

GlobFit: Consistently Fitting Primitives by Discovering Global Relations

Yangyan Li* Xiaokun Wu*,† Yiorgos Chrysathou Andrei Sharf Daniel Cohen-Or Niloy J. Mitra

†SIAT, China Zhejiang Univ. Cyprus Univ. Ben-Gurion Univ. TAU KAUST




Figure 1: Starting from a noisy scan, our algorithm recovers the primitive faces along with their global mutual relations, when are then used to produce a final model (all lengths in mm).

Abstract

Given a noisy and incomplete point set, we introduce a method that simultaneously recovers a set of locally fitted primitives along with their global mutual relations. We operate under the assumption that the data corresponds to a man-made engineering object consisting of basic primitives, possibly repeated and globally aligned under common relations. We introduce an algorithm to directly couple the local and global aspects of the problem. The local fit of the model is determined by how well the inferred model agrees to the observed data, while the global relations are iteratively learned and enforced through a constrained optimization. Starting with a set of initial RANSAC based locally fitted primitives, relations across the primitives such as orientation, placement, and equality are progressively learned and conformed to. In each stage, a set of feasible relations are extracted among the candidate relations, and then aligned to, while best fitting to the input data. The global coupling corrects the primitives obtained in the local RANSAC stage, and brings them to precise global alignment. We test the robustness of our algorithm on a range of synthesized and scanned data, with varying amounts of noise, outliers, and non-uniform sampling, and validate the results against ground truth, where available.

Keywords: 3D scanning, RANSAC, global relations, data fitting, symmetry relations.

1 Introduction

Mechanical parts mostly consist of simple primitives arranged together while adhering to precise global inter-part relations that naturally arise from design and fabrication considerations. Common design tools facilitate realizations involving regular arrangements and snapping to existing parts; analogously, functional requirements, fabrication constraints, and restricted budget considerations favor objects with relations among distant parts forming repeated subcomponents. Such relations not only manifest as orthogonal or parallel faces, but also as *precise* equality of attributes across primitives, both neighboring and distant, resulting in aligned placements, equality among subtended angles and encompassed lengths. Thus, seemingly complex man-made objects may have low information content consisting of primitive parts conforming to global relations (see Figure 2). In noisy, possibly incomplete, scanned data such relations, which are critical to the functionality of the original objects, are easily subdued and lost. Precise recovery of such relations remains challenging for low fidelity scans, especially with increased popularity of cheap, yet imprecise, acquisition devices such as Handyscan 3D[®] or Microsoft Kinect[®].




Figure 2: Man-made objects commonly consist of primitive faces conforming to various global relations.

A popular strategy in reverse engineering involves locally fitting primitives like planes, cylinders, cones using state-of-the-art RANSAC based methods [Schnabel et al. 2009]. Such a local approach, by itself, can be unreliable, especially in regions of biased noise or incomplete data leading to globally inconsistent reconstructions, and hence form poor proxies for the corresponding mechanical parts. We argue that unlike local relations, global ones are less easily disturbed. Further, such relations being non-local span a wider extent of the object, and thus are more robust to local inconsistencies. In this paper, we present a framework to learn and conform to global relations (see Figures 1 and 4).

Existing approaches typically make use of *smoothness priors* to process incomplete and noisy data. Alternately, Gal et al. [2007] use *local priors* to fit primitive shapes like boxes, cylinders, cones to the scanned data. While the strategy can produce sharp features using those inherited from the primitive shapes, the method being local fails to conform to global relations, which constitute essential characteristics of mechanical parts. Further, such approaches typically necessitate committing to a partitioning of the input early on, posing an additional challenge.

In this paper, we use global relations for recovering exact relations and constraints from imperfect acquisitions of man-made engineer-

* The work was primarily done while the first and second authors were visiting students at KAUST.




Figure 3: Our algorithm transforms the noisy input point data (left) to the final reconstructed model (right) through several iterations of RANSAC fitting and constrained optimization. The optimization achieves global alignment in three sequential stages, with each enforcing a different class of global relations (orientation, placement, or equality), while always accounting for the underlying data, and respecting relations conformed to in the previous stages. The final result is visualized with the fitted planes clipped by neighboring ones.

ing parts, possibly with structured noise, or misaligned scan parts. The challenge is to robustly infer potential relationship structures, possibly across distant parts, and refine them based on the input. Our algorithm directly couples the local and global aspects of the problem. The local fit of the model is determined by how well the inferred model agrees to the observed data, while the global relations are iteratively learned, and enforced via a constrained optimization. We favor alignment-relations involving orientation, placement, and equality across the primitives. The recovered global priors boost the local primitive fitting to a noise-resilient global primitive fitting (see Figure 12).

Primitives extracted by local RANSAC are unstable specially in regions of limited support. A naive sequential alignment to dominant primitives is insufficient as such an approach ignores the inter-relations across the relations themselves (see Figure 4). We take a global approach to constrain and optimize the local RANSAC based primitives. The problem is challenging since we simultaneously learn and infer valid relations among uncertain global constraints, while avoiding conflicts. We show that balancing between the data driven fitting error and the regularization effects of the inferred relationship structure, allows us to capture both local and global information directly from incomplete and noisy scans. In each stage, we extract a set of valid relations and use a nonlinear data error optimizer minimizing subject to the inferred constraints to simultaneously align all the primitives. Our algorithm works from coarse to fine scale, wherein relations from the coarse scale help position the unstable primitives extracted in the fine scale. We demonstrate the robustness of our algorithm on a variety of synthetic and scanned models subject to a wide range of data contaminations.

2 Related Work

Surface reconstruction from a scanned, unstructured point cloud is a difficult and ill-posed problem (see also [Dey 2007]). Usual strategies require high quality input data, restrict the problem space by focusing on specific object classes, or use prior knowledge about the scanned objects.

The pioneering work of Hoppe et al. [1992] computes a signed distance field to the underlying surface using local tangent plane construction, and extracts an isosurface from the distance field using the marching cubes algorithm. Over the years, other implicit formulations have been proposed, including radial basis function [Carr et al. 2001], Moving Least Squares (MLS) [Alexa et al. 2003], thin plate splines [Brown and Rusinkiewicz 2004], Poisson reconstruct-

tion [Kazhdan et al. 2006], and recently global stochastic signing of distance functions [Mullen et al. 2010]. These methods employ a *smoothness prior* to regularize the solution by favoring smooth reconstructions. Such an assumption is unsuited in many man-made and CAD objects, which often are piecewise smooth and contain characteristic sharp features.

