be highlighted. The versatility of this approach, resulting in varied visual interpretations based on the areas of emphasis, is illustrated in Figure 1 (a) and (c), Figure 4, Figure 5, Figure 6, and Figure 8.

### 4.1. Style Variations

Painting techniques and strategies lead to different styles, and our algorithm can output a set of stylistic outputs, ranging from realism to abstraction and painterly styles.

**Realism.** We begin by demonstrating a style that precisely reproduces the target image. Previous methods cannot achieve accurate image reproduction through optimization; this limits their ability to define new styles, since many styles cannot be accurately optimized, e.g., fine-scale details cannot be preserved. While most painterly styles involve greater degrees of abstraction, this style illustrates the superior performance of our sequential patch-based optimization, which can achieve accurate image reproduction.

For this style, we use a single layer ( $K = 1$ ), without semantic segmentation, and run four coarse-to-fine painting passes ( $P = 4$ ). The number of strokes  $s$  per patch increases with each layer, following the formula  $s_p =$

$(p + 5)^2$ . Stroke thicknesses are halved across consecutive passes, given by the formula:  $a_p = 2^{1-p}a_1$ , where  $a_p$  and  $a_1$  are the  $p$ -th and thicknesses. This approach ensures that the painting captures finer details, rendering a realistic appearance. Paintings employing this methodology are showcased in Figure 1 (c), Figure 3, Figure 5 (b), Figure 6 (c) top, and Figure 7 (fourth row). We use perceptual losses along with pixel losses to capture higher frequency areas. That is, we optimize the stroke parameters following Equation (1), where  $\alpha = 1$  and  $\beta = 0.01$ . This style achieves superior reconstruction details than previous methods as seen in Section 4.2, and in Figure 1 (c), and Figure 7 (fourth row). Analyzing these paintings, high-frequency areas are captured with higher definition than previous methods, that is, strokes follow edges and shapes more accurately, resulting in higher-fidelity paintings.

**Painterly Style.** The goal of this style is to generate a more loose representation of the input image, yet capturing enough details so that the original scene is represented. Leveraging semantic segmentation, our approach mirrors the human painting process, working piece-by-piece, focusing on one region or aspect of a painting at a time, and generating different painterly styles. Instead of a uniformly detailed canvas, semantic control offers refined governance over stroke distribution, tailored to distinct semantic zones. Our dynamic attention maps, inspired by visual working memory, are configured to a selective mode, honing in on areas necessitating enhanced detail. The stroke parameters are optimized solely based on pixel loss, that is, we set  $\beta = 0$  in Equation (1). For each painting pass  $p$ , we found enough to have a stroke budget of 16 strokes per patch, which yields a distinctive painterly style. This approach prioritizes artistic expression over mere replication, yielding a rendition that captures the essence rather than the exact likeness of the input image, as demonstrated in Figure 1 (b), Figure 4, Figure 5 (a), Figure 6 (a, bottom), (b, top), and Figure 8 (b,c). In these figures, we can see that this style has the loose quality of typical impressionistic paintings, yet forms are defined and the subject matter is clear.

**Semantic Abstraction.** Unlike the previous painterly style, which still prioritizes a good balance between accuracy and expression, this style leans with a heavier emphasis on expression. By adjusting the optimization procedure over painting passes  $P$  and semantic layers  $K$ , we can generate abstract representations, offering a novel painting synthesis. To do this, the user associates different optimization parameters such as  $P$ , stroke budget, and number of selective dynamic attention maps  $V$  to different semantic layers  $k$ . Since this style does not focus on high-fidelity, stroke parameters are optimized based on  $L_1$  loss, setting  $\beta = 0$  in Equation (1). This generates organic and abstract repre-




Figure 5. Different painterly styles given the same input. (a) Painterly style achieved by the use of our segmentation pipeline and selective dynamic attention maps. (b) Realistic style achieved by using uniform dynamic maps.



Figure 6. Demonstration of control over the painting. (a) The painting focal point falls in the buildings while leaving the vegetation abstracted. (b) The motorcycle is emphasized by adding more detail than the background, which is coarsely painted.

sentations of the object of interest, focusing on expression over representation, and can be applied universally or selectively to accentuate specific semantic zones. See how in Figure 1 (a), and Figure 8 (a), the abstract paintings loosely define the subject matter, prioritizing heavy abstract expressionism, in contrast to the painterly style, which captures the subject matter, but still leaving room to the viewer’s interpretation. A practical illustration: buildings can be rendered in finer detail to spotlight their significance, whereas adjacent elements like streets, pedestrians, or foliage are abstracted, adding an artistic bias to the scene (refer to Figure 4 and Figure 6). As a result, the painting looks more intentional, organic and fluid, similar to how human artists paint. Figure 1 (a) and Figure 8 (a) are excellent exam-

	Control	Style Var.	Large-Scale	High-low Freq.
PT	✗	✗	✓	✗
SNP	✗	✓	✓	✗
L2P	✗	✗	✗	✓
Ours	✓	✓	✓	✓

Table 1. Comparison with previous methods on attributes like painting control, style variations, large-scale parametric painting, and high-low frequency capture.

ples of the human-like qualities of this style, with brushwork being loose and gestural, creating a rich tapestry of dynamic shapes. Other examples using this style are shown in Figure 7 (bottom row), where semantic layers are entirely abstracted from their context. For further details on the parametrization of this style, please refer to supplementary material.

