Supplemental Document for Style Transfer for Headshot Portraits

1 Description of this document

In the title of each paragraph, we put the section number referenced in the main paper.

Additional comparisons to related work (Sec. 4.1) Figure 1 shows the comparisons on an extreme and low-key style. Without adaptation to the face mask, all the global methods fail in this case, since the background in the input is brighter than the foreground, but vice versa in the example. For fair comparison, we adapted the related methods to face mask. We replaced the input with the example background for Bae et al. [2006] and PhotoShop MatchColor. We limit the transfer in the face region defined by the mask for Sunkavalli et al. [2010], Pitié et al. [2005], and Reinhard et al. [2001]. For Sunkavalli et al., we started by their setup demonstrated on face portraits, and tested a few options. We found that disabling noise matching produces the best result. For Pitíe et al. [2005], we ran 30 iterations. We also tried HaCohen et al. [2011], but their implementation reports that no matching is found.

Figure 2 shows the comparison on a nearly all-black-and-white style. Our method transfers the right amount of details and brightness without being over-exposed or under-exposed. We used the same adaptation for the related methods. We also tried HaCohen et al. [2011], but their implementation again reports that no matching is found in this case.

Figure 3 shows the comparison on a color style to two methods adapted to face mask, as described above. The comparisons with the unadapted methods are in the main paper.

Figure 4 shows a close-up comparison between our result and the example. Our result matches well on lighting, color, facial details in all scales.

Figure 5 compares to HaCohen et al. [2011] on an example that their method finds non-empty matchings. The inset in Fig. 5(d) shows the matching area. Among all the examples used in our project, this example has the largest matching region.

Comparison to reference image (Sec. 4) Figure 6 shows comparison to a reference image. We use the reference as example. This is to show the ideal situation that the database is sufficiently large such that we can find an example almost identical to the input.

Comparison on makeup transfer (Sec. 4.2) Figure 7 shows the comparison on our extension to makeup transfer. Fig. 7(c) shows our original method before modification. The green eye shadow is bled to the sclera (the white of the eye). Our adapted method automatically transfers the sclera from the input to fix the problem, as shown in Fig. 7(d). The rest of Fig. 7 compares our result with two state-of-the-art methods designed for makeup transfer [Tong et al. 2007; Guo and Sim 2009]. All three methods achieve plausible results. Tong et al. require the before image of the example makeup, which is not shown here. Their results are directly taken from Guo and Sim's paper.

Additional results on automatic selection algorithm (Sec. 4) Figure 8 shows the style transfer results using the top four examples selected by our automatic algorithm. We show three styles.

User correction (Sec. 4) Our dense matching using computer vision techniques often produces satisfactory results. However, there are cases where matching is challenging, such as matching long hair to short hair in Figure 9. In this case, we provide users a manual correction work flow by using an user-created constraint map. Then our algorithm re-run the transfer, but this time we assign the energy gains of each pixel in the red region by the average of the gains in the green region. This process can be repeated as needed for additional corrections. To avoid discontinuities, we filter the gain map with a small Gaussian kernel after applying the constraint map. Figure 9d shows the successful result after user correction. In our results on Flickr data set, 5 out of 94 are corrected in such a way. All the results in the main paper are generated automatically; we did not correct them.

The artifacts due to transferring the example identity Figure 10 shows a failure case that the identity of the example is transferred to the input. This may occur when the example identity has different genders or very different skin colors.

2 Accompanying web pages

Massive results using Flickr data set (Sec. 4)
In `results_on_flickr_1.html`, we use inputs downloaded from an online web site, Flickr, on three different styles. The data collection workflow is described in the main paper (Sec. 4). The data set contains 94 images with various facial contents, expressions, under arbitrary lighting conditions. All inputs are under creative commons license.

3 Accompanying Video (Sec.4.2)

`supplemental.mp4` shows our video style transfer extension. We test two different inputs with moderate motion and extreme facial expressions, using three different styles. No audio in the video.

