Improved Reconstruction of Deforming Surfaces by Cancelling Ambient Occlusion

Supplemental Material

Thabo Beeler[1,2], Derek Bradley[1], Henning Zimmer[2], and Markus Gross[1,2]

[1]Disney Research Zurich – {thabo.beeler,derek.bradley}@disneyresearch.com
[2] ETH Zurich – {hzimmer,grossm}@inf.ethz.ch

1 Table of Contents

This supplemental document contains the following information:

- **Section 2** shows visualized flow fields similar to Figure 6 in the paper, extended to include all flow algorithms and multiple different temporal spacings, as well as more detailed analysis of the flow Endpoint Error's through percentile plots.
- **Section 3** lists the optical flow parameter values used in our experiments.
- **Section 4** shows an experiment that illustrates the convergence and robustness of our shape optimization approach by starting from a plane and reconstructing forehead wrinkles.

2 Extended Evaluation of Results

Here we show extended results of our optical flow evaluation. The ground truth dataset is shown in Figure 1, including the original images and the result after cancelling ambient occlusion. Three different temporal spacings are used for flow computation, as shown in the figure.

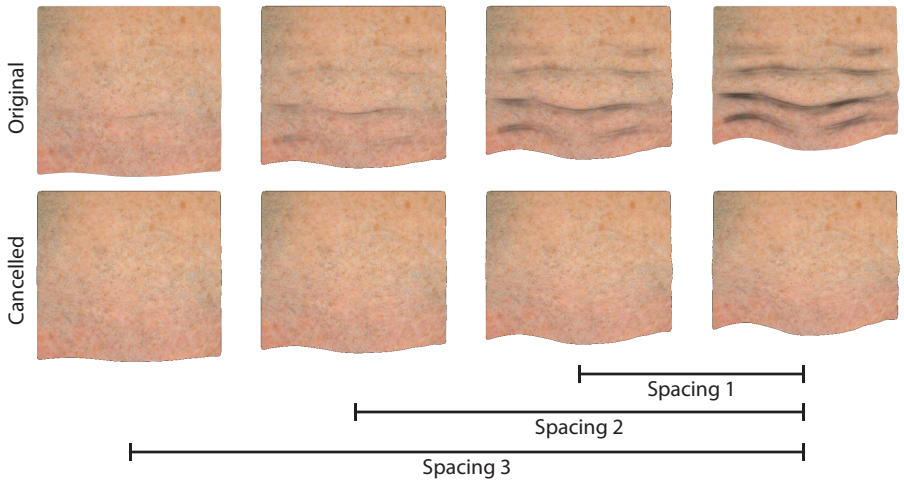

Fig. 1. The ground truth data for flow evaluation is simulated skin wrinkles. The top row shows the original images and the second row is the images after cancelling ambient occlusion. Three temporal spacings are used for the evaluation in this document.

