

# Content Masked Loss: Human-Like Brush Stroke Planning in a Reinforcement Learning Painting Agent

Peter Schaldenbrand<sup>1</sup>, Jean Oh<sup>2</sup>

<sup>1</sup> Human-Computer Interaction Institute, Carnegie Mellon University

<sup>2</sup> Robotics Institute, Carnegie Mellon University  
pschalde@andrew.cmu.edu, jeanoh@nrec.ri.cmu.edu

## Abstract

The objective of most Reinforcement Learning painting agents is to minimize the loss between a target image and the paint canvas. Human painter artistry emphasizes important features of the target image rather than simply reproducing it. Using adversarial or  $L_2$  losses in the RL painting models, although its final output is generally a work of finesse, produces a stroke sequence that is vastly different from that which a human would produce since the model does not have knowledge about the abstract features in the target image. In order to increase the human-like planning of the model without the use of expensive human data, we introduce a new loss function for use with the model’s reward function: Content Masked Loss. In the context of robot painting, Content Masked Loss employs an object detection model to extract features which are used to assign higher weight to regions of the canvas that a human would find important for recognizing content. The results, based on 332 human evaluators, show that the digital paintings produced by our Content Masked model show detectable subject matter earlier in the stroke sequence than existing methods without compromising on the quality of the final painting. Our code is available at <https://github.com/pschaldenbrand/ContentMaskedLoss>.

## Introduction

While there are numerous differences between the methods used by a skilled human painter and a machine printer, the end products may look similar. As is the case with photo-realistic portraits, a human painter hopes to recreate a target image as closely as possible with paint. A printer, though not using paint and a brush, would have a similar objective, however, would go about it in a completely different fashion. A printer may use a simple heuristic to render the image such as starting from the top left and working left-to-right and top-to-bottom, but a human painter would likely take an abstract approach. As a painter paints using the commonly used “blocking in” (Scott 2017) technique, you can recognize the subject of the painting early in the painting process, and the fidelity of the rendering improves with time. In this way, painting can be considered an anytime algorithm: If the painting process is interrupted at any point, a valid solution should be available, but the solution becomes better the longer the process executes for.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The goal of Stroke Based Rendering (SBR) is to find methods to reproduce a target image using visual elements, such as paint brush strokes. Many SBR methods (Hertzmann 2003) were deterministic algorithms which placed strokes onto a canvas that would minimize the loss between the next canvas and the target image. Modern SBR methods (Ganin et al. 2018; Huang, Heng, and Zhou 2019; Nakano 2019; Xie, Hachiya, and Sugiyama 2012) use Reinforcement Learning models to generate the stroke sequences. All of these RL methods use an adversarial or  $L_2$  loss in their reward functions which are designed to minimize the difference between the canvas and the target image. These loss functions are oblivious to abstract content of the image and the stroke sequence can follow any method that the model finds on its way to minimizing the loss.

With existing SBR RL, it is often not clear what the models are painting until the last set of brush strokes are performed as seen in Figure 3. In this way, the SBR RL models are similar to printers and deterministic SBR models. In generative art such as GANs (Elgammal et al. 2017), only the final output image is used for evaluating models. In machine painting, however, the whole performance needs to be evaluated in addition to the final canvas. In this paper, we view painting as a plan that is a sequence of actions (or brush strokes) towards an end output. In this context, the proposed idea is to improve the quality of a painting plan according to human-like painting techniques where a subject of the painting becomes incrementally more recognizable as brush strokes are added.

One approach for guiding a RL painting model to paint

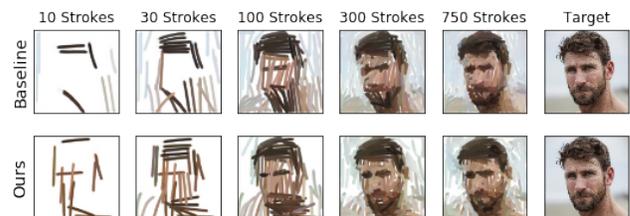


Figure 1: Our model was trained using a clipped  $L_1$  loss with Content Masking in its reward function, while the baseline used adversarial loss.

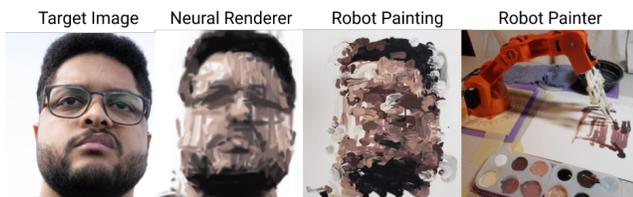


Figure 2: The Reinforcement Learning Painter generated brush strokes to paint the target image. The strokes can be rendered into an image by feeding them to the Neural Renderer or the robot painting agent as seen above. The robot painter is a TinkerKit Braccio robotic arm. 10 paint colors, that are chosen by running k-means clustering on the target image, are mixed by a human and provided to the robot. The arm has a single paint brush and can clean it using water and a towel. The robot is fully autonomous while painting.

more like a human would be to collect brush stroke level data from real human painters. This method would require collecting a large number of samples, each of which would take a long time to produce. The samples would need to be diverse so that the model generalizes on multiple different types of paintings such as portraits or landscapes. All-in-all it would be an expensive, time-consuming process.

