

Exploring Objects By Invariant-Based Tangential Viewpoint Control

Kiriakos N. Kutulakos

kyros@cs.rochester.edu

Martin Jägersand

jag@cs.rochester.edu

Department of Computer Sciences
University of Rochester
Rochester, NY 14627-0226 USA

Abstract

What viewpoint control strategies are important for exploring the shape of unknown, curved objects? We argue that strategies that control viewpoint on the tangent plane of automatically-selected points projecting to the occluding contour can be very useful for local, global, qualitative, as well as quantitative shape recovery. This paper studies the design and implementation of an uncalibrated hand-eye system for tangential viewpoint control. We show that tangential viewpoint control can be accurately performed in real time using an uncalibrated camera to control a robot manipulator with at least three degrees of freedom. The only requirement is that at least four points can be identified on the robot's end-effector and can be tracked across frames.

1 Introduction

In this paper we use an active observer [1, 2] that *purposefully* controls its viewpoint in order to explore an unknown curved object, i.e., extract geometrical properties of the object's surface from the images obtained. Our focus is on the development of an efficient, robust, and provably-correct viewpoint control strategy that requires no calibration of the camera or the manipulator that controls viewpoint and can be used as a basic tool for performing a variety of exploration tasks for curved objects. The goal of the strategy we consider in this paper can be stated as follows: Given a point projecting to the object's occluding contour, control viewpoint on the tangent plane of the point. We show that tangential viewpoint control can be accurately performed in real time using a single camera to control a manipulator capable of changing an object's orientation with three degrees of freedom. Furthermore, we show that this task enables an invariant-based analysis, can be provably achieved without requiring information about the camera's calibration parameters or the camera's relationship to the robot. The only requirement is that at least four points can be identified on the robot's end-effector and tracked across frames.

Little work has been published on the use of simple and efficient viewpoint-control mechanisms for reconstruction, exploration or recognition tasks. However, the few recent

approaches have shown that the tight coupling of viewpoint control and visual processing has a significant impact on the visual tasks that can be performed and on the computations needed to perform them [3–7]. Our previous work has shown that vision-guided viewpoint control allows the use of weaker camera calibration assumptions, simplifies shape computations, and enables exploration of complex curved objects. Example tasks we have studied include occluding contour detection [8], local surface curvature estimation [9], affine-invariant surface reconstruction [10], as well as global model-building [11].

A key result of our previous research is that even though the exploration tasks we considered have different objectives, being both local and global, and even though they rely on different principles for formulating the viewpoint control process, the mechanisms for controlling viewpoint always operate in one of two simple ways: Either by moving on the *tangent plane* of the viewed surface at a point automatically selected during the viewpoint control process, or by moving on a normal plane at such a point. As a consequence, the camera's motion is always planar and *locally-controlled* even though the task itself may require exploration of an object's global geometry (e.g., for global model-building or surface inspection).

Motivated by these results, this paper focuses on the development of a system for real-time tangential viewpoint control. The approach is inspired by research on uncalibrated visual servoing (e.g., see [12, 13]) and invariance-based hand-eye coordination [14–16]. These approaches aim to minimize dependence on calibration while performing object tracking or manipulation tasks. A key characteristic of these approaches is that both the goal of the manipulation and tracking tasks as well as the strategy for achieving them can be formulated entirely in terms of a few stable features on the object being viewed. These approaches avoid sensitivity to camera and robot calibration errors, and are applicable even when the calibration parameters change (e.g., in an active vision system) or are hard to obtain (e.g., when the manipulator is mounted on a mobile robot exploring the contents of a room). Unfortunately, such approaches are not directly applicable to the exploration of curved objects with few or no markings or discontinuities because stable features that can be tracked across frames cannot generally be found on such objects.

In our framework, visual exploration of curved objects

The support of the National Science Foundation under Grant No. CDA-9503996, of the University of Maryland under Subcontract No. Z840902, and of Honeywell under Research Contract No. 304931455 are gratefully acknowledged.

is performed using the *occluding contour* [17] both to control viewpoint and to recover shape information. The occluding contour is the projection of the visible rim, the one-dimensional set of visible surface points at which the line of sight is tangent. The occluding contour becomes a rich source of visual information during the exploration of curved objects because its shape is highly constrained by the surface giving rise to it [17]. Furthermore, several approaches have been suggested for the detection of the occluding contour (e.g., [8, 18]). With the contour at the center of our approach, viewpoint control is formulated as an on-line, geometric planning process whose goal is to achieve a well-defined 3D affine relationship between the object, its occluding contour, the manipulator, and the camera.

