Supplementary Document

Fig. 2 illustrates our approach with an overview figure. Fig. 3 illustrates our optimization process for estimating reflectance and sun visibility at sparse reconstructed 3D points. Fig. 4 shows input and results of our image-guided propagation step to estimate the intrinsic decomposition at all pixels.

We compare our results with a decomposition computed directly using the reconstructed geometry in Fig. 5, and with the results of a user-assisted and two automatic approaches in Fig. 7.

We show more results of our rich intrinsic decomposition in Fig. 6. The supplementary video shows examples of advanced image manipulations made possible by our approach in image editing software.

Details on our approach are given in the attached research report.

References

BOUSSEAU, A., PARIS, S., AND DURAND, F. 2009. User-assisted intrinsic images. *ACM Trans. Graph. 28*, 5.

SHEN, L., TAN, P., AND LIN, S. 2008. Intrinsic image decomposition with non-local texture cues. In *CVPR*.

SHEN, J., YANG, X., JIA, Y., AND LI, X. 2011. Intrinsic images using optimization. In *CVPR*.

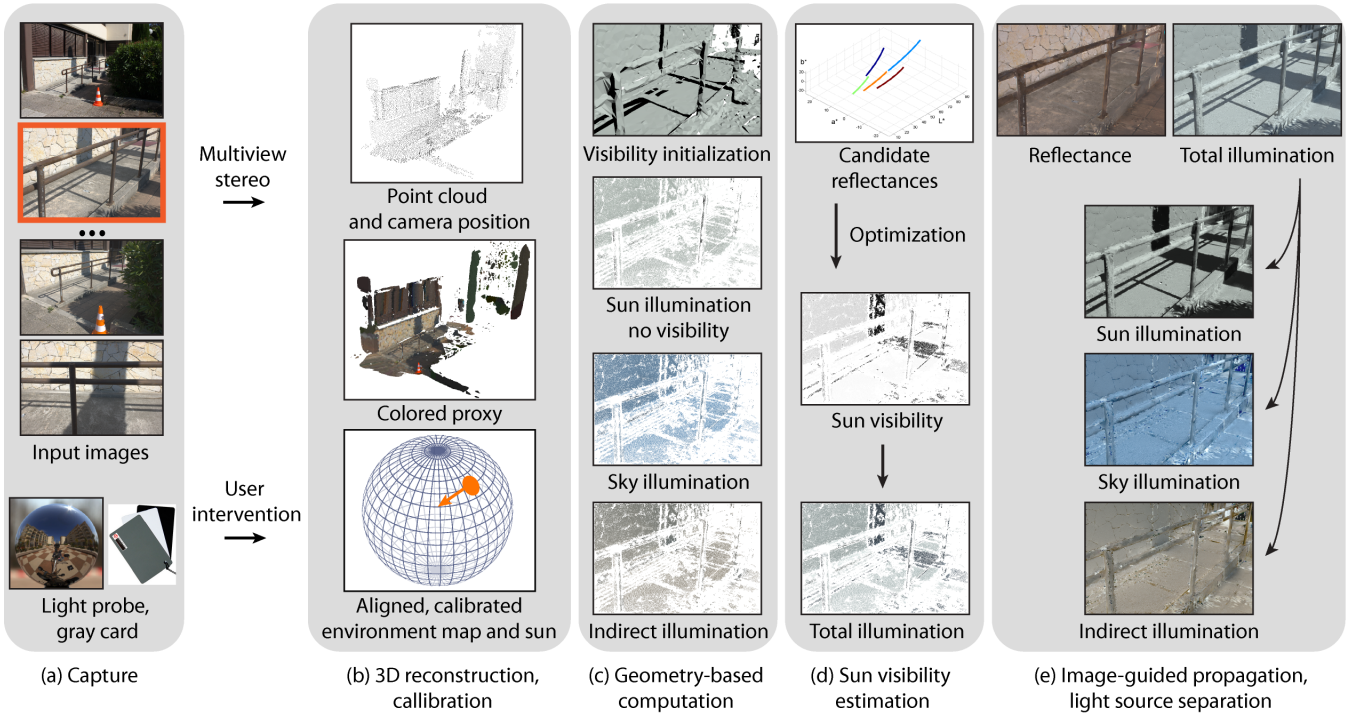

Figure 2: *Overview of our approach. Users capture a small set of pictures of the scene, along with an environment map and two pictures of a gray card in sun light and in shadow (a). We illustrate our intrinsic image decomposition with the picture highlighted in orange. We use multi-view stereo to reconstruct a point cloud of the scene and a coarse proxy geometry (b). Users align the environment map and the sun with this point cloud and use the gray card to calibrate their intensity. Once this simple manual calibration is performed, all the remaining steps are automatic. We use the reconstructed 3D geometry to compute sun, sky and indirect lighting over the point cloud (c). We also compute an initial guess of the sun visibility using the coarse proxy. These illumination values give us the necessary information to compute a set of candidate reflectances for each 3D point. The candidate