Global Surface Reconstruction
By Purposive Control of Observer Motion

Kiriakos N. Kutulakos
kyros@cs.wisc.edu

Charles R. Dyer
dyer@cs.wisc.edu

Computer Sciences Department
University of Wisconsin
Madison, Wisconsin 53706 USA

Technical Report #1141
April 1993

Abstract

What real-time, qualitative viewpoint-control behaviors are important for performing *global* visual exploration tasks such as searching for specific surface markings, building a global model of an arbitrary object, or recognizing an object? In this paper we consider the task of *purposefully* controlling the motion of an active, monocular observer in order to recover a global description of a smooth, arbitrarily-shaped object.

We formulate global surface reconstruction as the qualitative task of controlling the motion of the observer so that the visible rim slides over the maximal, connected, reconstructible surface regions intersecting the visible rim at the initial viewpoint. We show that these regions are bounded by a subset of the visual event curves defined on the surface.

By studying the epipolar parameterization, we develop four basic behaviors that allow reconstruction of a surface patch around any point in a reconstructible surface region. These behaviors control viewpoint to achieve and maintain a well-defined geometric relationship with the object's surface, rely only on information extracted directly from images (e.g., tangents to the occluding contour), and are simple enough to be executed in real time. We then show how global surface reconstruction can be provably achieved by (1) appropriately integrating these behaviors to iteratively "grow" the reconstructed regions, and (2) obeying four simple rules.

The support of the National Science Foundation under Grant Nos. IRI-9022608 and IRI-9220782 is gratefully acknowledged.

Contents

1 Introduction

Psychologists have long advocated the importance of simple behaviors in humans and biological organisms in general, whose purpose is the active acquisition of specific information about the environment (e.g., getting closer to an object if it is too far away, or picking up an object to determine its weight) [1–4]. In the context of visual information processing, this view suggests that the observer should be active, operating with a purpose within the environment and interacting with it in real time, in order to perform a specific task. It further suggests that visual information should be *actively obtained* rather than *imposed* [5], e.g., through an *a priori* determined collection of images.

Computer vision research has only recently started to investigate this approach to visual processing. However, it is becoming increasingly evident that a *behavior-based* [6], *animate* [7], and *purposive* [8] approach to visual processing holds great promise for developing abilities comparable to those of biological organisms; numerous recent results have shown that by using and combining simple behaviors that control the camera parameters (e.g., the direction of gaze), real-time and robust solutions to navigation, reconstruction and tracking tasks, for example, can be found [8–16].

To date, a behavior-based approach for visually exploring an unknown, arbitrarily-shaped object has not been reported, although the significance of such an ability is clear [2, 6]. In this paper we present an approach that fills this gap. We ask the following question: What behaviors are important for performing global, geometric tasks such as searching for specific surface markings (e.g., manufacturer's identification), building a global model of an object, or recognizing an object?

A general approach to any of these tasks requires the ability to control the point of observation based on the appearance of the viewed object. Behaviors implementing viewpoint control will correspond to object manipulations [4, 17, 18], where an observer purposefully rotates a hand-held object based on its appearance, or to observer motions, where the observer moves around an object. The development of viewpoint-control behaviors for performing the above tasks is the subject of this paper.

Before attempting to answer the question of what behaviors are important for solving these visual exploration tasks, it is important to first ask how we are going to evaluate the behaviors used. There are two issues here: *Efficiency*, i.e., how efficiently the viewing parameters can be controlled to solve the task, and *correctness*, i.e., whether the task can always be accomplished by an appropriate combination of behaviors.

Like many others (e.g., [8]), we are only interested in behaviors that can be executed in real time and that tightly couple sensing and action. This is because our purpose is to develop behaviors that exploit the ability to quickly interact with the environment to achieve simple and real-time solutions for a given task, rather than rely on the availability of large amounts of computational power [19].

Unfortunately, the correctness of a behavior or a collection of behaviors for performing a given task is much harder to evaluate. One way is to conduct an experimental validation, i.e., implement them and see how they perform under various conditions [20]. Instead, before building and testing an experimental system, we take a theoretical approach to this problem. We prove that the behaviors used will *always* perform the task. In this respect, we believe that a major contribution of our approach is a methodological one; although the behaviors we present are qualitative, fairly intuitive, and implementable in real time, they have provable properties and their success can always be evaluated in terms of the geometry of the surfaces being explored. When trying to perform tasks that depend on the appearance of an object, the use of such provably-correct behaviors is critical: The appearance of objects can drastically change depending on the observer's viewpoint, making *ad hoc* behaviors unpredictable, incomplete, and even non-terminating.

So how can one guarantee that a collection of behaviors will always perform a given task? The answer is to first isolate each behavior and study its effects (e.g., how each behavior controls a viewing parameter) and then, even more importantly, to study their interactions [14, 21–25]. We believe that it is in this long-range interaction and cooperation between the component behaviors that the solution to global geometric tasks lies. We put particular emphasis on this point, studying how each one of the developed behaviors affects each other, and how in unison they solve the task.

Based on the above general principles, in this paper we develop a collection of behaviors for deriving a global, three-dimensional description of an object. We assume the object is stationary and the observer is able to freely move on a sphere around the object. The behaviors we describe solve the following problem: How should the motion of the observer be controlled in order to generate a dense sequence of images that maximizes the area of the reconstructed regions on the object? We call this the *global surface reconstruction task*. We use a shape-from-motion module for extracting surface shape information [26, 27]. We consider this task for smooth surfaces of arbitrary shape; the viewed object is unknown, can be non-convex, and can self-occlude.

Our goal is to identify a collection of general behaviors that are important not only for performing this reconstruction task, but also for performing other more qualitative tasks requiring visual inspection of an object's surface (e.g., searching for surface markings) that do not require surface reconstruction. The connection of the global surface reconstruction task with these qualitative tasks is strong because an important underlying issue in all these tasks is how to control viewpoint so that previously-occluded points on the object's surface become visible.

In order to ensure that the developed behaviors are applicable to tasks in which surface reconstruction is not required, we pay particular attention to the kind of visual information used to control the motion of the observer. Our viewpoint-control behaviors rely on qualitative surface shape information that can be directly extracted from image features (specifically, the occluding contour) and do not rely on the quantitative three-dimensional information provided by the shape-from-motion module. Consequently, we believe that the behaviors developed here are of interest for performing visual exploration tasks that do not presuppose the ability to reconstruct the object's surface.

Very little work has been published on the use of real-time viewpoint-control behaviors for reconstruction, exploration or recognition tasks. However, the few recent approaches taking advantage of viewpoint-control behaviors demonstrate their importance and generality. We showed in an earlier paper [28] that the shape recovery problem for smooth surfaces becomes considerably simplified if the observer uses a simple viewpoint-control behavior to move to a special viewpoint, which for the case of surfaces of revolution corresponds to their side view. The work of Wilkes and Tsotsos [29] illustrates how the ability to purposefully change viewpoint can simplify the task of object recognition in a simple world of Origami objects. Grosso and Ballard [30] are currently designing a head-eye system for implementing viewpoint-control behaviors. Recent work by Blake *et al.* [31] showed that shorter paths can be achieved in robotic navigation tasks if the shape of the obstacles is taken into account during navigation.

Apart from the above approaches, viewpoint control for performing various tasks has been treated as a complex and computationally-intensive optimization problem (i.e., "where to look next"), where the best next viewpoint is searched for within the space of all possible viewpoints [32]. In tasks involving surface reconstruction, the viewpoint-control mechanisms assumed that a three-dimensional representation of the visible surfaces can be recovered independently from each viewpoint [33–35], ruling out the applicability of these mechanisms in more qualitative visual exploration tasks. Another disadvantage of these approaches to surface reconstruction is that they do not consider how the global geometry of the surface (e.g., self-occlusions) affects the correctness of their reconstruction algorithms, relying instead on heuristic optimization criteria for selecting the observer's next viewpoint. This makes the surface regions that are reconstructed unpredictable. Moreover, in the context of animate and purposive vision, their major drawback is their inability to exploit the real-time control of the observer's viewpoint to simplify the procedure for solving the task (e.g., by utilizing shape-from-motion approaches, which have been shown to considerably simplify shape recovery computations [26, 27, 36]), as well as to simplify the viewpoint control process itself.

1.1 Global Reconstruction from the Occluding Contour

The behaviors we develop in this paper control the observer's motion so that surface shape information can be recovered from a dense sequence of images using a shape-from-motion module. We assume that the shape-from-motion module uses the occluding contour for recovering this information. The occluding contour is the projection of the visible rim, the one-dimensional set of visible points at which the line of sight is tangent to the surface (Figure 1). It is well known that the occluding contour is a rich source of shape information [37–46]. Furthermore, contrary to approaches that use sophisticated sensing mechanisms to reconstruct the scene from a single viewpoint or a small number of viewpoints [22, 33, 34], recent results demonstrate that the occluding contour can be reliably detected in edge images [27, 47–49], and that quantitative shape information (e.g., curvature) can be efficiently and accurately recovered from the occluding contour [27, 31].

An important property of both the occluding contour and the visible rim is that they depend on the viewpoint of the observer and the shape of the surface: Under continuous observer motion the visible rim slides over the surface and may change its connectivity, affecting the geometry and topology of the occluding contour (Figures 1(b),(c)), and revealing shape information for the parts of the surface over which the visible rim slides. The deformations of the occluding contour have been studied by Giblin and Weiss [50], Cipolla and Blake [27], Vaillant and Faugeras [47], and Koenderink [51, 52], while the contour's topological changes have been an active research topic for more than a decade (e.g., [52–54]). Much of the analysis we present in this paper builds on this work.

The assumption that the deforming occluding contour is used for recovering surface shape information allows us to perform global surface reconstruction by answering the question: How should the observer's motion be controlled so that the area of the surface regions over which the visible rim slides is maximized?

This formulation of the global surface reconstruction task is particularly important because it allows us to separate the issue of controlling the viewpoint of the observer from the issue of reconstructing the surface, i.e., processing the images produced during the observer's motion. It does not presuppose the ability to reconstruct the surface, and allows us to develop qualitative viewpoint-control behaviors that do not require measuring distances to points on the viewed surface or recovering the surface shape at such points. This allows the behaviors described in this paper to be used for other visual exploration tasks.

To illustrate the qualitative nature of the behaviors we use, consider the following simple example. Suppose an observer is level with a point on a hill top; then this point will belong to the visible rim, which separates the visible points on the hill from the occluded points. By moving up or down, the observer will see just over the hill or not quite to the top, respectively. These motions cause the visible rim to slide over a neighborhood of that point, allowing reconstruction of the surface in its neighborhood. Clearly, this behavior is qualitative in the sense that it does

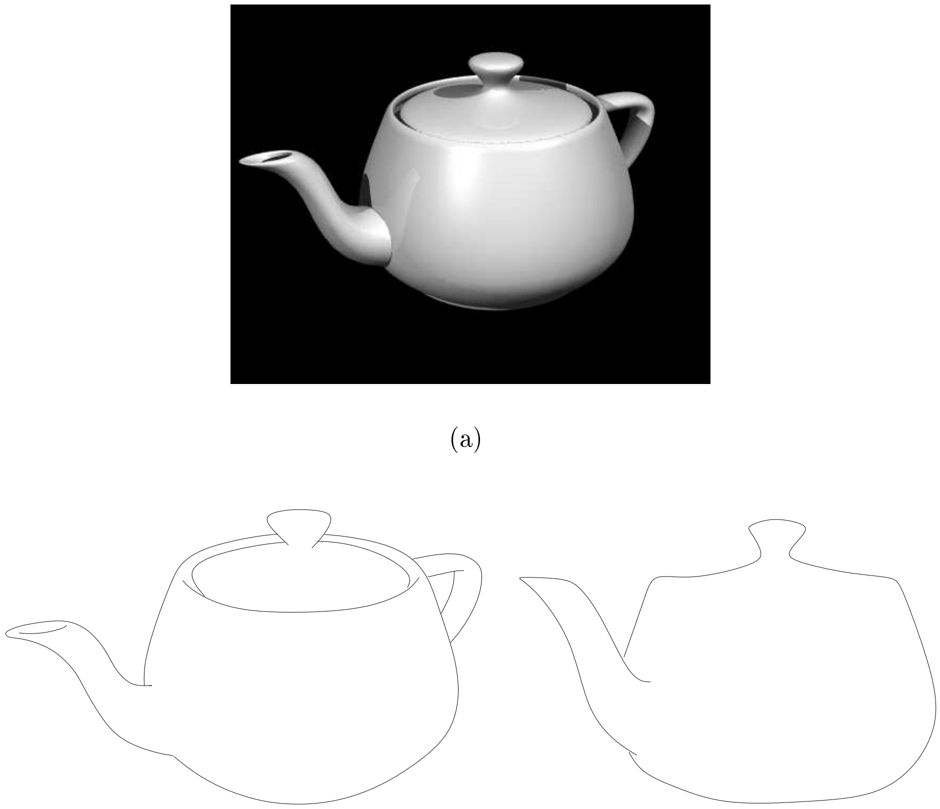
(a) (b) (c)

Figure 1: (a) Because of its geometry, part of the teapot's surface is not reconstructible. (b) The occluding contour corresponding to the view of the teapot in (a). (c) The occluding contour corresponding to a side view of the teapot. Note that the topology (i.e., connectivity) of the occluding contour is different from that in (b).

not require any knowledge of the observer's distance to the hill or of the curvature of the surface near the hill top. However, after processing the images obtained during the observer's motions these quantities can be computed.

Before we begin our analysis of the global surface reconstruction task and the behaviors we use to solve it, it is important to understand the kinds of results we expect. Suppose we want to reconstruct the surface of the teapot shown in Figure 1(a) by moving on a large sphere that encloses it. Clearly, the interior surface of the spout cannot be completely reconstructed because a portion of that surface will never be visible. This visibility constraint puts an upper limit on the set of *reconstructible* points, i.e., it tells us what is the most we can expect from any behavior that controls the observer's motion on a sphere around the teapot. Depending on what surface features are used for recovering surface shape information, geometric constraints other than simple point visibility may render additional regions of the surface unreconstructible, and the reconstructible points may form disconnected sets. For example, when the occluding contour is used (as we do), the shape of concave regions on a surface cannot be reconstructed because the visible rim will never slide over such regions. Given these geometric constraints, the global surface reconstruction task requires controlling the observer's motion, starting from some initial position, so that the visible rim slides over all reconstructible regions on the object's surface that intersect the visible rim at the initial position. We study this specific formulation of the global surface reconstruction task in the rest of the paper.

Despite these constraints, a substantial fraction of the surface of common objects is reconstructible. For example, the surface of a torus and the exterior surface of the teapot in Figure 1(a) are both completely reconstructible. In general, the constraints imposed by surface self-occlusions and the use of the occluding contour define a collection of surface curves that bound the unreconstructible regions on the surface. Examples of such regions are surface concavities, which are bounded by parabolic curves on the surface. These curves belong to a special class of surface curves, the *visual event curves* [53], which we discuss in detail later.

How easy is it to develop viewpoint-control behaviors that can provably perform the global surface reconstruction task? One would hope that the simple behavior used for reconstructing the surface around a hill top could be easily extended to globally reconstruct surfaces of arbitrary shape. Unfortunately, this is not the case. Intuitively, the reason for this is that although the observer has some control over the motion of the visible rim over the surface, this control is not complete; the motion of the visible rim also depends on the shape of the surface itself. In addition, the topology of the visible rim may change, further complicating this global reconstruction process.

To illustrate the difficulties involved in global surface reconstruction, consider the surface in

Figure 2(a), an ovoid with a dent. Suppose the observer wants to reconstruct the surface in the vicinity of the visible rim segment β. This is similar to the previous example of the hill top. Now suppose the observer wants to use the same behavior for reconstructing the whole surface. This means that the surface shape for all points along the dark line drawn on the surface must be recovered. As the observer continues to move upward, however, the visible rim slides on the surface to the right and eventually stops overlapping the dark line (Figure 2(c)). This means that the surface shape for points on the dark line will no longer be reconstructed if the observer continues to move upwards. A different viewpoint-control behavior is now required to continue the reconstruction of the surface along the dark line. A similar difficulty also occurs when the torus in Figures 2(d)-(e) is reconstructed, although in this case the difficulty arises from changes in the topology of the occluding contour.

The difficulties illustrated in these examples show that a number of different behaviors are needed for performing global surface reconstruction, and the observer must, under certain circumstances, switch between these behaviors. It is therefore necessary to ask how many such behaviors are required, how many times these behaviors will be executed, and even whether the global reconstruction process will always terminate. These questions are precisely the reason why provable correctness of the developed behaviors is necessary: Since the answers to these questions are not evident even for the geometrically simple surfaces in the examples above, provable correctness of the developed behaviors is required if one hopes to apply them for reconstructing a real object.

In order to guarantee the correctness of the behaviors we develop in this paper, we require that the reconstruction process is non-terminating only if the shape of the surface does not permit global reconstruction in a finite number of steps. Furthermore, we require that the behaviors are *convergent* and *complete*, i.e., the set of reconstructed points on the surface converges to a well-defined set when the reconstruction process does not terminate, and it is equal to the reconstructible surface regions intersecting the visible rim at the initial viewpoint.

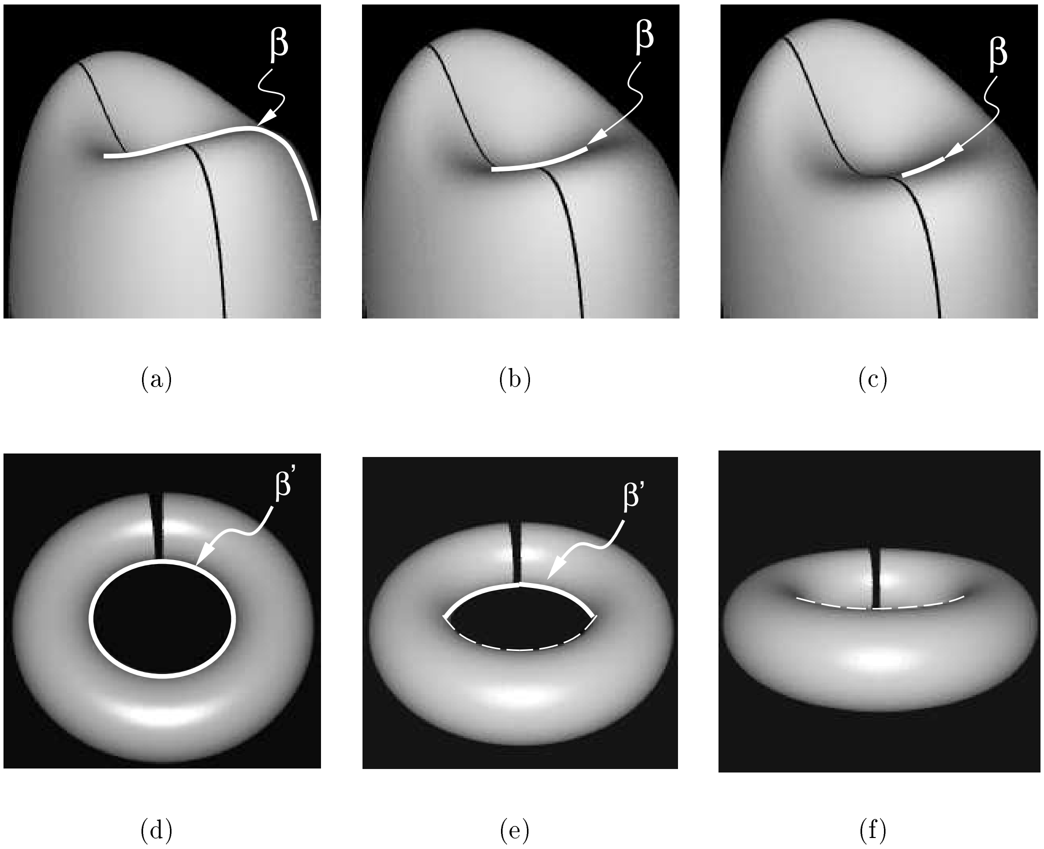

Figure 2: Difficulties involved in globally-reconstructing a dimple-shaped surface and a torus. (a)-(c) The observer is moving in an upward direction in order to reconstruct the surface in the vicinity of the tracked visible rim segment β. (d)-(f) The observer is moving in a downward direction in order to reconstruct the surface in the vicinity of the tracked segment β'. This segment has already disappeared when the observer reaches the viewpoint corresponding to (f); further downward motion cannot be used to reconstruct the surface in the vicinity of the dark line drawn on the surface (the dark line corresponds to a parallel [55] of the torus).

