

Global Optimization for Shape Fitting

Victor Lempitsky

Yuri Boykov

Abstract

This work proposes a global optimization framework for 3D shape reconstruction from noisy sparse 3D measurements frequently encountered in range scanning, sparse feature-based stereo, and shape-from-X. Firstly, we suggest a novel surface-fitting energy for sparse or incomplete noisy data. Our flux-based functional maximizes the number of data points contained by a surface while allowing for some measurement errors. Combined with regularization, this geometric functional significantly improves data alignment.

We also show that typical local or “banded” optimization techniques do not recover from initialization errors while standard global optimization methods are prohibitive due to sheer size of 3D grids required for high-resolution surface reconstruction. To overcome these limitations, we propose a first graph-cut algorithm that guarantees global optimality on huge 3D grids while working on automatically adjusted small sub-graph. Our global optima surface-fitting results for noisy point clouds or sparse/incomplete multi-view disparity maps are robust to noise, large missing parts, and varying sampling density.

1. Introduction

Surface fitting to sparse points is widely studied in vision and graphics. The main two contributions of this work are as follows. First, we address the surface fitting problem via a novel data-fit functional directly derived from quality of feature-matching at data points. Second, we propose a new memory-efficient global optima algorithm for high-resolution surface fitting to sparse data that alleviates problems of standard local minima and narrow-band methods.

Our novel surface fitting functional directly enforces geometric proximity to data points. Intuitively, it counts data points contained by a surface while allowing for localization errors. Data points may have different weights reflecting confidence level that a point is not an outlier. We also assume that each data point comes with some estimate of surface orientation - a vector that we use to softly constrain (outward) normal of a surface that fits the point. Unlike many previous methods using estimated surface normals, our approach requires orientations only for observed (sparse) data points and is robust to orientation errors, see

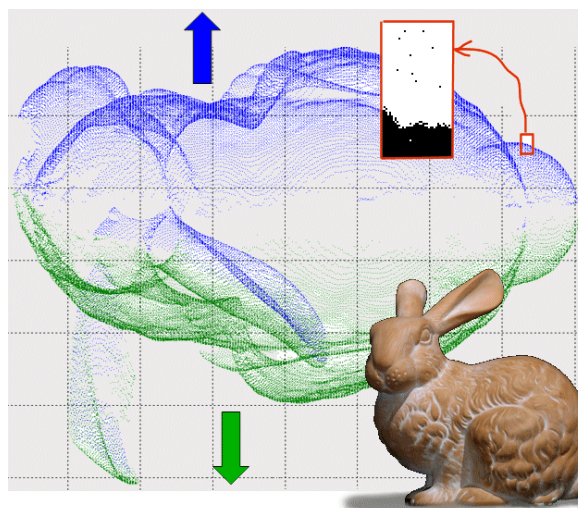


Figure 1. An example of input data for our method showing two (out of 10) range scans registered in 3D. Points in each scan come with weak orientation estimates (blue and green arrows) sufficient for our method. Typical noise and outliers are shown in a close-up (from a different and more complete set of 112 scans). We also show one image of the scanned object (not a part of our method’s input). This data is courtesy of the Stanford 3D repository.

Fig. 12. In practice, we produce good results even using directions towards the sensor/camera which are often known for data points (blue or green arrows in Fig. 1).

Our data-fit functional does not bias a reconstructed surface to any particular shape but it can be combined with appropriate shape priors (smoothness/regularization), volumetric occupancy data, or other terms. For example, our discrete data fit energy can be combined with a separate photoconsistency term for constraining the surface in featureless regions.

In contrast to many earlier regularization methods for fitting a surface to sparse data, global optimization of our cost functional does not result in a trivial (empty) solution. Thus, our energy is well suited for global minima algorithms which are typically more robust in practice. Unlike standard local or banded optimization methods, our reconstruction results do not depend on initial solution.

Global optimization methods for surfaces, e.g. graph-cut [2], typically require a lot of memory for high resolu-




Figure 2. Global vs. local/banded optimization in 2D shape reconstruction. We minimize a functional similar to (I) or (II). The data term of the energy is based on pixels’ potentials (a): dark – negative, bright – positive, gray – zero. The regularization term is geometric length of the shape’s boundary. Image (c) shows results of local optimization (level sets) from initial solution (b). Global optimization (graph-cut on a full-size 2D grid) is in (d). Our method detects such global minima using a small subgraph (e) automatically grown from a narrow-band around the same initial shape (b).

tion volumes. Standard hierarchical or narrow band techniques do not guarantee global minima and generally behave as local optimization methods, see Fig.2,3. In contrast, our new memory-efficient graph-cut algorithm guarantees global minima. It can compute global minima surfaces at fairly high resolutions on regular PCs. In the context of surface fitting problem, our technique computes regularized watertight surfaces demonstrating good alignment to data points without over-fitting to outliers. Our memory efficient graph-cut algorithm may potentially be generalized to other optimization problems in computer vision.

1.1. Global vs. local optimization of surfaces

There are many existing techniques for fitting 3D shapes to discrete data points. This paper follows a minimal surface (regularization) approach where output optimizes some energy functional encoding available data and prior assumptions about the shape. We choose to work with geometrically motivated functionals that typically include terms like (non-Euclidean) surface area, flux, and regional bias (volumetric potentials). Minimal surface methods typically yield consistent and predictable watertight surfaces, run reasonably fast, and, most importantly, are robust to noise, outliers, and large areas of missing data.

Commonly used geometric functionals are non-convex and may have many local extrema. Consequently, the particular choice of optimization method is important. One popular surface optimization approach is based on variational methods using deformable meshes or implicit level set representation of the shape. However, variational methods find local minima which may depend on initialization and on numerical implementation details (Figure 2a-c).

Recently it was demonstrated that many continuous geometric functionals can be optimized via combinatorial graph-cut algorithms on grids [2, 19]. This approach is motivated by integral geometry. It approximates continu-




Figure 3. A cloud of data points (left). **Banded graph-cut** from unoriented points [16] (middle): incorrect estimate of initial shape may lead to significant reconstruction errors. **Global graph-cut** approach proposed in this work (right) avoids local minima problem. We also show that even rough estimates of data points’ orientation may significantly improve alignment with the data.

ous surface functionals using regular N-neighborhood grids with specific edge weights. The optimal surface is implicitly represented by the minimum cut. As grid resolution and connectivity increase, this discrete approximation becomes more accurate. In contrast to local variational methods, graph-cut approach yields a global minimum (Figure 2d) and it is not sensitive to initialization. On the negative side, current min cut algorithms require a lot of memory when used on high resolution grids.

The objective of this paper is to develop a global optimization framework for 3D shape fitting problems. To develop this framework, we first propose a novel continuous geometric functional for surface-to-data fit (see Sec.3). Our flux-based functional is amenable to global (e.g. graph-cut) or local (e.g. level-set) optimization methods. Essentially, we maximize the number of points that lie on a surface constrained by a shape prior. Yet, local optimization of our functional is sensitive to initialization (Fig.2b,c) while global optima solution is consistent (Fig.2d).

We also study the problem of computing globally optimal cuts (surfaces) on high resolution volumes while taking advantage of high concentration of data points in a very

small part of the volume. We would like to restrict computations to a subgraph (e.g. narrow band), *i.e.* some part of the grid surrounding the actual data points. Yet, existing narrow-band strategies [22, 16] do not guarantee global optimality making the result depend on the actual shape of the narrow band and on initial guess. Essentially, previous narrow-band techniques converted graph-cut into a local optimizer similar to level-sets or other variational methods causing local minima problems, e.g Fig.3 (middle).