References

BAE, S., PARIS, S., AND DURAND, F. 2006. Two-scale tone management for photographic look. In *ACM Trans. Graphics*.

GUO, D., AND SIM, T. 2009. Digital face makeup by example. In *IEEE Conf. Computer Vision and Pattern Recognition*.

HACOHEN, Y., SHECHTMAN, E., GOLDMAN, D. B., AND LISCHINSKI, D. 2011. Non-rigid dense correspondence with applications for image enhancement. In *ACM Trans. Graphics*, vol. 30, ACM, 70.

PITIÉ, F., KOKARAM, A. C., AND DAHYOT, R. 2005. N-dimensional probability density function transfer and its application to color transfer. In *IEEE Conference on Computer Vision*.

REINHARD, E., ADHIKHMIN, M., GOOCH, B., AND SHIRLEY, P. 2001. Color transfer between images. *IEEE Computer Graphics and Applications 21*, 5, 34–41.

SUNKAVALLI, K., JOHNSON, M. K., MATUSIK, W., AND PFISTER, H. 2010. Multi-scale image harmonization. *ACM Trans. Graphics 29*, 4, 125.

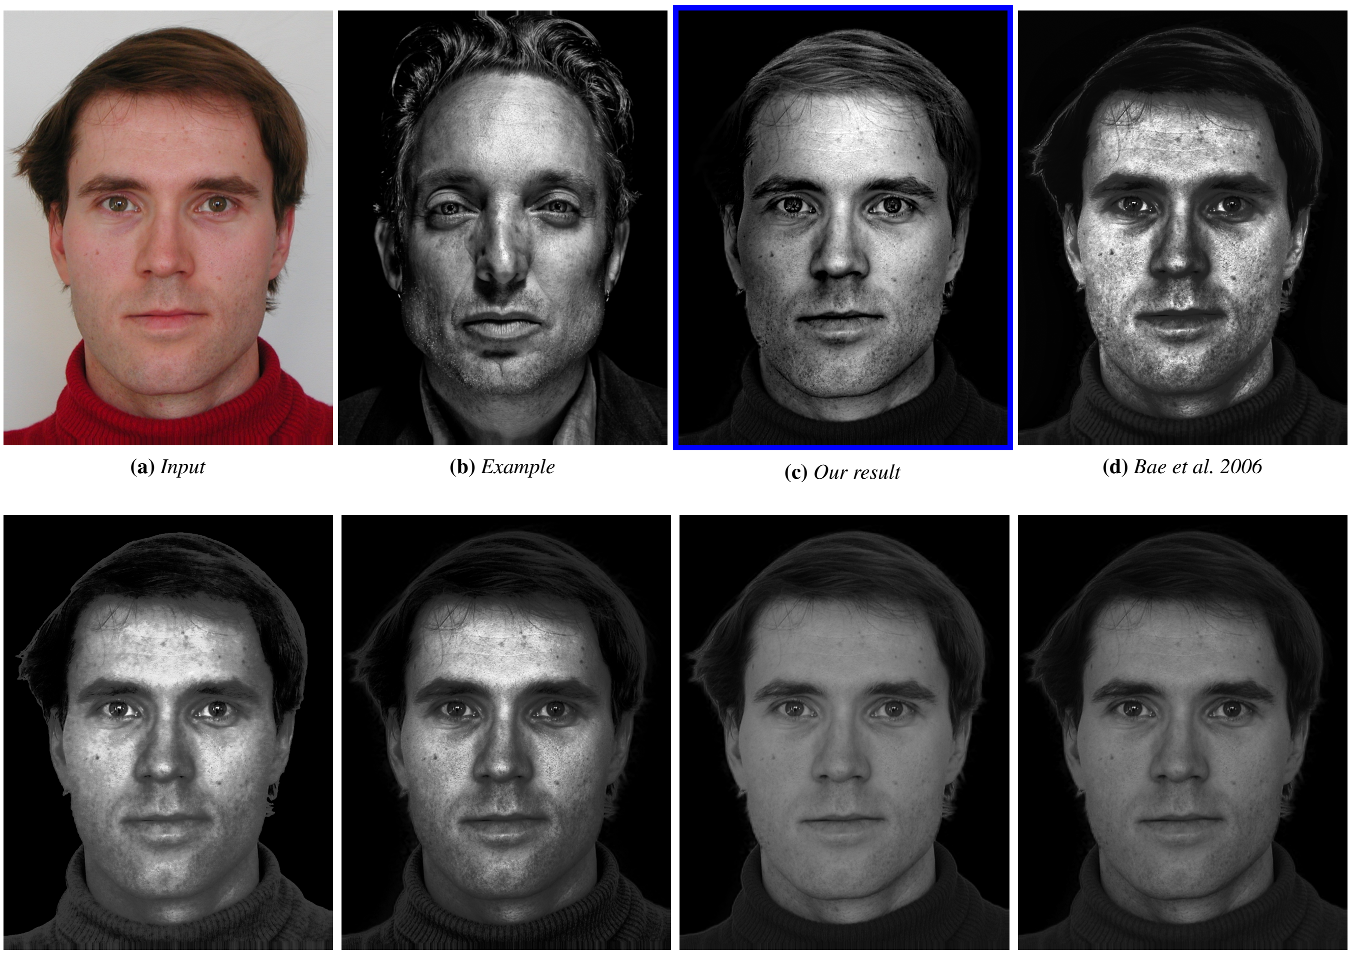

(a) *Input* **(b)** *Example* **(c)** *Our result* **(d)** *Bae et al. 2006*

(e) *Sunkavalli et al. 2011* **(f)** *Pitié et al. 2005* **(g)** *Reinhard et al. 2001* **(h)** *Photoshop Match Color*

Figure 1: *We show comparisons on an extreme style, using related methods adapted to face mask. We replaced the input with the example background for Bae et al. and PhotoShop MatchColor. We limit the transfer in the face region defined by the mask for Sunkavalli et al. Pitié et al. , and Reinhard et al.. Our method captures the smoothly fall-off lighting on the forehead and details on the face.*