1.2 Paper Overview

We study global surface reconstruction by first developing a collection of behaviors that perform simpler tasks, and then generalizing them to address the complete task (Table 1). Using this paradigm, we develop our behaviors in the context of three increasingly more general reconstruction tasks: The *local surface reconstruction task*, where the observer must control viewpoint to reconstruct a patch around a selected surface point; the *incremental surface reconstruction task*, where the observer must control viewpoint to iteratively "grow" a reconstructed region on the surface; and the *global surface reconstruction task*, where an entire reconstructible region on the surface must be reconstructed.

Section 2 reviews basic terminology and discusses the constraints imposed on the shape-recovery capabilities of the observer when the occluding contour is used for reconstruction and when the observer's motion cannot be controlled. These constraints are captured in four Epipolar Reconstructibility Constraints defined in Section 2.1, and lie at the heart of our analysis. Section 3 considers the local surface reconstruction task. We show how the Epipolar Reconstructibility Constraints give rise to four instances of the local surface reconstruction task, and then we describe four behaviors to perform this task, each of which applies to one specific instance of the task. These behaviors are considered in isolation, by temporarily ignoring the ways in which the behaviors can interact. The behaviors are discussed at a fairly high level, purposefully omitting steps that are not pertinent to the local task of reconstructing a small patch on the surface.

The local analysis of Section 3 provides us with a characterization of the maximal reconstructible regions on the surface. This characterization is considered in Section 4. It tells us what is the most we can expect from any approach that uses the occluding contour for reconstruction, and makes precise the goal we wish to achieve in performing global surface reconstruction. Having characterized the reconstructible regions, Section 5 analyzes how the four behaviors performing the local surface reconstruction task should be integrated. We consider this issue in the context of the incremental surface reconstruction task.

Finally, we consider the global surface reconstruction task in Section 6. We show that in order to use the behaviors developed in the previous sections to perform global surface reconstruction, the observer must obey a number of "rules" when executing these behaviors. We specify these rules with the help of three increasingly more general global reconstruction tasks: The *semi-global curve reconstruction task* and the *global curve reconstruction* task, which require the reconstruction of a curve drawn on the surface, and the *global surface reconstruction task*. The main result of the paper is Theorem 7, presented in Section 6.3, which shows that global

Reconstruction Task	**Section considered**	**Behaviors used**	**Rules obeyed**
Local surface reconstruction task	3	Patch Reconstruction (4 behaviors)	—
Incremental surface reconstruction task	5	Incremental Reconstruction Patch Reconstruction	—
Semi-global curve reconstruction task	6.1	Incremental Reconstruction Patch Reconstruction	Rules 1-3
Global curve reconstruction task	6.2	Incremental Reconstruction Patch Reconstruction	Rules 1-4
Global surface reconstruction task	6.3	Incremental Reconstruction Patch Reconstruction	Rules 1-4 (generalized)

Table 1: The tasks considered in the paper.

surface reconstruction (i.e., all reconstructible regions intersecting the visible rim at the initial viewpoint are reconstructed) is always achieved when the observer uses the developed behaviors and obeys the associated rules.

2 Surface Shape from Occluding Contour

We begin by specifying the class of surfaces considered in this paper. We assume the viewed surface S is smooth, opaque, oriented, and bounds an open volume V in $\Re^3$. Since our ultimate goal is global reconstruction of the surface, we constrain V to be finite and S to be finitely complicated (e.g., S has a finite number of undulations), a notion to be made more precise later (Sections 3.4 and 6.3). Finally, to simplify our exposition we will assume that the viewed surface is *generic*. Informally, generic surfaces exemplify the notion of non-degeneracy: They are surfaces whose topological and geometrical characteristics (e.g., parabolic curves, the Gauss map) are not affected by infinitesimal perturbations of the surface. Generic surfaces have been the subject of active research in singularity theory [56]; their precise definition is rather technical and beyond the scope of this paper. Although in this paper we focus on the global reconstruction of generic surfaces, our results generalize directly to non-generic surfaces.

We assume that S is viewed under spherical projection, and that the observer moves on a *motion sphere* surrounding the surface. In this case, the space of viewpoints is identical to the observer's motion sphere.[1] Visible points on S project to a spherical image that can be modeled as the unit sphere centered at the observer's position (Figure 3). Since S is opaque, only a subset of S will be visible from any given viewpoint: A point $p \in S$ is visible from viewpoint c if and only if the line segment connecting p and c does not intersect V; otherwise p is occluded.

Suppose $\mathbf{x}(s,t)$ is a local parameterization of S in the neighborhood of p. The partial derivatives $\mathbf{x}_s(p), \mathbf{x}_t(p)$ of $\mathbf{x}$ with respect to s and t define $T_p(S)$, the plane tangent to S at p. The *rim* of S at viewpoint c is the set of surface points for which $T_p(S)$ contains the line segment connecting p and c. For these points the projection mapping is singular. The *visible rim* consists of the rim points that are visible. The *occluding contour* of S is the projection of the visible rim on the image (Figure 3).

The occluding contour is a collection of open and closed smooth curves for almost all positions of the observer. The endpoints of open occluding contour curves are either cusps or T-junctions [51]. The shape and topology of this collection of curves depends on S and the observer's position.

[1] In the following we use the terms "viewpoint" and "observer position" interchangeably.

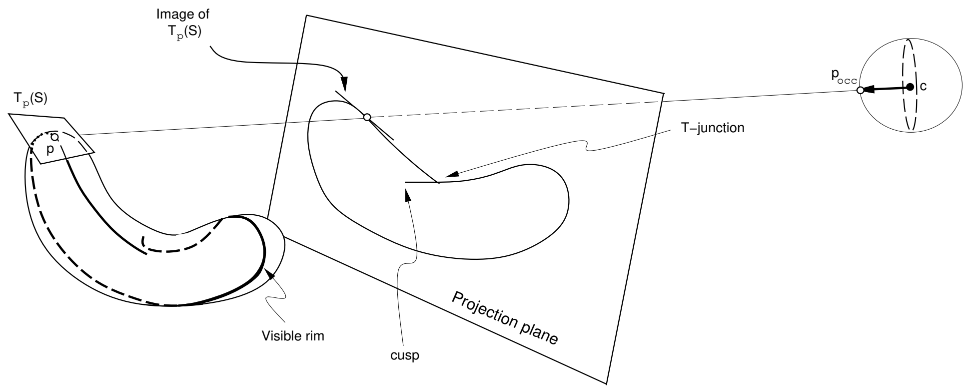

Figure 3: Viewing geometry. The projection, p_{occ}, of a point p on the spherical image can be thought of either as a point on the unit sphere or as a unit vector in the direction of the ray passing through p and the observer's viewpoint, c. The visible rim and the occluding contour corresponding to the projection of a bean-shaped surface are shown (adapted from [50]). The occluding contour consists of a single curve whose endpoints are a T-junction and a cusp. For simplicity, we show the visible rim projected to a planar image perpendicular to p_{occ}.

Giblin [50] and Cipolla [27] showed that under continuous observer motion along a curve $c(t)$, the changes in the geometry of the occluding contour can be used to completely describe the local shape of the surface (i.e., curvature) at points on the visible rim.

Specifically, suppose the observer is moving along a smooth curve $c(t)$ and p is a point lying on the visible rim when the observer is at position $c(t_0)$. The local shape of the surface at p can be described by finding a suitable parameterization $\mathbf{x}$ for the surface in a neighborhood Π of p and calculating the first and second fundamental forms of S with respect to $\mathbf{x}$ [52]. Furthermore, if such a parameterization can be found, the first and second fundamental forms for all points in Π can be computed. This allows the recovery of the local shape of S for a whole *patch* of points containing p. A suitable parameterization that relates the deformation of the occluding contour with the local shape of the surface at p is the *epipolar parameterization* [27, 50, 57], discussed below.

2.1 The Epipolar Parameterization

Intuitively, the epipolar parameterization captures the idea that under continuous motion of the observer (and when the topology of the occluding contour does not change), the set of points comprising the visible rim consists of smooth curves that "slide" over the surface. This allows the non-concave parts of the surface to be considered as a collection of patches, each of which is a family of visible rim curves.

In particular, suppose that the observer changes viewpoint in a continuous fashion by tracing a smooth curve $c(t), t \in [t_l, t_u]$ and that point p belongs to the visible rim when the observer is at some intermediate position $c(t_0)$. The epipolar parameterization is a local parameterization describing a surface patch Π in the neighborhood of p in terms of the visible rim and the instantaneous direction of motion $\mathbf{v}(t) = c'(t)$ of the observer. This patch is defined by a region R of $\Re^2$ with $R = \{(s,t) | t \in [t_l, t_u], s \in (s_l(t), s_u(t))\}$, and a mapping $\mathbf{x} : R \longmapsto \Pi$ such that (Figure 4):

- $\mathbf{x}(s_0, t_0) = p$,
- $\mathbf{x}(s, t_0 + \Delta t)$ is a smooth visible rim curve[2] at viewpoint $c(t_0 + \Delta t)$, and
- $\mathbf{x}(s_0 + \delta s, t)$ is a smooth curve[3] whose tangent at t is along the vector $\mathbf{x}(s_0 + \delta s, t) - c(t)$ and whose normal belongs to the *epipolar plane* defined by the vector product $[\mathbf{x}(s_0 + \delta s, t) - c(t)] \wedge \mathbf{v}(t)$. This curve defines a correspondence between a visible rim point at $c(t)$ and at $c(t + \Delta t)$.

See [27, 50] for details on how the epipolar parameterization can be used to recover the first and second fundamental forms of the surface for all points in Π.

The crucial point in the definition of the epipolar parameterization is that the epipolar parameterization imposes strong constraints on the observer's shape-recovery capabilities. We outline the four constraints that a surface point p must satisfy below.

[2] For Δt fixed and s such that $(s, t_0 + \Delta t) \in R$.
[3] For δs fixed and t such that $(s_0 + \delta s, t) \in R$.

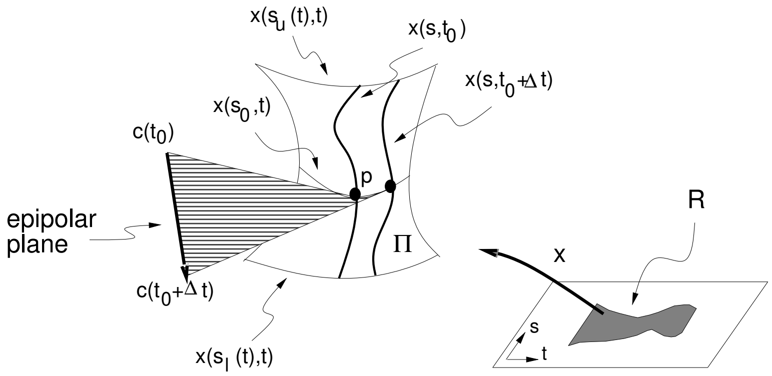

Figure 4: Epipolar parameterization for a surface patch Π around point p. The parameterization $\mathbf{x}$ defines a mapping from R to Π. Curves $\mathbf{x}(s, t_0)$ and $\mathbf{x}(s, t_0 + \Delta t)$ are curves on the visible rim of the surface corresponding to positions $c(t_0)$ and $c(t_0 + \Delta t)$, respectively. The tangent to the curve $\mathbf{x}(s_0, t)$ for $t = t_0$ is along the line through $c(t_0)$ and p. The curve's normal is in the epipolar plane defined by the direction of motion, $\mathbf{v}(t)$, and the line connecting $c(t_0)$ and p.

Epipolar Reconstructibility Constraints

C0: p must be visible from some viewpoints on its tangent plane. Otherwise, p will never belong to the visible rim.

C1: If $c(t)$ is the observer's viewpoint and p is on the visible rim, p must not be the endpoint of a visible rim curve. This is because it was assumed that $\mathbf{x}^{-1}(p) \in R$.

C2: If $c(t)$ and $\mathbf{v}(t)$ are the observer's viewpoint and velocity, respectively, and p is on the visible rim, $T_p(S)$ must not contain $\mathbf{v}(t)$. This is because in that case $\mathbf{x}_t = 0$ (i.e., point p remains on the visible rim under an infinitesimal viewpoint change [28, 57]) and consequently $\mathbf{x}_s \wedge \mathbf{x}_t$ does not define the tangent plane of S at p.

C3: The line segment connecting p and the observer's position must not be a ruling of the surfaces bounding the volume of viewpoints where the topology of the occluding contour and the visible rim (i.e., their connectivity) is constant. This is because the topology of the visible rim curve containing p must not change in the neighborhood of p under an infinitesimal motion of the observer. Only a finite collection of curves on the surface cannot satisfy this constraint. These curves bound the surface points not satisfying constraint C0. They are a subset of the *visual event curves*, which are discussed in Section 2.2.

The Epipolar Reconstructibility Constraints show that the epipolar parameterization cannot be used to describe the surface in the neighborhood of any visible rim point. Consider, for example, the situation shown in Figure 3. If the view of the bean-shaped surface from a position along the observer's path is the one shown in the figure, the epipolar parameterization cannot be used to define a surface patch that surrounds the endpoints of the visible rim curve; the points projecting to the cusp and the T-junction can only be contained in the *boundary* of such a patch. Hence, the local shape of the surface in a neighborhood of those points cannot be recovered.

The above constraints also show that the surface patch Π depends on how the visible rim curve $\mathbf{x}(s, t_0 + \Delta t)$ slides over the surface when Δt varies continuously. Consequently, the dynamics of the visible rim curves determine the patches reconstructed. These dynamics depend on the local and global shape of the surface as well as the motion of the observer. Therefore, if the observer's motion cannot be controlled, the parts of the surface that are reconstructed from the deformation of the occluding contour cannot be controlled.

The Epipolar Reconstructibility Constraints allow us to characterize the surface regions that are reconstructible on the surface, i.e., they tell us what is the most we can expect from any viewpoint-control behavior that uses the epipolar parameterization for surface reconstruction. In particular, we have the following characterization of the reconstructible points on the surface:

Reconstructible surface regions: The reconstructible regions on a surface are the maximal connected sets of surface points for which all four Epipolar Reconstructibility Constraints can be simultaneously satisfied.

We will see that the reconstructible surface regions are bounded by surface points that can satisfy constraint C0 but not constraint C3. Constraint C3 applies only to surface points that belong to visual event curves. In the next section we briefly describe these curves and their relationship to the topological changes of the occluding contour and the visible rim.

2.2 Visual Events and their Associated Visual Event Curves

The topology of the occluding contour of a smooth surface is stable for almost all viewpoints of the observer. That is, for almost all viewpoints, the contour's topology does not change when these viewpoints are infinitesimally perturbed (e.g., see [52, 58]). Results from singularity theory show that the space of viewpoints can be partitioned into a collection of maximal connected cells within which the topology of the occluding contour remains constant. Visual events occur when the observer's viewpoint belongs to the boundaries of these maximal cells. An infinitesimal perturbation of such a viewpoint results in changes in the topology of the occluding contour. A catalog of the visual events was given by Kergosien [56] for the case where a surface is smooth, generic and transparent (i.e., the occluding contour is considered to be the projection of all rim points, not just the points on the visible rim), and is observed under orthographic projection.

In the spherical projection model and when the observer is able to freely move in $\Re^3$, the cells for which the occluding contour's topology is constant occupy volumes in three-dimensional space. When the observer is constrained to move on a sphere, the viewpoints of constant contour topology correspond to the intersections of the motion sphere with these three-dimensional cells. The boundaries of these cells are manifolds of codimension one, two and three, i.e., they are surfaces, curves and points, respectively. Visual events of codimension one occur when the observer's viewpoint belongs to a surface bounding a cell of constant contour topology. Higher codimension visual events occur at the intersections of surfaces associated with events of codimension one. In the following we only use results from the analysis of codimension one events for transparent generic surfaces; see [54] for a detailed discussion of visual events of higher codimension and the geometry of viewpoint space partitioning in the neighborhood of such viewpoints both for transparent and opaque surfaces.

The cell boundaries corresponding to codimension one events are *ruled* surfaces, i.e., they can be represented as a 1-parameter family of lines called *rulings*. Each visual event is associated with such a ruled surface. The surface rulings are lines that contact the surface at multiple points, or have a high order contact with the surface at one or more points.[4] These lines touch the surface

[4] A line is said to have n-th order contact with a surface at a point p when all directional derivatives at p along the line up to (but not including) order n are zero.

at points on certain characteristic curves associated with the visual event [51, 53]. When a visual event occurs, the observer's viewpoint is contained in a line l of the ruling and there is a point on the rim for which the line l connecting it to the observer's viewpoint has a high order contact with the viewed surface or touches the surface at multiple points.

Visual events are classified into *local* and *multilocal* events, depending on whether the line l touches the surface at one or more points. *Local events* occur when l has high order contact with the surface at exactly one point. In this case the observer's viewpoint is contained in a ruled surface that touches the surface along a single curve characterizing the event. *Multilocal events* occur when the line l contacts the surface at at least two points. The observer's viewpoint in this case belongs to a ruled surface that touches the surface along two or more curves. Figure 5 shows the topological changes corresponding to all possible local and multilocal events of codimension one.

The local events are the *swallowtail*, *beak-to-beak* and *lip* events (Figure 5(a)). These events occur when there is a visible rim point p for which the line l connecting it to the observer's viewpoint has fourth order contact with the surface. The swallowtail event occurs when p is on a flecnodal curve [51], while the lip and the beak-to-beak events occur when p is on a parabolic curve.

The multilocal events of codimension one are the *triple-point*, *tangent-crossing* and *cusp-crossing* events [53] (Figure 5(b)). Triple-point events occur at viewpoints where there is a triplet of collinear rim points whose supporting line l passes through the observer's viewpoint. These three points project to a single point on the occluding contour. The line l has second order contact with each of the points. The ruled surface associated with a triple-point event contains rulings that touch the surface at three distinct points; it is formed by sweeping l while maintaining three-point contact with the surface. These points trace three curves on the surface, $\tau_1(t), \tau_2(t), \tau_{mid}(t)$ $(0 \leq t \leq T)$, such that for all $t \leq T$ the line through $\tau_1(t), \tau_2(t), \tau_{mid}(t)$ is tangent to the surface at the three points and $\tau_{mid}(t)$ corresponds to the middle point of contact.