In order to retain global optimality, in Section 4.1 we prove a lemma that relates a minimum cut on a subgraph to the minimum cut on the complete graph. Based on this lemma, we propose a memory efficient graph-cut algorithm that automatically adjusts the shape of the subgraph until its minimum cut is guaranteed to coincide with the minimum cut in the whole volume. This allows computation of a globally optimal minimum cut using a subgraph (Fig. 2d,e). Our surface fitting experiments showed that our algorithm can compute globally optimal cuts using an order of magnitude less memory than what normally would be required for a full resolution graph. These memory savings resulted in a leap forward in resolution at which globally optimal surface can be computed significantly increasing accuracy of shapes generated by local optimization (see Fig.2) or by standard narrow-band methods (see Fig.3).

Experiments show that our shape fitting method reconstructs high-resolution watertight surfaces in presence of noise, outliers, large missing parts, and varying sampling density. Global optimality of our surface fitting solutions implies no initialization issues, better robustness, consistency, and possibly wider scope of applications.

1.2. Surface fitting for image-based data

In many applications surface measurements (sparse 3D data points) come from triangulation of corresponding features in calibrated image sensing (passive light) and/or image projecting (active light) devices. Since our new surface-to-data fit functional can be directly related to the quality of matching between image features, it is possible to place our surface fitting method in the context of existing approaches to image-based shape reconstruction.

Many image-based reconstruction methods generate continuous 3D shapes based on surface photoconsistency. Yet, photoconsistency does not imply accuracy for surface reconstruction in textureless regions where lack of features prevents accurate triangulation (see Fig.4). Typically, reconstruction of textureless parts is based on assumptions about the surface, e.g. smoothness/regularity in a deformable model approach or maximal size in space carving.

An alternative standard approach to image-based shape reconstruction is to generate sparse 3D points by matching discriminant features. Typically, feature-based sparse stereo methods are fast and localize surface points with good pre-

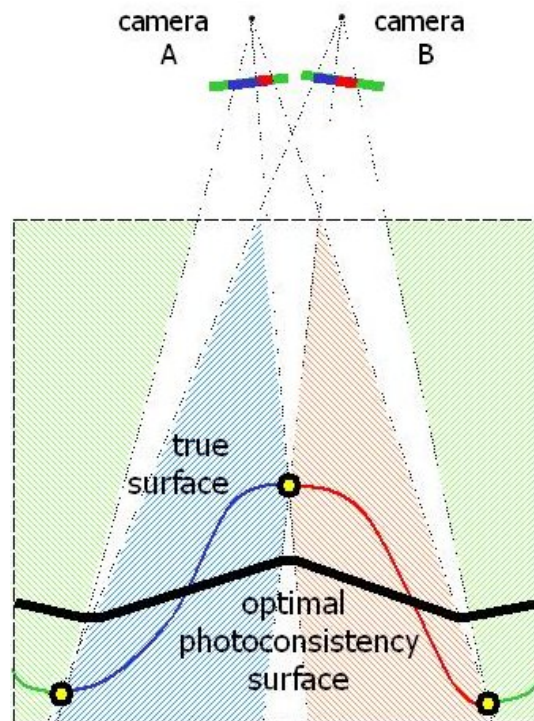


Figure 4. Photoconsistency vs. feature-based stereo. Three (yellow) data points are accurately localized by triangulation of the corresponding image features (intensity edges). Shaded areas show regions of good photoconsistency between green, blue, and red pixels of cameras A and B. Minimal photo-inconsistency surface may be far from the feature points. The problem is the “shrinking bias” typical for regularization methods and the fact that all points in the shaded area have the same cost even though three feature points are the most useful in finding the true surface position. In practice, the problem could be even worse because shaded regions may not even “touch” at feature points. Indeed, corresponding pixels near intensity edges may be photo-inconsistent due to sampling issues and partial voluming effects causing some erosion of high photoconsistency (shaded) regions.

cision¹. Yet, outliers could be a problem. Also, many applications need continuous shapes. Typically, continuous surfaces are fit to sparse data points in a separate step where some shape priors are introduced in order to deal with outliers, noise, and large gaps.

Our new surface fitting method optimizes an energy combining two terms: a shape prior (regularization) and a data-fit cost functional related to the quality of the matching between image features. Unlike many existing approaches to feature-based stereo, we combine feature matching costs and prior shape (regularization) constraints in a single optimization step. Thus, our approach can be directly compared to photoconsistency-based regularization methods for stereo combining image data and a shape prior in one en-

¹Laser scanning can be seen as an example of feature-based stereo.

ergy. In contrast to standard photoconsistency functionals, however, our flux-based data-fit cost at each point is not bounded below by zero (indicating “consistency”) and points with discriminant features can have large negative values actively attracting minimal surfaces.

Fundamentally, both photoconsistency-based and discrete data approaches to 3D surface reconstruction rely on two similar basic factors: the accuracy comes from relatively sparse features such as intensity edges, see Fig. 4, or silhouette edges [13] while the continuity comes from explicit or implicit shape priors. Yet, as illustrated in Fig. 4, optimal photoconsistency of a surface may contradict geometric alignment with feature points which are the main indicators of the true surface position. Thus, directly enforcing geometric proximity of a surface to 3D data points may improve accuracy. Another advantage of using discrete measurements could be efficiency gained from sparsity.

Computation of depth (disparity) can be seen as a 1D version of a more general *optical flow* (2D motion) problem. Interestingly, the basic problem illustrated in Fig. 4 connects to some known properties of standard optical flow methods. As widely known [1, 6], sparse (Lucas-Kanade) techniques that focus on points with sufficiently discriminant features can be more accurate and robust to noise than dense (Horn-Schunck) techniques that enforce color-consistency and regularization. Similarly to Fig. 4, the Horn-Schunck approach can oversmooth important data points because color-consistency (or optical flow constraint) does not distinguish points with reliable measurements from color-consistent points in textureless parts. Ideally, dense optical flows should be constrained to geometrically fit motion detected at reliable data points. Unfortunately, many standard techniques for dense optical flows (including [6]) do not use such geometric constraints.

2. Related work on surface fitting

There is a large amount of work on fitting continuous shapes to 3D sparse/incomplete data and we can reference only a fraction of existing publications. The majority of standard methods can be grouped according to their approach to numerical representation of surfaces (see Sec. 2.1). Due to commonly present noise, outliers, and large missing parts, energy-based methods are particularly appropriate for robust surface fitting. In Sec. 2.2 we review existing regularization-based surface fitting techniques and the surface functionals that they use.

2.1. Explicit and implicit surface representation

First, there is a group of methods using explicit representation of surfaces via triangle mesh [9, 11]. This approach allows optimization of geometrically-motivated energies using gradient descent but it can be highly dependent

on good initialization. Also, if the shape undertakes significant changes during gradient descent evolution, mesh-based techniques have to use different heuristics for handling topological changes. Mesh quality control is another significant practical issue.

Many methods for surface fitting avoid mesh-related numerical problems by representing surfaces implicitly. An earlier group of such methods, including widely cited work by Curless and Levoy [10], used level-set functions to represent surfaces but did not employ any surface regularization functionals. To avoid any confusion with the widely known *level-sets* technique of Osher and Sethian [24], it should be noted that [10] and other similar methods do not compute the gradient descent evolution of a surface. They use data points only to estimate some dense “interior” function in 3D space and reconstruct a continuous shape as its isosurface (level-set). For example, the VRIP algorithm of Curless and Levoy [10] computes a weighted sum of signed distance functions of individual range scans and thresholds it (at level zero) to recover a continuous surface. Another group of methods (e.g. [7]) compute an interior function as a combination of radial-basis functions. The Poisson surface reconstruction algorithm [17] considers oriented points as samples of gradients of the interior (signed distance map) function and reconstructs it using the Poisson equation.

