



Automatic Procedures to Create CAD Models from Measured Data

Tamas Varady

Geomagic, Inc., varady@geomagic.com

ABSTRACT

Digital shape reconstruction is the process of creating digital models from physical parts represented by 3D point clouds. The ideal process is expected to provide a boundary representation that is likely to be identical or similar to the original design intent of the object, and requires minimal user assistance. This paper discusses alternative state-of-the-art approaches, where emphasis is put on automatic methods (i) to create complete and consistent topological structures over polygonal meshes; and (ii) extract accurate and properly aligned surface features that yield complete, trimmed CAD models with fillets and corner patches. Problems and recommended solutions will be presented through case studies using industrial parts.

Keywords: shape sampling, RE, segmentation, functional decomposition, redesign over meshes.

DOI: 10.3722/cadaps.2008.577-588

1. INTRODUCTION

Digital Shape Sampling and Processing (DSSP) is a rapidly advancing technology [6, 11] that enables the conversion of physical objects into a digital shape representation. There are two basic technology components. Using non-contact 3D scanners, today it is possible to collect millions of points from the surface of objects with high performance and high accuracy. Using the latest software systems users can capture the shape and structure of complex objects. These systems run on high performance graphic computers, and it is possible to accomplish jobs within a reasonable time, that would have been almost impossible a few years ago. DSSP opens a wide range of applications including design, analysis, inspection and manufacturing of engineering parts; medical applications, such as digital dentistry or custom prosthetics, and other areas, such as preserving the historical and cultural heritage of mankind.

DSSP is considered as a complementary technology of CAD/CAM. CAD systems apply prescriptive modeling methods to create and combine 2D and 3D entities through a sequence of operations. They convert user specifications into an idealized digital form, but are limited when existing physical objects need to be represented. DSSP systems offer descriptive modeling methods to extract shape information from measured data and produce high-quality digital models in the virtual world.

Digital Shape Reconstruction (DSR) deals with the computational geometry aspects of converting point clouds into a digital representation. (Formerly, the term reverse engineering was used.) Objects can be classified into three basic categories. *Organic objects* are formed by the laws of nature and show an amazing variety in structure and shape. This group also includes parts of the human body that need to be replaced or matched - prosthetic devices, dentures, ear plugs, etc. *Artistic objects* have been created by humans, and show an incredible variety, as well. They are judged primarily by appearance and aesthetic values; for restoration they may need to be re-fabricated or for the entertainment industry they need to be placed into a virtual environment. *Engineering objects* have been created by engineers, and satisfy different requirements of form and function; obey physical laws (e.g. turbine blades); meet complex mechanical constraints, (e.g. multi-axis gear boxes) and satisfy aesthetic requirements (e.g. car bodies or consumer goods). These objects are strongly linked to manufacturing processes, which determine a well-defined coverage of bounding surfaces and a global structure, in which these surfaces are organized. Engineering objects are

typically bounded by relatively large, *primary surfaces* that are connected by highly-curved *transitional features*, such as fillets or free-form steps.

The focus of this paper is to present state-of-the-art methods that automate the process of reconstructing engineering objects and meet the requirements of downstream applications in CAD/CAM/CAE. It is useful to classify these objects: (i) *Prismatic* objects are typically created by Boolean combinations of volumetric primitives, each bounded by simple analytic surfaces or created by sweeping 2D profile curves. (ii) *Free-form* objects are typically created by parametric surfaces with more design freedom and styling variety; these are built together by intersections and other complex operations. (iii) *Hybrid* objects are created by mixing conventional and free-form geometries, and typically combine volumetric and surfacing operations. This is not a clear-cut classification, but it may help characterizing the dominant surface geometry and the scope of operations. Three examples are shown in Fig. 1.




Fig. 1: (a) Prismatic object, (b) Free-form object, (c) Hybrid object.

There are a few issues which will reoccur during our discussion. The *quality of measured data* - noise, completeness and sampling density - strongly influences the whole reconstruction process. The *amount of manual work* is a critical factor for creating complex parts. While automatic procedures are more efficient, they may lack perfection in structures and surface qualities. Restoring the *original design intent* is another strong demand in various applications; in principle, DSR is supposed to generate exact mathematical surfaces with well-defined dependencies and a faithful, topological structure, being conform to standard data formats.

Integration with CAD is also an important issue. There are systems that can produce a full CAD model independently, and act as a preprocessor for CAD/CAM applications. Other approaches are fully integrated into a CAD and utilize the system's powerful operations. A third option is to extract the basic geometric information in the course of DSR, but finalize surface trimming, filleting and stitching only within a CAD system. Finally, it is also possible to transfer information in both ways DSR and CAD, utilizing the best of each.

The paper is structured in the following way. Five different approaches leading to different digital representations are analyzed in Section 2, including functional decomposition and redesign over meshes. In Section 3 the process of functional decomposition is described in more detail, showing automatic steps to reconstruct structure and surface geometry. In Section 4, a few "hot" topics will be raised, leading DSR towards its ultimate goal of generating high quality shapes in a fully automatic manner.

