EUROGRAPHICS 2018 / D. Gutierrez and A. Sheffer
(Guest Editors)

Volume 37 (2018), Number 2

From Faces to Outdoor Light Probes
Supplementary Material

Dan A. Calian[1,5], Jean-François Lalonde[2], Paulo Gotardo[5], Tomas Simon[3], Iain Matthews[3] and Kenny Mitchell[4,5]

[1]University College London [2]Université Laval, Québec City [3]Carnegie Mellon University [4]Edinburgh Napier University [5]Disney Research

1. Implementation Details: Optimization

We have implemented both our autoencoder and light estimation algorithms in Python using the numerical computation library Theano (`http://deeplearning.net/software/theano`), which provides automatic symbolic differentiation and transparent use of the GPU amongst its main features. Our autoencoder is also based on Lasagne (`http://lasagne.readthedocs.io`), a lightweight library to build and train neural networks in Theano.

At test time, given an input face image, the initial step is face detection and precomputation of 3D geometry (including normals), required to build the transport matrix $\mathbf{T}$; this step was implemented in MATLAB and C++. In the second step, we initialize the unknowns in our albedo and light probe model as follows: initial albedo ρ is spatially constant and equal to the mean of the best matching mode of the albedo GMM (via exhaustive search); the ground albedo ρ_{gnd} is initialized to a small constant ($\rho_{gnd} = [.2\ .2\ .2]$ in our tests). We then initialize the sun azimuth ϕ_{sun} from the optimal linear direction obtained via an initial fit of a low-dimensional spherical harmonics probe [Gre03]. Finally, the latent vector $\mathbf{z}$ is initialized to the mean over all probes in the training set.

Our optimization procedure then iteratively improves on the current solution in a coarse-to-fine strategy, based on image pyramids for both face albedo and light probe. In each pyramid level, it computes quasi-Newton, L-BFGS updates as to minimize our objective function using the method in [SvdBFM09], with automatic differentiation. At convergence, the current solution is upsampled and optimization resumes at the next (finer) pyramid level. Empirically, we have found that our current initialization method and 4 pyramid levels provided quite robust recovery of spherical environment maps at 64×128 resolution (in latitude-longitude format), as demonstrated in our experimental results.

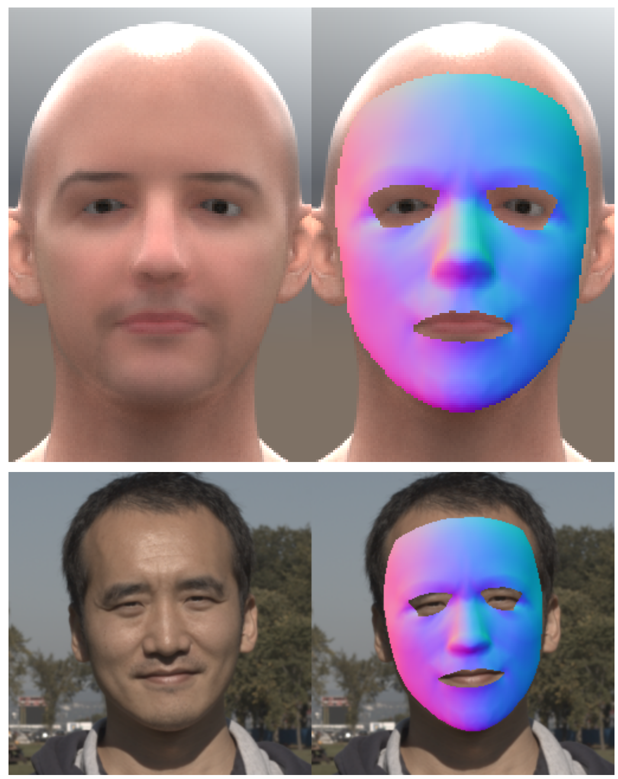

Figure 1: *Overlay of detected face mesh on synthetic (top) and real (bottom) images in our database showing the face region whose pixels are used for inverse lighting. Note that even though we did not synthesize hair for our synthetic images, the bald head is not used for inverse lighting, as in the real application scenario.*

2. Synthetic Face Database

Figure 2 shows an overview of our synthetic images rendered using the popular, physically-based Cycles render engine (`www.blender.org`). These were rendered with different combinations of 3D faces and illumination conditions.

