Cross-modal Deep Face Normals with Deactivable Skip Connections
Supplementary Material

Low-cost depth enhancement: We can use our model to enhance the appearance of the noisy depth data coming from low-cost RGB-D sensors, e.g. Kinect. We show an example of this using the FaceWarehouse dataset [1], where we use the accompanying RGB image to predict normals with our method, and append these normals to the raw depth image pixel-wise using normal mapping [2], thus rendering enhanced geometric shading. In Fig. 1 we show the RGB images in the first row, the raw depth in the second, and the same depth enhanced with our model's predictions in the last one. The ability to recover accurate normals allows to enhance the depth appearance significantly.

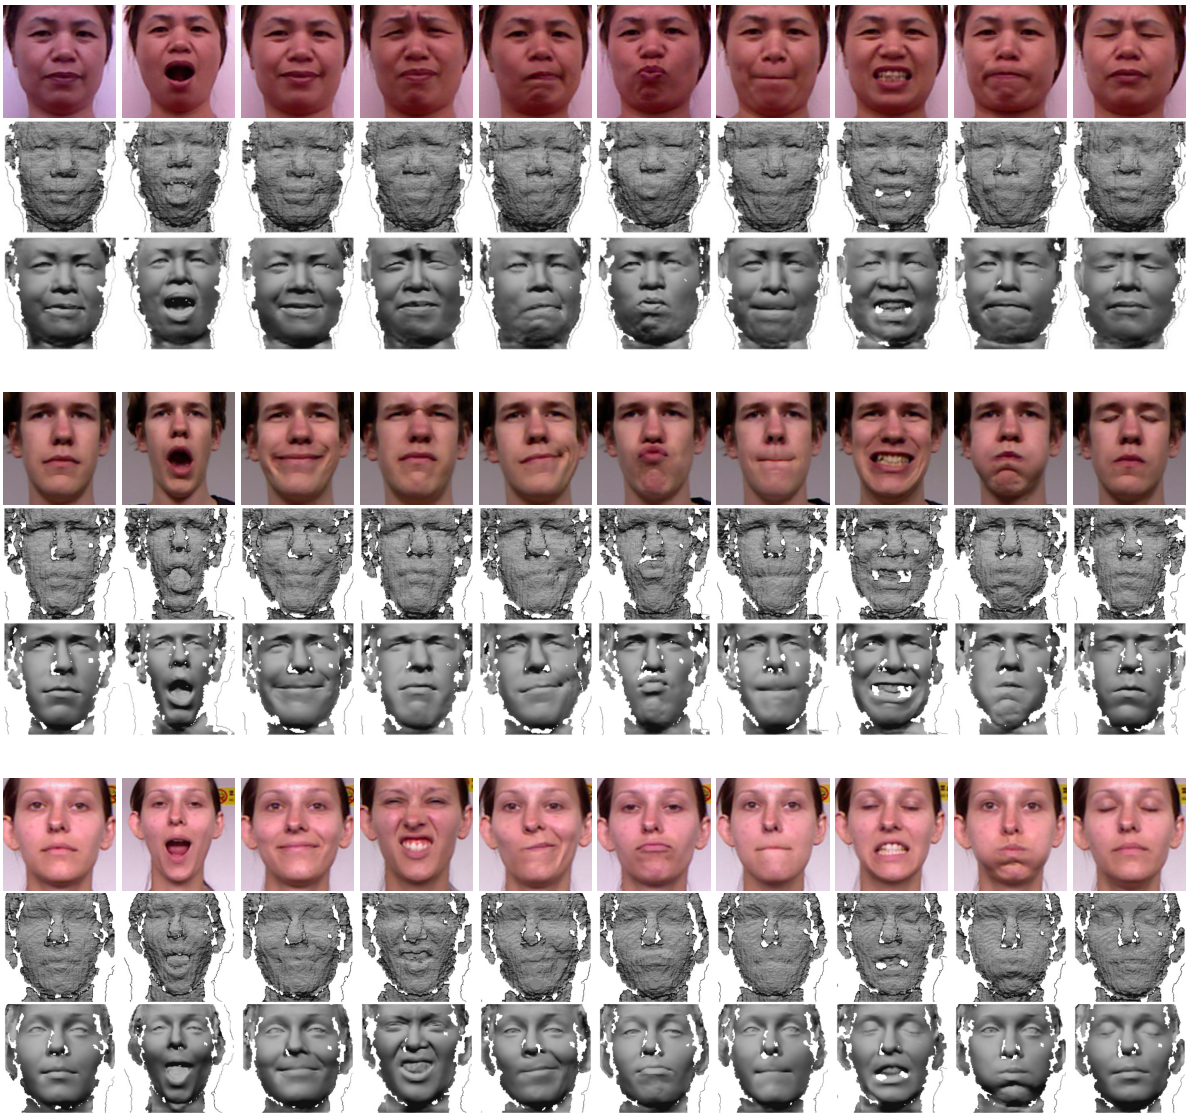

