Robust Poisson Surface Reconstruction

V. Estellers[1], M. Scott[2], K. Tew[2], S. Soatto[1]

[1]Univeristy of California, Los Angeles, US.
[2] Brigham Young University, Provo, US.

Abstract. *We propose a method to reconstruct surfaces from oriented point clouds with non-uniform sampling and noise by formulating the problem as a convex minimization that reconstructs the indicator function of the surface's interior. Compared to previous models, our reconstruction is robust to noise and outliers because it substitutes the least-squares fidelity term by a robust Huber penalty; this allows to recover sharp corners and avoids the shrinking bias of least squares. We choose an implicit parametrization to reconstruct surfaces of unknown topology and close large gaps in the point cloud. For an efficient representation, we approximate the implicit function by a hierarchy of locally supported basis elements adapted to the geometry of the surface. Unlike ad-hoc bases over an octree, our hierarchical B-splines from isogeometric analysis locally adapt the mesh and degree of the splines during reconstruction. The hierarchical structure of the basis speeds-up the minimization and efficiently represents clustered data. We also advocate for convex optimization, instead isogeometric finite-element techniques, to efficiently solve the minimization and allow for non-differentiable functionals. Experiments show state-of-the-art performance within a more flexible framework.*

1 Introduction

New challenges to surface reconstruction from measurements emerge as datasets grow in size but lose in accuracy. The reduction in accuracy appears as sensors evolve from short to long range, low-cost commodity scanners become widely available, and computer vision is increasingly used to infer 3D geometry from point sets. As a result, surface reconstruction methods must be robust to noise and outliers, and scale favorably in terms of computation and memory usage. This impacts the parametrization of the surface and the inference techniques.

We propose a robust but simple algorithm to reconstruct a water-tight surface from an oriented point cloud. We formulate the reconstruction as a convex optimization that recovers the indicator function of the interior of the surface. Our objective function penalizes deviations in the normal orientation with a Huber loss function to robustly recover the topology of the surface and allow for sharp corners; this makes our model more robust to noise and avoids the "shrinking bias" of least-squares models [1,2]. Our minimization exploits the convexity of the objective with an efficient first-order algorithm that is easy to parallelize and scales well with the size of the point cloud. This is our first contribution.

Our second contribution is to merge state of the art isogeometric analysis and surface reconstruction. Isogeometric analysis [3,4] is a generalization of finite element analysis which improves the link between geometric design and analysis. The isogeometric paradigm is simple: the smooth spline basis used to define the geometry is used as the basis for analysis. As a result, exact geometry is introduced into the analysis. The smooth basis can then be leveraged in analysis [5,6,7] and lead to innovative approaches to model design [8,9,10], analysis [11,12,13,14], optimization [15], and adaptivity [16,17,18,19,20].

The underlying implicit function is represented by an adaptive spline forest [18] developed in isogeometric analysis [3]. An isogeometric spline forest is a hierarchical spline representation capable of representing surfaces or volumes of arbitrarily complex geometry and topological genus. Spline forests can accommodate arbitrary degree and smoothness in the underlying hierarchical basis as well as non-uniform knot interval configurations. They accommodate efficient h, p, k-refinement and coarsening algorithms which we utilize in this work. In h-adaptivity the elements are subdivided or merged, in p-adaptivity the polynomial degree of the basis is changed, and in k-adaptivity the smoothness of the basis is changed. In all cases, the adaptive process remains local and preserves exact geometry at the coarsest level of the discretization. In the context of surface reconstruction, a hierarchical spline forest basis efficiently represents functions in three dimensions by their spline coefficients and provides analytic expressions for their derivatives. Our reconstruction exploits this local adaptivity to efficiently represent complex surfaces with sharp corners.

2 Related Methods and Choice of Representation

Surface reconstruction methods can be first classified by their surface representation: parametric or implicit. Parametric techniques represent the surface as a topological embedding of a 2D parameter domain into 3D space. Among them, approaches based on computational geometry partition the space into Voronoi cells from the input samples and exploit the intuitive idea that eliminating facets of Delaunay tetrahedra provides a triangulated parametrization of the surface [21,22,23,24,25,26,27]. The reconstructed surface thus interpolates most of the input samples and requires post-processing to smooth the surface and correct the topology. Parametric methods generally require clean data because they assume the topology of the surface to be known, while implicit methods are designed to reconstruct surfaces from noisy point clouds with unknown topology.

Implicit representations both reconstruct the surface and estimate its topology, but increase the dimension of the problem by representing the surface as the zero-level set of a volumetric function. Their accuracy is thus limited by the resolution of the grid, with efficient representations requiring non-uniform grids.