Alternate strategies have been proposed to recover *sharp features*. Methods that rely on an explicit representation of sharp features [Guennebaud and Gross 2007] classify the input samples into piecewise linear components to indirectly model sharp features. A more challenging task is to detect and reconstruct the sharp features present in noisy point clouds [Fleishman et al. 2005; Lipman et al. 2007; Oztireli et al. 2009]. These methods are typically local and ignore global characteristics of the subject.

Man-made objects are often composed of an assembly of basic primitive shapes congregated using global relations such as regularity and intra-symmetry relations. This is particularly true for architectural, CAD and mechanical models, which are predominantly made of repetitive basic structures to facilitate easy and economic fabrication. Surface reconstruction involving *local fitting* of primitive structures has long been the standard in reverse engineering [Benko et al. 2001] and also in the graphics community (cf. [Schnabel et al. 2007] and references therein). These methods support reconstruction of shapes consisting of sharp features through local fitting of surface geometry. In contrast, we take a global approach which accounts for the inter-primitive relations, both local and global.

In computer vision and reverse engineering, Fisher and colleagues (see [Werghe et al. 2002; Fisher 2004] and references therein) demonstrate the effectiveness of domain knowledge of standard shapes and relationships for improved architectural reconstruction, better parameter estimation, and data completion using non-local relations (cf. [Thrun and Wegbreit 2005]). The series of work uses prior knowledge of specific buildings, placement constraints, or extracted features like principal directions for reconstructing buildings from range images. Guided by a similar philosophy, we also make use of global relations and placement constraints, which are automatically extracted by our algorithm.

Starting from an input scan, Gal et al. [2007] use multi-scale partial matching to fit a small set of basic shapes to local neighborhoods as *local priors*. To facilitate reconstruction, the scan is augmented with noise-free samples, quality normals, and sharp features from the matched primitive shapes. Schnabel et al. [2009] present an interesting hole-filling algorithm that is guided by primitive detection

in the input point cloud. Since direct detection of good geometric primitives is difficult due to noise and missing data, they minimize an energy function defined over a guiding vector field and employ a graph-cut based approach to partition the data. In contrast, we detect global relations in the input scan, encode them in a dynamically constructed relationship graph structure, and conform to the learned relations using a data driven (non-linear) constrained optimization.

Our work is inspired by other efforts that are not directly related to reconstruction: Mitra et al. [2007] use symmetrization to bring shapes into canonical poses in an effort to enhance the (approximate) Euclidean symmetries present in the inputs. However, reliable detection of global relations in presence of significant noise and outliers remain a challenging problem (cf. [Pauly et al. 2008; Li et al. 2010] and references therein). Gal et al. [2009] propose an intuitive manipulation framework for man-made objects by analyzing their typical features and characteristics. They use wires to encode local geometric characteristics as well as global relationships across the wires, which are greedily optimized during deformations. The wires, in turn, are augmented with geometric attributes allowing coupled manipulation of the input. Our goal is to extract structural relations from uncertain data, while [Gal et al. 2009; Mehra et al. 2009] focus on using such relations. Detected symmetries and relations have also been used for inverse procedural modeling [Mitra and Pauly 2008; Bokeloh et al. 2010]. In our work, we also incorporate local reconstruction coupled with an extracted *relationship* graph in order to encode the intra-part relations in the input man-made shape.

3 Algorithm

Given a noisy scanned point set, we simultaneously recover a set of fitted basic primitives along with their global mutual relations. The algorithm partitions the data, extracts the fitted primitives, and simultaneously learns and enforces the mutual relations, among both near and distant parts. While current approaches focus primarily on




Figure 4: Given a noisy scan (top-left), a RANSAC based approach outputs a set of fitted basic primitives (top-right), while ignoring inter-primitive global relations. A greedy alignment of primitive relations produces suboptimal output (bottom-left). We use a constrained optimization to simultaneously position the primitives to enforce their global relationship structure (bottom-right). For comparison, we overlay our results (axes) on the RANSAC and greedy results (in black). See also Figure 12.

local primitive fitting, and ignore the global alignments due to the global relations, our method balances between the two.

We focus on three classes of commonly encountered relations, also referred to as *alignments*: (i) *orientation* relations like parallelism, orthogonality, regular angles, (ii) *placement* relations particularly coplanarity and coaxial relations between primitives, and (iii) *equality* relations among primitives. We solve a global problem to simultaneously learn which constraint relations to enforce without introducing conflicts, and optimize for the parameters of each primitive to meet the learned constraints, while conforming to the input (see Figure 3). The problem is challenging since we have to choose among uncertain relations. Based on the sensitivity of the various relations to noise, we detect them sequentially, starting with the most robust one. We perform global alignments one after another, without breaking already learned and satisfied relations.

State-of-the-art solutions, e.g., Schnabel et al. [2007], typically use a local RANSAC based approach to decompose the input points into subsets associated with fitted primitive shapes and a set of unclaimed points. To account for global relations, a naive option is to take the local primitive fits and progressively fix them, satisfying relations with the primitives having higher confidence. In each step of such an approach, a single primitive is positioned, and this placement is fixed during the subsequent steps. Further, primitives with lower confidence do not influence those with higher confidence. Such a greedy approach can be unreliable and fragile as one is forced to commit early to uncertain global relations (see Figures 4 and 12). Instead, we take a more global view leading to a robust primitive fitting, while simultaneously balancing among multiple constraints.

Let $P := \{\mathbf{p}_1, \mathbf{p}_2, \dots\}$ be the input pointset, with each point \mathbf{p}_i equipped with a normal \mathbf{n}_i and a confidence score w_i estimated using a standard local covariance analysis. Note that we do not require consistently oriented normals. To bootstrap the algorithm, we partition the input into sets $P = \cup_i P_i \cup P^*$, where each point group P_i has an associated primitive, say χ_i , along with initial estimates of its respective parameters, and P^* denotes the remaining points [Schnabel et al. 2007]. Later, we repartition the inputs based on the learned relations (see Figure 10).

For each of the alignment classes, we first construct a *relationship graph*. For example, in case of one type of orientation constraints, we check the angles between the normals of all pairs of primitives for candidate orthogonal or parallel relations. A confidence score is associated with each such relation. Let, C be the potential relations ordered according to their decreasing scores. We progressively extract a *maximal* subset of relations $C^* (\subset C)$ without any *redundant* constraints, and more importantly without any *conflicts*. By maximality, we refer to a *complete set* of constraints without redundancy, i.e., $\max |C^*|$, over all possible non-conflicting subsets. Such a non-redundant relation set significantly improves the efficiency and robustness of the constrained optimization (compare Figure 12). Since the exact minimal graph extraction problem is hard (even the boolean SAT problem is NP-complete), we use an iterative approximation. For other types of relations, we also build corresponding relationship graphs, but there each graph node represents a pair of primitives, as explained later.