Note that even though we did not synthesize hair for our synthetic images, the bald head is not used for inverse lighting, as in the real application scenario. Figure 1 show overlays of detected face mesh on synthetic and real images from our database, indicating the face region whose pixels are used for inverse lighting.

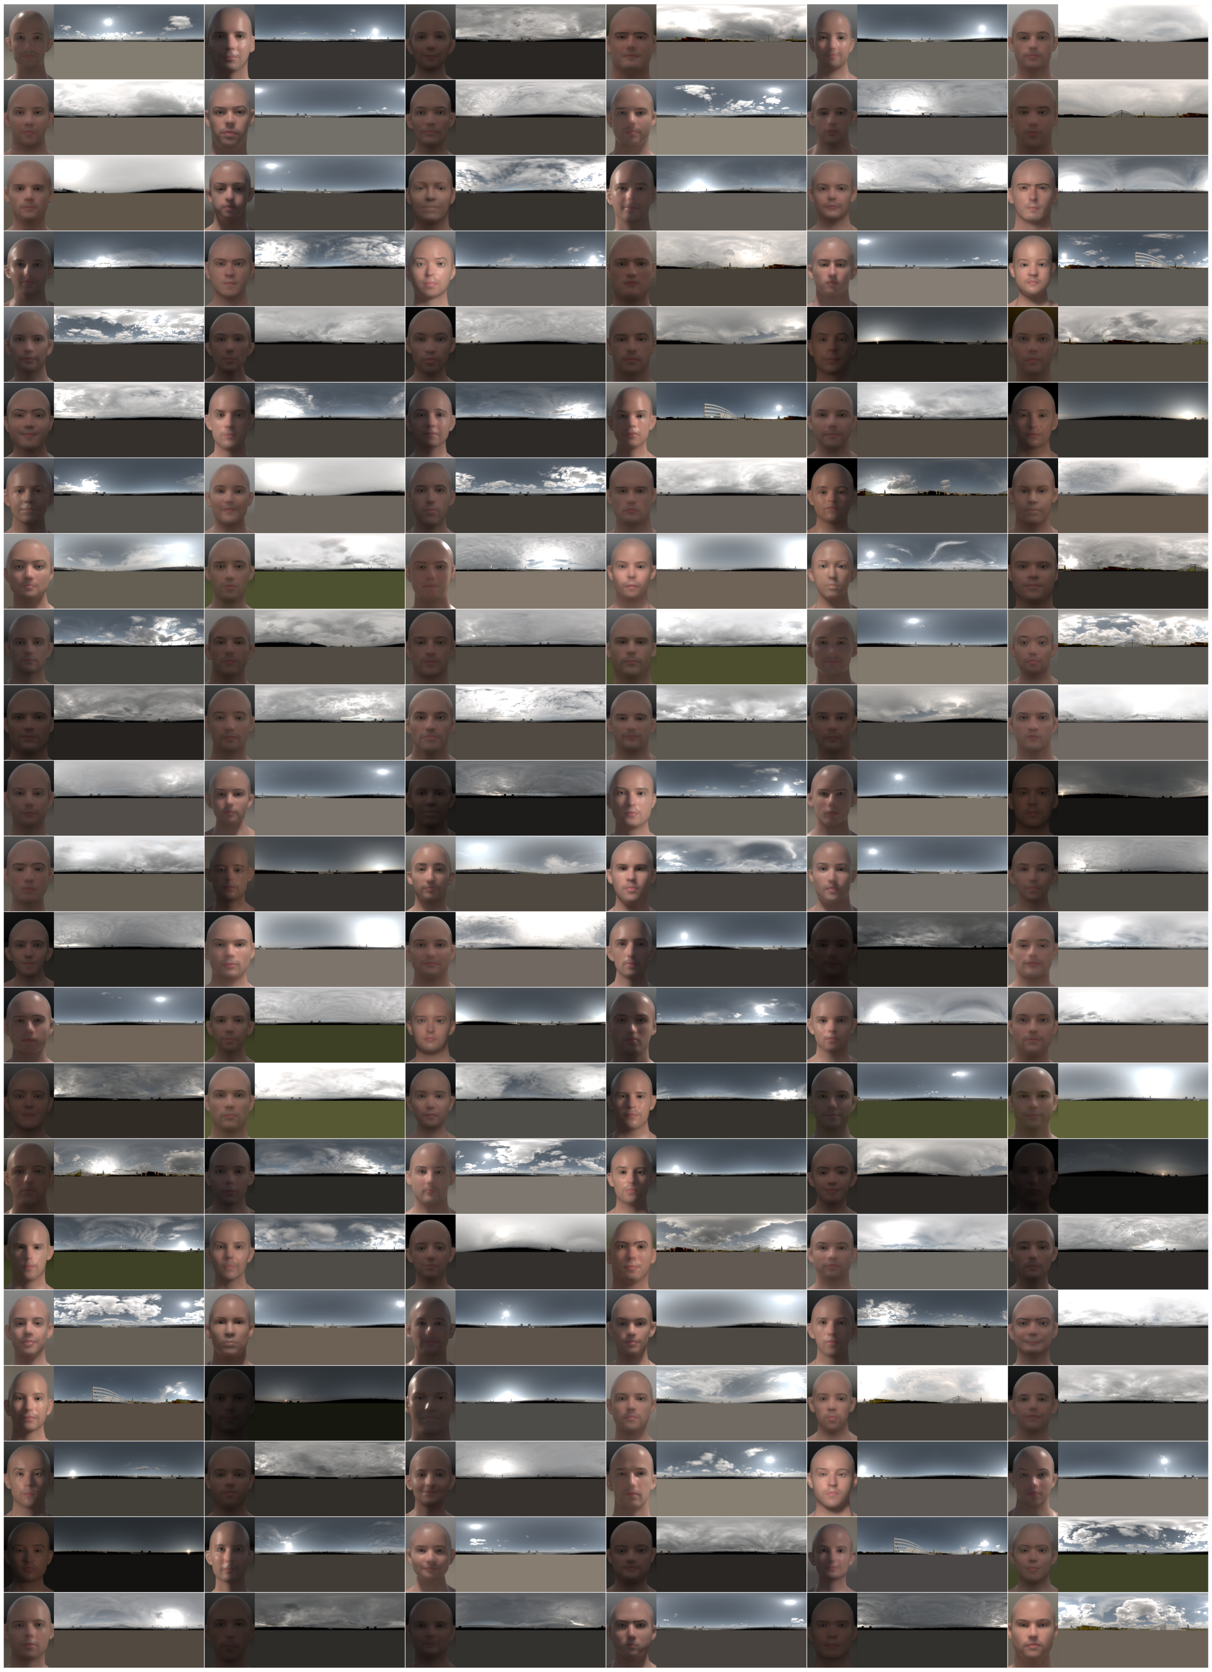

Figure 2: *Overview of our synthetic face database. Each randomly generated face model was rendered with the outdoor light probe on its right. These images have been tonemapped for display ($\gamma = 2.2$).*

Figure 3: *Overview of our real database: captured face images and corresponding outdoor light probes (better seen on the electronic version of this manuscript).*

