

Effective Face Frontalization in Unconstrained Images

Tal Hassner¹, Shai Harel¹ †, Eran Paz¹ † and Roei Enbar²


¹The open University of Israel. ²Adience.



Figure 1: **Frontalized faces.** Top: Input photos; bottom: our frontalizations, obtained without estimating 3D facial shapes.

“Frontalization” is the process of synthesizing frontal facing views of faces appearing in single unconstrained photos. Recent reports have suggested that this process may substantially boost the performance of face recognition systems. This, by transforming the challenging problem of recognizing faces viewed from unconstrained viewpoints to the easier problem of recognizing faces in constrained, forward facing poses. Previous frontalization methods did this by attempting to approximate 3D facial shapes for each query image. We observe that 3D face shape estimation from unconstrained photos may be a harder problem than frontalization and can potentially introduce facial misalignments. Instead, we explore the simpler approach of using a *single, unmodified*, 3D surface as an approximation to the shape of *all input faces*. We show that this leads to a straightforward, efficient and easy to implement method for frontalization. More importantly, it produces aesthetic new frontal views and is surprisingly effective when used for face recognition and gender estimation.

Observation 1: For frontalization, one rough estimate of the 3D facial shape seems as good as another, demonstrated by the following example:



The top row shows surfaces estimated for the same query (left) by Hassner [2] (mid) and DeepFaces [6] (right). Frontalizations are shown at the bottom using our single-3D approach (left), Hassner (mid) and DeepFaces (right). Clearly, both surfaces are rough approximations to the facial shape. Moreover, despite the different surfaces, all results seem qualitatively similar, calling to question the need for shape estimation for frontalization.

Result 1: A novel frontalization method using a single, unmodified 3D reference shape is described in the paper (illustrated in Fig. 2).

Observation 2: A single, unmodified 3D reference shape produces aggressively aligned faces, as can be observed in Fig. 3.

Result 2: Frontalized, strongly aligned faces elevate LFW [5] verification accuracy and gender estimation rates on the Adience benchmark [1].

Conclusion: On the role of 2D appearance vs. 3D shape in face recognition, our results suggest that 3D shape estimation may be unnecessary.

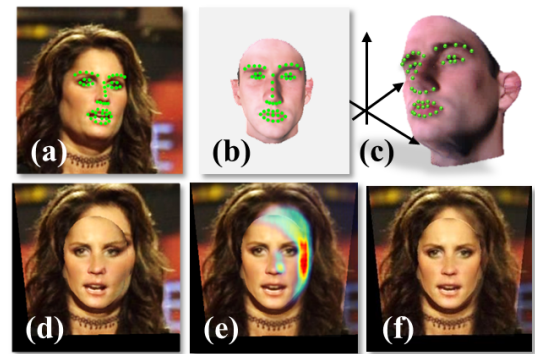


Figure 2: **Frontalization process.** (a) facial features detected on a query face and on a reference face (b) which was produced by rendering a textured 3D, CG model (c); (d) 2D query coordinates and corresponding 3D coordinates on the model provide an estimated projection matrix, used to back-project query texture to the reference coordinate system; (e) estimated self-occlusions shown overlaid on the frontalized result (warmer colors reflect more occlusions.) Facial appearances in these regions are borrowed from corresponding symmetric face regions; (f) our final frontalized result.





Figure 3: Average faces produced from the 41 Laura Bush and 530 George W. Bush images in the LFW set. Columns represent different alignments; from left to right: Funneling [4], LFW-a images (available only in grayscale) [7], deep-funneled images [3] and our own frontalized faces. The crisp details of our results testify to the strong facial feature alignment.

For full text, updated results, code and data please see our project webpage: goo.gl/RAZU67



- [1] Eran Eiding, Roei Enbar, and Tal Hassner. Age and gender estimation of unfiltered faces. *TIFS*, 9(12):2170–2179, 2014. Available: www.openu.ac.il/home/hassner/Adience/data.html.
- [2] Tal Hassner. Viewing real-world faces in 3D. In *ICCV*, 2013. Available: www.openu.ac.il/home/hassner/projects/poses.
- [3] Gary Huang, Marwan Mattar, Honglak Lee, and Erik G Learned-Miller. Learning to align from scratch. In *NIPS*, pages 764–772, 2012.
- [4] Gary B Huang, Vidit Jain, and Erik Learned-Miller. Unsupervised joint alignment of complex images. In *ICCV*. IEEE, 2007.
- [5] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. University of Massachusetts, Amherst, TR 07-49, 2007.
- [6] Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. Deepface: Closing the gap to human-level performance in face verification. In *CVPR*, 2013.
- [7] L. Wolf, T. Hassner, and Y. Taigman. Similarity scores based on background samples. In *ACCV*, 2009.