Implicit representations can be formulated as either global or local. Local methods consider subsets of nearby points one at a time and handle large datasets efficiently. Earlier methods [28,29] estimate tangent planes from the nearest neighbors of each sample and parametrize the surface by the signed dis-

tance to the tangent plane of the closest point in space. Moving least squares (MLS) techniques [30,31,32,33] reconstruct surfaces locally by solving an optimization that finds a local reference plane and fit a polynomial to the surface. The least-squares fit of MLS, however, is sensitive to outliers and smooths out small features; for this reason variants robust to outliers [34,35] and sharp features [36,37] appeared. [38] also constructs implicit functions locally but blends them together with partitions of unity. Common to these methods is their locality –partitioning into neighborhoods and merging local functions– that makes them highly scalable but sensitive to non-uniform sampling and point-cloud gaps.

Global methods define the implicit function as the sum of basis functions (RBFs[39], splines[1,2], wavelets [40]) and consider all the data at once without heuristic partitioning. Kazhdan et al. [1] solve a Poisson problem that aligns the gradient of the indicator function to the normals of the point cloud with a least-squares fit, not robust to outliers. Manson et al. [40] similarly approximate the indicator function with wavelets efficiently designed to compute basis coefficients with local sums over an octree. Calakli and Taubin [41] use instead a signed-distance function to represent the surface, but also rely on least squares to fit the normals and include screening and regularization terms. While the Hessian regularization introduces derivatives of higher order and limits the basis functions, the screening improves accuracy by fitting the input points to the zero-level set of the implicit function and has been lately adopted by [2]. For this reason, our model includes a screening term together with a robust Huber penalty to fit the normal field and allow for sharp edges. Existing methods account for sharp features by explicit representations [42,43,36] or anisotropic smoothing [44,45,46]; they are fast but depend on local operators that do not seek a global optimum.

Our reconstruction combines benefits of global and local schemes. It is global and does not involve heuristics on neighborhoods, while the basis functions are locally supported and adapt to the surface through the local refinement techniques of the spline forests [18]; this can be viewed as a generalization of well-known uniform splines [47,48,49]. Our minimization algorithm, however, departs from standard isogeometric finite element formulations and is instead inspired by [50].

3 Variational Model

The reconstruction of a surface S from oriented points can be cast as a minimization problem to estimate the indicator function χ of the interior of the surface. Let $\{x_k, \boldsymbol{n}_k\}_{k=1}^N$ be the oriented point cloud, with $x_k \in \mathbb{R}^3$ the point location and $\boldsymbol{n}_k \in \mathbb{S}^2$ its associated normal; we estimate $\chi\colon \mathbb{R}^3 \to \mathbb{R}$ such that $S = \{x\colon \chi(x) = 0\}$, $\chi < 0$ in the interior enclosed by S and $\chi > 0$ outside.

We reconstruct S by observing that each point in $\{x_k, \boldsymbol{n}_k\}_{k=1}^N$ is a sample of the gradient of the indicator function, that is, $\nabla\chi(x_k) = \boldsymbol{n}_k$. Given a continuous field $\boldsymbol{n}$ that approximately interpolates these samples[1], $\boldsymbol{n}(x_k) = \boldsymbol{n}_k$, we can reconstruct S by finding the scalar function whose gradient best matches this

[1] We approximate $\boldsymbol{n}$ with a standard ℓ_2 projection into a linear B-spline basis defined over the same grid as χ. We omit here the details to focus on reconstruction.

field. To account for noise in the data, we formulate the reconstruction as a minimization, instead of interpolation, problem:

$$\min_{\chi} \int_{\mathbb{R}^3} f(\boldsymbol{n} - \nabla\chi) + \frac{\alpha}{2} \sum_{k=1}^{N} \chi(x_k)^2 \quad \text{with} \quad f(\boldsymbol{v}) = \begin{cases} \frac{1}{2}|\boldsymbol{v}|_2^2 & |\boldsymbol{v}|_2 < \epsilon \\ \epsilon(|\boldsymbol{v}|_2 - \frac{\epsilon}{2}) & |\boldsymbol{v}|_2 \geq \epsilon \end{cases}, \quad (1)$$

where $\alpha > 0$ is a model parameter. The Huber loss function f is a convex and differentiable penalty that avoids two artifacts of least squares: shrinkage of thin structures and smoothing of sharp edges as the square norm over-penalizes outliers. It overcomes these limitations by using different penalties for small errors and outliers, but results in a minimization harder to solve than the Poisson problem of a least-squares fit. The second term in (1) sets the points as soft interpolation constraints for the zero-level set of χ and is a generalization of the screening term of [51], but defined over a sparse set of points rather than $\mathbb{R}^3$.

Since χ only contains surface information in its zero-level set, we approximate it with an adaptive spline forest basis defined over a non-uniform grid. Each basis function is a (smooth) piecewise polynomial inferred from the local knot structure of the underlying spline forest. For a complete description of spline forest basis functions we refer the reader to [18]. In one dimension, a spline forest is a hierarchical B-spline as illustrated in Figure 1, where the basis functions that span the hierarchical spline space are indicated by solid lines. The multi-level adaptive nature of the basis is evident. Similarly, in higher dimensions the first level of the hierarchy is as a NURBS or T-spline.