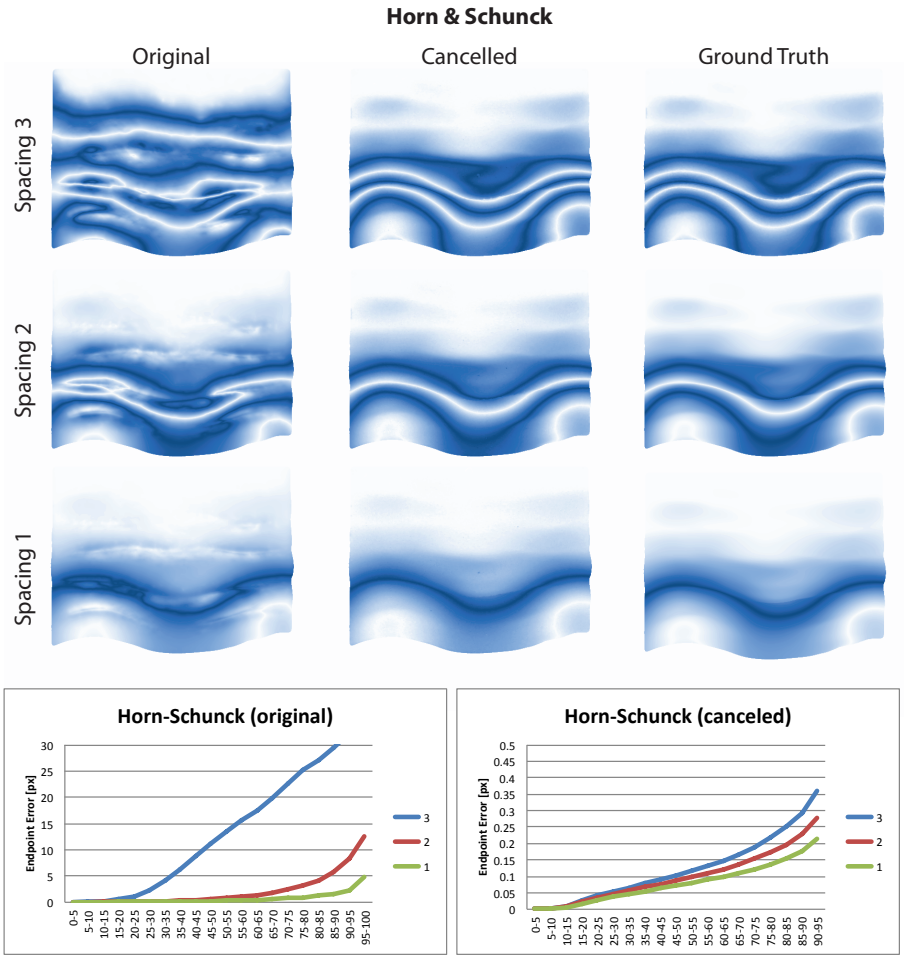

Fig. 2. This figure shows an extended evaluation of the improved flow for Horn and Schunck, before and after cancelling ambient occlusion. [TOP]: Flow visualizations for 3 temporally different spacings. Spacing 3 is the largest offset, and corresponds to Figure 6 in the paper. [BOTTOM]: Percentile plots for the Endpoint Error before and after cancellation for all 3 spacings. Notice the improvement in the upper percentiles, corresponding to the wrinkle region. Also, please note the large scale difference of the plots (60x).

