

Example-Based 3D Scan Completion

M. Pauly*
Stanford University

N. J. Mitra†
Stanford University

J. Giesen‡
ETH Zürich

L. Guibas§
Stanford University

M. Gross¶
ETH Zürich

Abstract

Optical acquisition devices often produce noisy and incomplete data sets, due to occlusion, unfavorable surface reflectance properties, or geometric restrictions in the scanner setup. We present a novel approach for obtaining a complete and consistent 3D model representation from such incomplete surface scans, using a database of 3D shapes to provide geometric priors for regions of missing data. Our method retrieves suitable context models from the database, warps the retrieved models to conform with the input data, and consistently blends the warped models to obtain the final consolidated 3D shape. We define a shape matching penalty function and corresponding optimization scheme for computing the non-rigid alignment of the context models with the input data. This allows a quantitative evaluation and comparison of the quality of the shape extrapolation provided by each model. Our algorithms are explicitly designed to accommodate uncertain data and can thus be applied directly to raw scanner output. We show on a variety of real data sets how consistent models can be obtained from highly incomplete input. The information gained during the shape completion process can be utilized for future scans, thus continuously simplifying the creation of complex 3D models.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—curve, surface, solid, and object representations

Keywords: shape completion, non-rigid alignment, 3D acquisition, hole filling, surface reconstruction

1 Introduction

3D shape acquisition has become a major source for the generation of complex digital 3D models. Numerous scanning systems have been developed in recent years, including low-cost optical scanners that allow 3D data acquisition on a large scale. Obtaining a complete and consistent 3D model representation from acquired surface samples is still a tedious process, however, that can easily take multiple hours, even for an experienced user. Significant manual assistance is often required for tasks such as scan path planning, data cleaning, hole filling, alignment, and model extraction. Knowledge about the acquired shape gained in this process is typically

not utilized in subsequent scans, where the same time consuming procedure has to be repeated all over again.

Our goal is to simplify the model creation process by exploiting previous experience on shapes stored in a 3D model database. This allows the generation of clean and complete 3D shape models even from highly incomplete scan data, reducing the complexity of the acquisition process significantly. The main idea is to mimic the experienced-based human approach to shape perception and understanding. Humans have the intuitive capability to quickly grasp a perceived 3D shape, even though only a fraction of the actual geometric data is available to the eye. This is possible because we make extensive use of prior knowledge about shapes, acquired over years of experience. When seeing an object, we immediately put it into context with other similar shapes that we have previously observed and transfer information from those shapes to fill missing parts in the perceived object. In digital 3D shape acquisition, we are faced with a similar problem: Most optical acquisition devices will produce incomplete and noisy data due to occlusions and physical limitations of the scanner. How can we obtain a complete and consistent representation of the 3D shape from this acquired data?

One approach is to apply low-level geometric operations, such as noise and outlier removal filters [Taubin 1995], [Jones et al. 2003], [Fleishman et al. 2003], [Weyrich et al. 2004] and hole-filling techniques based on smooth extrapolations, e.g., [Davis et al. 2002], [Verdera et al. 2003], [Liepa 2003], [Clarenz et al. 2004]. These methods are successful in repairing small deficiencies in the data, but have difficulties with complex holes or when large parts of the object are missing. In such cases, trying to infer the correct shape by only looking at the acquired sample points quickly becomes infeasible. A common way to address this ill-posed problem is to use an explicit prior in the form of a carefully designed template model. The prior is aligned with the acquired data and holes are filled by transferring geometric information from the warped template. The cost of designing the template model is quickly amortized when a whole set of similar objects is digitized, as has been demonstrated successfully with human heads [Blanz and Vetter 1999], [Kähler et al. 2002], [Blanz et al. 2004] and bodies [Allen et al. 2003].

We extend this idea to arbitrary shapes by replacing a single, tailor-made template model with an entire database of 3D objects. This allows shape completion by *combining* geometric information from different context models. To successfully implement such a system, however, we need to address the following issues: How can we extract models from the database that provide a meaningful shape continuation in regions of missing data? How can we compute and consistently evaluate local shape deformations that align the context models with the acquired data? How can we select among multiple database models the ones that provide the most adequate shape completion in different regions of the scan? And finally, how can we blend geometric information from different models to obtain a complete and consistent representation of the acquired object?

In this paper we mainly focus on the last three aspects, using a database retrieval method that relies on existing techniques adapted to our application. We define a shape similarity metric and corresponding optimization scheme for aligning context models with the acquired data set and present a method for segmenting and blending contributions from different models, while preserving appropriate continuity constraints. Our algorithms are designed to work

*mapauly@stanford.edu

†niloy@stanford.com

‡giesen@inf.ethz.ch

§guibas@cs.stanford.edu

¶grossm@inf.ethz.ch

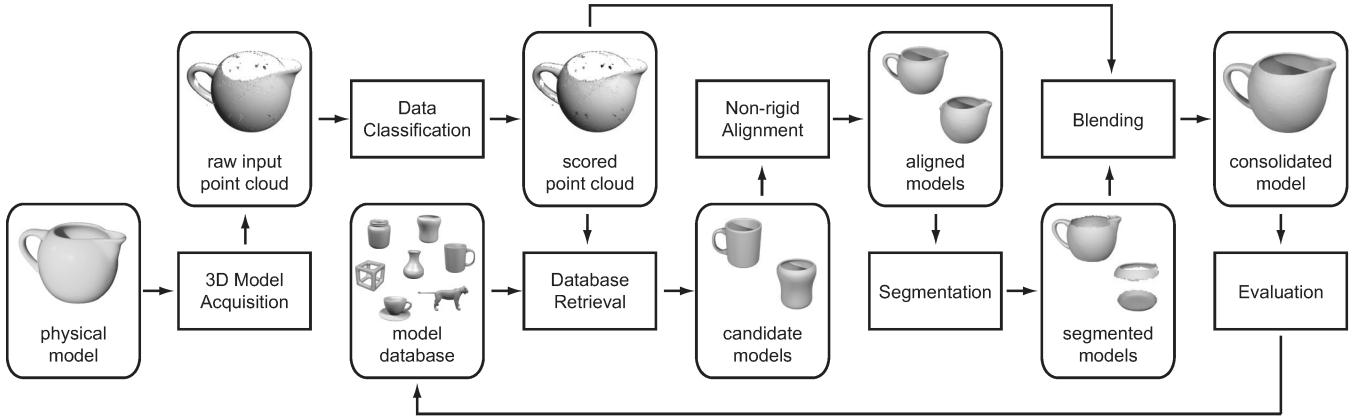


Figure 1: High-level overview of our context-based shape completion pipeline.

with uncertain data and can thus be applied directly to raw scanner output. We show how these methods can be integrated to yield a complete context-based shape completion pipeline.

