

Consolidation of Unorganized Point Clouds for Surface Reconstruction

Hui Huang¹

Dan Li¹

Hao Zhang²

Uri Ascher¹

Daniel Cohen-Or³

¹University of British Columbia

²Simon Fraser University

³Tel-Aviv University

Abstract

We consolidate an unorganized point cloud with noise, outliers, non-uniformities, and in particular interference between close-by surface sheets as a preprocess to surface generation, focusing on reliable normal estimation. Our algorithm includes two new developments. First, a *weighted locally optimal projection* operator produces a set of denoised, outlier-free and evenly distributed particles over the original dense point cloud, so as to improve the reliability of local PCA for initial estimate of normals. Next, an iterative framework for robust normal estimation is introduced, where a priority-driven normal propagation scheme based on a new priority measure and an orientation-aware PCA work complementarily and iteratively to consolidate particle normals. The priority setting is reinforced with front stopping at thin surface features and normal flipping to enable robust handling of the close-by surface sheet problem. We demonstrate how a point cloud that is well-consolidated by our method steers conventional surface generation schemes towards a proper interpretation of the input data.

1 Introduction

Surface reconstruction from point cloud data has been an extensively studied problem in computer graphics [Turk and Levoy 1994; Carr et al. 2001; Cazals and Giesen 2006; Ohtake et al. 2003; Kazhdan et al. 2006]. Typically acquired by a laser scanner, the raw input points are often unorganized, lacking inherent structure or orientation information. Oriented normals at the points play a critical role in surface reconstruction, as they locally define the reconstructed surface to first order and identify the inside/outside and hence topology of the underlying shape. Although photometric stereo may be applied to estimate normals from captured images, such estimates are not always reliable due to less than ideal acquisition conditions such as specular reflections, material artifacts, and shadowing [Sun et al. 2007]. Indeed, surface normal acquisition is a delicate process [Ma et al. 2007] requiring a well-controlled environment and careful calibration with the process of point acquisition.

We take as input an unorganized point cloud which may contain outliers, noise, and non-uniformities in thickness and spacing, due to acquisition errors or misalignment of multiple scans. Based on point positions alone, we *consolidate* [Alexa et al. 2003] the point cloud. This preprocessing phase for surface reconstruction includes denoising, outlier removal, thinning, orientation, and redistribution of the input points. During the process, we defer and avoid any surface generation, a phase that is highly susceptible to various data artifacts. Decoupling the two phases can effectively avoid premature and erroneous decisions in surface reconstruction.

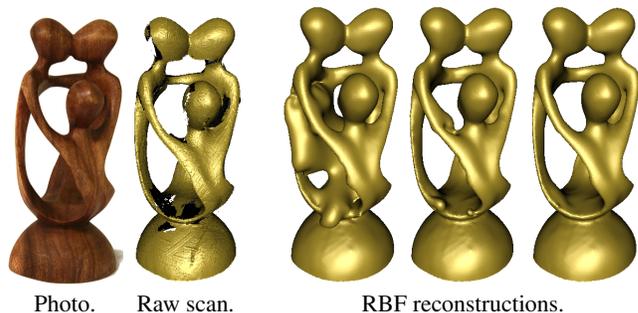


Figure 1: Data consolidation, especially accurate normal estimation, from a noisy, unorganized, raw point cloud is crucial to obtaining a correct surface reconstruction. The right-most result is produced after applying our point cloud consolidation scheme.

A central task to point consolidation is normal estimation. The classical scheme for estimating unsigned normal directions is principal component analysis (PCA), which can be unreliable due to thick point cloud, non-uniform distribution, or close-by surfaces, as shown in Figure 2. The most widely applied approach to consistent normal orientation [Hoppe et al. 1992] is via normal propagation, where propagation between close-by points whose unsigned normal directions make a small angle is given priority. However, under difficult scenarios such as the presence of close-by surfaces, propagation errors do occur, as shown in Figure 3(a-b). Close-by surface sheets also challenge sharp feature detection, a problem relevant to normal propagation. As shown in Figure 3(c), a *thin surface feature*, i.e., a sharp feature delimiting close-by surfaces, does not admit a bi-modal distribution of unsigned normal directions or a good fit using multiple surfaces. Thus, previous approaches to sharp feature detection, e.g., [Page et al. 2002; Fleishman et al. 2005], while generally robust, are not designed to handle such cases.

We address the above issues by combining two techniques. First, to make local PCA more robust, we denoise, remove outliers, and down-sample the otherwise dense input point cloud to obtain a thinned and evenly distributed set of points, called *particles* [Pauly et al. 2002] to distinguish from the input points. In this first step, we modify and extend the locally optimal projection (LOP) operator of Lipman et al. [2007] to deal with non-uniform distributions common in raw data. We then estimate particle normals via an iterative *predictor-corrector* scheme (see [Ascher and Petzold 1998] for origins and analogies for this term). The predictor uses PCA to predict unsigned normal directions. This is followed by particle orientation via a priority-driven normal propagation scheme. The obtained particle orientations are then utilized to correct or consolidate estimates of normal directions via an *orientation-aware PCA*, and the normal propagation scheme is re-applied. A novel contribution in the orientation scheme is a distance measure which prioritizes normal propagation and triggers proper normal flipping. The new measure combines Euclidean and angular distances with *propagation directions* to robustly handle the close-by surface sheet problem. The iterative approach to normal estimation by way of normal correction is also new and shown to be effective and necessary.

We demonstrate that the result of our algorithm, a clean and uniformly distributed point set endowed with reliable normals, leads

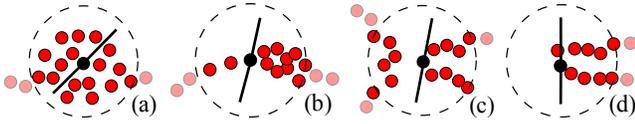


Figure 2: Classical PCA leads to inaccurate estimates of unsigned normal directions (black lines) near a thick point cloud (a), non-uniform distribution (b), or close-by surfaces (c-d).

