Fusing Structured Light Consistency and Helmholtz Normals for 3D Reconstruction

Michael Weinmann
mw@cs.uni-bonn.de
Roland Ruiters
ruiters@cs.uni-bonn.de
Aljosa Osep
osep@cs.uni-bonn.de
Christopher Schwartz
schwartz@cs.uni-bonn.de
Reinhard Klein
rk@cs.uni-bonn.de

Institute of Computer Science II,
University of Bonn,
Germany

For obtaining highly accurate 3d reconstructions, several methods combine positional information with normal information (e.g. [2], [1], [4]). Whereas triangulation-based 3D reconstruction techniques such as structured light or laser scanners typically suffer from noise or over-smoothing, reconstruction techniques based on normal information are capable of preserving high-frequency surface details but are prone to low-frequency drift due to the accumulation of errors. Fusing both types of information overcomes the individual problems.

In this paper, we propose a combination of normals estimated via Helmholtz stereopsis with structured light. In contrast to photometric stereo techniques, using Helmholtz normals has the advantage that it is largely BRDF-invariant. This is also true for structured light and, hence, our approach can be a applied to a wide range of materials. Furthermore, the structured light directly provides information about occlusion and shadowing that can be utilized in the Helmholtz stereopsis. We use a turntable-based setup which is capable of acquiring reciprocal image pairs as well as the structured light scans in an efficient and automated way.

Usually, when employing a turntable, the triangulations are performed independently for each turntable configuration. This way, not all cameras which have seen a certain location on the object surface at the same time are combined to compute one consistent point. Instead, several possibly contradicting point clouds have to be unified in the final surface reconstruction. To overcome this limitation and obtain one globally consistent reconstruction integrating the information over all rotations, we use a variational approach which combines a consistency term derived from the structured light with the Helmholtz normals. Here, the structured light consistency term allows us to combine several structured light measurements although the object was moved with respect to the projector.

This variational problem for the reconstruction of the object surface δV depends on a vector field of normals $\mathbf{H}$ and three scalar fields defined on the continuous volume $\mathbb{R}^3$: the consistency measure c, the outside count o and the visibility count v. At each point $\mathbf{x} \in \mathbb{R}^3$ and for all combinations (i, j, k) of rotations, cameras and projectors, we perform an independent classification. Utilizing the structured light information, we count the number of times the point is consistent ($c(\mathbf{x})$), lies before the surface ($o(\mathbf{x})$) and is thus outside of the object and the total number $v(\mathbf{x})$ of configurations in which it was visible from the camera. From these, we derive visibility-normalized versions $\hat{c} = \frac{c}{v}$ and $\hat{o} = \frac{o}{v}$.

Furthermore, the normal field $\mathbf{H}(\mathbf{x})$ is estimated via the Helmholtz principle [5]. We exploit the almost symmetrical mounting of the light sources around the cameras in our turntable-based setup. To compensate for the non-perfect symmetry between light source positions and camera positions in the setup, we relax the assumption of perfect Helmholtz stereopsis. Assuming that the BRDF is locally smooth enough allows us to use barycentric interpolation between three neighbouring light samples to approximate the brightness at the ideal position.

The resulting consistency-weighted vector field $c\,\mathbf{H}$ has a large magnitude in the vicinity of the surface, is aligned perpendicular to the surface and diminishes further away. To find the object interior $V \subset \mathbb{R}^3$, we seek to solve the following problem

$$\min_V \quad E(V) = -\lambda_1 \underbrace{\int_{\delta V} \langle c\,\mathbf{H}, \mathbf{n}\rangle \, \mathrm{d}A}_{E_1} + \lambda_2 \underbrace{\int_V \hat{o} \, \mathrm{d}V}_{E_2} + \lambda_3 \underbrace{\int_{\delta V} (\alpha - \hat{c}) \, \mathrm{d}A}_{E_3}, \quad (1)$$

where λ_1, λ_2 and λ_3 are relative weights of the individual terms and $\alpha > 1$ denotes a constant determining the minimum regularization strength within a consistent region. The first term E_1 considers the flux of the vector field $c\,\mathbf{H}$ through the object surface. This term is minimized by a surface running perpendicular to the reconstructed Helmholtz normals $\mathbf{H}$ and in regions with a high consistency c. The second term E_2 is used as an outside constraint to penalize regions of large values $\hat{o}$. This prevents the algorithm from short-cutting through concavities. The last term E_3 represents a regularization term and enforces a minimal surface. This penalty is weighted with the consistency $\hat{c}$ obtained from the structured light.

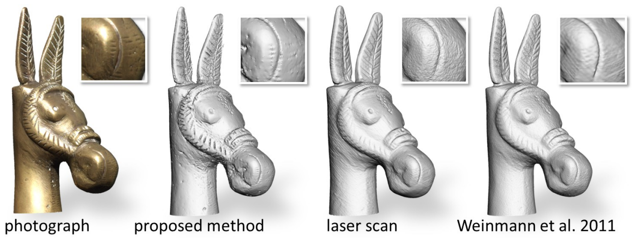

Figure 1: **Comparison of results on a glossy brass figurine.** The laser-scan was created with a high-precision line laser-scanner mounted on a measuring arm with a total accuracy of about $60\mu m$. Additionally, the results from the structured light based reconstruction exploiting projector super-resolution [3] are shown. Our reconstruction shows considerably more fine surface details.

We solve this optimisation problem via an octree-based continuous min-cut framework which is memory efficient and alleviates metrification errors. To compensate for the discretisation artefacts from the min-cut, a smooth signed distance function is then computed from the resulting binary labelling, again taking the reconstructed normals into account. Finally, the reconstruction result is derived from this smooth signed distance function obtained from the min-cut and the Helmholtz normals.

Further implementation details of the proposed approach are described in the paper. We demonstrate that the combination of structured light scanning with Helmholtz normal estimation enables the reconstructions of high-quality 3D models with a considerable amount of fine surface details. In addition, using Helmholtz normals instead of photometric stereo for normal recovery allows to handle objects with a significantly larger range of complex surface reflectance behaviour.

[1] C. Hernandez Esteban, G. Vogiatzis, and R. Cipolla. Multiview photometric stereo. *TPAMI*, 30(3):548–554, 2008.

[2] D. Nehab, S. Rusinkiewicz, J. Davis, and R. Ramamoorthi. Efficiently combining positions and normals for precise 3d geometry. *ACM Trans. Graph.*, 24:536–543, 2005.

[3] M. Weinmann, C. Schwartz, R. Ruiters, and R. Klein. A multi-camera, multi-projector super-resolution framework for structured light. In *3DIMPVT*, pages 397–404, 2011.

[4] Y. Yoshiyasu and N. Yamazaki. Topology-adaptive multi-view photometric stereo. In *CVPR*, pages 1001–1008, 2011.

[5] T. Zickler, P. N. Belhumeur, and D. J. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. In *IJCV*, pages 869–884, 2002.