Figure 1: Raw Kinect depth enhancement using our normals on the Facewarehouse dataset [1].

Supplementary qualitative comparisons: Figs. 2 and 3 show additional predictions from our model in comparison to competing methods on the 300-W dataset [3] in both the normal and geometry domains.

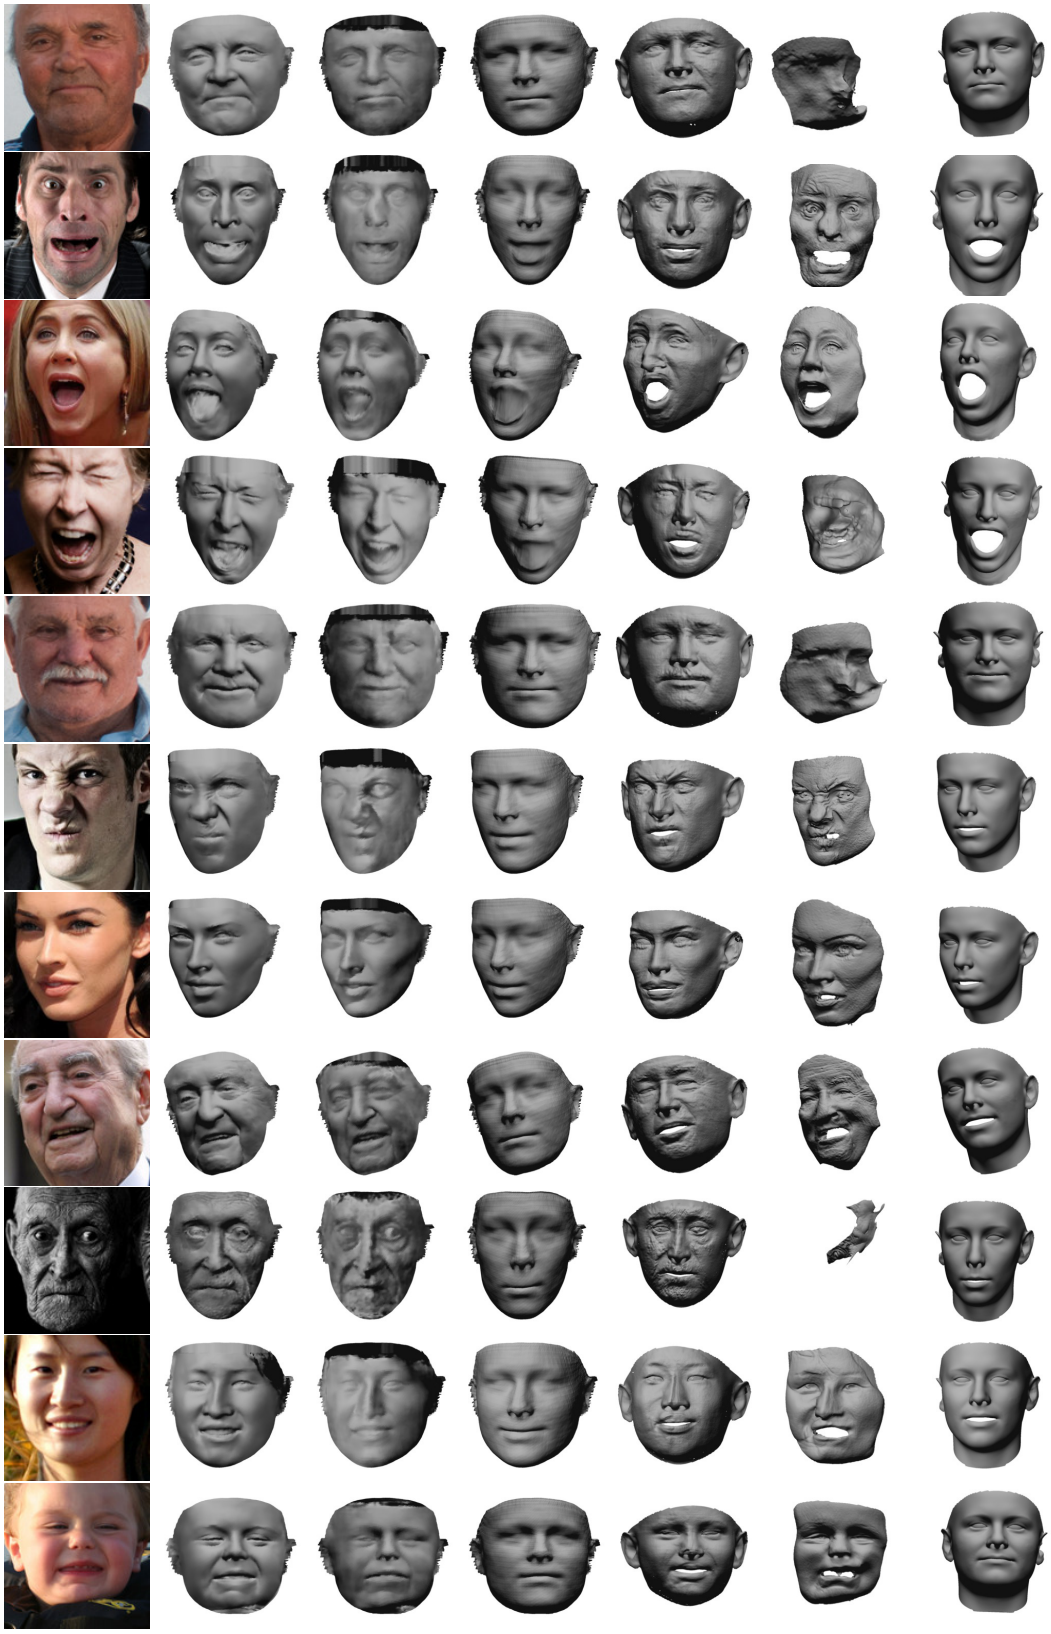
Input
Ours+PRN
SfSNet+PRN
PRN
Extreme
Pix2V
3DDFA