Our research goal in this paper is to make the brush stroke planning of a reinforcement learning painter model more human-like. We focus on the common painting technique of “blocking-in” where the subject of a painting is recognizable early in the painting process. We introduce Content Masked Loss which gives more reward to painting regions of the canvas that are important for recognizing the subject matter of the painting. An object recognition model is used to extract which regions of the image are important for recognition. Amazon Mechanical Turk workers found that the subject matter of the paintings produced from our models that used Content Masked Loss was apparent far earlier in the brush stroke planning than baseline models using adversarial or  $L_2$  loss.



Figure 3: Row (a) is the (Huang, Heng, and Zhou 2019) model and row (b) is our baseline model trained using a constrained Neural Renderer. The constraints assume that the brush strokes are fixed width, opaque, and do not exceed a maximum length. Our baseline model also assumes a white starting canvas whereas (Huang, Heng, and Zhou 2019) starts with a black canvas.

## Related Work

Stroke-based Rendering (SBR) is the process of layering individual elements, such as shapes or brush strokes, onto a digital canvas in order to reproduce a given image. The goal of SBR may be to perfectly recreate an image with a given set of tools, or it may be to add creative abstractions while rendering.

Recent methods attempt to tackle SBR using deep learning (Ganin et al. 2018; Huang, Heng, and Zhou 2019; Nakano 2019; Xie, Hachiya, and Sugiyama 2012), though the field has existed long before these tools became popular (Hertzmann 2003). “Artist Agent” (Xie, Hachiya, and Sugiyama 2012) was an approach to modelling brush strokes on paper using a reinforcement learning agent. The agent designed brush strokes and received feedback as to whether the strokes were believably made by a human or not. The individual brush strokes produced by Artist Agent were believably human but in order to render a whole painting, the stroke planning was manually produced by humans.

Other SBR work attempts to create an agent that can produce whole paintings rather than single brush strokes. SPIRAL (Ganin et al. 2018) used reinforcement learning to generate brush strokes that would render into a painting that minimized the difference between the canvas and the given image. Work that builds off of SPIRAL, such as (Nakano 2019; Zheng, Jiang, and Huang 2019; Huang, Heng, and Zhou 2019), often use the world model concept (Ha and Schmidhuber 2018) in order to speed up training. In the case of painting, the world model is called a Neural Renderer. The Neural Renderer is implemented with a differentiable neural network so that the loss can be back-propagated through it. (Huang, Heng, and Zhou 2019) achieves the best results at decomposing an image into a series of strokes. The objective of their reinforcement learning algorithm is to minimize the difference between the generated painting and the given painting. The authors find that using WGAN loss known as *Wasserstein-1* distance, or *Earth-Mover* distance, generates better results than the  $L_2$  distance. The results from this paper are very successful, for with enough strokes, the model is able to almost completely recreate the input image, but the planning process shows no indication of what the subject matter of the painting is.

Existing datasets that have paired images with strokes lack quality for painting portraits. For instance, the QuickDraw (Ha and Eck 2017) dataset, which contains cartoon sketches that users create while trying to render a given image or noun, contain both the brush strokes and image data, however, these sketches are too simple for learning real painting. SketchRNN (Ha and Eck 2017) does use paired data, but it can only draw objects that have been labelled in its training data. This dataset encodes the planning method for how humans sketch, but we need the method of how humans paint.

Reinforcement Learning models such as (Huang, Heng, and Zhou 2019) and SPIRAL (Ganin et al. 2018) learn through trial and error how to paint images. The reward functions are designed to make the model learn which strokes will decrease the loss between the painting and the target image. Because these models learn without human

data guidance, they do not encode how a human painter would paint. In this paper, we introduce Content Masked Loss which increases the human-like brush stroke planning abilities of the RL model without the use of human painter data.

The brush strokes used in the planning of (Huang, Heng, and Zhou 2019) are not capable of being reproduced by a standard paint brush, as seen in Figure 3. Because these brush strokes are not reproducible by a real paint brush, we cannot evaluate if the brush stroke planning is human-like or not. Constraining the model to use only realistic brush strokes will make it possible to evaluate whether the planning is human-like or not. In this paper, we constrain our Reinforcement Learning method of Stroke-Based Rendering to produce brush stroke instructions that are realistic and can easily be fed to a painting agent to be painted accurately. An example of feeding these paint brush instructions can be seen in Figure 2.