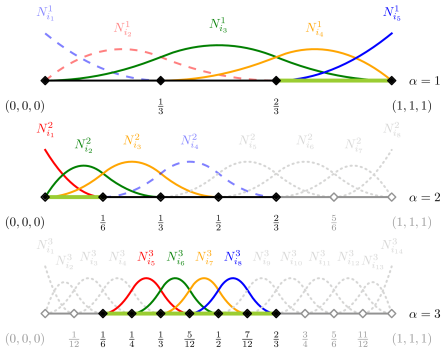

Fig. 1: Basis functions for a three-level quadratic hierarchical spline space, where N_i^l denotes the i-th spline function at level l. Basis functions are indicated by solid colored lines, functions that are linearly dependent on higher-level functions by dashed colored lines, and inactive functions by grey dotted lines. The finite elements correspond to the cells in green.

Our method is related to the Poisson reconstruction of [2], but our representation of the surface, model, and minimization technique are different. In terms

of the model, we propose a robust Huber penalty on the normals to be resilient to outliers, instead of the least-squares penalty of [2]. We represent χ with spline forests from isogeometric analysis – instead of uniform splines over an octree not defining a basis – to use h- and p-refinement to adapt our representation to the complexity of the surface, not to the noisy point cloud. Finally, our minimization exploits the convexity of (1) to develop an efficient primal-dual algorithm, instead of finite-element methods that cannot handle the Huber loss function.

4 Minimization Algorithm

To exploit the convexity of the functional in the minimization, we discretize the integral in (1) with standard quadrature rules used with non-uniform splines:

$$\int_{\mathbb{R}^3} f(\nabla\chi - \boldsymbol{n}) \approx \sum_{i=1}^{Q} w_i f(\nabla\chi(p_i) - \boldsymbol{n}(p_i)), \tag{2}$$

where $\{p_i, w_i\}_{i=1}^{Q}$ are the quadrature points and their weights. Finding the location and value of $\{p_i, w_i\}_{i=1}^{Q}$ that best approximates the integral is a classical problem with exact solution for polynomials. By restricting χ to the span of the hierarchical basis $\{N_A\}_{A=1}^{n}$, we confine the minimization to coefficients $c \in \mathbb{R}^n$:

$$\min_c \sum_{i=1}^{Q} w_i f(\sum_A c_A \nabla N_A(p_i) - \boldsymbol{n}(p_i)) + \frac{\alpha}{2} \sum_{k=1}^{N} [\sum_A c_A N_A(x_k)]^2. \tag{3}$$

We solve (3) efficiently with the primal-dual algorithm [52] by providing closed-form solutions for each proximal update. We choose a first-order method because the size of the problem makes second-order methods unfeasible. The convexity of each term in the objective allows us to re-formulate the minimization as a saddle-point problem that is separable and easy to solve in each variable.

Let $P \in \mathbb{R}^{3Q \times n}$ be the matrix with block components $P_{ij} = \nabla N_j(p_i)$ and

$$F(V) = \sum_{i=1}^{Q} w_i f(V_i - \boldsymbol{n}(p_i)), \quad G(c) = \frac{\alpha}{2} \sum_{k=1}^{N} [\sum_A c_A N_A(x_k)]^2,$$

we can write (3) as the constrained minimization

$$\min_{c,V} F(V) + G(c) \;\text{ s.t. }\; V = Pc \tag{4}$$

and use convex analysis to find the equivalent saddle-point problem with dual variable $\lambda \in \mathbb{R}^{3Q}$, F^* the convex conjugate of F, and the inner product $\langle \cdot, \cdot \rangle$:

$$\max_{\lambda} \min_{c} \; -F^*(\lambda) + G(c) + \langle \lambda, Pc \rangle. \tag{5}$$

This formulation allows us to apply the primal-dual algorithm of Chambolle and Pock [52], which solves (5) by iteratively solving the following sub-problems:

$$\lambda^{n+1} \leftarrow \min_{\lambda} \ \sigma F^*(\lambda) + \frac{1}{2}\|\lambda - \lambda^n - \sigma P\bar{c}^n\|^2 \tag{6}$$

$$c^{n+1} \leftarrow \min_{c} \ \tau G(c) + \frac{1}{2}\|c - c^n + \tau P^*\lambda^{n+1}\|^2 \tag{7}$$

where $\bar{c}^{n+1} = c^{n+1} + \theta(c^{n+1} - c^n)$, P^* is the adjoint of P, and τ, σ, θ are algorithm parameters. As G is strongly convex, our algorithm can be further accelerated by updating parameters τ, σ, θ according to Alg. 2 in [52].

The efficiency of the proposed algorithm comes from the spatial separability of F^* and from the ability to find closed-form solutions for each of the minimization problems. The derivation of closed-form solutions is detailed next, and Algorithm 1 summarizes the resulting updates, which are easy to parallelize.