2. FIVE APPROACHES FOR DIGITAL SHAPE RECONSTRUCTION

2.1 Polygonal Meshes

Modern 3D scanners provide point clouds with high density and high accuracy. The typical workflow is that partial scans are aligned and merged, and a triangular polygonal mesh is generated as the first, piecewise linear approximation of the object. Polygonal meshes are ideal for various operations including mesh repair, noise reduction and hole filling. Polygonal meshes can be appropriately simplified (decimated) to reduce size; while the details of the shape can also be preserved. Polygonal meshes can be smoothed, deformed and enhanced by special "perfecting" operations, such as sharpening or aliasing along highly curved features. The majority of segmentation procedures are

also performed over meshes. There is an enormous literature on mesh processing; important algorithms have been collected in [1,5].

The following four DSR approaches in this section all use polygonal mesh input to extract surface geometry and build CAD models. At the same time, the polygonal representation in itself is an adequate form for many applications, in particular, when there is no need to build high-quality surface models or recover the design intent. Triangular meshes are uniform and simple, and can be efficiently processed. One basic rationale to apply them is often quoted: rendering, FE analysis, and manufacturing utilize preprocessed discrete data, so why waste time and energy to build complex surfaces in a laborious way and then switch back to discretized data again.

Polygonal meshes are widely used in applications where organic and artistic shapes need to be processed, but show their limitations when segmented engineering objects need to be extracted, redesigned and manufactured.

2.2 Quadrilateral Patch Layouts by Automatic Surfacing

Automatic surfacing procedures create complex quadrilateral patch layouts over polygonal meshes with high computational efficiency and minimal user input. After extracting a few, characteristic subdividing feature lines on the mesh a quadrilateral curve network is generated in a fully automatic manner. Each quadrilateral is approximated by a NURBS surface, and the neighbors are stitched together, enforcing tangent plane continuity (G1). The strength of this approach is the uniformity of the representation and the existence of watertight connections, which is heavily exploited in applications, such as finite element analysis. The amount of user effort is only a fraction of the more advanced methods, so this approach is attractive when a continuous surface representation without segmentation is needed in a short time (Fig. 2). This approach aims at capturing the exact shape versus perfecting the shape, as the next three procedures.




Fig. 2: Quadrilateral patch layout for a hybrid object.

The weakness of this representation is mainly due to the uniform subdivision by quadrilaterals, but the "natural" faces of engineering objects are non-four-sided and contain hole loops. Superfluous subdivisions lead to a loss of internal continuity and limit surface quality. Even the simplest analytic surfaces are represented as a collection of approximating NURBS patches. The quadrilaterals are generally not aligned to the highly-curved features of the object, and undesirable surfacing artifacts may occur when single NURBS patches need to approximate highly-curved and flat surface portions simultaneously.

It is a recent advance in automatic surfacing that tiling is preceded by a feature-based segmentation [8]. Then the quadrilaterals are properly aligned with the features of the object, which leads to a more natural, CAD-like structure and optimized NURBS patches. Highly curved parts are covered by a densely distributed row of quadrilaterals, flat parts by a smaller number of patches. At the same time computational efficiency, moderate user involvement and watertight representation are all preserved.

The quadrilateral surface representation is frequently used to process organic and artistic objects, when smooth surfaces are important for manufacturing or other reasons. In mechanical engineering, quadrilaterals are also well-suited to reconstruct prismatic, free-form and hybrid objects, in particular, when fast, automatic solutions supersede requests to recover design intent and high quality surfaces. It must be emphasized that automatic surfacing does work for very complex objects as well, where manual approaches may fail or require an extremely long time. (Representative commercial systems include Geomagic Studio [8], Rapidform XOR [13], or PolyWorks [12].)

2.3 Manually Segmented Surfaces

Manual segmentation has been and still is the classical way of approaching the reverse engineering problem. Users draw a segmenting curve network over the mesh, and approximate the created surface regions step by step. The construction and the exact positioning of these curves are critical, as mislocation of the boundary curves may destroy surface quality. In this approach mostly, but not exclusively four-sided surfaces are created and surface elements get stitched in a pair-wise manner enforcing different degrees of continuity. Advanced feature curve extracting tools can help the process, but overall, the creation of complex free-form curve networks is very laborious, and requires much iteration. This approach is still widely used in industry, since the recently emerging automatic segmentation systems are not yet widely distributed, and the automatically created networks are often incomplete or not sufficiently accurate for complex free-form objects.

Manual surfacing is laborious, but users gain a full confidence that they create models exactly the way they want, and thus design intent is generally recaptured. Most advanced systems offer a special fairing-smoothing toolkit to perfect surface geometries. The process may take several days (or weeks) of manual tuning, but eventually it is possible to create perfect surfaces, often called Class A, which are needed in the automotive industries or elsewhere where aesthetics and perfect reflection line distribution are fundamental. (Representative commercial systems include Alias [2], Icem SURF [9] or Tebis [15].)