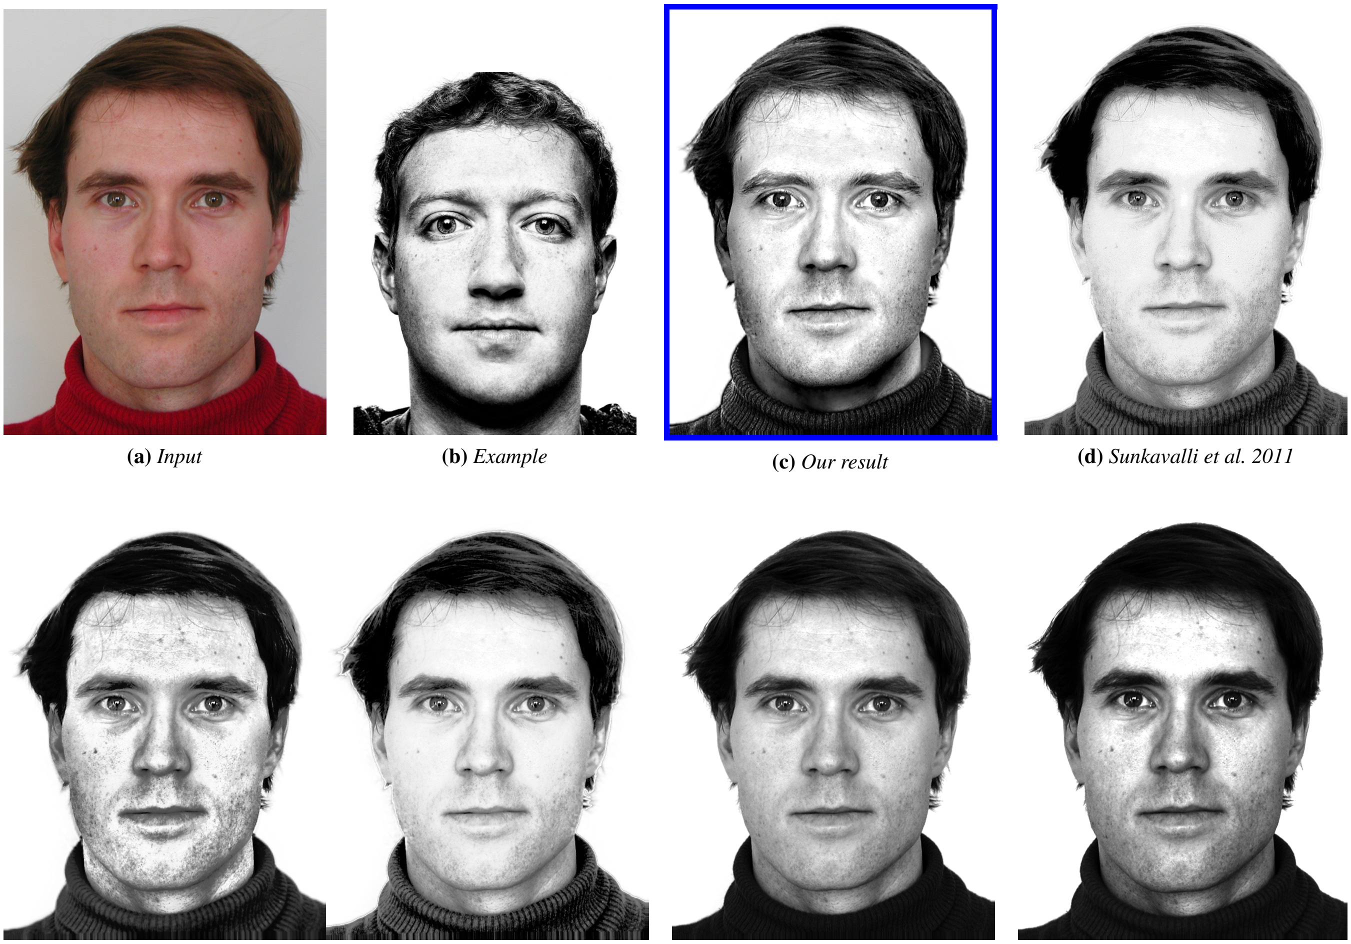

(a) *Input* **(b)** *Example* **(c)** *Our result* **(d)** *Sunkavalli et al. 2011*

(e) *Bae et al. 2006* **(f)** *Pitié et al. 2005* **(g)** *Reinhard et al. 2001* **(h)** *PhotoShop Match Color*

Figure 2: *We show a comparison on a nearly all-black-and-white style. Our method captures the right amount of exposure and details on the face and hair.*