1.1 Related Work

Spawned by a proliferation of commercially available 3D acquisition devices, surface reconstruction from point samples has become an active area of research over the last few years. Early methods based on signed distance field interpolation have been presented by [Hoppe et al. 1992] and [Curless and Levoy 1996]. Voronoi-based approaches with provable guarantees were introduced by [Amenta et al. 1998] and [Dey and Goswami 2003], who extended their work in [Dey and Goswami 2004] to handle noisy input data. While these techniques can handle small holes and moderately under-sampled regions in the input data, they are less suited for data sets with large holes that often occur in 3D data acquisition. To address this problem, various methods for model repair based on smooth extrapolation have been proposed. [Carr et al. 2001] and [Ohtake et al. 2003] use radial basis functions to extract manifold surfaces from incomplete point clouds. [Davis et al. 2002] presented an algorithm that applies volumetric diffusion to a signed distance field representation to effectively handle topologically complex holes. An extension of this idea has been proposed by [Verdera et al. 2003], who use partial differential equations to evolve the distance field, similar to inpainting techniques used for images. [Sharf et al. 2004] presented a system that preserves high-frequency detail by replicating local patches within the acquired 3D data set, as an extension of the 2D method proposed in [Drori et al. 2003]. Other approaches for shape completion are the triangulation-based method proposed by [Liepa 2003] and the system based on finite elements presented in [Clarenz et al. 2004].

The underlying assumption in all of these methods is that an appropriate shape continuation can be inferred from the acquired sample points only, using generic smoothness or self-similarity priors for missing parts of the model. Since the surface reconstruction problem is inherently ill-posed, the use of explicit template priors has been proposed by various authors. [Ramamoorthi and Arvo 1999] presented a method to recognize and fit a parametric spline surface to acquired surface data. Template-based hole filling has also been used in [Blanz and Vetter 1999], [Kähler et al. 2002], [Blanz et al. 2004] and [Allen et al. 2003], where input data and morphable template model were represented as triangle meshes. These methods are well-suited for object classes with well-defined shape variability, where a single template model can be adjusted to fit the entire acquired data set. Our approach differs in that we are not assuming

a priori knowledge of the underlying shape space, but try to infer automatically how to combine different context models that are retrieved from a database for the specific input data produced by the scan.

A central component of our system is a method for computing non-rigid alignments of database models with the acquired input data. [Allen et al. 2003] and [Sumner and Popovic 2004] presented alignment algorithms similar to ours that use an optimization framework to compute a smooth warping function. We extend this scheme to allow a quantitative comparison of the quality of the alignment across different models, which has not been a concern in previous methods. An interesting alternative has been proposed by [Anguelov et al. 2004], who introduced a probabilistic scheme for unsupervised registration of non-rigid shapes.

Our system bears some resemblance to the shape modeling system proposed by Funkhouser et al. [2004], where the user can create new models by cutting and pasting rigid parts of existing shapes retrieved from a shape database. The focus of their work is on creative design and interaction, while we concentrate on model repair and shape completion of surface scans.

1.2 Overview

Figure 1 gives an overview of our shape completion pipeline. 3D acquisition devices typically produce a set P of possibly noisy point samples $\mathbf{p}_i \in \mathbb{R}^3$ that describe (parts of) a 2D boundary surface of a 3D object. We assume that P is equipped with approximate normals, which are commonly provided by the scanner, or can be estimated directly from the point samples [Hoppe et al. 1992]. This input point cloud is pre-processed using multi-scale analysis to obtain a scalar confidence estimate that quantifies the consistency of each sample point with its neighboring samples. Subsequent stages of the pipeline will take these confidence weights into account to adapt the processing to the uncertainty in the acquired data. In the next stage, we retrieve a small set of candidate models from the database using a combination of multiple retrieval methods. The candidate models are then warped to match the shape of the input point cloud. To compute this non-rigid alignment we use an optimization process that balances geometric error, distortion of the deformation, and semantic consistency defined by a small set of feature correspondences. We segment the warped models into parts that best correspond to the input data based on a local shape similarity metric. Context information is propagated into regions of missing data, while continuously updating the alignment to ensure consistency between different context models. The segments are then combined using a geometric stitching technique that blends



Figure 2: Left: Acquisition setup with Cyberware Desktop 3D Scanner 15 and physical model, right: Raw point cloud obtained from six range images in a single rotational scan. The spiky outliers are artifacts caused by specular reflections.

adjacent parts from different models to avoid visual discontinuities along the seams. If a successful shape completion has been obtained, the final model is entered into the database for future use as a context model.

The following sections discuss these individual stages in more detail, using the data set shown in Figure 2 to illustrate the complete shape completion pipeline. Note that in this model almost half of the surface geometry is missing due to occlusions, including the bottom and interior of the pot, as well as parts of the handle. Hole-filling techniques based on extrapolation would not be able to recover the correct shape from this data, but would require a significantly more complex scanning procedure with multiple scans of different poses of the object.

2 Data Classification

As illustrated in Figure 2, the acquired sample set P is inherently unreliable and cannot be treated as ground truth. Noise and outliers introduce uncertainty that needs to be considered when trying to reconstruct a consistent model. If available, we use confidence information provided by the scanning device to estimate the saliency of each sample point. We found it necessary, however, to include additional local consistency information to make the subsequent stages more robust. We compute per-point confidence estimates as a combination of two local geometry classifiers that analyze the distribution of samples within a small sphere centered at each sample point. The first classifier $c_i^\lambda \in [0, 1]$ measures the quality of fit of a local tangent plane estimate at $\mathbf{p}_i \in P$, while the second classifier $c_i^\sigma \in [0, 1]$ analyzes the uniformity of the sampling pattern to detect hole boundaries (see Appendix). The combination of both classifiers yields the confidence estimate $c_i = c_i^\lambda \cdot c_i^\sigma \in [0, 1]$, which we evaluate at multiple scales by varying the size of the local neighborhood spheres. Similar to [Pauly et al. 2003a], we look for distinct local maxima of c_i across the scale axis to automatically determine the appropriate scale at each sample point (see also Lindeberg's paper [1998] on automatic scale selection). For all examples in this paper we use ten discrete scale values, uniformly distributed between $2h$ and $20h$, where h is the minimum sample spacing of the scanning device. Figure 3 shows the results of this multi-scale classification.