A tangent-crossing event occurs at viewpoints where there is a pair of rim points with a common tangent plane whose supporting line l passes through the observer's viewpoint. The ruled surface associated with this event touches the surface along two curves, $\gamma_1(t), \gamma_2(t)$ $(0 \leq t \leq T)$,

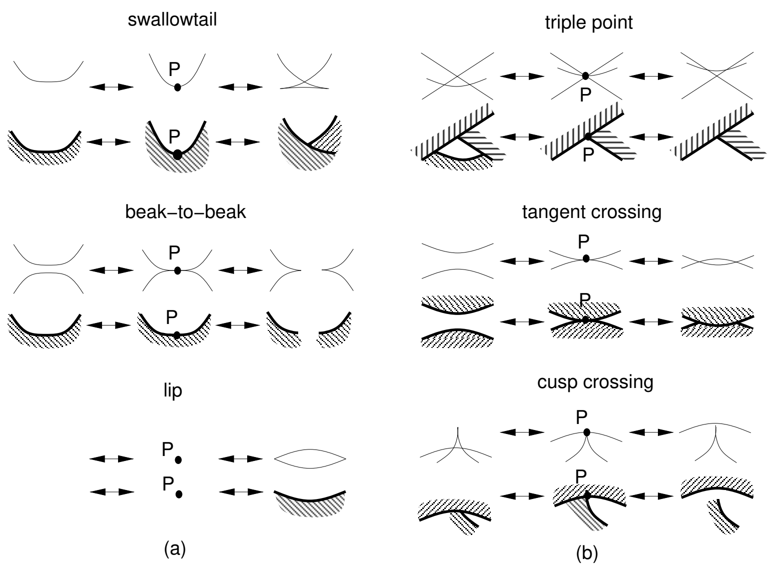

Figure 5: Visual events (adapted from [59]). Intermediate views of the occluding contour correspond to viewpoints on the boundaries of regions of constant contour topology. For these viewpoints, the line l through the observer's viewpoint and the rim point(s) projecting to P has a high order contact with the surface, or touches the surface at multiple points. The visible rim point projecting to P farthest away from the observer's viewpoint is a point for which the Epipolar Reconstructibility Constraint C3 is not satisfied. The events represent changes in the topology of the occluding contour of a transparent surface. Also shown are examples of how these events appear when the viewed surface is opaque. (a) Local events. In a swallowtail event the occluding contour develops a singularity and then it breaks off into three segments forming two cusps and a T-junction. In a beak-to-beak event two occluding contour curves (of which only one is the projection of a visible rim curve) meet at a point and then split off, generating two cusped contours. In a lip event a cusped contour appears out of nowhere. (b) Multilocal events. In a triple-point event, points on three occluding contour segments project to a single point. In a tangent-crossing event two contours meet creating a pair of T-junctions. Finally, in a cusp-crossing event three occluding contour segments connected by two T-junctions split off with one of the segments ending with a cusp.

such that for all $t \leq T$, the line through $\gamma_1(t)$ and $\gamma_2(t)$ is tangent to the surface at those points. Finally, a cusp-crossing event occurs at viewpoints where there is a rim point projecting to a cusp of the occluding contour and another rim point whose supporting line l passes through the observer's viewpoint. In this case, l has third order contact with the rim point projecting to the cusp and second order contact with the other one. Again, the ruled surface associated with a cusp-crossing event is created by sweeping l while maintaining one second-order and one third-order contact with the surface. It touches the surface along two curves, $\sigma(t)$ and $\sigma_{cusp}(t)$ ($0 \leq t \leq T$), such that for all $t \leq T$ the line through $\sigma(t)$ and $\sigma_{cusp}(t)$ is tangent to the surface at $\sigma(t)$ and has third-order contact at $\sigma_{cusp}(t)$.

From the description of the visual events given above, it follows that each visual event is associated with a collection of curves on the surface, called the *visual event curves*. In particular, we can define the following visual event curves on the surface:

- the parabolic curves of the surface
- the flecnodal curves of the surface
- the curves $\tau_1(t), \tau_2(t), \tau_{mid}(t)$, associated with triple-point events
- the curves $\gamma_1(t), \gamma_2(t)$, associated with tangent-crossing events
- the curves $\sigma(t), \sigma_{cusp}(t)$, associated with cusp-crossing events

The analysis of the visual events associated with generic surfaces generalizes to non-generic surfaces in the following sense. Visual events for non-generic surfaces also occur when the line connecting the observer's viewpoint and a point on the surface has high order of contact with the surface, or when that line touches the surface at multiple points. Consequently, the same catalog of visual events is still valid. However, degeneracies may also occur: The visual event "curves" defined above may in fact become two-dimensional regions on the surface (e.g., all points on a cylinder are parabolic). See [59] for a discussion of this issue in the case of algebraic surfaces (which are not generic), and the computation of visual event curves for these surfaces.

A subset of the visual event curves defined above bounds the maximal region on the surface that can be reconstructed by *any* viewpoint-control behavior. We characterize this boundary precisely in Section 4.

3 Behavior-based Reconstruction of Local Surface Patches

In this section we consider the local surface reconstruction task: Suppose the observer is at position c, and let p be a visible rim point on the object's surface that is identified by its projection, p_{occ}, on the occluding contour. The task of the observer is to continuously control viewpoint, starting from point c, in order to recover the local shape of the surface for all points in *some* neighborhood Π of p using the epipolar parameterization. To perform this task we pay closer attention to the four Epipolar Reconstructibility Constraints that the point p must satisfy so that the epipolar parameterization can be used to describe a neighborhood of that point. In particular, we use the following three observations:

- If p is not the endpoint of a visible rim curve and the visible rim's topology does not change in a neighborhood Π of p if the observer's viewpoint is infinitesimally perturbed, the observer's motion can be controlled so that the surface in Π can be described by the epipolar parameterization. This observation allows us to control the observer's motion so that Epipolar Reconstructibility Constraint C2 is satisfied for p.

- If p is the endpoint of a visible rim curve, the epipolar parameterization cannot describe the surface in the neighborhood of p. However, there are other viewpoints on p's tangent plane at which p is not the endpoint of a visible rim curve. This observation allows us to control the observer's motion so that Epipolar Reconstructibility Constraint C1 is satisfied for p.

- The point p and the observer's viewpoint may be such that the occluding contour's topology changes in the neighborhood of p under an infinitesimal viewpoint perturbation. For all points p except those lying on a subset of the visual event curves, the observer can move to other viewpoints on p's tangent plane at which the contour's topology does not change in the neighborhood of p if these viewpoints are infinitesimally perturbed. Points for which Epipolar Reconstructibility Constraint C3 cannot be satisfied cannot have their neighborhood reconstructed by *any* viewpoint-control behavior.

Based on these observations, for any given viewpoint we distinguish four types of visible rim points that indicate the position of those points within the visible rim curve containing them and the Epipolar Reconstructibility Constraints that need to be satisfied: *Ordinary points*, which are not endpoints of a visible rim curve and that satisfy constraint C3; *cusp points*, which are visible rim endpoints projecting to a cusp on the occluding contour and that satisfy constraint C3; *T-junction points*, which are visible rim endpoints projecting to a T-junction on the occluding contour and that satisfy constraint C3; and *degenerate points*, which are visible rim points not satisfying constraint C3. These four classes of visible rim points are exhaustive and define four instances of the local surface reconstruction task.

To perform the local surface reconstruction task we use a basic behavior that controls the motion of the observer in order to deal with the case where p is an ordinary visible rim point. This behavior is a generalization of the example behavior discussed in the Introduction for reconstructing a surface patch around a hill top. The other three cases are treated by (1) controlling the position of the observer in order to reach a viewpoint where p is an ordinary visible rim point, and (2) using the basic behavior in order to recover the shape of the surface in a neighborhood of that point. Table 2 summarizes these behaviors.

Visible rim point type	**Constraints that need to be satisfied**	**Behavior used to satisfy constraints**	**Section presented**
Ordinary	C2	Ordinary Patch Reconstruction Behavior	3.1
Cusp	C1,C2	Cusp Patch Reconstruction Behavior	3.2
T-junction	C1,C2	T-junction Patch Reconstruction Behavior	3.3
Degenerate	C1,C2,C3	Degenerate Patch Reconstruction Behavior	3.4

Table 2: Behaviors used for performing the local surface reconstruction task.

3.1 Behavior-based Reconstruction Around Ordinary Points

In this section we assume that the observer is initially positioned at point c and has selected an ordinary visible rim point p. The task of the observer is to move in a way that allows the surface shape for all points in a neighborhood of p to be recovered.

Since p lies on the visible rim, any neighborhood of p will contain three types of points, namely points that are occluded, points that lie on the visible rim, and points that are visible but do not lie on the visible rim. Consequently, in order to recover the shape of the surface for all points in a neighborhood of p, the observer must control its motion so that neighboring points that are either occluded or visible from position c become part of the visible rim. Intuitively, this means that the task of recovering the shape of the surface for all points in a neighborhood of p using the epipolar parameterization is equivalent to the task of controlling the motion of the observer so that the smooth visible rim curve containing p "slides" over all points in that neighborhood.

The task of forcing the rim to slide over all points in a neighborhood of p can be cast as the task of inducing the visibility or the occlusion of all points in this neighborhood. Under continuous motion the observer traces a path that is a continuous curve. From any viewpoint along this path, the points belonging to the visible rim are in a transitional visibility state. Some visible rim points will become occluded under an infinitesimal change of viewpoint along the path, and some will remain visible but leave the rim. Hence, the task of forcing all points in a neighborhood of p to lie on the visible rim for some position along the observer's path corresponds to the tasks of (1) inducing the visibility of all points in a neighborhood of p that are occluded when the surface is viewed from the observer's initial position, and (2) inducing the occlusion of all points in a neighborhood of p that are visible from the initial position.

Suppose the observer changes viewpoint by tracing a smooth curve $c(t)$ with $c(0) = c$, and let $\mathbf{v}(t) = c'(t)$ be the instantaneous direction of motion. Given a segment $\beta(t)$ of the visible rim at viewpoint $c(t)$, the epipolar parameterization allows us to define the segment $\beta(t + \delta t)$ of the visible rim at $c(t + \delta t)$ that corresponds to $\beta(t)$. Theorem 1 shows that we can get a qualitative characterization of the motion of the visible rim over the surface by looking at the surface normal:

Theorem 1 (Visibility transition dynamics) *Suppose* $\mathbf{n}(q)$ *is the outward surface normal at* q *and that the observer's position is* $c(t)$*. If* $\beta_1^+(t), \ldots, \beta_n^+(t)$ *and* $\beta_1^-(t), \ldots, \beta_m^-(t)$ *are the (open) smooth segments of the visible rim that contain ordinary points* q *satisfying* $\mathbf{n}(q) \cdot \mathbf{v}(t) > 0$ *and* $\mathbf{n}(q) \cdot \mathbf{v}(t) < 0$*, respectively, then*

1. *all points* $\beta_i^+(t + \delta t)$*,* $i = 1, \ldots, n$*, are occluded from position* $c(t)$*.*

2. *all points* $\beta_i^-(t+\delta t)$, $i = 1, \ldots, m$, *are visible from position* $c(t)$.

3. *All ordinary visible rim points satisfying* $\mathbf{n}(q) \cdot \mathbf{v}(t) = 0$ *will be contained in the visible rim at* $(t + \delta t)$.

The proof of Theorem 1 can be found in Appendix A.

When the observer continuously changes viewpoint along a smooth curve $c(t)$ ($t_{start} \leq t \leq t_{end}$), the visible rim segments will slide over the surface. If $\beta_i(t)$ is the segment containing the selected point p, $\beta_i(t)$ will trace a patch Π on the surface around p that can be described using the epipolar parameterization. The boundary of this patch consists of the segments $\beta_i(t_{start}), \beta_i(t_{end})$ contained in the visible rim at viewpoints $c(t_{start}), c(t_{end})$, respectively, and the traces of the endpoints of $\beta_i(t)$. The endpoints of $\beta_i(t)$ will either be points satisfying $\mathbf{n}(p) \cdot \mathbf{v}(t) = 0$, or will be the endpoints of a visible rim curve. The following simple, qualitative viewpoint-control behavior can now be used to solve the local reconstruction task around p (Figure 6):

Ordinary Patch Reconstruction Behavior

Step 1: Select a point p on the visible rim that is not the endpoint of a visible rim curve. This selection is done indirectly by selecting p's projection, p_{occ}, on the occluding contour. Point p_{occ} must not be the endpoint of an occluding contour curve.

Step 2: Compute the surface normal at p. The normal is given by $\mathbf{T} \wedge p_{occ}$, where $\mathbf{T}$ is the tangent to the occluding contour at p_{occ} [51].

Step 3: (Reconstructing the occluded points near p.) Select a direction $\mathbf{v}_1$ for moving on the motion sphere that satisfies the inequality $\mathbf{n}(p) \cdot \mathbf{v}_1 > 0$. Change viewpoints along $\mathbf{v}_1$ while continuously monitoring the deformation of the occluding contour curve that initially contains p_{occ}.

Step 4: (Reconstructing the visible points near p.) Move back to the initial viewpoint and reapply Step 3 by selecting another direction of motion $\mathbf{v}_2$ that satisfies the inequality $\mathbf{n}(p) \cdot \mathbf{v}_2 < 0$.

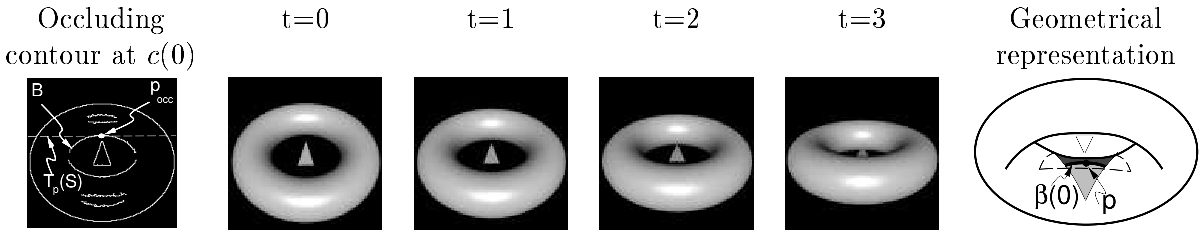

Figure 6: Reconstructing a patch around an ordinary visible rim point on a torus. The leftmost image shows the edges detected in the initial view, $t = 0$. The small triangle in the middle of the torus points toward the direction of the line connecting the initial viewpoint of the observer and the center of the torus. The point selected is point p, shown in the rightmost figure, in which the torus is viewed from below. The point is selected by selecting its projection, p_{occ}, on the occluding contour from the initial viewpoint, $c(0)$. The figure shows the views of the surface as the observer executes Step 3 of the Ordinary Patch Reconstruction Behavior. The tangent to the occluding contour at p_{occ} is horizontal and, hence, the projection of the surface normal at p in the image is a vertical line. The observer changes viewpoint by moving vertically downward. B is the projection of the visible rim segment $\beta(0)$ that contains p. The endpoints of $\beta(t)$ are T-junction points. Since $\beta(t)$ disappears during the observer's motion, after the execution of Step 3 the patch reconstructed on the surface is bounded by the curves traced by the endpoints of $\beta(t)$ and by $\beta(0)$ (i.e., a triangle-like patch). The patch is shown as the lightly-shaded area on the rightmost figure. Step 4 completes the reconstruction process around p by reconstructing a patch on the other side of $\beta(0)$ (shown as the darkly-shaded area on the rightmost figure).

It should be noted that two components of the above behavior are purposely left unspecified. In particular, the choice of directions $\mathbf{v}_1$ and $\mathbf{v}_2$ in Steps 3 and 4 is required to satisfy a particular inequality, but no exact value is given. Furthermore, no precise stopping condition is provided for terminating the observer's motion in these two steps. This leaves considerable freedom for the observer to make a choice that satisfies additional requirements [60, 61] (e.g., distance to the surface). In Section 6, where we consider the global surface reconstruction task, we show that when executing the Ordinary Patch Reconstruction Behavior the observer must obey a number of rules that "ground" these steps.

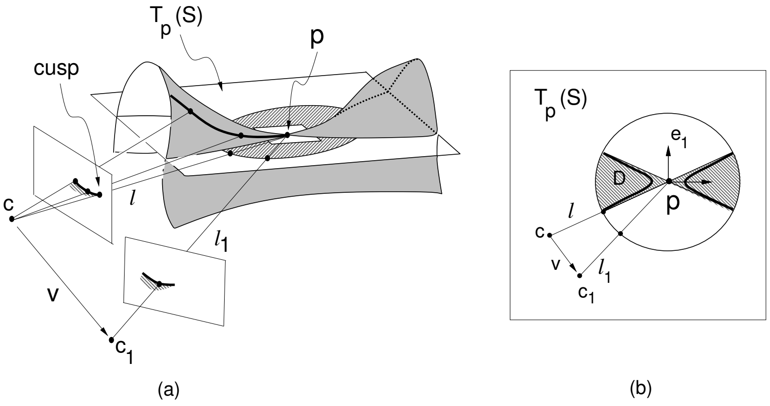

Figure 7: (a) Moving on the tangent plane of a cusp point p. The bold curve on the surface is the visible rim curve containing p. c is the initial position of the observer. $\mathbf{v}$ is a motion direction on $T_p(S)$ satisfying the conditions of Theorem 2: p is an ordinary visible rim point when viewed from viewpoint c_1. (b) A top view of the tangent plane at p is shown. e_1 is the first principal direction of the surface at p. D is the Dupin indicatrix at p. The line segment l connecting c and p is along an asymptote of D. p is occluded when l lies in the shaded areas.

3.2 Behavior-based Reconstruction Around Cusp Points

In this section we consider the local reconstruction task around cusp visible rim points by studying the geometrical changes on the occluding contour as the observer moves on the selected point's tangent plane. The main idea is similar to the one used in [28]: If the observer moves along specific directions on the tangent plane of the selected point p, p will remain on the visible rim but will cease to be the endpoint of the visible rim curve containing it. Hence, after the observer performs such a change in viewpoint, the local reconstruction task around p is transformed to the reconstruction task around an ordinary visible rim point.

In particular, let p_{occ} be the projection of p and let $C(s)$ be the occluding contour curve containing p_{occ}, with $p_{occ} = C(s_0)$. The following result characterizes the special directions for changing the observer's viewpoint that force p to become an ordinary point (Figure 7(a)):

Theorem 2 *Let p_{occ} be the projection of a cusp point p, and let $\mathbf{T}$ be the tangent to the occluding contour at p_{occ} defined as the limit $\lim_{s \to s_0} C'(s)$. If the observer performs an infinitesimal viewpoint change along a direction $\mathbf{v}$ in $T_p(S)$ such that $\mathbf{v} \cdot \mathbf{T} > 0$, p will be an ordinary visible* *rim point at the new viewpoint.*

The proof of Theorem 2 can be found in Appendix A.

Theorem 2 immediately suggests a qualitative behavior for changing viewpoint in order to recover the shape in a neighborhood of p. We only need to answer two additional questions to completely specify such a behavior: how to move and when to stop. The only computation necessary for selecting the direction of motion is determining the tangent to the occluding contour at p_{occ}. This decision is particularly simple since the observer is constrained to move on a sphere. In this case, the motion decision corresponds to deciding whether to move clockwise or counterclockwise on the unit circle centered at p. Although p will become an ordinary point after an infinitesimal viewpoint change, such a change can leave p arbitrarily close to the visible rim's endpoint. We temporarily leave the stopping condition of this behavior unspecified; as in the case of the Ordinary Patch Reconstruction Behavior, this step will be grounded in Section 6 where the observer must obey a number of rules when executing the Cusp Reconstruction Behavior. This analysis results in the following simple viewpoint-control behavior (Figure 8):

Cusp Patch Reconstruction Behavior

Step 1: Compute the tangent T to the occluding contour at the selected occluding contour cusp point p_{occ}.[5]

Step 2: Compute the normal $\mathbf{n}(p)$ of the surface at the corresponding visible rim point p.