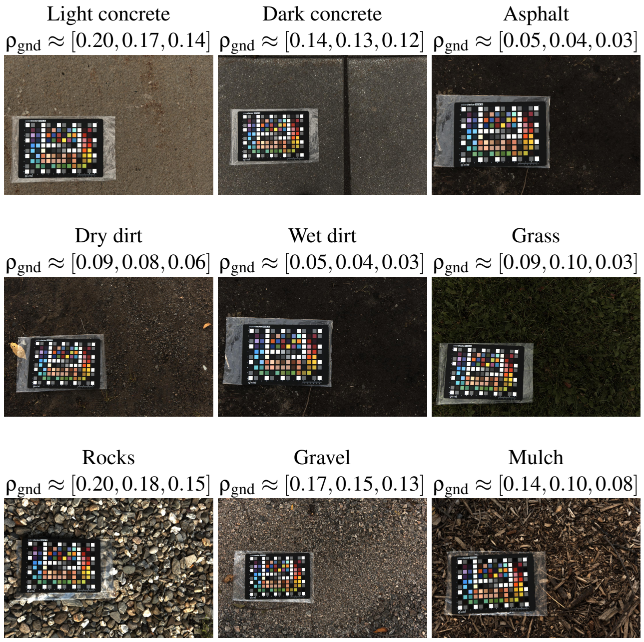

Figure 4: *Calibrated outdoor ground albedos. We use the Xrite™ color chart to calibrate RAW photographs of various ground textures, and extract a mean albedo by computing the mean pixel value over a large region of the image. These images have been rescaled and tonemapped for display.*

Since the sky probe database captures only the sky hemisphere, we synthesized an infinite Lambertian ground plane in the bottom hemisphere. To obtain realistic and calibrated ground albedos, we captured photographs of "typical" outdoor ground scenes alongside an Xrite™ color chart. The ground photographs were shot in RAW mode to extract linear data, and lens vignetting was corrected by capturing flat field images. Finally, calibrated ground albedos are obtained by computing the mean pixel color over a large region of the ground visible in the photograph. We include examples from our calibrated ground albedo database in figure 4.

3. Real Face Database

Figure 3 shows an overview of our real face database (better seen on electronic version of manuscript). In total, 9 subjects were recruited for this task and were asked to be photographed with a mostly neutral expression. They were photographed under 25 different lighting conditions. There were 8 male and 1 female subjects, most with fair skin. Subjects had varying amounts of facial hair, ranging from none to full beards.

4. Sensitivity experiments

We present in figures 5, 6 and 7 results obtained on all error metrics for the sensitivity analysis presented in section 5.5 and figure 8 of the main paper.

References

[Gre03] Green R.: Spherical harmonic lighting: The gritty details. In *Archives of the Game Developers Conference* (2003), vol. 5. 1

[SvdBFM09] Schmidt M., van den Berg E., Friedlander M., Murphy K.: Optimizing costly functions with simple constraints: A limited-memory projected quasi-newton algorithm. In *Proc. AISTATS* (2009). 1