3 Database Retrieval

To transfer geometric information from the knowledge database to the acquired object, we need to identify a set of candidate models M_1, \dots, M_n that are suitable for completing the input data P . Database retrieval of 3D objects has gained increasing attention in recent years and a variety of shape descriptors have been proposed to address this problem (see [Tangelder and Veltkamp 2004] for a recent

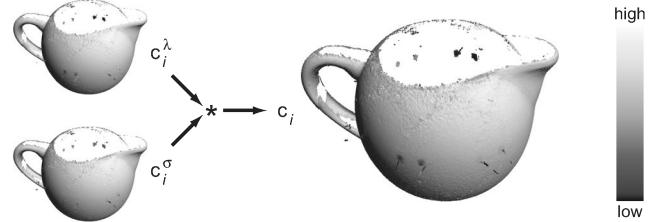


Figure 3: Quality of fit estimate c_i^λ and local uniformity estimate c_i^σ are combined to yield the final confidence estimate c_i .

survey). In our case the retrieval problem is particularly difficult, since we are dealing with noisy, potentially highly incomplete data. We thus rely on a combination of textual search and shape-based signatures, similar to [Funkhouser et al. 2004]. We first confine the search space using a few descriptive keywords provided by the user. On this restricted set of models, we compute a similarity measure based on point-wise squared distances. We use PCA to factor out global scaling and estimate an initial pose. Then we apply the alignment method proposed in [Mitra et al. 2004] to optimize the rigid part R of the transform by minimizing the sum of squared distances $E(M, P)$ between acquired shape P and a database model M given as

$$E(M, P) = \sum_{i \in P} c_i \|R\mathbf{p}_i - \mathbf{q}_i\|^2, \quad (1)$$

where \mathbf{q}_i is the closest point on M from $R\mathbf{p}_i$. Note that the squared distances are weighted by the confidence estimate c_i to avoid false alignments due to outliers or noisy samples. Equation 1 can be evaluated efficiently by pre-computing the signed distance field of each database model as described in detail in [Mitra et al. 2004]. In our current system, database models are represented as triangle meshes, though any other surface representation that allows closest point queries can be used also.

The residual of the optimization is used to rank the retrieved context models. Objects that align well with the acquired data are likely candidates for a successful shape completion and will thus be given a high score. Figure 4 shows the results of the database retrieval for the coffee creamer example.

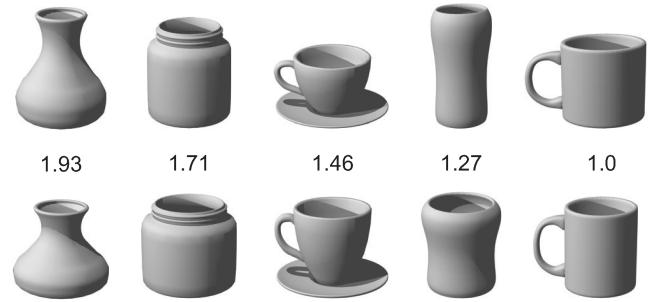


Figure 4: Models retrieved by a combination of geometric and textual query ordered from left to right according to decreasing relative alignment error. The top row shows the original models, the bottom row shows the aligned models after re-scaling.

4 Non-rigid Alignment

The global similarity transform computed in the retrieval stage will in general not align the extracted context models exactly with the acquired data. We thus need to deform each model M before transferring shape information from M to P . The goal is to find a

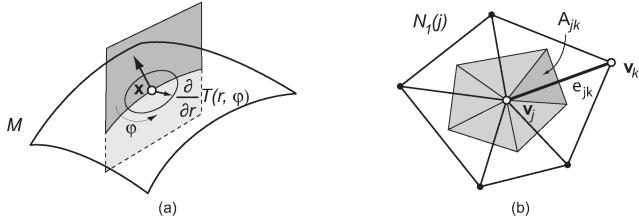


Figure 5: Measuring distortion for a continuous surface (a) and in the discrete setting (b). The shaded region in (b) shows the area A_j of the restricted Voronoi cell of \mathbf{v}_j .

smooth warping function $T : M \rightarrow \mathbb{R}^3$ such that the deformed model $M' = T(M)$ matches P . At the same time we want the distortion of M induced by T to be as small as possible. The idea is that if only a small deformation is necessary to align a model with the acquired data set, then this model is more likely to provide a meaningful continuation of the shape in regions of missing data. We capture this intuition by defining a shape matching penalty function Ψ that combines the distortion of the transform and the geometric error between warped model and input data. The optimal warping function T can then be determined by minimizing this function [Allen et al. 2003], [Sumner and Popovic 2004].

Similar to the rigid alignment computed in Section 3, we use the residual of the optimization to evaluate the quality-of-fit of each database model. Hence the shape penalty function should be compatible across different context models to allow a quantitative comparison between the models. Additionally, we need to be able to do this comparison locally, so that we can determine which model best describes the acquired shape in a certain region of the object.

4.1 Distortion Measure

To meet the above requirements and make the penalty function independent of the specific discretization of a context model, we derive the distortion measure for discrete surfaces from the continuous setting (see also [Levy 2001]. Let S be a smooth 2-manifold surface. We can measure the distortion $\Phi(S, T)$ on S induced by the warping function T as

$$\Phi(S, T) = \int_S \int_\varphi \left(\frac{\partial}{\partial r} T_\mathbf{x}(r, \varphi) \right)^2 d\varphi d\mathbf{x}, \quad (2)$$

where $T_\mathbf{x}(r, \varphi)$ denotes a local parameterization of T at \mathbf{x} using polar coordinates (r, φ) . The inner integral measures the local distortion of the mapping T at \mathbf{x} by integrating the squared first derivative of the warping function in each radial direction (see Figure 5 (a)).

Since we represent database models as triangle meshes, we approximate T as a piecewise linear function by specifying a displacement vector \mathbf{t}_j for each vertex $\mathbf{v}_j \in M$. The angular integral in Equation 2 is discretized using a set of normal sections defined by the edges $\mathbf{e}_{jk} = \mathbf{v}_j - \mathbf{v}_k$, where $k \in N_1(j)$ with $N_1(j)$ the one-ring neighborhood of vertex \mathbf{v}_j . We approximate the first derivative of T using divided differences, which yields the discrete version of the distortion measure $\Phi(M, T)$ as

$$\Phi(M, T) = \sum_{j \in M} \sum_{k \in N_1(j)} A_{jk} \left(\frac{\mathbf{t}_j - \mathbf{t}_k}{|\mathbf{e}_{jk}|} \right)^2. \quad (3)$$

As shown in Figure 5 (b), A_{jk} is the area of the triangle defined by \mathbf{v}_j and the Voronoi edge dual to \mathbf{e}_{jk} in the Voronoi diagram of P restricted to $N_1(j)$. Note that $A_j = \sum_{k \in N_1(j)} A_{jk}$ is the area of the

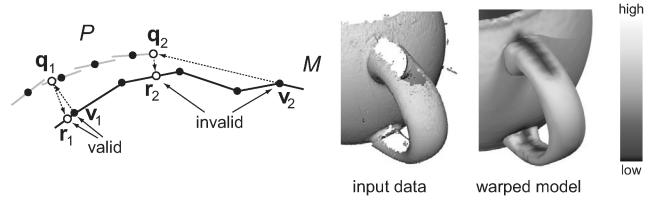


Figure 6: Correspondence weights v_j are determined using a bidirectional closest point search.

Voronoi cell of \mathbf{v}_j restricted to M , hence the surface area of M is given as $A_M = \sum_j \sum_k A_{jk}$.