Step 3: Determine whether a clockwise or counterclockwise direction of motion $\mathbf{v}$ on $T_p(S)$ satisfies the inequality $\mathbf{v} \cdot \mathbf{T} > 0$.

Step 4: Perform a small viewpoint change on $T_p(S)$ in the direction selected at Step 3 while fixating at p.

Step 5: Use the Ordinary Patch Reconstruction Behavior to reconstruct a surface patch around p.

[5]It should be noted that an accurate calculation of T is not always necessary. Inaccuracies in this calculation will result in choosing a motion plane that is not the tangent plane at p. However, since the behavior allows a *neighborhood* of the selected point to be reconstructed, small errors in the observer's motion will not have an adverse effect if the patch reconstructed during Step 5 of the behavior is large.

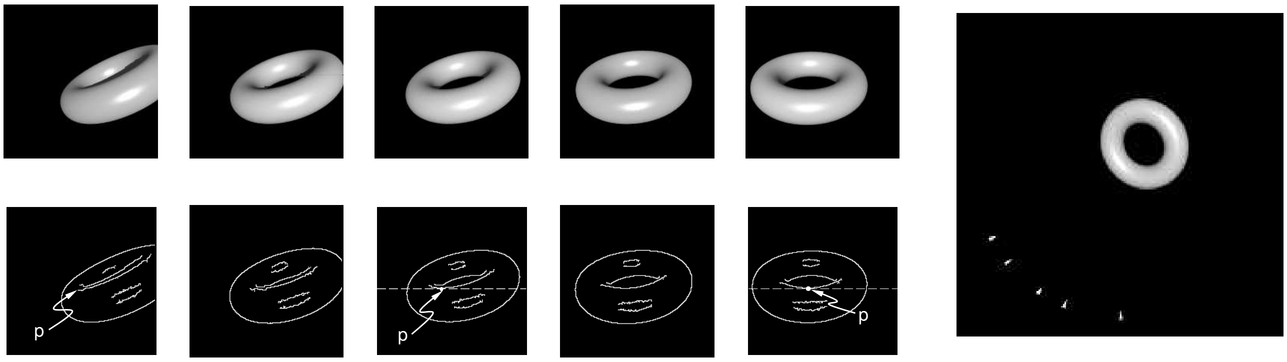

Figure 8: Forcing a cusp point on the visible rim of a torus to become a ordinary point. The upper-left view corresponds to the initial position of the observer. The lower-left view was derived by applying an edge detector to the image above it. The tangent to the selected cusp point p is horizontal. The observer moves on the tangent plane of the selected point, which in this case is a horizontal plane perpendicular to the plane of the page. The sequence of images shows views of the surface from consecutive positions along the observer's path. The rightmost image shows a top view of the tangent plane at p (i.e., "looking down" on the torus), and the five viewpoints where the observer is moving, from left to right. Point p remains on the rim throughout the motion and can be tracked by tracking the occluding contour point whose tangent is horizontal (middle view). The sequence clearly shows that p becomes an ordinary visible rim point after the viewpoint adjustments.

3.3 Behavior-based Reconstruction Around T-junction Points

The behavior for reconstructing the surface in the neighborhood of a T-junction point is similar to the one used for cusp points. The observer changes viewpoint on the tangent plane of the selected point p in order to force p to become an ordinary visible rim point.

In particular, let C_1 be the projection of the visible rim curve ending at p, and let p_{occ} be the projection of p (Figure 9). Furthermore, let β be the visible rim curve whose projection, C_2, intersects C_1 at p_{occ}. Since a T-junction is formed at p_{occ}, there is a point q on β that also projects to p_{occ}. The viewpoint-control behavior is based on the following theorem, which relates the motion of the observer to the motion of the T-junction in the image during the observer's motion:

Theorem 3 *Let p be a T-junction visible rim point and let q be the visible rim point whose projection coincides with the projection of p. Let $\mathbf{n}(q)$ be the outward surface normal at q. If the observer performs an infinitesimal viewpoint change along a direction $\mathbf{v}$ in $T_p(S)$ such that $\mathbf{v} \cdot \mathbf{n}(q) > 0$, p will be an ordinary visible rim point at the new viewpoint.*

The proof of this theorem is omitted for brevity. The idea behind the proof is to apply Theorem 1 to control the motion on $T_p(S)$: By moving on $T_p(S)$ the observer ensures that p will remain on the visible rim; by choosing the direction of motion to satisfy $\mathbf{v} \cdot \mathbf{n}(q) > 0$, the observer ensures that if C_1 lies to the left of p_{occ} in the neighborhood of p_{occ}, the projection of the T-junction moves to the *right* of p_{occ} (Figure 9). Following a similar analysis as the one presented at the end of Section 3.2, a simple viewpoint-control behavior for recovering the shape at points in a neighborhood of a T-junction point can be formulated:

T-junction Patch Reconstruction Behavior

Step 1: Let C_1, C_2 be the occluding contour segments creating the T-junction corresponding to the selected point p, where C_1 is the projection of the visible rim curve ending at p. Compute the tangent T to C_1 at the projection, p_{occ}, of p.

Step 2: Let q be the second visible rim point projecting to p_{occ}. Compute the surface normal, $\mathbf{n}(q)$, at q.

Step 3: Determine whether a clockwise or counterclockwise motion $\mathbf{v}$ on $T_p(S)$ satisfies the inequality $\mathbf{v} \cdot \mathbf{n}(q) > 0$.

Step 4: Perform a small viewpoint change on $T_p(S)$ in the direction determined in Step 3 while fixating at p.

Step 5: Apply the Ordinary Patch Reconstruction Behavior to reconstruct a surface patch around p.

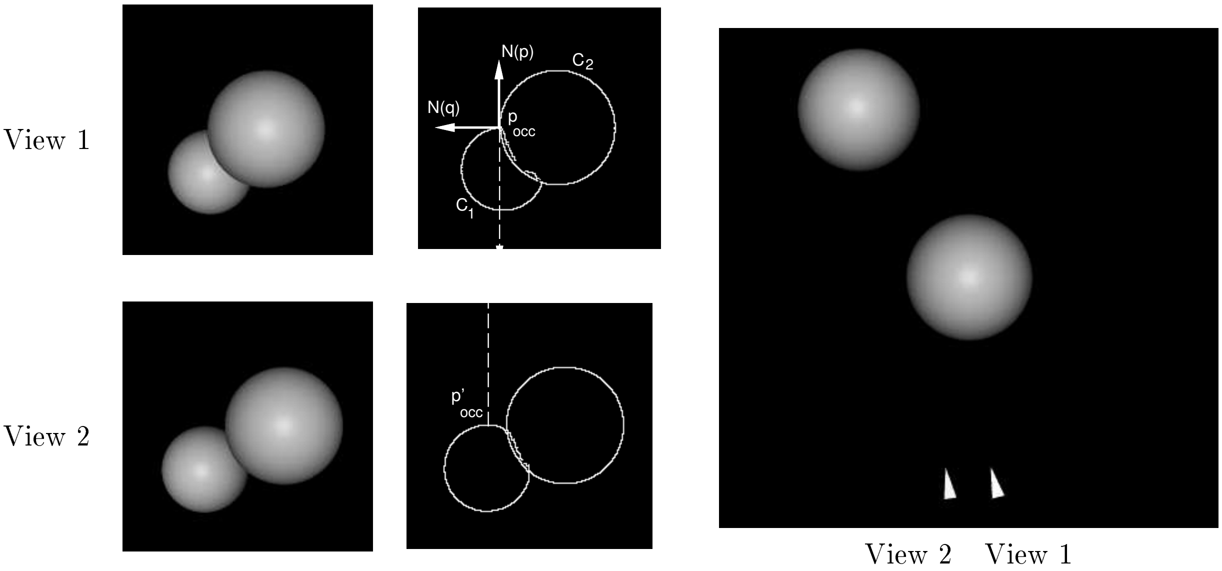

Figure 9: Forcing a T-junction point to become ordinary. In this example a sphere partially occludes another sphere. View 1 corresponds to the initial viewpoint. View 2 corresponds to the view of the spheres after changing viewpoint on the tangent plane of the selected point. Here, this plane is horizontal and perpendicular to the plane of the page. The rightmost view corresponds to a top view of this plane, i.e., "looking down" at the two spheres. $\mathbf{N}(p), \mathbf{N}(q)$ are the projections of the surface normals $\mathbf{n}(p), \mathbf{n}(q)$ at points p and q projecting to p_{occ}, respectively. Theorem 3 requires the viewpoint to move *left*, causing the T-junction to move right. p'_{occ} is the projection of p on the occluding contour corresponding to View 2. The tangent to the occluding contour at p'_{occ} is horizontal.

3.4 Behavior-based Reconstruction Around Degenerate Points

The viewpoint-control behavior for reconstructing a surface patch in the neighborhood of degenerate visible rim points is similar to the behaviors used for reconstructing patches around cusp and T-junction points. The observer again changes viewpoint on the tangent plane of the selected point p in order to force p to become ordinary.

In particular, suppose that the observer is moving along a curve $c(t)$, and that the topology of the visible rim changes in the vicinity of the visible rim point p at viewpoint $c = c(t_0)$. As discussed in Section 2.2, this can happen only if the line connecting c and p has a high order

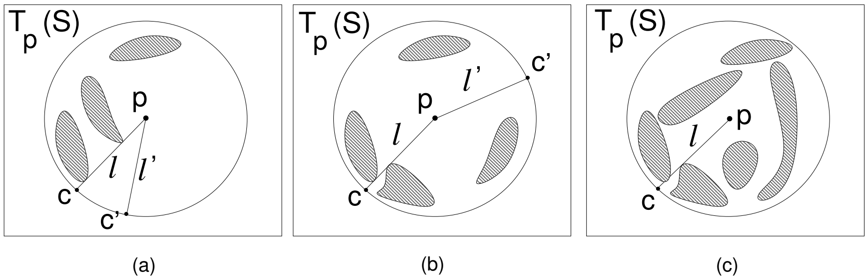

Figure 10: Reconstructing a patch around a degenerate point p. A top view of the tangent plane of p is shown. Shaded regions correspond to the intersections of $T_p(S)$ with the object. In this example, p belongs to a visual event curve associated with a triple-point event: the line through p and the observer's viewpoint, c, touches the surface at three points. (a) A small viewpoint change on $T_p(S)$ makes p ordinary. (b) The geometry of the intersection $T_p(S) \cap S$ forces p to become occluded when small viewpoint changes are performed. However, there are viewpoints on $T_p(S)$ at which p is ordinary. (c) The geometry of the intersection $T_p(S) \cap S$ forces p to be occluded at all viewpoints except c.

contact with the surface or if it contacts the surface at multiple points. Furthermore, if the observer makes an infinitesimal viewpoint change to a new viewpoint, c', on $T_p(S)$, the line connecting c' and p will either have lower order contact with the surface or will touch the surface at fewer points. Hence, if p is visible from the new viewpoint, p will become an ordinary visible rim point (Figure 10(a)). Unfortunately, p may no longer be visible (Figure 10(b)). In this case, in order to make p ordinary the observer must move to distant viewpoints on $T_p(S)$ from which p is visible. We therefore need to specify how the observer should move and when to stop. As in the reconstruction of cusp and T-junction patches, the first question can be answered by moving either clockwise or counterclockwise on a circle in $T_p(S)$ around p. The direction of motion on this circle is not important.

The observer should stop when p becomes ordinary. It is easy to see that if there is an open arc of viewpoints on the observer's motion circle from which p is visible, any viewpoint on that arc guarantees that p is an ordinary visible rim point at that viewpoint. To completely specify the observer's motion it remains to give a way of detecting when p becomes visible again. A quantitative approach to this problem is to first determine the three-dimensional coordinates

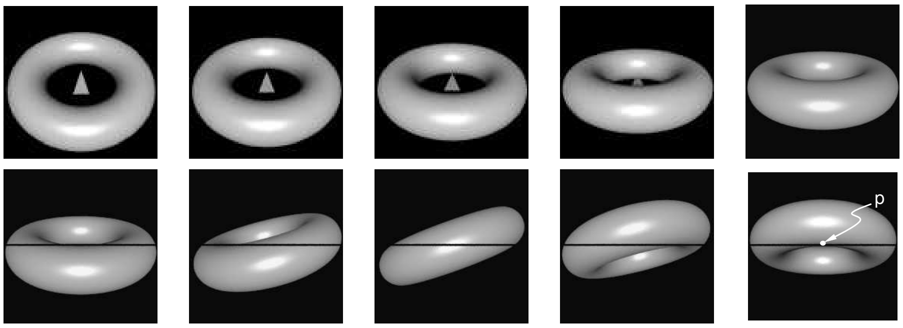

Figure 11: Reconstructing a patch around a degenerate point on the torus. *Top row:* The observer moves downward until the visible rim segment pointed by the triangle shrinks to a point p and disappears. The tangent plane at p is horizontal and perpendicular to the plane of the page. The visual event corresponding to the disappearance of that segment is a tangent-crossing event. Due to this event, the surface in the neighborhood of p cannot be reconstructed by performing a small viewpoint change. *Bottom row:* Moving on $T_p(S)$ in order to make p ordinary. The black horizontal line is the projection of $T_p(S)$ in the image. The observer performs a clockwise viewpoint change on $T_p(S)$ until an ordinary visible rim point with tangent plane identical to $T_p(S)$ is detected. After performing a 180° rotation, such a visible rim point is found; in this case the point is p. The observer can now use the Ordinary Patch Reconstruction Behavior to reconstruct a surface patch around p.

of p, and then continuously check, during the observer's motion, if any visible rim point with tangent plane coincident to $T_p(S)$ matches those coordinates. Instead, in order to keep the observer's behavior qualitative, we simply assume that *all* visible rim points with tangent planes coincident to $T_p(S)$ are identical to p. The above considerations lead to the following qualitative viewpoint-control behavior for reconstructing the surface around p (Figure 11):

Degenerate Patch Reconstruction Behavior

Step 1: Let $p(t_0 - \delta t)$ be the visible rim point at position $c(t_0 - \delta t)$ that is matched to p by the epipolar parameterization. Compute the tangent plane at p as the limit $\lim_{\delta t \to 0} T_{p(t_0 - \delta t)}(S)$.

Step 2: Perform a small counterclockwise motion on $T_p(S)$. If p remains visible, stop. Otherwise, return to the initial viewpoint, c.

Step 3: Perform a small clockwise motion on $T_p(S)$. If p remains visible, stop. Otherwise, return to the initial viewpoint, c.

Step 4: Move clockwise on a circle around p on $T_p(S)$ while continuously monitoring the occluding contour, until either c is reached or there is an ordinary visible rim point q whose tangent plane coincides with $T_p(S)$.

Step 5: If the initial viewpoint is reached, stop. Otherwise, apply the Ordinary Patch Reconstruction Behavior to reconstruct a patch around q, and continue with Step 4.

The Degenerate Patch Reconstruction Behavior can be thought as a generalization of the Cusp and T-junction Patch Reconstruction Behaviors. Instead of moving to a *single* viewpoint on p's tangent plane and then executing the Ordinary Patch Reconstruction Behavior just once, as in the case of those two behaviors, the observer moves to multiple viewpoints on p's tangent plane and executes the Ordinary Patch Reconstruction Behaviors at each of them.

The repetitive nature of the Degenerate Patch Reconstruction Behavior brings us to the first "finiteness" restriction we impose on the surface S: The connected sets of viewpoints on p's tangent plane from which p is visible must be finite. Given this restriction, the Degenerate Patch Reconstruction Behavior will not permit reconstruction of a patch around p if and only if p is occluded from all but a finite set of viewpoints on its tangent plane. Such points never become ordinary during the observer's motion on $T_p(S)$ and, consequently, the surface around these points cannot be reconstructed (Figure 10(c)). This is not a limitation of the Degenerate Patch Reconstruction Behavior; Epipolar Reconstructibility Constraint C3 cannot be satisfied, and there are simply no observer motions that force the visible rim to slide over a neighborhood of such a point to allow reconstruction.

4 The Reconstructible Surface Regions

The analysis of the local surface reconstruction task gives us a way to characterize the reconstructible regions on the surface by characterizing their boundaries. In particular, the Patch Reconstruction Behaviors allow us to reconstruct a surface patch around all surface points except for (1) points that are never visible from viewpoints on their tangent plane, and (2) points on visual event curves that are visible only from a finite number of viewpoints on their tangent plane. For these points, there is no path the observer can follow that forces the visible rim to slide over their neighborhood. This leads directly to the following characterization of the reconstructible regions on the surface:

> **Reconstructible surface regions:** The reconstructible regions on the surface are the maximal connected sets of surface points that are visible from a one-dimensional set of viewpoints on their tangent plane.

Unless the surface is entirely reconstructible, each connected reconstructible surface region $\mathcal{R}$ forms an open set on the surface. Its boundary contains the segments of visual event curves whose points are visible only from a finite number of viewpoints on their tangent plane.[6] By checking which visual event curves can contain such segments, it is easy to show the following:

Theorem 4 (Reconstructible region boundaries) *A point p is on the boundary of a reconstructible surface region if and only if it belongs to either*

- *a parabolic curve bounding a surface concavity,*
- *a curve $\tau_1(t)$ or $\tau_2(t)$ associated with a triple-point event,*
- *a curve $\gamma_1(t)$ or $\gamma_2(t)$ associated with a tangent-crossing event, or*
- *a curve $\sigma_{cusp}(t)$ associated with a cusp-crossing event,*

and is visible from only a finite number of viewpoints on its tangent plane.

[6] To keep things simple, we slightly abuse terminology here: These segments comprise the set of *limit points* of $\mathcal{R}$; in addition to $\mathcal{R}$'s boundary, this set may contain one-dimensional "spines" that protrude from $\mathcal{R}$'s boundary.

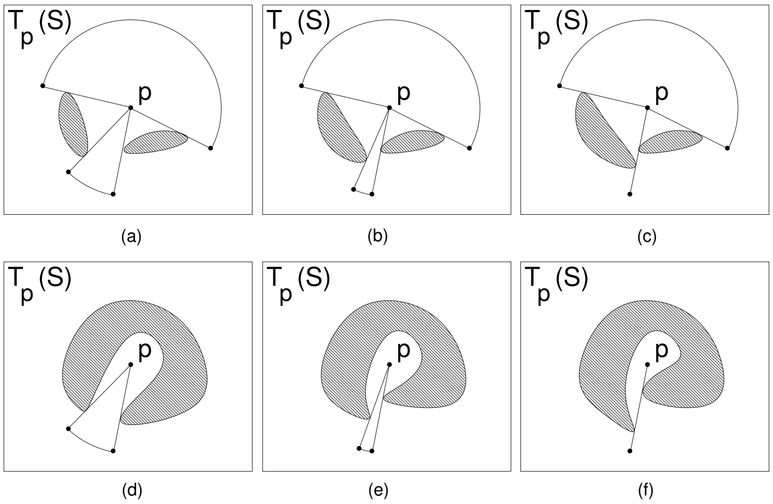

Figure 12: The visibility arcs of a point p. A top view of the tangent plane of p is shown. Shaded areas correspond to the intersections of $T_p(S)$ with the object. (a)-(c) Approaching a visual event curve τ_1 associated with a triple-point event. Point p has two visibility arcs. As p approaches τ_1, one of the visibility arcs of p degenerates to a point. In this case, the point p in (c) belongs to τ_1, but is not contained in the boundary of a reconstructible surface region; the neighborhood around p can be reconstructed by moving to a viewpoint in the remaining visibility arc of p. (d)-(f) Approaching a visual event curve τ_1 associated with a triple-point event. Point p now has one visibility arc. As p approaches τ_1, the visibility arc of p degenerates to a point. In this case, p asymptotically approaches the boundary of a reconstructible surface region.