Initialize $c = 0$, $\boldsymbol{\lambda} = \mathbf{0}$, $\bar{c} = c$. Choose $\tau, \sigma > 0$, $\theta \in [0, 1]$.
while $\|c^{n+1} - c^n\| > 1^{-4}$ **do**
$\lambda_i^{n+1} = \epsilon w_i \frac{\hat{\lambda}_i^n - \sigma n(p_i)}{\max(\ \epsilon(w_i+\sigma), |\hat{\lambda}_i^n - \sigma n(p_i)|_2\ )}$ with $\hat{\lambda}_i^n = \lambda_i^n + \sigma \sum_A \bar{c}_A^n \nabla N_A(p_i)$
$c^{n+1} = [I_n + \alpha\tau M]^{-1}\hat{c}$ with $\hat{c}_A = c_A^n - \tau \sum_{i=1}^Q \nabla N_A(p_i) \cdot \lambda^{n+1}$
$\bar{c}^{n+1} = c^{n+1} + \theta(c^{n+1} - c^n)$
end

Algorithm 1: Primal-dual minimization algorithm.

Minimization in dual variable: Let $\hat{\lambda} = \lambda + \sigma P\bar{c}$, we solve the minimization in λ (6) through Moreau's identity [53]:

$$\lambda \leftarrow \min_{\lambda} \ \sigma F^*(\lambda) + \frac{1}{2}\|\lambda - \hat{\lambda}\|^2 \iff \lambda = \hat{\lambda} - \sigma V^*, \ V^* \leftarrow \min_{V} \ F(V) + \frac{\sigma}{2}\|V - \frac{\hat{\lambda}}{\sigma}\|^2.$$

The minimization in V is decoupled spatially in each quadrature point as follows:

$$\min_V \sum_{i=1}^Q \underbrace{w_i f(V_i - \boldsymbol{n}(p_i)) + 0.5\sigma(V_i - \sigma^{-1}\hat{\lambda}_i)^2}_{h_i(V_i)} = \min_{V_1,\dots,V_Q} \sum_{i=1}^Q h_i(V_i). \tag{8}$$

It is thus solved by independently minimizing $h_i(V_i)$ in each each $V_i \in \mathbb{R}^3$. Due to the convexity and differentiability of the Huber norm, the optimality conditions are obtained by differentiating the objective function with respect to V_i; i.e.,

$$\epsilon w_i \frac{V_i^* - \boldsymbol{n}(p_i)}{\max(\epsilon, |V_i^* - \boldsymbol{n}(p_i)|_2)} + \sigma V_i^* - \hat{\lambda}_i = 0. \tag{9}$$

After some algebra, we obtain the closed-form update for the dual variable,

$$\lambda_i = \epsilon w_i \frac{\hat{\lambda}_i - \sigma \boldsymbol{n}(p_i)}{\max(\ \epsilon(w_i + \sigma), |\hat{\lambda}_i - \sigma \boldsymbol{n}(p_i)|_2\ )}. \tag{10}$$

Minimization in primal variable Let $\hat{c} = c^n - \tau P^* \lambda^{n+1}$, the minimization in c (7) is the least-squares problem

$$\min_c \ \alpha\tau \sum_{k=1}^{N} [\sum_A c_A N_A(x_k)]^2 + \sum_A (c_A - \hat{c}_A)^2, \tag{11}$$

whose optimality conditions lead to the sparse linear system: $c + \alpha\tau M c = \hat{c}$. Indeed, differentiating with respect to c_B and equating to 0 gives

$$\alpha\tau \sum_A [\ \sum_{k=1}^{N} N_A(x_k) N_B(x_k)\ ] c_A + c_B = \hat{c}_B \iff c_B + \alpha\tau \sum_A M_{AB} c_B = \hat{c}_B,$$

with coefficients $M_{AB} = \sum_k N_A(x_k) N_B(x_k)$, $\hat{c}_A = c_A^n - \tau \sum_{i=1}^{Q} \nabla N_A(p_i) \cdot \lambda^{n+1}$. As M is fixed, we only decompose it once for all the primal-dual iterations.

5 Refinement

We start our reconstruction with first-order splines over a coarse uniform grid and refine the basis in two stages: we first refine the resolution of the basis by partitioning cells with large values of the objective function or interpolation error, and we then increase the degree of basis elements over cells where first order splines have large objective values. Refinement is illustrated in Figure 3(d).

5.1 h-Refinement: minimizing point-cloud and surface errors

For every cell Ω_i in our grid, we compute the error between the continuous approximation to the normal field and the input normal at the sampled points, $e(\Omega_i) = \sum_{x_k \in \Omega_i} |\boldsymbol{n}(x_k) - \boldsymbol{n}_k|^2$, and refine the cells whose error is larger than the average mean error over the grid. The size of the cells is thus determined by the input point cloud, not by the reconstructed surface, and is sensitive to noise. This first refinement assures an accurate approximation of the normal field $\boldsymbol{n}$ from the samples and is used to obtain a first approximate surface.

Given this first estimate of χ, we refine the mesh based on the error in the reconstructed surface. To be robust to outliers, we measure this error with the Huber penalty, that is, $e_s(\Omega_i) = \sum_{x_k \in \Omega_i} f(\nabla\chi(x_i) - \boldsymbol{n}_i)$, and refine cells whose error is larger than the mean error over the grid. Compared to h-refinement based on the interpolation error, this new criterion is robust to noise in the point cloud and implicitly measures the complexity of the surface's geometry over each cell.