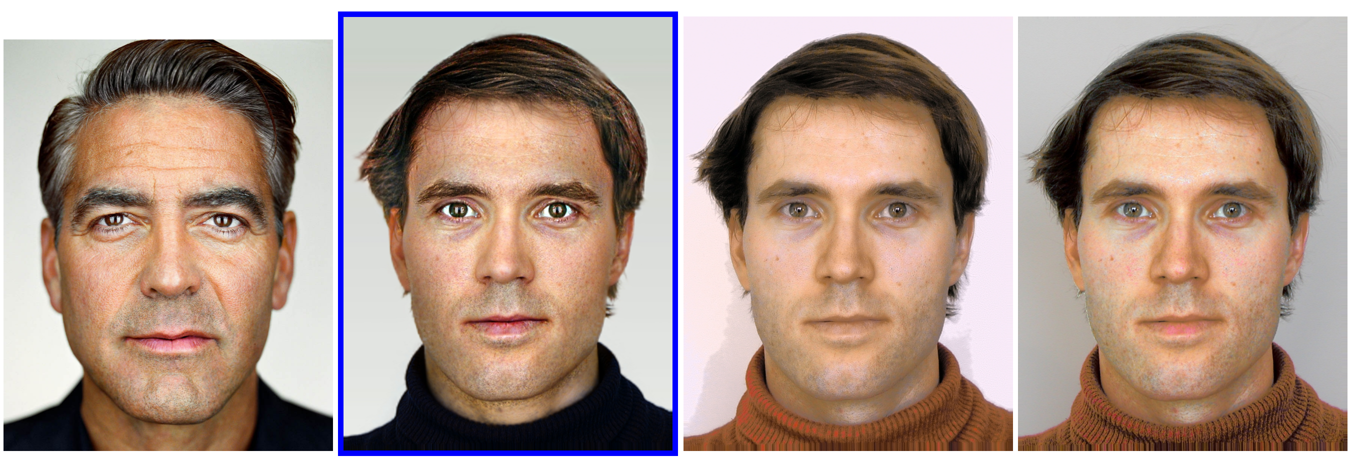

(a) *Example* **(b)** *Our result* **(c)** *Sunkavalli et al. 2011* **(d)** *Pitíe et al. 2005*

Figure 3: *Using the input in Fig. 2, we compare two methods adapted to face mask on a color style. The comparison to the unadapted methods are in the main paper.*

(a) *Example* **(b)** *Our result*

Figure 4: *Close-up comparison to the example, using the input in Fig. 2. Our result matches well on lighting, color, facial details in all scales.*

(a) *Input* **(b)** *Example* **(c)** *Our result* **(d)** *HaCohen et al. 2011*

Figure 5: *We compare to HaCohen et al. on a case that their method finds matching region, shown in the inset in (d). This example is has the largest matching region among all examples used in the paper.*

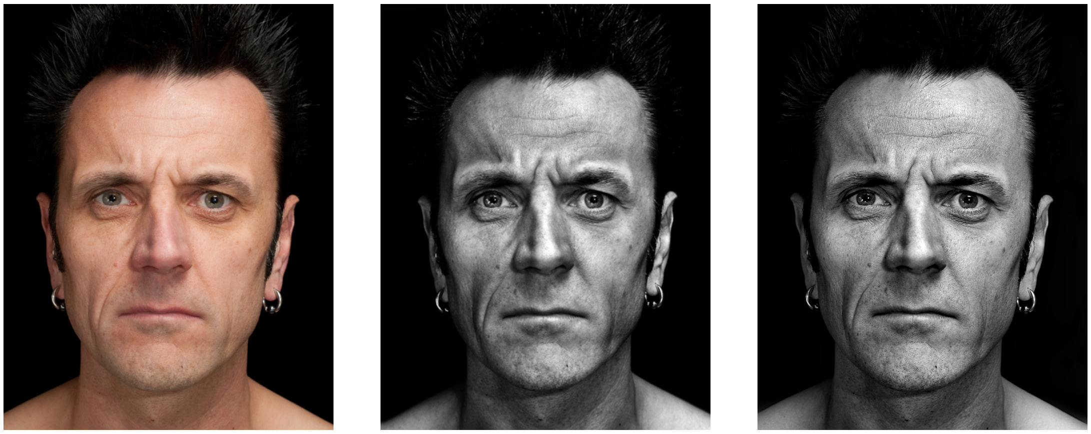

(a) *Input (before editing)*

(b) *Example (after editing), reference image.*

(c) *Our result*

Figure 6: *We test the "upper bound" of our method by using a pair of before/after editing images in (a) and (b) as input and example. Our result (c) is visually close to (b). This is to simulate the ideal situation that we can find an example subject whose look is very close to the input.*

TONG, W.-S., TANG, C.-K., BROWN, M. S., AND XU, Y.-Q. 2007. Example-based cosmetic transfer. In *IEEE Pacific Graphics*.