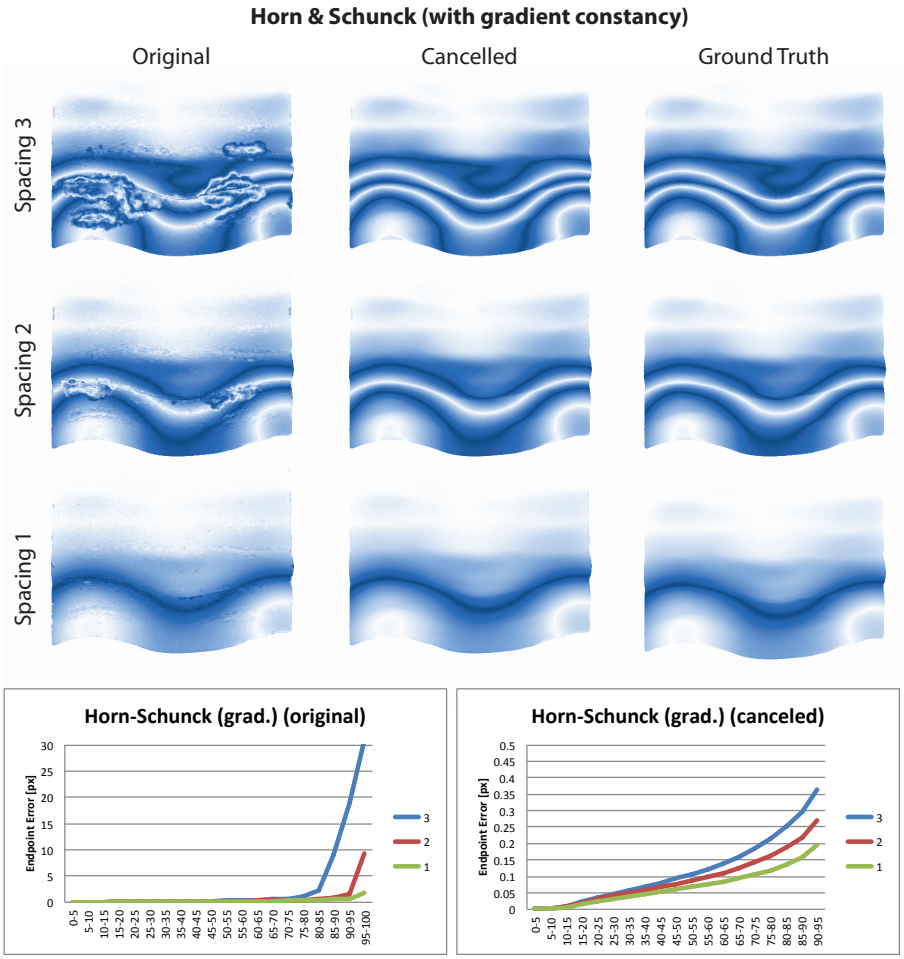

Fig. 3. This figure shows an extended evaluation of the improved flow for Horn and Schunck with added gradient constancey, before and after cancelling ambient occlusion. [TOP]: Flow visualizations for 3 temporally different spacings. Spacing 3 is the largest offset, and corresponds to Figure 6 in the paper. [BOTTOM]: Percentile plots for the Endpoint Error before and after cancellation for all 3 spacings. Notice the improvement in the upper percentiles, corresponding to the wrinkle region. Also, please note the large scale difference of the plots (60x).