5.2 p-Refinement: Degree Elevation

Elevating the degree of the spline basis allows to better approximate smooth regions and represent higher-order geometric properties of the surface. Higher-order splines, however, lead to computationally more demanding models that smooth out corners; for this reason it is critical to only increase the degree of the spline basis over cells where the surface is smooth. We thus increase the degree of the spline basis over cells already refined, but with large residual objective.

6 Experimental Results

We perform experiments with two kinds of data: synthetic data with ground truth, and point clouds obtained from structured-light scanning with a Kinect camera [54]. Kinect data suffers from non-uniform noise, large scanning gaps and artifacts. We use the synthetic point clouds for quantitative evaluation and the noisy ones to test reconstruction with data with real noise and artifacts.

Our first experiment does not include refinement to focus on our first contribution, the use of a robust penalty for the normal field. Figure 2 compares our model to a least-squares fit equivalent to the Poisson reconstruction of [2] without the octree structure. Our model reconstructs a cube with sharp corners and avoids shrinkage artifacts in the horns and legs of the cow. In all our experiments, we hand-picked the best parameters for each model.

Our second experiment investigates the effects of refinement in the accuracy of the reconstruction, see Figure 3. We again compare our model to the least-squares model of [2] –where refinement is defined by point density through the octree– implemented with our spline forest to focus on the refinement strategies. For synthetic data both models perform similarly, showing how a correctly refined basis compensates for the use of a non-robust penalty. For noisy data, however, our method reconstructs more accurately the topology of the cleaning spray of Figure 4, while refinement based on point density replicates the artifacts in the noisy point cloud. The effects of p-refinement, which is only possible in our isogeometric framework, are analyzed in Figure 3(c). This kind of refinement allows us to obtain both smooth regions and sharp edges within a single surface.

A third set of experiments shows that model and implementation lead to reconstructions comparable to state-of-the-art techniques [1,2,41]. Table 1 presents a quantitative comparison of the methods with synthetic or high quality point clouds that have been perturbed, by sub-sampling or adding noise to the point locations, and evaluates the Haussdorf distance to the reconstructed surface. All the methods lead to reconstructions with similar mean Haussdorf distance, but the average reconstruction time of our algorithm – with a resolution comparable to the octree-based methods – is an order of magnitude larger than [1,2,41] with the current non-optimized code. Finally, Figure 5 shows a qualitative comparison with a noisy point cloud scanned with Kinect: our method is better at recovering shape and topology from noisy data, e.g., the holding hands of the statue.

7 Conclusions and Future Work

We reconstruct surfaces from corrupted point clouds by formulating the problem as a convex minimization that is robust to outliers, avoids the shrinking bias of least squares, and is able to recover sharp corners as well as smooth regions. For an efficient parametrization, we approximate the implicit function with a hierarchical B-spline forest that locally adapts its resolution and smoothness to the surface. This allows us to dynamically refine the reconstruction guided by the surface, instead of the input point cloud, and be robust to noise and artifacts.

(a) Least-squares (b) Our model (c) Least-squares (d) Our model

Fig. 2: Reconstruction in uniform grid: Poisson reconstruction rounds the cube's corners and shrinks the cows horns because least-squares fits overpenalize outliers, our robust model overcomes avoids this limitation with a Huber penalty.

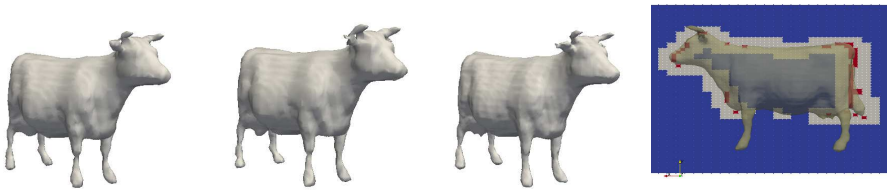

(a) Least-squares (b) Our model (c) p-ref, ours (d) Hierarchy

Fig. 3: Refinement: 3(a) h-refinement of least-squares model 2(c); 3(b) h-refinement of our model 2(d); 3(c) p-refinement of 3(b). The Hierarchical mesh 3(d) with first-level elements in blue, second level in gray, and third level in red.

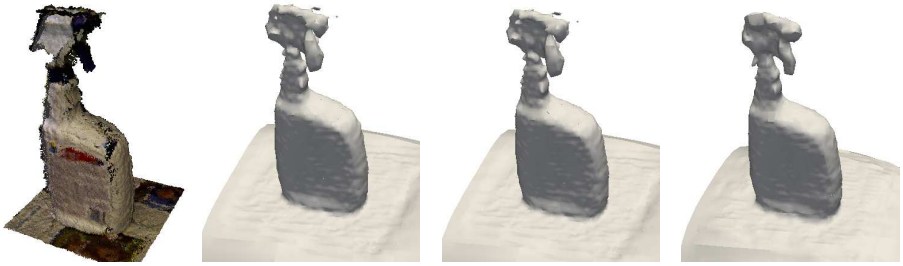

(a) Point cloud (b) Point density (c) Interpolation err. (d) Reconstruct. err.