Having extracted the non-redundant and non-conflicting relations C^* , along with the relations learned in the previous stages, we solve a nonlinear optimization over the parameters of primitives χ_i to minimize data error while exactly conforming to the constraint set C^* . In this stage, all the parameters of the primitives are simultaneously optimized for, while keeping the point association to the respective sets P_i fixed.




Figure 5: Progressive construction of conflict-free reduced orientation graph C_o^* for a cube dataset. The set of detected relations in C_o are sorted in decreasing order of confidence score. In each step, we show the next to be considered edge relation as purple dotted line.

In an outer loop, we identify the input points that are explained by the current set of learned primitives, remove them from the input set P , and repeat the algorithm for the remaining unclaimed points. Thus, we repartition the data based on the learned relations, and in the process improve the data partitioning that is difficult to perform just based on local reasoning. We stop once the remaining point set size is small, the potential relationship graph is empty, or when we reach a maximum number of iterations, 3 in our experiments (see Figure 10).




Figure 6: Types of basic primitives used in our method: plane, cylinder, cone, and sphere, respectively along with their attributes.

In our implementation, we use planes, cylinders, spheres, and right circular cones as basic primitives. We parameterize each primitive by an associated orientation direction and attribute(s) as follows: A plane is represented in the normal-intercept format using a direction \mathbf{n} and a signed distance from origin d ; a cylinder is represented by its axis \mathbf{n} , a point on the axis \mathbf{p} , and a radius r ; a cone is represented by its axis \mathbf{n} , its apex \mathbf{p} , and an apex angle ϕ ; a sphere by its center \mathbf{p} and radius r . Each such primitive χ_i has an associated set of points P_i that are used to clip the respective primitive. We clip a plane using the convex hull of its projected member points, we clip a cylinder or a cone by bounding the projected member points using plane(s) aligned along its respective primitive axis, and a sphere by a pair of parallel planes (see Figure 6). We next describe how we detect and align to global orientation relations by optimizing over the primitive parameters, while conforming to the input data.

3.1 Orientation alignment

a) Parallel and orthogonal relations. We now proceed as follows: (i) identify candidate relations C_o between the primitives χ_i , (ii) extract a maximal set of non-conflicting relations $C_o^* \subset C_o$, and (iii) use a constrained nonlinear optimization to align the primitives to C_o^* and the data.

i) *Candidate relations.* We create an undirected graph $G_o := (V_o, E_o)$ with each primitive χ_i represented as a node $v_i \in V_o$. We classify edge $e_{ij} \in E_o$ based on the angle $\theta_{ij} = \cos^{-1}(|\mathbf{n}_i \cdot \mathbf{n}_j|)$ between primitives χ_i and χ_j . If $\theta_{ij} \leq \pi/4$, then e_{ij} is a *parallel* relation-edge with a score $s_{ij} = -\theta_{ij}$; else, e_{ij} is an *orthogonal* relation-edge with a score $s_{ij} = (\theta_{ij} - \pi/2)$. We leave out weak edges based on a default threshold, $\pi/12$ for moderate noise level. We thus have a set of candidate relations $C_o := \{c_1, c_2, \dots\}$ sorted according to their decreasing scores.

ii) *Extracting C_o^* .* A graph is said to be *biconnected* if removing a node, called *cut-vertex*, increases the number of connected components. In the context of the graph G_o , cut-vertices represent the outlier primitives with chance alignment to isolated primitives, since outliers in this context are primitives with comparatively few relations. Unless removed, they can adversely influence subsequent optimization. We decompose G_o into bi-connected components [Hopcroft and Tarjan 1973] and operate relations in each connected component individually, henceforth referred to as C_o . Note that although there are no explicit relations across the components, the primitives are held together by a data fitting term.

Any constraint set C_o may possibly contain conflicts making it infeasible to optimize the primitives such that all the relations in C_o are simultaneously fulfilled, i.e., the aligned primitives *exactly* satisfy the relations. Since the input data is imprecise, accumulated error can lead to erroneous relations or conflicts, e.g., in course of reduction three nodes can be wrongly connected by an orthogonal and two parallel relations (see Figure 7). One could use a weighted least squares solution to approximately align to relations in C_o , but the result can easily be unsatisfactory since there might be outlier constraints, none of the relations will be truly satisfied but only approximated (see Figure 12). Instead, we directly construct a subset $C_o^* \subset C_o$, without redundancy and free of conflicts, e.g., preventing multi-edges and self-loops in the reduced relationship graph. Our construction is based on the following properties of the maximal subset C_o^* :



Figure 7: In course of reduction of graph G_o , conflicting edges can be created. In such a case, we remove the edge with lower confidence, in this example BC (3), after parallel collapse $B \rightarrow A$.

- (i) No two nodes are connected by a parallel-relation edge, i.e., all edges are orthogonal-relation edges, otherwise we can collapse nodes joined by the parallel-relation edge, and thus remove redundancy.
- (ii) A *clique* in the graph C_o^* can have a maximal size of three, since any orthogonal basis in 3D has only three elements.
- (iii) If two graph loops share an edge, say ABC and ACD, then edge BD must necessarily be a parallel edge, otherwise they are in conflict. Note this requirement is a stronger condition compared to the non-existence of a 4-clique.

These properties motivate the following algorithm: Starting with $C_o^* \leftarrow \emptyset$, we progressively add relations $c_k \in C_o$ to C_o^* according to the decreasing scores of c_k based on reduction steps (see Figure 5):

- For any parallel edge $c_k := e_{ij}$, we collapse node $v_j \rightarrow v_i$, and update the edge connectivity accordingly.
- If adding an orthogonal edge c_k results in 3-length loops ABC and ACD, then we check whether a parallel edge BD exists in the remaining C_o set. If parallel BD edge is present, we add c_k and BD to C_o^* and trigger a parallel edge collapse. Otherwise, we have a conflict and remove the lowest confidence edge involved.
- If the candidate edge c_k is in conflict with an existing edge, which is formed by result collapses in previous steps, then we ignore c_k .