Our approach combines the paradigms of purposive viewpoint control and uncalibrated visual servoing to achieve accurate and provably-correct motions on the tangent plane of selected points on an object’s surface. The approach is based on the following simple observation. Suppose the object is placed on a horizontal turntable. Then, tangential viewpoint control can be performed accurately for any horizontal tangent plane by simply rotating the turntable. To perform tangential viewpoint control on an arbitrary tangent plane we use a manipulator with at least three degrees of freedom and a strategy analogous to (1) re-orienting the object so that it becomes parallel to the turntable’s plane, and (2) controlling viewpoint on the desired tangent plane by rotating the turntable. We show that this object re-orientation process can be provably performed without imposing any Cartesian calibration requirements, and without requiring the camera’s parameters to remain fixed during the viewpoint control process. Accuracy is achieved by formulating the object re-orientation process within an affine-invariant framework; provable correctness is achieved by formulating robot control as a joint minimization problem of two affinely-computable functions whose joint minima are reachable using local minimization methods; insensitivity to camera parameter changes is achieved by using an adaptive control scheme [19, 20] that constantly updates the image Jacobian estimate.

The significance of our method lies in the use of purposive observer motion to achieve a well-defined geometric relationship with respect to a 3D shape prior to recognition. Its main contributions are: (1) Use of *tangential viewpoint control* as a basic strategy for controlling viewpoint during the exploration of unknown curved objects, (2) design and implementation of a system for real-time tangential viewpoint control, and (3) development of a provably-correct viewpoint control strategy that operates entirely in affine space. The system can also be used for viewpoint control on any plane whose intersection with an image is known.

2 Exploration By Tangential Viewpoint Control

When viewpoint is controlled in a continuous fashion, the occluding contour deforms. Tangential viewpoint control becomes a key motion in the exploration of an object because it can be used to dynamically structure the occluding

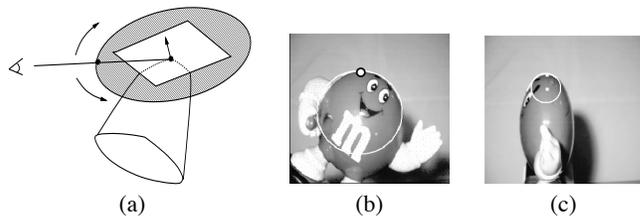


Figure 1: (a) Moving on the tangent plane, $T_p(S)$, of a point p . N is the surface normal at p . (b) Surface curvature for a point projecting to the occluding contour becomes trivial to compute from a special viewpoint which, for point q , corresponds to the side view of the toy [9]. The viewpoint is reached from (b) by rotating viewpoint on the tangent plane of the point projecting to q (which is horizontal and perpendicular to the plane of the page), *until* the curvature of the occluding contour at its uppermost point is maximized, as shown in (c).

contour’s shape and the visible rim’s motion in a way that factors out some of their viewpoint-dependent properties; what remains is *viewpoint-independent* information about the local and global shape of the object being explored.

In particular, tangential motion allows us to compensate for three viewpoint-dependent properties of the occluding contour and visible rim: (1) The one-dimensional set of points comprising the visible rim will change under any infinitesimal change in viewpoint, (2) the shape of the occluding contour (i.e., tangents and curvatures) may also change under such an infinitesimal viewpoint change, and (3) the visible rim’s motion and its connectivity depend on the local and global shape of the surface, the initial viewpoint, and the direction of the infinitesimal viewpoint change. Below we briefly outline how the dependency on viewpoint of each one of these properties can be factored out through tangential viewpoint control. See [8–11] for details.

Forcing visible rim point stationarity: Suppose p is a point on the visible rim at a given viewpoint. p will not remain on the visible rim if that viewpoint is perturbed in an arbitrary fashion. However, if we start changing viewpoint on p ’s tangent plane, p will remain on the visible rim for as long as it does not become occluded by points that are closer to the camera [21]. Consequently, even though in general the visible rim can be thought of as sliding over the surface when viewpoint changes continuously, we can force the visible rim to remain “stationary” at specific points by controlling viewpoint on their tangent plane. By appropriately choosing these points we can detect the occluding contour [8] and recover affine representations of unknown curved objects without the need for extrinsic camera calibration or viewpoint motion measurements [10].

Factoring out the dependence of the occluding contour’s shape on viewpoint: In general, when a surface point p is viewed from different viewpoints on its tangent plane, the local shape of the occluding contour (i.e., its curvature at p ’s projection) will be different. Nevertheless, the manner in which the occluding contour’s shape changes

around p 's projection as viewpoint changes on p 's tangent plane is completely determined by the local shape of the surface at p . This allows us to reach special viewpoints on p 's tangent plane at which the contour's shape depends only of the local shape of the surface [9] (Figures 1(b),(c)).