Any surface fitting method that lacks geometrically motivated regularization is prone to problems when the data contains large gaps and outliers. Some of these methods are also sensitive to the accuracy of point orientations estimates and varying sampling density. To counteract such problems, another group of algorithms (e.g. [28, 29, 25]) combine implicit surface representation with regularization via the well-known *level-sets* technique of Osher and Sethian [24]. This variational method can compute the gradient descent evolution of an implicitly represented surface with respect to a given cost functional. Level-sets cope seamlessly with topological changes and avoid re-meshing issues. Yet, surface fitting methods based on level-sets converge to a local minimum of the corresponding functionals. Thus, they depend on proper initialization. Some regularization functionals for surface fitting are reviewed in Sec. 2.2.

Our work uses a different approach where surfaces are represented as binary partitionings (s/t cuts) on discrete grids [2, 5]. This implicit representation technique is based on integral geometry [2] and allows global optimization of a large class of surface functionals [19] via efficient combinatorial graph-cut algorithms. Robust global graph-cut algorithms have already demonstrated a strong potential for many problems in computer vision, but they have not been applied to surface fitting yet. One exception is a very recent method of Hornung and Kobbelt [16], but they compute an optimal surface in a fixed narrow band (a crust) and the shape of this crust directly affects the optimized energy.

Similarly to variational methods, their results depend on the initialization, *i.e.* on the particular shape of the constructed crust. Using banded (non-global) optimization may lead to gross errors in reconstruction, while discarding orientation information may lead to oversmoothing (Figure 3). Our graph-cut based algorithm is a first surface fitting method that guarantees global optima in the whole volume. We also designed a novel geometric surface-to-data fit functional suited for global optimization.

2.2. Geometric surface fitting functionals

Regardless of numerical approach to surface representation, different regularization-based surface fitting techniques can be compared by the energy they use in order to cope with noise, outliers, and large gaps. Since we use a geometric approach, we mainly concentrate on geometric regularization functionals used for surface fitting in the past.

To merge incomplete range scans obtained from multiple view points Whitaker [28] uses functional

$$E(S) = \int_{\text{interior}(S)} D(v)dv + \int_S ds \quad (1)$$

where $D(v) = \sum_j D_j(v)$ is a potential function based on a combination of signed distance maps of independent scans obtained from multiple view points. Function $D(v)$ is very similar to the interior function used by Curless and Levoy [10]. Both [28] and [10] assume that each j -th scan is dense which allows the corresponding signed distance function to be defined in a straightforward fashion. In fact, minimizing the first term in (1) is equivalent to extracting a zero-level isosurface of $D(v)$ which is exactly the algorithm in [10]. The main contribution of [28] was to combine the data-fit energy as in [10] with the second regularization term in (1) thereby improving robustness to outliers and noise.

Zhao, Osher, and Fedkiw [29] proposed a different regularization-based surface fitting functional applicable to an arbitrarily sparse set of points

$$E(S) = \int_S d_P^n(s) \cdot ds, \quad 1 \leq n \leq \infty \quad (2)$$

where $d_P(s)$ is a distance from point s on surface S to the nearest data point in set P . Function $d_P(\cdot)$ is an unsigned distance map that is well defined for any set of points P . Thus, functional (2) can be applied to a wider class of surface fitting problems than (1) as the set of points P can combine arbitrarily sparse data from different views and no estimates of surface orientation are required.

Functional (2) combines surface-to-data proximity with smoothing/regularization into one term. Similarly to geodesic active contours [8], functional (2) can be seen as the area of surface S under a Riemannian metric d defined by the data. As discussed in [29], optimization of (2) via

variational (level-set) approaches will converge to a local minima of this functional: a surface composed of facets from Delaunay triangulation for a given sparse set of points. The particular local solution will depend strongly on good initialization.

Hornung and Kobbelt [16] apply graph-cut surface optimization [2] to a sparse data fitting functional very similar to (2). However, functional (2) is not appropriate for global optimization since its global minima is a trivial (null) surface. Thus, [16] compute their solution in a (fixed) narrow band reducing graph-cut framework to local optimization essentially equivalent to variational approach in [29]. Connection between narrow-band graph-cut methods and variational optimization (e.g. level-sets) was shown in [4].

Savadjiev, Ferrie, and Siddiqi [25] formulate surface fitting as a problem of estimating a dense vector field $\{v_p\}$ of surface normals from sparse data points. Then, a continuous surface can be recovered from a dense field $\{v_p\}$ by optimizing a *flux* functional

$$E(S) = \int_S \langle v_s, n_s \rangle \cdot ds \quad (3)$$

where n_s is a surface normal. To estimate a dense vector field $\{v_p\}$, [25] finds a local surface patch of the least square fit at each sparse data point and use regularization framework to propagate local patch information around. In particular, they use *variational relaxation* on a discrete 3D grid to enforce curvature consistency over local (quadric) patch models. Once local patch models are estimated, normal vectors $\{v_p\}$ are determined with ± 180 degree ambiguity due to lack of global surface orientation in the formulation of [25]. Thus, [25] cannot recover a continuous surface by optimizing flux (3) with respect to the obtained $\{v_p\}$. They propose some heuristic to overcome the ambiguity of their vector field orientation but it causes geometric artifacts (e.g. “thickness” [25]) and reconstruction errors.

We believe that it is possible to resolve the ambiguity of an estimated dense vector field $\{v_p\}$ in [25] by enforcing global orientation consistency. At each grid node p one can select one of two possible vectors $\pm v_p$ in a spatially consistent way according to a binary MRF formulation with pairwise smoothness between neighboring grid points. Such MRFs posterior energies (possibly super-modular) can be optimized using graph-cut, QPBO methods [20], and/or other powerful combinatorial optimization techniques. Yet, in this work we study a different approach to surface fitting. In contrast to [25], we do not want to brake the surface fitting problem into two (artificially?) separated steps: computing a dense field $\{v_p\}$ of normals by propagating surface patches estimated in sparse data points and then computing a surface that fits this dense field of normals.

2.3. Our geometric approach vs. prior art

Our insight is that flux-based data fit functionals like (3) do not have to rely on dense and accurate estimates of surface normals $\{v_p\}$. Unlike [25], we skip attempts to accurately estimate a dense field of surface normals and use only some weak estimates of global surface orientation at sparse data points $p_i \in P$. We assume a sparse vector field

$$\{v_i \mid p_i \in P\}$$

where each vector v_i softly constrains true surface normal n_i at data point p_i . In particular, we expect large positive values of the cosine of the angle between vectors v_i and n_i . Yet, we show that even fairly coarse orientation estimates at data points $p_i \in P$ work well in practice. For example, the experimental results in this paper use the direction from a data point p_i to the corresponding sensor². We also note that v_i does not have to be a unit vector and its length may indicate confidence of data point p_i and/or its orientation.

Before computing flux (3), we convert our sparse vector field $\{v_i \mid p_i \in P\}$ to a semi-dense field by “blurring” each vector v_i in a small neighborhood of point p_i using a Gaussian (or other) kernel corresponding to the measurement error or uncertainty of data point p_i , see Fig. 5b. In Section 3.1 we show that flux with respect to such semi-dense vector field can be directly related to the quality of feature matching. Despite blurring, such semi-dense vector fields may still have large gaps due to data sparsity. Note that our flux based data-fit functional is linear w.r.t. combinations of multiple data points (or scans).

