

Q: HOW DO WE GET NEURAL NETS TO PAINT?



A: GIVE IT A DIFFERENTIABLE PAINTBRUSH!

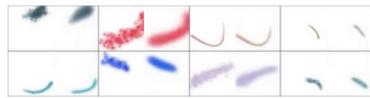
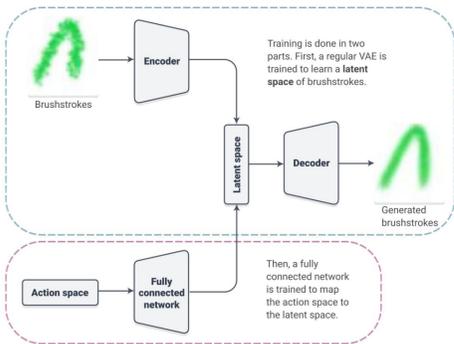


Neural Painters: a learned differentiable constraint for generating brushstroke paintings by Reiichiro Nakano

STEP 1: CREATE A PAINTBRUSH BY TRAINING A DIFFERENTIABLE MODEL OF BRUSHSTROKES

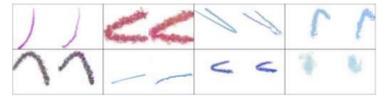
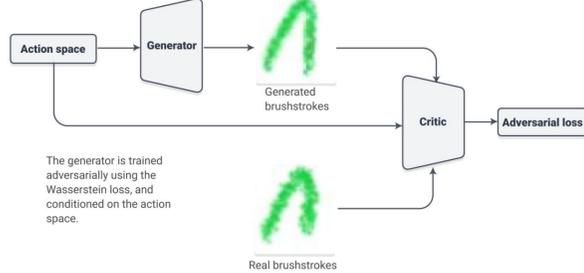
There are numerous painting programs, but most are *non-differentiable*. Since it's difficult to work with non-differentiable things in the context of deep learning, we'll train a neural network (the **neural painter**) to *simulate* an external painting program (MyPaint). A neural painter is a **learned differentiable painting program**. We explore 2 ways to train a neural painter.

VAE Neural Painter



Left: Real Right: Generated
A VAE neural painter results in brushstrokes that are slightly blurred and smudged.

GAN Neural Painter



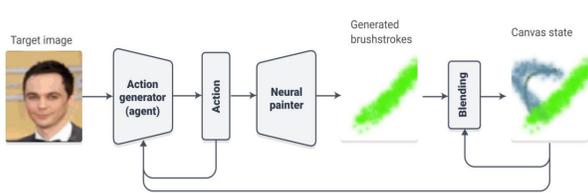
Left: Real Right: Generated
A GAN neural painter results in brushstrokes that emulate the original texture and look more realistic.

STEP 2: USE THE NEURAL PAINTER TO PAINT! (MULTIPLE WAYS)

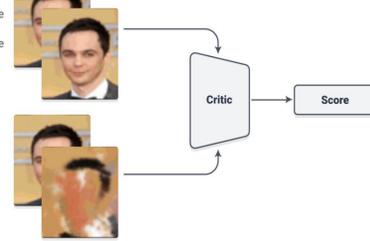
2A: TRAIN A SEPARATE AGENT TO RECONSTRUCT ANY TARGET IMAGE WITH BRUSHSTROKES

One way to use the neural painter is to train a separate network (the agent) to reconstruct any given target image using brushstrokes. We use an adversarial training framework, with a (1) **generator/agent** conditioned to reconstruct an input image with brushstrokes from a neural painter, and a (2) **critic** deciding how well the generated painting matches the target. Training is fully differentiable and done with backpropagation, but at evaluation, the brushstrokes generated by the agent can be transferred back to the real non-differentiable painting program. In this way, the method can be seen as a form of model-based reinforcement learning with a learned world model.