An intuitive description of Theorem 4 can be given as follows. To each point p in a reconstructible surface region we can associate a collection of *visibility arcs*. These arcs are simply the connected one-dimensional sets of viewpoints on p's tangent plane from which p is visible (Figure 12). When p asymptotically approaches one of the above visual event curves, the length of at least one of its visibility arcs decreases, diminishing to zero (Figure 12(a)-(c)); this can only happen for the visual event curves listed above. Now, if the length of *all* visibility arcs of p diminishes to zero, the visual event curve approached by p belongs to the boundary of the reconstructible regions on the surface (Figure 12(d)-(f)).

The visual event curves listed in Theorem 4 are therefore *potential boundaries* of a reconstructible surface region. They bound such regions only if they contain points with no visibility arcs. Figure 13 shows the reconstructible regions for two objects.

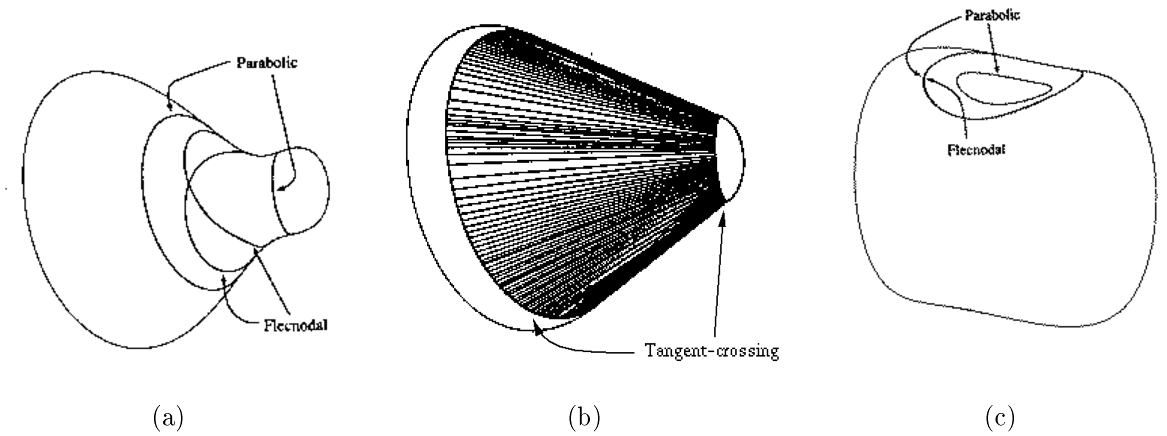

Figure 13: Applying Theorem 4 to some surfaces studied by Petitjean *et al.* [53] and Koenderink [52]. The figures are taken from [53]. The squash-shaped surface in (a) and (b) is completely reconstructible: No cusp-crossing or triple-point visual events can occur, and the surfaces do not have concavities. Furthermore, for each point on the two visual event curves corresponding to a tangent-crossing event (shown in (b)), we can associate at least one visibility arc. The dimple-shaped surface in (c) has one reconstructible region. This region is bounded by the parabolic curve bounding the concavity on the surface.

5 Integrating Behaviors for Incremental Reconstruction

The goal of the global surface reconstruction task is to reconstruct the reconstructible surface regions that intersect the visible rim at the initial viewpoint. To achieve this, we integrate the four Patch Reconstruction Behaviors. We consider the integration of these behaviors in this section in the context of the incremental surface reconstruction task: The observer must integrate these behaviors so that the resulting motion of the observer allows an incremental expansion of the set of reconstructed points on the surface. We consider the composite behavior that achieves this task below.

5.1 The Incremental Reconstruction Behavior

An incremental reconstruction behavior must control viewpoint so that new patches on the surface are successively reconstructed. In order to achieve this we need to answer two questions: (1) How can the observer force points on the boundary of the already-reconstructed patches to lie on the visible rim, and (2) how can the observer control viewpoint so that new patches around those points can be recovered?

The first question can be answered by considering the fact that the boundaries of the already-reconstructed patches were points on the visible rim from previous viewpoints. Hence, it suffices for the observer to move back to the viewpoint where a given boundary point belongs to the visible rim. This can be achieved by saving, along with each occluding contour image, the viewpoint corresponding to that image during the application of the Ordinary Patch Reconstruction Behavior. Since there is a correspondence between the points on the reconstructed patch boundaries and the images they project to, this information is sufficient to guide the observer to the position where a particular boundary point was on the visible rim. Furthermore, since any patch boundary point in a reconstructible surface region can be forced to become a visible rim point, and the behaviors developed in Section 3 can be used to reconstruct a surface patch around any visible rim point for which this is possible, the second question is easily answered by the behaviors already presented. These considerations lead to the following composite behavior (Figure 14):

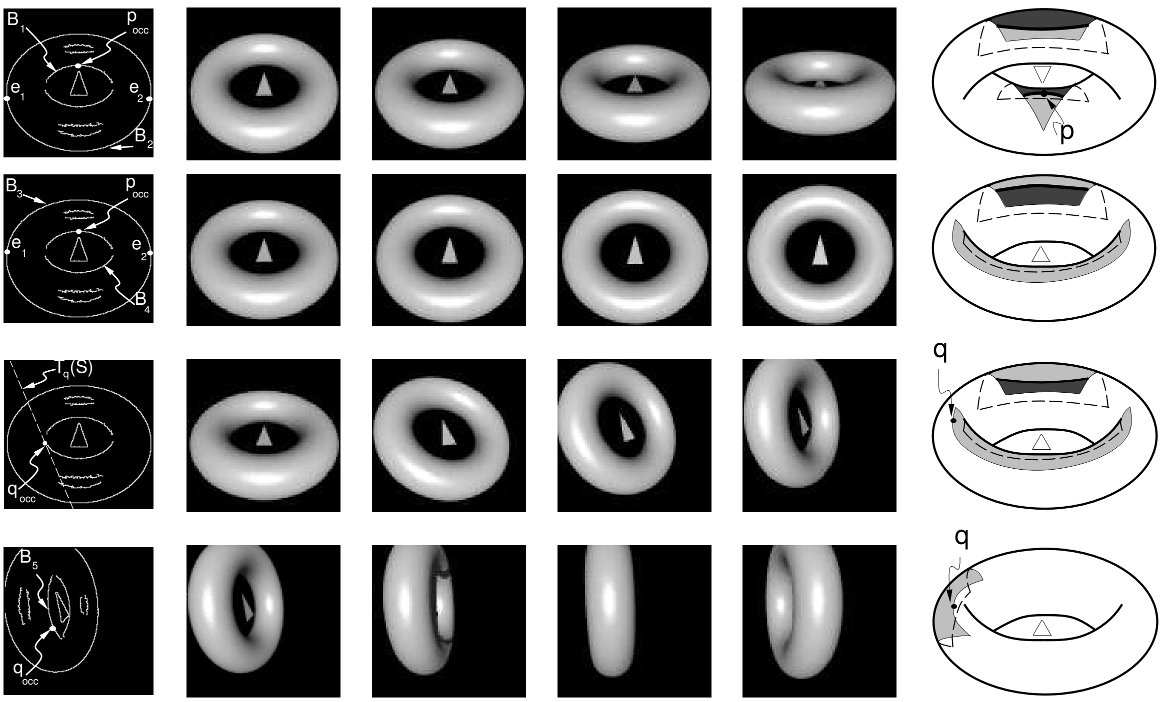

Figure 14: Incremental reconstruction of the surface of a torus. The point selected initially is point p, which is an ordinary visible rim point. The small triangle in the middle of the torus points toward the direction of the line connecting the initial viewpoint of the observer and the center of the torus. *Rows 1 & 2:* The first row of images shows consecutive views of the torus as the observer executes Step 3 of the Ordinary Patch Reconstruction Behavior for p. The second row of images shows the views of the torus as Step 4 of the behavior is executed. There are four visible rim segments sliding on the torus as the observer moves. Two of the segments, whose projections are B_1 and B_2, slide over previously-occluded parts of the surface during Step 3. The other two segments, whose projections are B_3, B_4, slide over previously-visible parts of the surface. Points e_1, e_2 are the endpoints of both B_2 and B_3. The figures to the right of these views show the patches traced by these segments. Lightly-shaded patches correspond to the patches reconstructed during Step 3. Darkly-shaded patches correspond to the patches reconstructed during Step 4. *Row 3:* The observer selects a point q on the boundary of the reconstructed strip. The point is a cusp point corresponding to the viewpoint, c', of the surface in the first row, second image from the left. The observer executes Steps 1-4 of the Cusp Patch Reconstruction Behavior to force q to become ordinary. *Row 4:* The observer executes the Ordinary Patch Reconstruction Behavior to reconstruct a patch for the visible portion of the neighborhood of q, as seen from c'. The figure on the right shows the patch traced by the visible rim segment projecting to B_5.

Incremental Reconstruction Behavior

Step 1: If there exists a portion of the surface that has not been reconstructed, select a point p on its boundary and let c be the viewpoint at which p projected to the occluding contour.

Step 2: Move to c.

Step 3: Apply one of the four behaviors for reconstructing a patch around p according to Table 2, and continue with Step 1.

The specific algorithm for selecting the points p on the boundary of the already-reconstructed surface region is not important for guaranteeing its successive expansion. However, in order to perform the global surface reconstruction task the observer must obey an additional rule when doing this selection; we discuss this rule in Section 6.

6 Global Surface Reconstruction

What kinds of behaviors can perform global surface reconstruction? In the Introduction we motivated the need for behaviors that have three properties:

- *Finite termination:* The reconstruction process does not terminate (i.e., at least one of the component behaviors is executed an infinite number of times) if and only if there is no finite-length path the observer can follow that allows reconstruction of all reconstructible regions on the surface that intersect the visible rim at the initial viewpoint.

- *Convergence:* If the reconstruction process does not terminate, the set of points reconstructed on the surface must converge to a well-defined limit as the number of times the basic behaviors are executed goes to infinity.

- *Completeness:* If the reconstruction process terminates, the reconstructed points must be the union of the reconstructible surface regions intersecting the visible rim at the initial viewpoint. On the other hand, if the reconstruction process does not terminate, the reconstructed points must asymptotically approach that set.

In this section we show that global surface reconstruction can be achieved by (1) using the Incremental Reconstruction Behavior and the four Patch Reconstruction Behaviors, while (2) obeying a number of simple rules that constrain how these behaviors are executed. Motivated by our characterization of the reconstructible regions and the difficulties involved in the reconstruction of surfaces such as those shown in Figures 2 and 15, we develop these rules by considering the following three increasingly more general global reconstruction tasks:

- *Semi-global curve reconstruction task:* Suppose a curve is drawn on the surface so that it intersects the visible rim at the initial viewpoint (Figure 15(a)). The task of the observer is to reconstruct the segments of this curve that are connected, reconstructible, intersect the visible rim at the initial viewpoint, and terminate on one of the visual event curves listed in Theorem 4 (i.e., visual event curves that *potentially* bound a reconstructible surface region, as in Figure 12(c)).

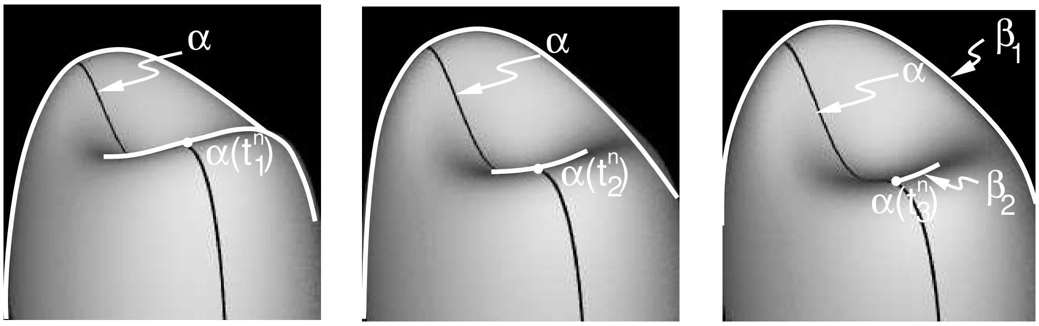

(a) (b) (c)

Figure 15: Difficulties involved in globally reconstructing a dimple-shaped surface and a torus. (a)-(c) In the n-th iteration of the Incremental Reconstruction Behavior the observer is moving in an upward direction in order to reconstruct points in the neighborhood of the point $\alpha(t_1^n)$ on α, which lies on the visible rim. The visible rim eventually slides to the right, making the observer's upward motion ineffective for reconstructing the surface in the vicinity of $\alpha(t_3^n)$.

- *Global curve reconstruction task:* Suppose a curve is drawn on the surface so that it intersects the visible rim at the initial viewpoint. The task of the observer is to reconstruct the segments of this curve that are connected, reconstructible, intersect the visible rim at the initial viewpoint, and terminate on the boundary of a reconstructible surface region (i.e., on precisely those visual event curves listed in Theorem 4 that bound a reconstructible surface region, as in Figure 12(f)).

- *Global surface reconstruction task:* Global surface reconstruction can be considered a generalization of the global curve reconstruction task in the following sense. It is equivalent to the task of reconstructing for every surface curve that intersects the visible rim at the initial viewpoint, a connected, reconstructible segment that terminates on a visual event curve bounding a reconstructible surface region.

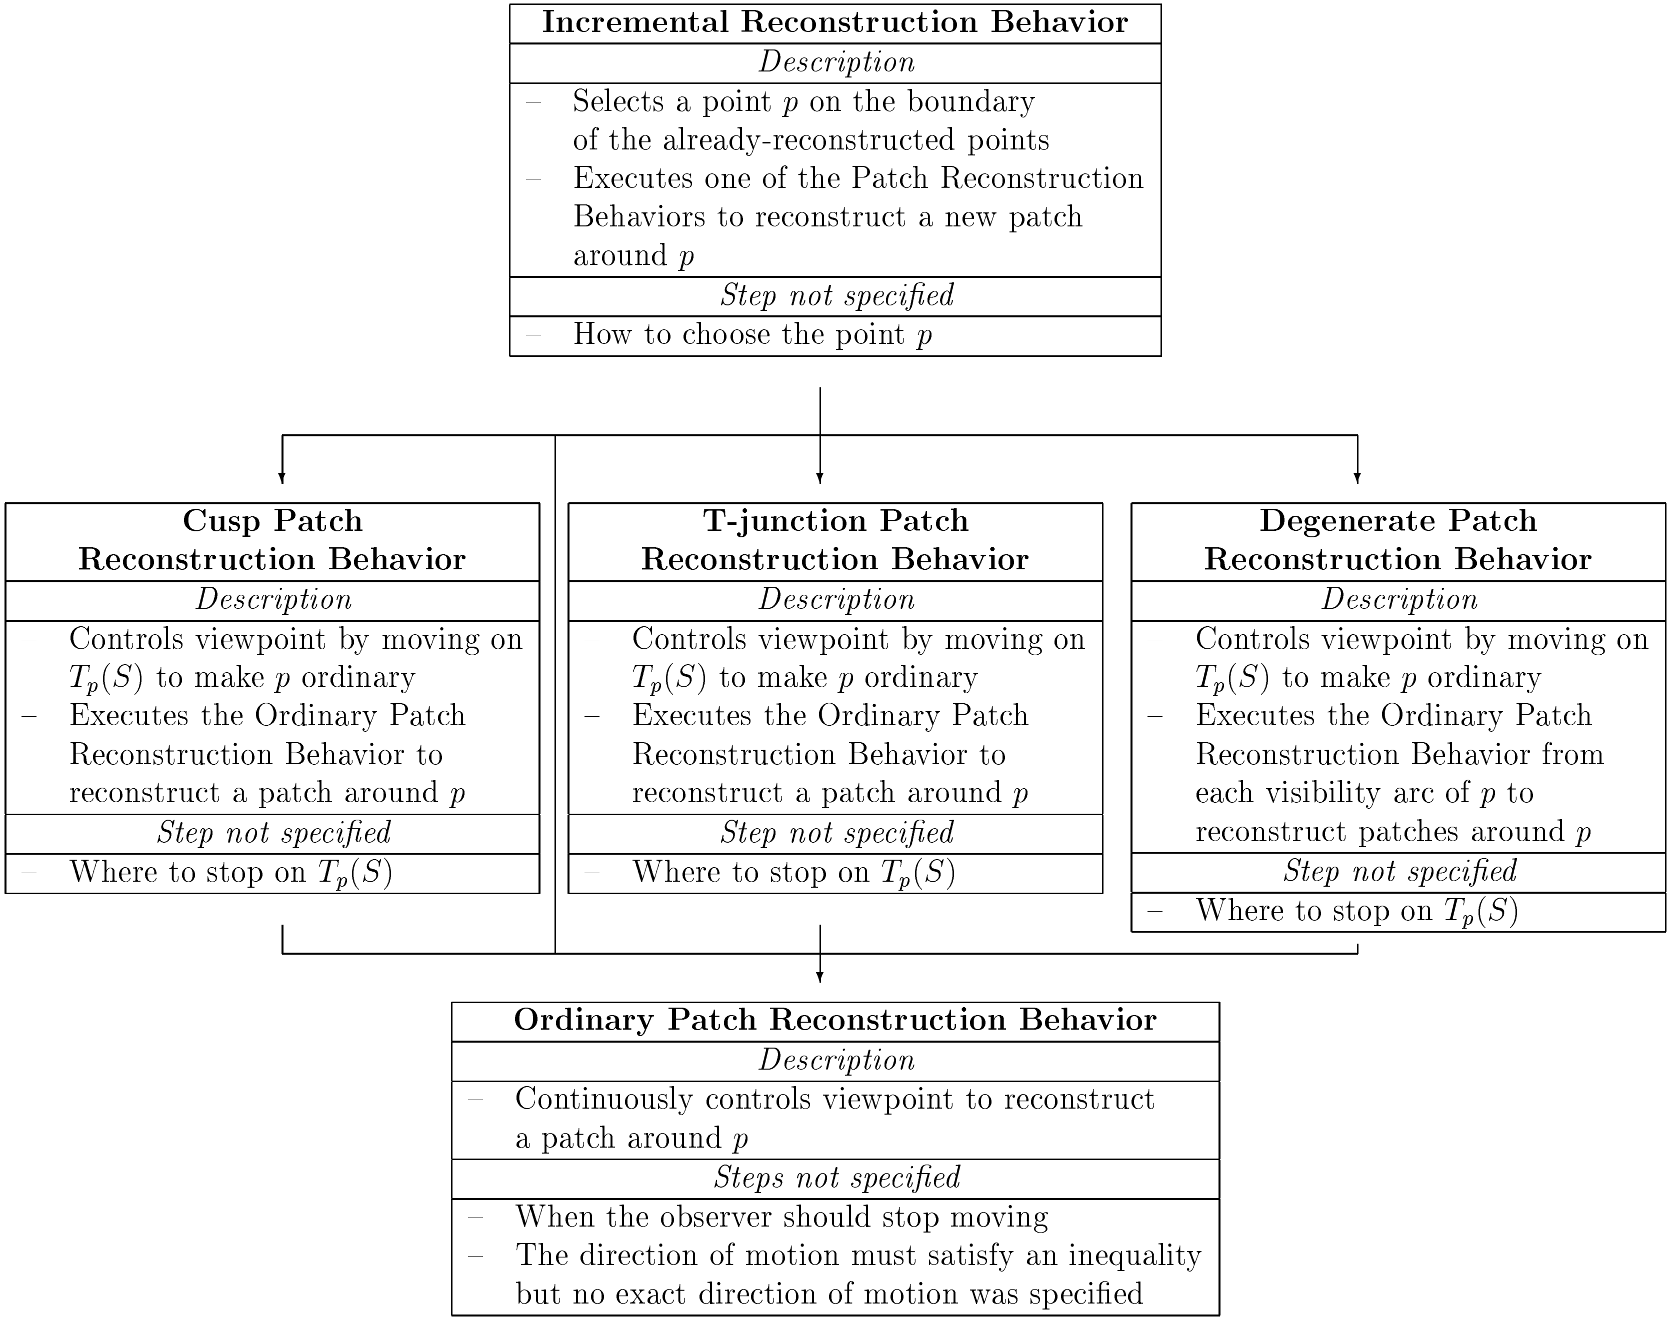

Figure 16: Summary of the viewpoint-control behaviors used to perform the incremental reconstruction task. Arrows indicate the dependencies between the behaviors. Also shown are the steps whose description was omitted in Sections 3 and 5 because they did not affect the observer's ability to incrementally expand the already-reconstructed surface region.