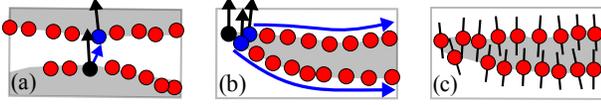


Figure 3: Close-by surfaces cause erroneous normal propagation when prioritized only by Euclidean and angular distances: (a) between particles on opposite surfaces; (b) through a thin surface feature. (c): A thin surface feature does not admit a bi-modal distribution of normal directions or a good fit using multiple surfaces.

to quality up-sampling and Delaunay-based surface reconstruction [Amenta et al. 2001; Dey and Giesen 2001]. Above all, it enables conventional surface generation schemes which rely on point normals, such as radial basis function (RBF) [Carr et al. 2001] and Poisson [Kazhdan et al. 2006] techniques, to obtain a fine interpretation of the input data in various challenging situations. In Figure 1, the two erroneous RBF reconstructions are obtained from a raw scan after preprocessing by LOP (30 and 100 iterations, respectively) and normal estimation via classical PCA and normal orientation [Hoppe et al. 1992]. The final image shows RBF result from the same input after applying 100 iterations of our improved LOP operator and then our normal estimation scheme.

2 Background and related work

The literature on surface reconstruction is vast. Delaunay techniques [Cazals and Giesen 2006] typically produce a mesh which interpolates the input points but contains rough geometry when the points are noisy. These methods often provide provable guarantees under prescribed sampling criteria [Amenta and Bern 1998] or noise models [Dey and Goswami 2006] that are generally not realizable by real-world data. Additional assumptions, such as even point distribution, may also be required in other point set processing which claim guarantees [Mitra et al. 2004]. Approximate reconstruction works mostly with implicit surface representations followed by iso-surfacing. Most notable are methods which carry out implicit modeling via tangent plane estimates [Hoppe et al. 1992], RBF [Carr et al. 2001], or Poisson fields [Kazhdan et al. 2006], all of which require oriented normals. Accurate normal estimates, even point distribution, and sufficient sampling density (via proper up-sampling) can all be achieved by point consolidation.

Deriving a new point set from a given point cloud has been considered in the context of defining point set surfaces. Well-known definitions include moving least squares (MLS) [Alexa et al. 2003] and extremal surfaces [Amenta and Kil 2004]. In practice, these definitions can be applied to smooth or down-sample a raw point cloud. There are also algorithms for point cloud smoothing [Lange and Polthier 2005] and simplification [Pauly et al. 2002] guided by local geometry analysis, such as curvature estimation. To better deal with outliers and delicate surface structures, Lipman et al. [2007] develop a highly effective, parameterization-free projection operator (LOP). However, we have observed that LOP can fail to converge, oscillating near a solution instead, and it may not work well when the distribution of the input points is highly non-uniform.

Point normals are essential for surface reconstruction as they provide local first-order surface approximations and inside/outside directions. Most normal estimation schemes rely on PCA in some form [Hoppe et al. 1992; Pauly et al. 2002; Alexa et al. 2003; Mitra et al. 2004; Lange and Polthier 2005]. Classical PCA relies on Euclidean distances between points. More recently, directional information has been taken into account when computing an improved centroid of a set of points in a local neighborhood [Amenta and Kil 2004; Lehtinen et al. 2008], replacing Euclidean distances by Mahalanobis distances. Nehab et al. [2005] combine positional and normal measurements via photometric stereo to produce a surface which conforms to both. As PCA normals are un-oriented, computing a consistent normal orientation requires additional work. This problem turns out to be surprisingly difficult [Hoppe et al. 1992; Mello et al. 2003], and challenges due to close-by surfaces (Figure 3) have not been specifically addressed in previous works.

Aside from estimating point normals via purely geometric means, acquisition mechanisms such as photometric stereo [Woodham 1980; Nehab et al. 2005] are also possible but they are often subject to error caused by surface or illumination artifacts. If a scanner can return the “outward direction” at each point, a vector from the point to the scanner “head”, such a direction may be used to orient the normal towards the “outside” by insisting that the two directions make an acute angle. However, inaccurately estimated normal directions via classical PCA or near orthogonality between the normal and outside directions can be sources for error, where the latter is likely to occur near close-by surface sheets. We believe that the normal orientation problem is one which requires a global consistency evaluation and is not entirely solvable only through purely local considerations. In practice, the outward directions are still not widely available from current acquisition devices and they can be a source of significant noise, especially for hand-held or other scanners which continually change head positions.

3 Improved weighted LOP (WLOP)

The LOP operator [Lipman et al. 2007] takes as input a noisy point cloud, possibly with outliers, and outputs a new point set which more faithfully adheres to the underlying shape. LOP operates well on raw data without relying on a local parameterization of the points or on their local orientation. Given an unorganized set of points $P = \{p_j\}_{j \in J} \subset \mathbb{R}^3$, LOP defines a set of projected points $X = \{x_i\}_{i \in I} \subset \mathbb{R}^3$ by a fixed point iteration where, given the current iterate X^k , $k = 0, 1, \dots$, the next iterate X^{k+1} is to minimize

$$\sum_{i \in I} \sum_{j \in J} \|x_i - p_j\| \theta(\|\xi_{ij}^k\|) + \lambda_i \sum_{i' \in I \setminus \{i\}} \eta(\|x_i - x_{i'}^k\|) \theta(\|\delta_{ii'}^k\|),$$

with $\xi_{ij}^k = x_i^k - p_j$ and $\delta_{ii'}^k = x_i^k - x_{i'}^k$. In practice, $n = |I|$ is often significantly smaller than $m = |J|$. Intuitively, LOP distributes the points by approximating their l_1 medians to achieve robustness to outliers and data noise. Here $\|\cdot\|$ is the 2-norm, $\theta(r) = e^{-r^2/(h/4)^2}$ is a rapidly decreasing smooth weight function with support radius h defining the size of the influence neighborhood, and $\eta(r)$, the *repulsion term*, is another decreasing function penalizing points x_i that get too close to other points in X . The balancing terms $\{\lambda_i\}_{i \in I}$ vary at each point but depend on just one parameter, $0 \leq \mu < .5$, controlling the repulsion force. Throughout our experiments, we set $\mu = 0.45$ and h may be adjusted, but the default value of $h = 4\sqrt{d_{bb}/m}$, where d_{bb} is the diagonal length of the bounding box of the input model, generally works very well.