In contrast to [25], we combine flux of a semi-dense field of “rough surface normals” with regularization (or shape prior³) in a single geometric functional and compute its global minima. The regularization term propagates information and resolves ambiguities while flux enforces alignment with the observed data points and counteracts the “shrinking bias” of standard regularization (length/area) functionals. It is known that combining flux with regularization allows reconstruction of elongated structures, narrow protrusions, or other fine surface details [26, 18, 19, 5].

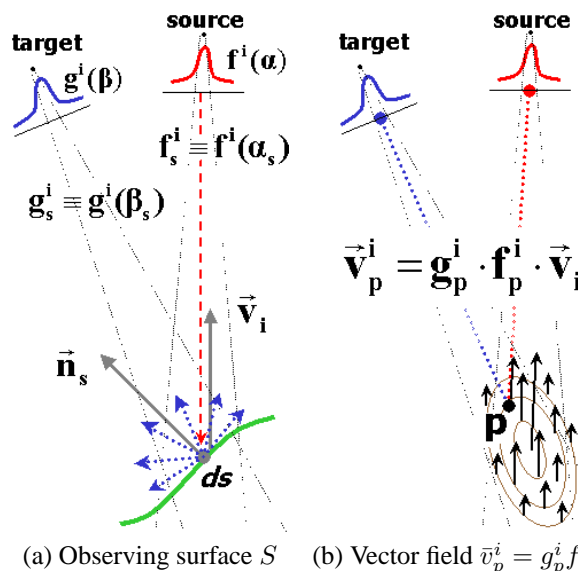
3. Our energy formulation

Discrete measurements are common in range scanning, feature-based sparse stereo or shape-from-X. Dense but incomplete data⁴ can also be seen as a collection of discrete data points. In all cases, data points come with some measurement errors or other forms of uncertainty. We derive a

²If sparse data points $p_i \in P$ come with absolutely no surface orientation information, a sparse vector field can be obtained by locally fitting plane patches into points $p_i \in P$ and by removing ± 180 orientation ambiguity via global consistency optimization (as suggested earlier) on a binary MRF with nearest k-neighbor or Delaunay-graph connectivity.

³E.g. elastic membrane (low area) or stiff membrane (low curvature).

⁴As in multiview reconstruction from multiple disparity maps



(a) Observing surface S (b) Vector field $\vec{v}_p^i = g_p^i f_p^i \vec{v}_i$

Figure 5. Data points from active or passive reconstruction techniques (e.g. range scanning or feature-based stereo) are not as discrete as they may appear. For example, the i -th data point from a laser scanner (a) corresponds to noisy non-deterministic measurements of an illuminated patch which could be smaller or larger depending on the beam’s width and the patch’s orientation. In laser scanning, “source” image $f^i(\alpha)$ is a probability of emitting a photon in direction α and “sensor” image $g^i(\beta)$ measures radiance of surface points viewed at angle β . Statistically motivated functionals like (6) suggest that sparse/discrete surface measurements can be represented via flux of continuous vector fields (b).

novel geometrically motivated surface-to-data fit functional (Sec.3.1) and combine it with surface regularization or other generic forms of shape prior (Sec. 3.2).

3.1. Quality of Fit to Data

Ideally, if all discrete data points have absolute precision then the quality of a surface fit can be measured by the number of data points that lie on it. Reconstruction of a continuous/dense surface would require imposing a shape prior (e.g. regularization) even for perfect discrete data but this factor is separate from the quality of fit to data.

In practice, it is well known that even range scanning measurements are not exact, not to mention accuracy of feature-based sparse stereo. How should the quality of a fit be measured in the presence of noisy data? In order to clarify this question it is very instructional to analyze the uncertainty of a single data point in laser scanning which is often considered to be the most accurate approach to surface reconstruction. Our basic analysis can be extended to other active or passive light acquisition methods and laser scanning is chosen mainly as the simplest example illustrating our main approach.

Consider the active light surface measurement process




Figure 6. Maximization of flux through vector field $\bar{v}_p = \sum_{i=1}^N \bar{v}_p^i$ can be seen as maximization of the number of data points contained by a surface while allowing measurement uncertainty (a). This is equivalent to maximizing divergence of vector field $\{\bar{v}_p\}$ inside surface interior (b). Points with non-zero (positive or negative) values of scalar function $div(\bar{v}_p)$ are indicated by + or -.

shown in Figure 5 (a). The amount of light energy reaching surface patch ds from the source in a unit of time is

$$ds \cdot f_s^i \cdot \langle \bar{v}_i, \bar{n}_s \rangle \quad (4)$$

where \bar{n}_s is surface normal, \bar{v}_i gives direction to the source, and $f_s^i := f^i(\alpha_s)$ for angle α_s at which the source views given point s in space. The Lambertian assumption about surface S implies that *radiance* of points on patch ds is the same in all angular directions. Since image sensor measures radiance of surface points then the total energy of light reflected by patch ds in a unit of time is

$$ds \cdot \pi \cdot g_s^i \quad (5)$$

where $g_s^i := g^i(\beta_s)$ for angle β_s at which the sensor observes point s in space. There is a connection between (4) and (5): if surface *albedo* is ρ then $\pi \cdot g_s^i = \rho \cdot f_s^i \cdot \langle \bar{v}_i, \bar{n}_s \rangle$. Functions (4) and (5) are defined for points s on any surface S . However, the linear relationship

$$g_s^i \propto f_s^i \cdot \langle \bar{v}_i, \bar{n}_s \rangle \quad \forall s \in S$$

only holds if points s and normals n_s belong to the (unknown) surface that reflected the light. It is natural to estimate surface S by maximizing some measure of similarity between the two functions on S .

For example, maximization of the dot product

$$\max \leftarrow \int_S g_s^i \cdot f_s^i \cdot \langle \bar{v}_i, \bar{n}_s \rangle \cdot ds \quad (6)$$

could be a reasonable approach to align (4) and (5). Alternatively, one can maximize *covariance*⁵ of (4) and (5) as functions (random variables) over $\Omega = S$ with respect to Lebesgue measure

$$\max \leftarrow \int_S g_s^i \cdot f_s^i \cdot \langle \bar{v}_i, \bar{n}_s \rangle \cdot ds - \int_S c \cdot g_s^i \cdot ds \quad (7)$$

⁵ $Covar(X, Y) = E(X - \bar{X})(Y - \bar{Y}) = E(XY) - \bar{X}\bar{Y}$.

where constant $c = \int f^i(\alpha) \cdot d\alpha$ (total energy of the source) is independent of S . Large covariance indicates that two functions “vary together.” Note that functionals (6) and (7) are essentially the same except for an additional smoothness term in (7). Since surface regularization is imposed in a separate shape-prior functional (see Sec.3.2), we will concentrate on *flux*-based data fit functionals like (6).

Note that Lambertian assumption above is non-essential for our main argument as long as there is a correlation between the irradiance on patch ds and the reflected light radiance in the direction to the sensor.

