Feature-Based Metrics for Exploring the Latent Space of Generative Models

Samuli Laine
NVIDIA
slaine@nvidia.com

Introduction

- Given a generative model, how is motion in latent space $\mathcal{Z}$ related to changes in output space $\mathcal{X}$?
  - How to interpolate generated images in a perceptually meaningful way?
- *Naïve solution:* Linearly interpolate in latent space $\mathcal{Z}$.
- *Previous work:* Find path in $\mathcal{Z}$ such that path length in $\mathcal{X}$ is minimized.
  - I.e., find shortest path in $\mathcal{X}$ that is on generator's output manifold.
  - **Problem:** Euclidean L_2 metric in pixel space $\mathcal{X}$ is a bad measure of perceptual differences.
  - With L_2, the "best" solution would be a cross-fade between images.

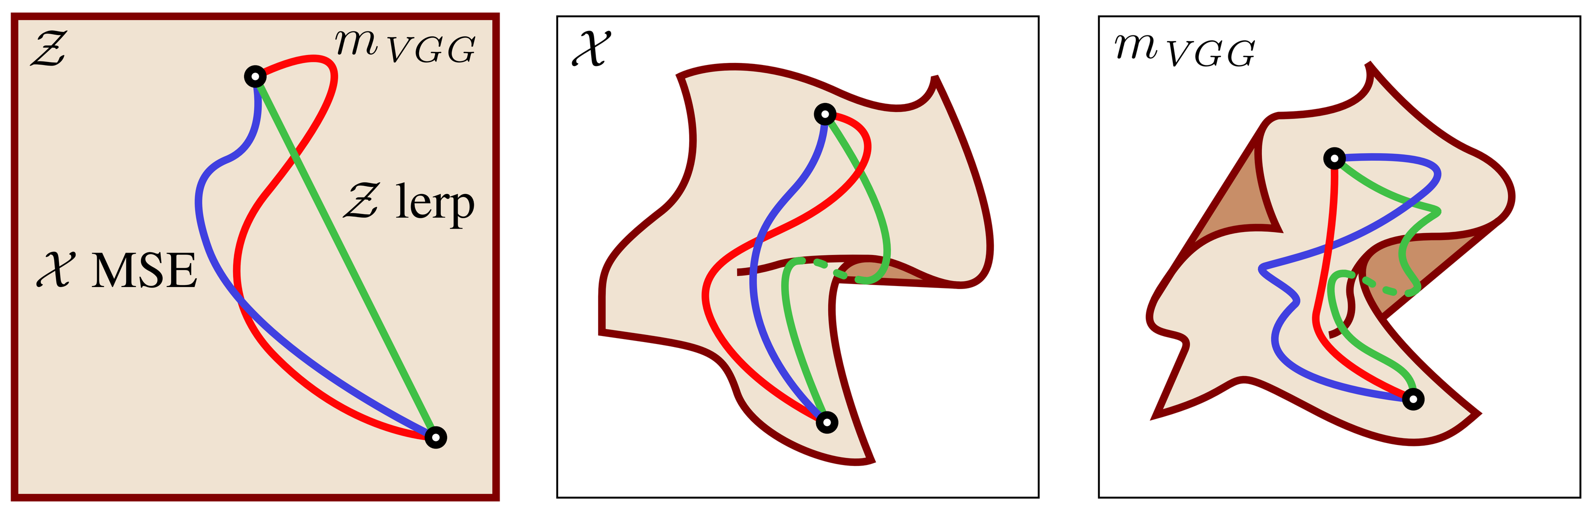

Latent space $\mathcal{Z}$ | Pixel space $\mathcal{X}$ | Space induced by metric m_{VGG}

Contributions

- We replace L_2 metric in $\mathcal{X}$ by a **VGG-19 -based feature-space metric**.
  - This yields paths in $\mathcal{Z}$ that minimize perceptual changes in output images.
- To prevent a failure mode where image gets darker at the middle of the path, we **equalize brightness and contrast** prior to evaluating the metric.
  - Denoted as $\widetilde{\mathcal{X}}$ MSE and $\widetilde{m}_{VGG}$ in the images on the right.
- **Progressive path subdivision** allows finding minimal paths efficiently.
- Experiments using a state-of-the-art GAN show that the proposed method results in more consistent interpolations.