New repulsion term LOP often works well, but we have found that the original repulsion function $\eta(r) = 1/(3r^3)$ may drop too quickly to guarantee sufficient penalty when r is large. This could

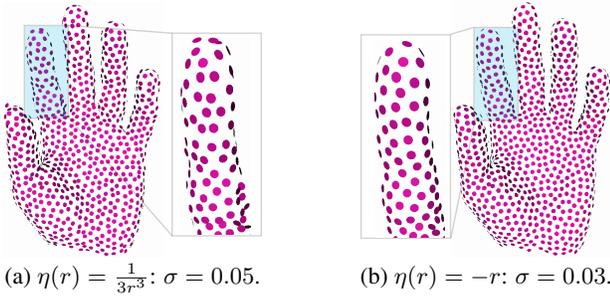


Figure 4: Particle distributions after LOP with different repulsions. For this illustration, all particles are properly oriented with back-face culling. Visually and from the σ measure, we see that the new repulsion term $\eta(r) = -r$ produces a more regular distribution.

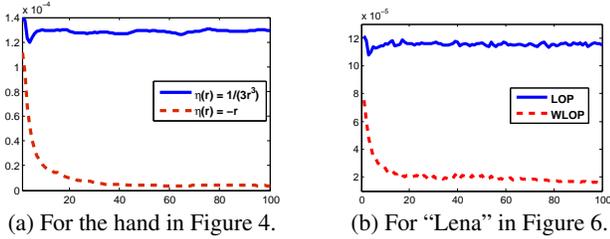


Figure 5: Plots of distances $\|X^{k+1} - X^k\|/n$ between consecutive iterates to illustrate convergence behavior. (a) Without density weights, the old repulsion term leads to oscillation near a solution and the new term results in smooth convergence. (b) With density weights, WLOP apparently retains such convergence property.

lead to a lack of clear-cut convergence and an undesirably irregular point distribution, especially when $n \ll m$. To this end, we propose to use the new repulsion $\eta(r) = -r$, which decreases more gently and penalizes more at larger r , yielding both better convergence and a more locally regular point distribution, as shown in Figures 4(b) and 5(a). As a rough quantitative regularity measure for point distributions, we use the *variance of distances to nearest neighbors* at the points, which we denote by σ throughout.

Density weights The first term in the optimization criteria above for LOP is closely related to the multivariate median, also referred to as the ℓ_1 median, which leads to projection points moving toward the *local distribution center*. If the given point cloud is highly non-uniform, as in the example given by Figure 6(a), projection by LOP tends to follow the trend of such non-uniformity, no matter what initial set X^0 we choose. This may be desirable in certain cases, e.g., to allow higher point densities near shape features. In other cases, e.g., normal estimation, one may prefer uniform point distribution everywhere. To achieve this, we propose to incorporate locally adaptive density weights into LOP, resulting in WLOP.

Let us define the weighted local densities for each point p_j in P and x_i in X during the k th iteration by $v_j = 1 + \sum_{j' \in \mathcal{J} \setminus \{j\}} \theta(\|p_j - p_{j'}\|)$ and $w_i^k = 1 + \sum_{i' \in \mathcal{I} \setminus \{i\}} \theta(\|\delta_{ii'}^k\|)$, $k = 0, 1, 2, \dots$. Then the projection for point x_i^{k+1} finally becomes

$$x_i^{k+1} = \sum_{j \in \mathcal{J}} p_j \frac{\alpha_{ij}^k / v_j}{\sum_{j \in \mathcal{J}} (\alpha_{ij}^k / v_j)} + \mu \sum_{i' \in \mathcal{I} \setminus \{i\}} \delta_{ii'}^k \frac{w_{i'}^k \beta_{ii'}^k}{\sum_{i' \in \mathcal{I} \setminus \{i\}} (w_{i'}^k \beta_{ii'}^k)},$$

where $\alpha_{ij}^k = \frac{\theta(\|\xi_{ij}^k\|)}{\|\xi_{ij}^k\|}$ and $\beta_{ii'}^k = \frac{\theta(\|\delta_{ii'}^k\|) \eta'(\|\delta_{ii'}^k\|)}{\|\delta_{ii'}^k\|}$. Thus, the attraction of point clusters in the given set P is relaxed by the

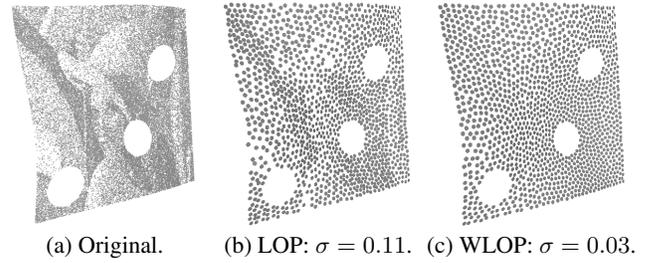


Figure 6: WLOP vs. LOP: (a) The “Lena” image is mapped onto a curved surface with three holes to produce a point set with point densities proportional to image intensities. Then, randomly taking 1/20 of the points in (a) as initial set, the results of LOP and WLOP projections are shown in (b) and (c), respectively. Non-uniformity of the LOP result is manifested by traces of “Lena” in (b).