The construction stops once all the relations in C_o have been considered. Figure 5 shows a typical reduction sequence.

iii) *Aligning to C_o^* .* Given a point cluster P_i along with its associated primitive χ_i , we measure the data fitting error as,

$$E_d(P_i, \chi_i) := \sum_{\mathbf{p} \in P_i} w_{\mathbf{p}} d^2(\mathbf{p}, \chi_i),$$

where $d(\mathbf{p}, \chi_i)$ measures the distance of point \mathbf{p} , with weight $w_{\mathbf{p}}$, to the primitive χ_i [Kanatani 2008]. Then the accumulated data error is simply given by $\sum_i E_d(P_i, \chi_i)$. Our goal is to minimize the data error, while aligning to the relations in C_o^* . Since we have only orthogonality constraints at this stage, alignment to any such constraint $c \in C_o^*$ between primitives χ_i, χ_j implies $g_c := \mathbf{n}_i^T \mathbf{n}_j = 0$. Additionally, for each primitive χ_i , the normal direction satisfies $\mathbf{n}_i^T \mathbf{n}_i = 1$. Thus we have the non-linear optimization with equality constraints,

$$\min_{\{\chi_i\}} \sum_i E_d(P_i, \chi_i), \text{ such that, } g_c = 0 \forall c \in C_o^*, \mathbf{n}_j^T \mathbf{n}_j = 1 \forall j. \quad (1)$$

We use a trust region method based on interior point nonlinear programming to solve the optimization [Byrd et al. 2000; Ziena 2010]. At the end of this stage, we have learned and aligned to orthogonal and parallel relations among the primitives.

b) Equal angle relations. Regular structures are common in man-made objects. Such regularity often results in equal angle constraints across primitive parts. Unlike parallel or orthogonal constraints, equal angles, however, are not known in advance, and are only recovered in course of optimization. A possible approach is to first cluster angles between primitive pairs, and then rank the primitive pairs based on the deviation of their mutual angle with the cluster peaks, similar to the parallel/orthogonal case. However, we found this approach to be unreliable due to presence of outliers, and spurious peaks in the angle space, e.g., histogram in Figure 8 shows why it is challenging to select clustering parameters in our setting. Instead, we directly detect and align to equal angle relations

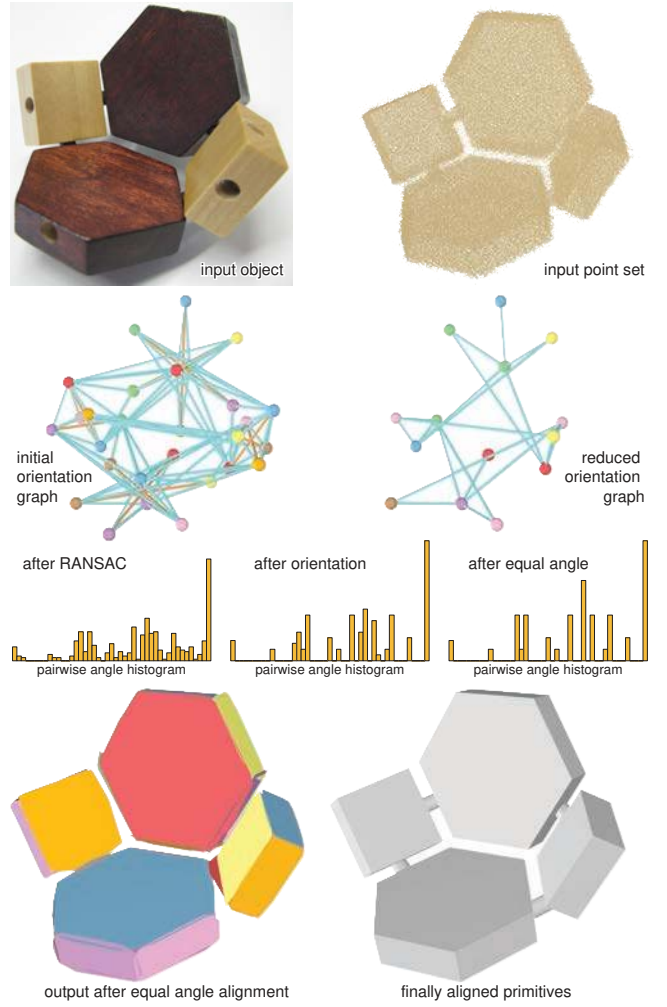


Figure 8: Starting with primitives obtained using RANSAC, orientation alignment first conforms to parallel and orthogonal relations. In addition, in the next stage, equal angle relations are learned and aligned to. The histogram of angles between primitive pairs (range $[0 - \pi/2]$) progressively gets cleaned along the stages of the algorithm.

without committing to a desired angle value, while still conforming to existing parallel/orthogonal relations.

Any angle relation involves a pair of primitive pairs, i.e., four primitives. Thus, we build a graph G_e , where each vertex represents an unordered pair of primitives. In noisy data, spatially distant primitives easily pollute the space with spurious relations. Hence, if primitives (χ_i, χ_j) are farther than 10% of the larger bounding box length, we remove its corresponding node. As a result in most cases graph G_e has only $O(m)$ nodes, instead of $O(m^2)$. For any pair of such vertices, say involving primitive pairs (χ_i, χ_j) and (χ_k, χ_l) , we add an edge based on the similarity of their pairwise angles. Note that primitives (χ_i, χ_j) and (χ_k, χ_l) can be at arbitrary distances. The corresponding relation c is assigned a score $s_c = -|\angle \mathbf{n}_i \mathbf{n}_j - \angle \mathbf{n}_k \mathbf{n}_l|$ with the associated constraint $g_c := (\mathbf{n}_i^T \mathbf{n}_j)^2 - (\mathbf{n}_k^T \mathbf{n}_l)^2 = 0$. We leave out edges if the subtended angle is more than a threshold, $5 - 10^\circ$ in our experiments. The threshold depends on the noise margin against the separation among the modes in the ground truth data, i.e., the regularity of the object. We collect the relations as $C_e = \{c_1, c_2, \dots\}$. Now starting with a set $C_e^* \leftarrow \emptyset$, we progressively add edges $c \in C_e$, based on decreasing confidence. Exact




Figure 9: Various global relations extracted and conformed to, while aligning to the scanned data of a clutch-part. In the top-right region, the primitives remain separated in absence of enough datapoints to justify a connection. Extracted relations and lengths were validated with respect to the physical model.

equality relation being transitive, the relation graph should be free of cycles. Hence, we only add $c \rightarrow C_e^*$ if no cycles are formed. This step significantly decreases the number of graph edges, while retaining the confident ones. Set C_e^* typically consists of $O(m)$ edges.

The set C_e^* can still contain conflicts, and hence we use interior point nonlinear programming to detect such cases. We test if the optimization of the data energy (Equation 1) has any feasible solution subject to the constraints $C_e^* \cup C_o^*$. If no feasible solution exists, we take out the relation $c \in C_e^*$ with the lowest score, and test for feasibility again with the remaining set of relations. Typically 1-2 such relations were dropped in our experiments.

Figure 8 shows a typical scenario for orientation optimization. The initial graph G_o has 84 edges, which gets reduced to 21 relations in C_o^* . In the equal angle stage, 33 initial candidate relations are recovered and conformed to while minimizing the data error. A feasible solution was found without discarding any relation of C_e . Total time taken was 680 seconds with most time spent in the constrained solver.