Fig. 4: Reconstruction with different refinement criteria: refinement based on point density, (4(b)) like [1,2,41,40], or on the interpolation error (4(c)) reconstruct the artifacts present in the point cloud, while refinement based on the reconstruction error recover the right shape in the cleaner's head. (4(d)).

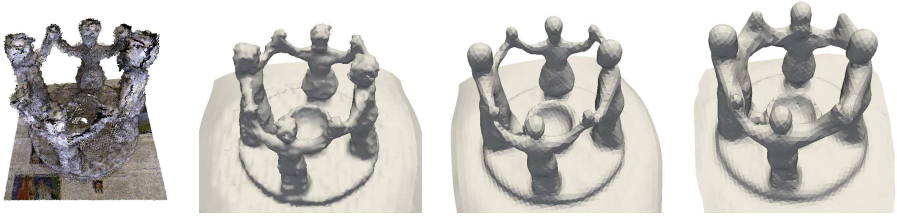

(a) Point cloud (b) Ours (c) Method [2] (d) Method [41]

Fig. 5: Comparison of our model to state-of-the-art methods [2,41]. Our model better recovers the shape of the faces and hands of the statues: avoids disconnecting the joined hands ([2]) or their merging with the person's arms ([41]).

Table 1: Average reconstruction time and Haussdorf distance $d_{\mathcal{D}}$ (10^{-2} distance units) between the point cloud and the reconstructed surface.

$d_{\mathcal{D}}$	subsampled point cloud				perturbed point-cloud				time (s)
	bunny	cow	horse	cube	bunny	cow	horse	cube	
[1]	0.074	4.85	0.237	6.25	0.086	3.04	0.232	4.65	2.7
[41]	0.065	1.72	0.107	1.16	0.069	1.80	0.108	2.37	1.2
[2]	0.073	4.84	0.232	3.94	0.337	3.07	0.249	4.58	2.8
ours	0.0708	2.92	0.143	2.37	0.0768	3.38	0.137	3.05	476

Hierarchical spline forests and isogeometric analysis are powerful tools for vision and imaging and open a new line of future work, from more efficient implementations that exploit the capabilities of GPU to new isogeometric models designed for the noisy data and unknown geometries common in computer vision.

Acknowledgments. This research is supported by the Swiss National Science Foundation grant P2ELP2_148890, Air Force Office of Scientific Research FA9550-14-1-0113, FA9550-12-1-0364:P00002 and N00014-11-1-0863:P00010.

References

1. M. Kazhdan, M. Bolitho, and H. Hoppe, "Poisson surface reconstruction," *Eurographics Symposium on Geometry Processing*, pp. 61–70, 2006.
2. M. Kazhdan and H. Hoppe, "Screened poisson surface reconstruction," *ACM Trans Graph*, vol. 32, no. 3, pp. 1–13, Jun. 2013.
3. T. Hughes, J. Cottrell, and Y. Bazilevs, "Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement," *Comput.Methods in Appl.Mech.Eng.*, vol. 194, no. 39-41, pp. 4135–4195, Oct. 2005.
4. J. A. Cottrell, T. J. R. Hughes, and Y. Bazilevs, *Isogeometric analysis: Toward Integration of CAD and FEA*. Chichester: Wiley, 2009.
5. J. A. Evans, Y. Bazilevs, I. Babuška, and T. J. R. Hughes, "n-widths, sup-infs, and optimality ratios for the k-version of the isogeometric finite element method," *Comput.Methods in Appl.Mech.Eng.*, vol. 198, no. 21-26, pp. 1726–1741, 2009.
6. T. J. R. Hughes, J. A. Evans, and A. Reali, "Finite element and NURBS approximations of eigenvalue, boundary-value, and initial-value problems," *Comput.Methods in Appl.Mech.Eng.*, vol. 272, pp. 290 – 320, 2014.
7. J. A. Cottrell, T. J. R. Hughes, and A. Reali, "Studies of refinement and continuity in isogeometric analysis," *Comput.Methods in Appl.Mech.Eng.*, vol. 196, pp. 4160–4183, 2007.
8. E. Cohen, T. Martin, R. M. Kirby, T. Lyche, and R. F. Riesenfeld, "Analysis-aware modeling: Understanding quality considerations in modeling for isogeometric analysis," *Comput.Methods in Appl.Mech.Eng.*, vol. 199, no. 5-8, pp. 334–356, 2010.
9. W. Wang, Y. Zhang, M. A. Scott, and T. J. R. Hughes, "Converting an unstructured quadrilateral mesh to a standard T-spline surface," *Computational Mechanics,*, vol. 48, pp. 477 – 498, 2011.
10. L. Liu, Y. Zhang, T. J. R. Hughes, M. A. Scott, and T. W. Sederberg, "Volumetric T-spline construction using boolean operations," in *Proceedings of the 22nd International Meshing Roundtable*, J. Sarrate and M. Staten, Eds. Springer, 2014, pp. 405–424.