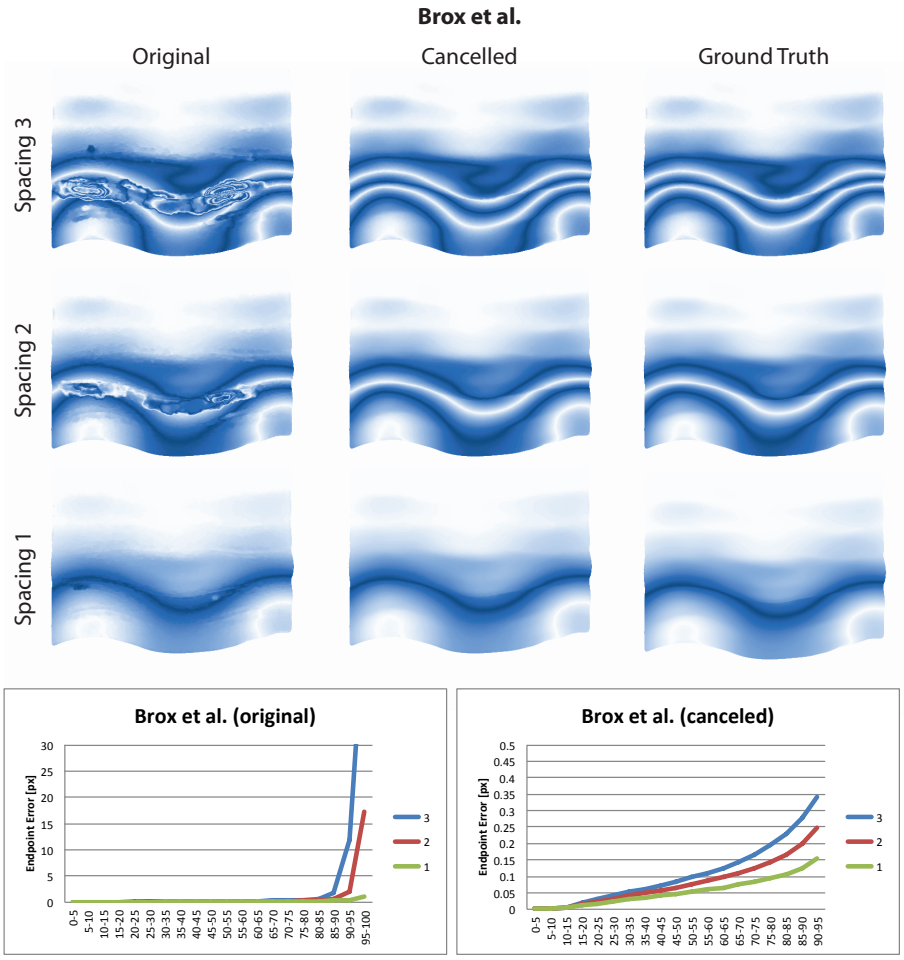

Fig. 4. This figure shows an extended evaluation of the improved flow for Brox et al., before and after cancelling ambient occlusion. [TOP]: Flow visualizations for 3 temporally different spacings. Spacing 3 is the largest offset, and corresponds to Figure 6 in the paper. [BOTTOM]: Percentile plots for the Endpoint Error before and after cancellation for all 3 spacings. Notice the improvement in the upper percentiles, corresponding to the wrinkle region. Also, please note the large scale difference of the plots (60x).

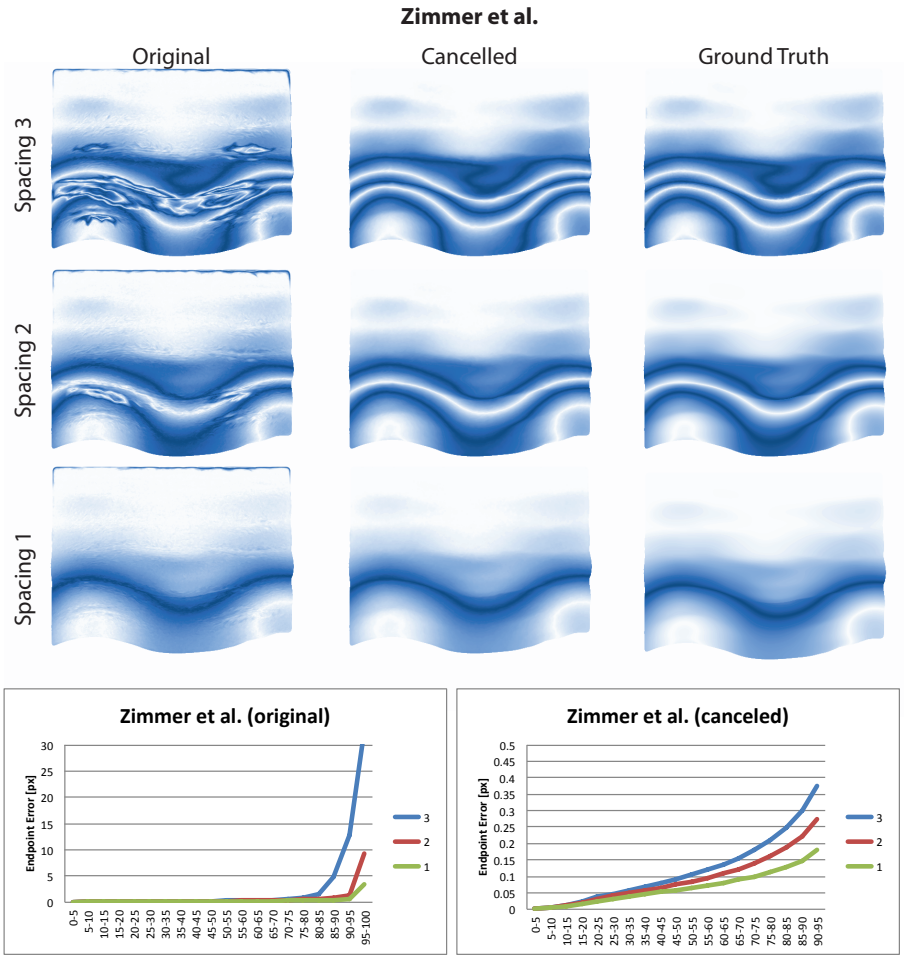

Fig. 5. This figure shows an extended evaluation of the improved flow for Zimmer et al., before and after cancelling ambient occlusion. [TOP]: Flow visualizations for 3 temporally different spacings. Spacing 3 is the largest offset, and corresponds to Figure 6 in the paper. [BOTTOM]: Percentile plots for the Endpoint Error before and after cancellation for all 3 spacings. Notice the improvement in the upper percentiles, corresponding to the wrinkle region. Also, please note the large scale difference of the plots (60x).

