

Surface Reconstruction using Local Shape Priors

Ran Gal¹ Ariel Shamir² Tal Hassner³ Mark Pauly⁴ Daniel Cohen-Or¹

¹Tel-Aviv University, Israel

²The Interdisciplinary Center, Israel

³The Weizmann Institute of Science, Israel

⁴ETH Zurich, Switzerland

Abstract

We present an example-based surface reconstruction method for scanned point sets. Our approach uses a database of local shape priors built from a set of given context models that are chosen specifically to match a specific scan. Local neighborhoods of the input scan are matched with enriched patches of these models at multiple scales. Hence, instead of using a single prior for reconstruction, our method allows specific regions in the scan to match the most relevant prior that fits best. Such high confidence matches carry relevant information from the prior models to the scan, including normal data and feature classification, and are used to augment the input point-set. This allows to resolve many ambiguities and difficulties that come up during reconstruction, e.g., distinguishing between signal and noise or between gaps in the data and boundaries of the model. We demonstrate how our algorithm, given suitable prior models, successfully handles noisy and under-sampled point sets, faithfully reconstructing smooth regions as well as sharp features.

1. Introduction

Since reconstructing a continuous surface from discrete samples is an inherently ill-posed problem, any reconstruction method must rely on prior assumptions. These can be assumptions on the sampling process, or on the nature of the 3D surface or model sampled. In a Bayesian framework these assumptions are converted to some prior probability which constrains the shape being reconstructed. Such a Bayesian approach has been used successfully in numerous works on images, medical-data, and surface reconstruction before [MCSM99, WK01, FWZ03, DTB, JWB*06].

However, there is a limit to what can be achieved in reconstruction algorithms using general priors (e.g. that the surface is smooth, that the noise is Gaussian distributed), and without using the correct context of the shape. Recently, works such as [PMG*05, KS05] use explicit priors from a database of example shapes to assist reconstruction. However, these methods search for a *global* match for the shape, which can be difficult to find, potentially gives low confidence in the matching, or requires non-rigid alignment.

In many cases, the scanned shape consists of a composition of parts of previously known shapes. In these situa-

tions, instead of using global matching with low confidence, we advocate using as-large-as-possible local similarities but with very high confidence. Our method allows specific regions in the scan to match the most relevant prior that fits best. The higher the confidence in the match is, the more relevant the shape prior is to the reconstruction process. Using such an approach can help resolve many ambiguities that come up during reconstruction, e.g., distinguishing between signal and noise, between sharp features and noise, or between gaps in the data and boundaries of the model.

In this work we present an example-based approach that uses explicit *local shape priors* that can provide correct context for complex shape reconstruction. We concentrate on cases where digital models similar to the reconstructed one could be found. This can be done by searching for similar objects or by extracting similar parts from previously scanned models to create a database of examples for the reconstruction. These form a basis to learn typical shapes and details by creating local shape priors in the form of enriched-patches. These are point-set neighborhoods sampled from the database models that include additional information such as normals and point feature classification (e.g. if a point lies on an edge, a corner, or other features). The patches serve as




Figure 1: Surface reconstruction with local shape priors. Left: some examples of the local shape priors defining the context for reconstructing the Lego model. Right: the augmentation process finds shape priors matching local neighborhoods in the scan, aligns them, and places them in the scan. As a consequence, the scan is augmented with noiseless samples, quality normal information, and sharp feature classifications (green points).

a type of training-set for the reconstruction process, and define its context to assist the reconstruction of similar, newly scanned shapes and objects (Figure 1).

The basic idea for utilizing such patches is to add an *augmentation* phase of the scanned point-set before applying the actual reconstruction algorithm. In this phase the local patches are used to augment, de-noise and repair parts of the input data. Hence, the final reconstructed model will be affected by a blend of the original point-set and the enriched local patches from the shape prior training set. Consequently, the shape priors define the *context* of the reconstruction.

Effective augmentation requires efficient partial shape-matching of local neighborhoods of the scanned point-set with local patches of the shape priors. As a pre-process we build a library of context patches by sampling the database models to gather local point-set neighborhoods at multiple scales. Local shape descriptors are used to efficiently compute a set of matching candidate patches, which is then refined using local registration. Patches with best matching scores, i.e. highest confidence, are used to augment the scan by adding, after proper alignment, the patch's enriched points to the scan. These points reduce the signal to noise ratio in the scan, fill gaps and holes in the original data, and carry additional information such as normals and feature classification.