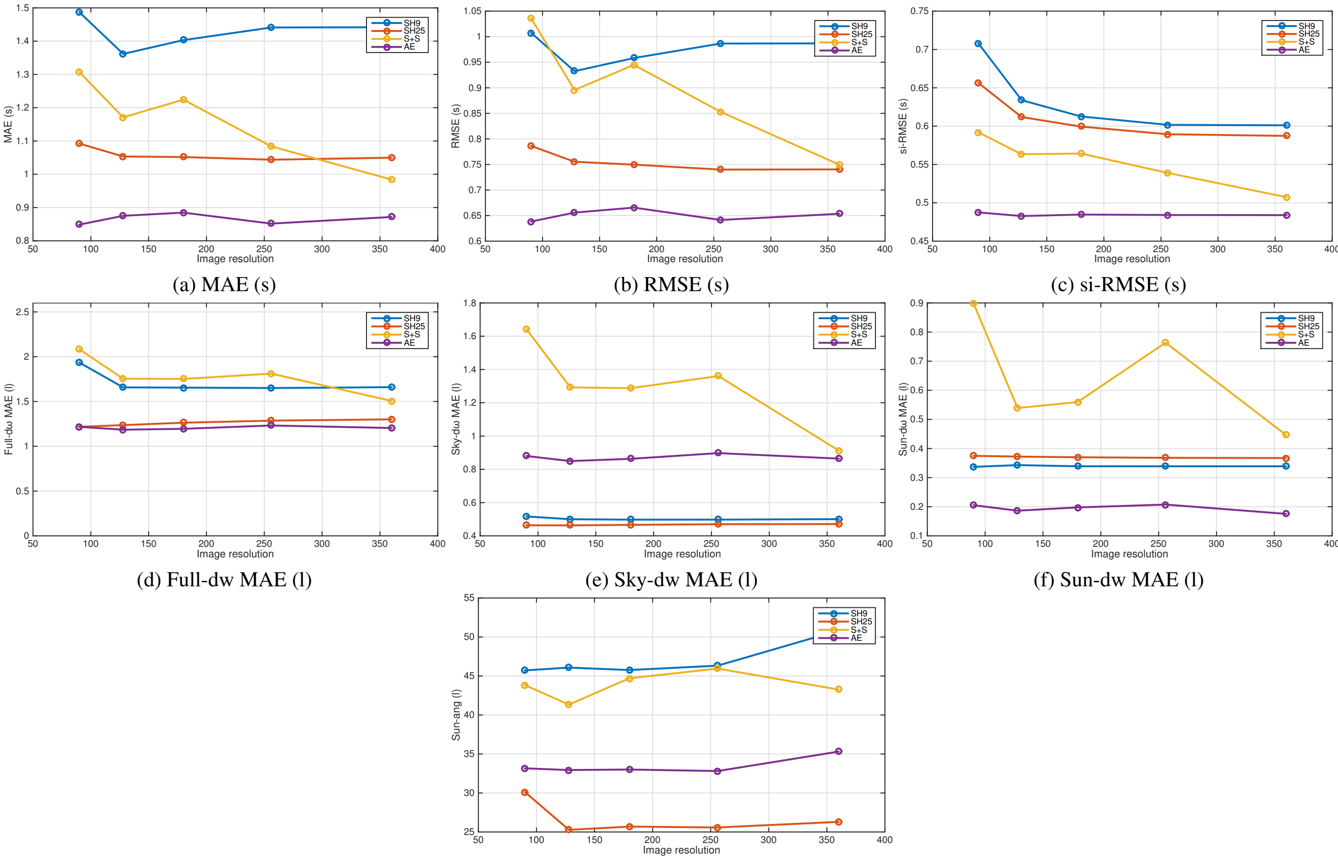

(a) MAE (s) (b) RMSE (s) (c) si-RMSE (s)

(d) Full-dw MAE (l) (e) Sky-dw MAE (l) (f) Sun-dw MAE (l)

(g) Sun-ang (l)

Figure 5: *Mean sensitivity of the methods as a function of the image resolution ranging from 90×90 to 360×360 for all metrics to complement figure 8 of the main paper.*