Figure 2: Qualitative comparisons with geometries in the 300-W dataset [3].


Input
Ours+PRN
SfSNet+PRN
PRN
Extreme
Pix2V
3DDFA

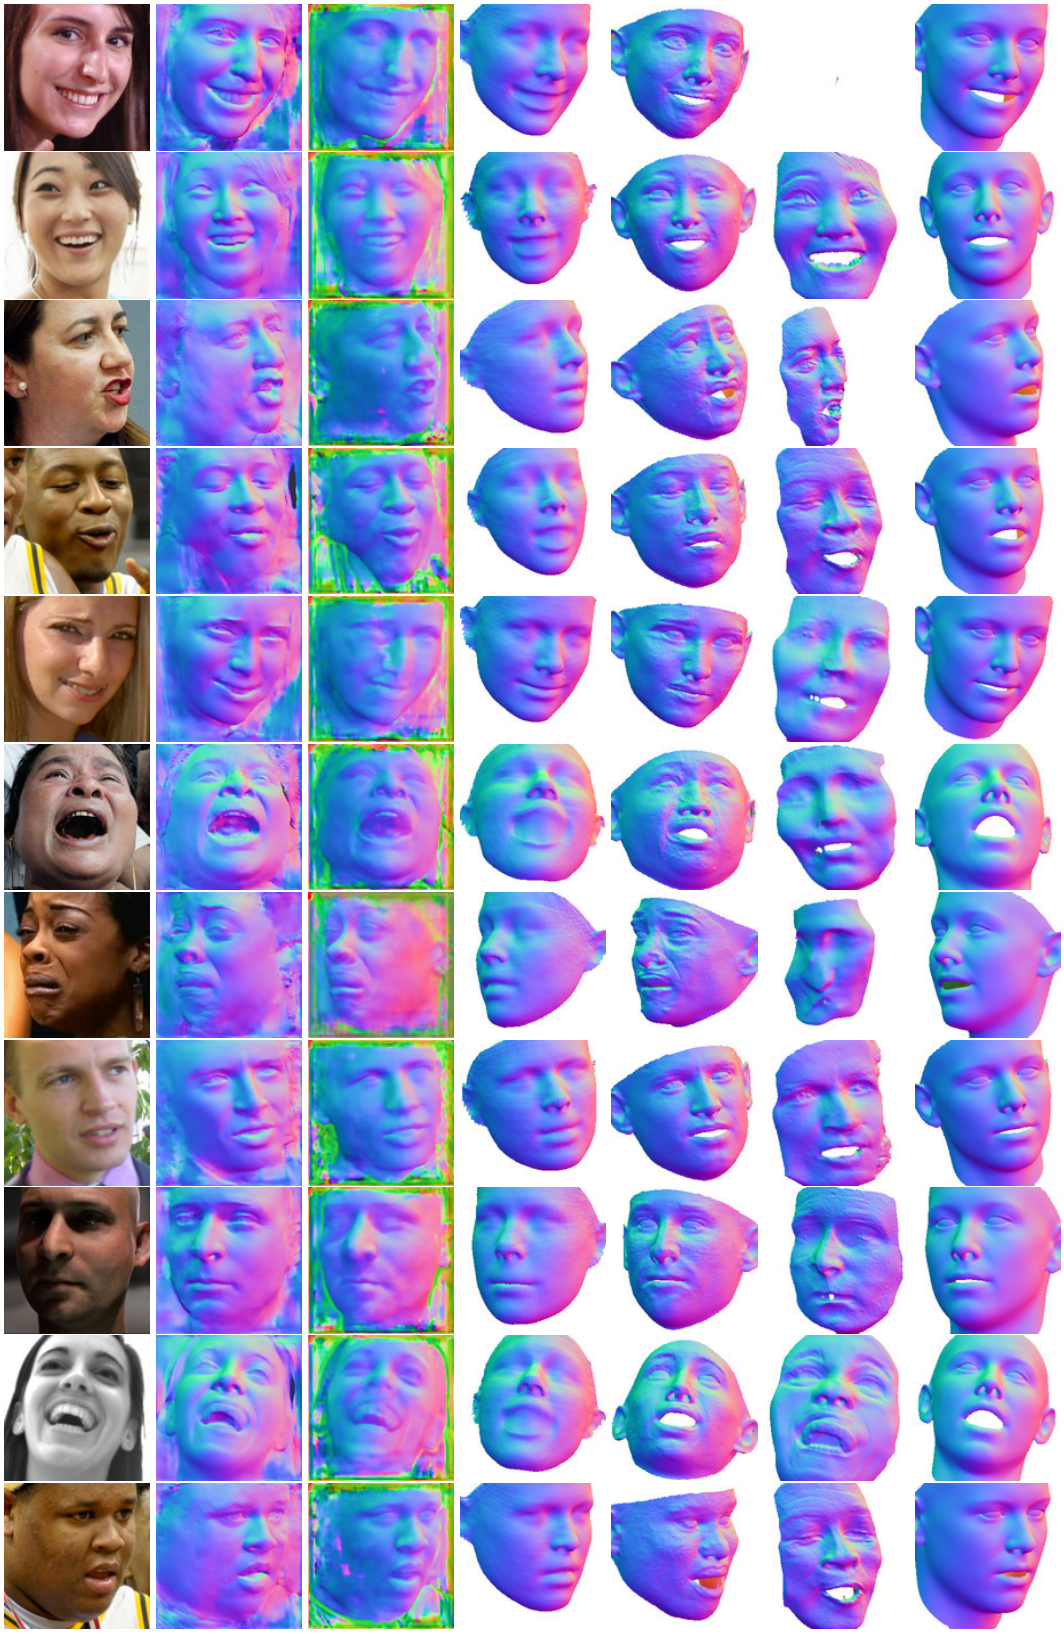

Figure 3: Qualitative comparisons with normals in the 300-W dataset [3].

References

[1] Chen Cao, Yanlin Weng, Shun Zhou, Yiying Tong, and Kun Zhou. Facewarehouse: A 3d facial expression database for visual computing. *IEEE Transactions on Visualization and Computer Graphics*, 2013. 1

[2] Jonathan Cohen, Marc Olano, and Dinesh Manocha. Appearance-preserving simplification. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques*, 1998. 1

[3] Christos Sagonas, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 faces in-the-wild challenge: The first facial landmark localization challenge. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2013. 1, 3, 5