One of the main benefits of this example-based approach is flexibility and adaptivity to different application scenarios. For example, in certain applications sharp features might be desirable, while other domains favor rounded features of a certain curvature radius. Instead of tuning parameters or adding new specialized algorithms to accommodate such shape properties, our approach allows the user to define the

context of the reconstruction explicitly by providing suitable example shapes.

The rest of the paper is composed as follows. After reviewing related work we describe the construction of the shape prior library in Section 3. Next, we present the scan augmentation process in Section 4, and the definition of the reconstructed augmented surface in Section 5. We show results in Section 6 and conclude.

2. Related work

Reconstructing a surface from an unstructured point cloud is a difficult, ill-posed problem, which has received considerable attention over the years. The pioneering work of [HDD*94] suggest a solution to this problem based on interpolating signed distance fields. Computational geometry based approaches attempt to reconstruct the surface using Voronoi diagrams and offer provable sampling conditions under which the reconstructed surface is homomorphic to a smooth compact 2-manifold [ABK98, DG03, DG04, DG06]. These methods, however, are limited when faced with large holes or under-scanned point clouds. A different class of solutions thus attempt to first reconstruct a mesh [CL96, TL94] and then apply smoothing and stitching to remove noise and seal holes [Tau95, DMSB99, CDR00]. Alternatively, surface normals and viewing directions can be used to consolidate points that were scanned multiple times [CL96, WSI98]. These methods all make the implicit assumption that missing information in the scanned data can be inferred from its immediate surroundings. This is often true, but can fail with detailed objects, or large holes spanning geometric features.

A popular means to overcome the ill-posed nature of the reconstruction problem, is to make general assumptions on the nature of the surface or model being reconstructed. Typically, these include smoothness assumptions, and assumptions on the noise model (e.g. that the noise




Figure 2: An illustration of the reconstruction process in the area of sharp features: (a) The input point set, (b) The result of our augmenting process (green points with normals), (c) The normal clustering results near the sharp features (d) The final projection result.

can be modeled by a Gaussian distribution). These assumptions are effective for both noise removal and gap filling [FDCO03, SACO04]. However, such assumptions often contradict the faithful reconstruction of delicate surface details and sharp features. More advanced techniques which try to preserve sharp features by post-processing, such as anisotropic noise removal [CDR00, TWBO02, FDCO03, BX03], robust statistics [FSCO05], or Bayes’ rule for geometry inference [JWB*06] are also limited in their ability to reconstruct details.

Recently, a number of methods have suggested using class based information for the purpose of surface reconstruction. This information is often captured by explicit template priors for the shape of the object. Unlike previous methods, that are oblivious to the particular object being processed, these methods are “data-aware” in the sense that they carry information specific to the class of the object being reconstructed. An early example is the work of [RA99] which attempts to fit known parametric spline surfaces to acquired surface data. Methods such as [ACP03, APD*05, BMVS04, KS05] use a global template which is morphed to fit the scanned data. These work best when the scanned data is similar in global structure to the prior template. Finally, [SACO04] suggest a method for sealing holes in surfaces by reusing surface patches from around the hole. These patches serve as priors to any missing information. To reconstruct a complete surface from scanned data [PMG*05] form global matches between the point cloud being reconstructed and a set of database objects (the priors). Their method performs best when the database contains objects which are similar to the one being reconstructed. In this work we take a similar data-aware approach. However, unlike [PMG*05], we use local surface patches as priors, and not global parts. Thus, our method seeks to match the input point cloud with many local surface patches, each with high confidence but each carrying less information.

© The Eurographics Association 2007.

3. The Shape Prior Patch Library

We define the context for example-based reconstruction by building a shape patch library from a given set of database models. For ease of exposition we assume only a single database model given as a triangle mesh M . To create the point-set patches we first sample the mesh uniformly and create a sampled model point set $\mathbb{P}(M)$ [Tur90]. We sample points on all mesh elements: vertices, edges, and faces, store their normals, and signify points belonging to sharp features as such.

To support multi-scale processing, we compute for each point $q \in \mathbb{P}(M)$ a set of n point patches $\{P_i(q)\}_{i=1}^n$ such that

$$P_i(q) = \{\mathbb{P}(M) \cap \mathbb{S}(r_i, q) \mid q \in \mathbb{P}(M), r_i = \frac{i}{n}R\},$$

where $\mathbb{S}(r, q)$ is a solid sphere of radius r centered at q , R is a predefined basic radius, and n the number of scales we use for this model. For each patch $P_i(q)$ we compute its canonical position using *weighted* PCA, such that the center of mass of the patch is translated to the origin, and the principal axes coincide with the coordinate axes, in decreasing order of principal components. The weights for each point in $P_i(q)$ are computed similar to [PKKG03], as the average distance to it’s k -neighbors (we use $k = 16$). This increases the contribution of sparsely sampled areas and accounts for differences in the sampling rate. Canonical scaling is achieved by uniformly scaling all patches so that the length of the largest principal component becomes 1.