In case of multiple data points $1 \leq i \leq N$, our measure of surface fitness sums functionals (6) for all points i . This is analogous to counting the number of points contained by S in the case of absolutely precise data. We obtain functional

$$\max \leftarrow \int_S \langle \bar{v}_s, \bar{n}_s \rangle \cdot ds \quad (8)$$

which is *flux* for vector field $\bar{v}_p = \sum_{i=1}^N \bar{v}_p^i$ representing all (uncertain) data points, see Figures 5(b) and 6(a). Note that Gauss-Ostrogradsky (a.k.a. *divergence*) theorem

$$\int_S \langle \bar{v}_s, \bar{n}_s \rangle \cdot ds = \int_{inter(S)} div(\bar{v}_p) \cdot dp \quad (9)$$

implies that (8) is equivalent to maximizing the integral of vector field’s *divergence* in the interior of S , see Figure 6(b).

We showed that flux-based functionals like (9) are well justified data-fit measures accounting for uncertainties or noise in the data points. This approach easily generalizes beyond laser scanning. In the majority of existing active or passive light methods for acquiring discrete/sparse surface measurements, each i -th data point naturally comes with direction to the source \bar{v}_i or even better estimate of surface orientation. Moreover, instead of a single data point p_i most methods can report some distribution function $\rho_i(p)$ describing probability that i -th measurement corresponds to a surface patch located at point $p \in \mathcal{R}^3$. The specific form of $\rho_i(\cdot)$ depends on particularities of the specific acquisition method. In many cases this distribution is a Gaussian with a given mean (data point p_i) and some covariance matrix specific to each technique. In general, we suggest flux-based functional (9) as a generic surface-to-data fit quality measure where vector field $\{\bar{v}_p\}$ representing data is

$$\bar{v}_p = \sum_i \rho_i(p) \cdot \bar{v}_i \quad \forall p \in \mathcal{R}^3 \quad (10)$$

In the special case of laser scanning we had $\rho_i(p) = g_p^i \cdot f_p^i$ which in most cases is a Gaussian whose variance depends on the laser beam width. In feature-based sparse stereo, variance depends on the size of the image features.




Figure 7. Surface fitting with different sampling density and priors. a) Optimization of the data-fit term in combination with a membrane prior produces consistent reconstruction for densely sampled points. b) The same fails when the sampling density is low. c) Widening the support of each point (ad-hoc “low curvature” prior) allows to obtain correct reconstruction.

3.2. Adding Shape Prior

In the presence of large gaps and outliers in the data, imposing some shape prior is essential for surface reconstruction. One simple approach is to augment functional (8) with an area-based regularization term (*elastic membrane* prior) giving the following minimization problem:

$$\min \leftarrow \int_S \lambda \cdot dS - \int_S \langle \bar{v}, \bar{n}_s \rangle dS \quad (11)$$

where $\{\bar{v}_p\}$ is a vector field representing data, as in (10).

The exact value of λ determines the strength of the membrane prior and may be chosen according to the sampling density. If sampling density varies significantly over different parts of the true surface, a better approach is based on spatially varying $\lambda = \lambda(p)$ proportional to some estimate of local data sampling density. In our experiments, however, constant λ gave good result and we used constant λ in all experiments (unless noted otherwise).

The optimization of the functional (11) maximizes the number of collected data points while minimizing the area, thus handling noise and outliers in the initial data, see Figure 7(a). However, if data points are sampled at distances much sparser than each point’s “support”, see Figure 7(b), then the membrane prior may produce inconsistent results for all choices of fixed λ . It is unclear if spatially varying λ can solve the problem. Potentially, one may try combining our flux-based data-fit term with a regularization functional similar to (2) where value of spatially varying regularization parameter λ could be set according to the distance to the nearest facet on the data points’ Delaunay graph.

An alternative and probably better approach is to use a low-curvature (*stiff membrane*) prior instead of a low-surface (*elastic membrane*) prior. The corresponding variational models are known. In the context of global graph-cut optimization it might be possible to incorporate low curvature priors using N-cliques but, to the best of our knowl-

edge, an exact construction is unknown. Yet, we propose a simple heuristic for low curvature shapes. Note that each observed data point p_i with estimated orientation \bar{v}_i may correspond to a small surface patch of certain size Δ . The smaller expected surface curvature is, the larger Δ gets. Then, the i -th data point support function $\rho_i(p)$ that represents the likelihood of point p to be on a surface with orientation \bar{v}_i (see (10)) can be widened by Δ in the direction orthogonal to \bar{v}_i . Thus, the low curvature prior may justify the use of support functions $\rho_i(p)$ wider than the span of the measurement error model. Figure 7(c) demonstrates the effect of this ad-hoc low-curvature prior.

3.3. Occupancy data

In many cases, there is some information about the scene’s geometry that comes not in the form of surface points but in the form of *spatial occupancy*. It can be defined by a volumetric function $O(p)$ where positive values indicate that point p is likely to be inside the surface and vice versa, while absolute values of $O(p)$ correspond to the certainty. Such occupancy data usually comes from line-of-sight information⁶ or silhouette intersection and may be easily incorporated into functional (11). As a result, in the most general case the optimization problem has the form:

$$\min \leftarrow \int_S \lambda \cdot dS - \int_S \langle \bar{v}, \bar{n}_s \rangle dS - \int_{inter(S)} O(p) \cdot dp. \quad (12)$$

4. Energy optimization

In [19], it was shown that global minima of geometric surface functionals combining area, flux, and volumetric potentials as in (12) can be efficiently computed via a simple graph-cut. A typical graph construction is shown in Figure 8(a). Normally, neighboring pixels are connected with n -links. The cost of n -links severed by a cut represents length (or area in 3D) of the boundary. Pixels are also connected to the terminals S and T via t -links. The cost of t -links represents volumetric potentials. Divergence of a vector field in (9) is an example of volumetric potential that can be represented by weighted t -links either to the source (positive divergence) or to the sink (negative divergence).

4.1. Global optimization via subgraphs

As shown in Figure 9(a), the data (or non-zero t -links) may concentrate in a very small subset of the domain of interest. It may look good enough to consider optimization in a band containing all the data, but this converts graph-cut into a local optimization technique (like level-sets). As was shown in Figure 2, this makes surface reconstruction very

⁶For each range scan, the space between the scan surface and the scanner is likely to be empty (*lines of sight*).

sensitive to initial guess. Yet, in cases like one in Figure 9(a) it seems ridiculous to do global optimization by building a huge 3D grid in the whole domain. The natural question is whether is it possible to guarantee global optimality while doing computations on a small subgraph. The lemma below demonstrates a condition that has such a guarantee. Our optimization algorithm is a natural consequence of this lemma. We compute an optimal cut on a subgraph and then adjust the subgraph until the conditions in the lemma are met.

We will use the following definitions. Consider subgraph (a band) G_{sub} of graph G that contains only a subset of the nodes on G and both terminals S and T . Edges connecting any pair of vertices in G_{sub} are the same as edges connecting them in G . We call two subgraphs *disjoint* if their node sets do not intersect (except for terminal nodes S and T). We also say that two disjoint subgraphs (or simply non-intersecting node subsets) G_1 and G_2 are *connected*, if there is a pair of nodes $n_1 \in G_1$ and $n_2 \in G_2$ connected with an edge in G .

Global optimality lemma: Let G_S , G_T , and B be disjoint subgraphs of G such that their union contains all nodes in G . Assume that:

1. Nodes in G_S do not have (non-zero) t-links to T .
2. Nodes in G_T do not have (non-zero) t-links to S .
3. G_T and G_S are not connected (*i.e.* separated by B).

Consider a minimal cut $[B_S, B_T]$ on B separating its nodes into set B_T connected to sink and set B_S connected to source. If B_S is disconnected from G_T and B_T is disconnected from G_S (see Fig.8b), then $[G_S \cup B_S, G_T \cup B_T]$ is a minimal cut on G (*thus, the minimal cut on G is found directly from the computations done within its subgraph B*).