4.2 Geometric Error

If P was a perfect sample, we would want to find the warping function T that minimizes $\Phi(M, T)$ under the constraint that $T(M)$ interpolates P . However, since we are dealing with uncertain data obtained in a physical measurement, enforcing interpolation can lead to highly distorted context models that do not describe the acquired shape adequately. We thus drop the interpolation constraint and instead define a geometric penalty function Ω that measures the deviation of the deformed model from the input sample. For two smooth surfaces S_1 and S_2 , we can define the squared geometric distance of S_1 to S_2 as

$$\Omega(S_1, S_2) = \int_{S_1} d(\mathbf{x}, S_2)^2 d\mathbf{x}, \quad (4)$$

where $d(\mathbf{x}, S_2)$ is the distance of a point $\mathbf{x} \in S_1$ to the closest point in S_2 . To discretize this equation we represent the surface S_P defined by P as a collection of tangent disks attached to each sample $\mathbf{p}_i \in P$. The orientation of the disk is given by the normal at \mathbf{p}_i and its radius is determined from the size of a local k -neighborhood, similar to [Pauly et al. 2003b]. We can then approximate the geometric error by summing up the area-weighted squared distance of each transformed vertex $\mathbf{v}_j + \mathbf{t}_j$ of M to the closest compatible point \mathbf{q}_j on S_P , leading to

$$\Omega(P, M, T) = \sum_{j \in M} \omega_j A_j \|\mathbf{v}_j + \mathbf{t}_j - \mathbf{q}_j\|^2. \quad (5)$$

By compatible we mean that the normal of the tangent disk of \mathbf{q}_j deviates less than 90° from the normal of \mathbf{v}_j to avoid matching front- and back-facing parts of both surfaces. ω_j is an additional weight defined as the product of two terms: The confidence estimate c_j of the sample point in P associated with \mathbf{q}_j , and a correspondence weight v_j that quantifies the validity of the correspondence between \mathbf{v}_j and \mathbf{q}_j . Since P can be incomplete and M might contain parts that do not match with the acquired object, we need to discard vertices of M from contributing to the geometric error, if no valid correspondence with the samples of P can be established. We define the correspondence weight using a simple, but effective heuristic: Let \mathbf{r}_j be the closest point on the surface of model M from \mathbf{q}_j . If \mathbf{v}_j and \mathbf{r}_j are close, then we have a strong indication that the correspondence is valid. We thus set $v_j = e^{-\|\mathbf{v}_j - \mathbf{r}_j\|^2/h^2}$, where h is the average local sample spacing of P (see Figure 6).

4.3 Optimization

We combine the distortion metric Φ and the geometric error Ω to define the shape matching penalty function Ψ as

$$\Psi(P, M, T) = \alpha \cdot \Phi(M, T) + (1 - \alpha) \cdot \Omega(P, M, T), \quad (6)$$

where α is a parameter that allows to balance distortion and geometric error. The warping function T is then computed iteratively

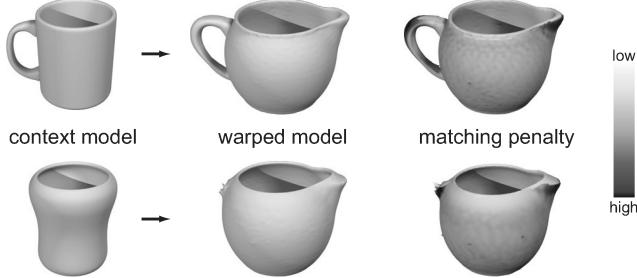


Figure 7: Non-rigid alignment. Points without valid correspondence are colored in gray in the images on the right.

by minimizing Ψ with respect to the unknown displacement vectors \mathbf{t}_j . This yields a sparse linear system of size $3n \times 3n$, where n is the number of vertices in the mesh, that we solve using a sparse matrix conjugate gradient solver.

We use a multi-level optimization scheme similar to the method proposed by [Allen et al. 2003] and later adopted by [Sumner and Popovic 2004]. There are a number of differences that we believe are interesting to discuss. In the previous methods the deformation is represented by specifying an affine transform per vertex, whereas we approximate the warping function directly using displacement vectors. This leads to significant performance improvements, since linear system is reduced from $12n$ to $3n$ unknowns. More importantly, defining a compatible distortion metric directly on the unknowns is straightforward using the discretization of the area integral of the local stretch measure (Equation 3). In the affine case this is more involved, since the same deformed mesh can be obtained via multiple different local transforms. On the other hand using affine transforms has the advantage of better handling rotations and scaling. Since we factor out global rotations and scaling during the initial alignment, we found this to be less critical, however. Figure 7 shows the two warped context models for the coffee creamer example.

Feature Correspondences. To avoid local minima in the optimization of Equation 6, we adapt the geometric penalty to include a small set $F \subset M$ of user-specified feature points. The user explicitly defines this set by selecting vertices of M together with corresponding points on S_P . The influence of each feature vertex $\mathbf{v}_j \in F$ can be controlled by scaling the weight ω_j in Equation 5. Explicit feature points are crucial for models for which the correct correspondence cannot be derived with the purely geometric approach of Section 4.2. An example is shown in Figure 10, where the semantics of each part of the models is clearly defined and needs to be observed by the warping function. Feature points also provide a mechanism for the user to control the non-rigid alignment for difficult partial matches as shown in Figure 11. Similar to [Allen et al. 2003], we start the optimization with high feature weights and strong emphasis on the smoothness term to obtain a valid initial alignment on a low resolution model. At higher resolutions we decrease the influence of the feature points and steadily increase α to about 0.9 so that the geometric error dominates in the final alignment.

5 Segmentation

After non-rigid alignment, we now need to determine how to compose the final model from different parts of the warped context models. In particular, we need to decide which model provides the most adequate shape continuation in regions of missing data. This decision is based on the matching penalty Ψ computed during the alignment stage, as it provides a measure of how well the deformed

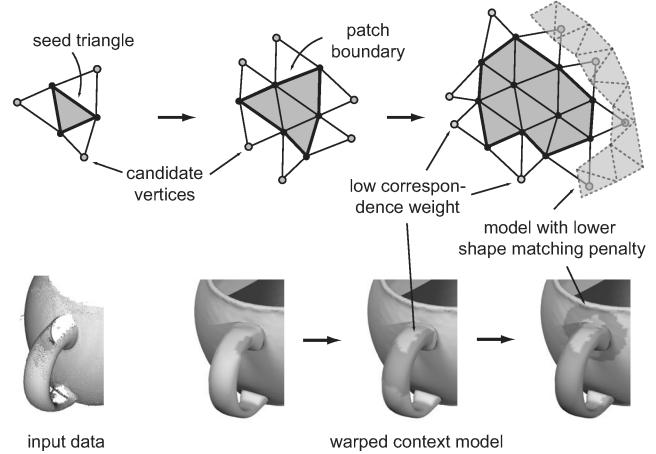


Figure 8: Incremental region growing for segmenting the context models.

database models approximate the shape of the acquired object. We first compute a segmentation of the context models into patches that cover the input point cloud in regions of high data confidence. In the next section we will describe how to extrapolate geometric information from these patches to consistently fill in missing regions and obtain a complete model representation.