Manual segmentation is not well-suited to reconstruct prismatic objects, where the shape is dominated by simple analytic surfaces and their intersections, which cannot be easily and accurately reconstructed over the mesh. Prismatic objects may have several hundred faces; and for such complex topology the manual approach would mean disproportional effort. For reconstructing free-form geometry the use of manual segmentation is still justified, as objects may need to be subdivided by many smooth, segmenting edges, which often cannot be easily extracted through surface-surface intersections.

2.4 Functionally Decomposed Surface Models

The functional decomposition paradigm utilizes the dominant structure of the vast majority of engineering parts, where objects are being bounded by a well-defined hierarchy of surfaces. *Primary* surfaces are independent geometric entities, which define large portions of the shape as determined by functional or aesthetic requirements. Between adjacent primaries special *secondary*, or connecting surfaces can be found, such as fillets or free-form steps. Secondary surfaces are "glued" onto the primaries and are connected to them by tangent or curvature continuity, thus these are dependent surface elements. Finally, at the junctions where connecting surfaces come together, *corner patches* need to be inserted, which are smoothly connected to the neighboring secondary surfaces.

It is also true that the majority of connecting surfaces represent highly curved surface portions in contrast to primaries, which are relatively flat. This is the basis of *automatic segmentation* methods, which detect the primary regions and - at the same time - create a network of connecting features and corner patches. The power of functional decomposition it always processes a topologically valid structure of regions, which is likely to correspond - one-to-one - to the face structure of the final CAD model. The segmented structure may need to be manually enhanced, when the mesh is noisy or not sufficiently dense. In other cases to obtain smooth internal subdividing edges some manual assistance may be needed, since these cannot be located by the sudden change of curvature.

The complete topological model over the mesh drives the creation of B-rep models. After region classification (analytics, swept and free-form surfaces) first primary surfaces are fitted, followed by the automatic generation of connecting surfaces and corner patches. There is no need to touch them one by one: neighbors are given and the related parameters, such as connection width or filleting radius can be extracted from the underlying mesh. In the course of segmentation, regions with well-defined edge loops have been created. The edge loops of the corresponding trimmed primary surfaces will be determined through the fitting of connecting surfaces, whose boundaries will serve as

trimming curves. In this way, the final trimming and stitching is also automatic, and the CAD model reflecting the segmented structure is ready. (This technology has been developed and released in Geomagic Studio 10 [8], further details of the process will be given in Section 3.)




Fig. 3: Structured CAD model using functional decomposition: primaries (gold), connections (red).

The main benefit of functional decomposition is its efficiency to create complete, trimmed B-rep models. It is applicable for fairly complex prismatic and free-form objects. Untrimmed surfaces, profile curves, and other auxiliary elements are created as a by-product of the process, and if needed, these entities can be perfected, including the alignment of planes or rounding a radius. One limitation of the approach, however, is that it represents a global structure with automatically extracted entities. There is no history file and no global parameters are associated with the CAD model, so modifying the model is limited to local changes, which basically retain the original segmented structure. This deficiency can be overcome by transferring the extracted information to external CAD packages.

2.5 CAD Models Redesigned Over Meshes

A new integrated approach for shape reconstruction has been suggested recently exploiting the full functionality of CAD packages (Rapidform XOR [13], Solidworks' ScanTo3D [14]). The basic idea is to perform exactly the same CAD sequence as one would perform in prescriptive CAD, but here geometric information is extracted from the underlying mesh. As an analogy, imagine a hand sketched drawing, which is redrawn over the sketch using compasses and rulers.

The method is particularly useful for complex prismatic objects, which are typically created in three consecutive phases in the course of conventional CAD. First volumetric primitives are defined, second these are combined by Boolean operations, and finally sharp intersection edges are filleted. In our context, the majority of volumetric primitives are simple extrusions or surfaces of revolution. Intersecting the mesh by datum planes yields various polylines, which can be converted into straight-arc profiles using the profile editing operations of the CAD system. Adding further parameters, such as the direction of the extrusion or the rotational axis, completes the definition of these primitive solids which will be united or subtracted in the forthcoming steps. One disadvantage of this approach is that every single operation must be explicitly performed; profiles must be edited and constrained; volumes need to be parameterized, added or subtracted; blending radii need to be numerically specified for each edge or edge loop. At the same time it is very valuable that this sequence creates a history file with prescribed parameters for the individual entities. This can help later to create modified objects, which are derivatives of the original measured one; having the history file it is possible to create similar objects with different parameters, as one would do in ordinary CAD.

The redesign approach is not limited to volumetric build up, and it is possible to access and use all important surfacing operations, as well, including surface fitting over a point region, or lofting through extracted section curves. The main point is that the structure and the geometric entities are literally created *by the user and the Boolean engine* step by step. The process can be laborious and error prone; in particular, for objects with many faces and many filleted edges, but the history file gives great flexibility. The redesign approach is much harder when complex free-form objects with

several smooth subdividing edges need to be reconstructed; in this case the volumetric approach is quite tedious, and the Booleans cannot be applied to create the best structure and parameterization for the object.