reflectances form curves in color space parametrized by the sun visibility (d). We introduce an iterative optimization that identifies the reflectance of each 3D point from these candidates, along with a precise estimation of the sun visibility. The final step of our method consists in propagating the illumination values computed at 3D points to every pixel in the image (e). We decompose the propagated illumination into the sun, sky and indirect lighting components.*

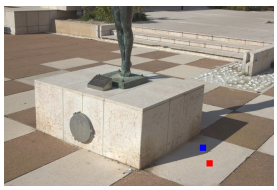

(a) Selected reconstructed points

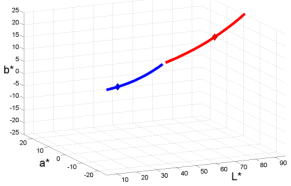

(b) Candidate reflectance curves

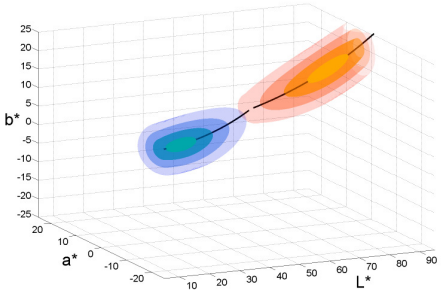

(c) Regions of influence at initialization

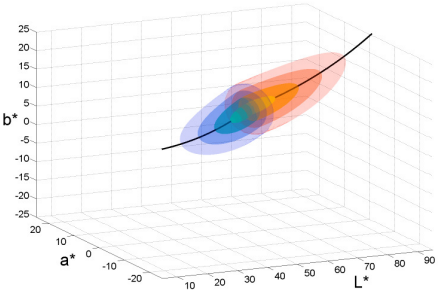

(d) Regions of influence after one iteration

Figure 3: *Illustration of our optimization process which estimates reflectance and sun visibility at a sparse set of points. Multiple reconstructed points sharing the same reflectance will generate intersecting candidate curves in color space. (a) We selected two reconstructed points with the same reflectance but different illuminations (red and blue squares). (b) The corresponding candidate reflectance curves (nearly) intersect at one end, which corresponds to the reflectance of both points. Diamond markers on the curves correspond to the initial guess for visibility (randomly set in this example). (c) Each curve affects a region of the surrounding color space. Regions of influence at initialization are illustrated as isosurfaces with varying opacity. Regions closer to the curve and the current visibility estimate are more affected. (d) After one iteration, the visibility estimates have moved towards the intersecting end of the curves, increasing the overlap between the regions of influence. Our algorithm has detected that both reconstructed points share a similar reflectance.*

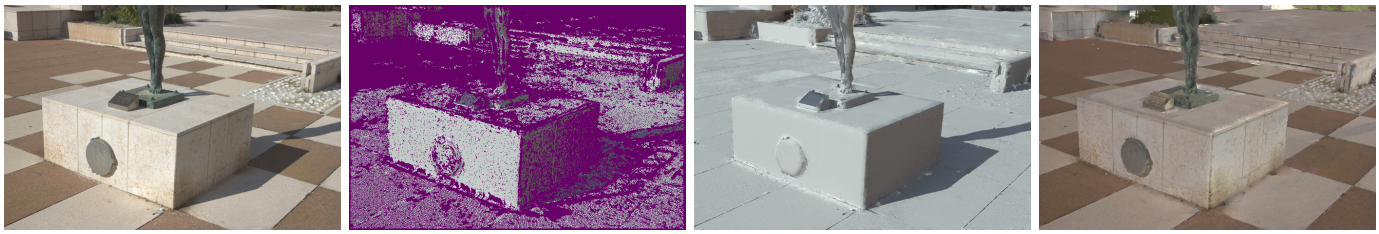

(a) Input photograph (b) Illumination constraints (c) Estimated total illumination (d) Estimated reflectance