The initial segmentation is computed using an incremental region growing process as shown in Figure 8. Starting from a seed point $\mathbf{p}_i \in P$, we determine which model best matches the acquired data in the vicinity of that point by evaluating the matching penalty on a small patch around \mathbf{p}_i . The model M_k with the smallest local penalty will be our candidate for this region. We then successively expand this patch by adding triangles adjacent to the patch boundary. The incremental growth is stopped wherever we encounter a different model M_l with a smaller matching penalty, indicating that this model provides a better representation of the acquired shape in that region. To evaluate this criterion we require a mapping between candidate models, which we establish using the correspondence computed in the alignment stage. For each new candidate vertex $\mathbf{v}_j \in M_k$ adjacent to the current patch boundary, we look at the corresponding point $\mathbf{q}_j \in S_P$ used to compute the geometric error in Equation 5. We then find all vertices in M_l that were mapped to points in the vicinity of \mathbf{q}_j and compare the matching penalty of these vertices with the one at \mathbf{v}_j . If a vertex with a smaller value is found, the triangle of \mathbf{v}_j will not be added to the patch. We also discard this triangle if the correspondence weight v_j (see Section 4.2) is low, indicating that we have reached a hole boundary in the data. The growth of a patch terminates as soon as no more candidate vertices can be added, as illustrated in Figure 8.

We seed the patch creation by maintaining a priority queue of all samples $\mathbf{p}_i \in P$ with high confidence (we use the top 5 percent of samples in all our examples) that have not yet been visited. The queue is sorted according to decreasing confidence c_i such that samples with high confidence will be used first to initiate a new patch. The region growing is terminated once the queue is empty.

6 Blending

The segmentation of the warped context models provides a suitable representation of the scanned object in regions of high data confidence. To fill in parts where no reliable samples could be acquired, we need to extrapolate geometric information from the context patches. Filling holes is straightforward if only one candidate

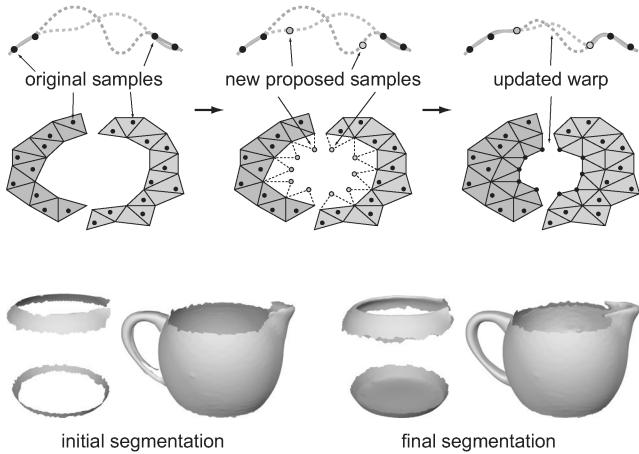


Figure 9: Blending. Top row from left to right: Two patches from different models meet at a hole boundary, new sample points are added from the model with lower shape matching penalty, both models are re-aligned with the enhanced point cloud and patches are enlarged. The top image shows a 2D illustration of the warped context models, the bottom image shows the current patches. Bottom row: segmentation before and after blending. Back-facing triangles are colored.

model covers the entire boundary of a hole. For example, the hole on the top of the creamer's handle is entirely enclosed by a single patch from the warped cup model (see Figure 8). We can thus simply copy and paste the corresponding surface part of that model.

The situation is more complicated when two or more models meet at a hole boundary. Copying and pasting parts of each model will not yield a consistent surface, since the candidate models do not agree away from the input data. Even if a cut along an intersection curve can be found, an unnatural crease might be created that causes visual artifacts. To address this issue we propose an incremental blending method illustrated in Figure 9. Starting from the initial patch layout computed in the segmentation stage, we successively add samples to the input data by copying vertices from the patch boundaries of the segmented context models. These newly added sample points represent the continuation of the data surface at a hole boundary, as suggested by the best matching model in that region. We then re-compute the warping function for all database models to conform with this enhanced point set. Since the previous alignment provides a very good initial guess, only a few iterations of the optimization are required. After updating the alignment, we enlarge the context patches using the region growing algorithm described in the previous section. We repeat this procedure until the patch growing terminates, indicating that all the holes have been closed.

Stitching. The patch layout now provides the necessary pieces to compose the final model. We enlarge each patch by adding triangles along the patch boundary to create a smooth and seamless transition between adjacent patches. We achieve this blend by applying the same optimization as in the non-rigid alignment stage described in Section 4, except that we do not warp the models towards the input point cloud, but towards each other. Consider the example shown in Figure 13. As shown on the left, the two patches from the vase and the cup do not match exactly in the region of overlap. We therefore compute a warping function T_1 that aligns the vase with the cup and a warping function T_2 that aligns the cup with the vase, and apply half of each transform to the corresponding model. A few iterations of this process create a conforming overlap region



Figure 13: Stitching, from left to right: initial configuration, two intermediate steps of alignment, final stitched model.

between the two patches. We then use the stitching method of [Turk and Levoy 1994] to obtain a single manifold surface.

7 Results and Discussion

We have tested our model completion pipeline on a number of acquired data sets with significantly different shape characteristics. All examples contain large, complex holes due to occlusions, grazing angles, or specular reflections. Repairing these models without context information from the database would require substantial manual intervention using geometric modeling tools, since model completion techniques based on smooth extrapolation would not be able to create a consistent model.

Figure 14 shows the final reconstruction of the coffee creamer example. Note how the characteristic features of the model are faithfully recovered and different parts of the two database models are blended in a natural way without any visual seams. The deformation of the context models even captures the spout, which is not present initially in any of the two models. However, in regions of insufficient input data, e.g., around the rim or at the top of the handle, the reconstructed model clearly exhibits characteristics of the context models. Apart from specifying optimization parameters and keywords for the textual search, this example requires no further user interaction. In particular, no feature points need to be specified to guide the alignment process. This leads to an overall processing time of less than two minutes.



Figure 14: Reconstructed coffee creamer, from left to right: Physical model, acquired data set, reconstructed model.

On the other hand, precise matching of feature points is crucial in the example of Figure 10. Since the semantics of every part of the animals' bodies are well defined, even slight misalignments would be clearly noticeable. The acquired giraffe data set has been completed with parts of the horse, camel, and lion. After computing the non-rigid alignment using 40 manually specified feature correspondences, the automatic segmentation and blending methods create a faithful reconstruction of the giraffe model. This example clearly demonstrates the advantages of combining context information from different models, since a satisfactory shape completion could not be obtained from any of the deformed context models alone.

Figure 15 illustrates how shape completion is continuously simplified by enriching the database with already acquired and consolidated models. The two giraffes are completed using the model of Figure 10 as a context model. Even though the input data is noisy and consists of multiple, imperfectly aligned scans, a high-quality reconstruction is obtained. Note that due to size restrictions

of the scanner, no data has been acquired for the head of the tall giraffe, hence its orientation is not captured correctly. This example demonstrates how our system can be used in connection with low-cost scanning devices to quickly create detailed 3D models that can readily be used in interactive applications such as games.