11. D. Schillinger, L. Dedé, M. A. Scott, J. A. Evans, M. J. Borden, E. Rank, and T. J. R. Hughes, "An isogeometric design-through-analysis methodology based on adaptive hierarchical refinement of NURBS, immersed boundary methods, and T-spline CAD surfaces," *Comput.Methods in Appl.Mech.Eng.*, vol. 249 – 252, pp. 116 – 150, 2012.
12. M. A. Scott, R. N. Simpson, J. A. Evans, S. Lipton, S. P. A. Bordas, T. J. R. Hughes, and T. W. Sederberg, "Isogeometric boundary element analysis using unstructured T-splines," *Comput.Methods in Appl.Mech.Eng.*, vol. 254, pp. 197 – 221, 2013.
13. R. Schmidt, R. Wüchner, and K.-U. Bletzinger, "Isogeometric analysis of trimmed NURBS geometries," *Comput.Methods in Appl.Mech.Eng.*, vol. 241–244, pp. 93 – 111, 2012.
14. D. J. Benson, Y. Bazilevs, E. De Luycker, M. C. Hsu, M. A. Scott, T. J. R. Hughes, and T. Belytschko, "A generalized finite element formulation for arbitrary basis functions: From isogeometric analysis to XFEM," *International Journal for Numerical Methods in Engineering*, vol. 83, pp. 765–785, 2010.
15. W. A. Wall, M. A. Frenzel, and C. Cyron, "Isogeometric structural shape optimization," *Comput.Methods in Appl.Mech.Eng.*, vol. 197, pp. 2976–2988, 2008.
16. Y. Bazilevs, V. M. Calo, J. A. Cottrell, J. A. Evans, T. J. R. Hughes, S. Lipton, M. A. Scott, and T. W. Sederberg, "Isogeometric analysis using T-splines," *Comput.Methods in Appl.Mech.Eng.*, vol. 199, no. 5-8, pp. 229–263, 2010.
17. M. Dörfel, B. Jüttler, and B. Simeon, "Adaptive isogeometric analysis by local h-refinement with T-splines," *Comput.Methods in Appl.Mech.Eng.*, vol. 199, no. 5–8, pp. 264–275, 2009.
18. M. A. Scott, D. C. Thomas, and E. J. Evans, "Isogeometric spline forests," *Comput.Methods in Appl.Mech.Eng.*, vol. 269, pp. 222–264, 2014.
19. E. J. Evans, M. A. Scott, X. Li, and D. C. Thomas, "Hierarchical T-splines: Analysis-suitability, Bézier extraction, and application as an adaptive basis for isogeometric analysis," *Comput.Methods in Appl.Mech.Eng.*, vol. 284, pp. 1–20, 2015.
20. D. C. Thomas, M. A. Scott, J. A. Evans, K. Tew, and E. J. Evans, "Bézier projection: a unified approach for local projection and quadrature-free refinement and coarsening of NURBS and T-splines with particular application to isogeometric design and analysis," *submitted*, 2014.
21. N. Amenta, M. Bern, and M. Kamvysselis, "A New Voronoi-Based Surface Reconstruction Algorithm," *ACM SIGGRAPH*, pp. 415–421, 1998.
22. T. Dey and S. Goswami, "Tight cocone: a water-tight surface reconstructor," *ACM symposium on Solid modeling and applications*, pp. 127–134, 2003.
23. N. Amenta, S. Choi, and R. K. Kolluri, "The power crust," *ACM symposium on Solid Modeling and Applications*, pp. 249–266, 2001.
24. J. Podolak and S. Rusinkiewicz, "Atomic Volumes for Mesh Completion," *Symposium on Geometry Processing*, pp. 33–41, 2005.
25. J.-D. Boissonnat and S. Oudot, "Provably good sampling and meshing of surfaces," *Graphical Models*, vol. 67, no. 5, pp. 405–451, 2005.
26. J. R. Shewchuk and J. F. O. Brien, "Spectral Surface Reconstruction from Noisy Point Clouds," *ACM SIGGRAPH*, vol. 14, pp. 11–21, 2004.
27. P. Labatut, J.-P. Pons, and R. Keriven, "Robust and efficient surface reconstruction from range data," *Computer Graphics Forum*, vol. 28, no. 8, pp. 2275–2290, 2009.
28. H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle, "Surface reconstruction from unorganized points," *ACM SIGGRAPH*, pp. 71–78, 1992.
29. B. Curless and M. Levoy, "A Volumetric Method for Building Complex Models from Range Images Volumetric integration," *ACM SIGGRAPH*, pp. 303–312, 1996.
30. M. Alexa, J. Behr, D. Cohen-or, S. Fleishman, D. Levin, and C. T. Silva, "Computing and rendering point set surfaces," *IEEE Transactions on Visualization and Computer Graphics*, vol. 9, no. 1, pp. 3–15, Jan. 2003.

