Deep Feature Interpolation for Image Content Changes: Supplemental

1. Changing Face Attributes

In this section we provide details on the test images. We use attributes predicted by a machine learning model to perform attribute changes on aligned LFW faces. The attributes are in the form of scalar decision values which we convert into boolean attributes by assigning *True* to the top two-thirds of images whose decision value is positive. The LFW images are aligned using the same code used by AEGAN to facilitate comparison.[1] We first determined whether each attribute was more commonly associated with men or women, since the majority of images in LFW are of men. In a purely data-driven fashion we found 38 test images randomly selected according to attribute constraints, allowing us to apply multiple transformations to the same face. For example, if we wish to transform a test image to have a mustache and separately to have glasses, the test image would ideally not already have a mustache or glasses. Additional test results are shown in Figure 2. Output resolution is 200×200.

The 38 test images are: Melchor Cob Castro 0001, John Stockton 0001, Ralf Schumacher 0005, Charlton Heston 0002, Tom Ridge 0032, Silvio Berlusconi 0023, Brad Johnson 0001, Kaoru Hasuike 0001, George W Bush 0156, George W Bush 0471, George W Bush 0334, Gian Marco 0003, Tony Blair 0093, Lachlan Murdoch 0001, George W Bush 0151, Juan Carlos Ferrero 0017, Hugo Colace 0001, George W Bush 0433, Daniel Bruehl 0001, Robinson Stevenin 0001, Takahiro Mori 0001, George W Bush 0190, Ben Broussard 0001, Tomomi Morita 0001, Arturo Gatti 0001, Charles Tannok 0001, Melinda Czink 0001, David Coulthard 0001, Bob Hope 0007, Masaru Hayami 0001, George W Bush 0123, George W Bush 0453, Will Self 0001, Scott McClellan 0001, Tony Blair 0098, George W Bush 0211, Bill Nelson 0001, and Ben Affleck 0001.

2. High Resolution Face Editing

In Figure 4 and Figure 5 we show additional results on high resolution face images for the tasks of adding facial hair and aging.

Figure 1. Perceptual study tutorial test. Workers are instructed to pick the image that adds facial hair to the original image.

3. Perceptual Study Tutorial

In this section we describe the tutorial used to train workers to evaluate face attribute chages. Workers were instructed to pick the image which has the target attribute and preserves the unique characteristics of the original image. They were instructed to use quality only as a tie-breaker.

Workers were required to pass a test before participating. The test image was a picture of Robert Downey Jr. (RDJ) and the 4 pictures to chose from were: A. a RDJ lookalike with facial hair (Jeffrey Dean Morgan), B. the original RDJ picture without facial hair, C. a different picture of RDJ with facial hair, and D. the original RDJ picture edited to add facial hair (Figure 1). If workers picked the wrong image then they were shown an explanation of why it was wrong and were not allowed to proceed until they selected image D.

4. Inpaint Without Attributes

In Figure 3 we show additional inpainting results (random test images). Source and target sets are found by taking the $K = 100$ nearest neighbors (by cosine distance) in VGG-19 `pool5` space. Output resolution is 200×200.