Efficient local matching of patches is facilitated by computing a signature for each patch $P_i(q)$. We use geometric moments as our shape descriptor. Moments have been used for matching and recognition of 2D shapes in images [Hu62, AZHH88, SF01] and for global shape matching in 3D as well [ETA00, ETA01]. The p, q, r -moment of a shape S in 3D is a scalar defined by:

$$M_{p,q,r}(S) = \int_{\partial S} x^p y^q z^r dx dy dz.$$

For uniformly sampled shapes (as we can assume is our case in $P_i(q)$) one can replace the integral with a sum over all the sample points. We define the descriptor of a patch P as the vector moments up to some order d :

$$V(P) = (M_{0,0,1}(P), M_{0,0,2}(P), \dots, M_{1,1,1}(P), \dots, M_{d,0,0}(P)),$$

such that $i + j + k \leq d$. It has been shown [ETA01] that for geometrically similar shapes P and Q , the descriptor vectors $V(P)$ and $V(Q)$ are also close, namely $\|V(P) - V(Q)\|$ is small. In our implementation we use $d = 5$, creating a descriptor vector of 55-dimensions. It is important to stress that many different shape descriptors have been developed for shape matching and 3D search engines (see e.g. comparison in [SMKF04]). Our choice of geometric moments was based on the fact that they are fast and easy to compute, they do not require additional information such as reliable normals, they can be used on un-organized point-sets (the last two are

important since we shall use the same descriptor for neighborhoods of the scanned point-set), and they support partial matching of neighborhoods.

Lastly, we insert all the patches $\{P_i(q)\}_{i=1}^n$ for all $q \in \mathbb{P}(M)$ into a k -nearest-neighbors search data structure [MA97] based on their 55-dimensional descriptor vectors $V(P_i(q))$. This structure defines the shape prior patch library \mathcal{L} that supports fast matching queries based on the descriptor. We also store the transformation to canonical position for each patch, which will later be used to correctly orient the elements during augmentation.

4. The Augmentation Procedure

When a new scanned point-set \mathcal{C} is received, we define its context by choosing a set of shape priors and creating the corresponding patch library \mathcal{L} . To augment the scanned point-set \mathcal{C} we must search for patches in the library similar to local neighborhoods in \mathcal{C} . Hence, we create local neighborhoods of the scan in a similar manner as the procedure described above for the shape priors. We pick a sub-set $\mathbb{C} \subset \mathcal{C}$ of points from the scan point-set uniformly at random. For each point $p \in \mathbb{C}$ we build n neighborhood patches $C_i(p)$ for n scales around p and calculate their descriptors $V(C_i(p))$. Before calculating the descriptors we transform them to a canonical placement using the weighted PCA based positioning described before. Using these n descriptors we retrieve the best matching shape prior patch by querying \mathcal{L} .

After this initial matching stage, each point $p \in \mathbb{C}$ is associated with a set of n prior-patches retrieved from \mathcal{L} , each one matching a neighborhood in growing scales around p . Although it seems it would be best to use the largest possible prior around each point p , note there may be a large variance in the quality of the matches. Some of them may indeed assist future reconstruction, but some may harm it by adding noise or outliers. Instead of taking the largest possible neighborhood, we would like to choose the ones with the highest confidence. To find these reliably, we do not rely on the shape descriptor distance only, but devise a more elaborate matching refinement procedure.

Let P be the shape prior patches returned by querying \mathcal{L} with the descriptor of the neighborhood $C_i(p)$ around $p \in \mathbb{C}$. To actually benefit from placing P onto \mathcal{C} and to define an accurate confidence measure for the match, P must be aligned carefully onto \mathcal{C} . We first apply the inverse of the transformation of $C_i(p)$ to its canonical position on P . This brings P to correct scale and a good initial position. Next, this position is refined using the Iterative Closest Point (ICP) algorithm [RL01]. To achieve more reliable alignment results, we use a slightly larger neighborhood than $C_i(p)$ (around 10% to 50% larger) for aligning P to \mathcal{C} in the ICP algorithm.