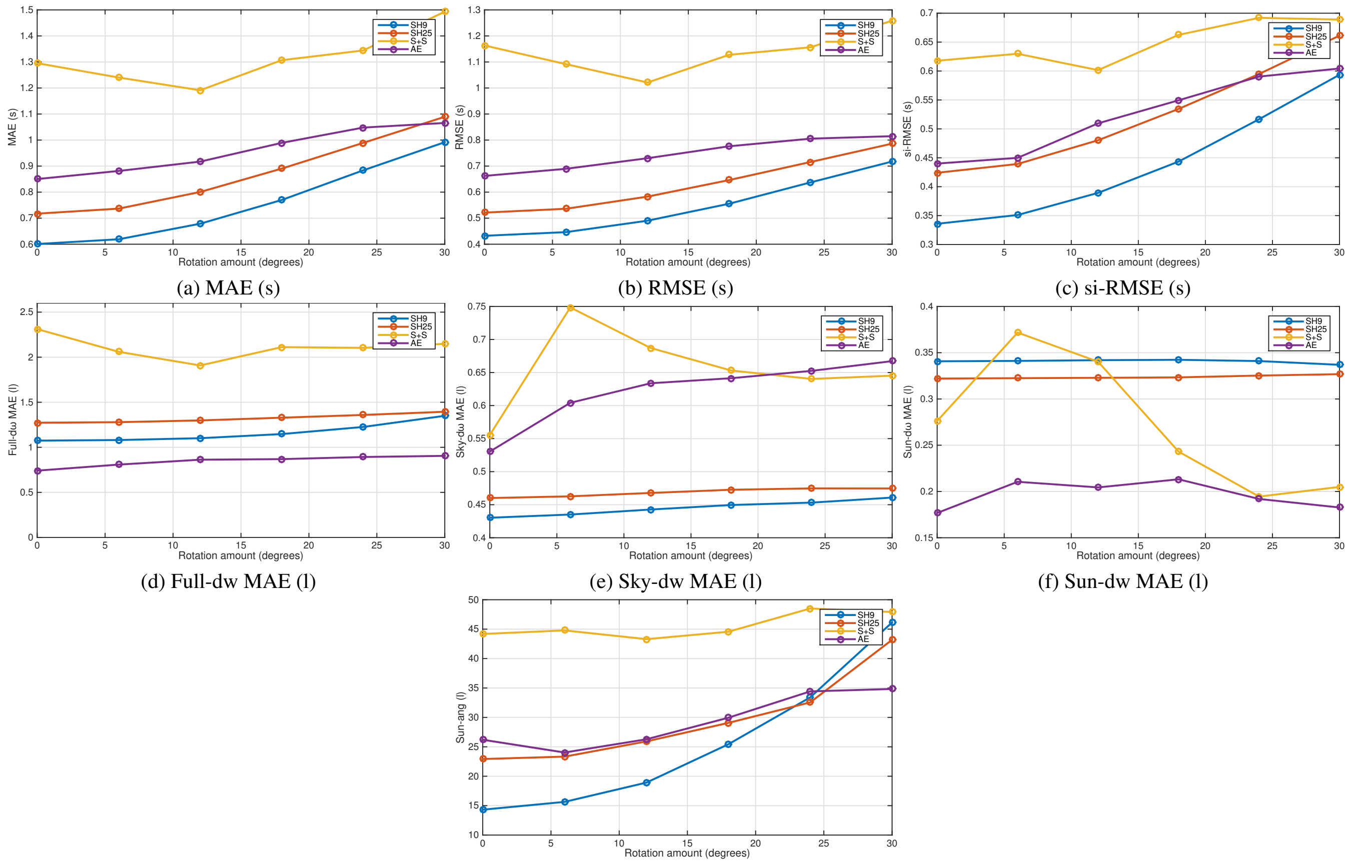

(a) MAE (s) (b) RMSE (s) (c) si-RMSE (s)

(d) Full-dw MAE (l) (e) Sky-dw MAE (l) (f) Sun-dw MAE (l)

(g) Sun-ang (l)

Figure 6: *Mean sensitivity of the methods as a function of the out-of-plane rotation from 0° to 30° for all metrics to complement figure 8 of the main paper.*

