Multiview Reconstruction of Complex Organic Shapes

Jasenko Zivanov
jasenko.zivanov@unibas.ch
Thomas Vetter
thomas.vetter@unibas.ch

Department of Mathematics and Computer Science
University of Basel
Basel, Switzerland

We propose a novel narrow baseline multiview stereo surface reconstruction method that is specifically aimed at complex shapes of biological origin that show many thin protrusions and curved occluding contours. Our method is built around fitting local quadrics to pixel-precise occluding contours and it avoids any planarity assumptions that are common to other state of the art methods. We describe a complete pipeline that begins with a sequence of calibrated noisy images and produces a final watertight surface. Although quadrics have been used before to model geometry ([1, 2, 4]), our work is to our knowledge the first to use them as part of a complete stereo reconstruction pipeline. The pipeline is sketched out in figure 1.

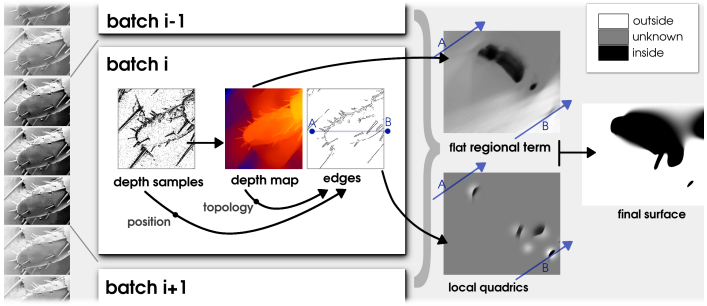

Figure 1: An overview over our pipeline. A detailed explanation is given in the text below.

Shapes of biological origin often feature curved contours that do not correspond to static curves in space - instead, those curves drift along the surface as the vantage point changes. For this reason, our method estimates the positions of occluding contours only from small subsets of the entire image sequence. By only looking at short batches of frames, we can determine the momentary positions of such contours, which later enable us to fit local quadrics to them. This separation into batches also makes our method very robust to non-Lambertian reflectiveness and to changes in illumination. Our method can even be applied to sequences where the illumination is fixed to the observer and not the scene.

For every batch, we compute a sparse depth map corresponding to the center image of the batch. The depth values are determined independently for every pixel, which allows us to reconstruct very small features. If we were to compare image windows consisting of multiple pixels, we would be making an implicit planarity assumption. In spite of considering only single pixels, our depth estimation process is still very robust to noise, since the optimal depths and the optimal colors of each pixel are determined simultaneously. This results in both a sparse depth map and a dense, denoised version of the image. An example of a denoised image is shown in figure 2.

We then apply a process similar to [3] to estimate a dense depth map that matches the sparse depth map and that conforms to the structure of the denoised image. This can be viewed as a process of reasoning that propagates information from the observed edges into the smooth areas of the image. This dense depth map is then used to detect the occluding contours of that batch while the original sparse depth map is used to estimate their locations and orientations.

Next, we use the contours detected in all the batches to estimate local quadrics in a discretized volume. Using quadrics as a surface representation offers the advantage that the smoothness of the quadric parameters does not translate into planarity of the surface. Instead, arbitrarily curved shapes can be represented by smooth quadric coefficients. At every voxel v, we determine a real, symmetrical 4×4 matrix C_v that holds the coefficients of the local quadric, where the quadric is defined as the zero level set of $f(x) = x^T C_v x$, where $x = (x_1, x_2, x_3, 1)^T$ is a homogeneous position vector. The coefficients of C_v are determined such that the quadric is tangential to the contours detected in the area of space surrounding v.

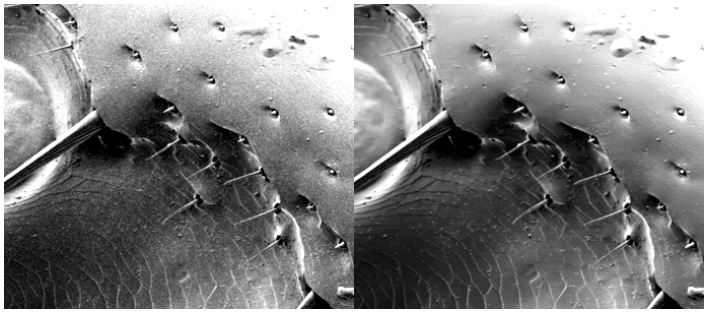

Figure 2: The denoising effect of our depth estimation process. The original image (left) and the denoised version (right). The corresponding batch contained 50 images. The contrast of the displayed images has been enhanced to emphasize the noise.

Finally, the estimated quadrics are combined with a more traditional regional term to obtain a final watertight surface using a volumetric optimization process. The final surface is required to fit the quadrics in areas close to detected contours and to be consistent with the dense depth maps otherwise.

We have tested our method on a sequence of 1400 images of a cat flea captured by a scanning electron microscope using secondary electron imaging. The results can be seen in figure 3.We show that very thin features can be reconstructed correctly by our method, even in the absence of strong contrast and even if the illumination is constant with respect to the viewer and not to the scene.

Figure 3: Reconstruction of a cat flea (ctenocephalides felis) from 1400 scanning electron microscope images. From left to right: one of the input images; enlarged area of the input image; a reconstruction of the enlarged area.

[1] Roberto Cipolla, Bjorn Stenger, Arasanathan Thayananthan, and Philip HS Torr. Hand tracking using a quadric surface model and bayesian filtering. In *Mathematics of Surfaces*, pages 129–141. Springer, 2003.

[2] Geoffrey Cross and Andrew Zisserman. Quadric reconstruction from dual-space geometry. In *Computer Vision, 1998. Sixth International Conference on*, pages 25–31. IEEE, 1998.

[3] Qi Shan, Brian Curless, Yasutaka Furukawa, Carlos Hernandez, and Steven M Seitz. Occluding contours for multi-view stereo. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 4002–4009. IEEE, 2014.

[4] Hui Xie, Jianning Wang, Jing Hua, Hong Qin, and Arie Kaufman. Piecewise c1 continuous surface reconstruction of noisy point clouds via local implicit quadric regression. In *Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*, page 13. IEEE Computer Society, 2003.