

Merge2-3D: Combining Multiple Normal Maps with 3D Surfaces

Presenter Name
 Organization
 Email

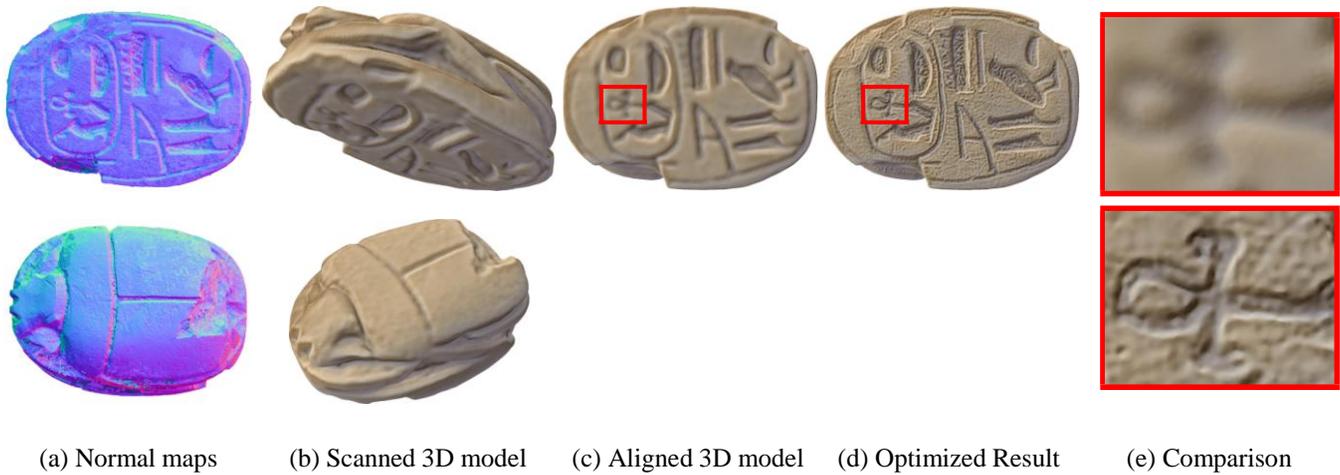


Figure 1: (a) 2D normal maps of an Egyptian scarab from the back and front. (b) Coarse 3D model of the scarab obtained with a laser scanner. (c) Coarse 3D model rendered from views matching the normal maps. (d) Our result of combining the 3D model and the normal maps. (e) Closeup images of the rough (top) and enhanced models (bottom).

Abstract

In this work, we propose an approach to enhance rough 3D geometry with fine details obtained from multiple normal maps. We begin with rough alignment between the normal maps and geometry, and automatically optimize the alignments through iterative registration. We then map the normals onto the surface, correcting and seamlessly blending them together. Finally, we optimize the geometry to produce high-quality 3D models that incorporate the high-frequency details from the normal maps. We demonstrate that our algorithm improves upon the results produced by some well-known algorithms: Poisson surface reconstruction [Kazhdan et al. 2006] and the algorithm proposed by Nehab et al. [2005].

Very clear upfront on results

Tagged this appropriately and identified audience

AUDIENCE:[Computer Graphics/Vision], [3D reconstruction], [3D/2.5D image alignment], [Photometric Stereo].

[Advanced Technical Talk]

1 Introduction

In this paper, we present a complete system to enhance rough geometry with multiple detailed normal maps. The pipeline consists of four main steps:

- Acquisition of 3D rough geometry, as well as extraction of 2D normal maps from images under varying light directions with respect to the object;
- Alignment of the normal maps to the rough geometry, beginning with a user-provided rough registration;
- Blending of data from multiple normal maps to produce a seamless normal field over the surface; and
- Optimization of the 3D model to incorporate high-frequency details from the 2D normal maps.

Our pipeline takes two types of inputs (multiple 2D normal maps and coarse 3D geometry), and produces an enhanced 3D model, as shown in Figure 1. The advantage of our system is that it requires neither a perfect initial alignment nor resolution/precision compatibility between different data types. This way, we can scan large objects such as 2–3 meter-high statues and enhance details with several normal maps of small patches of the surface (obtained, perhaps, using a digital SLR with a macro lens); or we can scan very small objects (~2–3cm) and improve their geometry, which cannot be acquired at high quality with typical laser scanners.