[1] DeepPy: https://github.com/andersbll/deeppy

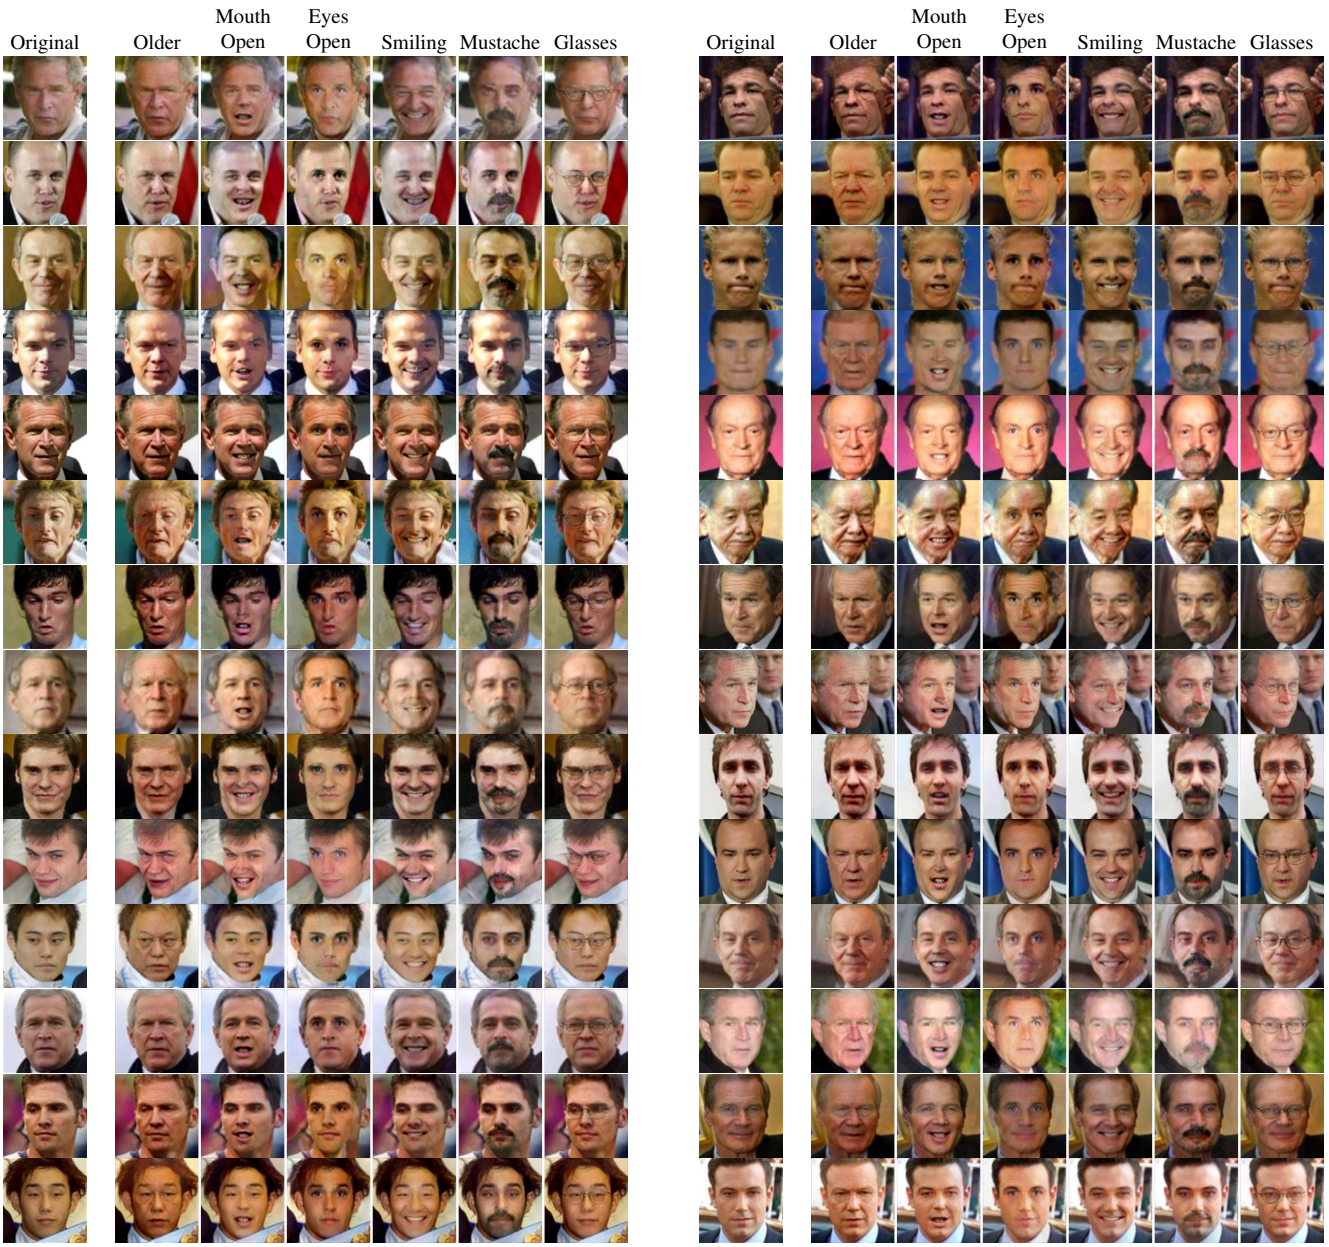

Figure 2. **(Zoom in for details.)** Attribute changes on aligned faces with Deep Feature Interpolation. Looking down a column the face should have the target attribute. Looking across a row the image should preserve the unrelated attributes of the original image as much as possible (e.g., the person's identity, clothing and background).

Figure 3. Image inpainting with DFI. The missing regions are filled with plausible pixel values although partial objects are not completed (e.g., eyeglasses). These results were produced without supervised attributes.

Figure 4. **(Zoom in for details.)** High resolution results for the task of adding facial hair. Each row shows a different image. Each column represents a different value of α. In each row, the first column is the original image.

Figure 5. **(Zoom in for details.)** High resolution results for the task of aging faces. Each row shows a different image. Each column represents a different value of α. In each row, the first column is the original image.