3.2 Placement alignment

In most man-made objects, coplanarity and coaxial relations carry important relation cues about the object parts (see also [Gal et al. 2009]). We conform to placement relations after orientation alignment, while preserving the orientation relations already aligned to.

Coaxial relations are simple to detect and enforce. Two primitives χ_i and χ_j , e.g., two cylinders, or a cylinder and a cone, are potentially coaxially aligned if their axes are *exactly* parallel, since we already have orientation alignment. We assign a score to such a potential primitive pair based on the distance between their axes, namely, $s_p = -\|(\mathbf{p}_i - \mathbf{p}_j)^T \cdot \mathbf{n}_i - \|\mathbf{p}_i - \mathbf{p}_j\|\|$ with $\mathbf{p}_i, \mathbf{p}_j$ being points on their respective axes. The relations are collected in a set C_p with each of the form $g_p := ((\mathbf{p}_i - \mathbf{p}_j)^T \cdot \mathbf{n}_i)^2 - \|\mathbf{p}_i - \mathbf{p}_j\|^2 = 0$. We prune edges when deviation lengths exceed 2 – 5% of point set bounding box length. Similarly, for aligning a sphere with axis of a cone or a cylinder, we simply constrain the sphere center to lie on the axis of the other primitive.

Coplanar relations exist between primitive pairs, potentially non-local ones. Since we clip cylinders and cones by respective axis




Figure 10: Our technique detects and aligns even smaller features. In the first pass, RANSAC identifies only the large primitives, which are then globally aligned. In the next pass, the points agreeing to the aligned primitives are removed, allowing RANSAC to identify the smaller primitives, which are aligned using their inter-relations, and also using the relations established earlier.

aligned planes, all coplanar relations in our framework reduce to detecting and conforming to coplanarity between two planes. Thus for a primitive pair χ_i and χ_j , with *exactly* parallel normals, we add a relation $d_i = d_j$, if the normals point in the same direction, and $d_i = -d_j$, if the normals point in opposite directions. In both cases, we assign a confidence score $s_p = -\|d_i - |d_j|\|$. The relations are appended to the set C_p . Such relations, even if they are few, group initially independent RANSAC primitives, e.g., plane primitives, potentially distant in space.

We sort the relations by their decreasing confidence scores and similar to the case of equal angle relations, extract a subset $C_p^* \subset C_p$ by aligning to the relations $C_p^* \cup C_e^* \cup C_o^*$, while minimizing the data error. During the extraction, we progressively remove low confidence relations, if necessary, until we obtain a feasible solution.

3.3 Equality alignment

After orientation and placement alignments, we detect and conform to equality relations. In our case, two types of equalities are detected: (i) For a pair of cylinder primitives χ_1 and χ_2 , with comparable radii, we add a potential constraint $g_l := r_1 - r_2 = 0$ with confidence score $s_l = -|r_1 - r_2|$. (ii) Relations between plane primitives and also between clipping planes, are more complicated as they involve relations between a pair of primitive pairs, similar to the equal-angle relations. For a pair of parallel plane pairs, $\{(\mathbf{n}_{12}, d_1), (\mathbf{n}_{12}, d_2)\}$ and $\{(\mathbf{n}_{34}, d_3), (\mathbf{n}_{34}, d_4)\}$, we introduce a constraint $g_l := (d_1 - d_2)^2 - (d_3 - d_4)^2 = 0$ with confidence score $s_l = -\|d_1 - d_2 - |d_3 - d_4|\|$ (for each g_l , we adjust signs of d_i so that all the normals point to same direction). We collect both types of relations into a set C_l , and as previously with equal angle case use transitivity to discard redundant relations, and extract a feasible set C_l^* such that the data error is minimized, while simultaneously conforming to all the constraints $C_l^* \cup C_p^* \cup C_e^* \cup C_o^*$.

3.4 re-RANSAC

In our framework, the optimization proceeds from a coarse to fine scale. Once the initial set of primitives are aligned using the detected global relations, we reassign the points in P to the aligned primitives. The points with high data error to all the primitives are collected as unclaimed points in a set P^* . Additional checks for spatial consistency are not necessary since the primitives are

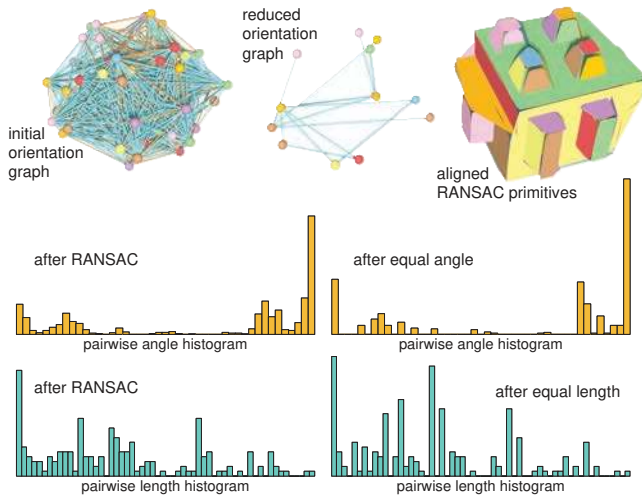


Figure 11: Result of orientation, placement, and equality, of both angle and length, on the foam-box model. Extracted relations and angles were validated against the physical object (see also Figure 1).

already clipped and bounded at this stage. We iterate through the pipeline starting with RANSAC on the unclaimed points P^* , with two modifications: (i) we lower the threshold on the allowed size of the smallest detected primitive based on the extent of connected components in P^* ; (ii) we start with a relationship set of $C_1^* := C_l^* \cup C_p^* \cup C_e^* \cup C_o^*$, instead of an empty set. Note that for global alignment, the primitives from the previous step, and their associated data points are also incorporated. For example, in Figure 10 the small cylinders are detected in the second iteration, which are missed in the first pass due to the points from the dominant surrounding regions. Although the initial parameters for the small cylinders are unstable due to their limited support, they are stabilized by aligning to the first iteration relations C_1^* .

4 Evaluation

We evaluated our framework on a variety of synthetic and scanned datasets. Where precise models were available, we compared the reconstructed models with the ground truth, both in terms of recovered face normals as well as distance between parallel primitive pairs. Where precise ground truth was unavailable, we measured the physical model for comparison.