Figure 4: *Illustration of the Image-guided propagation step: separation of reflectance and total illumination. Our optimization procedure enforces illumination constraints at reconstructed points (b) and propagates them to all pixels, in order to separate an input photograph (a) into total illumination (c) and reflectance (d).*

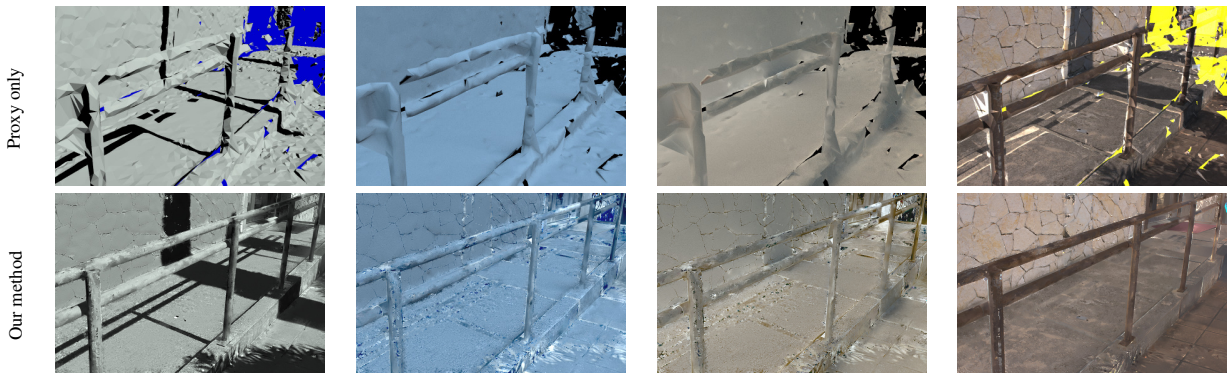

Sun Illumination Sky Illumination Indirect Illumination Reflectance

Figure 5: *Comparison between our results after lighting separation (bottom row) and the decomposition estimated from the geometric proxy (top row), where the illumination at each pixel is computed directly using the reconstructed geometry. The proxy reflectance is obtained by dividing the input image by the sum of the illumination components. Holes and inaccuracies in the proxy translate to artifacts and residual shadows in the reflectance. In constrast, our method uses geometric cues to estimate illumination only in regions that are well reconstructed, and uses image-guided propagation to propagate the decomposition to all pixels. This allows us to accurately detect shadow boundaries and obtain a shadow-free reflectance image.*

