2019 IEEE ICRA

Mobile Robotic Painting of Texture

Majed El Helou[1,4], Stephan Mandt[2], Andreas Krause[3], Paul Beardsley[4]

[1] School of Computer and Communication Sciences, EPFL; [2] Computer Science, UCI;
[3] Computer Science, ETH Zurich; [4] Disney Research Zurich

1 Introduction

We present a self-supervised deep learning approach to take an input ink map of a desired texture, and infer robotic paint commands to produce that texture, for instance with a quadrotor [1]. We analyze the trade-offs between reconstruction quality and ease of execution. Our method is general for different kinds of robotic paint delivery systems, with an emphasis on spray painting. More generally, the framework can be viewed as an approach for solving a specific class of inverse imaging problems.

2 Method overview

Learning to paint presents multiple scalability advantages relative to optimization with simulators and constant feedback mechanisms. To learn spray commands, we design a differentiable spray simulator that mimics the behavior of robotic painting systems. We discretize the space of spray patterns as well as the space of possible spraying locations to make a convolutional-type spray simulator possible. This (non-trained) differentiable simulator is enclosed in a deep autoencoder [2] CNN as the decoder part, and the encoder's output thus generates the spray commands driving the robot, as internal network activations.

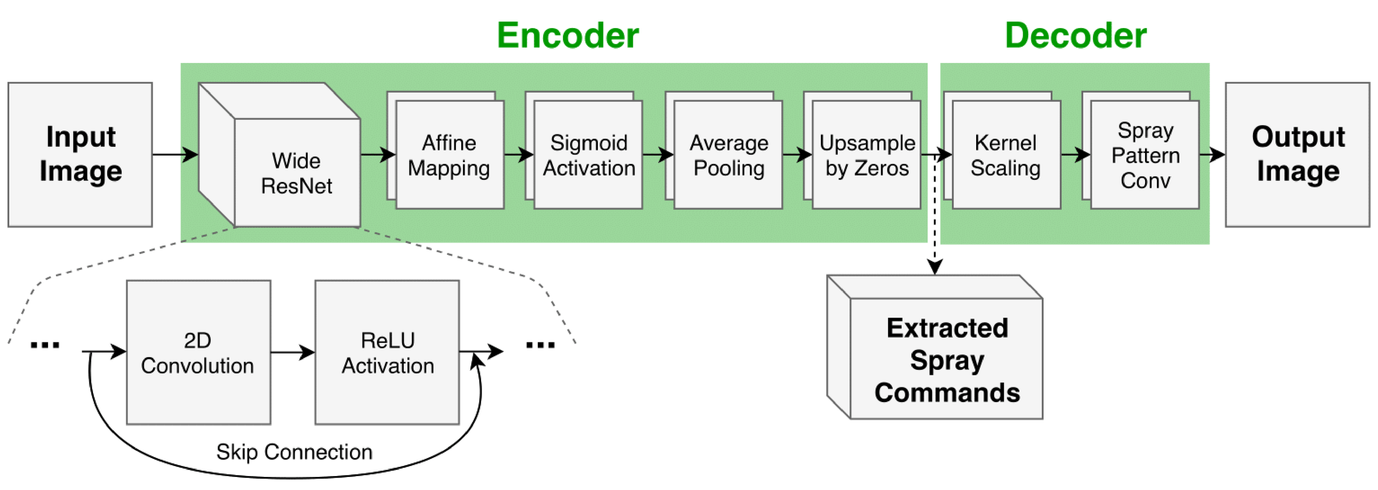

Fig. 1. The encoder is made up of a wide ResNet architecture containing a sequence of res-blocks. The blocks are constituted of 2D convolutions and ReLU activation layers with skip connections. The output is then passed through an affine mapping to shift the distribution of activation values similar to a batch normalization. This affine mapping allows for a varied output in the range [0,1] after passing through the sigmoid activation. The average pooling followed by upsampling with zero filling is only used when sparsity is to be imposed. The decoder scales the spray pattern magnitudes and applies our version of a spray simulator. The largest magnitude can be mapped in the physical world to the longest spray duration before droplets are formed.

3 Results and discussion

Imposing an overall training loss that incorporates both the reconstruction quality of the autoencoder as well as TV smoothness loss [3] on the intermediate feature maps representing the spray commands allows for a trade-off between quality (evaluated with MSE and SSIM [4]) and ease of execution in terms of command smoothness (**Fig. 2, 3**). We also assess ease of execution in terms of command sparsity for different scales of spray patterns (**Fig. 4**).