## Approach

A robot painting problem is informally defined as follows. Given an input image and a blank canvas, the goal is to generate a sequence of brush strokes that will turn the canvas into a recreation of the input image, following a human-like painting style. This section describes an RL formulation of the task, focusing on how the proposed reward functions induce human-like plans.

### Reinforcement Learning Painter Model

The proposed work builds on the deep reinforcement learning framework from (Huang, Heng, and Zhou 2019). The original framework is based on several unrealistic assumptions; We modify the framework to address realistic constraints. First, the brush strokes are allowed to be unrealistically large, e.g., more than half the width of the canvas as seen in Figure 3. We constrain what is possible for the brush strokes based on what our robot painting agent is able to do with a single paint brush. In this way, we can easily feed the brush stroke instructions to a robot such as the one used in our experiments shown in Figure 2. Specifically, we constrained them to be fixed width and only 5% of the width of the canvas which is consistent with the robot’s brush width. Second, it was assumed that the blank starting canvas is black which is not consistent with the white paper or canvas that a painter would generally use; we changed this assumption to be white in our model. Third, a paint brush can only hold so much paint, so we add the maximum length of one half of the canvas width to the brush strokes based on the robots attempts at painting on a 20cm piece of paper. Additionally, since the acrylic paint that the robot utilizes is opaque, we add this as a constraint for the purpose of our study.

We formulate painting as a reinforcement learning task, following the problem formulation in (Huang, Heng, and Zhou 2019). Here, states represent the canvas from blank to completion. The action space describes the brush strokes in the form of Bézier curves with values for three position coordinates, thickness of the brush stroke at both ends, opacity

$$Reward = mean [L(c_t, y) - L(c_{t+1}, y)] \quad (1)$$

#### Baseline

$$L_{GAN}(c, y) = 1 - Discriminator(c, y) \quad (2)$$

$$L_{L2}(c, y) = (c - y)^2 \quad (3)$$

#### Ours

$$L_{CL}(c, y) = (VGG_l(c) - VGG_l(y))^2 \quad (4)$$

$$L_{CM+L2}(c, y) = (c - y)^2 * norm(VGG_l(y)) \quad (5)$$

$$L_{L1^*}(c, y) = min(|c - y|, \lambda) \quad (6)$$

$$L_{CM+L1^*}(c, y) = min(|c - y|, \lambda) * norm(VGG_l(y)) \quad (7)$$

Figure 4: The reward function is the per-pixel average of the loss between the canvas at time  $t$  and the canvas with an additional stroke at time  $t + 1$ . The various loss functions used in this paper follow, where  $c$  is the canvas and  $y$  is the target image. The *Discriminator* is a model used to differentiate paintings of an image and the actual image, and *norm()* is a 0-1 normalizing function.

of the paint at both ends, and color. The Neural Renderer is the transition function which paints a brush stroke in this action space onto a given canvas.

As the action space is continuous, our model adopts the Deep Deterministic Policy Gradient (DDPG) (Lillicrap et al. 2016). The reward function is calculated by computing the difference between the losses at time  $t$  and  $t + 1$  as seen in Equation 1. Huang *et al.* uses an adversarial model to compute the loss function (Equation 2). The model based DDPG critic predicts an expected reward function:

$$V(c_t) = r(c_t, a_t) + \gamma V(c_{t+1})$$

where  $r(c_t, a_t)$  is the reward of performing action  $a_t$  on canvas  $c_t$ , and  $\gamma$  is a penalty parameter which controls how influential future value is.  $c_{t+1}$  can be rendered using the Neural Renderer to apply the brush stroke  $a_t$  to the canvas  $c_{t+1}$ , seen below as  $trans(c_t, a_t)$ . The actor in the DDPG,  $\pi(c_t)$ , is trained to maximize the expected reward:

$$r(c_t, \pi(c_t)) + \gamma V(trans(c_t, \pi(c_t)))$$

### Reward Functions

The (Huang, Heng, and Zhou 2019) model used the adversarial loss as seen in Equation 2 and is used in the baseline model in this paper. The *Discriminator* model outputs 1 if the canvas and target image are the same and 0 if they are completely different. The *Discriminator* is only trained to determine visual differences between the paintings and the target image so it does not factor in content and abstract detail from the images.

Painters take an abstract approach when painting and consider the subject of the target image. Human painters use techniques such as “blocking-in” (Scott 2017) which ensure

that the subject of the image is recognizable early in the painting process. We took two different approaches to increasing the recognizability of the painting strokes: Content Loss Reward and Content Masked Reward.