By obeying the rules we develop in this section, the observer can "ground" the steps in the Incremental Reconstruction Behavior and the Patch Reconstruction Behaviors that we left unspecified in their earlier presentation (Figure 16). In the following we keep our analysis at a fairly intuitive level, working through specific examples to motivate the rules used. The interested reader can find formal proofs of correctness in Appendix B.

6.1 Semi-Global Curve Reconstruction

The question we ask in this section is: How can we use the Incremental Reconstruction Behavior to provably perform the semi-global curve reconstruction task? Recall that during the execution of the Incremental Reconstruction Behavior the observer selects a point p on the boundary of the already-reconstructed regions and applies one of the Patch Reconstruction Behaviors to reconstruct a new patch around p (Figure 16). To achieve semi-global curve reconstruction, the length of the curve segment reconstructed at each iteration must diminish if and only if it asymptotically approaches a visual event curve *potentially* bounding a reconstructible surface region. The following theorem gives the three rules the observer must obey in order to perform the semi-global curve reconstruction task. See Appendix B for the proof.

Theorem 5 (Semi-global curve reconstruction rules) *Let α be a finite-length curve drawn on the surface. If α intersects the visual event curves associated with tangent-crossing events at most a finite number of times, and the following three rules are obeyed by the observer, the observer will provably perform the semi-global curve reconstruction task:*

Rule 1: *When choosing the point p on which to apply the Patch Reconstruction Behaviors, always select a point of intersection of α with the visible rim.*

Rule 2: *Always execute the Ordinary Patch Reconstruction Behavior after first moving to a viewpoint c corresponding to the middle of a visibility arc of p.*

Rule 3: *When executing the Ordinary Patch Reconstruction Behavior to reconstruct a patch around p starting from an initial viewpoint c, move around the surface on the normal plane* *at p and stop only if c is reached again, or if the endpoint of the segment of α being reconstructed coincides with a cusp, T-junction or degenerate visible rim point.*

Rules 1 and 3 are obvious. For example, consider the semi-global curve reconstruction task for the surface in Figure 15(a). The n-th iteration of the Incremental Reconstruction Behavior requires selecting a point on the visible rim in order to reconstruct a patch in its neighborhood. Rule 1 simply states that the point selected should be $\alpha(t_1^n)$; the utility of this rule is obvious. Now suppose that the observer starts moving in an upward direction according to the Ordinary Patch Reconstruction Behavior in order to reconstruct a patch around $\alpha(t_1^n)$. Rule 3 states that the observer should move upward until the cusp endpoint of the visible rim coincides with the dark curve at $\alpha(t_3^n)$, as shown in Figure 15(c). Again, there is no reason for the observer to stop before reaching that viewpoint, and there certainly is no reason for continuing its upward motion since reconstruction of a larger piece of the dark curve will not be possible.

The utility of the "finiteness" condition in Theorem 5 and of Rule 2 is not as obvious, although they are crucial for achieving the semi-global curve reconstruction task. Rule 2 constrains the long-range behavior of the observer. The finiteness condition simply states that in order to reconstruct the curve drawn on the surface in a finite number of steps, both the curve and the surface cannot be infinitely complicated (e.g., the curve cannot have an infinite number of inflections and the surface cannot have an infinite number of undulations). In Appendix B we show that this condition ensures the visibility arcs of $\alpha(t_1^n)$ are well-defined in the limit.

To see why Rule 2 is necessary for the semi-global reconstruction of α, suppose the viewpoints corresponding to Figures 15(a)-(c) are $c(t_1^n)$, $c(t_2^n)$ and $c(t_3^n)$, respectively, and the line through $c(t)$ and $\alpha(t)$ is $l(t)$. To achieve semi-global curve reconstruction, the length of the segment between $\alpha(t_1^n)$ and $\alpha(t_3^n)$ must diminish if and only if $\alpha(t_1^n)$ asymptotically approaches a visual event curve that potentially bounds a reconstructible surface region. Now consider Figure 17. Since $\alpha(t_3^n)$ is a cusp point, the line $l(t_3^n)$ is along an asymptote at $\alpha(t_3^n)$. Therefore, if $\psi(t)$ is the angle between $l(t)$ and the corresponding asymptote at $\alpha(t)$, we can conclude that a necessary and sufficient condition for the curve point $\alpha(t_3^n)$ to become a cusp visible rim point is that $\psi(t)$ goes to zero as t approaches t_3^n.

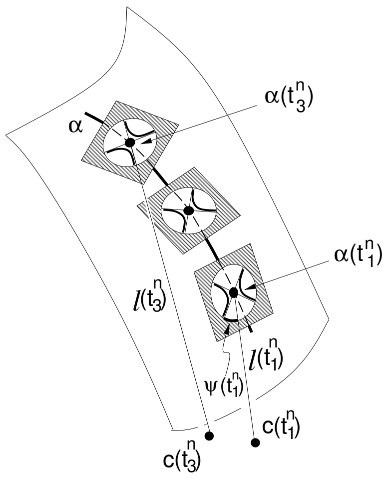

Figure 17: Geometry of the reconstruction of a segment of curve α in Figure 15. Viewpoints $c(t_1^n), c(t_3^n)$ correspond to Figures 15(a) and (c), respectively. The tangent plane and Dupin's indicatrix of points $\alpha(t_1^n)$ and $\alpha(t_3^n)$ is also shown.

Clearly, if $\psi(t_1^n)$ is large, the length of the segment between $\alpha(t_1^n)$ and $\alpha(t_3^n)$ will also be large. It is therefore necessary to require $\psi(t_1^n)$ to be large. But how large can we make $\psi(t_1^n)$? If $\psi(t_1^n)$ is too large, the line $l(t_1^n)$ may approach the other asymptote at $\alpha(t_1^n)$; the best we can do is to ensure that $c(t_1^n)$ is in the *middle* of the visibility arc, which in this case is bounded by the two asymptotes at $\alpha(t_1^n)$. At that viewpoint, $\psi(t_1^n)$ will form equal angles with both asymptotes at $\alpha(t_1)$.

Obeying Rule 2 is quite easy: The observer determines the extent of the visibility arc containing $c(t_1^n)$, and then moves to the middle of that arc. To measure the extent of the visibility arc, the observer can simply move on $T_{\alpha(t_1^n)}(S)$ first in a clockwise and then in a counterclockwise direction, terminating its motions when a cusp or T-junction is formed at the projection of $\alpha(t_1^n)$.

By following the above rules, the observer is now able to perform the semi-global curve reconstruction task: The finiteness condition ensures that the visibility arcs of $\alpha(t_1^n)$ are well-defined as n goes to infinity, and that the distance between $\alpha(t_1^n)$ and $\alpha(t_3^n)$ will diminish if and only if the visibility arc at $\alpha(t_1^n)$ degenerates to a point. This occurs only when $\alpha(t_1^n)$ approaches one of the visual event curves listed in Theorem 4.

6.2 Global Curve Reconstruction

In the global curve reconstruction task the observer must reconstruct the whole curve drawn on the surface if this is possible, or a segment of that curve whose endpoints either lie on or asymptotically approach the boundary of a reconstructible surface region. This task is harder to perform than semi-global curve reconstruction because if the whole curve cannot be reconstructed, the endpoints of the reconstructed segment must lie on visual event curves that are *actual*, not just *potential*, boundaries of a reconstructible surface region. To achieve this task the observer must obey the three rules guaranteeing semi-global reconstruction of a curve. The following theorem shows that in addition to these rules, a fourth rule must also be obeyed. Its proof follows from the proof of Theorem 5 and is omitted.

Theorem 6 (Global curve reconstruction rules) *Let α be a finite-length curve drawn on the surface. If α intersects the visual event curves associated with tangent-crossing events at most a finite number of times, and the following four rules are obeyed by the observer, the observer will provably perform the global curve reconstruction task:*

Rule 1: *When choosing the point p on which to apply the Patch Reconstruction Behaviors, always select a point of intersection of α with the visible rim.*

Rule 2: *Always execute the Ordinary Patch Reconstruction Behavior after first moving to a viewpoint c corresponding to the middle of a visibility arc of p.*

Rule 3: *When executing the Ordinary Patch Reconstruction Behavior starting from an initial viewpoint c, move on a great circle around the surface and stop only if c is reached again, or if the endpoint of the segment of α being reconstructed coincides with a cusp, T-junction or degenerate visible rim point.*

Rule 4: *In order to reconstruct a patch around the selected point p, always apply the Degenerate Patch Reconstruction Behavior independently of whether p is an ordinary, cusp, T-junction or degenerate visible rim point.*

Rules 1-3 are identical to those used for performing the semi-global reconstruction task. To see why the fourth rule is necessary for achieving global curve reconstruction, suppose the observer obeys only the first three rules to perform the task for the curve in Figure 15. Then, the length of the segment reconstructed at the n-th iteration of the Incremental Reconstruction Behavior may diminish even when these endpoints asymptotically approach a point on a visual event curve that does not bound a reconstructible surface region (Figure 12(c)). These points have at least one visibility arc.

Rule 4 requires the observer to reconstruct several patches around the selected point p, by moving to the middle viewpoints of *all* the visibility arcs of p and then executing the Ordinary Patch Reconstruction Behavior starting at each one of those viewpoints. By obeying this rule, the length of the reconstructed segment can diminish only if *all* visibility arcs at p diminish. Since this occurs only when p approaches a point bounding a reconstructible surface region, Rules 1-4 guarantee that global curve reconstruction is always achieved.

Enforcing Rule 4 implies that the whole surface reconstruction process is simplified: The observer needs to use only two of the four Patch Reconstruction Behaviors (i.e., the Ordinary Patch Reconstruction Behavior and the Degenerate Patch Reconstruction Behavior) to perform global curve reconstruction. The disadvantage of this rule, however, is that by executing the Degenerate Patch Reconstruction Behavior the observer will perform larger motions than those dictated in the Cusp and T-junction Patch Reconstruction Behaviors.

6.3 Global Surface Reconstruction

In this section we consider the global surface reconstruction task. The observer must now reconstruct not only points lying on a single curve drawn on the surface that intersects the visible rim at the initial viewpoint, but must also reconstruct points lying on *every* such curve that can be drawn on the surface. The following theorem shows how this task can be performed using the Incremental Reconstruction Behavior and the Patch Reconstruction Behaviors (Figure 18).

Theorem 7 (Global surface reconstruction rules) *If the surface is such that only a finite number of visual event curves associated with tangent-crossing events intersect at a single point, and the following four rules are obeyed by the observer, the observer will provably perform the global surface reconstruction task:*

Rule 1: *Always choose the point p on which to apply the Patch Reconstruction Behaviors so that the boundary of the already-reconstructed surface region expands uniformly.*

Rule 2: *Always execute the Ordinary Patch Reconstruction Behavior after first moving to a viewpoint c corresponding to the middle of a visibility arc of p.*

Rule 3: *When executing the Ordinary Patch Reconstruction Behavior starting from an initial viewpoint c, move on a great circle around the surface and stop only if c is reached again or the visible rim segment initially containing p (and all visible rim segments splitting from or merging with it) disappears.*

Rule 4: *In order to reconstruct a patch around the selected point p, always apply the Degenerate Patch Reconstruction Behavior independently of whether p is an ordinary, cusp, T-junction, or degenerate visible rim point.*

See Appendix B for the proof. The "finiteness" condition in the theorem simply states that global surface reconstruction can only be achieved if the surface is not infinitely complicated. Rules 2 and 4 are identical to those used to perform the global curve reconstruction task. Rule 3 is a generalization of the corresponding rule used in the global curve reconstruction task in the following sense. When performing the global curve reconstruction task, as in the example in Figure 15(a)-(c), the observer was required to stop only after an endpoint of the visible rim "slid over" the curve drawn on the surface. In the global surface reconstruction task, the same rule must hold for *every* curve that we can draw on the surface. This would require the observer to move upward until the visible rim segment in Figure 15(a) containing $\alpha(t_1^n)$ disappears or, equivalently, the two segments β_1, β_2 in Figure 15(c) disappear), or until the initial viewpoint is reached again.

Rule 1 is also a generalization of the corresponding rule used in the global curve reconstruction task. It requires the observer to reconstruct a patch in the neighborhood of every curve that can be drawn on the surface, intersects the visible rim at the initial viewpoint, and intersects the boundary of the reconstructed points after a finite number of iterations of the Incremental Reconstruction Behavior. The need for such a rule should be clear since if the reconstruction

Incremental Reconstruction Behavior	
Rules that must be obeyed	
Rule 1:	Select points uniformly on the boundary of the reconstructed region
Rule 4:	Always execute the Degenerate Patch Reconstruction Behavior independently of the type of p

Degenerate Patch Reconstruction Behavior	
Rule that must be obeyed	
Rule 2:	Always move to the middle of a visibility arc

Ordinary Patch Reconstruction Behavior	
Rule that must be obeyed	
Rule 3:	Move on a great circle around the surface and stop when either the initial position is reached again, or the visible rim segment initially containing p (and all segments splitting from it or merging with it) disappear

Figure 18: Behaviors used to perform the global surface reconstruction task. By obeying Rule 4 of Theorem 7, the observer only needs to use two of the four Patch Reconstruction Behaviors resulting in a simpler control structure and larger observer motions than those implied by Figure 16. Also shown are the rules that must be obeyed when each of the behaviors are executed. These rules "ground" the steps summarized in Figure 16.

process does not terminate and the reconstructed region is expanded in only one direction, some pieces of the boundary of the already-reconstructed points will never be expanded.[7]

[7] To obey Rule 1, the Incremental Reconstruction Behavior is executed in *cycles*. At the beginning of cycle N, the observer selects enough points on the boundary of the reconstructed points so that their distance is less than $\epsilon/2^N$, where $\epsilon > 0$ is an *a priori* defined constant.

By obeying Rules 1-4, the observer is able to guarantee global surface reconstruction. The Incremental Reconstruction Behavior terminates after a finite number of steps precisely when the whole surface is reconstructible. Otherwise, the set of points reconstructed will converge to the reconstructible surface regions intersecting the visible rim at the initial viewpoint. Furthermore, in this case, it is easy to show that *no behavior* allows global surface reconstruction to be achieved in a finite number of steps: When the visible rim touches a visual event curve on the boundary of a reconstructible surface region, it touches it at exactly one point, making it impossible to reconstruct the surface in every neighborhood of such a curve in a finite number of steps.

7 Concluding Remarks

We have demonstrated that an active, monocular observer can use a collection of simple viewpoint-control behaviors to recover a global description of a smooth, non-convex object from the occluding contour. The regions that are reconstructed on the object's surface can be accurately characterized, and depend only on qualitative, global shape properties of the surface.

The use of an active observer produces a qualitative change in the way vision algorithms are designed and analyzed. The observer's ability to *purposefully* control viewpoint enables the use of a set of local, visually-guided behaviors that are simple enough to be executed in real time and permit *global* reconstruction when they are appropriately combined. The reason is that viewpoint control is not used merely to change the shape of the occluding contour in an arbitrary manner (as in existing approaches), but it is used to change it in a well-defined way. This allows precise statements to be made about the progress of the global surface reconstruction process, and allows its outcome to be accurately controlled.

The current limitations of our approach come from (1) the use of an observer that moves on a sphere around the object, and (2) the lack of optimizations in controlling the observer's motion. By controlling distance to the object it might be possible to reconstruct larger regions on the surface. On the other hand, given the minimal surface shape information on which our behaviors rely, any performance improvements (e.g., minimizing the observer motion required to reconstruct a specific local patch on the surface) would only have a local effect. Generalizing these behaviors to control distance to the object and to take into account efficiency considerations is an important future direction of research. We are also in the process of implementing these viewpoint-control behaviors using a camera mounted on a Hitachi M6100 robot arm.

We believe that our active approach of moving toward viewpoints that are closely related to the geometry of the object is a very important and general one. For example, the ability to "draw" on the object's surface the visual event curves listed in Theorem 4 and determine their type is a simple by-product of the execution of the behaviors we developed (e.g., by keeping track of the extent of the visibility arcs of the points selected by the Incremental Reconstruction Behavior). Another by-product of this process is the discovery of the supporting planes of the object, a piece of information that can be very useful for either directly matching the viewed surface with a stored model, or for doing so after first moving to a "standard viewpoint" related to the discovered supporting planes.

How can the behaviors we developed generalize to perform more qualitative tasks? We are currently investigating the task of visually exploring the surface of an object in order to find a specific surface feature (e.g., your cup does not have a blue handle but *mine* does [62]). This task is not as constrained as the surface reconstruction task; what is needed is a way to control the motion of the *occlusion boundary* (i.e., the boundary of the visible points on the surface), which is a superset of the visible rim. The four local Patch Reconstruction Behaviors developed in this paper are a step in this direction: They control the motion of the visible rim and do not rely on any three-dimensional information. The crucial issue that needs to be addressed is their integration. For this, it is necessary to find ways of representing the boundary of the surface already made visible so that patches around points on this boundary can be made visible by an appropriate viewpoint-control behavior. Although this, in general, requires the recovery of some three-dimensional information, a simple and qualitative representation of this boundary can be derived for surfaces with non-uniform texture by representing the texture of the surface in the boundary's vicinity.

References

[1] J. J. Gibson, *The Ecological Approach to Visual Perception.* Houghton Mifflin, 1979.

[2] E. J. Gibson, "Exploratory behavior in the development of perceiving, acting, and the acquiring of knowledge," *Ann. Rev. Psychol*, vol. 39, pp. 1–41, 1988.

[3] A. L. Yarbus, *Eye Movements and Vision.* Plenum Press, 1967.