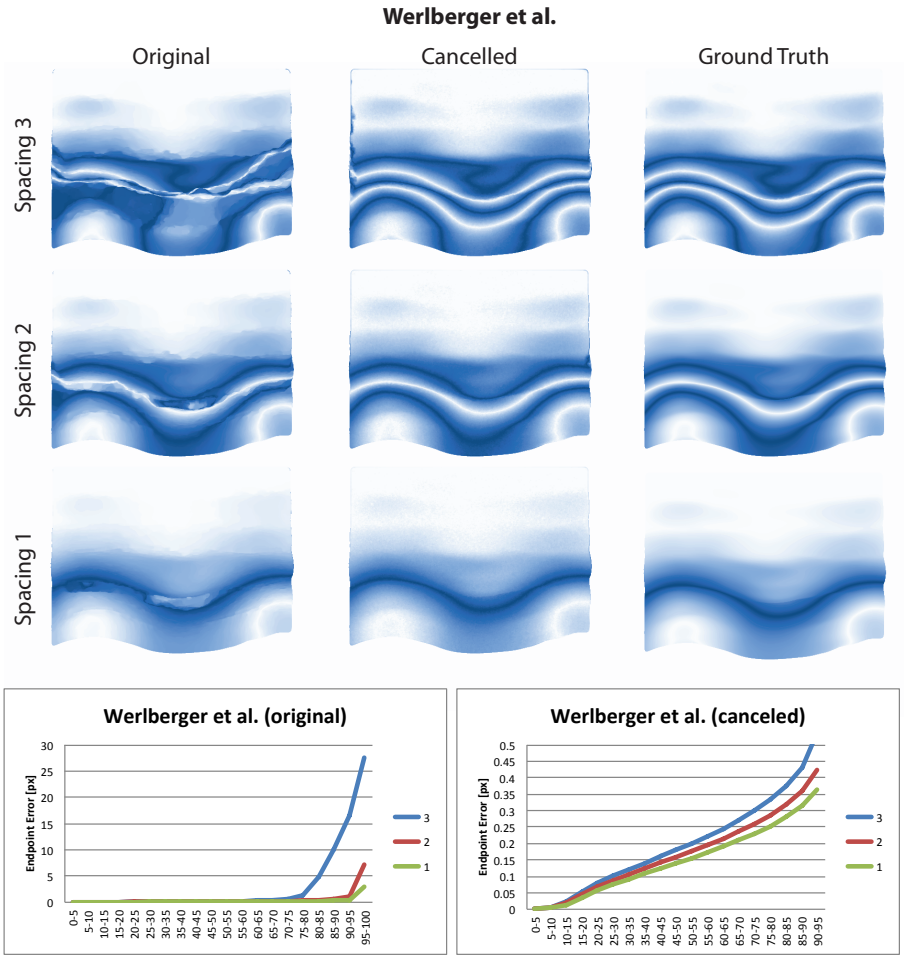

Fig. 6. This figure shows an extended evaluation of the improved flow for Werlberger et al., before and after cancelling ambient occlusion. [TOP]: Flow visualizations for 3 temporally different spacings. Spacing 3 is the largest offset, and corresponds to Figure 6 in the paper. [BOTTOM]: Percentile plots for the Endpoint Error before and after cancellation for all 3 spacings. Notice the improvement in the upper percentiles, corresponding to the wrinkle region. Also, please note the large scale difference of the plots (60x).