**Content Loss Reward** In Content Loss (Equation 4), we used the  $L_2$  loss on the features extracted from the canvas and target images instead of the images themselves. We used the VGG-16 (Simonyan and Zisserman 2015) model to extract the features. The VGG-16 parameters were fixed. The intuition behind Content Loss is that if a painting and the target image have similar features extracted from an object detection model, then they should recognizably be paintings of the same object.

The VGG-16 (Simonyan and Zisserman 2015) object detection model was trained on photographs, so it is not expected to generalize to paintings. We attempt to use Content Loss but recognize that a feature extractor that is robust to both photographs and paintings may be necessary to make this loss function work properly.

**Content Masked Loss Reward** Content Masked Loss was designed to increase the weight of the per-pixel loss in “important” regions of the canvas. The “importance” of a region is determined by extracting features using the VGG-16 (Simonyan and Zisserman 2015) model. As seen in Figure 5, the Content Masking can be added to the  $L_2$  distance to weight the difference between a canvas and the target image. The features are computed by feeding the target image into the VGG-16 model, taking the output of an intermediate layer, averaging the image across the channels, and then 0-1 normalizing it. We experimented using the output from different layers in the VGG-16 model in Figures 8 and 9.

**L1 with Loss Clipping** Certain biases of the painting model may cover up the effects of Content Masking. All of the loss functions except Content Loss rely on measures of distance between pixels in the canvas and the target image. Since the images are in RGB format, the largest distance would be between a white and black pixel. The canvas is initially white, so the model will have a bias towards painting darker colors, since it can get more reward with them. We reduced this bias by experimenting with  $L_1$  instead of  $L_2$  losses. We also set a constant maximum value for the per-pixel loss. Loss Clipping used in a loss function can be seen in Figure 4 with  $L_1$  ( $L1*$ ) and in conjunction with Content Masked Loss ( $CM + L1*$ ).

$$\text{Content Masked Loss} = \left( \begin{array}{c} \text{Canvas} \\ \text{Target} \end{array} - \begin{array}{c} \text{Target} \\ \text{Canvas} \end{array} \right)^2 * \begin{array}{c} \text{Features} \\ \text{Target} \end{array}$$

Figure 5: Content Masking added to the  $L_2$  loss is defined as the difference between a canvas and the target image weighted by the features extracted from the target image.



Figure 6: On the left is the target image. The paintings were created by the model generating 125 strokes looking at the whole canvas and then 125 strokes in each of 25 evenly distributed subsections of the canvas. (a) was created using the (Huang, Heng, and Zhou 2019) model as is while (b) used constraints on the brush’s abilities.

## Results

The effects of the brush stroke constraints (fixed brush width, maximum brush stroke length, fully opaque strokes, and starting with white canvas) can be seen in Figure 3. Our baseline model is able to reproduce the target image but not with as high fidelity as the (Huang, Heng, and Zhou 2019) model. The strokes are now fixed width and smaller than what the (Huang, Heng, and Zhou 2019) model allows, so it takes more strokes to reproduce the target image in our baseline model. The fixed width strokes are too big to make out details such as mouth and eyes. In order to paint with a smaller brush size, the canvas is divided up into subsections, then the model is run over each subsection. So if we divide an image into quadrants and run our baseline model over each, this is equivalent to painting with a brush half the width of fixed width constraint.

The last canvases in Figure 3 show a significant difference in fidelity between our baseline model and the (Huang, Heng, and Zhou 2019) model. One of the most impressive aspects of the (Huang, Heng, and Zhou 2019) model was that it could nearly replicate an image given enough strokes. We tested whether our baseline model could also achieve this accuracy. In Figure 6, the models generated 125 strokes using the entire image as input, then 125 strokes for each non-overlapping 1/25th of the image. Our baseline model produces a painting with a similar loss to the (Huang, Heng, and Zhou 2019) model’s painting.

## Content Loss Reward

We tried to use this content loss (Equation 4) reward function, but the model would only go to a local minimum and produce blank canvases no matter what parameters we used. The main reason why the Content Loss Reward function was unsuccessful was the feature extraction process does not translate well to the paintings. The VGG-16 model (Simonyan and Zisserman 2015) was trained on photographs. The canvases (especially with only a few strokes) are so vastly different from photographs that the extracted features appear random. As seen in Figure 7, even when the painting is extremely similar to the target image, the features extracted are very different from those extracted from the pho-

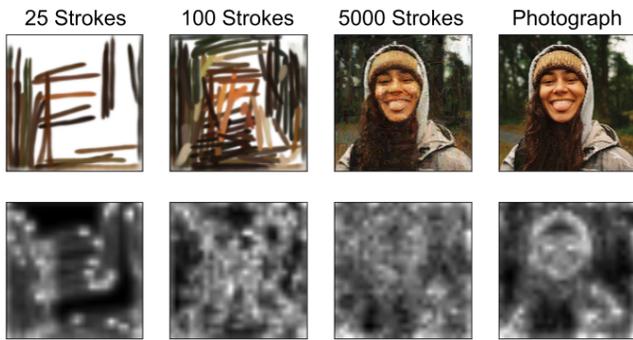


Figure 7: The top left three images are output from the model painting the target image in the top right. The second row are the features extracted from their corresponding image above them using the first 17 layers of the VGG-16 model.

tograph. The features extracted from the photographic target image make semantic, abstract sense while the features from the paintings are unintelligible. In future work we will attempt to use feature extraction models that are robust to texture variance, such as (Geirhos et al. 2019) who minimized the texture bias in an ImageNet (Deng et al. 2009) classifier model.