Proof. The proof uses the standard terminology from network flows theory and the Ford-Fulkerson theorem (minimal cut/maximum flow duality) [12]. We will also use the corollary of the theorem stating that if all edges that are cut by some cut are saturated by some (feasible) flow than the cut is minimal and the flow is maximal.


Consider a cut $[G_S \cup B_S, G_T \cup B_T]$ in graph G and a cut $[B_S, B_T]$ in subgraph B . Due to adjacency conditions, they cut the same set of edges, which we denote C . Let us prove, that the cut $[G_S \cup B_S, G_T \cup B_T]$ is minimal.

Consider the maximal flow F in B . Due to the Ford-Fulkerson theorem this flow saturates all edges in C . Also, this flow is feasible in G (since B is a subgraph in G).

Therefore, for cut $[G_S \cup B_S, G_T \cup B_T]$ in G all severed edges are saturated by a feasible flow. The aforementioned corollary implies that the cut is minimal. \square


Our lemma suggest an efficient algorithm that can take advantage of the banded structure of the data (t-links) without loosing the global optimality of the cut.

Algorithm: assume that we are given a network graph G with some initial segmentation/partitioning into sets G_S ,



(a) Grid graph G (b) Partitioning of nodes

Figure 8. A typical grid graph in computer vision (a). Pixels correspond to nodes. See lemma formulation for other details. Global optimality lemma assumes partitioning of graph nodes into sets G_T , G_S , and $B = B_S \cup B_T$ (b). Note that G_T does not contain t-links to source and G_S does not contain t-links to sink.



(a) Divergence of $\{\bar{v}_p\}$ (b) Global minima in \mathcal{R}^3

Figure 9. Our algorithm can take advantage of narrow concentration of data points. Image (a) shows divergence of a vector field representing "Bunny" data. While our algorithms guarantees global optima in the whole domain, it can often find it while working on a relatively small subgraph B in (b).

G_T , and subgraph(narrow band) B . Alternatively, if our graph is embedded in 3D space and the initialization is given by some shape S_0 , one may create G_T from the nodes exterior to S_0 , G_S from the nodes interior to S_0 , and B from the nodes near the shape surface. The following algorithm computes the globally minimal cut on G using such initialization. (The final result is independent of it!)

On the first step, we add to B all nodes in G_S with non-zero t-links to the sink and all nodes in G_T with non-zero t-links to the source. Such subgraphs G_S , G_T , and B meet the first three conditions of the lemma. We compute a minimum s/t cut in B . If the min-cut in B satisfied the last condition of the lemma, we have found the globally minimal cut and, hence, stop. If obtained components B_S and B_T have nodes connected to the "wrong" inactive component (G_T for B_S and G_S for B_T) then the subgraph B is dilated near these nodes. Then we iterate until a global minima is found. Since B grows monotonically at each step, the process converges.

The number of iterations may be large (*e.g.* dozens). It is important, therefore, that the maximum flow computed in

B on the previous iteration is a feasible flow for the next iteration. Thus, each iteration has to do only small updates of the flow to get the maximum flow, which is much faster than computing maximum flow from scratch.

This algorithm always yields the globally minimal cut. However, the size of B and the number of iterations depends on the accuracy of the initialization. This behaviour is different from the local optimization methods, where both result and consumed resources depend on the accuracy of initialization, see Figure 2. Note, although we used level sets to get the results in Figure 2(c), any other local optimization method (inc. snakes, banded graph-cut) would perform similarly.


Hierarchical algorithm: the proposed algorithm can be used within a hierarchical strategy. Indeed, the result computed using downsampled data can be used as a good initialization on the higher level of hierarchy, which is a common approach in computer vision. What is also important, is that the shapes of the subgraph B , which our algorithm creates at different resolutions are also similar.

We assume that we can compute t-link and n-link capacities for the grids at different resolutions. One way to do this is to recall that these capacities are motivated by the continuously-formulated geometric functional. A more general and simple strategy, which we actually use, is to downsample the t-links and n-links by summation, i.e. construct a pyramid of n-links and t-links.

The following algorithm computes the globally minimal cut on the grid of resolution $nx \times ny \times nz$:

1. Compute the shape $S_{1/4}$ as a result of minimal cut at full grid at resolution $nx/4 \times ny/4 \times nz/4$.
2. Compute the minimal cut at the resolution $nx/2 \times ny/2 \times nz/2$, using the shape $S_{1/4}$ as initialization for our narrow band algorithm. Let $G_{S_{1/2}}$, $G_{T_{1/2}}$, and $B_{1/2}$ be the network graph components on convergence.
3. Compute the minimal cut at full resolution $nx \times ny \times nz$, using the “upsampled” versions of $G_{S_{1/2}}$, $G_{T_{1/2}}$, and $B_{1/2}$ as initialization for our narrow band algorithm.

Typically, we observe 10–20 iterations on step 2, and 2–3 iterations on step 3. However, the vast majority of time is spent at single graph-cut computation on the starting iteration of step 3. Thus, such algorithm has a very small computation overhead compared with a single minimal cut computation in a narrow band. In other words, the algorithm acts as follows: **a)** automatically chooses the shape of the narrow band (steps 1,2) **b)** computes the minimal cut in a narrow band (beginning of step 3), and **c)** makes few addi-



(a) Low-res. graph-cut (b) Sub-pixel smoothing

Figure 10. Extracting subvoxel resolution isosurfaces via level-sets (b) from binary segmentation on a grid (a). To emphasize the smoothing effect, we used a low-resolution grid in this example.

tional iterations updating the narrow band and the minimal cut (rest of step 3). Here, a) and c) takes very small amount of time compared to b).

5. Experimental evaluation

This section presents our experimental results based on hierarchical algorithm from the previous section, which employs the popular min-cut/max-flow algorithm from [3] as a subroutine. In general, our global optimization algorithm produced watertight meshes of low genus. In the cases when our method produced several connected components, we left only the largest, due to the prior knowledge of shape connectivity. To extract isosurfaces from binary segmentation of discrete grids without aliasing artifacts, we used a modification of the method in [27] where *constrained level-sets* smooth a surface on subvoxel level for better 3D rendering, see Fig. 10. Our modification applies such constrained level-sets approach to our own functional (12)⁷.

The numerical details of our algorithm’s performance in the experiments are summarized in table 1. The peak RAM allocation size for all models was 1.5-2.5 GB. Note, that most of these experiments would demand prohibitively large amount of memory being run at the same resolutions with full min-cut/max-flow.

5.1. Fitting to range data

In this subsection, we present the result of our method on the range scan datasets from Stanford 3D Scanning Repository. These well known datasets contain 10-112 registered range scans per dataset. Each range scan was treated simply as a collection of 3D points. To demonstrate the robustness of our method to orientation estimation errors, we used a single orientation vector (corresponding to scan viewing direction) for all points in the scan. We also used coarse line-of-sight information near the legs of Buddha for correct hole

⁷Using our own functional was suggested by Olivier Juan




Figure 11. Our results on the range scan datasets from the Stanford Repository using coarse orientation estimates (one direction per scan).




Figure 12. Comparison of our method with Poisson surface reconstruction algorithm [17]. Both methods use surface orientation information. Two images on the right show that both methods work well when points come with accurate estimates of surface normals (here we used surface normals generated by VRIP reconstruction [10]). Two images on the left show the results obtained when data points had only weak orientation information (direction towards the scanner) clearly demonstrating that our flux-based approach is much less sensitive to accuracy of surface orientation information.