A more complex example is shown in Figure 11. The input data is a single range image that contains large, complex holes due to occlusion. The two pillars shown in 3 and 4 are used as context models to repair the highly incomplete lower sections of the wall. The panels on the ceiling are completed successively using multiple iterations of our pipeline. The first panel on the lower left is repaired using a simple plane as a geometric prior. The consolidated panel is then used to fix the other panels in this arch. Once the whole arch is completed, it can be extracted to be used as a context model for the right arch. Note that the panels are not exact copies of each other, so simple copy and paste operations will not yield adequate results. User assistance is required to select appropriate parts in the data that can be used as context models for other regions, and to provide an initial alignment for those parts using four feature correspondences per piece. Interaction time for a trained user is less than half an hour, compared to multiple hours that would be required with standard modeling tools. This example also illustrates how large scale model acquisition can be significantly simplified if the acquired scene contains repetitive structures. An element acquired with high accuracy can be used as a context model for similar parts, which can then be scanned at much lower resolution.

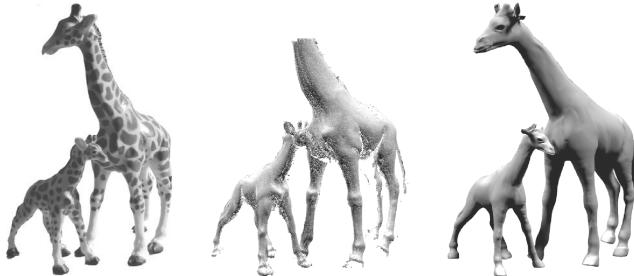


Figure 15: Shape reconstruction from low-quality data, from left to right: Physical model, acquired data set, reconstructed model.

Additional Constraints. The shape matching penalty defined in Section 4 only considers low-level geometric properties to determine the warping function for non-rigid alignment. However, many models have specific high-level semantics that are not considered in this measure. For example, certain models exhibit symmetries that should be preserved by the warping function. As shown in Figure 12, we can adapt the alignment by adding appropriate constraints in the optimization. Another typical example is articulated models, where deformations that describe rotations around joints should be penalized significantly less than ones that result in a bending of rigid parts of the skeletal structure. This can be achieved by using a full kinematic description of the context models to adjust the matching penalty function accordingly.

Evaluation. A distinct advantage of our method is that it not only provides a final consolidated surface mesh, but also allows a local evaluation of the quality of the reconstructed model. We can easily identify regions where no adequate shape completion can be obtained, either because no valid correspondence between input data and context models can be established, or because the distortion of the warping function is too high to provide a meaningful shape prior for the acquired data. The zoom of the giraffe’s head shown in Figure 16 depicts a case where our method does not recover a semantically correct shape, since the horns of the giraffe are not

present in any of the context models and the data set is incomplete in this region. In such cases, the user either needs to acquire more data, enrich the database by providing more suitable context models, or manually edit the final model.

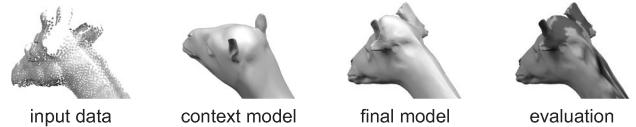


Figure 16: Evaluating the final completed shape. The color-coding in the right image shows the shape matching penalty, where red color indicates insufficient surface completion due to invalid correspondence between input data and context models.

Limitations. Our context model retrieval relies on textual queries, which requires a well annotated shape database. This is particularly important for models that only provide partial completions in a certain region of the input data, but disagree greatly in other parts. Pre-segmentation of database models can simplify the retrieval of partially matching shapes, but requires a substantially more involved database search. In our current implementation, context models have to be manifold surfaces to ensure correct front-propagation during segmentation. Many existing 3D shapes available in public databases are non-manifold models that would require pre-processing to transform them into manifolds, e.g., using the method proposed in [Shen et al. 2004].

Similar to [Allen et al. 2003] and [Sumner and Popovic 2004] we control the distortion of the warping function when computing the non-rigid alignment, not the shape of the deformed model. Thus we can make no guarantees that the warped model is free of self-intersections. We discard a context model when we detect such a case, yet constraining the deformation to prevent self-intersections might be a more adequate solution.

The distortion measure that controls the smoothness of the warping function is isotropic, i.e., penalizes distortion equally in all radial directions. If the acquired model has a high-frequency detail, e.g., a sharp crease, that is not present in the context model, the weight on the distortion measure needs to be low (i.e., α has to be close to one in Equation 6), so that the warped context model can be aligned to this geometric feature. This, however, will also pick up noise present in the input data, as can be observed in Figure 11. A solution could be to design an anisotropic shape matching penalty that locally respects that characteristics of the input geometry, similar to anisotropic low-pass filters used in data smoothing.

The blending method of Section 6 requires consistent topology of the context models in regions where two or more models are blended. We detect topological mismatches from inconsistencies in the correspondence between different models, and exclude the model with higher shape matching penalty from the blending stage in this region. We can give no guarantees, however, that this heuristic always produces the correct shape topology. We thus also allow the user to manually discard individual models, if the topology is inconsistent, which provides more explicit control of the semantics of the consolidated shape.

8 Conclusion and Future Work

We have presented an example-based shape completion framework for acquired 3D surface data. Central to our method is the ability to combine context information from different geometric priors

retrieved from a 3D model database. For this purpose we have introduced a normalized shape matching penalty function and corresponding optimization scheme for non-rigid alignment that allow a quantitative comparison between different context models. This facilitates an adaptive segmentation of the warped context models, which can then be blended consistently using incremental patch growing and continuous re-alignment, to yield the final consolidated shape representation. Our method is robust against noise and outliers and provides a quantitative evaluation of the quality of the produced output model. We achieve efficient reconstruction of complete and consistent surfaces from highly incomplete scans, thus allowing digital 3D content creation with significantly simplified acquisition procedures.

Possible extensions of our system include more powerful, anisotropic shape similarity measures, enhanced semantic constraints to control the deformation of context models, and the use of additional model attributes such as surface texture to improve the retrieval, alignment, and segmentation stages of our pipeline. Ultimately, we want to minimize user intervention and completely automate the database retrieval and non-rigid alignment stages of the pipeline. We believe that this poses interesting research challenges in automatic feature extraction, semantic shape segmentation, inter-model correspondence, and efficient encoding of shape variability.

A Appendix

The quality-of-fit estimate c_i^λ is derived from weighted covariance matrix

$$\mathbf{C}_i = \sum_j (\mathbf{p}_j - \mathbf{p}_i)(\mathbf{p}_j - \mathbf{p}_i)^T \phi_i(\|\mathbf{p}_j - \mathbf{p}_i\|),$$

where the weight function ϕ_i is the compactly supported fourth order polynomial

$$\phi_i(r) = \begin{cases} 1 - 6r^2 + 8r^3 - 3r^4 & r \leq 1 \\ 0 & r > 1 \end{cases}$$

with $r = (\|\mathbf{p}_j - \mathbf{p}_i\|)/h_i$. The support radius h_i defines the geometric scale at which the data is analyzed. Comparable results are obtained using truncated Gaussians, or similar positive, monotonously decreasing weight functions.