2.6 Summary

To select the most suitable DSR approach is a difficult task. A high-level comparison is given below:

	<i>Surface quality</i>	<i>Amount of manual work</i>	<i>Design intent</i>
<i>Polygonal meshes</i>	planar approximation	moderate	no
<i>Quadrilateral patch layouts</i>	medium (but watertight)	minimal	no
<i>Manually segmented surfaces</i>	(very) high	very high	yes (mostly)
<i>Functional decomposition</i>	high	moderate	yes
<i>Redesign over meshes</i>	high	high	yes + history

Requirements often contradict, and users must carefully balance the significance of surface quality, the amount of required manual work and the final model structure. Readers may want to make their own choice concerning the examples shown in Fig. 4.




Fig. 4: (a) Part of a prismatic object, (b) Part of a free-form object.

3. THE FUNCTIONAL DECOMPOSITION PARADIGM

Functional decomposition coupled with automatic segmentation is a new, emerging technology, and it is worth discussing the related workflow of creating trimmed surface models. We are going to review the whole process and point at some difficult geometric problems that must have been solved. Our input is a preprocessed mesh carrying the measured point cloud information. We assume that noise has been filtered, artifacts have been removed, and the mesh has been properly decimated to be able to recover relatively small geometric features as well.

We identify 10 basic steps in the workflow, as follows: 1. Region detection, 2. Contour Detection, 3. Surface structure generation, 4. Region classification, 5. Primary surface fitting, 6. Primary surface enhancement, 7. Transitional surface fitting, 8. Quality assurance, 9. Final trimming and stitching, 10. Export to CAD. The various stages of the process are illustrated in Fig. 5.

Step 1: Region detection. Our automatic segmentation method is based on Morse theory that studies smooth functions over manifolds. A Morse complex partitions the manifold into a collection of regions by a network of curves that connect non-degenerate critical points (minima, maxima, and saddles) of a given function. We use a special curvature indicator, whose values are estimated at each vertex of the polygonal mesh, and create a piecewise linear function that highlights the highly curved parts of the shape (Fig. 5a). Utilizing concepts of persistence and prioritization [7], a hierarchy of topologically simplified segmentations is obtained. Each particular segmentation is determined by a sensitivity value, and represents a valid topological structure of "monotonic" regions bounded by a network of polylines. These polylines run on connected triangle edges and are typically ragged. We perform additional mesh

operations to cover these ragged boundaries by strips of triangles – called *separator sets*, and simultaneously shrink the original monotonic regions [18]. Fig. 5b shows the segmented mesh: the relatively flat primary regions (in light colors) are separated by triangle strips associated with highly curved transitions (in red). If necessary, separator sets can be manually edited using paintbrush operations.

Step 2: Contour Detection. We automatically construct a smooth curve network over the mesh that runs in the middle of separator sets. Special curve tracing algorithms are applied to determine trajectories of a local translational vector field [18]. The network is called a *feature skeleton*, and it reflects the feature structure of the object (Fig. 5c). The edges are often called *contours* in this context, and are classified as either *smooth* or *sharp*. Smooth contours correspond to wide separator sets and will guide the generation of connecting features between adjacent primaries, such as fillets. Sharp contours typically occur within narrow separator sets and will be converted into single edges determined by surface-surface intersections. We also classify the vertices of the feature skeleton; in particular, we identify *collinear* contour pairs that approach a vertex from opposite directions with tangential continuity; for example, a T-node within the feature skeleton is degree-3 vertex with one pair of collinear edges. If necessary, the feature skeleton network can be manually edited, modified or smoothed using interpolating B-splines over the mesh.

Step 3: Surface structure generation. In this phase we create the final topology of the connecting features and corner patches (Fig. 5d). First, we replace each smooth contour by a pair of longitudinal boundaries using a special tracing process, that is guided by the middle contour and produces "quasi parallel" curves on the mesh that smoothly approximate the separator sets. The algorithm guesses whether the current feature is of constant or variable width, and traces accordingly [18].

Corner patches are to be represented by *setback vertex blends*, which have been introduced a long time ago to assure smooth transitions for complex junctions with uneven radii and uneven geometric configurations, e.g. [16]. Unlike an "ordinary" suitcase corner where 3 fillets are connected with an exactly 3-sided vertex blend, in the case of setback vertex blends, we "set back" the termination of the connections from the common vertex, and insert a larger, more flexible connecting patch that may have 6 ($=2n$) sides. The loop of such a vertex blend may consist of an alternating sequence of profile curves and spring curves. *Profile curves* terminate the connections, *spring curves* connect the pairs of adjacent connection boundaries lying on the same primary region.