[4] R. L. Klatzky and S. J. Lederman, "The intelligent hand," *The Psychology of Learning and Motivation*, vol. 21, pp. 121–151, 1987.

[5] J. J. Gibson, *The Senses Considered as Perceptual Systems.* Houghton Mifflin Company, 1966.

[6] R. Bajcsy and M. Campos, "Active and exploratory perception," *Computer Vision, Graphics, and Image Processing: Image Understanding*, vol. 56, no. 1, pp. 31–40, 1992.

[7] D. H. Ballard and C. M. Brown, "Principles of animate vision," *Computer Vision, Graphics, and Image Processing: Image Understanding*, vol. 56, no. 1, pp. 3–21, 1992. Special Issue on Purposive, Qualitative, Active Vision.

[8] Y. Aloimonos, "Purposive and qualitative active vision," in *Proc. Int. Conf. on Pattern Recognition*, pp. 346–360, 1990.

[9] A. L. Abbott, "Selective fixation control for machine vision: A survey," in *Proc. IEEE Conf. Syst. Man Cybern.*, 1991.

[10] D. H. Ballard and A. Ozcandarli, "Eye fixation and early vision: Kinetic depth," in *Proc. 2nd Int. Conf. on Computer Vision*, pp. 524–531, 1988.

[11] D. H. Ballard, "Reference frames for animate vision," in *Proc. Int. Joint Conf. Artificial Intelligence*, pp. 1635–1641, 1989.

[12] R. C. Nelson and J. Aloimonos, "Obstacle avoidance using flow field divergence," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 11, no. 10, pp. 1102–1106, 1989.

[13] T. J. Olson and D. J. Coombs, "Real-time vergence control for binocular robots," *Int. J. Computer Vision*, vol. 7, no. 1, pp. 67–89, 1991.

[14] K. Pahlavan and J.-O. Eklundh, "A head-eye system — analysis and design," *Computer Vision, Graphics, and Image Processing: Image Understanding*, vol. 56, no. 1, pp. 41–56, 1992. Special Issue on Purposive, Qualitative, Active Vision.

[15] E. Krotkov, "Focusing," *Int. J. Computer Vision*, vol. 1, pp. 223–237, 1987.

[16] R. Sharma and Y. Aloimonos, "Visual motion analysis under interceptive behavior," in *Proc. Computer Vision and Pattern Recognition*, pp. 148–153, 1992.

[17] R. Bajcsy and C. Tsikos, "Perception via manipulation," in *Robotics Research 4* (R. C. Bolles and B. Roth, eds.), pp. 237–244, MIT Press, 1988.

[18] C. J. Tsikos and R. K. Bajcsy, "Segmentation via manipulation," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 7, no. 3, pp. 306–319, 1991.

[19] D. H. Ballard, "Animate vision," *Artificial Intelligence*, vol. 48, pp. 57–86, 1991.

[20] L. E. Wixson, "Exploiting world structure to efficiently search for objects," Tech. Rep. 434, University of Rochester, July 1992.

[21] C. Brown, "Prediction and cooperation in gaze control," *Biological Cybernetics*, vol. 63, pp. 61–70, 1990.

[22] A. L. Abbott and N. Ahuja, "Active surface reconstruction by integrating focus, vergence, stereo, and camera calibration," in *Proc. 3rd Int. Conf. on Computer Vision*, pp. 489–492, 1990.

[23] E. Krotkov, *Active Computer Vision by Cooperative Focus and Stereo.* Springer-Verlag, 1989.

[24] R. D. Rimey and C. M. Brown, "Task-oriented vision with multiple bayes nets," in *Active Vision* (A. Blake and A. Yuille, eds.), pp. 217–236, MIT Press, 1992.

[25] R. A. Brooks, "A robot that walks: Emergent behaviors from a carefully evolved network," in *Proc. IEEE Robotics Automat. Conf.*, pp. 692–696, 1989.

[26] Y. Aloimonos, I. Weiss, and A. Bandyopadhyay, "Active vision," in *Proc. 1st Int. Conf. on Computer Vision*, pp. 35–54, 1987.

[27] R. Cipolla and A. Blake, "Surface shape from the deformation of apparent contours," *Int. J. Computer Vision*, vol. 9, no. 2, pp. 83–112, 1992.

[28] K. N. Kutulakos and C. R. Dyer, "Recovering shape by purposive viewpoint adjustment," in *Proc. Computer Vision and Pattern Recognition*, pp. 16–22, 1992.

[29] D. Wilkes and J. K. Tsotsos, "Active object recognition," in *Proc. Computer Vision and Pattern Recognition*, pp. 136–141, 1992.

[30] E. Grosso and D. H. Ballard, "Head-centered orientation strategies in animate vision," Tech. Rep. 442, Univeristy of Rochester, October 1992.

[31] A. Blake, A. Zisserman, and R. Cipolla, "Visual exploration of free-space," in *Active Vision* (A. Blake and A. Yuille, eds.), pp. 175–188, MIT Press, 1992.

[32] S. A. Hutchinson and A. Kak, "Planning sensing strategies in a robot work cell with multi-sensor capabilities," *IEEE Trans. Robotics Automat.*, vol. 5, no. 6, pp. 765–783, 1989.

[33] P. Whaite and F. P. Ferrie, "From uncertainty to visual exploration," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 13, no. 10, pp. 1038–1049, 1991.

[34] J. Maver and R. Bajcsy, "Occlusions and the next view planning," in *Proc. IEEE Robotics Automat. Conf.*, pp. 1806–1811, 1992.

[35] C. I. Connoly, "The determination of next best views," in *Proc. IEEE Robotics Automat. Conf.*, pp. 432–435, 1985.

[36] C. Tomasi and T. Kanade, "Shape and motion from image streams under orthography: A factorization method," *Int. J. Computer Vision*, vol. 9, no. 2, pp. 137–154, 1992.

[37] H. G. Barrow and J. M. Tenenbaum, "Interpreting line drawings as three-dimensional images," *Artificial Intelligence*, vol. 17, pp. 75–116, 1981.

[38] M. Brady, J. Ponce, A. Yuille, and H. Asada, "Describing surfaces," *Computer Graphics and Image Processing*, vol. 32, pp. 1–28, 1985.

[39] D. J. Kriegman and J. Ponce, "On recognizing and positioning curved 3-d objects from image contours," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 12, no. 12, pp. 1127–1137, 1990.

[40] D. Marr and H. K. Nishihara, "Visual information processing: Artificial intelligence and the sensorium of sight," *Technology Review*, vol. 81, pp. 2–23, 1978.

[41] J. Malik, "Interpreting line drawings of curved objects," *Int. J. Computer Vision*, vol. 1, pp. 73–103, 1987.

[42] J. Ponce, D. Chelberg, and W. B. Mann, "Invariant properties of straight homogeneous generalized cylinders and their contours," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 11, no. 9, pp. 951–966, 1990.

[43] W. Richards, J. J. Koenderink, and D. D. Hoffman, "Inferring 3d shapes from 2d silhouettes," in *Natural Computation* (W. Richards, ed.), pp. 125–136, Cambridge, MA: MIT Press, 1988.

[44] W. Richards, B. Dawson, and D. Whittington, "Encoding contour shape by curvature extrema," in *Natural Computation* (W. Richards, ed.), pp. 83–98, Cambridge, MA: MIT Press, 1988.

[45] W. B. Seales and C. R. Dyer, "Viewpoint from occluding contour," *Computer Vision, Graphics, and Image Processing: Image Understanding*, vol. 55, no. 2, pp. 198–211, 1992.

[46] F. Ulupinar and R. Nevatia, "Using symmetries for analysis of shape from contour," in *Proc. 2nd Int. Conf. on Computer Vision*, pp. 414–426, 1988.

[47] R. Vaillant and O. D. Faugeras, "Using extremal boundaries for 3-d object modeling," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 14, no. 2, pp. 157–173, 1992.

[48] W. B. Thompson, "Structure-from-motion by tracking occlusion boundaries," in *Proc. Workshop on Visual Motion*, pp. 201–203, 1989.

[49] W. B. Thompson and J. S. Painter, "Qualitative constraints for structure-from-motion," *Computer Vision, Graphics, and Image Processing: Image Understanding*, vol. 56, no. 1, pp. 69–77, 1992. Special Issue on Purposive, Qualitative, Active Vision.

[50] P. Giblin and R. Weiss, "Reconstruction of surfaces from profiles," in *Proc. 1st Int. Conf. on Computer Vision*, pp. 136–144, 1987.

[51] J. J. Koenderink, "An internal representation for solid shape based on the topological properties of the apparent contour," in *Image Understanding 1985-86* (W. Richards and S. Ullman, eds.), pp. 257–285, Norwood, NJ: Ablex Publishing Co., 1987.

[52] J. J. Koenderink, *Solid Shape.* MIT Press, 1990.

[53] S. Petitjean, J. Ponce, and D. J. Kriegman, "Computing exact aspect graphs of curved objects: Algebraic surfaces," *Int. J. Computer Vision*, vol. 9, no. 3, pp. 231–255, 1992.

[54] J. H. Rieger, "The geometry of view space of opaque objects bounded by smooth surfaces," *Artificial Intelligence*, vol. 44, pp. 1–40, 1990.

[55] M. P. D. Carmo, *Differential Geometry of Curves and Surfaces.* Englewood Cliffs, NJ: Prentice-Hall Inc., 1976.

[56] Y. L. Kergosien, "La famille des projections orthogonales d'une surface et ses singularités," *C. R. Acad. Sc. Paris*, vol. 292, pp. 929–932, 1981.

[57] P. Giblin and M. G. Soares, "On the geometry of a surface and its singular profiles," *Image and Vision Computing*, vol. 6, no. 4, pp. 225–234, 1988.

[58] J. J. Koenderink and A. J. van Doorn, "The singularities of the visual mapping," *Biological Cybernetics*, vol. 24, pp. 51–59, 1976.

[59] J. Ponce, S. Petitjean, and D. Kriegman, "Computing exact aspect graphs of curved objects: Algebraic surfaces," in *Proc. Second European Conference on Computer Vision*, 1992.

[60] C. K. Cowan and P. D. Kovesi, "Automatic sensor placement from vision task requirements," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 10, no. 3, pp. 407–416, 1988.

[61] K. Tarabanis and R. Y. Tsai, "Computing viewpoints that satisfy optical constraints," in *Proc. Computer Vision and Pattern Recognition*, pp. 152–158, 1991.

[62] M. J. Swain and D. H. Ballard, "Color indexing," *Int. J. Computer Vision*, vol. 7, no. 1, pp. 11–32, 1991.

[63] E. Grosso, M. Tistarelli, and G. Sandini, "Active/dynamic stereo for navigation," in *Proc. Second European Conference on Computer Vision*, pp. 516–525, 1992.

[64] J. J. Koenderink and A. J. van Doorn, "The shape of smooth objects and the way contours end," in *Natural Computation* (W. Richards, ed.), pp. 115–124, Cambridge, MA: MIT Press, 1988.

[65] J. J. Koenderink and A. J. van Doorn, "A description of the structure of visual images in terms of an ordered hierarchy of light and dark blobs," in *Proc. Second Int. Visual Psychophysics and Medical Imaging Conf.*, pp. 173–176, 1981.

Appendix A: Proofs for Section 3 Theorems

A.1 Proof of Theorem 1

Let p be an ordinary visible rim point. Consider the epipolar plane, Λ, defined by the vector $\mathbf{v}(t)$ and the line segment connecting p and $c(t)$, and the intersection α of S with Λ in the neighborhood of p.

Suppose that the observer does not move on the tangent plane of the surface at p, i.e., $\mathbf{n}(p) \cdot \mathbf{v}(t) \neq 0$. In this case, $\Lambda \neq T_p(S)$ and α is a regular curve in the neighborhood of p [51]. Since p is an ordinary point, p must be a convex point of α, and the open line segment connecting p and $c(t)$ does not intersect the surface (Figure 19). This implies that p does not become occluded by a distant point of S under an infinitesimal viewpoint change on Λ.

The visibility of p (and of points on α close to p) in this case is determined by the sign of $\mathbf{n}(p) \cdot [p - c(t)]$ [51]. Therefore, changes in the visibility state of p under infinitesimal observer motion occur due to changes in the sign of this dot product. Since p is a visible rim point, this dot product is zero at position $c(t)$. Therefore, the visibility of p under an infinitesimal viewpoint change depends on the sign of the derivative $\{\mathbf{n}(p) \cdot [p - c(t)]\}'$. We have:

$$\frac{d}{dt}\{\mathbf{n}(p) \cdot [p - c(t)]\} \; = \; \mathbf{n}(p) \cdot \frac{d}{dt}[p - c(t)] + [\frac{d}{dt}\mathbf{n}(p)] \cdot [p - c(t)] \; = \; -\mathbf{n}(p) \cdot \mathbf{v}(t)$$

If the observer moves on p's tangent plane, p may become occluded by points in the neighborhood of p but will always remain on the rim. It will remain visible unless p is hyperbolic and the line connecting $c(t)$ and p is along an asymptotic direction of the surface at p [28]. This, however, cannot occur since p is ordinary.

A.2 Proof of Theorem 2

First note that point p must be hyperbolic [64]. Furthermore, since p projects to a cusp point on the occluding contour, the line segment l connecting p and c does not intersect the surface elsewhere. l is contained in $T_p(S)$ and runs along an asymptotic direction of S at p.

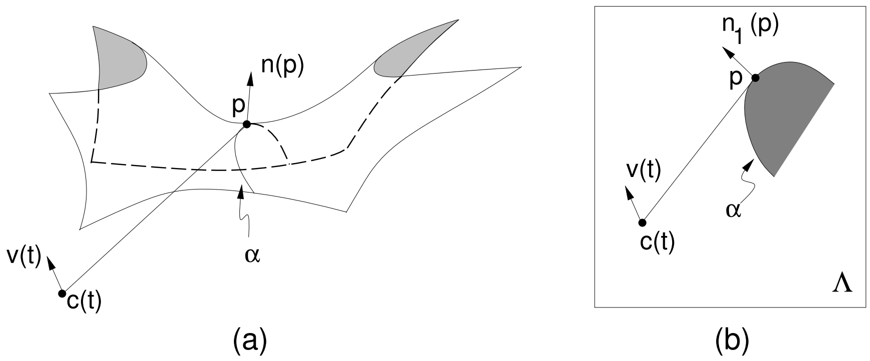

Figure 19: Inducing the visibility of points in a neighborhood of a hyperbolic point p. (a) The curve α is the curve of intersection of S with the epipolar plane Λ. If $\mathbf{v}(t) \cdot \mathbf{n}(p) > 0$, p will be matched to a currently-occluded point on α under an infinitesimal viewpoint change along $\mathbf{v}(t)$. (b) A face-on view of the plane Λ. The outward normal $n_1(p)$ of α at p is the projection of $\mathbf{n}(p)$ on the plane Λ. The geometry of the intersection of S with the epipolar planes corresponding to visible rim points *close* to p is also similar to the one shown.

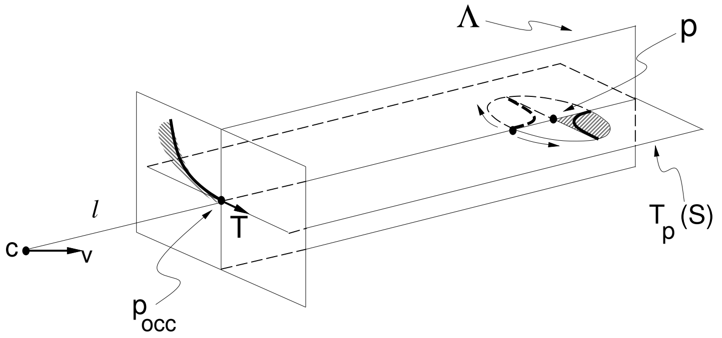

Figure 20: Visibility-preserving motions on $T_p(S)$. For readability, the surface in the neighborhood of p is not shown. Λ is the normal plane at p in the direction of l. The hyperbola on $T_p(S)$ is Dupin's indicatrix at p. The plane Λ separates the surface points around p that are above $T_p(S)$ into two sets: When the surface is viewed along l, all points to the right of Λ (i.e., in the direction of $\mathbf{T}$) are visible from p; there are both visible and occluded points to the left of Λ. The line segment l is an asymptote of the component of Dupin's indicatrix that lies to the *left* of Λ.

In [28] we showed that if l is along an asymptotic direction, infinitesimal viewpoint changes on $T_p(S)$ that move the observer away from l can be classified into two categories depending on whether or not they force p to become occluded. From the analysis in [28], it follows that if p remains visible at the new viewpoint, the curvature of the occluding contour at p becomes finite, and p remains on the visible rim after the change in viewpoint (i.e., p becomes an *ordinary* visible rim point at the new viewpoint). Figure 20 shows why the directions $\mathbf{v}$ ensuring the visibility of p after the viewpoint change satisfy $\mathbf{v} \cdot \mathbf{T} > 0$.

Appendix B: Proofs for Section 6 Theorems

B.1 Proof of Theorem 5

Let $\alpha(t^n)$ be the point on α selected at the n-th iteration of the Incremental Reconstruction Behavior. First, note that if α is not reconstructed in a finite number of steps, the limit, $\alpha(t^\infty)$, exists: Since the points $\alpha(t^n)$ belong to a curve of finite length, and the length of the reconstructed portion of the curve is an increasing function, the sequence $(\alpha(t^n))_n$ has a limit point.

To prove the proposition we study how the visibility arcs of $\alpha(t^n)$ change as n goes to infinity. The finiteness assumption of Theorem 5 ensures that the visibility arcs of $\alpha(t^n)$ for large n are well-defined:

Lemma 1 *If α intersects the visual event curves associated with tangent-crossing events a finite number of times, and if N is sufficiently large, the topology of the intersection $S \cap T_{\alpha(t^n)}(S)$ does not change for $n > N$.*

Proof: Consider the intersection $S \cap T_{\alpha(t)}(S)$, for $0 \leq t \leq T$. Its topology will change only if the number of contact points of S with $T_{\alpha(t)}(S)$ changes. It suffices to consider the case in which the number of contact points of $T_{\alpha(t)}(S)$ with S for $t < t_c$ and $t > t_c$ is one less than those for $t = t_c$. Then, the plane $T_{\alpha(t_c)}(S)$ touches the surface at two distinct points, $\alpha(t_c)$, and another point, p. Consequently, $\alpha(t_c)$ belongs to a visual event curve, γ_1 or γ_2, associated with a tangent-crossing event. Now note that $\alpha(t_c \pm \epsilon)$ ($\epsilon > 0$ small) does not belong to these curves since $T_{\alpha(t_c \pm \epsilon)}(S)$ contacts the surface at exactly one point. Hence, $\alpha(t_c)$ is a point of intersection of α with either γ_1 or γ_2. Since there is only a finite number of such intersection points, all $\alpha(t^n)$ will fall in a single interval between these points, if $n > N$ and N is sufficiently large. □

The boundaries of the visibility arcs of $\alpha(t^n)$ belong to two types of lines on $T_{\alpha(t^n)}(S)$: *Bitangent lines* through $\alpha(t^n)$, i.e., lines that touch the surface at at least two distinct points, one of which is $\alpha(t^n)$, and asymptotes of $\alpha(t^n)$. To show that the semi-global curve reconstruction task is achieved for α, we consider the angles formed by the bitangent lines and asymptotes at $\alpha(t^n)$ as n goes to infinity. The main idea of the proof is to observe that we can determine these angles from the observer's motions during the n-th iteration of the Incremental Reconstruction

Behavior.