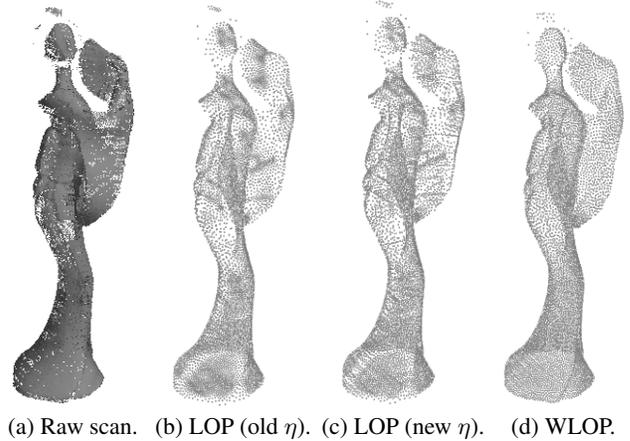


Figure 7: WLOP vs. LOP on the raw scan of a Japanese lady (a). The quantitative measure of point regularity takes on values: (b) $\sigma = 0.24$; (c) $\sigma = 0.18$; (d) $\sigma = 0.09$, indicative of improvement.

weighted local density v in the first term, and the repulsion force from points in dense areas is strengthened by the weighted local density w in the second term.

LOP vs. WLOP Note that LOP with the new repulsion term is a special case of WLOP by setting all density weights to 1. In this case, it is possible to show contraction of the fixed point iteration near an assumed solution. For the more general WLOP with adaptive weights, empirically, we have consistently obtained error plots that are indicative of convergence; see Figure 5(b) for an example. In addition to the synthetic “Lena” example, we also show improved point regularity provided by the new repulsion and density weights on a raw scan example in Figure 7.

4 Normal estimation and consolidation

After WLOP, we obtain a thinned, outlier-free, and uniformly distributed set of particles, denoted by x_1, \dots, x_n . For the next step, normal estimation, we start with the predictor step based on initial unsigned normal directions estimated via classical weighted PCA [Pauly et al. 2002]. The neighborhood size h and weight function θ for PCA are the same as those for WLOP. For subsequent corrector iterations, we employ an orientation-aware PCA (Section 4.2) to consolidate the normals, where particle orientations are obtained via normal propagation (Section 4.1).

4.1 Normal propagation

We wish to find an optimal assignment of particle orientations to maximize a certain *consistency* criterion. Hoppe et al. [1992] use the sum of $1 - |\langle \mathbf{v}_i, \mathbf{v}_j \rangle|$ over all pairs of particles that are sufficiently close to model consistency, where \mathbf{v}_i and \mathbf{v}_j are the particle normal directions. We refer to it as the *traditional* propagation scheme and adopt the same priority-driven propagation strategy while introducing a new priority measure for reliable propagation under problems such as those arising from close-by surfaces.

Overall scheme First, a source particle is selected, where a reliable normal direction can be obtained. Then we perform a conservative check to identify certain particles at thin surface features, where the advancing front of normal propagation is forced to stop. Then the orientation, starting from the source, is propagated, as permitted, via a priority-driven traversal of the particles. Specifically, once a particle is oriented, its k -nearest neighbors (k NNs), where $k = 6$ by default, are added into a priority queue. Potential orientation errors may happen possibly due to the greedy approach or erroneous propagation near undetected thin surface features. Thus an additional error check is performed which may trigger one or more *normal flips*. This is followed by another propagation pass, and these may be iterated until no more orientation changes.

Source selection Typically, an extremal particle, e.g., one with the maximum x coordinate, is chosen [Hoppe et al. 1992]. However, it is not unusual for such an extremal particle to be at a sharp feature and cause erroneous results, e.g., see Figure 3(b). We propose to pick a source over a *flat region* — a particle whose unsigned direction has the least angular variation from those of its k NNs. As the particles have been denoised and evenly distributed by WLOP, a desirable source can be reliably found. The orientation chosen at the source is less important, as a flipping of all the consistent normals, if deemed necessary, is simple to carry out at the end.

Distance measure Previous considerations for the distance or priority measure which drives normal propagation include both Euclidean and angular distances but could still lead to error, e.g., see Figure 3(a). What has been missing is the *direction of propagation*. The intuition here is that a correct normal propagation should less likely be along the local normal direction — it should be along the local *tangential direction*, which, for two close-by particles, we approximate simply by the vector connecting the particles.

Let x_i and x_j be two particles with associated directions \mathbf{v}_i and \mathbf{v}_j (their orientations do not play a role in the following analysis), respectively. Consider four points x_i^1, x_i^2, x_j^1 and x_j^2 that are unit distance away from x_i and x_j along these directions, as shown in Figure 8. Let m_{rs} be the midpoint of the line segment $\overline{x_i^r x_j^s}$, $r, s \in \{1, 2\}$, and o_{rs} the perpendicular projection of m_{rs} onto the estimated tangent line $\overline{x_i x_j}$ or its extensions. Note that if x_i and x_j coincides ($\overline{x_i x_j}$ is undefined), we simply let $o_{rs} = x_i$. We define the normalized distance to prioritize normal propagation by

$$\mathcal{D}_{ij} = 1 - |\langle \mathbf{v}_i, \mathbf{v}_j \rangle| \cdot \frac{\max_{r,s \in \{1,2\}} \|m_{rs} - o_{rs}\|}{1 + \|x_i - x_j\|}. \quad (1)$$

Note that $\mathcal{D}_{ij} \in [0, 1]$; it combines Euclidean distance (the denominator), angular distance $|\langle \mathbf{v}_i, \mathbf{v}_j \rangle|$, and a third term $d_{ij} = \max_{r,s \in \{1,2\}} \|m_{rs} - o_{rs}\|$, which is designed to weigh in propagation direction. See Figure 8 for an analysis of d_{ij} and \mathcal{D}_{ij} and Figure 9 for a comparison to the traditional scheme.