4.1 Synthetic datasets

In the first experiment, we consider three Platonic solids, namely octahedron, dodecahedron, and icosahedron. In each case, we perturb the normals of the polyhedrons' faces and then sample the model while adding noise to each sample independently. Note that the normal perturbation mimics misalignment among multiple scans (see later discussion on slippage). For each point set, we align the initial RANSAC primitives using our global optimization. Then, we compare face normals and distance between parallel primitive pairs with those in the ground truth model; we do the same for the model using initial RANSAC primitives. For a consistent comparison, we retain only the top relevant number of RANSAC planes in each case, e.g., for octahedron, we retain the top 8 RANSAC planes. Table 1 demonstrates the stability of our method across varying amount of noise.




Figure 12: In presence of structured noise in an input scan, our method favors alignment to global priors. We compare results obtained with initial RANSAC primitives, a greedy strategy, a least squares solution, and our global alignment, respectively. We overlay ground truth (black boundary) on the results for visual evaluation.

In the second experiment, we used a standard CAD tool to create two models with regular angles and repeated parts. We scanned the objects using a virtual scanner, added noise, and compared the quality of the result with ground truth, with and without our global alignment. In Figures 3 and 12 we show the corresponding results, and for the boxy model, the recovered angle/length as compared to ground truth are listed in Table 1. In Figure 12, we compare the trimmed models obtained using initial RANSAC primitives, a greedy approach, a least squares solution, and our global alignment, respectively. In the greedy approach, starting from the best fit primitive, we progressively aligned the primitive with the next strongest parallel/orthogonal edges to the aligned ones, while keeping previously aligned primitives unchanged. In the least squares solution, we obtained a solution combining the data term with penalty energies for all the parallel/orthogonal candidate edges, weighted by their confidence. Note that, in presence of structured (biased) noise or misaligned scans, our method prefers alignment to global relations (regular configurations).

4.2 Scanned datasets

We tested our method on real-world objects of varying complexity. In absence of a precise virtual model as ground truth, we used physical measurements of the objects for comparison. We also validated orientation, coplanarity, coaxial, and equality relations detected by




Figure 13: Cylinder, plane, and sphere primitives are aligned using extracted coaxis, coplanar, parallel/orthogonal axes, equal angle as well as equal length constraints. They converge to the final model after two iterations of RANSAC fitting and constraint optimization. We overlap the final result on the initial RANSAC results for comparison; the histograms demonstrate the effect in the primitive pair angle space.

our system on the physical models, as possible.

In our simplest example of the wheel model, Figure 14, all the major axes of the planes, cylinders, and cones are detected parallel, and the reduced orientation graph is trivial in absence of any orthogonal relations. The dataset is interesting for the large number of coaxial and coplanar relations, and also for the equality of attributes for cones and cylinders. We detected all the relations (see histograms), thus producing a natural segmentation of the object based on spatial clustering of various equality relations, i.e., the corresponding subparts of the object have a common subgraph structure in the primitive pair relation graph. Note that the final model, automatically allows exploded views of the subparts, along the recovered axis directions.

Next, we take a packing foam-box (see Figures 1 and 11) with small parts sharing a variety of relations. Here the detected lengths and angles were in good agreement with the physical model. Interestingly, a protrusion on the box (right feature in Figure 1) was in the range of 0.1mm to the other flanks on (front side of) the model. All these protrusions were grouped together and made identical.

	$\theta = 2$	$\theta = 4$	$\theta = 6$	$\theta = 8$	$\theta = 10$
octahedron (nor:8)	3/8	2/8	2/8	0/8	0/8
(len:4)	2/4	0/4	0/4	0/4*	0/4*
dodecahedron (nor:12)	5/12	4/12	3/12	2/12	2/4
(len:6)	1/6	0/6	0/6	0/6*	0/2
icosahedron (nor:20)	11/20	7/20	5/20	5/10	4/7
(len:10)	5/10	3/10	1/10*	1/10*	0/10*
boxy (nor:36)	19/36	10/36	7/36	5/30	4/15
(len:66)	49/66	34/66	24/66	17/36	9/16

Table 1: Comparison of angle and length of the primitives, against the ground truth. The number of correctly detected relations across the RANSAC based primitives, in red, and among the optimized primitives, in blue, are shown. Ground truth data are also shown, e.g., nor:8 refers to 8 different normal directions in the ground truth. A number like 8 denotes 8 correctly (in 0.1° agreement or 0.5% of primitive bounding box length with ground truth) recovered normals. A number like 4* means 4 groups of primitives converged to equal lengths, but with converged value different from the ground truth; added noise was in $U[-\theta^\circ, \theta^\circ]$ for normals, and results are averaged over 10 runs.

This demonstrates our global prior that if primitives are in near alignment, they are likely to be in exact alignment.

The clutch model, Figure 9, is challenging given the small features, and the poor quality of the scan, due to the dark material color (we were not allowed to paint the model white). Even on this sparse dataset, our method correctly detected various relations, hinting at the ubiquity of relations and repetitions in man-made engineering objects. Note that in one part the sample points are almost non-existent and the method had little evidence to propagate and connect up the primitives.

In the final, but most complicated example, we scanned a machine part (see Figure 13). Various relations were identified and conformed to in the final model. The angle pair histogram clearly

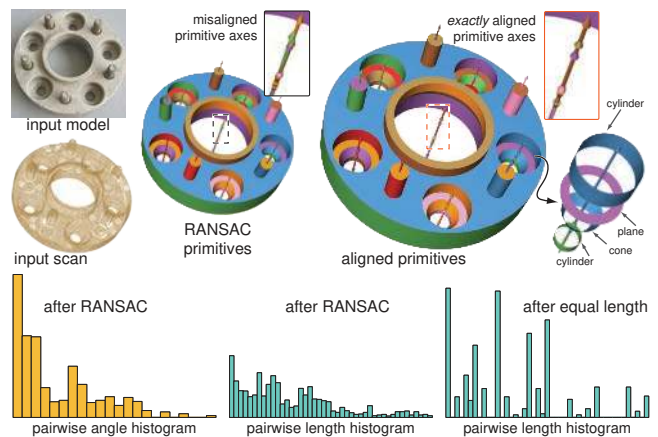


Figure 14: Global alignment of primitives involving parallel axes, coplanarity, coaxial, and equal length constraints on the wheel model, while fitting to the scanned data. In this case, the coaxial and equality relations, lead to a hierarchy based on the mutual relations and spatial proximity among the primitives. We show an exploded view of one of the subcomponents along its common axis. RANSAC primitives have pairwise angle histogram in range $[0, \pi/6]$, while after global alignment all pairwise angles vanish; the pairwise plane distance histogram also significantly improves.




Figure 15: Global relations can help register multiple scans, A and B in this example, which are otherwise difficult to align simply using local registration. We first individually align each of the scans using our framework to yield aligned models A and B . Locally registering the aligned models using ICP leaves ambiguity due to rotational slippage. We remove this ambiguity using global constraints, orthogonality between yellow-blue cylindrical hole pair in this example, learned and optimized by our method.

demonstrates the improvement in the pairwise angle relations. RANSAC failed to fit initial primitives, within fitting tolerance, for parts of the top of the model; we place a MLS surface fitted to the unclaimed point region for comparison (in slightly darker shade).