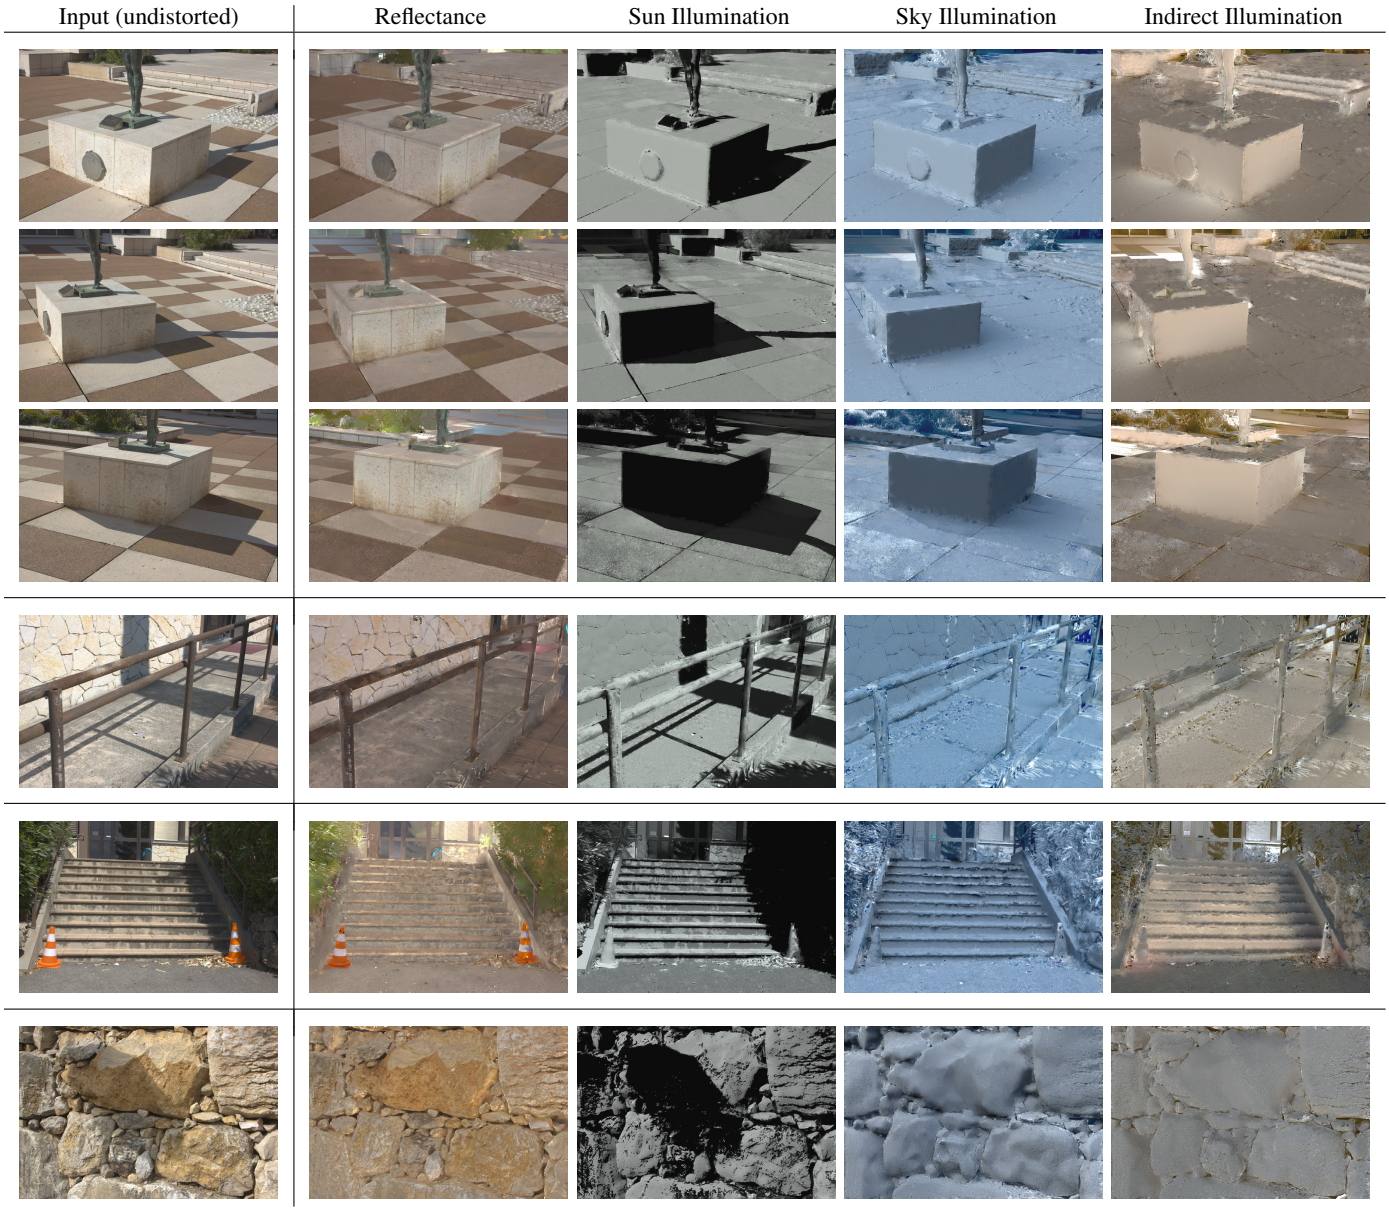

Figure 6: *Results of our decomposition after image guided propagation and lighting separation. We adjusted the brightness of each image for illustration purposes. For each scene, the sun illumination is usually much more intense than sky illumination on average, and the sky illumination is more intense than indirect illumination on average.*