Scene	Grid size	Band size	Time
Bunny	551x544x428(6)	3.72%	61.7
Dragon	602x425x269(6)	9.61%	62.9
Dragon(VRIP)	901x637x403(6)	3.06%	47.5
Armadillo	601x635x501(6)	2.23%	70.7
Buddha	377x914x377(6)	10.75%	1269.9
Temple Full	452x704x332(6)	10.37%	359.5
Temple Ring	451x707x331(6)	11.91%	640.7
Temple Sparse	275x427x202(26)	21.92%	1116.1
Dino Full	252x300x253(26)	20.59%	629.3
Dino Ring	251x300x253(26)	21.31%	1038.2
Dino Sparse	252x300x253(26)	21.68%	667.3

Table 1. Minimal cut stats for the models in subsections 5.1,5.2. Neighborhood system size is shown in brackets. The fourth column contains the percentage of all nodes that were included in the subgraph B on convergence. Time is given in seconds.

filling of a large hole. Otherwise, occupancy data were not used. The results for all datasets are shown in Fig. 11.

We also compared the results of our method with Poisson surface reconstruction recently proposed in [17]. There, it was demonstrated to compare favourably to several other state-of-the-art methods. For a less challenging dataset with accurate normal orientations it produced the results equally good to ours (Figure 12, right two images). However, it was unable to handle original raw scans dataset with coarse orientation estimates, whereas our method produced a reasonable results (Figure 12, left two images).

We also tested the ability of our approach to handle large variations in sampling density. To do that, we removed 98% of points from the right half of Armadillo and then applied our method. With the use of non-uniform Euclidean regularization, our method was able to handle such 50-to-1 difference in sampling density gracefully (Figure 13).




Figure 13. Our results (right) on the pruned Stanford Armadillo range scans data (left). 98% of points were removed from the left part, thus creating 50-to-1 difference in density.

5.2. Fitting to incomplete stereo data

Recently, a number of multiview reconstruction methods have been proposed that produce highly-accurate "quasi-dense" output, leaving holes where the stereo correspondence cannot be established [21, 14, 15]. Our method can post-process such outputs turning them into consistent, watertight meshes, while preserving their high level of details.

Thus, we considered incomplete meshes produced by multiview reconstruction method [14] on Middlebury multiview stereo page datasets [23]. We used the vertices as input points to our algorithms with the normals estimated from local mesh structure. In order to perform the hole filling in ambiguous cases, we used the occupancy information based on the coarse silhouette intersection.

We performed the evaluation on 6 datasets (a group of three for dino model, and another group of three for temple model). Originally, within each group the multiview datasets differed by the number of images. Therefore, the meshes produced by [14] differ by the size of the holes and by the amount of noise and outliers. Also, the sil-