4.3 Global relations for scan alignment

Multiple scans of an object are typically aligned using local refinement methods like point-to-point and point-to-plane iterative closest pair (ICP) algorithms. A commonly encountered problem in scan alignment is *slippage*, i.e., in absence of appropriate local features, the scans slip over each other, instead of locking to a well defined configuration [Gelfand and Guibas 2004]. For example, in Figure 15, aligning source and target scans, S_1 and S_2 , using ICP results in rotational slippage of the target. We use global relations to remove the ambiguity. First, for each individual scan we learn their global relations $C_{S_1}^*$ and $C_{S_2}^*$, respectively while conforming to the underlying scans. Then, in an outer loop we extract global relations on the combined data set $S_1 \cup T(S_2)$, while optimizing over the unknown transformation T that is *restricted* to the applicable slippage transforms, in this case, rotation about the axis of S_2 . The algorithm alternates between learning relations, and solving for the aligning transform T . In each iteration, we detect the top relation candidate c , use c to compute T that minimizes data cost on $S_1 \cup T(S_2)$ subject to constraint c along with the established constraints. We update $S_2 \leftarrow T(S_2)$, and proceed with learning further relations. Finally, we perform one round of global relation alignment on the aligned data, which mainly helps in conforming to placement and equality relations in the data sets.

Performance

Table 2 lists the performance of our algorithm on various scanned models and the simulated scan on the ‘boxy’ model presented in the paper. Note that in the wheel model since all the directions get aligned to the same direction by parallel/orthogonal alignment, subsequently no equal angle relations are detected. Further, the coplanarity relations already result in equal length (trimming planes) for this example, and thus no equal length relations are detected. Note that the running time depends more on the complexity of the relational structure of the model, rather than the sampling resolution. Our graph reduction for redundant relations makes the nonlinear solver (‘fmincon’ Matlab function) robust, and also significantly faster, e.g., for the boxy model the optimization time just for parallel/orthogonal without reduction takes 164 seconds and with our

reduction only 64 seconds. For denser relationship, without redundancy removal, e.g., pipe-model with 696 original relations, the optimization may run in order of hours.

Discussion

Although model reconstruction is not our primary goal, the aligned primitives can easily be used to extract a clean model for many man-made machine parts (see Figures 1 and 12). Using the relationship graph and the underlying data, we extrapolate and compute pairwise primitive intersections as applicable. In our current implementation, for a primitive pair involving complex primitives, e.g., cylinder with cone, we use a simple numerical method to extract a polygonal intersection curve. Our method, even under irregular sampling, yields precise models along with sharp features, which is otherwise difficult with existing alternatives. Note that by aligning to global relations in conjunction with local fitting, we demonstrate that it is possible to extract the global relations even in noisy and sparse datasets. The recovered relations are naturally a sparse encoding of the input model, where we make use of the fact that such models are of low information content. Our algorithm can be seen as precisely recovering such key relations.

Unlike our method, most scan reconstruction methods use local priors, while ignoring global relations that are dominant in man-made structures. As in any shape analysis with unreliable data, there is always a tradeoff between approximation level and extent of regularity detected. Our method allows aligning relations even at the expense of slightly increased data fitting cost. However, the user can require our method to disallow large increase in the data error, and dropping the corresponding relations. It will be interesting to explore Pareto-optimal solutions in this setting. Note that aligning to regular configurations can also help in identifying defects in scanned models, and encoding shapes as deviations from regular configurations.

A possible improvement is to consider higher-order relations involving pair of triplets (or more) of primitives, but the computational complexity quickly makes a simple generalization impractical. We leave this to future research.

Limitations

For certain types of relations, we investigate pairs of primitive pairs. This can potentially lead to $O(m^4)$ relations for m primitives, especially in regular configurations, closely packed in space, resulting in significant computation times. Also when the angle or length space histograms have closely spaced values, increased noise level can submerge such differences making it difficult to correctly recover




Figure 16: (Left) Aligning to global relations corrects for significant misaligned RANSAC primitives obtained using local fits. However, when the data quality is sufficiently poor, our algorithm may fail to recover all global relations, and hence the reconstruction is only partially correct (right).

model	point #	primitive #	parallel/ortho.	eq. angle	placement	eq. length	eq. radii	time (in sec)
wheel	283K	39	741→0(38)	N/A	359→75	N/A	66→18	227.3
foam	382K	56	544→14(58)	128→37	25→18	68→20	N/A	875.6
clutch	841K	21	205→8(25)	26→17	27→18	8→3	4→3	573.1
pipe	529K	43	696→7(44)	40→21	90→43	222→19	88→27	838.5
boxy	100K	26	121→13(28)	250→44	11→9	20→8	7→4	319.3

Table 2: Performance statistics on a laptop with Intel(R) Core(TM) 2.53GHz with 2GB, for models presented in the paper. We show the number of relations before and after reduction for the various alignment stages.

the relations in our framework (see Figure 8 and Table 1). Such errors can potentially accumulate across primitives.

Chance alignments involving multiple outlier primitives may produce high confidence but at the same time incurring possible undesirable relations. If such groups are large then biconnected components cannot isolate them, and subsequently these global ‘relations’ may pollute the real constraints. Finally, our approximate solution of relation group extraction, in certain cases, can miss global relations, or even introduce outlier relations that happen to have high confidence score (see Figure 16). Although we do not expect to extract the best feasible subset given the relation to the NP-complete satisfiability problem, in practice large data error can indicate such problem. In our experiments though, we did not encounter such a situation.

5 Conclusions and Future Directions

We presented a method for incorporating global relations to reliably recover relations among both local and distant parts of man-made objects. We demonstrated that for regular objects, i.e., objects with low information content, various global relations can be reliably extracted from noisy (structured) and incomplete datasets, even when the local signals are ambiguous and unreliable. Starting with a set of initial RANSAC based locally fitted primitives, our algorithm progressively learns and infers orientation, placement, and equality relations. In each stage, a set of feasible relations are extracted among the candidate relations, and then aligned while best fitting to the input data.

In the context of low complexity, procedurally modeled man-made objects, learning global priors is likely to significantly assist in shape understanding and sparse encoding. Moreover, man-made models, e.g., Figure 2-left, commonly involve hierarchical relations. Such a grouping naturally arises from manufacturing ease and reuse options. Inferring such hierarchical grouping, however, is challenging since it will involve solving a consistent segmentation across the noisy input.