### Content Masked Reward

The VGG-16 model has 13 convolutional layers each with an activation function and 5 pooling layers. These 31 layers can be used in the feature extraction portion of the architecture. To compute the feature mask, the output of a layer is averaged across channels and 0-1 normalized. The output of each of the first 29 layers of the pretrained VGG-16 model are shown in Figure 8. The initial layers extract features such as edges and texture, while the later layers extract higher level information such as eyes and mouths.

We trained the model using the Content Masked Loss (Equation 5) in the reward function using various layers from the VGG-16 model. After the models were trained, the painting process for each model was tested in Figure 9. The lower level layers produced models that focused more broadly on the subjects face while higher layers were focused on abstract features such as eyes. All models performed very similar in their final results, however, the intermediate strokes varied. The models that used latter layers performed more human-like since they focused on the high level features of the image.

Layer 17 was chosen for future experiments since it was a good balance between using high and low level features.

### Final Painting Quality

We wanted to test whether there was a difference in quality of paintings amongst the models trained with different reward functions. We used both the Fréchet Inception Distance (FID) (Heusel et al. 2017) and Inception Score (Salimans et al. 2016). These measures have known issues (Barratt and Sharma 2018), e.g., Inception Score, which uses a network pre-trained on ImageNet (Deng et al. 2009), can

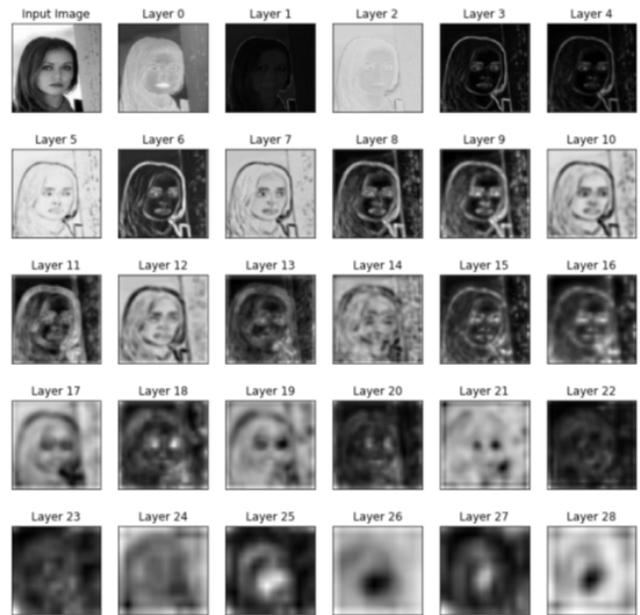


Figure 8: Features extracted from a given image at various layers within the VGG-16 (Simonyan and Zisserman 2015) model.

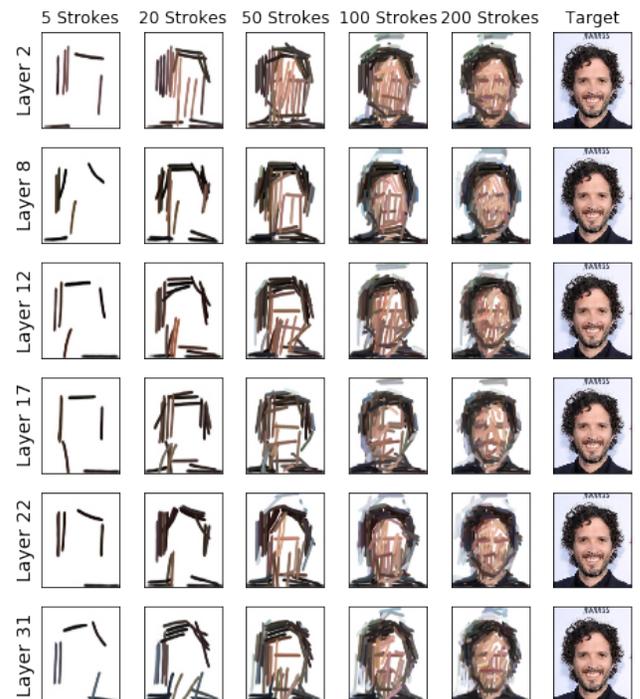


Figure 9: Models were trained using Content Masked Reward where the features were generated using various layers in the VGG-16 model.