Fig. 6 illustrates how setback vertex blends work. The first picture shows a feature skeleton with 3 edges, including a collinear edge pair. An ordinary vertex blend would have only 3-sides, but would distort the natural cross-sectional flow of the transitional surfaces. Inserting a spring curve due to collinearity (bottom left), and then inserting another spring curve (bottom right), create more natural structures with 4 or 5 sides, respectively, and will enhance the overall surface quality of our CAD model when these corner patches are fitted. As in the previous steps, this extended curve net structure is automatically created; if the algorithm creates imperfect results or the user would like to reconfigure the network, this is possible using manual operations.

Step 4: Region classification. The structure of the model with related dependencies has been completed in the previous phase. The next step is to assign a likely surface type to each region (Fig. 5e). This classification is based on methods published earlier [3] using local normal vector estimations and analyzing their distribution on the Gaussian sphere. A sequence of hypothesis is set up, moving from simple to more complex types, and as a result, regions will be classified as planes, cylinder, cones, spheres, extrusions, drafted extrusions or surfaces of revolution. The remaining ones, which do not show any geometric regularity, will be classified as free-form.

Classification is also automatic, and can be overruled, if the user wishes to do so. Typically this happens (i) when the segmenting boundaries are poorly located, (ii) the data set is very noisy, or (iii) the user wants to enforce a particular surface type by his preference.

Step 5: Primary surface fitting. This can be performed in one go for the whole object or in selected groups by type or location. Least-square fitting of analytic surfaces is a well-known topic in the literature; however, the fitting of simple swept surfaces (extrusions, drafted extrusion and rotational surfaces) is more complex. For example, take the extruded region at the bottom of the blade object shown in Fig. 7. The direction of extrusion has already been computed during classification. For surface creation a profile curve needs to be determined in a working plane orthogonal to the sweeping direction.