In the future, we want to consider directly integrating this stage with the global alignment leading to a robust RANSAC. More precisely, partially recovered global relations can guide an importance sampling based RANSAC fitting to robustly extract primitives from locally inconsistent and sparsely sampled regions. Bringing in similar global relations to surface reconstruction remains to be explored.

Acknowledgements

We are grateful to Ruwen Schnabel for making code from [Schnabel et al. 2007] publicly available, Ran Gal and Suhil Alsaran for proofreading the paper, KAUST car garage for lending machine parts for scanning, and AIM@Shape for the joint model used for the simulated scans in Figure 16. Niloy thanks Tanveer Alam for initial experiments for this project; Pierre Alliez, Cengiz Oztireli for running comparison tests using their state-of-the-art reconstruction algorithms; and Reinhard Klein for inspiring discussions.

The work was partially supported by KAUST international visiting-student scholarships for Li and Wu, NSFC grants 60902104, 61025012, 61003190, CAS 100 scholar program, Shenzhen Science and Technology Foundation (JC201005270340A) for Li and Wu, CAS Visiting Professorship for Senior Int’l Scientists for Cohen-Or, CAS Fellowship for Young Int’l Scientists for Sharf, Cyprus RPF GrantDiethnis/Stochos/0609 for Chrysanthou, Israel Science Foundation for Cohen-Or, and European FP7 under grant agreement 276982 for Sharf.

References

- ALEXA, M., BEHR, J., COHEN-OR, D., FLEISHMAN, S., LEVIN, D., AND SILVA, C. T. 2003. Computing and rendering point set surfaces. *IEEE TVCG* 9, 1, 3–15.
- BENKO, P., MARTIN, R. R., AND VÁRADY, T. 2001. Algorithms for reverse engineering boundary representation models. *Computer-Aided Design* 33, 11, 839 – 851.
- BOKELOH, M., WAND, M., AND SEIDEL, H.-P. 2010. A connection between partial symmetry and inverse procedural modeling. *ACM SIGGRAPH* 29, 104:1–104:10.
- BROWN, B. J., AND RUSINKIEWICZ, S. 2004. Non-rigid range-scan alignment using thin-plate splines. In *3DPVT*, 759–765.
- BYRD, R., GILBERT, J. C., AND NOCEDAL, J. 2000. A trust region method based on interior point techniques for nonlinear programming. *Mathematical Programming* 89, 1, 149–185.
- 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.
- DEY, T. 2007. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*, 1 ed. Cambridge University Press.
- FISHER, R. 2004. Applying knowledge to reverse engineering problems. *Computer-Aided Design* 36, 501–510.
- FLEISHMAN, S., COHEN-OR, D., AND SILVA, C. T. 2005. Robust moving least-squares fitting with sharp features. *ACM SIGGRAPH* 24, 544–552.
- GAL, R., SHAMIR, A., HASSNER, T., PAULY, M., AND COHEN-OR, D. 2007. Surface reconstruction using local shape priors. In *Symp. on Geometry Processing*, 253–262.
- GAL, R., SORKINE, O., MITRA, N. J., AND COHEN-OR, D. 2009. iWIRES: An analyze-and-edit approach to shape manipulation. *ACM SIGGRAPH* 28, 3, #33, 1–10.
- GELFAND, N., AND GUIBAS, L. J. 2004. Shape segmentation using local slippage analysis. In *Symp. on Geometry Processing*, 214–223.

- GUENNEBAUD, G., AND GROSS, M. 2007. Algebraic point set surfaces. *ACM SIGGRAPH* 26, 3, 23.
- HOPCROFT, J., AND TARJAN, R. 1973. Algorithm 447: efficient algorithms for graph manipulation. *Comm. of the ACM* 16 (June), 372–378.
- HOPPE, H., DEROSE, T., DUCHAMP, T., McDONALD, J., AND STUETZLE, W. 1992. Surface reconstruction from unorganized points. In *Proc. of ACM SIGGRAPH*, 71–78.
- KANATANI, K. 2008. Statistical optimization for geometric fitting: Theoretical accuracy bound and high order error analysis. *Int. Journal of Computer Vision* 80, 167–188.
- KAZHDAN, M., BOLITHO, M., AND HOPPE, H. 2006. Poisson surface reconstruction. In *Symp. on Geometry Processing*, 61–70.
- LI, G., LIU, L., ZHENG, H., AND MITRA, N. J. 2010. Analysis, reconstruction and manipulation using arterial snakes. *ACM SIGGRAPH ASIA* 29, 152:1–152:10.
- LIPMAN, Y., COHEN-OR, D., LEVIN, D., AND TAL-EZER, H. 2007. Parameterization-free projection for geometry reconstruction. *ACM SIGGRAPH* 26, 3, 22.
- MEHRA, R., ZHOU, Q., LONG, J., SHEFFER, A., GOOCH, A., AND MITRA, N. J. 2009. Abstraction of man-made shapes. *ACM SIGGRAPH ASIA*, 137:1–137:10.
- MITRA, N. J., AND PAULY, M. 2008. Symmetry for architectural design. In *Advances in Architectural Geometry*, 13–16.
- MITRA, N. J., GUIBAS, L., AND PAULY, M. 2007. Symmetrization. In *ACM SIGGRAPH*, vol. 26, #63, 1–8.
- MULLEN, P., GOES, F. D., DESBRUN, M., COHEN-STEINER, D., AND ALLIEZ, P. 2010. Signing the unsigned: Robust surface reconstruction from raw pointsets. In *Symp. on Geometry Processing*, 1733–1741.
- OZTIRELI, C., GUENNEBAUD, G., AND GROSS, M. 2009. Feature preserving point set surfaces based on non-linear kernel regression. *Computer Graphics Forum* 28, 2, 493–501.
- PAULY, M., MITRA, N. J., WALLNER, J., POTTMANN, H., AND GUIBAS, L. 2008. Discovering structural regularity in 3D geometry. *ACM SIGGRAPH* 27, 3, #43, 1–11.
- SCHNABEL, R., WAHL, R., AND KLEIN, R. 2007. Efficient ransac for point-cloud shape detection. *Computer Graphics Forum* 26, 2, 214–226.
- SCHNABEL, R., DEGENER, P., AND KLEIN, R. 2009. Completion and reconstruction with primitive shapes. *CGF Eurographics* 28, 2, 503–512.
- THRUN, S., AND WEGBREIT, B. 2005. Shape from symmetry. In *ICCV*, 1824–1831.
- WERGHI, N., FISHER, R., ASHBROOK, A., AND ROBERTSON, C. 2002. Shape reconstruction incorporating multiple nonlinear geometric constraints. *Constraints* 7, 117–149.
- ZIENA, 2010. Knitro optimization software, Dec.