Now we can define the basic matching confidence score

$D(P)$ of P as the mean squared error (MSE) of the distance between all points $q \in P$ and their closest match in \mathcal{C} :

$$D(P) = \frac{1}{k} \sum_k \min_{q_k \in P, c \in \mathcal{C}} (\|q_k - c\|^2)$$

The larger $D(P)$ is, the smaller the confidence in the match. For each point $p \in \mathbb{C}$ we use the D scores of the patches to choose the best matching patch from the n possible matches of the n scales neighborhoods of p . However, $D(P)$ still relies heavily on small neighborhood information in \mathcal{C} . To enlarge the effect of the global structure of \mathcal{C} , we propagate information from neighboring patches to strengthen or weaken the confidence score $D(P)$.

Each local shape prior patch P can overlap several other matched and aligned patches. To increase the coherency of the matches, we build a graph G where each node represents one shape prior patch P that was fetched in the querying process, and the edges of G connect two patches iff they overlap geometrically. We use this graph to propagate consensus information in larger neighborhoods in \mathcal{C} . The graph is traversed in a breadth first manner, visiting each edge and applying a simple procedure for each pair of connected patches. In the overlapping region of the two patches, we compare the average difference of the normals of pairs of closest points, one from each patch. If this average difference deviates by less than 30° , the two patches are coherent and we reduce their D score by a certain factor (i.e. we increase their match confidence).

After visiting all edges in the graph we arrive at the final confidence score for each match that combines both local information (the MSE of the points) and patch neighborhood coherency information (the deviation of normals). At the last stage we remove all matched patches with score above some threshold (i.e. with low confidence), and use only the remaining patches to create the augmented point set \mathcal{A} .

5. Augmented Projection

Similar to numerous recent implicit reconstruction methods [ABCO*03, AK04, CFS02], our work is based on the moving least-square approach [Lev03]. One of the key factors in the success of these methods is the definition of consistent normals for the point set being reconstructed. Better normals lead to better reconstruction results, in particular in regions of high geometric detail and sharp features, where the variance of normals is high. In this context, using the augmented point-set \mathcal{A} offers the following advantages:

- improved normal estimation for smooth neighborhoods in the scan,
- different piecewise smooth regions in the scan can be distinguished more robustly,
- sharp features such as edges and corners in the scan can be

directly identified instead of determined by methods such as [FCOS05].

In general, our algorithm classifies the points in the augmented point-set \mathcal{A} into a number of subsets, each corresponding to a smooth region of the surface, and to sharp features such as edges and corners. This classification allows us to project points on a locally piecewise smooth region rather than a globally smooth surface, thus defining a surface with the ability to construct sharp features. Consequently, each of those regions will be reconstructed independently according to its characteristic and the resulting mesh will be the combinations of these reconstructed elements.

Given a point $p \in \mathcal{A}$, we first analyze its neighborhood S . If S does not contain enough information from prior patches (more than 50% of the points in S are from the original point cloud), we use the standard MLS operator or robust approaches such as defined in [FCOS05]. If the neighborhood S of p contains enough prior information, i.e. points from the prior patch library \mathcal{L} , we continue in the following manner.

Recall that these points from \mathcal{L} have high quality normals, since they originated from a surface model. Hence, we leverage this information to create better normal estimation everywhere in S . Beginning at p , we traverse the same graph G of patches built for refining the confidence measure of the patches, and construct coherent normal directions for the neighborhood S of p . This enhances MLS projection results even for smooth regions of the scan.

In addition the points from \mathcal{L} contain feature classification information, i.e. some of them are classified as edge points (intersection of two smooth surfaces) and some are classified as corner points (intersection of more than two smooth surfaces). This prior information resolves in a simple manner the difficult and often ambiguous distinction between smooth regions, i.e. S does not contain any feature points, and regions that contain sharp features. Our goal in projection is to locally fit a number of polynomials to the points in S . If no feature points are included in the neighborhood S , then S is smooth and we fit a single polynomial to the points in S . Otherwise, we search for several piecewise smooth regions in this neighborhood, and fit multiple polynomials for subsets of points in S (Figure 2). To find these subsets, or smooth regions, we assume that each of them can be approximated by a plane locally near the edge. Therefore, we use normal clustering to find these regions. We first find the normal direction that agrees (up to an angle threshold α degrees) with all normals in the neighborhood and define it as the first cluster of normals. We then remove the points with such normals and continue to define the next cluster the same way. We consider only subsets with more than 10% of the points in the neighborhood. If S contains points from sharp features, we use $\alpha = 10$, otherwise we use $\alpha = 75$. When we identify two or more subsets, the surface is defined as the intersection of the smooth surfaces defined by




Figure 3: Lower dimension features projection. The only shape prior in the database for this reconstruction is a straight L-shaped corner (shown at the top). Hence, the connection curve between the box and the cylinder is augmented by piecewise straight lines (a). Nevertheless, since we know these points are classified as sharp edge feature points, we use a constrained lower-dimension projection (b), and arrive at the final smooth connecting curve (c).

the different subsets. We then continue and project p to its position on the correct surface as in [FCOS05].