Our main contributions are threefold:

- An algorithm that aligns 2D normal maps to a 3D surface by iteratively minimizing projected dissimilarity of normals.
- Seamless mapping of the aligned 2D normal maps to the 3D geometry.
- A method for combining the original 3D positions with the mapped normals.

2 Data Acquisition

In this work, we experiment with both synthetic and acquired data. The synthetic dataset begins with a ground-truth 3D model of the armadillo, and acquires the following:

- 2D normal maps: We render high-resolution normal maps from multiple views. We save the viewing directions to enable ground-truth evaluation of our alignment algorithm.
- Gaussian filter with $\sigma = 4 \times (\text{average edge length})$.

Real data acquisition uses two separate capture stages:

- 2D normal maps: To obtain each normal map we capture a series of digital-SLR images of the object, with fixed camera position. We move a hand-held flash around the object, as

*e-mail

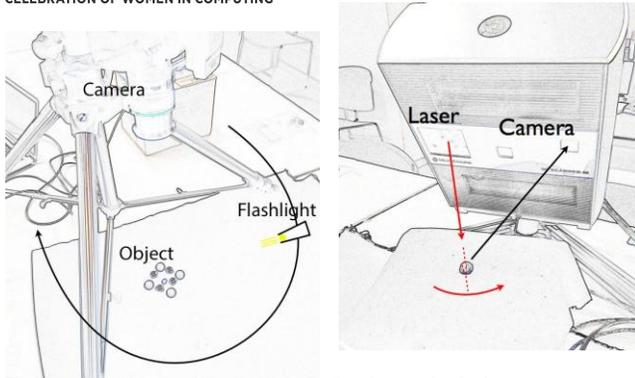


Figure 2: Acquisition setups. Left: the multi-light capture setup for photometric stereo. Right: NextEngine laser scanner used for 3D acquisition.

shown in Figure 2 left, to obtain different light directions. We place several mirror and white diffuse spheres next to the object, enabling us to estimate the light directions and intensities on each image. Finally, we use a variant of the photometric stereo algorithm [Woodham 1980] to estimate surface normals from these images. The same capture setup was also used to estimate surface normals by Toler-Franklin et al. [2007].

- 3D geometry acquisition: Of the many sensors available on the market for 3D acquisition, including laser scanners, Microsoft Kinect, time-of-flight scanners, etc., we experiment with two very different ones. First, we use a relatively high-resolution NextEngine scanner (Figure 2, right), which is based on laser stripe triangulation. We use the same setup as the one used in [Brown et al. 2008], which scans the object from n evenly distributed angles for both back and front sides, and aligns all the partial scans with multi-way ICP. Later, Poisson surface reconstruction [Kazhdan et al. 2006; Kazhdan and Hoppe 2013] is used to obtain a watertight surface from the point cloud of all aligned partial scans.

To explore even lower-resolution geometric input, we have also experimented with Kinect [Microsoft, Inc. 2010], with the Skanect [Occipital, Inc. 2013] software used to reconstruct the surface from the captured data.

3 Method

Our algorithm consists of three main steps:

- Alignment: Inspired by ICP, we propose an iterative alignment algorithm that begins with a rough alignment between a 2D normal map and 3D model and improves the registration by minimizing an energy function, which measures the dissimilarity between corresponding geometric and photometric normals. To optimize this energy function we need to transform the model and render the normal map, making the problem non-linear. We therefore use the Nelder-Mead algorithm [Nelder and Mead 1965], which initializes a simplex in the 6-dimensional rotation-translation space and heuristically grows, moves, and shrinks the simplex.
- Blending: Inspired by Poisson image editing [Pérez et al. 2003], we propose blending between the overlapping normal maps in a way that preserves their gradients. To do this we optimize an energy function with two terms -with linear least squares-, a vertex term that pushes the result towards a con-

sensus (weighted average) among samples, and an edge term that preserves differences between neighboring normals.

- Enhancement: Given high-quality normals and rough vertex positions, we approximate the true surface by performing a non-linear optimization, minimizing an energy function consisting of two terms: position and normal terms.