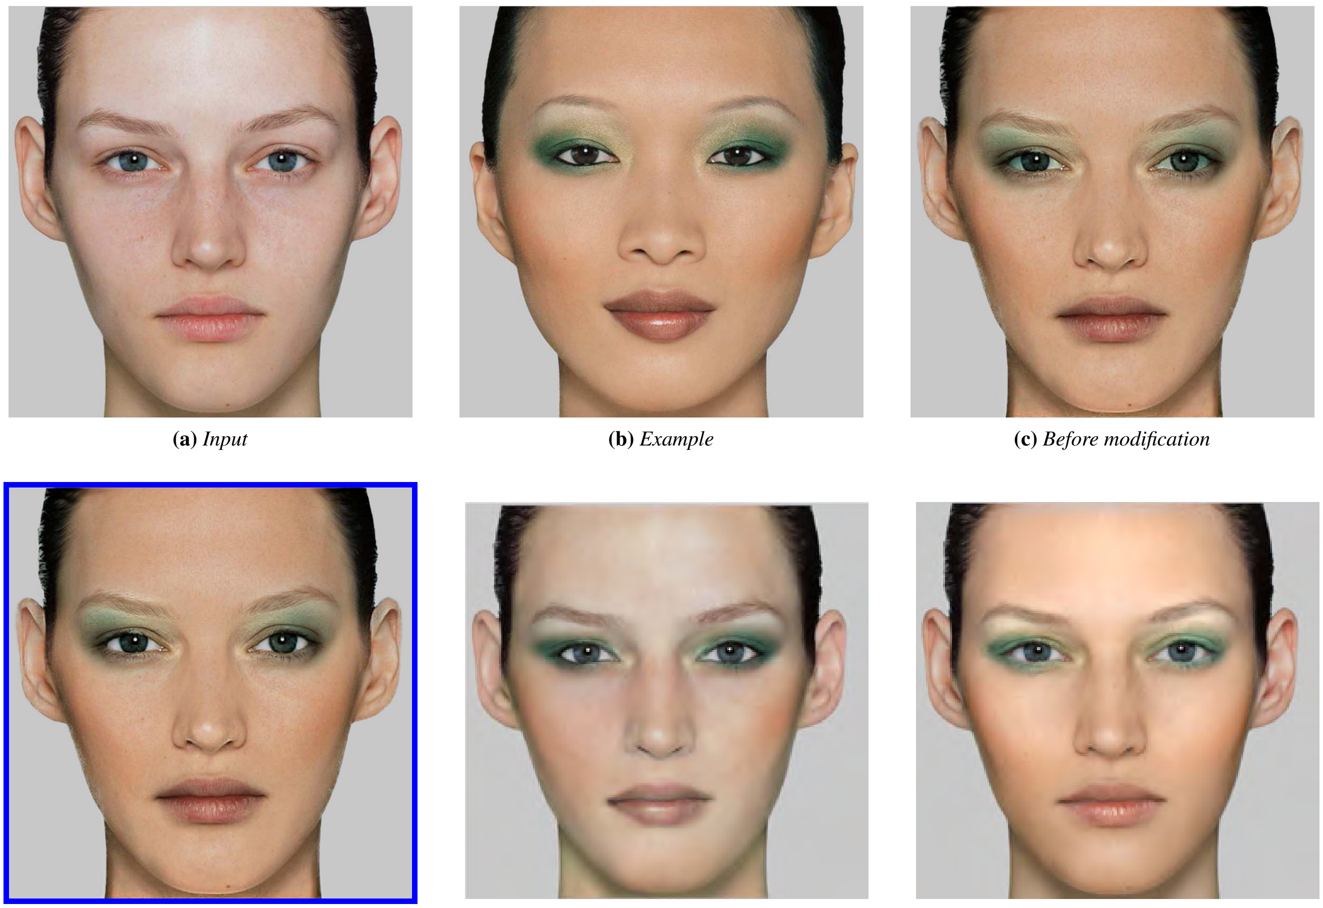

(a) *Input* **(b)** *Example* **(c)** *Before modification*

(d) *Our final result* **(e)** *Tong et al. 2007, require before image of (b)* **(f)** *Guo et al. 2009*