We remark that our distance \mathcal{D} is not a metric and neither should it be. From Figure 3(b), we see that propagation from the black

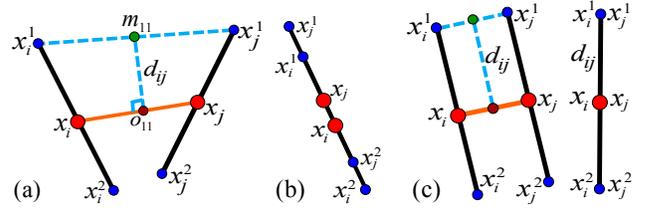


Figure 8: Propagation distance \mathcal{D}_{ij} (1) between particles x_i and x_j with unsigned normal directions. (a) Maximal projected distance d_{ij} from midpoint m_{rs} , $r, s \in \{1, 2\}$, to $\overline{x_i x_j}$ captures propagation direction information. (b) When normal directions $\overline{x_i^1 x_i^2}$ and $\overline{x_j^1 x_j^2}$ coincide with $\overline{x_i x_j}$, signifying a propagation along normal direction, we have $d_{ij} = 0$ and $\mathcal{D}_{ij} = 1$ at its maximum. (c) As normal directions become more aligned and perpendicular to $\overline{x_i x_j}$, d_{ij} increases. It attains maximum value 1 when these conditions hold exactly, signifying a propagation along tangential direction. Further, if x_i and x_j are coincident then \mathcal{D}_{ij} is minimized at 0.

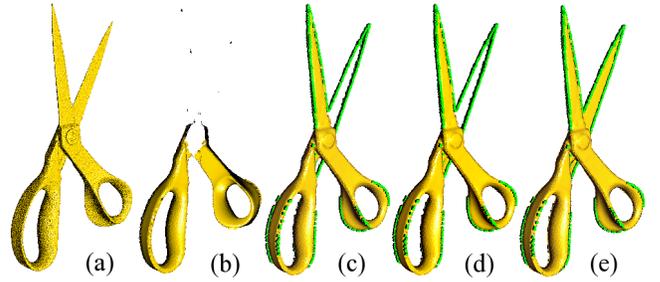


Figure 9: Steps of our normal propagation scheme. (a) Raw data with thin features. (b) Result of traditional scheme. (c) With new distance measure and awareness of thin features (green), better results but errors still remain. (d) Adding normal flipping fixes some errors (under the pivot). (e) Three corrective iterations with orientation-aware PCA lead to final successful orientation.

particle to each blue particle is encouraged (small \mathcal{D}), but not between the blue particles (large \mathcal{D}) as they belong to opposite surfaces; such a distance configuration violates the triangle inequality. Finally, taking direction information into account when measuring distances is not new. The Mahalanobis distance is defined between a point and an oriented point with a stretch factor characterizing the elliptical field around the oriented point. Our distance avoids such a free parameter and is an integrated measure defined on two unsigned directions associated with particles.

Thin surface features and normal flipping Although \mathcal{D} can by and large avoid propagation between particles residing on close-by and opposite surfaces, it does not prevent propagation through a thin surface feature, one which separates two such surfaces, as shown in Figure 3(b). We design a simple and *conservative* method to detect such features. The k NNs of particle x_i are projected onto its tangent plane, which is determined by the current unsigned normal at x_i . If the projection of x_i lies outside the convex hull of its k NN projections, then x_i is deemed to be at a thin feature; see Figure 10(a). Particles at a thin surface feature can be oriented, but are not allowed to propagate their orientations.

Note that the above detection mechanism is only specialized to handle thin surface features: it is not a generic sharp feature detector. Moreover, it cannot distinguish between a flat neighborhood and a thin feature whose crease is a concave curve; see Fig-

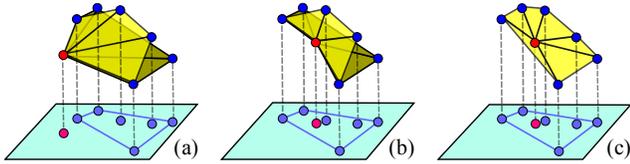


Figure 10: Particle projection (red) lying outside the convex hull of projection of k NNs implies a thin surface feature (a). This test does not distinguish between a flat neighborhood (c) and a case where the particle lies on a concave crease curve (b). Triangles are used only to aid visualization; they are not part of the data.

ure 10(b-c). As a remedy, we execute a check during the normal propagation to detect and reverse orientation between close-by surface sheets. Specifically, for a propagated particle pair x_i and x_j , if both $|\cos(\angle(\mathbf{n}_i, \vec{x_i x_j}))|$ and $|\cos(\angle(\mathbf{n}_j, \vec{x_i x_j}))|$ exceed a threshold (set to 0.8 throughout), signifying a potential propagation along normal direction, we *flip* the normal orientation at particle x_j . Then, the priority-driven propagation continues.

4.2 Orientation-aware PCA

The normal propagation scheme described above works on a fixed set of unsigned normal directions. Despite all the care taken so far, orientation errors may occasionally persist due to error in the normal directions computed by the classical, orientation-oblivious PCA. By making PCA orientation-aware, unsigned normal directions and orientation estimations can complement each other and fix errors within a corrector loop.

In our implementation, when performing local weighted PCA at a particle x_i , we exclude from a Euclidean h -ball centered at x_i those particles whose oriented normals are opposite (negative dot product) to the normal at x_i . In other words, the considered neighbors are now all those *facing* the same way as x_i . The unsigned normals recomputed in this way would not be expected to change much on flat, correctly oriented regions, but they may well vary and become more accurate near thin structures or areas of surface interference. Thus, the errors in subsequent orientation sweeps via normal propagation may be reduced. We apply such corrector iterations until normal orientations no longer change. Figures 9 and 15 show orientation errors corrected via orientation-aware PCA.