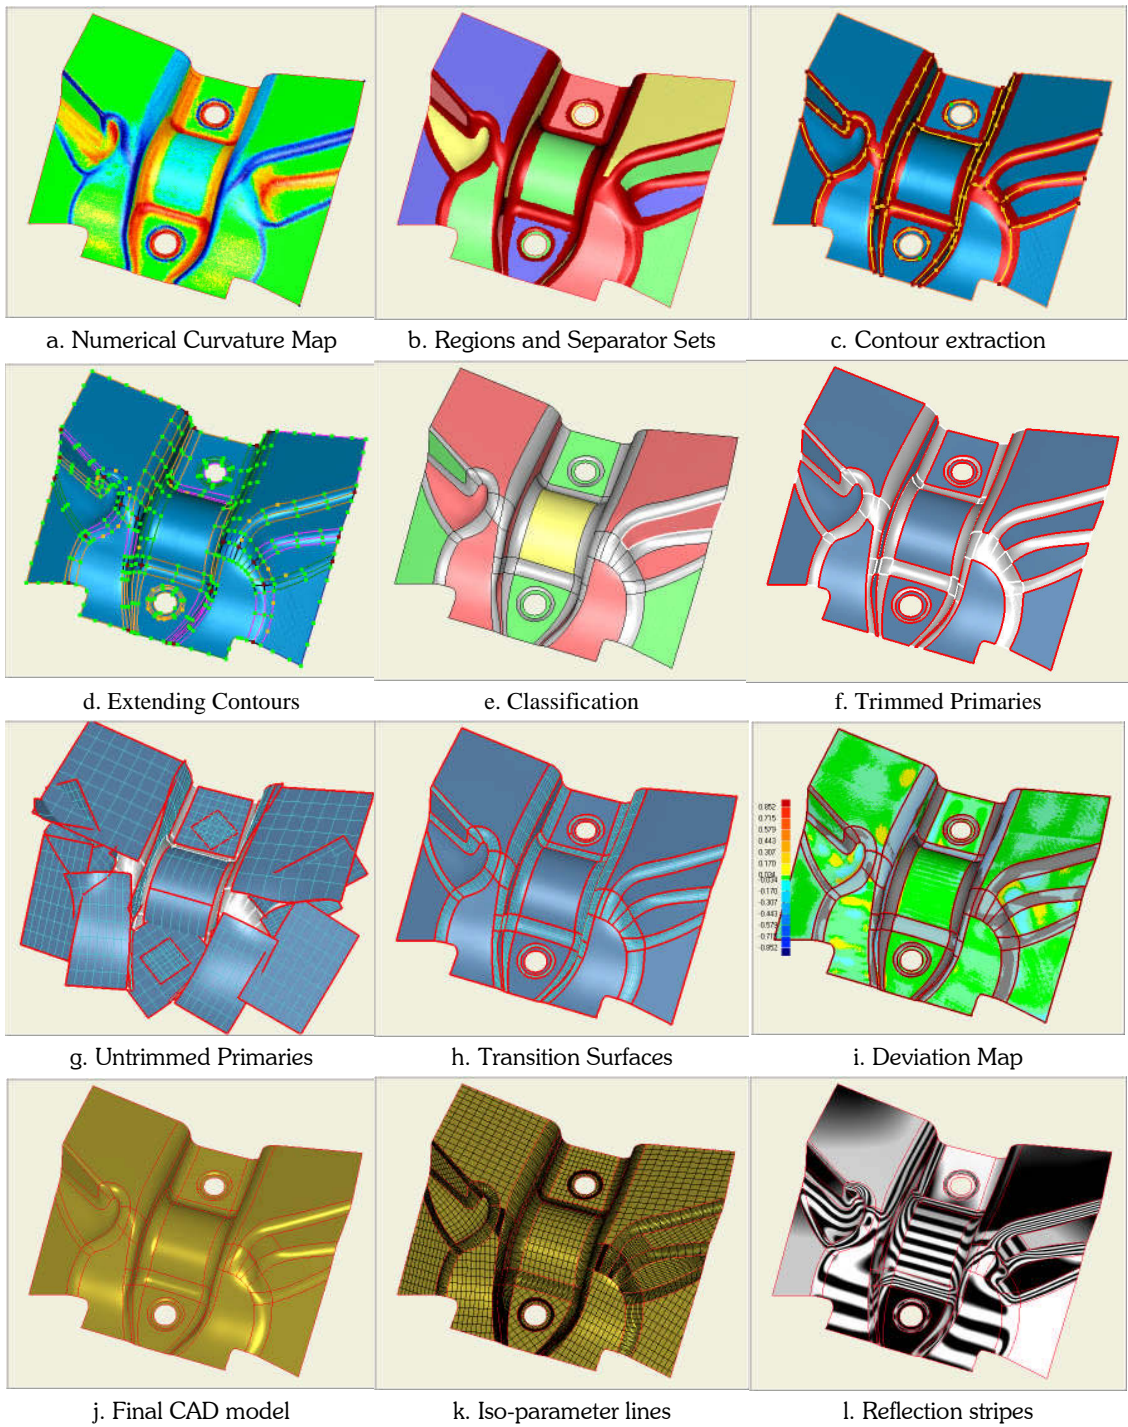


Fig. 5: Stages of the Trimmed Surface Creation workflow.

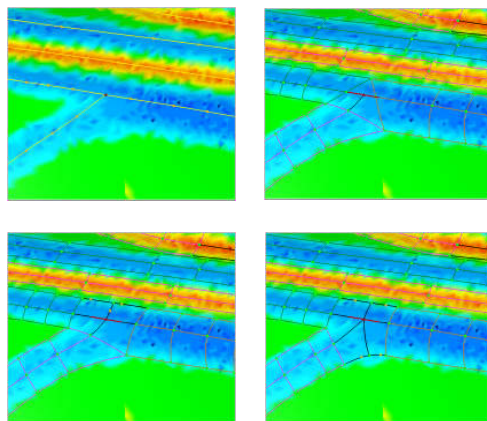


Fig. 6: Corner patch variants using setbacks.

The profile can either be a free-form curve or a composite line-arc curve, typically constrained to be tangent continuous. The system first suggests a segmentation, i.e. computes a sequence of marker points that correspond to the endpoints of consecutive segments, then creates a curve approximating the underlying segmented data points. This can be done by well-known constrained fitting methods [4], where the roughly positioned markers (Fig. 7a) will be shifted to their exact position (Fig. 7b), and the profile is ready to create a final extruded surface.

The method of creating swept surfaces is another good example to illustrate the difference between functional decomposition and the redesign approach. In redesign a particular planar section is taken, the user defines lines and circles over a polyline; then constrains them and the system determines the best, accurate profile from these abstract elements. In functional decomposition, an automatic algorithm finds a good, possible structure of the profile, and the goal of automatic fitting is to *simultaneously* satisfy deduced constraints and well approximate the underlying data points. Moreover, not only a particular sectional polyline, but all data points within the given region need to be approximated.

Surface fitting for free-form surfaces may depend on various parameters, such as the number of control points, smoothness weight, tolerance, outlier percentage or enlargement ratio. The paradigm requests to always offer an automatically computed surface candidate for free-form surfaces, as well. The system computes an initial parameterization, and being guided by a prescribed tolerance value it optimizes the number of control points and the smoothing factor to get a good surface approximation. As before, if the user is not satisfied with the result, he can manually adjust the parameters and refit.

The result of this phase is a collection of primary surfaces, which are available in their trimmed form (Fig. 5f) or untrimmed form, which extends beyond the related regions of the mesh (Fig. 5g)




Fig. 7: (a) Segmented profile polyline, (b) Fitted profile curve, (c) Extruded surfaces.