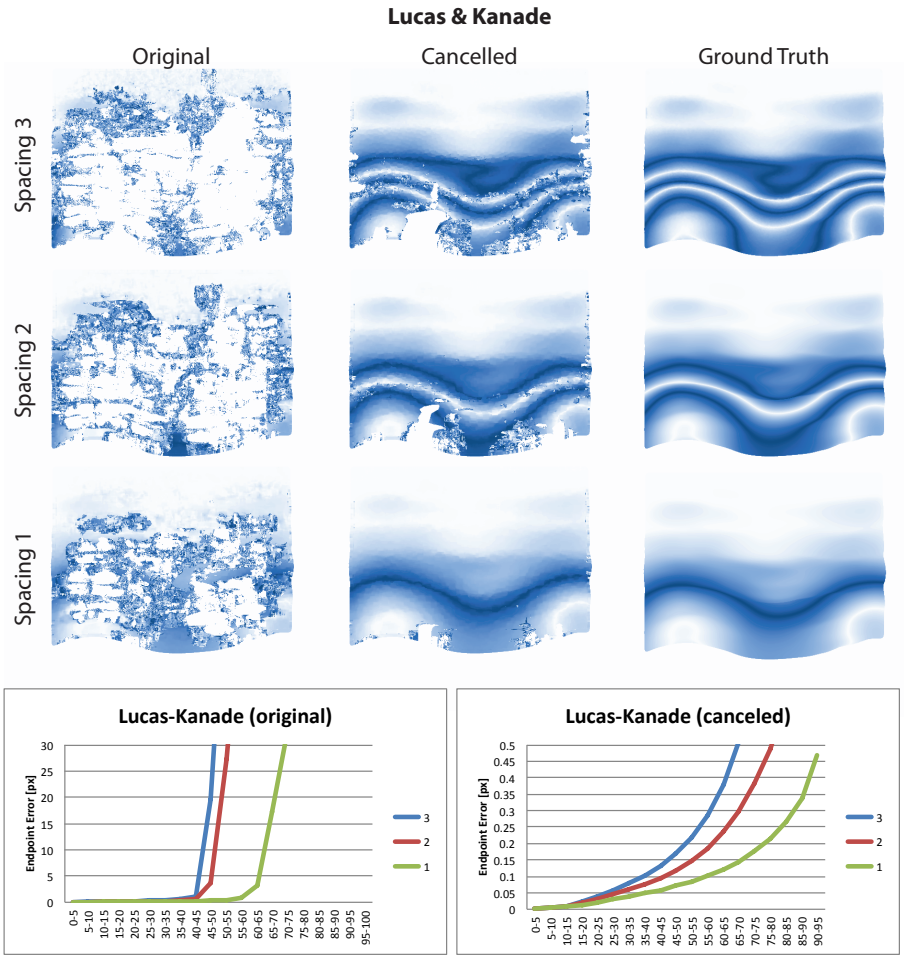

Fig. 7. This figure shows an extended evaluation of the improved flow for Lucas and Kanade, before and after cancelling ambient occlusion. [TOP]: Flow visualizations for 3 temporally different spacings. Spacing 3 is the largest offset, and corresponds to Figure 6 in the paper. [BOTTOM]: Percentile plots for the Endpoint Error before and after cancellation for all 3 spacings. Notice the improvement in the upper percentiles, corresponding to the wrinkle region. Also, please note the large scale difference of the plots (60x).

3 Flow Algorithm Parameters

As mentioned in Section 5.1 we chose different parameters for original and refined sequences (where the ambient occlusion has been cancelled) to achieve the best results in both scenarios. The parameters we used in our experiments are listed below in Table 1.

Algorithm	Parameters	Original	Refined
Horn-Schunck	α	1500	10
Horn-Schunck (grad.)	α/γ	1500/200	150/5
Brox et al.	α/γ	20/5	20/0
Zimmer et al.	α/γ	75/20	75/5
Werlberger et al.	$\alpha/\beta/\lambda/\epsilon$	5/1.0/30/0.001	5/0.5/50/1.500
Lucas-Kanade	window/iterations	5/100	5/100

Table 1. Parameters of the benchmarked algorithms.