5 Results and applications

Point cloud consolidation cleans up raw input, removes a variety of data artifacts, and provides essential geometric attributes, in our case point normals, to facilitate subsequent processing. In this section, we demonstrate how a well consolidated point set via WLOP and normal estimation using our iterative framework can benefit such processing tasks as up-sampling and surface reconstruction.

Visualization of point sets is best achieved using splatting, based on points with normals or *surfels*. A frequently encountered operation during splatting is point cloud up-sampling, e.g., for a zoomed-in view or when the given point cloud was under-sampled during data acquisition or subsampled for efficient processing. We employ the fast dynamic algorithm of Guenebaud et al. [2004] for real-time point cloud refinement in our experiment. Accurate normals and regular point distributions are typical requirements to achieve quality for such up-sampling, and Figure 11 shows the kind of positive difference point consolidation can make.

Let us now show the necessity of a disciplined point consolidation step for quality surface reconstruction. Noisy input must first be

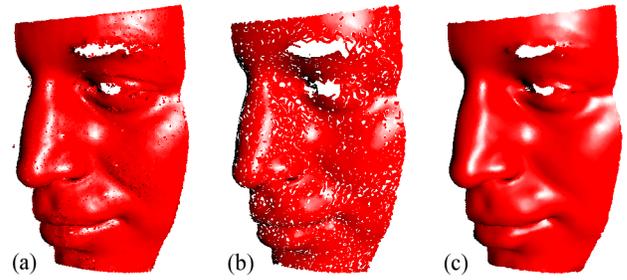


Figure 11: Effect of point consolidation on up-sampling. (a) Noisy input data with 84,398 points. (b) Result of up-sampling, to 95,863 points, from a down-sampled (2,814 points or 3.3%) point set obtained from (a), after data cleaning by LOP and normal estimation via classical PCA. (c) Up-sampling to 91,438 points after the same down-sampled point set is consolidated using our algorithm.

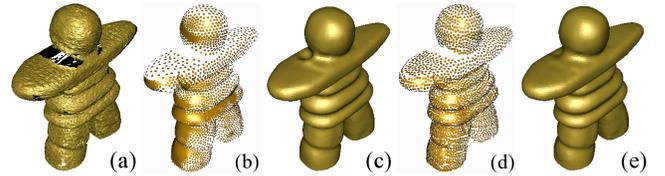


Figure 12: Effect of WLOP on RBF surface reconstruction. Raw scan (a) of an Inukshuk is cleaned by the original LOP (b) with the point regularity measure taking on value $\sigma = 0.2953$ and our WLOP (d) with $\sigma = 0.0987$. Observe the uniformity of the resulting particle distributions. After normal consolidation and up-sampling, RBF constructions, (c) from (b) and (e) from (d), show qualitative difference in hole filling, e.g., around the neck.

cleaned before surface generation, since imperfect point distribution or orientation, which occurs with existing schemes, can result in visible reconstruction error, as first shown in Figure 1.

We mainly draw comparisons with the use of classical PCA and the traditional normal propagation scheme due to Hoppe et al. [1992]. In addition, we also provide an example comparing our consolidation framework with a Delaunay-based one: NormFet+AMLS. These steps, combined with the well-known Cocone mesh generation [Dey and Giesen 2001], are a series of techniques developed by Dey and co-authors. In particular, NormFet [Dey and Giesen 2005] performs normal estimation in the presence of noise using the Delaunay ball technique and AMLS [Dey and Sun 2005] employs adaptive MLS for smoothing noisy point clouds based on normals and detected features from NormFet. In other cases, we choose RBF [Carr et al. 2001] and Poisson [Kazhdan et al. 2006] surface reconstruction for demonstration. The implementations are due to FarField Technology (FastRBF) and M. Kazhdan, respectively.

First, we show the effect of WLOP in Figure 12. The input is a raw scan with missing data. After WLOP, up-sampling and robust normal consolidation, RBF is able to successfully close holes and construct a quality surface. In contrast, with the original LOP operator, although noise and outliers are removed as well, the resulting irregular particle distribution (quantified by σ) may cause some defects on hole closure during surface generation.

Next, we show the effect of normal orientation, where all the input point clouds were first cleaned via WLOP and then subsequently up-sampled. For a Delaunay-based approach, we consider the (NormFet, AMLS, Cocone) combination. Our experiments show that errors arising from the traditional scheme or Delaunay-based

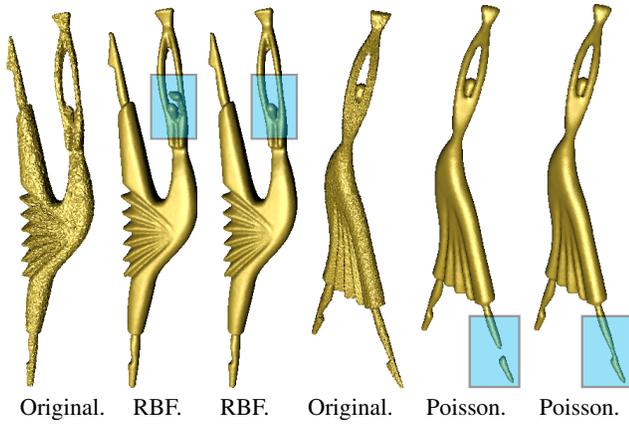


Figure 13: Effect of normal consolidation on surface reconstruction for models with thin structures. In each series, following the original, we show reconstruction results after normals are computed via the traditional scheme, and then results after our normal estimation algorithm. Highlighted areas show differences made by the latter.

