Deep Feature Interpolation for Image Content Changes

A new baseline for automatic high-resolution image content transformation

Paul Upchurch[1,2], Jacob Gardner[1,2], Geoff Pleiss[2], Robert Pless[3], Noah Snavely[2], Kavita Bala[2], Kilian Q. Weinberger[2]

[1] Authors contributed equally [2] Cornell University [3] George Washington University

Goal: Edit Image Content

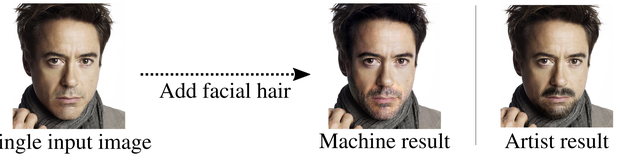

Deep Feature Interpolation (DFI)

Applies the feature space difference of two image sets to the input image.

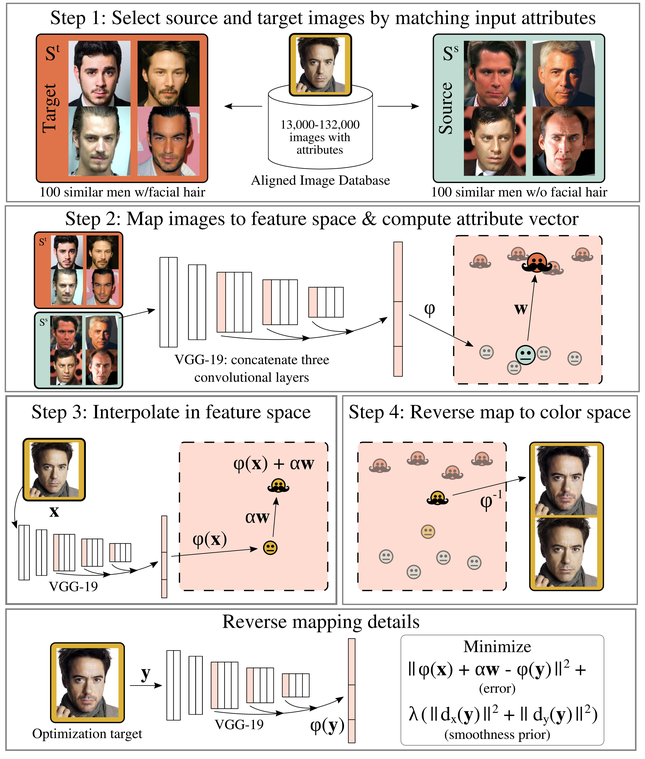

Code: https://github.com/paulu/deepfeatinterp/

Aging a Face, Changing Expression, Adding Moustache and Glasses

Model-free editing (no training step). Across a row: identity is preserved. Down a column: the desired attribute is added.

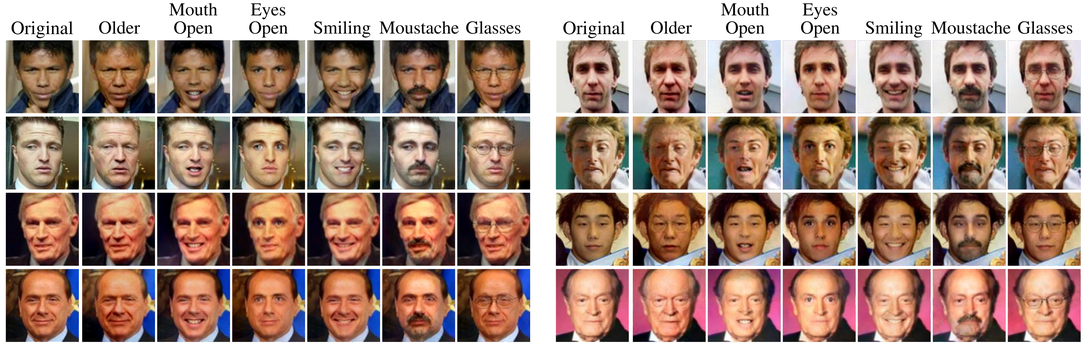

Inpainting Faces and Shoes

Missing region is filled with plausible content.

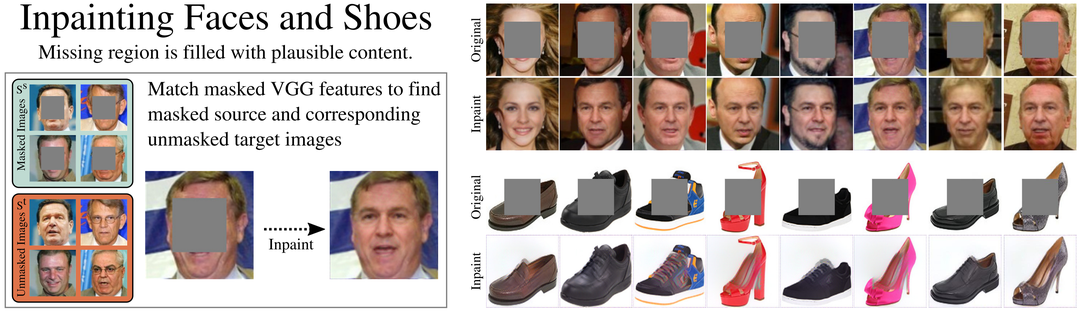

Megapixel-scale Editing

Convolutional features work with any input size. Memory and processing time scale linearly.

Age Progression

Interpolation steps are smooth and produce stable videos.

Data-driven Results

Attribute matching relates correlated attributes so young men get dark facial hair and seniors get gray facial hair.

Perceptual Study vs a Generative GAN

Both methods use linear interpolation in a feature space and the same image database. Numbers are the ratio by which DFI is preferred.

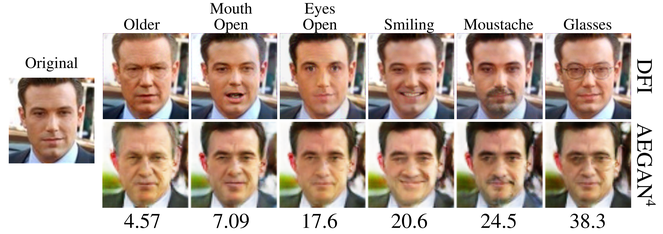

Conclusions

Attribute matching combined with a discriminative convolutional feature space is surprisingly good at editing the content of photographic images.

Acknowledgments

Thanks to support by National Science Foundation grants (III-1618134, III-1526012, IIS-1149882, IIS-1617861, IIS-1011919, IIS-1161645, IIS-1149393, IIS-1149393) and the Bill and Melinda Gates Foundation.

[4] Larsen, Sønderby, Larochelle, and Winther. Autoencoding beyond pixels using a learned similarity metric.