## 4.2. Comparison with State-of-the-Art

We evaluate our work against three notable parametric methods: an optimization-based method “Stylized Neural Painting” (SNP) [40], a learning-based method with Transformer “Paint Transformer” (PT) [23], and a RL method “Learning to Paint” (L2P) [15]. Our method stands out by supporting any resolution or aspect ratio, preserving artwork quality, unlike L2P and PT which resize outputs, often reducing quality. Moreover, SNP is limited to square formats, restricting its application. An overview of algorithmic attributes is shown in Table 1. While all methods can output paintings at high resolutions, L2P’s results are blurry and seams are visible. For a fair comparison, all paintings have been generated using the same number of strokes.

**Qualitative Comparison** A visual comparison of our method against previous methods is depicted in Figure 1, Figure 7 and Figure 8. For more comparisons at larger scale,

Method	Realism	Visual Appeal
Ours vs. PT [23]	81.82%	79.55%
Ours vs. SNP [40]	100%	70.45%
Ours vs. L2P [15]	95.45%	56.82%

Table 2. Qualitative evaluation based on user preference. This table shows a 2-way comparison between our method and previous methods.

refer to supplemental. Figure 1 shows how SNP and PT generate impressionistic-like paintings, although the brushwork is patchy and rigid, and both methods are unable to capture edges. This is further seen in Figure 7, where PT introduces highly contrasted brush sizes for different semantic regions, and noise in high-frequency areas. SNP’s rigid and uniform patches do not capture well finer details. Out of the previous methods, L2P captures micro and macro structures better, in a realistic style, although at large resolutions produces blurry results, and it is unable to generate other painting styles. This is further observed in Figure 9, where L2P cannot distinguish between semantic areas or objects, treating the entire canvas with the same emphasis.

**User Study.** We assessed our painting method’s effectiveness through a user study involving 30 participants, contrasting our outputs with three existing techniques. Participants judged the paintings based on realism and visual appeal across six painting sets, for a total of 12 questions, using a two-way comparison, and collecting 360 responses. As shown in Table 2, our method is preferred over existing methods in all tasks. Overall, our method substantially outperforms previous methods in realism, selected by 92.42% of the users. Our method is also preferred in the visual appeal task, voted by 68.94%. When evaluating our method on visual appeal, our strongest competitor is L2P, selected by 43.18% of the participants.

**Quantitative Comparison** To further substantiate our claims, we compare the performance of our model on reconstruction of the original input against previous methods, and use pixel loss  $L_1$  and perceptual loss  $L_{perc}$  as evaluation

	Ours		SNP		PT		L2P	
	$L_1 \downarrow$	$L_{perc} \uparrow$	$L_1 \downarrow$	$L_{perc} \uparrow$	$L_1 \downarrow$	$L_{perc} \uparrow$	$L_1 \downarrow$	$L_{perc} \uparrow$
Objects	<b>0.038</b>	<b>0.680</b>	0.052	0.554	0.085	0.480	0.053	0.586
Buildings	<b>0.026</b>	<b>0.720</b>	0.049	0.515	0.071	0.490	0.045	0.632
Streets	<b>0.031</b>	<b>0.709</b>	0.041	0.582	0.072	0.486	0.041	0.620
Animals	<b>0.016</b>	<b>0.745</b>	0.029	0.573	0.054	0.535	0.026	0.647

Table 3. Quantitative results on realistic style across different image domains. All methods use the same number of strokes (4000) on a 512x512 canvas.



Figure 7. Comparison of four parametric painting techniques—L2P, SNP, PT, and our approach—across three different scenarios. Our method’s outcomes show improved detail rendition and variability. Our realistic approach finely details both high and low-frequency elements, accurately rendering tree branches and window frames. Note L2P’s realistic yet blurry style, SNP’s broad painterly effect which struggles to capture fine details (tree branches or windows framing), and PT’s inconsistent detail and brushstroke scale. Zoom-in recommended for full appreciation of textural nuances.

tion metrics, reported in Table 3. We select 10 images from each domain, resize them to 512x512 and set the number of strokes to 4000. Our method outperforms previous methods in high-fidelity realistic paintings across all evaluated



Figure 8. Painterly styles with a different degree of abstraction. This style generates organic human-like paintings, biasing the painting towards different structures.