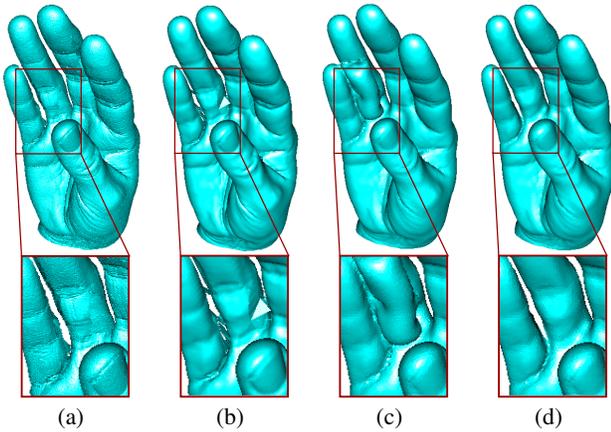


Figure 14: Effect of normal consolidation on surface reconstruction for a raw point cloud (a) with close-by surfaces and missing data. (b) Result from (NormFet, AMLS, Cocone). (c) RBF result after normal estimation via the traditional scheme. (d) RBF result after our normal estimation scheme.

approach may lead to various topological artifacts in the reconstructions. Such errors typically occur near thin surface structures (Figure 13) or close-by surface sheets (Figures 14 and 15), where our point cloud consolidation method succeeds. In particular, Figures 14 and 1 show that accurate normals can effectively compensate for missing data in a point cloud, allowing reconstructions, such as RBF, to infer the underlying shape correctly. Finally, we provide timing results for our algorithm in Table 1.

Limitations Although the strengths of our method lie in its handling of thin surface structures, failure cases can still occur in cases under extreme conditions such as severe noise or undersampling. For example, in Figure 16, we see that the ears of the horse are thin structures having extremely low sampling rate. Our algorithm treats each ear as a single sheet and the resulting reconstruction has obvious defects. Another such example leading to topological error can be observed between the feet of the mannequin in Figure 15. Ideally, we would like to obtain a theoretical guarantee for the correctness of our normal estimates under appropriate

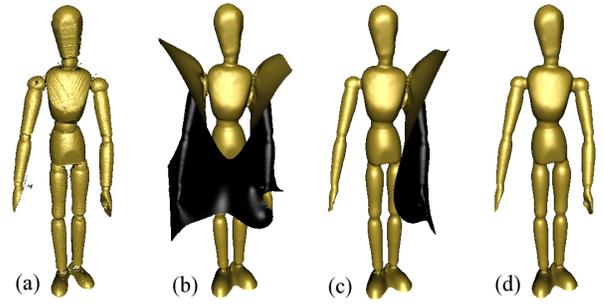


Figure 15: Effect of normal consolidation, in particular, the corrector iteration and orientation-aware PCA. (a) Raw scan. (b) RBF result after normal estimation via the traditional scheme. (c) RBF result based on normals oriented by one pass of our propagation scheme. (d) RBF result after further correction of orientation errors via iteration and orientation-aware PCA.

Table 1: CPU runtime for consolidation of several raw datasets. O-No: number of original points; P-No: number of projected particles; W-T: time for WLOP; N-T: time for normal estimation; U-T: time for up-sampling. Only the Face model in Figure 11 is up-sampled twice. All examples were run on an Intel Pentium 4, 3.2 GHz CPU with 1GB RAM and times are reported in seconds.

	O-No	P-No	W-T	N-T	U-T
Figure 1	634,386	21,147	443	28	166
Figure 11	84,398	8,440	47	7	492
Figure 12	206,002	10,301	112	8	73
Figure 14	204,068	20,407	69	35	138
Figure 15	261,011	13,051	50	46	111

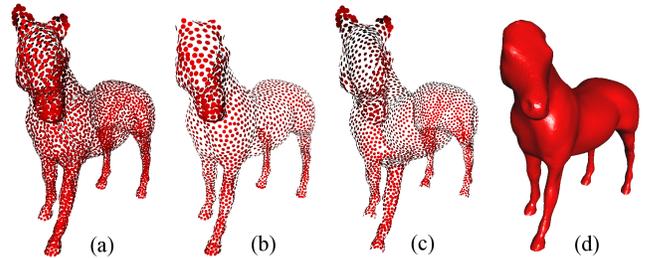


Figure 16: A failure case in the presence of extreme undersampling. (a) A horse point set consolidated using our algorithm; note severe undersampling near the ears. (b) Back face culling view. (c) Front face culling view. (d) Poisson surface reconstruction.

sampling conditions. Also on the theory front, we do not have a convergence proof for the iterative predictor-corrector scheme for normal estimation. In practice, we have not encountered a case of oscillation either. Like LOP, our point consolidation framework does not address the missing data problem. However, numerous examples highlight the importance of having accurate normals for surface completion schemes such as RBF and Poisson to succeed.

6 Conclusion and future work

Accurate estimation of normals is crucial to obtaining a correct interpretation of the input data. We show that the incorporation of propagation direction information into priority setting, as well as a coupled and iterative approach on normal orientation and orientation-aware PCA, provides consolidation of the data points

in various difficult settings. The prelude to all these is a necessary step for data clean-up, for which we develop WLOP, an improved locally optimal projector with weighting option for denoising and outlier removal from imperfect point data and producing an evenly distributed set of particles which faithfully adheres to the captured shape. With our point cloud consolidation, conventional surface reconstruction schemes can better infer the topology and geometry of the shape from raw input data in challenging situations. We believe that such consolidation of points should be a routine procedure applied to raw data similarly to common denoising procedures.

While our current consolidation algorithm has been shown to perform robustly and efficiently through numerous experiments, we next would like to seek a rigorous theoretical analysis of the predictor-corrector iteration. Also possible as future work is better handling of missing data, taking advantage of the reliable orientation information we can extract from the raw input. Finally, we would like to incorporate recovery and enhancement of sharp features into our point consolidation framework.