Let $\lambda_i^1 \leq \lambda_i^2 \leq \lambda_i^3$ be the eigenvalues of \mathbf{C}_i . The normalized weighted least squares error of the best tangent plane estimate can be derived as $\lambda_i = \lambda_i^1/(\lambda_i^1 + \lambda_i^2 + \lambda_i^3)$ [Pauly et al. 2004]. Since $\lambda_i = 0$ indicates a perfect fit and $\lambda_i = 1/3$ denotes the worst possible distribution, we define $c_i^\lambda = 1 - 3\lambda_i$.

We measure local sampling uniformity as the ratio $c_i^\sigma = \lambda_i^2/\lambda_i^3$. $c_i^\sigma = 0$ means that all samples in the support of ϕ_i lie on a line (see outliers in Figure 3), whereas $c_i^\sigma = 1$ indicates a uniform distribution of samples around \mathbf{p}_i . Points on the boundary will be assigned intermediate values.

References

ALLEN, B., CURLESS, B., AND POPOVIC, Z. 2003. The space of human body shapes: reconstruction and parameterization from range scans. *ACM Trans. Graph.* 22, 3, 587–594.

AMENTA, N., BERN, M., AND KAMVYSSELIS, M. 1998. A new voronoi-based surface reconstruction algorithm. In *Proceedings of SIGGRAPH '98*, ACM Press, 415–421.

ANGUELOV, D., SRINIVASAN, P., KOLLER, D., THRUN, S., PANG, H., AND DAVIS, J. 2004. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. *NIPS*.

BLANZ, V., AND VETTER, T. 1999. A morphable model for the synthesis of 3d faces. In *Proceedings of SIGGRAPH '99*, ACM Press/Addison-Wesley Publishing Co., 187–194.

BLANZ, V., MEHL, A., VETTER, T., AND SEIDEL, H. 2004. A statistical method for robust 3d surface reconstruction from sparse data. In *Int. Symp. on 3D Data Processing, Visualization and Transmission*.

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 *Proceedings of SIGGRAPH '01*, ACM Press, 67–76.

CLARENZ, U., DIEWALD, U., DZIUK, G., RUMPF, M., AND RUSU, R. 2004. A finite element method for surface restoration with smooth boundary conditions. *Comput. Aided Geom. Des.* 21, 5, 427–445.

CURLESS, B., AND LEVOY, M. 1996. A volumetric method for building complex models from range images. In *Proceedings of SIGGRAPH '96*, ACM Press, 303–312.

DAVIS, J., MARSCHNER, S., GARR, M., AND LEVOY, M. 2002. Filling holes in complex surfaces using volumetric diffusion. In *1st Int. Symp. on 3D Data Processing, Visualization, and Transmission*.

DEY, T. K., AND GOSWAMI, S. 2003. Tight cocone: a water-tight surface reconstructor. In *Proceedings of Solid Modeling '03*, 127–134.

DEY, T. K., AND GOSWAMI, S. 2004. Provable surface reconstruction from noisy samples. In *Proc. 20th Symp. Comp. Geom.*, 330–339.

DRORI, I., COHEN-OR, D., AND YESHURUN, H. 2003. Fragment-based image completion. *ACM Trans. Graph.* 22, 3, 303–312.

FLEISHMAN, S., DRORI, I., AND COHEN-OR, D. 2003. Bilateral mesh denoising. *ACM Trans. Graph.* 22, 3, 950–953.

FUNKHOUSER, T., KAZHDAN, M., SHILANE, P., MIN, P., KIEFER, W., TAL, A., RUSINKIEWICZ, S., AND DOBKIN, D. 2004. Modeling by example. *ACM Trans. Graph.*

HOPPE, H., DEROSE, T., DUCHAMP, T., McDONALD, J., AND STUETZLE, W. 1992. Surface reconstruction from unorganized points. In *Proceedings of SIGGRAPH '92*, ACM Press, 71–78.

JONES, T. R., DURAND, F., AND DESBRUN, M. 2003. Non-iterative, feature-preserving mesh smoothing. *ACM Trans. Graph.* 22, 3, 943–949.

KÄHLER, K., HABER, J., YAMAUCHI, H., AND SEIDEL, H.-P. 2002. Head shop: generating animated head models with anatomical structure. In *ACM SIGGRAPH/Eurographics Symp. on Comp. Animation*, 55–63.

LEVY, B. 2001. Constrained texture mapping for polygonal meshes. In *Proceedings of SIGGRAPH '01*, ACM Press, 417–424.

LIEPA, P. 2003. Filling holes in meshes. In *ACM SIGGRAPH/Eurographics Symp. on Geometry Processing*, Eurographics Association, 200–205.

LINDEBERG, T. 1998. Feature detection with automatic scale selection. *Int. Journal of Computer Vision* 30, 2.

MITRA, N. J., GELFAND, N., POTTMANN, H., AND GUIBAS, L. 2004. Registration of point cloud data from a geometric optimization perspective. In *ACM SIGGRAPH/EG Symp. on Geometry Processing*, 23–31.

OHTAKE, Y., BELYAEV, A., AND SEIDEL, H.-P. 2003. A multi-scale approach to 3d scattered data interpolation with compactly supported basis functions. In *Proceedings of Shape Modeling International 2003*.

PAULY, M., KEISER, R., AND GROSS, M. 2003. Multi-scale feature extraction on point-sampled models. In *Proceedings of Eurographics*.

PAULY, M., KEISER, R., KOBELT, L. P., AND GROSS, M. 2003. Shape modeling with point-sampled geometry. *ACM Trans. Graph.* 22, 3.

PAULY, M., MITRA, N., AND GUIBAS, L. 2004. Uncertainty and variability in point cloud surface data. In *Symposium on Point-Based Graphics*.

RAMAMOORTHI, R., AND ARVO, J. 1999. Creating generative models from range images. In *Proceedings of SIGGRAPH '99*, ACM Press/Addison-Wesley Publishing Co., 195–204.

SHAFR, A., ALEXA, M., AND COHEN-OR, D. 2004. Context-based surface completion. *ACM Trans. Graph.* 23, 3, 878–887.

SHEN, C., O'BRIEN, J. F., AND SHEWCHUK, J. R. 2004. Interpolating and approximating implicit surfaces from polygon soup. In *ACM Trans. Graph.*, ACM Press.

SUMNER, R. W., AND POPOVIC, J. 2004. Deformation transfer for triangle meshes. *ACM Trans. Graph.* 23, 3, 399–405.

TANGELDER, J., AND VELTKAMP, R. 2004. A survey of content based 3d shape retrieval methods. In *Shape Modeling International*.

TAUBIN, G. 1995. A signal processing approach to fair surface design. In *Proceedings of SIGGRAPH '95*, ACM Press, 351–358.