Figure 7: *We extend our method to makeup transfer. Directly using our algorithm results color bleeding on eyes (c). With minor modification that handles eye sclera (eye white), we can achieve better result (d). We show comparison with two state-of-art methods designed for makeup transfer. (e) requires before image of (b), which is not shown here. (f) explicitly models foundation, eye shadow and lip color. All results achieve plausible makeup transfer. (e) and (f) are directly taken from their papers, respectively.*

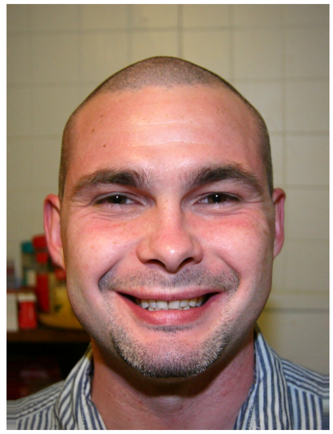

(a) *Input*

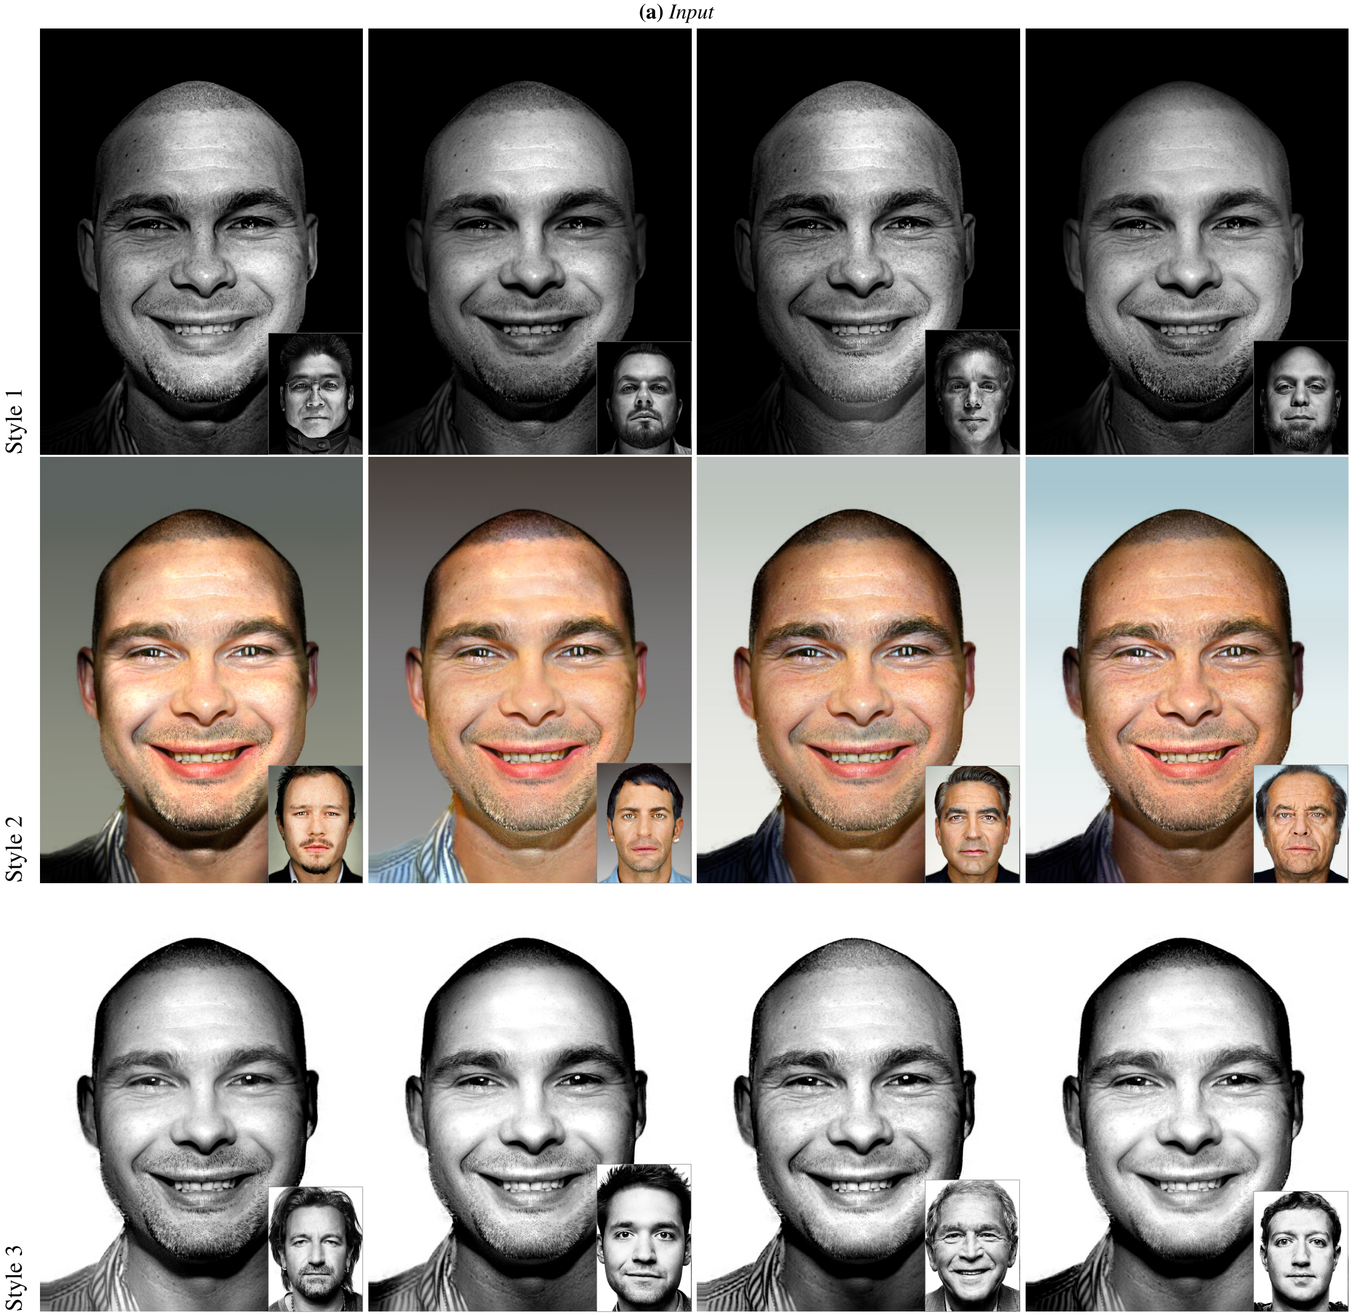

Figure 8: *We show style transfer results on the input in (a), using different styles in the three rows. We use the top four examples selected by our automatic selection algorithm, shown in the insets.*

(a) *Input* **(b)** *Example* **(c)** *Result failed on hair* **(d)** *Corrected result using a user-provided constraint mask in the blue box*

Figure 9: *Our transfer can fail if the input (a) and example (b) have very different hair styles, and cause artifacts on the hair in (c). We demonstrate that the user can fix this in (d) by providing a constraint map in the blue box. This map constrains that the gains of the red region to be the same as those of the green region*

(a) *Input* **(b)** *Failed output*

Figure 10: *A failure case that the identity of the example (inset in (a)) is transferred to the output.*