Factoring out dependence of the visible rim's motion on viewpoint: Under small viewpoint adjustments, the changes in the connectivity of the visible rim near a point p as well as the visible rim's motion near that point are completely determined by the relationship between the initial viewpoint and the endpoints of the arcs of viewpoints from which p is visible from its tangent plane. We show in [11] that we can force the visible rim's motion to depend only on the global shape of the surface by keeping the viewpoint's relationship to the endpoints of these arcs fixed (e.g., always at the middle of the arcs) before forcing the visible rim to slide over previously-unreconstructed points on the object's surface. This is sufficient to structure the visible rim's motion in a way that guarantees global reconstruction of arbitrarily-shaped generic objects [11].

3 Affine Point Representations

One of the key difficulties in controlling viewpoint relative to the tangent plane of a surface point p is that no image information about $T_p(S)$ is available when p is not on the visible rim. A basic step in our method for tangential viewpoint control involves an automatic "visual alignment" of the hand-eye system so that the viewpoint's motion can be accurately controlled in $T_p(S)$. This requires the construction of a representation for both the image plane of the camera and the tangent plane at point p . Affine point and camera representations are important because they can be constructed without knowing the viewpoint's motion, the camera's parameters and their relation to the hand-eye system, and without recovering any metric properties of the viewed surface. The basic principles behind this representation are briefly reviewed next. We assume the weak perspective projection model in the rest of the paper.

Let $p_1, \dots, p_n \in \mathbb{R}^3, n \geq 4$, be a collection of points at least four of which are not coplanar. An *affine representation* of those points is a representation that does not change if the same non-singular linear transformation (e.g., translation, rotation, scaling) is applied to all the points. Affine representations consist of three components: The *origin*, which is one of the points p_1, \dots, p_n ; the *affine basis points*, which are three points from the collection that are not coplanar with the origin; and the *affine coordinates* of the points p_1, \dots, p_n . We use the following two properties of affine point representations [22–25]:

Property 1 (Re-Projection Property) *When the projection of the origin and basis points is known along a viewing direction $\xi(t_m)$, we can compute the projection of a point p viewed under weak perspective projection from its affine coordinates using*

$$\begin{bmatrix} u_p^m \\ v_p^m \end{bmatrix} = \begin{bmatrix} u_{b_1}^m & u_{b_2}^m & u_{b_3}^m \\ v_{b_1}^m & v_{b_2}^m & v_{b_3}^m \end{bmatrix} \mathbf{A}_p + \begin{bmatrix} u_{p_o}^m \\ v_{p_o}^m \end{bmatrix} \quad (1)$$

where $[u_p^m \ v_p^m]$ is the projection of p , $[u_{p_o}^m \ v_{p_o}^m]$ is the projection of the origin, and \mathbf{A}_p is the vector of p 's affine coordinates.

Property 2 (Affine Reconstruction Property) *The affine coordinates of p_1, \dots, p_n can be computed using Eq. (1) when their projection along at least two viewing directions is known.*

Property 1 tells us that the projection process at any given viewpoint is completely determined by the 2×3 *projection matrix* Π collecting the projection of the affine basis points in Eq. (1). Property 2 shows that this process can be inverted if at least four non-coplanar 3D points can be tracked across frames as the viewpoint changes. This allows us to recover both the affine coordinates of the tracked points as well as the projection matrix. Eq. (1) leads directly to a generalization of the notions of the "viewing direction", "row," and "column" vectors:

Definition 1 (Affine Camera Parameters) *If $\Pi^{-1} = [\mathbf{R} \ \mathbf{C}]$ is the pseudoinverse of the projection matrix Π , (1) the vector \mathbf{R} is parallel to the image rows in the recovered 3D affine coordinate frame, (2) \mathbf{C} is parallel to the image columns, and (3) the vector Ξ representing the viewing direction is the vector product $\mathbf{R} \wedge \mathbf{C}$.*

Given these definitions, the image plane is spanned by the vectors \mathbf{R} and \mathbf{C} , and all points along Ξ will project to the same camera pixel. Once the affine camera's row and column vectors are determined, we can derive a representation for any plane intersecting the image along a known vector. Specifically, if $[u \ v]^T$ is the tangent of the occluding contour at the projection of a visible rim point $p, T_p(S)$ will be the plane spanned by vectors Ξ and $(u\mathbf{R} + v\mathbf{C})$. This observation extends Barrow and Tannenbaum's expression for $T_p(S)$ [26] to the case of an affine camera.

4 Uncalibrated Tangential Viewpoint Control

Tangential viewpoint control requires viewpoint to be controlled in a very constrained way, i.e., by moving on a plane corresponding to the tangent plane at a selected visible rim point. Below we describe one approach for implementing such constrained motions that (1) exploits the natural constraints imposed on the motion of orienting devices such as PUMAs and pan-tilt units, (2) changes viewpoint by controlling the orientation of the object, and (3) does not require calibration of the camera or the hand-eye system.