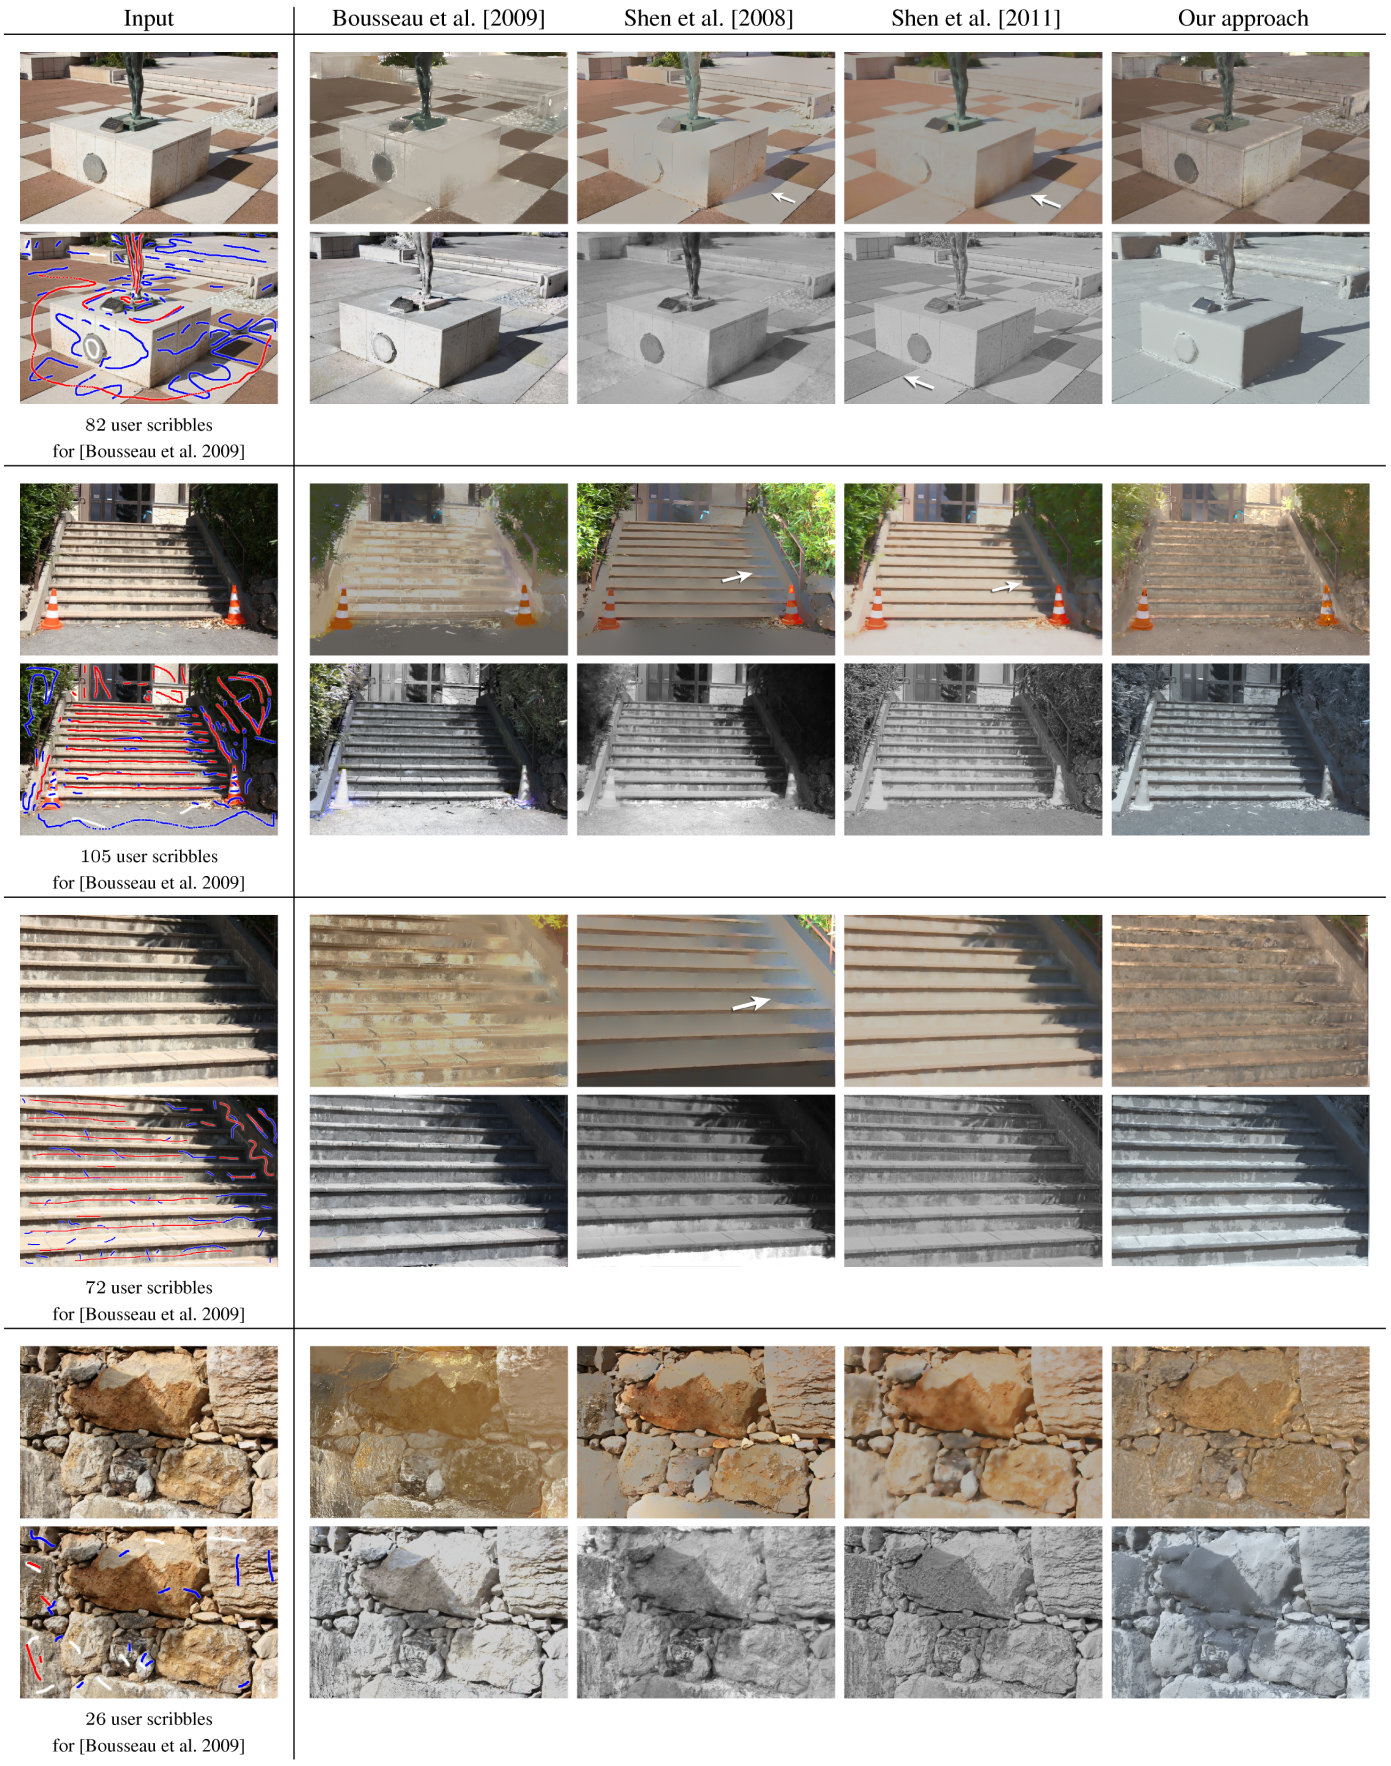

Figure 7: *Comparison of our method with a user-assisted approach [Bousseau et al. 2009] and two automatic algorithms [Shen et al. 2008; Shen et al. 2011]. In the first column, the user scribbles used for the method of Bousseau et al. [2009] are shown under the input image. In the next columns, the first row contains the estimated reflectance while the second one corresponds to the total illumination. Our results are shown in the last column. Our multiview approach outperforms single-image automatic algorithms and achieves results of comparable quality to the user-assisted approach with significantly less user intervention. In addition, our approach allows further decomposing the total illumination into sun, sky, and indirect illumination (not shown in this figure). White arrows point to residual reflectance or shading variations. The brightness has been adjusted for comparison.*