Figure 9. Stroke distribution comparison of Figure 7 middle column. SNP (a) and PT (c) exhibit a disperse clustering, while L2P (b) paints evenly in a scene-agnostic way. Our method (d) combines the selective focus of SNP with the clarity of PT, indicating a refined approach to applying strokes.

domains, which is consistent with the qualitative results obtained in our user study.

To visually quantify how each method places strokes on the canvas, Figure 9 compares the stroke distributions for the middle column paintings shown in Figure 7 and Figure 8. Our method (d) shows a more intentional stroke placement and clarity in stroke application. In contrast, SNP (a) shows a scattered, uniform treatment of scene elements, while our approach strategically focuses on two key areas, corresponding to the vegetation, enhancing the composition’s structure. L2P (b) applies strokes uniformly in a scene-agnostic way, lacking the focal emphasis of our method. PT (c) creates abrupt density transitions due to its multi-scale approach, unsuitable for gradual tonal shifts, and fails to disentangle high-density areas, not differentiating between scene elements like trees and roads. Our method, however, clearly delineates regions with varying stroke densities, offering a nuanced and intentional subject representation. See supplemental for additional details.

**Efficiency** We compare the efficiency of the different previous methods in a single RTX A6000 Nvidia GPU for a single image (see Table 4). While neural methods require

Method	Optimization-based		Neural-based	
	Ours	Optim [40]	Transf [23]	RL [15]
Res. 1024	<b>99.60s</b>	447.10s	<b>0.514s</b>	0.537s
Req. Training	No	No	Yes	Yes
Use Dataset	No	No	No	Yes

Table 4. Inference efficiency for different methods, measured in seconds for a single image at 1024x1024 and 1600 strokes.

training, they are faster at inference time than optimization methods. However, when comparing with previous optimization methods [40], our method is significantly faster (4x) for the same number of strokes (1600) and resolution size (1024x1024). This is specially important because optimization methods require an optimization process per image. For a fair comparison, we use the uniform setting for our dynamic maps, and set the style to realism.

## 5. Conclusion

This paper presents a semantic-based painting method that is more controllable, produces more visually appealing paintings, and achieves better realistic paintings than previous methods, as validated in the user studies and quantitative evaluations. Contrary to previous methods, our work offers a controllable framework that, besides realism, also aims for painterly and abstracted styles without the need for style-transfer like techniques. Furthermore, it is able to paint at any aspect ratio without further resizing steps. Future work includes incorporating a differentiable pasting module that enables general-canvas level loss, along with our patch loss strategy, for further exploration in stylization.