Loss Function	FID ↓	IS ↑
GAN	240.05 ± 0.43	3.05 ± .08
L2	241.02 ± 0.24	3.05 ± .14
CM + L2	242.67 ± 0.24	3.01 ± .04
L1*	241.58 ± 0.14	3.45 ± .09
CM + L1*	243.13 ± 0.20	3.47 ± .05

Table 1: A small Fréchet Inception Distance (FID shown in column 2) indicates that the generated paintings are visually similar to the photographs. The larger the Inception Score (IS shown in column 3), the more similar the distributions of the paintings and target images are.

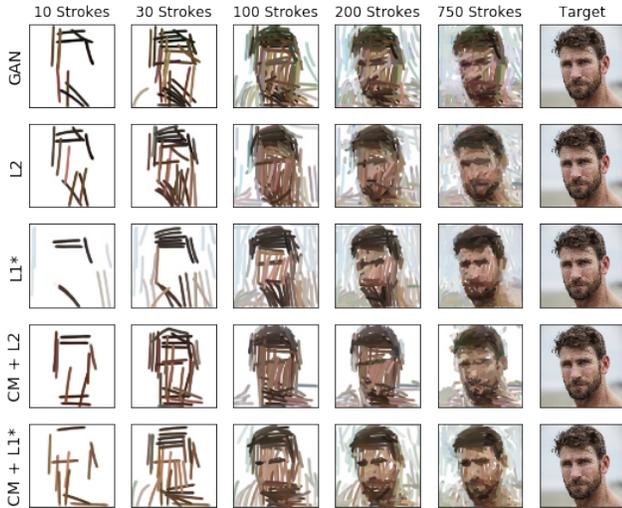


Figure 10: The painting process for models trained using different loss functions in their reward functions.

be misleading when used on the models trained on non-ImageNet datasets. Since we are merely comparing reward functions and the models have otherwise been trained in the exact same fashion, we can use these scores for comparison.

We computed the FID and Inception Scores for 2000 held out images from the CelebA dataset. The FID and Inception Scores are shown in Table 1. They were computed using 5 fold cross validation on 2000 held out images from the CelebA dataset. We note that both FID and IS scores from our experiments seem far outside the normal range of those scores found in GANs literature (Ostrovski, Dabney, and Munos 2018); the reason for this is that the generated output in our case are paintings that are visually very different from photorealistic input images as in other works. There are significant differences in the FID and IS values between certain models, however, there is no model that is significantly superior to the others in terms of both FID and IS. The lack of a significantly superior model indicates that the final painting quality of each model is roughly the same, and therefore, the intervention of loss functions does not significantly affect the final painting quality.

## Human-Like Planning

As a human painter paints, the subject of the painting should be apparent per the common “blocking-in” (Scott 2017) technique. As seen in Figure 10, the loss function used to train the painter model affects the intermediate outputs of the model. For instance, the GAN,  $L1^*$ , and  $L2$  Loss models show that the painting model is not influenced by the content in the target image. The CM + L2 and CM + L1\* models add incentive to paint regions of the image that are important for understanding the content of the image, and so observers can recognize what the model is painting early in the stroke sequence.

We hypothesize that our model, aiming at a human-like planning, produces stroke sequences that can be recognized from early strokes. To verify this hypothesis, we conduct two experiments: first, we use a facial detection AI model on the canvas as the model painted; and second, we ran a study on Amazon Mechanical Turk for human evaluation.

We note on the difference between AI and human perception. In general, humans tend to focus on shape-based features in recognizing objects, whereas Convolutional Neural Networks including FaceNet (Schroff, Kalenichenko, and Philbin 2015) used in our experiment exhibit strong texture biases (Geirhos et al. 2019). Due to the nature of binary classification used in the experiment, the two experiments are not designed in the identical format, i.e., binary classification vs. ranking; however, the results agree that at early stage of under 200 strokes the proposed approach generates paintings where the subject is more easily recognizable when compared to the baseline approaches.

We tested the following five models, each varying with only the loss function used to compute the reward:

- GAN - Adversarial Loss. Equivalent to (Huang, Heng, and Zhou 2019) but with brush stroke constraints (Equation 2).
- L2 -  $L_2$  Loss (Equation 3).
- L1\* -  $L_1$  with Clipping Loss (Equation 6).
- CM + L2 - Content Masked  $L_2$  Loss (Equation 5).
- CM + L1\* - Content Masked,  $L_1$  with Clipping Loss (Equation 7).