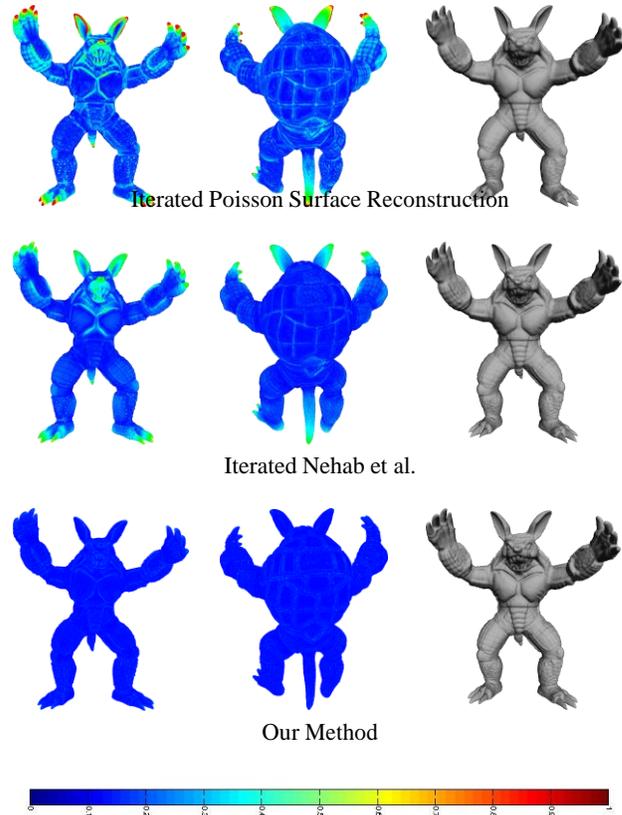


Figure 3: Positional errors of Poisson surface reconstruction, Nehab et al., and of our method with respect to the ground truth. Blue represents zero distance between the reconstructed surface and ground truth, while red is greater than or equal to 0.001 of the radius of the bounding sphere.

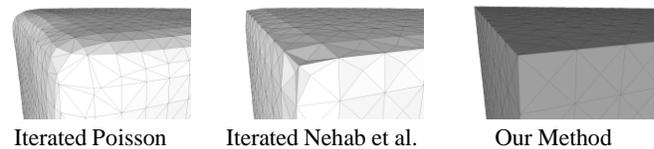


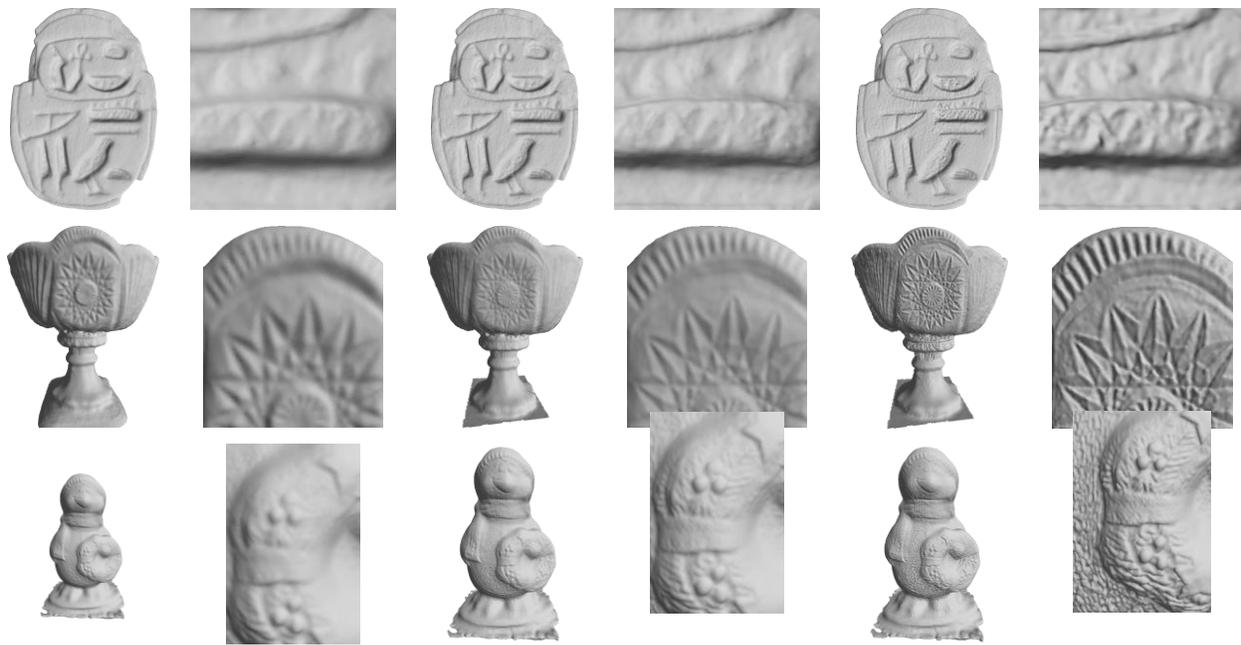
Figure 4: Comparison of results of Poisson surface reconstruction, Nehab et al., and our method on a cube. The ground-truth vertex positions and normals of the cube were used as input.