Lower-dimensional features such as corners and edges, that are marked as sharp features explicitly, are projected in a constrained manner at the last stage. These are projected either onto a one-dimensional edge (intersection of two smooth surfaces) or onto a corner (intersection of more than two surfaces) similar to [PKKG03]. In order to approximate the location of the points on the curve (sharp edge) we use a variation of the MLS projection operator. We project each point that is marked as sharp feature onto a one dimensional polynomial of degree 1 (a line) that best approximate the neighborhood of sharp features points. This stage reconstructs the sharp features explicitly in greater accuracy (Figure 3).

6. Results

We analyze the performance of our reconstruction using local-priors on a variety of raw scans acquired by different 3D scanners, and on some artificial examples. We focus on models that contain sharp features, which are difficult to scan, and include regions with under sampling as well as holes. The shape-prior patch library \mathcal{L} built from the prior models usually contains between 5,000 to 30,000 patches. In all our experiments, in the augmentation stage, we increase the scanned point set size by a factor of ten to fifteen. Our current implementation is MATLAB based and does not include optimization. The reconstruction time (including augmentation) of the paper-clip model ($\sim 24K$ points) is around 30 minutes, the remote control model ($\sim 60K$ points) about 45 minutes, the octopus leg ($\sim 60K$ points) around an hour, and the Lego ($\sim 47K$ points) and the pulley ($\sim 41K$ points) models just under 2 hours each.

Figure 4 shows a single laser-scan depth image of a remote control. Note that the sides of the buttons were not sampled at all. We generated our shape prior set from a 3D model




Figure 4: Reconstructing part of a remote control. In (a) we show the original object scanned (left) using a single laser scan from above, and the prior we used (right) - a remote control model from the web. (b) is the input point-cloud. Note the missing samples on the sides of the buttons. (c) is our reconstruction. The holes in the upper parts are due to the fact that our prior database does not contain information for this type of buttons. In (d) we show the standard MLS reconstruction.




Figure 5: Reconstructing a Lego model. Low-left: A photograph of the original model scanned using a structured-light scanner. Upper right: a comparison between Local-prior reconstruction (upper right) and the standard MLS method (bottom-right).

of a remote control we found on the web (Figure 4(a)), which has only little resemblance to the scanned remote control. As can be seen, we were able to reconstruct the missing parts in the buttons better than standard MLS reconstruction, that uses only a global smoothness prior.

In the example of Figure 5, we scanned a Lego model using a structured light scanner (Vialux Z-snapper camera). We use seven aligned scans as our input. The input data contains many holes and the sampling density is inhomogeneous across the model. In addition, the sampling rate is very low compared to the feature size in the model, which poses great difficulties for methods based on naive priors. The local priors that we used were extracted from one Lego brick model that we found on the web. In Figure 1 we show some examples of our priors and a visualization of some steps of the




Figure 6: Reconstructing a Lego model. On the bottom-right, we show the reconstruction result of an under-sampled area. Bottom-middle shows that our method fails to reconstruct missing parts. Bottom right shows areas that were smoothed out due to missing reliable information in our prior database.

partial matching process. Note that the sharp features of the reconstructed model are solely due to the use of priors that contain sharp features. The full model reconstruction with close-ups on some areas of the model can be seen in Figure 6.

Even when a specific similar prior does not exist, using very simple priors we can often improve the reconstruction result of existing scans. We used our method on a Pulley model from the AIM@SHAPE repository. In Figure 7 we show the original model, the reconstruction result, and a close-up of the final triangulation. Note that the resulting reconstruction is biased towards our simple priors, which emphasizes the importance of the relevance of prior set.




Figure 7: Reconstruction of a pulley model. Top-left is the input point-cloud extracted from the pulley model. Bottom-left is the result of our augmentation process using local shape priors. Middle-top is the triangulation result of the original point-cloud and middle-bottom is the triangulation result of our projected point-cloud. The left image shows details of a triangulation of our projected point-cloud.