**AI evaluation.** In our AI test, each model painted 2000 portraits from a subset of data held out from training. After each brush stroke, the painting was tested by the FaceNet facial detection model. The proportion of paintings with detected faces within a certain number of strokes is plotted in Figure 11. Adding content masking generally improved facial detection in early strokes (50-200), although Content Masking’s effect on L2 was not strong, likely due to the color bias concealing it. The proposed CM + L1\* model produced the most face detected paintings.

**Human evaluation.** Each model painted the images stroke by stroke. We showed the workers the canvases produced from each of the five models at various points in the painting process: 10, 30, 100, 200, and 750 strokes. An example of what the Mechanical Turk workers saw in the study

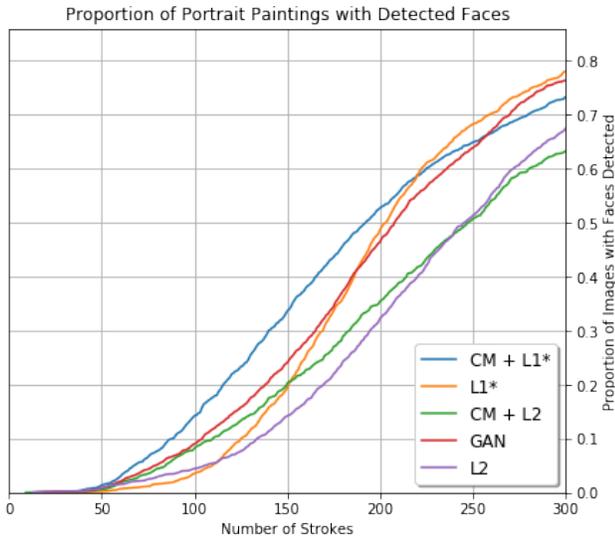


Figure 11: The proportion of paintings produced by each model that had faces detected by a model (Schroff, Kalenichenko, and Philbin 2015) within the specified number of brush strokes.

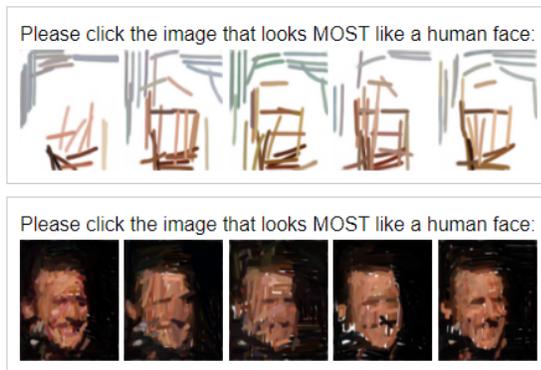


Figure 12: Example of questions in a HIT in our study carried out on Amazon MTurk. Workers viewed paintings by each of the models trained with different loss functions and selected the painting that looked most like a face.

is shown in Figure 12. The order of algorithms was randomized, as was the order of the number of strokes in each image. 332 unique MTurk workers evaluated images. 400 images were used for the paintings, and each question (sample questions are shown in Figure 12) was evaluated by 4 different workers.

The results of the study are shown in Figure 13. Adding Content Masking to L1\* loss improved the face-like qualities of the paintings that contained between 30 and 200 brush strokes. Adding Content Masking to L2 did not significantly improve the face-like qualities of the paintings, and we suspect that this is because of the dark color bias since L1, which is less affected by a color bias, with Content Masking was successful. Overall, Content Masking with L1 loss

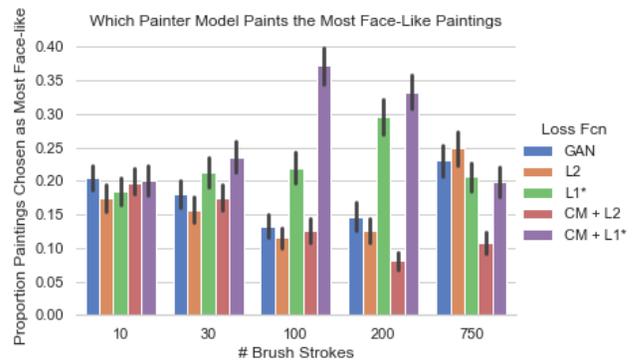


Figure 13: Results from the MTurk study. Workers chose which painting (performed using 10, 30, 100, 200, or 750 brush strokes) looks most like a face given five paintings from each model.

proves to create the painter model that produces the most face-like paintings through the painting process.

## Conclusion and Future Work

This paper presents a reinforcement learning approach towards a human-like brush stroke planning, focusing on the procedure of the painting in addition to the final quality. Experimental results from both AI and human evaluations show that adding Content Masking to the loss functions greatly improves the facial detection abilities of the model's stroke sequence, without sacrificing the quality of the final painting results. Content Masked Loss provides a data-efficient way to increase the human-like abilities of the model's brush stroke sequence without needing additional, expensive human labeling.