Generator



Critic



Results

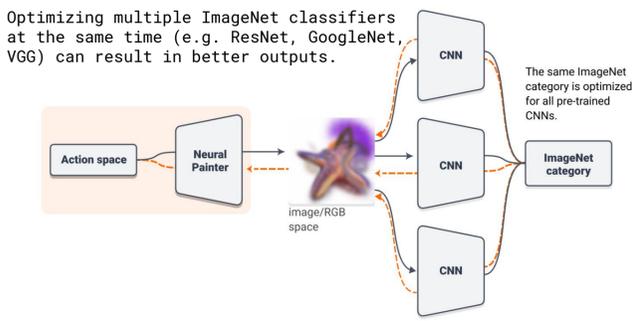


Left: Target image
Center: Generated painting with VAE neural painter-trained agent
Right: Generated painting with GAN neural painter-trained agent

2B: OPTIMIZE BRUSHSTROKES TO MAXIMIZE IMAGENET CLASSES IN A PRE-TRAINED IMAGE CLASSIFIER

Since the neural painter is differentiable, we can directly use gradient descent to find a set of brushstrokes that minimize some loss function. For example, we can use gradient ascent to find a set of brushstrokes that *maximizes* the probability of a particular class in a pre-trained ImageNet classifier. A fun way to interpret the results of maximizing the probability of the *panda* class is as the answer to the question:

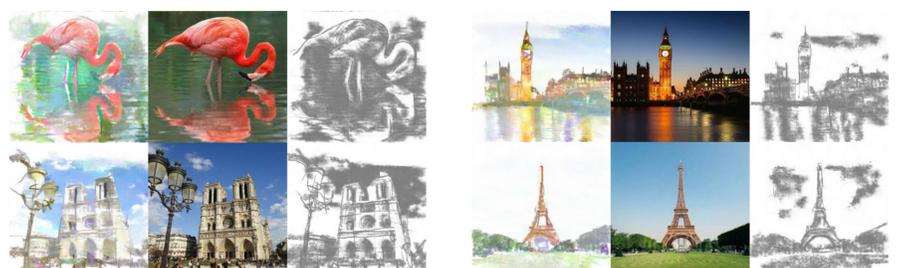
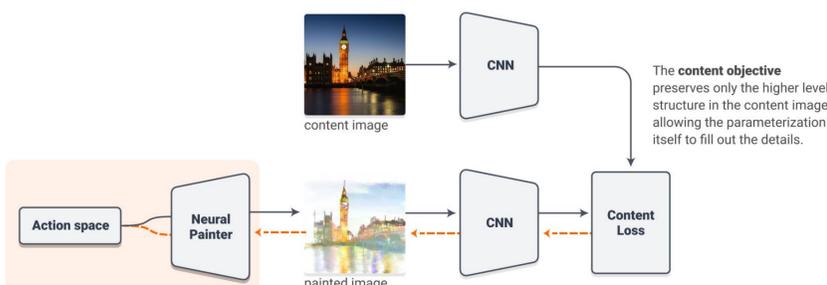
Given a brush, how would a pre-trained ImageNet classifier paint the optimal panda?



Other ImageNet class visualizations

2C: INTRINSIC STYLE TRANSFER: OPTIMIZE BRUSHSTROKES TO MINIMIZE CONTENT LOSS IN NEURAL STYLE TRANSFER

In intrinsic style transfer, we use gradient descent to find a set of brushstrokes that minimize neural style transfer's content loss for a target content image. Intuitively, this method finds brushstrokes that preserve only the meaningful higher-level content in the target image. The style is an **intrinsic** property dictated purely by the artistic medium, in this case, brushstrokes.



Intrinsic style transfer using a GAN neural painter for both colored and grayscale brushstrokes. Note that by manually changing the primitives of the brushstroke (color -> grayscale), we can achieve vastly different styles.

The strokes are generated on multiple overlapping grids. Although the neural painter is only designed to output 64x64 pixels on a canvas, we can stitch multiple canvases together to achieve an arbitrary resolution, limited only by GPU memory.

STEP 3: TRY IT YOURSELF! An easy-to-use library and sample notebooks are available at <https://github.com/reinakano/neural-painters-pytorch>