Here, the parameters have the following meanings:

- α is in general a smoothness weight that balances the contribution of the data term and the smoothness term (regularizer, prior) in the energy function to be minimized. A larger value of α results in a stronger regularization and thus yields a smoother flow field. In the method of Werlberger et al., there is no smoothness weight, but a weight λ is multiplied with the data term. Thus, the corresponding smoothness weight is given by $\alpha = 1/\lambda$. However, the parameter α is also used in the paper of Werlberger et al., but here it determines the shape of an image-driven weighting function that is used in the smoothness term.
- γ determines the weight of the gradient constancy assumption in comparison to the brightness constancy assumption in the data term.
- β is another parameter used to steer the shape of the image-driven weighting function in the method of Werlberger et al.
- ϵ steers the shape of the robust penaliser function in the smoothness term proposed by Werlberger et al.
- The *window size* parameter in the local Lucas-Kanade method determines the size of a window around each pixel in which the flow field is assumed to be constant. In this sense it serves a similar role as the smoothness weight in the other global methods.
- The *iterations* parameter for the Lucas-Kanade method determines the number of refinement steps performed. Here, the images are warped by the current flow and then a flow correction is computed to iteratively refine the solution.

4 Shape Optimization Analysis

In this section we evaluate our shape optimization technique using the ground truth synthetic skin example (right-most column of Figure 1). In order to also test the robustness of our optimization, we show the shape convergence starting from an initial configuration of a simple plane. The result is shown in Figure 8. The ground truth shape is illustrated from two views in the left column. The following three columns show the shape optimization after 0, 30 and 100 iterations of shape optimization. The bottom row indicates the reconstruction error, measured as the absolute angular difference of the surface normals. In this visualization, the error ranges from 0 (blue) to 54 degrees (red). This curvature-based error was chosen over absolute positional error since the result can differ from the ground truth by an arbitrary translation.

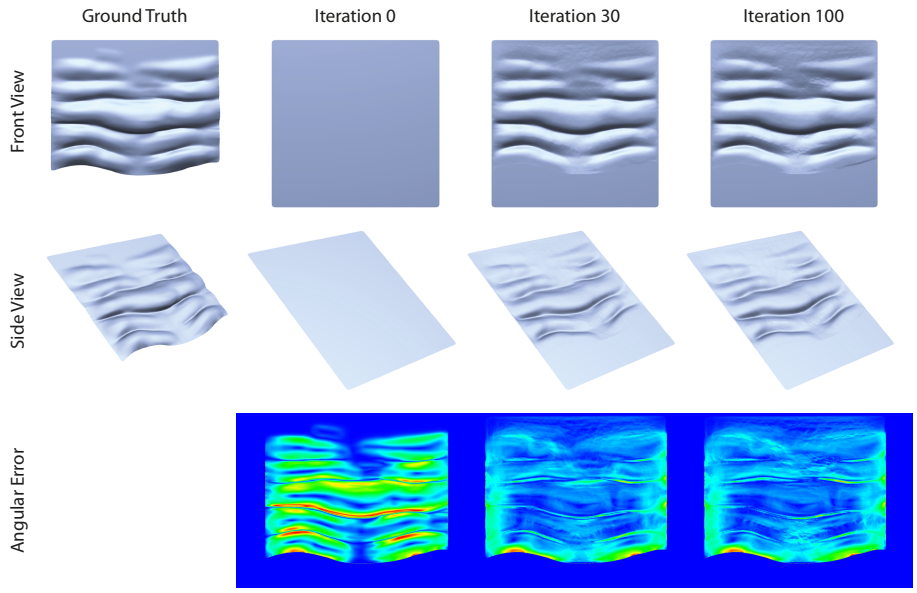

Fig. 8. Shape optimization analysis. Starting from a simple plane (Iteration 0), we converge to a wrinkle configuration (Iteration 100) corresponding to the ground truth (left column). Both front view and side view are shown. The angular error (bottom row) measures the quality of the reconstruction.

This experiment shows that our shape optimization procedure converges to a very similar shape as the ground truth, even in this pathological worst-case scenario of starting from a flat skin patch. We note that the error remains higher around the boundaries of the example patch, since we assume each vertex has a complete 360-degree neighborhood which will contribute to its ambient occlusion.