We found that the features generated by the VGG-16 object detection model correlate with facial features as seen in Figure 8 which allowed us to use these features to weight the loss in Content Masking. In future work, in order to address the texture bias issue prevalent in CNN-based perception models, we plan to explore using more robust object classifiers as tools for producing the weight matrix in Content Masking that may correlate better with human evaluation.

## Acknowledgements

We thank Zhewei Huang, Wen Heng, and Shuchang Zhou whose code (<https://github.com/hzwer/ICCV2019-LearningToPaint>) we modified for the reinforcement learning models in this paper.

Thank you to Nicholas Eisley (<http://neisley.squarespace.com/>) for the photographs used in Figures 1, 3, 6, 7, and 10.

## References

- Barratt, S.; and Sharma, R. 2018. A Note on the Inception Score. ArXiv preprint arXiv:1801.01973.
- Deng, J.; Dong, W.; Socher, R.; Li, L.; Kai Li; and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image

- database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Elgammal, A. M.; Liu, B.; Elhoseiny, M.; and Mazzone, M. 2017. CAN: Creative Adversarial Networks, Generating “Art” by Learning About Styles and Deviating from Style Norms. In *International Conference on Computational Creativity*. URL <http://arxiv.org/abs/1706.07068>.
- Ganin, Y.; Kulkarni, T.; Babuschkin, I.; Eslami, S. M. A.; and Vinyals, O. 2018. Synthesizing Programs for Images using Reinforced Adversarial Learning. *CoRR* abs/1804.01118. URL <http://arxiv.org/abs/1804.01118>.
- Geirhos, R.; Rubisch, P.; Michaelis, C.; Bethge, M.; Wichmann, F. A.; and Brendel, W. 2019. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *International Conference on Learning Representations*. URL <https://openreview.net/forum?id=Bygh9j09KX>.
- Ha, D.; and Eck, D. 2017. A Neural Representation of Sketch Drawings. *CoRR* abs/1704.03477. URL <http://arxiv.org/abs/1704.03477>.
- Ha, D.; and Schmidhuber, J. 2018. Recurrent World Models Facilitate Policy Evolution. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc. URL <https://proceedings.neurips.cc/paper/2018/file/2de5d16682c3c35007e4e92982f1a2ba-Paper.pdf>.
- Hertzmann, A. 2003. A Survey of Stroke-Based Rendering. *IEEE Computer Graphics and Applications* 23(4): 70–81. doi:<http://doi.ieeecomputersociety.org/10.1109/MCG.2003.1210867>. Special Issue on Non-Photorealistic Rendering.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; Klambauer, G.; and Hochreiter, S. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Nash Equilibrium. *CoRR* abs/1706.08500. URL <http://arxiv.org/abs/1706.08500>.
- Huang, Z.; Heng, W.; and Zhou, S. 2019. Learning to Paint with Model-based Deep Reinforcement Learning. *CoRR* abs/1903.04411. URL <http://arxiv.org/abs/1903.04411>.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2016. Continuous control with deep reinforcement learning. In *ICLR*. URL <http://dblp.uni-trier.de/db/conf/iclr/iclr2016.html#LillicrapHPHETS15>.
- Nakano, R. 2019. Neural Painters: A learned differentiable constraint for generating brushstroke paintings. *CoRR* abs/1904.08410. URL <http://arxiv.org/abs/1904.08410>.
- Ostrovski, G.; Dabney, W.; and Munos, R. 2018. Autoregressive Quantile Networks for Generative Modeling. *CoRR* abs/1806.05575. URL <http://arxiv.org/abs/1806.05575>.
- Salimans, T.; Goodfellow, I. J.; Zaremba, W.; Cheung, V.; Radford, A.; and Chen, X. 2016. Improved Techniques for Training GANs. *CoRR* abs/1606.03498. URL <http://arxiv.org/abs/1606.03498>.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. FaceNet: A unified embedding for face recognition and clustering. In *CVPR*, 815–823. IEEE Computer Society. ISBN 978-1-4673-6964-0. URL <http://dblp.uni-trier.de/db/conf/cvpr/cvpr2015.html#SchroffKP15>.
- Scott, D. 2017. Blocking In. URL <https://drawpaintacademy.com/blocking-in/>. [Last Accessed: August 2nd, 2020].
- Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *International Conference on Learning Representations*.
- Xie, N.; Hachiya, H.; and Sugiyama, M. 2012. Artist Agent: A Reinforcement Learning Approach to Automatic Stroke Generation in Oriental Ink Painting. *CoRR* abs/1206.4634. URL <http://arxiv.org/abs/1206.4634>.
- Zheng, N.; Jiang, Y.; and Huang, D. 2019. StrokeNet: A Neural Painting Environment. In *International Conference on Learning Representations*. URL <https://openreview.net/forum?id=HJxwDiActX>.