## References

- [1] Ardavan Bidgoli, Manuel Ladron De Guevara, Cinnie Hsiung, Jean Oh, and Eunsu Kang. Artistic style in robotic painting; a machine learning approach to learning brush-stroke from human artists. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 412–418, 2020. [2](#)
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 213–229. Springer, 2020. [4](#)
- [3] Rebecca Chamberlain and Johan Wagemans. The genesis of errors in drawing. *Neuroscience & Biobehavioral Reviews*, 65:195–207, 2016. [2](#)
- [4] Dale J Cohen. Look little, look often: The influence of gaze frequency on drawing accuracy. *Perception & psychophysics*, 67:997–1009, 2005. [2](#)
- [5] Yaroslav Ganin, Tejas Kulkarni, Igor Babuschkin, SM Ali Eslami, and Oriol Vinyals. Synthesizing programs for images using reinforced adversarial learning. In *International Conference on Machine Learning*, pages 1666–1675. PMLR, 2018. [2](#)
- [6] 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 (CVPR)*, 2016. [3](#)
- [7] Kuba Glazek. Visual and motor processing in visual artists: Implications for cognitive and neural mechanisms. *Psychology of Aesthetics, Creativity, and the Arts*, 6(2):155, 2012. [2](#)
- [8] Jörg Marvin Gülzow, Liat Grayver, and Oliver Deussen. Self-improving robotic brushstroke replication. *Arts*, 7(4), 2018. [2](#)
- [9] David Ha and Douglas Eck. A neural representation of sketch drawings. *arXiv preprint arXiv:1704.03477*, 2017. [3](#)
- [10] Paul Haerberli. Paint By Numbers: Abstract Image Representations. In *Proc. SIGGRAPH*, 1990. [2](#)
- [11] Aaron Hertzmann. Painterly Rendering with Curved Brush Strokes of Multiple Sizes. In *Proc. SIGGRAPH*, 1998. [3](#)
- [12] Aaron Hertzmann. Paint by relaxation. In *Proc. CGI*, 2001. [2](#)
- [13] Aaron Hertzmann. A Survey of Stroke-Based Rendering. *IEEE Computer Graphics & Applications*, 23(4), 2003. [2](#)
- [14] S. Hodge. The painting techniques of 5 famous artists, from botticelli to vermeer. 2019. [5](#)
- [15] 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 (ICCV)*, 2019. [1](#), [2](#), [4](#), [6](#), [7](#), [8](#)
- [16] Biao Jia, Jonathan Brandt, Radomir Mech, Byungmoon Kim, and Dinesh Manocha. Lpaintb: Learning to paint from self-supervision. In *Proc. Pacific Graphics*, 2019. [2](#)
- [17] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. [2](#), [3](#)
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [4](#)
- [19] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [3](#)
- [20] Tzu-Mao Li, Michal Lukáč, Michaël Gharbi, and Jonathan Ragan-Kelley. Differentiable vector graphics rasterization for editing and learning. *ACM Trans. Graph.*, 39(6), 2020. [2](#), [3](#)
- [21] Thomas Lindemeier, Jens Metzner, Lena Pollak, and Oliver Deussen. Hardware-based non-photorealistic rendering using a painting robot. *Computer Graphics Forum*, 34(2):311–323, 2015. [2](#)
- [22] Peter Litwinowicz. Processing Images and Video for an Impressionist Effect. In *Proc. SIGGRAPH*, 1997. [2](#)
- [23] 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 (ICCV)*, pages 6598–6607, 2021. [1](#), [2](#), [6](#), [7](#), [8](#)
- [24] Barbara J Meier. Painterly rendering for animation. In *Proc. SIGGRAPH*, pages 477–484, 1996. [2](#)
- [25] John F. J. Mellor, Eunbyung Park, Yaroslav Ganin, Igor Babuschkin, Tejas Kulkarni, Dan Rosenbaum, Andy Ballard, Theophane Weber, Oriol Vinyals, and S. M. Ali Eslami. Unsupervised doodling and painting with improved spiral, 2019. *arXiv:1910.01007*. [2](#)
- [26] Haoran Mo, Edgar Simo-Serra, Chengying Gao, Changqing Zou, and Ruomei Wang. General virtual sketching framework for vector line art. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2021)*, 40(4):51:1–51:14, 2021. [3](#)
- [27] Peter O’Donovan and Aaron Hertzmann. Anipaint: Interactive painterly animation from video. *IEEE Transactions on Visualization and Computer Graphics*, 18(3):475–487, 2012. [2](#)
- [28] Florian Perdreau and Patrick Cavanagh. Drawing experts have better visual memory while drawing. *Journal of Vision*, 15(5):5–5, 2015. [2](#)
- [29] Paul Rosin and John Collomosse. *Image and Video-Based Artistic Stylisation*. Springer, 2013. [2](#)
- [30] Peter Schaldenbrand and Jean Oh. Content masked loss: Human-like brush stroke planning in a reinforcement learning painting agent. In *Proc. AAAI*, 2021. [2](#)
- [31] Peter Schaldenbrand, James McCann, and Jean Oh. Frida: A collaborative robot painter with a differentiable, real2sim2real planning environment. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11712–11718, 2023. [2](#)
- [32] Johannes Schmid, Martin Sebastian Senn, Markus Gross, and Robert W. Sumner. Overcoat: An implicit canvas for 3d painting. In *Proc. SIGGRAPH*, 2011. [2](#)
- [33] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [4](#)

- [34] Jaskirat Singh and Liang Zheng. Combining semantic guidance and deep reinforcement learning for generating human level paintings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16387–16396, 2021. [2](#)
- [35] John Tchalenko. Segmentation and accuracy in copying and drawing: Experts and beginners. *Vision research*, 49(8):791–800, 2009. [2](#)
- [36] John Tchalenko and R Chris Miall. Eye–hand strategies in copying complex lines. *Cortex*, 45(3):368–376, 2009. [2](#)
- [37] Ning Xie, Hirotaka Hachiya, and Masashi Sugiyama. Artist agent: A reinforcement learning approach to automatic stroke generation in oriental ink painting. *IEICE TRANSACTIONS on Information and Systems*, 96(5):1134–1144, 2013. [2](#)
- [38] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. CVPR*, 2018. [2](#), [4](#)
- [39] Ningyuan Zheng, Yifan Jiang, and Dingjiang Huang. Strokenet: A neural painting environment. In *International Conference on Learning Representations*, 2019. [3](#)
- [40] Zhengxia Zou, Tianyang Shi, Shuang Qiu, Yi Yuan, and Zhenwei Shi. Stylized neural painting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15689–15698, 2021. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)