Figure 8: Reconstruction of a paper-clip using parts of a cylinder as shape priors. The input consists of two laser scans. In (a) we show the input point-cloud with some prior matches on the straight and the curved parts of the point-cloud. Note how the straight parts of the scan match larger neighborhood priors than on the curved part. In (b) we shows our reconstruction result (left) and a standard MLS result (right). We colored the surface according to the geodesic distance from one end of the paper-clip, using local neighborhoods of the reconstructed shape. Note that the MLS method dose not identify the correct topology of the object.

Using reliable priors, our method can also assist in resolving topology ambiguities and complete hidden surface parts. In Figure 8 we show a paper-clip that was scanned from only two views (top and bottom), using a high definition laser scanner (Optimet MiniConoscan 3000). Nevertheless, the original model contains regions of very close cylinder-like shapes that prohibit the scanner from acquiring reliable surface information. Using common reconstruction techniques or global smoothness priors, these two close cylinder parts will be merged into one surface.

By augmenting the scan using local shape priors extracted from a simple straight cylinder we can fill missing parts in the shape. More importantly, since we obtain high quality normal information from the prior patches, we can distinguish between the two cylinder parts in the shape. This is

illustrated in Figure 8 by coloring local geodesic neighborhoods on the surface of the shape starting from green in one side to blue in the other. When the cylinders are merged using previous methods, their colors are similar on both sides. Using our technique the different green and blue colors of the distinct parts are evident.

The enabling factor of our method is providing correct context for reconstruction. The reconstruction depends on the type of shapes used as priors and on the similarity between these shapes to the scanned model. In Figure 9 we compare the results of reconstruction using different shape priors. Using general shape priors such as a sphere and cylinder (Fig. 9(b)) can assist in lowering the noise on smooth areas. However, many points ($\sim 15\%$) in the scanned point-set \mathcal{C} do not match any prior patch. These points are posi-




Figure 9: Reconstructions with different priors context. (a) Without using priors. (b) With a general sphere and cylinder. (c) With a highly similar shape. (d) Adding a cone to represent the tips. See discussion in text.

tioned mainly around the surface details, which indeed are not reconstructed faithfully. Using a highly similar shape prior (Fig. 9(c)) leads to higher matching rate ($\sim 92\%$), and to better details preservation. Adding a general cone prior (Fig. 9(d)) pushes the matching rate even further ($\sim 95\%$) and reconstructs a smoother ending for the octopus’s leg tip. Using wrong priors can sometimes become a limitation since the results may be biased towards an erroneous surface (Figure 10). However, to some extent this limitation could be addressed by lowering the rejection threshold for the shape-prior matches.

7. Conclusion

We have demonstrated how the use of explicit priors can significantly improve the quality of surfaces reconstructed from discrete point sets. The benefits of our approach are particularly evident for under-sampled, noisy scans of objects containing sharp features and/or difficult to resolve local topology. For densely sampled, highly accurate laser scans, existing methods using implicit priors are probably more suit-




Figure 10: Reconstructions with wrong priors context. In (a) and (b) a cube model with noise of up to 5% of the bounding box size is reconstructed differently using different context shape priors (either sharp or rounded corners). Using a wrong prior (sphere) in (c) reduces the percent of database matches from 93% to 27%, and the results are still very noisy.

able, as they avoid the overhead of local shape matching and augmentation.

Instead of tuning threshold parameters of the surface reconstruction algorithm, our system allows the user to define the context for reconstruction in an intuitive way using example shapes. The reconstructed surfaces clearly depend on the set of prior shapes, hence the results can only be as accurate as the provided database models. We believe that in many application scenarios (in particular industrial applications) such example databases could be easily constructed, or they might even already exist from previous scans or designs.

We envision many directions for future work. Apart from re-using the geometric information of previously processed models, it would be interesting to investigate re-using information of reconstruction processes as well. For example, we could assign higher weights to priors they have been used frequently with high matching scores to build a likelihood function of the shape space defined by the database models. Our matching approach can potentially be used for automated shape recognition, context-aware model segmentation or context-sensitive compression. It would also be interesting to augment input scans with higher-level information derived from the priors such as symmetries or topological features.

Acknowledgements

This work was supported in part by the Israeli Ministry of Science and the Swiss National Science Foundation. Tal Hassner was supported by the European Community grant IST-2002-506766 Aim@Shape. The vision group at

the Weizmann Institute is supported in part by the Moross Foundation.