Step 6: Enhancing primary surfaces. Before starting transitional surface fitting, it is advised to enhance or perfect the primary surfaces. For analytic surfaces this may mean aligning the directions or axis lines to the principle coordinate axes, or rounding/snapping various values, such as angles or radii of cylinders. This can also be performed in an automatic manner driven by positional and angular tolerances. We have already discussed the possibility to enhance profile segmentation for simple sweeps. The enhancement of free-form surfaces typically requires more manual effort. Without going into details of all possible options, an example is enhancing surface quality through modifying its parameterization, which is one of the most crucial issues in free-form surface fitting. There are many practical situations where the boundaries of a given 3D region and the sides of the parametric domain can be aligned in a sensible manner. This process is called *labeling*; the initial labels are computed by the system, but they can be overruled by the user. In our example (Fig. 8), the boundaries of a region are labeled in CW order by green, blue (2 pieces), red and yellow (2 pieces). As a result, the parameterization of the data points will be such that the iso-parameter lines of the fitted surface will be aligned with the labeled boundaries, and in this way both accuracy and surface quality will be significantly improved.




Fig. 8: Labeling leads to enhanced surface quality.

Step 7: Transitional surface fitting. Having a set of primary regions, now it is possible to generate transitions. For each transitional region, the system checks whether its neighbors have already been created and if yes, fitting will be performed (Fig. 5h). The transitions between primary surfaces are always four-sided, the degree of continuity can be prescribed (tangent or curvature continuity) and constrained fitting will be automatically performed. Corner patches are setback type vertex blends, and their structure has been discussed earlier. Note, that along the profile curves these need to connect smoothly to the adjacent transitions, while along the spring surfaces they need to smoothly connect to adjacent primary surfaces. Depending on the data set a global setting to strengthen accuracy in contrast to smoothness can also be set. Having the transitional surfaces, it rarely happens that the user would like to overwrite the surface parameters. One special case, though, is when a more complex transitional surface, such as a free-form step, needs to be created; this may require increasing the number of cross-sectional control points to get a more accurate approximation.

Step 8: Quality assurance. In order to accept the reconstructed model the inspection of the quality and accuracy of surfaces is quite important. There are many well-known graphic rendering options to show curvature distributions and highlight potential surfacing artifacts, including texturing, isophotes or reflection lines (see Figs. 5k and 5l). Concerning accuracy, the standard approach is to create deviation maps (Fig. 5i), by means of which errors or suspicious areas to be improved can be located.

Step 9: Final trimming and stitching. As it was discussed earlier functional decomposition works over a global topological structure all the time, carrying all topological and geometrical information needed for building a boundary representation. Once the individual surfaces have been created, trimming and stitching them together is fairly straightforward (Fig. 5j).

Another note, at this time in favor of the redesign approach: for a functionally decomposed model, a complete underlying polygonal mesh is needed, since the whole topological structure is built over the mesh. Consequently, problems may arise when there are large portions of missing data which must be artificially

filled up. Redesign is not sensitive to this problem, since unknown areas with missing data can simply be defined by explicit CAD operations.

Step 10: Export to CAD. Having completed a functionally decomposed model, all related entities with the full topology can be exported in standard data exchange formats, or transferred through internal communication channels to various CAD systems in their specific data formats.

Note: There is an alternative workflow in functional decomposition to *extract untrimmed surfaces* that are to approximate only the primary regions of the segmented structure. In this case, all related information is transferred to a CAD system and trimming, filleting, stitching will be performed in CAD. This workflow is significantly simpler than the previous one, and four steps (2, 3, 7 and 9) are skipped: 1. Region detection; 4. Region classification; 5. Primary surface fitting; 6. Primary surface enhancement; 8. Quality assurance; 10. Export to CAD. The time and effort saved in this workflow is due to the fact that there is no need to create a feature skeleton and a full segmenting curve network. For primary fitting the user just selects regions obtained in the first segmentation step, and takes the points surrounded by the loops of the separator sets, but no transition surfaces are fitted. Classification, surface enhancement and quality assurance take place in the same way as before.

4. CONCLUSION AND TOPICS FOR FUTURE RESEARCH

DSSP technologies have progressed at an incredible pace in the last decade, and the latest 3D scanners and software systems can solve very complex and difficult problems of which we did not dare to think earlier. The range of applications is also rapidly widening, and DSSP penetrates to more and more areas of engineering, medical sciences, and art and entertainment. In spite of the significant advances, the basic demand to get better shape quality, in a shorter time and with minimal user assistance prevails. Amongst many others, we propose three topics for future research that will definitely contribute to reaching the above goals of digital shape reconstruction.

Models with sharp edges. In many applications users may want to get idealized objects with or without fillets. In other words, having a mesh and a reconstructed model with filleted and sharp edges, users would like to switch the individual entities, i.e. replace fillets by edges of intersection or vice versa. Though it may sound simple this is a very complex problem, as the sharpening operations must guarantee the geometric and topological consistency of the model. A simple example is a degree 4 vertex determined by 4 intersecting surfaces, where ideally 4 sharp edges meet in a single point. But we have measured data with noise, so it is unlikely that the four surfaces will go through a single point, and thus a very small, artificial edge needs to be incorporated into the data structure. One may say, no problem, let us apply a constrained fit enforcing a common vertex for the four surfaces, but just take the various curved surface types and see that even this is not so easy. Going further, imagine that there are many degree 4 vertices, alignments and smoothness constraints, which may lead to a huge, often over-constrained system of many surface elements. In fact, switching from sharp edges to fillets seems somewhat easier, as in this case the topology can locally be restructured by inserting a transition surface and two modified corner patches at the two ends. However, consistency problems may occur here, as well, when one transition surface overlaps with another and their interference needs to be handled.