4 Outcome/ Results and Discussion

Figure 3 shows a quantitative comparison of our method to Poisson surface reconstruction [Kazhdan et al. 2006] and the method of Nehab et al. [2005]. We

Presented results in context of other solutions out there.

found that the latter two methods exhibit less smoothing if the surface positions are close to the correct ones, and so we actually run these methods in an iterative manner.



Iterated Poisson surface reconstruction

Iterated Nehab et al.

Our Method

Figure 5: Comparison of results of Poisson surface reconstruction, Nehab et al., and our method. The 3D models in the first row is acquired with the NextEngine laser scanner, while the last two are obtained with MS Kinect.

In Figure 4, we used a tessellated cube to test an extreme case. We used the original positions and normals of the cube as input, in order to observe the behavior of the three algorithms in a controlled setting. The previous algorithms are unable to preserve the sharp corners of the cube, introducing significant smoothing. In contrast, our method neither introduces any artifacts nor smooths over the corners.

Figure 5 demonstrates results for the three algorithms, using source data from two different types of 3D acquisition techniques: moderate-resolution models acquired with the NextEngine laser scanner and coarse models obtained with MS Kinect. The left-most column in Figure 5 shows the results and close-up images for the iterative Poisson reconstruction; the middle column shows the results of the algorithm proposed by Nehab et al. [2005]; and finally the rightmost column shows our results. All of these used our proposed alignment and blending algorithms.

Our method preserves more of the high-frequency information from the normal maps, as compared to the other two methods. Of course, in some cases this may also mean that our method preserves more noise, since it does not smooth the normal maps.

PARTICIPATION STATEMENT

I will attend the conference if accepted.

BIO (small paragraph bio)

References

BROWN, B. J., TOLER-FRANKLIN, C., NEHAB, D., BURNS, M., DOBKIN, D., VLACHOPOULOS, A., DOUMAS, C.,

RUSINKIEWICZ, S., WEYRICH, T., AND WEYRICH, T. 2008. A system for high-volume acquisition and matching of fresco fragments: reassembling Theran wall paintings. *ACM Trans. Graphics* 27, 3, 84.

KAZHDAN, M., AND HOPPE, H. 2013. Screened Poisson surface reconstruction. *ACM Trans. Graphics* 32, 3, 1–13.

KAZHDAN, M., BOLITHO, M., AND HOPPE, H. 2006. Poisson surface reconstruction. In *Proc. Symposium on Geometry Processing*, 61–70.

MICROSOFT, INC., 2010. Kinect. <http://www.xbox.com/en-US/kinect>.

NEHAB, D., RUSINKIEWICZ, S., DAVIS, J., AND RAMAMOORTHY, R. 2005. Efficiently combining positions and normals for precise 3D geometry. *ACM Trans. Graphics* 24, 3, 536–543.

NELDER, J. A., AND MEAD, R. 1965. A simplex method for function minimization. *The Computer Journal* 7, 4, 308–313.

OCCIPITAL, INC., 2013. Skanect 1.5. <http://skanect.manctl.com/>.

PÉREZ, P., GANGNET, M., AND BLAKE, A. 2003. Poisson image editing. *ACM Trans. Graphics* 22, 3, 313–318.

TOLER-FRANKLIN, C., FINKELSTEIN, A., AND RUSINKIEWICZ, S. 2007. Illustration of complex real-world objects using images with normals. In *Proc. NPAR*, 111–119.

WOODHAM, R. J. 1980. Photometric method for determining surface orientation from multiple images. *Optical Engineering* 19, 1, 139–144.

Inclusion of references helps the reviewers distinguish a good idea from work that has been tried and tested.