In particular, after the observer selects point $\alpha(t^n)$, the following steps are then taken: (1) The observer moves to a previously-visited viewpoint, c^n_{vis}, from which $\alpha(t^n)$ was on the visible rim. (2) Since the observer obeys Rule 3, we can assume without loss of generality that $\alpha(t^n)$ is a cusp, T-junction or degenerate visible rim point at c^n_{vis}. In either case, Rule 2 forces the observer to move to the middle, c^n_{mid}, of a visibility arc of $\alpha(t^n)$. (3) The observer executes the Ordinary Patch Reconstruction Behavior to reconstruct a new segment of α, ending at $\alpha(t^n + \delta t^n)$. The Ordinary Patch Reconstruction Behavior terminates when $\alpha(t^n + \delta t^n)$ belongs to the visible rim, and Rule 3 allows us to assume without loss of generality that if c^n_{stop} is the observer's viewpoint when this occurs, $\alpha(t^n + \delta t^n)$ is either a cusp, T-junction or degenerate point.

We now use the following two observations:

1. The visibility arc containing c^n_{mid} can be of three types: Type AA, whose endpoints lie on the two asymptotes at $\alpha(t^n)$; type AB, whose endpoints lie on one asymptote of $\alpha(t^n)$ and one bitangent line through $\alpha(t^n)$ on $T_{\alpha(t^n)}(S)$; and type BB, whose endpoints lie on two bitangent lines through $\alpha(t^n)$ on $T_{\alpha(t^n)}(S)$.

2. The line through c^n_{stop} and $\alpha(t^n + \delta t^n)$ is either a bitangent line (B) through $\alpha(t^n + \delta t^n)$ on $T_{\alpha(t^n+\delta t^n)}(S)$, or an asymptote (A) at $\alpha(t^n + \delta t^n)$.

In all, there are six different combinations of the above cases, which we represent by the string **X-Y** where $\mathbf{X} = \{AA, AB, BB\}$, and $\mathbf{Y} = \{A, B\}$. Without loss of generality we prove the proposition for each of the above cases separately. Only three of the six cases are considered below. The remaining three cases can be treated in an identical manner.

B.1.1 Case AA-A

This case is shown in Figure 21(a).

The line $c^n_{stop}\alpha(t^n+\delta t^n)$ is an asymptote at $\alpha(t^n+\delta t^n)$. Without loss of generality assume that $\alpha(t^n)$ is also hyperbolic, and let l^n be the corresponding asymptote at $\alpha(t^n)$. Since $\lim_{n\to\infty} \alpha(t^n) = \lim_{n\to\infty} \alpha(t^n + \delta t^n)$, we get

$$(1) \qquad \lim_{n\to\infty} \angle(l^n, c^n_{stop}\alpha(t^n + \delta t^n) = 0$$

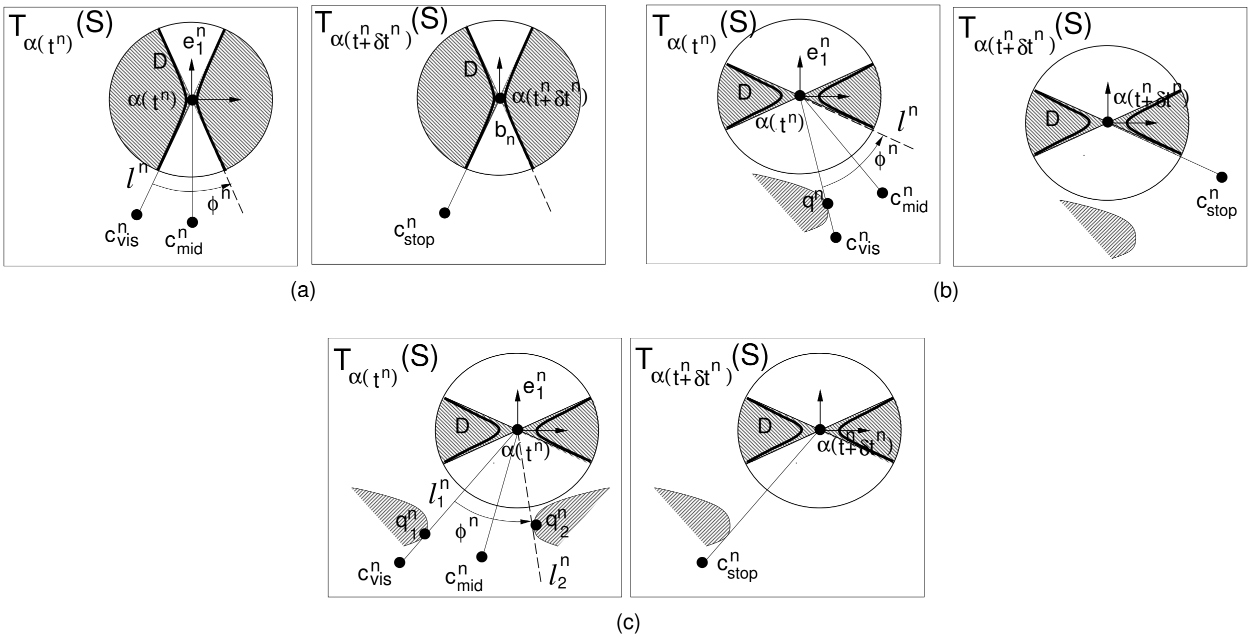

Figure 21: Representing the configurations of the asymptotes and bitangent lines through $\alpha(t^n)$ and $\alpha(t^n + \delta t^n)$. For each of the three cases, the left figure corresponds to a "top" view of the plane $T_{\alpha(t^n)}(S)$ while the right figure corresponds to a "top" view of the plane $T_{\alpha(t^n+\delta t^n)}(S)$. (a) Case AA-A. (b) Case AB-A. (c) Case BB-B.

Because the motion of the visible rim during the execution of the Ordinary Patch Reconstruction Behavior between c^n_{mid} and c^n_{stop} describes S in the neighborhood of $\alpha(t^n)$, it follows that[8]

(2) $$\lim_{n\to\infty} c^n_{stop} = \lim_{n\to\infty} c^n_{mid}$$

From Eqs. (1) and (2) we conclude that

(3) $$\lim_{n\to\infty} \angle(l^n, c^n_{mid}\alpha(t^n)) = 0$$

Since c^n_{mid} is the midpoint of a visibility arc of $\alpha(t^n)$ of type AA, we have

(4) $$\angle(l^n, c^n_{mid}\alpha(t^n)) = \phi^n/2$$

[8] To see this, note that the observer's motion from c^n_{mid} to c^n_{stop} is along a smooth curve, c. Since the segment $\alpha(t)$, $(t^n < t < t^n + \delta t^n)$ can be described by the epipolar parameterization, the observer's motion can be described by the curve $c(t), t^n < t < t^n + \delta t^n$ such that $c(t)$ is the observer's position when $\alpha(t)$ is on the visible rim. The mapping $m : \alpha \to c$ is smooth, and consequently $\|c^n_{stop} - c^n_{mid}\| \to 0$ if and only if $\|\alpha(t^n + \delta t^n) - \alpha(t^n)\| \to 0$.

where ϕ_n is the angle between the two asymptotes at $\alpha(t^n)$. Hence, from Eqs. (3) and (4) it follows that

$$\lim_{n\to\infty} \phi^n = 0 \tag{5}$$

This implies that the direction of the asymptotes tends toward the first principal direction, e_1^∞, of the surface at $\alpha(t^\infty)$ (Figure 21(a)).

The angle ψ^n between the asymptotes and the first principal direction, e_1^n, of the surface at $\alpha(t^n)$ is given by

$$\tan\psi^n = \sqrt{-\frac{k_1^n}{k_2^n}} = \tan(\phi^n/2) \tag{6}$$

where k_1^n, k_2^n are the first and second principal curvatures of the surface at $\alpha(t^n)$, respectively.

Since the surface is smooth, k_2^n is always bounded. Hence, from Eqs. (5) and (6) it follows that $\lim_{n\to\infty} k_1^n = 0$. For generic surfaces, this implies that $\alpha(t^n)$ is on a parabolic curve bounding a surface concavity.

B.1.2 Case AB-A

This case is shown in Figure 21(b).

The line $c_{stop}^n\alpha(t^n+\delta t^n)$ is an asymptote at $\alpha(t^n+\delta t^n)$. Without loss of generality assume that $\alpha(t^n)$ is also hyperbolic, and let l^n be the corresponding asymptote at $\alpha(t^n)$. Since $\lim_{n\to\infty}\alpha(t^n) = \lim_{n\to\infty}\alpha(t^n+\delta t^n)$, we get

$$\lim_{n\to\infty} \angle(l^n, c_{stop}^n\alpha(t^n+\delta t^n)) = 0 \tag{7}$$

As in Case AA-A, we have

$$\lim_{n\to\infty} \angle(l^n, c_{mid}^n\alpha(t^n)) = 0 \tag{8}$$

Since the visibility arc containing c_{mid}^n is of type AB, the following equations hold:

$$\angle(l^n, c_{mid}^n\alpha(t^n)) \geq \phi^n/2 \tag{9}$$

$$\angle(l_1^n, c_{mid}^n\alpha(t^n)) = \phi^n/2 \tag{10}$$

$$\angle(l_2^n, c_{mid}^n\alpha(t^n)) = \phi^n/2 \tag{11}$$

where l_1^n and l_2^n are the asymptote and bitangent line, respectively, bounding the visibility arc containing c_{mid}^n, and ϕ^n is the angle they form. From Eqs. (8) and (9) it follows that

$$\lim_{n\to\infty} \phi^n = 0 \tag{12}$$

Since l_2^n is a bitangent line, it contacts the surface at an additional point, q^n. Lemma 1 tells us that we can choose N large enough so that the topology of the intersection $T_{\alpha(t^n)}(S) \cap S$ does not change for $n > N$. We can therefore assume without loss of generality that for $n > N$, the points q^n belong to a single curve.[9] Consequently, the limit $\lim_{n\to\infty} q^n$ exists and is equal to some point q^∞. We now distinguish two cases:

- $q^\infty \neq \alpha(t^\infty)$. Then the line $q^\infty\alpha(t^\infty)$ touches the surface at two distinct points. Furthermore, from Eqs. (10)-(12) we can conclude that the line $q^\infty\alpha(t^\infty)$ coincides with the line $\lim_{n\to\infty} l_1^n$, which is an asymptote at $\alpha(t^\infty)$. Hence, $\alpha(t^\infty)$ belongs to the curve σ_{cusp} associated with a cusp-crossing visual event. Furthermore, the visibility arc containing c_{mid}^n diminishes as $\alpha(t^n)$ approaches $\alpha(t^n)$ (i.e., it approaches a visual event curve that is a *potential* boundary of a reconstructible surface region).

- $q^\infty = \alpha(t^\infty)$. Since l_1^n and l_2^n bound the region on $T_{\alpha(t^n)}(S)$ from which $\alpha(t^n)$ is visible, and their angle, ϕ^n, tends to zero, it follows that (1) the limit $\lim_{n\to\infty} l_2^n$ exists and is an asymptote at $\alpha(t^\infty)$, and (2) the angle between $\alpha(t^n)$'s asymptotes tends to zero. From the analysis of Case AA-A it now follows that $\alpha(t^n)$ belongs to a parabolic curve bounding a surface concavity.

B.1.3 Case BB-B

This case is shown in Figure 21(c).

The line $c_{stop}^n\alpha(t^n + \delta t^n)$ is a bitangent line at $\alpha(t^n + \delta t^n)$. Without loss of generality assume that the topology of the intersection, $T_{\alpha(t)}(S) \cap S$, does not change for $t^n < t < t^n + \delta t^n$, and let

[9] Since the topology of the intersection $T_{\alpha(t^n)}(S) \cap S$ does not change for $n > N$, the number of bitangent lines in $T_{\alpha(t^n)}(S)$ will be equal to some fixed constant k. Furthermore, these lines can be grouped into k continuous families of lines whose contacts with the surface trace continuous curves. The line l_2^n will belong to one of those k families. We can therefore partition the sequence $(\alpha(t^n))_n$ into k subsequences in which l_2^n always belongs to the same family. For each such subsequence, l_2^n contacts the surface at an additional point q^n that belongs to the continuous trace of contacts of a single family of bitangent lines.

l^n be the corresponding bitangent line at $\alpha(t^n)$. Since $\lim_{n\to\infty} \alpha(t^n) = \lim_{n\to\infty} \alpha(t^n + \delta t^n)$, we get

$$\lim_{n\to\infty} \angle(l^n, c^n_{stop}\alpha(t^n + \delta t^n)) = 0 \tag{13}$$

As in Case AA-A, we have

$$\lim_{n\to\infty} \angle(l^n, c^n_{mid}\alpha(t^n)) = 0 \tag{14}$$

Since the visibility arc containing c^n_{mid} is of type BB, the following equations hold:

$$\angle(l^n, c^n_{mid}\alpha(t^n)) \geq \phi^n/2 \tag{15}$$

$$\angle(l^n_1, c^n_{mid}\alpha(t^n)) = \phi^n/2 \tag{16}$$

$$\angle(l^n_2, c^n_{mid}\alpha(t^n)) = \phi^n/2 \tag{17}$$

where l^n_1 and l^n_2 are the two bitangent lines bounding the visibility arc containing c^n_{mid}, and ϕ^n is the angle they form. From Eqs. (14) and (15) it follows that

$$\lim_{n\to\infty} \phi_n = 0 \tag{18}$$

Since l^n_1 and l^n_2 are bitangent lines through $\alpha(t^n)$, they contact the surface at points q^n_1 and q^n_2, respectively, distinct from $\alpha(t^n)$. As in Case AB-A, we assume without loss of generality that the points q^n_1 and q^n_2 belong to the trace of two continuous curves on the surface. Hence, the limits $q^\infty_1 = \lim_{n\to\infty} q^n_1$ and $q^\infty_2 = \lim_{n\to\infty} q^n_2$ exist.

We now distinguish four cases:

- $\alpha(t^\infty), q^\infty_1, q^\infty_2$ are distinct. In this case the lines $\alpha(t^\infty)q^\infty_1$ and $\alpha(t^\infty)q^\infty_2$ are bitangent lines. From Eqs. (16)-(18) we conclude that the two lines are identical, and hence $\alpha(t^\infty)q^\infty_1$ is tangent to the surface at three distinct points. Furthermore, $\alpha(t^\infty)$ is not between points q^∞_1 and q^∞_2. Consequently, $\alpha(t^\infty)$ belongs to a visual event curve τ_1 or τ_2 associated with some triple-point visual event. Furthermore, the visibility arc containing c^n_{mid} diminishes as $\alpha(t^n)$ approaches that point (i.e., it approaches a visual event curve that is a *potential* boundary of a reconstructible surface region).

- $q_1^\infty = q_2^\infty \neq \alpha(t^\infty)$. This implies that the topology of the intersection $T_{\alpha(t^\infty)}(S) \cap S$ in the neighborhood of q_1^∞ is different from that of $T_{\alpha(t^n)}(S) \cap S$ for any large n. This occurs only when $T_{\alpha(t^\infty)}(S)$ is tangent to the surface at point q_1^∞ [65], and q_1^∞ is hyperbolic. Hence, $T_{\alpha(t^\infty)}(S)$ is tangent to the surface at two distinct points and, consequently, $\alpha(t^\infty)$ belongs to a visual event curve γ_1 or γ_2 associated with a tangent-crossing event. Furthermore, the visibility arc containing c_{mid}^n diminishes as $\alpha(t^n)$ approaches that point (i.e., it approaches a visual event curve that is a *potential* boundary of a reconstructible surface region).

- $\alpha(t^\infty) \in \{q_1^\infty, q_2^\infty\}$ and $q_1^\infty \neq q_2^\infty$. Suppose $\alpha(t^\infty) = q_1^\infty$. Since the lines l_1^n, l_2^n bound a region in which $\alpha(t^n)$ is visible and $q_1^\infty = \alpha(t^\infty)$, it follows that l_1^∞ is an asymptote at $\alpha(t^\infty)$. But Eq. (18) implies that l_1^∞ also touches the surface at q_2^∞. Hence, $\alpha(t^\infty)$ belongs to a visual event curve σ_{cusp} associated with a cusp-crossing visual event. Furthermore, the visibility arc containing c_{mid}^n diminishes as $\alpha(t^n)$ approaches that point (i.e., it approaches a visual event curve that is a *potential* boundary of a reconstructible surface region).

- $\alpha(t^\infty) = q_1^\infty = q_2^\infty$. Since the lines l_1^n, l_2^n bound a region in which $\alpha(t^n)$ is visible, it follows that l_1^∞, l_2^∞ are the two asymptotes at $\alpha(t^\infty)$. Eq. (18) and our analysis in Case AA-A imply that $\alpha(t^\infty)$ is on a parabolic curve bounding a surface concavity.

B.2 Proof of Theorem 7

Suppose Γ^n is the set of surface points reconstructed after n iterations of the Incremental Reconstruction Behavior. Furthermore, suppose the reconstruction process never terminates. Let $\Gamma^\infty = \cup_{n\to\infty}\Gamma^n$, and take q to be a limit point of Γ^∞. We show that q is a limit point of a reconstructible surface region.

Let β be a curve lying in the open set $\overline{\Gamma^\infty} - \Gamma^\infty$, connecting q to the visible rim at the initial viewpoint. By definition, any point on β except its endpoint, q, will be reconstructed after a finite number of iterations of the Incremental Reconstruction Behavior. We distinguish two cases, depending on whether or not q belongs to a curve γ_1 or γ_2 corresponding to a tangent-crossing visual event.

First, suppose q does not belong to such a curve. In this case, we can find for any $\epsilon > 0$, a neighborhood Π of q of radius less than ϵ, such that Π does not intersect any visual event curves γ_1, γ_2. Since the observer obeys Rule 1 of Theorem 7 and β always intersects the boundary of reconstructed points, it follows that after a finite number of iterations of the Incremental Reconstruction Behavior, a point p will be selected that is contained in Π. After selecting p, the observer executes the Degenerate Patch Reconstruction Behavior to reconstruct the surface around p.

Now, we can define a smooth curve $\gamma(t)$ $(0 < t < \epsilon)$, such that $\gamma(0) = p$, $\gamma(\epsilon) = q$, and $\gamma(t)$ is contained in Π for all t. By its definition, q will not be contained in the patch reconstructed by the observer. Let r be the first point on γ intersecting the boundary of this patch. If $r = \gamma(t_r)$, the topology of the intersection $T_{\gamma(t)}(S) \cap S$ is the same for all $0 < t < t_r$. Hence, we can use the arguments in Theorem 5 to conclude that q is a limit point of a reconstructible surface region.

Now suppose q belongs to some curve, γ_1 or γ_2, corresponding to a tangent-crossing visual event, and let $\epsilon > 0$. Only a finite number, K, of such curves can intersect at q. These curves partition every neighborhood Π of q of radius less than ϵ into K regions, in a star-shaped fashion. After a finite number of iterations of the Incremental Reconstruction Behavior, a point p will be selected that is contained in one of those K regions. Without loss of generality, we may assume that p does not belong to the boundaries of these regions. We can now define a curve γ connecting p to q such that its trace is contained in Π, and intersecting a visual event curve, γ_1 or γ_2, only at q. This reduces the theorem's proof to the previous case.