Acknowledgments

The authors would like to thank all the reviewers for their valuable comments. This work is supported in part by grants from NSERC (No. 84306 and No. 611370), the Israeli Ministry of Science, and the Israel Science Foundation. The hands, horse, dancer, and scissors data are from the AIM@SHAPE shape repository. The face model in Figure 11 is courtesy of Yaron Lipman. Our code is based on the VCG library from the Visual Computing Lab in Pisa, Italy. Thanks go to Federico Ponchio for the original LOP implementation and consultation on VCG.

References

- ALEXA, M., BEHR, J., COHEN-OR, D., FLEISHMAN, S., LEVIN, D., AND SILVA, C. T. 2003. Computing and rendering point set surfaces. *IEEE Trans. Vis. & Comp. Graphics* 9, 1, 3–15.
- AMENTA, N., AND BERN, M. W. 1998. Surface reconstruction by Voronoi filtering. In *Symp. on Comp. Geom.*, 39–48.
- AMENTA, N., AND KIL, Y. J. 2004. Defining point-set surfaces. *ACM Trans. on Graphics* 23, 3, 264–270.
- AMENTA, N., CHOI, S., AND KOLLURI, R. K. 2001. The power crust. In *ACM Symp. on Solid Modeling and Appl.*, 249–266.
- ASCHER, U., AND PETZOLD, L. 1998. *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*. SIAM, Philadelphia, PA.
- CARR, J. C., BEATSON, R. K., CHERRIE, J. B., MITCHELL, T. J., FRIGHT, W. R., MCCALLUM, B. C., AND EVANS, T. R. 2001. Reconstruction and representation of 3D objects with radial basis functions. In *Proc. of ACM SIGGRAPH*, 67–76.
- CAZALS, F., AND GIESEN, J. 2006. Delaunay triangulation based surface reconstruction. In *Effective Computational Geometry for Curves and Surfaces*. Springer, 231–276.
- DEY, T. K., AND GIESEN, J. 2001. Detecting undersampling in surface reconstruction. In *Symp. on Comp. Geom.*, 257–263.
- DEY, T. K., AND GIESEN, J. 2005. Normal estimation for point clouds: a comparison study for a Voronoi based method. In *Eurographics Symp. on Point-Based Graphics*, 39–46.
- DEY, T. K., AND GOSWAMI, S. 2006. Provable surface reconstruction from noisy samples. *Comp. Geom.: Theory & Appl.* 35, 1, 124–141.
- DEY, T. K., AND SUN, J. 2005. An adaptive MLS surface for reconstruction with guarantees. In *Symp. on Geom. Proc. (SGP)*, 43–52.
- FLEISHMAN, S., COHEN-OR, D., AND SILVA, C. T. 2005. Robust moving least-squares fitting with sharp features. *ACM Trans. Graph.* 24, 3, 544–552.
- GUENNEBAUD, G., BARTHE, L., AND PAULIN, M. 2004. Real-time point cloud refinement. In *Eurographics Symp. on Point-Based Graphics*, 41–49.
- HOPPE, H., DEROSE, T., DUCHAMP, T., McDONALD, J., AND STUETZLE, W. 1992. Surface reconstruction from unorganized points. In *Proc. of ACM SIGGRAPH*, 71–78.
- KAZHDAN, M., BOLITHO, M., AND HOPPE, H. 2006. Poisson surface reconstruction. In *Symp. on Geom. Proc. (SGP)*, 61–70.
- LANGE, C., AND POLTHIER, K. 2005. Anisotropic smoothing of point sets. *Comput. Aided Geom. Des.* 22, 7, 680–692.
- LEHTINEN, J., ZWICKER, M., TURQUIN, E., KONTKANEN, J., DURAND, F., SILLION, F., AND AILA, T. 2008. A meshless hierarchical representation for light transport. *ACM Trans. on Graphics* 27, 3, 37:1–37:9.
- LIPMAN, Y., COHEN-OR, D., LEVIN, D., AND TAL-EZER, H. 2007. Parameterization-free projection for geometry reconstruction. *ACM Trans. on Graphics* 26, 3, 22:1–22:6.
- MA, W.-C., HAWKINS, T., PEERS, P., CHABERT, C.-F., WEISS, M., AND DEBEVEC, P. 2007. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Eurographics Symp. on Rendering*, 183–194.
- MELLO, V., VELHO, L., AND TAUBIN, G. 2003. Estimating the in/out function of a surface represented by points. In *ACM Symp. on Solid Modeling and Appl.*, 108–114.
- MITRA, N. J., NGUYEN, A., AND GUIBAS, L. 2004. Estimating surface normals in noisy point cloud data. *Int. J. Comput. Geom. and Appl.* 14, 261–276.
- NEHAB, D., RUSINKIEWICZ, S., DAVIS, J., AND RAMAMOORTHY, R. 2005. Efficiently combining positions and normals for precise 3D geometry. *ACM Trans. on Graphics* 24, 3, 536–543.
- OHTAKE, Y., BELYAEV, A., ALEXA, M., TURK, G., AND SEIDEL, H.-P. 2003. Multi-level partition of unity implicit. *ACM Trans. on Graphics* 22, 3, 463–470.
- PAGE, D. L., SUN, Y., KOSCHAN, A., PAIK, J., AND ABIDI, M. A. 2002. Normal vector voting: Crease detection and curvature estimation on large noisy meshes. *Graphical Models* 64, 199–229.
- PAULY, M., GROSS, M., AND KOBBELT, L. P. 2002. Efficient simplification of point-sampled surfaces. In *Proc. of IEEE Visualization*, 163–170.
- SUN, J., SMITH, M., SMITH, L., AND FAROOQ, A. 2007. Examining the uncertainty of the recovered surface normal in three light photometric stereo. *Image Vis. Comput.* 25, 7, 1073–1079.
- TURK, G., AND LEVOY, M. 1994. Zippered polygon meshes from range images. In *Proc. of ACM SIGGRAPH*, 311–318.
- WOODHAM, R. J. 1980. Photometric method for determining surface orientation from multiple images. *Optical Engineering* 19, 1, 139–144.