31. D. Levin, "Mesh-Independent surface Interpolation," *Geometric modeling for scientific visualization*, pp. 37–49, 2004.
32. C. Shen, J. F. O'Brien, and J. R. Shewchuk, "Interpolating and approximating implicit surfaces from polygon soup," *ACM SIGGRAPH*, vol. 23, no. 3, pp. 896–904, 2004.
33. N. Amenta and Y. J. Kil, "Defining point-set surfaces," *ACM SIGGRAPH*, vol. 23, no. 3, pp. 264–270, 2004.
34. A. C. Öztireli, G. Guennebaud, and M. Gross, "Feature Preserving Point Set Surfaces based on Non-Linear Kernel Regression," *Computer Graphics Forum*, vol. 28, no. 2, pp. 493–501, Apr. 2009.
35. S. Fleishman, D. Cohen-or, and C. T. Silva, "Robust moving least-squares fitting with sharp features," *ACM SIGGRAPH*, vol. 44, no. 3, pp. 544–552, 2005.
36. Y. Lipman, D. Cohen-or, and D. Levin, "Data-dependent MLS for faithful surface approximation," *Eurographics Symposium on Geometry Processing*, pp. 59–67, 2007.
37. J. Daniels II, L. K. Ha, T. Ochotta, and C. T. Silva, "Robust Smooth Feature Extraction from Point Clouds," *IEEE International Conference on Shape Modeling and Applications (SMI'07)*, pp. 123 – 136, 2007.
38. Y. Ohtake, A. Belyaev, M. Alexa, G. Turk, and H.-P. Seidel, "Multi-level partition of unity implicits," *ACM Trans Graph*, vol. 22, no. 3, p. 463, Jul. 2003.
39. J. C. Carr, R. K. Beatson, and T. R. Evans, "Reconstruction and Representation of 3D Objects with Radial Basis Functions," *ACM SIGGRAPH*, pp. 67–76, 2001.
40. J. Manson, G. Petrova, and S. Schaefer, "Streaming surface reconstruction using wavelets," *Computer Graphics Forum*, vol. 27, no. 5, pp. 1411–1420, 2008.
41. F. Calakli and G. Taubin, "SSD: Smooth Signed Distance Surface Reconstruction," *Computer Graphics Forum*, vol. 30, no. 7, pp. 1993–2002, Sep. 2011.
42. A. Adamson and M. Alexa, "Point-sampled cell complexes," *ACM Trans Graph*, vol. 1, no. 212, pp. 671–680, 2006.
43. G. Guennebaud and M. Gross, "Algebraic Point Set Surfaces," *ACM Trans Graph*, vol. 26, no. 3, pp. 1–10, 2007.
44. U. Clarenz, U. Diewald, and M. Rumpf, "Anisotropic geometric diffusion in surface processing," *Proceedings of the conference on Visualization*, pp. 397–405, 2000.
45. T. Tasdizen, R. Whitaker, P. Burchard, and S. Osher, "Geometric surface smoothing via anisotropic diffusion of normals," *Proceedings of IEEE Visualization (VIS2002)*, pp. 125–132, 2002.
46. M. Chuang and M. Kazhdan, "Interactive and anisotropic geometry processing using the screened Poisson equation," *ACM SIGGRAPH*, vol. 30, no. 4, p. 57, 2011.
47. M. Unser, "Splines a Perfect Fit for Signal and Image Processing," *IEEE Signal Processing Magazine*, no. Nov, pp. 22–38, 1999.
48. M. Arigovindan, M. Sühling, P. Hunziker, and M. Unser, "Variational Image Reconstruction from Arbitrarily Spaced Samples: a Fast Multiresolution Spline Solution," *IEEE Trans. Image Process.*, vol. 14, no. 4, pp. 450–60, Apr. 2005.
49. G. Steidl, S. Didas, and J. Neumann, "Splines in Higher Order TV Regularization," *International Journal of Computer Vision*, vol. 70, no. 3, pp. 241–255, Dec. 2006.
50. J. Balzer and T. Morwald, "Isogeometric finite-elements methods and variational reconstruction tasks in vision—A perfect match," *International Conference on Computer Vision and Pattern Recognition*, pp. 1624–1631, 2012.
51. D. Nehab, S. Rusinkiewicz, J. Davis, and R. Ramamoorthi, "Efficiently combining positions and normals for precise 3D geometry," *ACM Trans Graph*, vol. 24, no. 3, p. 536, Jul. 2005.
52. A. Chambolle and T. Pock, "A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging," *Journal of Mathematical Imaging and Vision*, vol. 40, no. 1, pp. 120–145, Dec. 2010.
53. R. T. Rockafellar, *Convex analysis*. Princeton university press, 1997, no. 28.
54. J. Balzer, M. Peters, and S. Soatto, "Volumetric Reconstruction Applied to Perceptual Studies of Size and Weight," *IEEE WACV*, vol. 704, Nov. 2013.