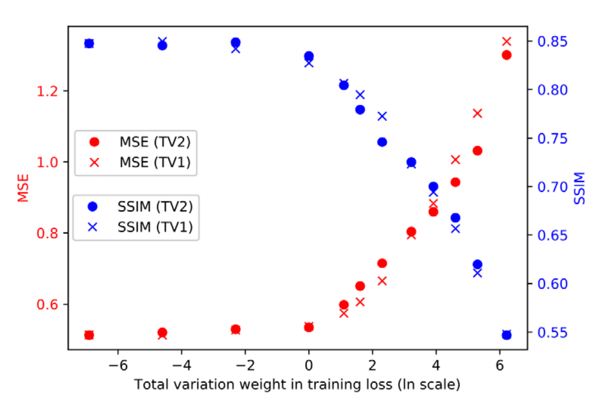

Fig. 2. Effect of increasing smoothness loss vertically (TV1) or horizontally (TV2) on reconstruction quality in terms of MSE and SSIM.

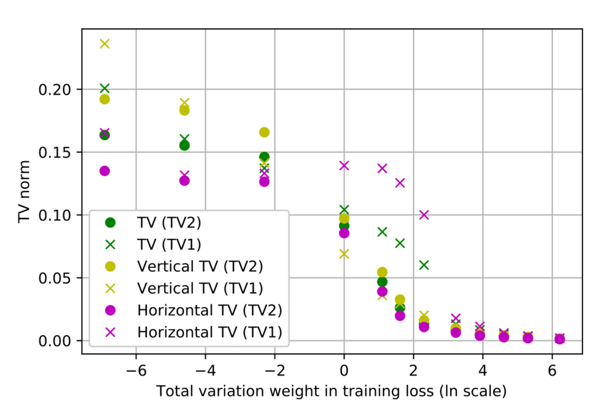

Fig. 3. Effect of increasing smoothness loss on the smoothness of painting commands, vertically, horizontally, and in average directions.

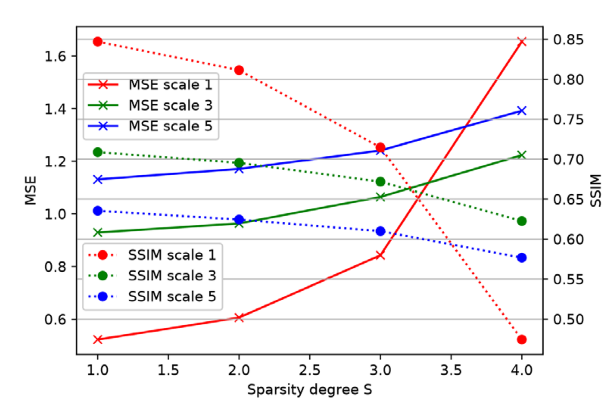

Fig. 4. Effect of sparsity in commands on the reconstruction quality, for different scales of spray pattern sizes.

4 Conclusion

We present a method for inferring spray paint commands to paint a desired texture, specified as an input image. Our self-supervised approach does not require training data beyond an easy-to-obtain unlabeled texture image dataset, and is generalizable to any paint application method. Only a one-time training is required for a new robotic painting system, and painting commands can be inferred for any image with no further processing.

Our approach could also be:
- applied to image inverse problems such as deblurring, deconvolution, or dehazing [5], if the simulation is made differentiable.
- used for domain adaptation [6], as the autoencoder produces a style transfer, generating not an image in a spray-painted style, but from the actual spray patterns of a pre-defined paint system.

5 References

[1] Anurag Sai Vempati, Mina Kamel, Nikola Stilinovic, Qixuan Zhang, Dorothea Reusser, Inkyu Sa, Juan Nieto, Roland Siegwart, and Paul Beardsley, "Paintcopter: An autonomous UAV for spray painting on three-dimensional surfaces," IEEE Robotics and Automation Letters, vol. 3, no. 4, 2018.
[2] Diederik P Kingma and Max Welling, "Auto-encoding variational Bayes," arXiv preprint arXiv:1312.6114, 2013.
[3] Curtis R Vogel and Mary E Oman, "Fast, robust total variation-based reconstruction of noisy, blurred images," IEEE transactions on image processing, vol. 7, no. 6, pp. 813–824, 1998.
[4] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli, "Image quality assessment: from error visibility to structural similarity," IEEE transactions on image processing, vol. 13, no. 4, pp. 600–612, 2004.
[5] Majed El Helou, Frederike Dumbgen, Radhakrishna Achanta, and Sabine Susstrunk, "Fourier-domain optimization for image processing," arXiv preprint arXiv:1809.04187, 2018.
[6] Amir Atapour-Abarghouei and Toby P Breckon, "Real-time monocular depth estimation using synthetic data with domain adaptation via image style transfer," in Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), vol. 18, 2018, p. 1.