TURK, G., AND LEVOY, M. 1994. Zippered polygon meshes from range images. In *Proceedings of SIGGRAPH '94*, ACM Press, 311–318.

VERDERA, J., CASELLAS, V., BERTALMIO, M., AND SAPIRO, G. 2003. Inpainting surface holes. In *Int. Conference on Image Processing*.

WEYRICH, T., PAULY, M., HEINZLE, S., KEISER, R., SCANDELLA, S., AND GROSS, M. 2004. Post-processing of scanned 3d surface data. In *Symposium on Point-Based Graphics*, 85–94.

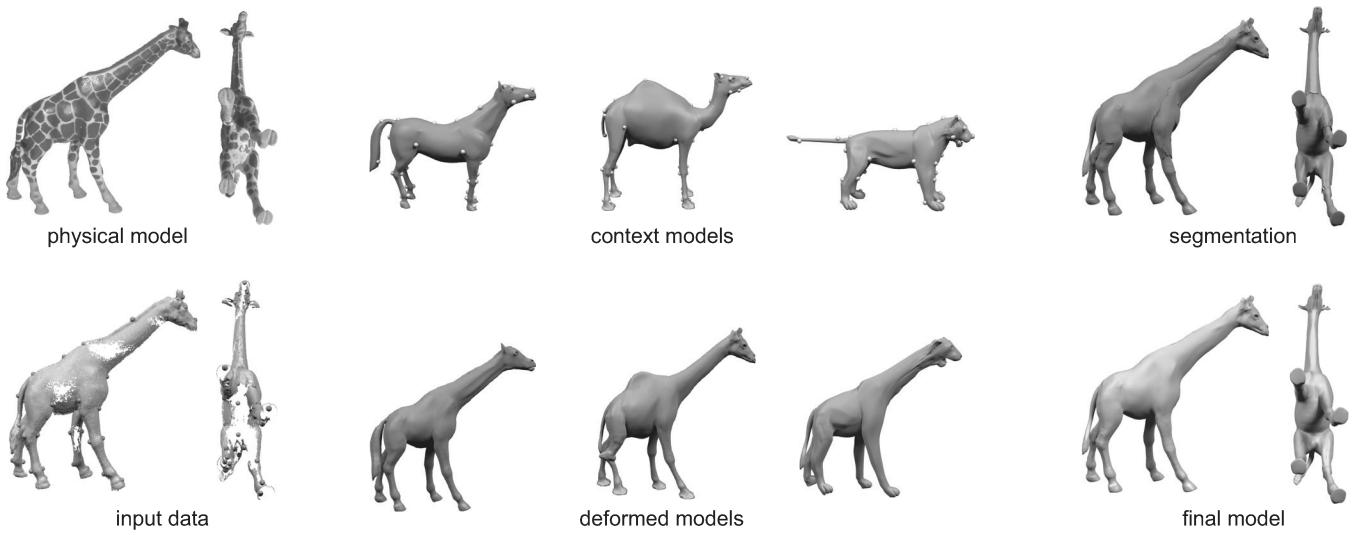


Figure 10: Shape completion zoo. Horse, camel, and lion are deformed, segmented and blended to yield the final shape of the giraffe.

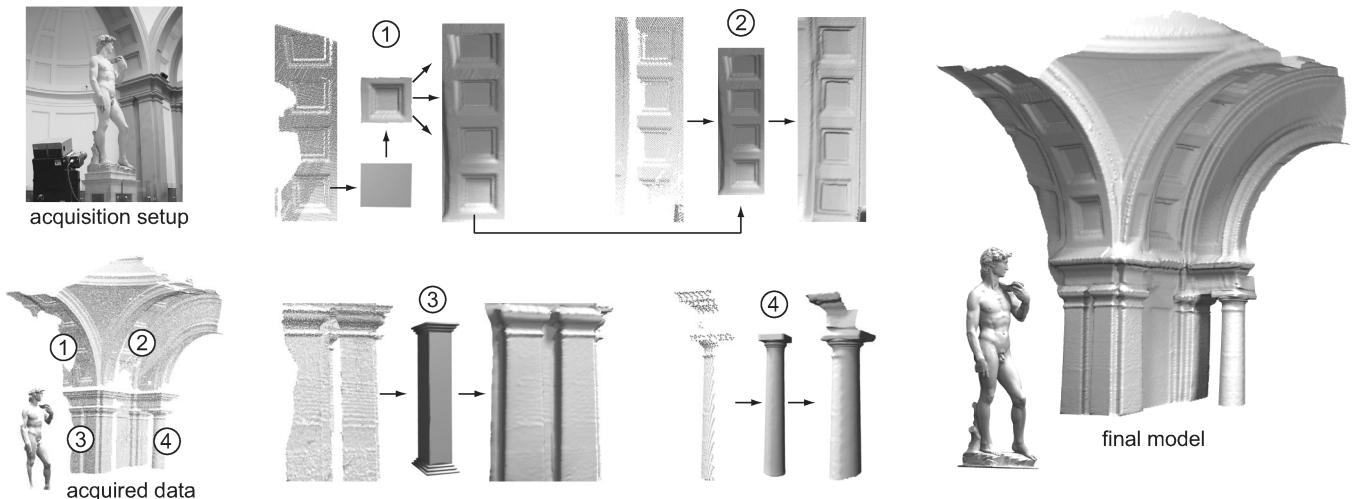


Figure 11: Completion of a single range image acquired in the Galleria dell'Accademia in Florence. Context models, shown in brown, are either retrieved from the database or extracted by the user from already completed parts of the model. The David model has been added for completeness.

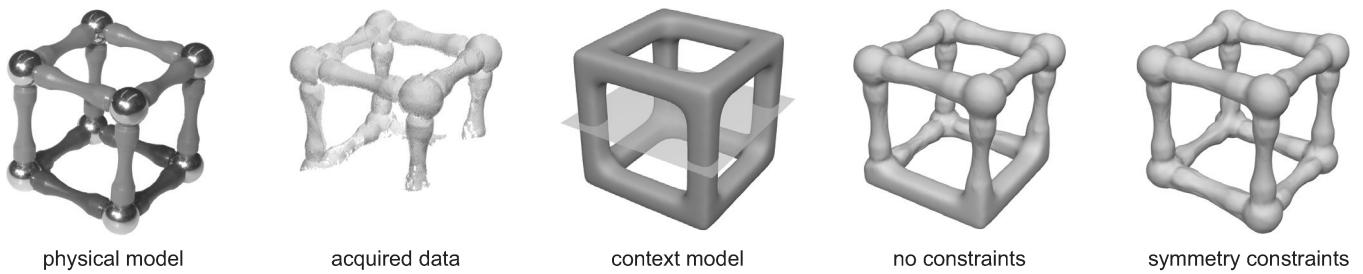


Figure 12: Symmetry constraints yield a semantically more adequate shape completion. The warping function for the model on the right has been constrained to be symmetric with respect to the semi-transparent plane shown in the center.