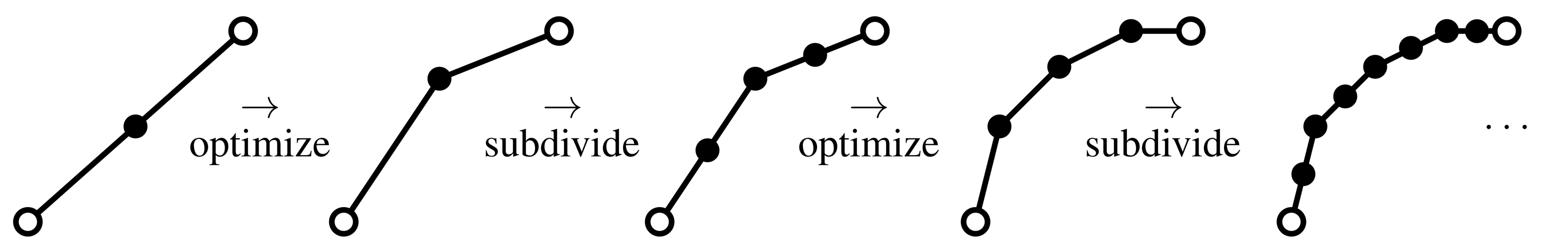

Results

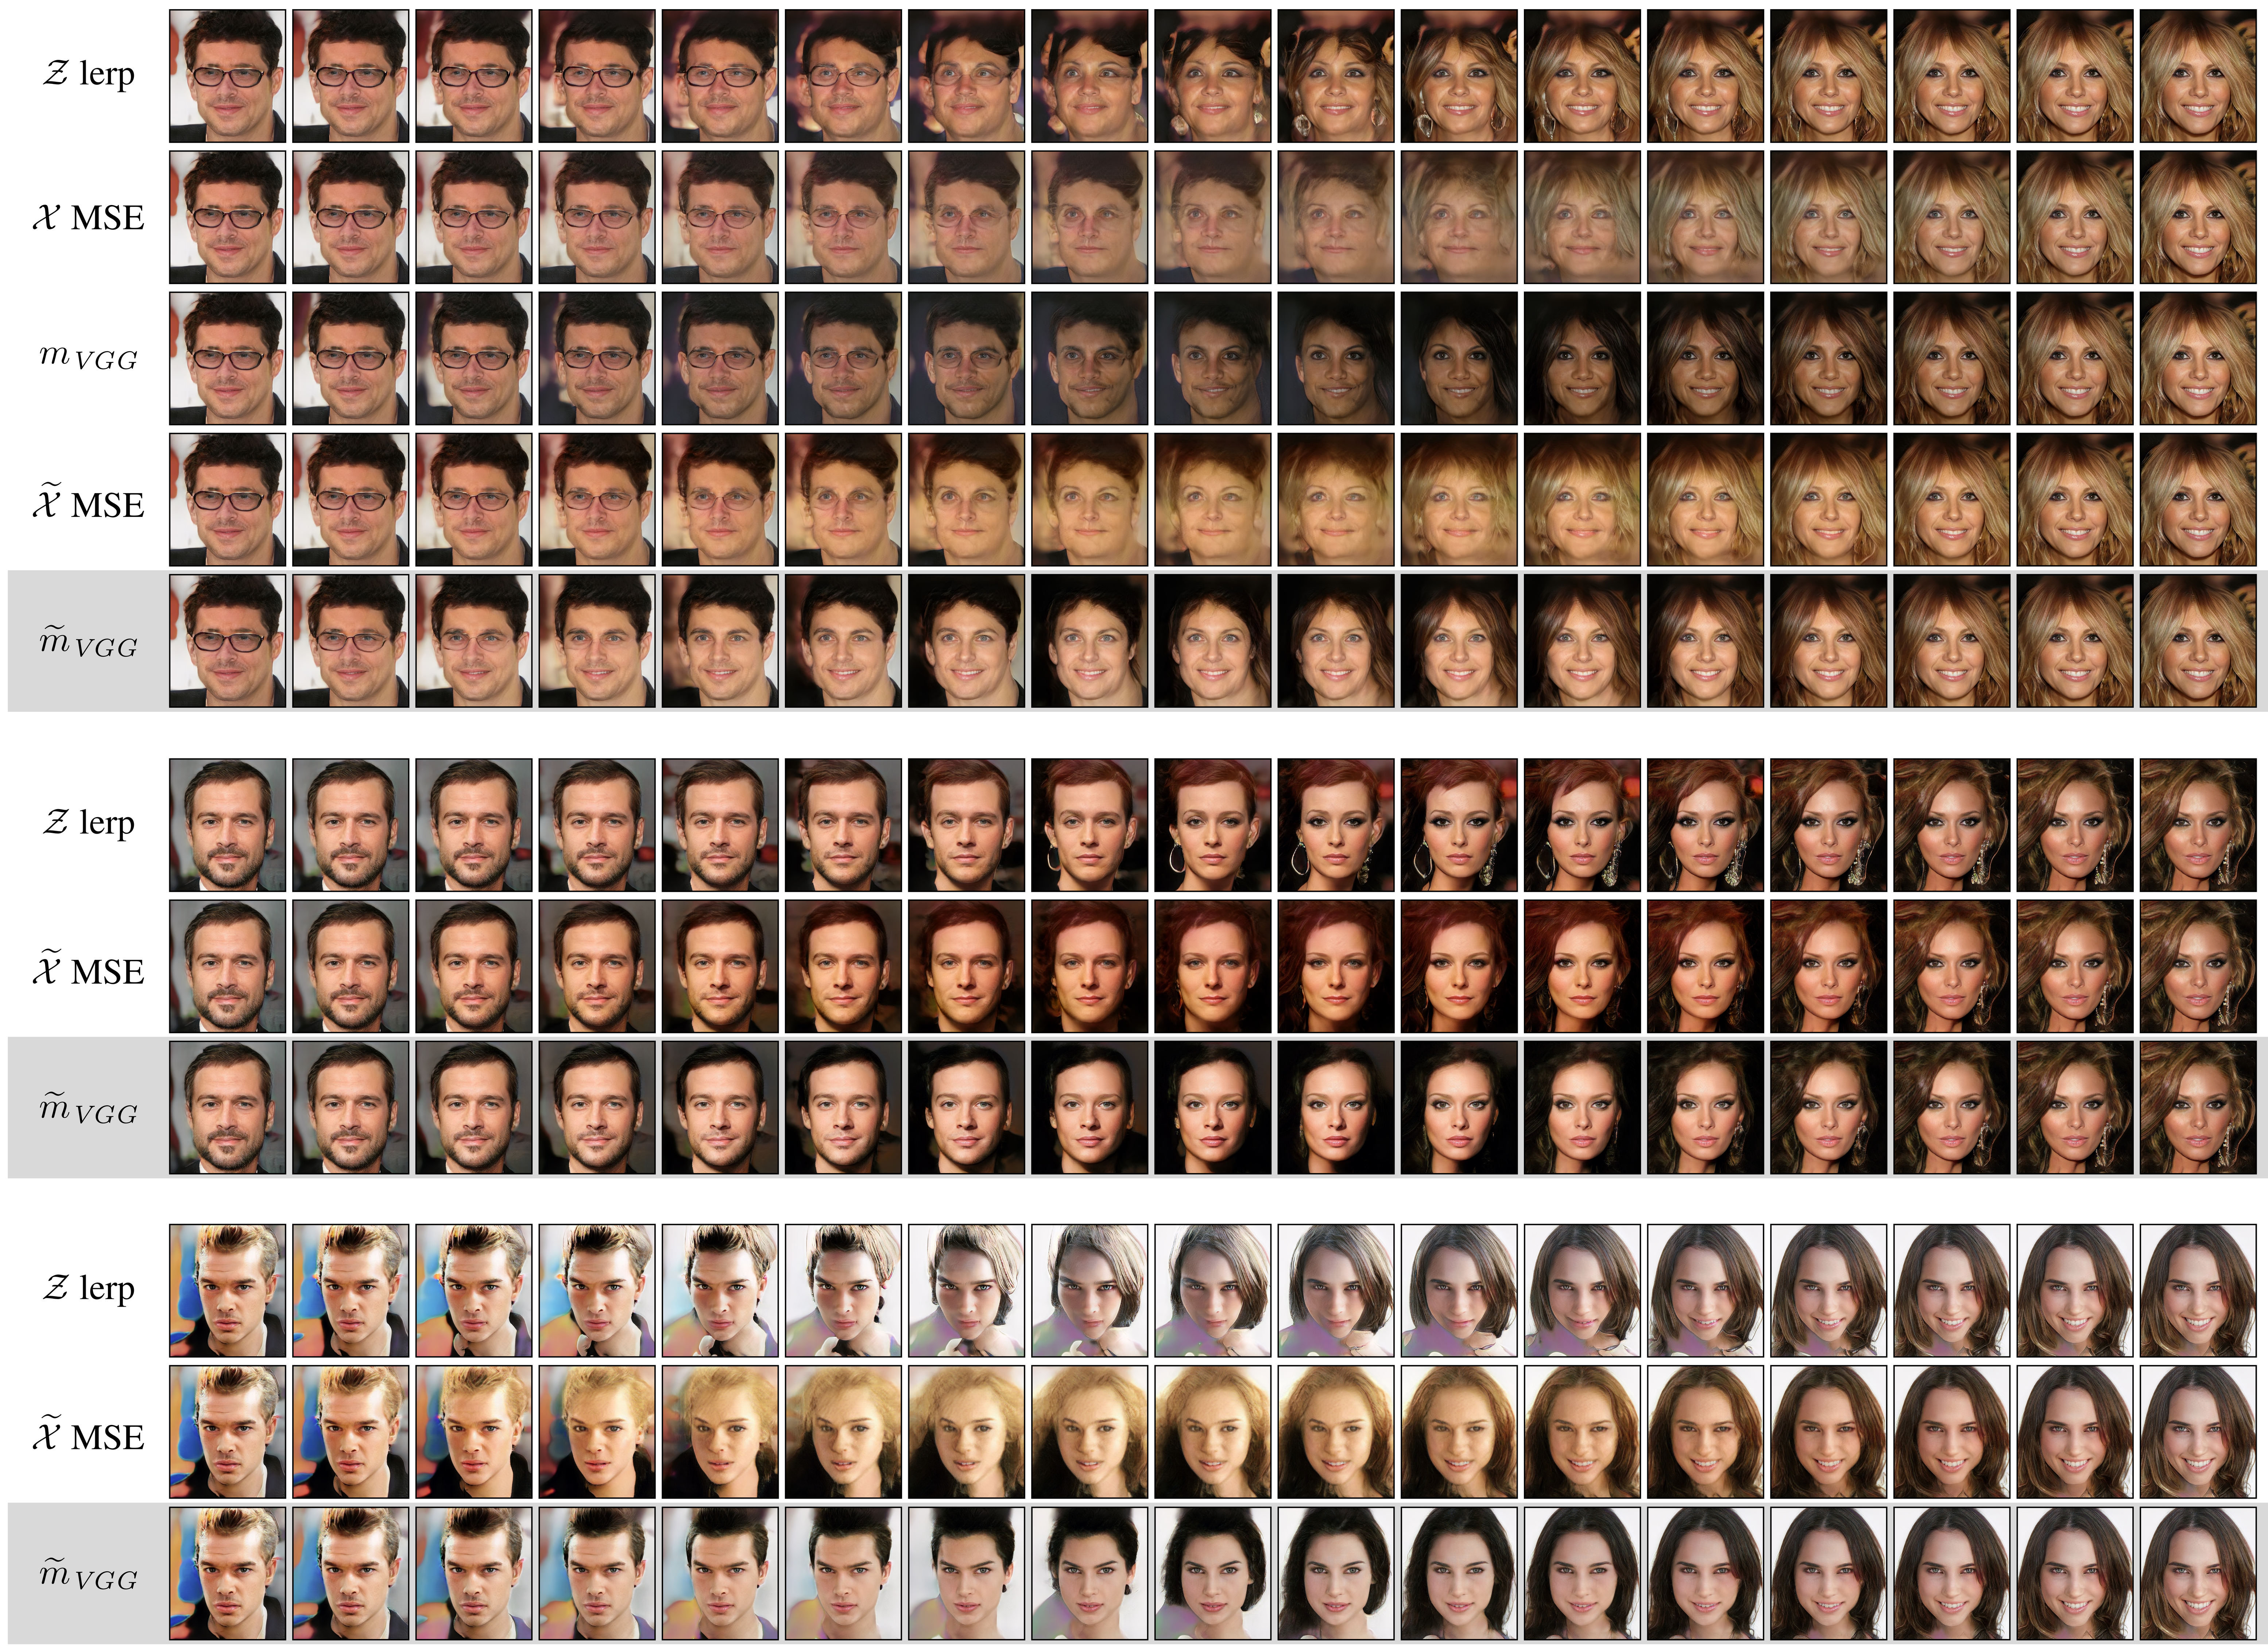

References

Georgios Arvanitidis, Lars Kai Hansen, and Søren Hauberg. Latent space oddity: on the curvature of deep generative models. *ICLR*, 2018.

Nutan Chen, Alexej Klushyn, Richard Kurle, Xueyan Jiang, Justin Bayer, and Patrick van der Smagt. Metrics for deep generative models. *CoRR*, abs/1711.01204, 2017.

Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. *ICCV*, 2017.

Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, 2016.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks. In *NIPS*, 2014.

Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. *ICLR*, 2018.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.

Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. *IJCV*, 2015.

Hang Shao, Abhishek Kumar, and P. Thomas Fletcher. The Riemannian geometry of deep generative models. *CoRR*, abs/1711.08014, 2017.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.

Paul Upchurch, Jacob Gardner, Geoff Pleiss, Robert Pless, Noah Snavely, Kavita Bala, and Kilian Weinberger. Deep feature interpolation for image content changes. *CVPR*, 2017.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. *CoRR*, abs/1801.03924, 2018.