References

- [ABCO*03] ALEXA M., BEHR J., COHEN-OR D., FLEISHMAN S., LEVIN D., SILVA C. T.: Computing and rendering point set surfaces. *IEEE Transactions on Visualization and Computer Graphics* 9, 1 (January 2003), 3–15.
- [ABK98] AMENTA N., BERN M., KAMVYSSELIS M.: A new voronoi-based surface reconstruction algorithm. In *Proceedings of SIGGRAPH 1998* (1998), ACM, ACM Press / ACM SIGGRAPH, pp. 415–422.
- [ACP03] ALLEN B., CURLESS B., POPOVIC' Z.: The space of human body shapes: reconstruction and parameterization from range scans. In *Proceedings of SIGGRAPH 2003* (2003), ACM, ACM Press / ACM SIGGRAPH, pp. 587–594.
- [AK04] AMENTA N., KIL Y.: Defining point-set surfaces. In *Proceedings of SIGGRAPH 2004* (2004), ACM Press, pp. 264–270.
- [APD*05] ANGUELOV D., P.SRINIVASAN, D.KOLLER, S.THRUN, RODGERS J., J.DAVIS: SCAPE: Shape completion and animation of people. In *Proceedings of SIGGRAPH 2005* (2005), ACM, ACM Press / ACM SIGGRAPH, pp. 408–416.
- [AZHH88] ABO-ZAID A., HINTON O. R., HORNE E.: About moment normalization and complex moment descriptors. In *Proceedings of the 4th International Conference on Pattern Recognition* (1988), pp. 399–409.
- [BMVS04] BLANZ V., MEHL A., VETTER T., SEIDEL H.-P.: A statistical method for robust 3D surface reconstruction from sparse data. *3DPVT 00* (2004), 293–300.
- [BX03] BAJAJ C. L., XU G.: Anisotropic diffusion of subdivision surfaces and functions on surfaces. *Transactions on Graphics (TOG)* 22, 1 (2003), 4–32.
- [CDR00] CLARENZ U., DIEWALD U., RUMPF M.: Anisotropic geometric diffusion in surface processing. *IEEE Visualization 2000* (2000), 397–405.
- [CFS02] CORRÊA W., FLEISHMAN S., SILVA C.: Towards point-based acquisition and rendering of large real-world environments. In *Proceedings of the 15th Brazilian Symposium on Computer Graphics and Image Processing* (2002), IEEE Computer Society, ACM Press / ACM SIGGRAPH, p. 59.
- [CL96] CURLESS B., LEVOY M.: A volumetric method for building complex models from range images. In *Proceedings of SIGGRAPH 1996* (1996), ACM, ACM Press / ACM SIGGRAPH, pp. 303–312.
- [DG03] DEY T., GOSWAMI S.: Tight cocone: A water tight surface reconstructor. In *Proc. 8th ACM Sympos. Solid Modeling Appl.* (2003), pp. 127–134.
- [DG04] DEY T., GOSWAMI S.: Provable surface reconstruction from noisy samples. In *symposium on Computational geometry* (2004), ACM, ACM Press, pp. 330–339.
- [DG06] DEY T., GOSWAMI S.: Provable surface reconstruction from noisy samples. *Computational Geometry Theory & Applications* 35 (2006), 124–141.
- [DMSB99] DESBRUN M., MEYER M., SCHRÖDER P., BARR A.: Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proceedings of SIGGRAPH 1999* (1999), ACM, ACM Press / ACM SIGGRAPH, pp. 317–324.
- [DTB] DIEBEL J., THRUN S., BRÜNING M.: A bayesian method for probable surface reconstruction and decimation. To appear.
- [ETA00] ELAD M., TAL A., AR S.: *Directed Search In A 3D Objects Database Using SVM*. Tech. rep., HP Laboratories Israel Technical Report, HPL-2000-20 (R.1), August 2000.
- [ETA01] ELAD M., TAL A., AR S.: Content based retrieval of VRML objects: an iterative and interactive approach. In *Proceedings of the 6th Eurographics workshop on Multimedia* (New York, NY, USA, 2001), Springer-Verlag New York, Inc., pp. 107–118.
- [FCOS05] FLEISHMAN S., COHEN-OR D., SILVA C. T.: Robust moving least-squares fitting with sharp features. *ACM Trans. Graph.* 24, 3 (2005), 544–552.
- [FDCO03] FLEISHMAN S., DRORI I., COHEN-OR D.: Bilateral mesh denoising. *ACM Trans. Graph. (SIGGRAPH 2003 Proceedings)* 22, 3 (2003), 950–953.
- [FSCO05] FLEISHMAN S., SILVA C., COHEN-OR D.: Robust moving least-squares fitting with sharp features. *ACM Trans. Graph. (SIGGRAPH 2005 Proceedings)* 24, 3 (2005).
- [FWZ03] FITZGIBBON A., WEXLER Y., ZISSERMAN A.: Image-based rendering using imagebased priors. In *Int. Conf. on Computer Vision* (2003), vol. 63, IEEE Computer Society, pp. 1176–1183.
- [HDD*94] HOPPE H., DE ROSE T., DUCHAMP T., HALSTEAD M., JIN H., McDONALD J., SCHWEITZER J., STUETZLE W.: Piecewise smooth surface reconstruction. In *Proceedings of SIGGRAPH 1994* (1994), ACM, ACM Press / ACM SIGGRAPH, pp. 295–302.
- [Hu62] HU M.-K.: Visual pattern recognition by moment invariants. *IRE Transactions on Information Theory IT-8* (February 1962), 179–187.
- [JWB*06] JENKE P., WAND M., BOKELOH M., SCHILLING A., ER W. S.: Bayesian point cloud reconstruction. *Computer Graphics Forum (Proceedings of Eurographics 2006)* 25, 3 (2006), 379–388.
- [KS05] KRAEVOY V., SHEFFER A.: Template-based mesh completion. In *Eurographics Symposium on Geometry Processing* (2005), pp. 13–22.