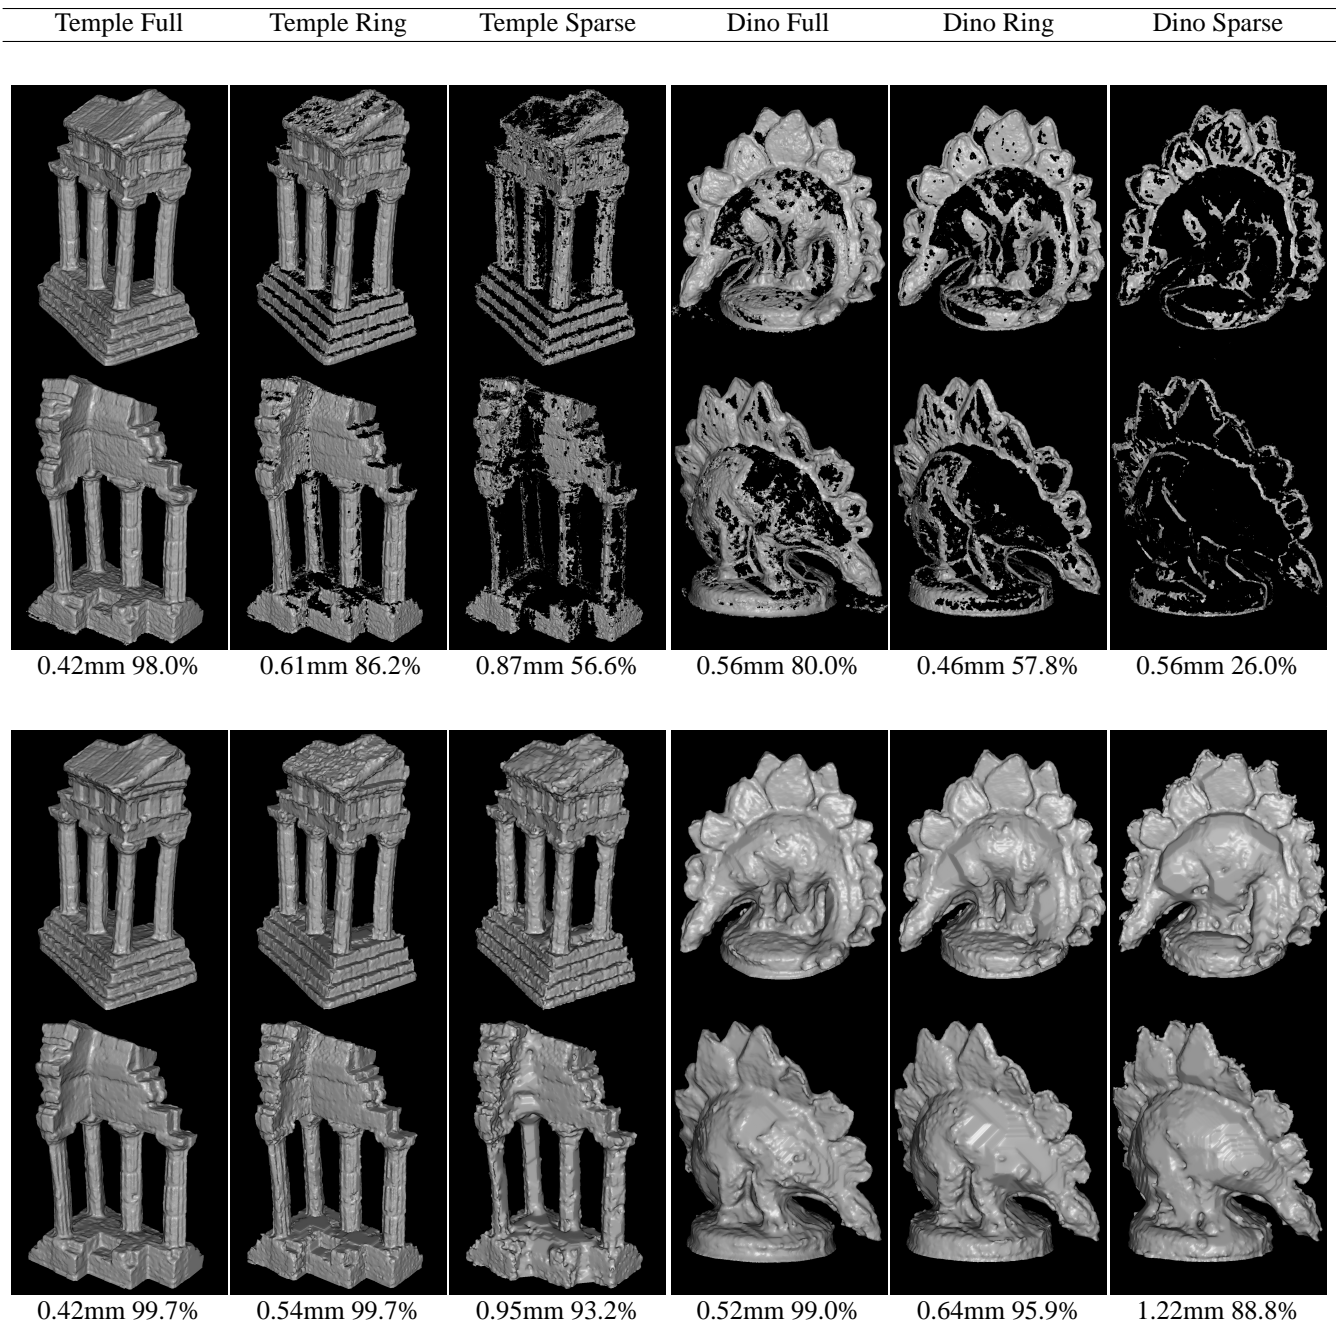


Figure 14. Shape fitting to incomplete meshes on six Middlebury multiview stereo datasets. Top – two views of the results of Goesele *et al* method [14], bottom – the same views for the watertight shapes fitted to those results using our method. First numbers are the accuracy measures (90% of the surface lies within this distance from the ground truth). Second numbers are completeness measures (percentage of the ground truth that lies within 1.25 mm from the surface).

houette intersection results for ‘sparse’ and ‘ring’ datasets are much coarser approximations to real surface than for “full” datasets. The input (results of [14]) and the output of our method, along with statistical measurements of similarity with ground truth are shown on Figure 14. As both visual inspection and statistical measurements verify,

our method significantly improved the completeness of the meshes, while the accuracy of the reconstructions remained on par with the input meshes.

Acknowledgements

Our research benefited from discussions with Alexander Hornung and Leif Kobbelt (RWTH Aachen) who shared their “monkey” dataset as well as their results. Michael Goesele (University of Washington) kindly provided the results of [14] (incomplete Temple and Dino models). Olivier Juan (UWO) suggested reusing the same surface functional at the last stage of our method when we run constrained level-sets to extract our surfaces with subvoxel accuracy. Discussions with Andrew Delong (UWO) and his suggestions helped to clarify many points. John Barron (UWO) helped by sharing his insights on related optical flow problems. We also thank Daniel Scharstein (Middlebury College), Brian Curless (University of Washington), Steven Seitz (University of Washington), and other maintainers of the Middlebury multiview stereo database for creating the datasets as well as for the statistical evaluation of our results and designing a custom-made web-page (<http://vision.middlebury.edu/mview/2006-lempitsky/>).

References

- [1] J. Barron, D. Fleet, and S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision (IJCV)*, 12(1):43–77, 1994. 4
- [2] Y. Boykov and V. Kolmogorov. Computing geodesics and minimal surfaces via graph cuts. In *ICCV’03(I)*, page 26. 1, 2, 4, 5
- [3] Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE TPAMI*, 26(9):1124–1137, 2004. 10
- [4] Y. Boykov, V. Kolmogorov, D. Cremers, and A. Delong. An integral solution to surface evolution PDEs via geo-cuts. In *European Conference on Computer Vision (ECCV)*, volume III, pages 409–422, Graz, Austria, May 2006. 5
- [5] Y. Boykov and V. Lempitsky. From photohulls to photoflux optimization. In *BMVC’06(III)*, pages 1149–1158. 4, 6
- [6] A. Bruhn, J. Weikert, and C. Schnorr. Lukas/Kanade meets Horn/Schunck: Combining local and global optic flow methods. *International Journal of Computer Vision (IJCV)*, 61(3):211–231, 2005. 4
- [7] J. C. Carr, R. K. Beatson, J. B. Cherrie, T. J. Mitchell, W. R. Fright, B. C. McCallum, and T. R. Evans. Reconstruction and representation of 3d objects with radial basis functions. In *SIGGRAPH ’01*, pages 67–76. 4
- [8] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International Journal of Computer Vision*, 22(1):61–79, 1997. 5
- [9] L. Cohen and I. Cohen. Finite element methods for active contour models and balloons from 2-D to 3-D. In *CVPR’92*, pages 592–598. 4
- [10] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *SIGGRAPH’96*, pages 303–312. 4, 5, 12
- [11] Y. Duan, L. Yang, H. Qin, and D. Samaras. Shape reconstruction from 3d and 2d data using pde-based deformable surfaces. In *ECCV’04(III)*, pages 238–251. 4
- [12] L. Ford and D. Fulkerson. *Flows in Networks*. Princeton University Press, 1962. 9
- [13] Y. Furukawa and J. Ponce. Carved visual hulls for image-based modeling. In *European Conference on Computer Vision (ECCV)*, volume I, pages 564–577, Graz, Austria, May 2006. 4
- [14] M. Goesele, B. Curless, and S. Seitz. Multiview stereo revisited. In *CVPR’06(II)*, pages 2402–2409. 12, 13, 14
- [15] M. Habbecke and L. Kobbelt. Iterative multi-view plane fitting. In *VMV*, pages 73–80, 2006. 12
- [16] A. Hornung and L. Kobbelt. Robust reconstruction of watertight 3d models from non-uniformly sampled point clouds without normal information. In *Eurographics Symposium on Geometry Processing’06*, pages 41–50. 2, 3, 4, 5
- [17] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Eurographics Symposium on Geometry Processing’06*, pages 61–70. 4, 12
- [18] R. Kimmel and A. M. Bruckstein. Regularized Laplacian zero crossings as optimal edge integrators. *International Journal of Computer Vision*, 53(3):225–243, 2003. 6
- [19] V. Kolmogorov and Y. Boykov. What metrics can be approximated by geo-cuts, or global optimization of length/area and flux. In *ICCV’05(I)*, pages 564–571. 2, 4, 6, 8
- [20] V. Kolmogorov and C. Rother. Minimizing non-submodular functions with graph cuts - a review. *Microsoft Research, Cambridge, UK, tech. rep. MSR-TR-2006-100*, 2006 (to appear in PAMI in 2007). 5
- [21] M. Lhuillier and L. Quan. Quasi-dense reconstruction from image sequence. In *ECCV*, pages 125–139, 2002. 12
- [22] H. Lombaert, Y. Sun, L. Grady, and C. Xu. A multilevel banded graph cuts method for fast image segmentation. In *ICCV’05*, pages 259–265. 3
- [23] Middlebury Multiview Stereo Page. <http://vision.middlebury.edu/mview/>. 12
- [24] S. Osher and J. Sethian. Fronts propagating with curvature-dependent speed: algorithms based on the Hamilton-Jacobi formulation. *Journal of Computational Physics*, 79:12–49, 1988. 4
- [25] P. Savadjiev, F. P. Ferrie, and K. Siddiqi. Surface recovery from 3d point data using a combined parametric and geometric flow approach. In *EMMCVPR’03*, pp. 325–340. 4, 5, 6
- [26] A. Vasilevskiy and K. Siddiqi. Flux maximizing geometric flows. *PAMI*, 24(12):1565–1578, December 2002. 6
- [27] R. Whitaker. Reducing aliasing artifacts in iso-surfaces of binary volumes. pages 23–32. 10
- [28] R. Whitaker. A level-set approach to 3d reconstruction from range data. *IJCV*, 29(3):203–231, 1998. 4, 5
- [29] H. Zhao, S. Osher, and R. Fedkiw. Fast surface reconstruction using the level set method. In *VLSM ’01*, p. 194. 4, 5