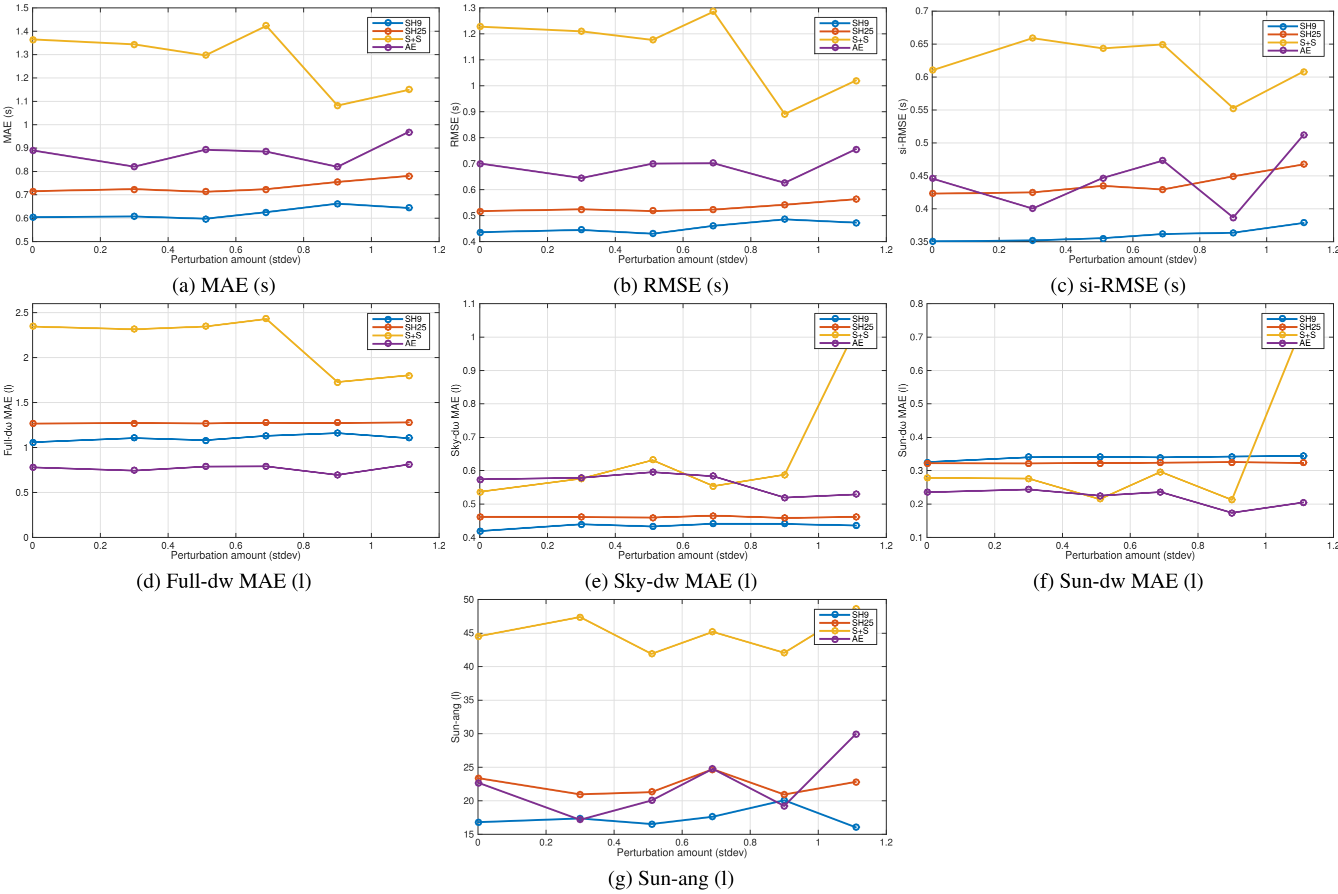

(a) MAE (s) (b) RMSE (s) (c) si-RMSE (s)

(d) Full-dw MAE (l) (e) Sky-dw MAE (l) (f) Sun-dw MAE (l)

(g) Sun-ang (l)

Figure 7: *Mean sensitivity of the methods as a function of the random Gaussian noise of standard deviation ranging from 0 to 1.1 applied to the blendshape coefficients for all metrics to complement figure 8 of the main paper.*