- [Lev03] LEVIN D.: Mesh-independent surface interpolation. *Geometric Modeling for Scientific Visualization* (2003), 37–49.
- [MA97] MOUNT D., ARYA S.: Ann: A library for approximate nearest neighbor searching, 1997.
- [MCSM99] MORRIS R., CHEESEMAN P., SMELYANSKIY V., MALUF D.: A bayesian approach to high resolution 3D surface reconstruction from multiple images. In *Proceedings of IEEE Signal Processing workshop on Higher-Order Statistics* (June 1999), pp. 140–143.
- [PKG03] PAULY M., KEISER R., KOBELT L. P., GROSS M.: Shape modeling with point-sampled geometry. *ACM Trans. Graph.* 22, 3 (2003), 641–650.
- [PMG*05] PAULY M., MITRA N., GIESEN J., GROSS M., GUIBAS L.: Example-based 3D scan completion. In *Symposium on Geometry Processing* (2005), pp. 23–32.
- [RA99] RAMAMOORTHI R., ARVO J.: Creating generative models from range images. In *Proceedings of SIGGRAPH 1999* (1999), ACM, ACM Press / ACM SIGGRAPH, pp. 195–204.
- [RL01] RUSINKIEWICZ S., LEVOY M.: Efficient variants of the icp algorithm. In *Third International Conference on 3D Digital Imaging and Modeling (3DIM)* (June 2001).
- [SACO04] SHARF A., ALEXA M., COHEN-OR D.: Context-based surface completion. *ACM Trans. Graph.* 23, 3 (2004), 878–887.
- [SF01] SUK T., FLUSSER J.: Features invariant simultaneously to convolution and affine transformation. In *CAIP '01: Proceedings of the 9th International Conference on Computer Analysis of Images and Patterns* (2001), pp. 183–190.
- [SMKF04] SHILANE P., MIN P., KAZHDAN M., FUNKHOUSER T.: The princeton shape benchmark. In *Shape Modeling International* (2004), pp. 167–178.
- [Tau95] TAUBIN G.: A signal processing approach to fair surface design. In *Proceedings of SIGGRAPH 1995* (1995), ACM, ACM Press / ACM SIGGRAPH, pp. 351–358.
- [TL94] TURK G., LEVOY M.: Zippered polygon meshes from range images. In *Proceedings of SIGGRAPH 1994* (1994), ACM, ACM Press / ACM SIGGRAPH, pp. 311–318.
- [Tur90] TURK G.: Generating random points in triangles. *Graphics gems* (1990), 24–28.
- [TWBO02] TASDIZEN T., WHITAKER R., BURCHARD P., OSHER S.: Anisotropic geometric diffusion in surface processing. In *Proceedings of IEEE Visualization* (2002).
- [WK01] WERMAN M., KEREN D.: A bayesian method for fitting parametric and nonparametric models to noisy data. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 5, 23 (May 2001), 528–534.
- [WSI98] WHEELER M., SATO Y., IKEUCHI K.: Consensus surfaces for modeling 3D objects from multiple range images. In *Int. Conf. on Computer Vision* (1998), IEEE Computer Society, pp. 917–924.