High-level segmentation. It is not only the robustness and efficiency of segmentation methods that need to be enhanced, but currently some basic functionality is missing. Feature based modeling is a fundamental concept in modern CAD/CAM, and we envision that digital shape reconstruction will move to this direction, as well. It means that not only individual regions will be detected, classified and fitted, but rather face groups forming a particular composite feature will be detected as a single entity; and its faces will be locally constrained and fitted simultaneously. The feature-based shape reconstruction is going to lead to more concise models, where design intent will be expressed not by means of individual surface entities, but at a high abstraction level.

Perfecting models. This is also an important area with many diverging problems to be solved. One area is to detect likely geometric constraints, and adjust the model accordingly. For example, if two planes are almost parallel, set them to perfect parallel. If an axis is almost perpendicular, set it to perpendicular. If a radius is almost a unit number set it to an accurate value. If there is a pattern of holes, that seem to be very similar to each other, set them uniformly. If a part shows symmetry, fit a surface in a special way, and enforce exact mirrored geometries meeting tangentially along their common seam line, etc. The examples show that detecting likely geometric constraints and perfecting models accordingly is a very complex issue. The other sort of perfection involves mainly free-form geometry. Currently,

"perfectly smooth" surface qualities can only be achieved through very tedious, iterative manual methods. The fairing literature deals only with single surface entities. Instead of this, high level surface fairing for segmented models, where the dependencies between the adjacent surfaces can also be taken into consideration, would mean a significant step forward.

5. ACKNOWLEDGEMENTS

The functional decomposition algorithm has been developed and implemented by Geomagic's international engineering team residing in North Carolina and Hungary, and the technology has been released in the product Geomagic Studio 10 (Shape and Fashion modules). All illustrations of this paper have been prepared using this system. The author would like to acknowledge many fruitful and thought provoking discussions with members of the whole team, in particular, with Dr. Mike Facello. The research efforts have been supported by two NSF-SBIR grants, namely Award #0450230, "Creating functionally decomposed surface models from measured data" and Award #0521838, "Applications of Morse theory in reverse engineering."

6. REFERENCES

- [1] M. Attene, M.; Katz, S.; Mortara, M.; Patane, G.; Spagnuolo, M.; Tal, A.: Mesh Segmentation - A Comparative Study, IEEE Int'l Conf. on Shape Modeling and Applications, (SMI'06), 2006.
- [2] Autodesk/Alias: <http://usa.autodesk.com/>
- [3] Benko, P.; Varady, T.: Segmentation methods for smooth point regions of conventional engineering objects, *Computer-Aided Design*, 36, 2004, 511–523.
- [4] Benko, P.; Kos, G.; Varady, T.; Andor, L.; Martin, R. R.: Constrained fitting, *Computer-Aided Geometric Design*, 19, 2002, 173–205.
- [5] Botsch, M.; Pauly, M.; Kobbelt, L.; Alliez, P.; Levy, B.: Geometric modeling based on polygonal meshes, Int'l Conf. on Computer Graphics and Interactive Techniques, ACM SIGGRAPH courses, 2007.
- [6] DSSP: <http://www.geomagic.com/en/dssp/resources/>
- [7] Edelsbrunner, H.; Letscher, D.; Zomorodian, A.: Topological persistence and simplification, *Discrete Computational Geometry*, 28, 2002, 511–533.
- [8] Geomagic Studio: <http://www.geomagic.com/>
- [9] Icem SURF: <http://www.icem.com/>
- [10] Imageware: http://www.plm.automation.siemens.com/en_us/products/nx/styling/reverse_eng.shtml
- [11] Marks, P.: Capturing a Competitive Edge through Digital Shape Sampling & Processing (DSSP), SME, Blue Book Series, 2005.
- [12] Polyworks: <http://www.rapidform.com/>
- [13] Rapidform XOR: <http://www.rapidform.com/>
- [14] Solidworks Scanto3D: <http://www.solidworks.com/>
- [15] Tebis: <http://www.tebis.com/>
- [16] Varady, T.; Hoffmann, C. M.: Vertex blending: problems and solutions, *Mathematical Methods for Curves and Surfaces II*, Eds: Daehlen, M.; Lyche, T.; Schumaker, L. L., Vanderbilt University Press, 1998, 501–527.
- [17] Varady, T.; Martin, R. R.: Reverse Engineering, Chapter 26, In: *Handbook of Computer Aided Geometric Design*, Eds.: Farin, G.; Hoschek, J.; Kim, M.-S., North-Holland, 2002, 651–681.
- [18] Varady, T.; Facello, M. A.; Terek, Zs.: Automatic extraction of surface structures in digital shape reconstruction, *Computer-Aided Design*, 39, 2007, 379–388.