Suppose that we have placed an object on a horizontal turntable and that both the viewing direction and the image rows are horizontal. Furthermore, suppose that the image rows are tangent to the object's occluding contour at the projection of a point p (Figure 2(a)). Then, because the viewing direction and the image rows define a horizontal plane, any rotation of the turntable will force the viewing direction to move strictly parallel to p 's tangent plane. In order to enforce planar viewpoint control on the tangent plane of an arbitrary visible rim point we develop a strategy that enforces this special geometry between the object, the turntable, and the camera.

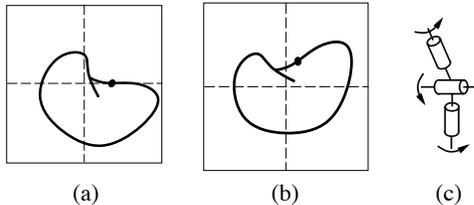


Figure 2: (a) Geometry of viewpoint control within $T_p(S)$. (b) Viewpoint control on $T_p(S)$ for an arbitrary visible rim point p is achieved by re-orienting the object to obtain a view identical to the one in (a) up to a translation in the image plane. (c) Object re-orientation can be achieved using an orienting device with three consecutive rotational joints (RRR unit). Joints z_1, z_2 correspond to a ‘‘pan’’ axis and a ‘‘tilt’’ axis, respectively.

In theory, object re-orientation amounts to a rotation about an axis parallel to the viewing direction. In practice, object re-orientation is limited by the configuration and the degrees of freedom of the manipulator used to re-orient the object. The questions we answer are: Given the projection of $T_p(S)$, (1) how can we re-orient the object so that $T_p(S)$ becomes horizontal, and (2) how can we ensure that the viewpoint remains unaffected by this process? Object re-orientation is achieved with a rotational stage mounted on pan-tilt unit or with the last three rotational joints of a PUMA robot arm (Figure 2(c)). In the following we assume that joint axis z_1 is perpendicular to the image rows.¹

4.1 Object Re-Orientation with an Affine Camera

Let p be a point on the visible rim. The objective of the object re-orientation strategy can be stated as follows:

Object re-orientation goal: Let $[u \ v]^T$ be the tangent of the occluding contour at p ’s projection, and let $\Xi(0)$ be the observer’s viewing direction. Control viewpoint according to $\Xi(t), t \geq 0$, until the plane spanned by $\Xi(t)$ and an image row at the current viewpoint becomes parallel to the plane of $\Xi(0)$ and $[u \ v]^T$.

We use the following theorem which gives us a way to perform this viewpoint control process using a simple minimization method. In particular, suppose $\Xi(0), \mathbf{R}(0), \mathbf{C}(0)$ are the viewing direction, row, and column vectors of the affine camera at the initial viewpoint, respectively, and let $\Xi(t), \mathbf{R}(t), \mathbf{C}(t)$ be the corresponding vectors at time t . Furthermore, let $\mathbf{Q} = u\mathbf{R}(0) + v\mathbf{C}(0)$ be the 3D affine vector corresponding to the tangent to the occluding contour at p ’s projection along $\Xi(0)$. Theorem 1 gives us a criterion that is based entirely on affinely-computable quantities for deciding whether re-orientation has been achieved:

¹This assumption is not crucial, however. Alignment of axis z_1 with an image row can be performed using a strategy identical to the one described in this paper. We are currently implementing such a strategy that assigns axes z_1, z_2, z_3 to axes 4-6, respectively, of a PUMA robot and aligns axis 4 with an image column by visually controlling axes 1 and 3 of the PUMA.

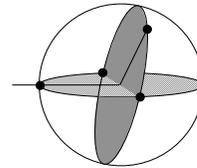


Figure 3: Representing the minima of functions C_1, C_2 over the sphere of orientations of $T_p(S)$ relative to the camera. The two joint minima of C_1, C_2 occur at the intersection of orientations where $T_p(S)$ contains $\mathbf{R}(t)$ and orientations where $T_p(S)$ contains $\Xi(t)$.

Theorem 1 (Object Re-Orientation Theorem) If

$[\kappa_i \ \lambda_i \ \mu_i \ \nu_i], i = 1, 2$, is a non-zero solution of the system

$$[\kappa_i \ \lambda_i \ \mu_i \ \nu_i] \cdot \begin{bmatrix} \mathbf{X}_i^T(t) \\ \mathbf{C}_i^T(t) \\ \Xi^T(0) \\ \mathbf{Q}^T \end{bmatrix} = 0 \quad (2)$$

with $X_1 = \mathbf{R}(t)$ and $X_2 = \Xi(t)$, respectively, the functions

$$C_i(t) = 1 - \frac{\kappa_i^2}{\kappa_i^2 + \lambda_i^2}, i = 1, 2 \quad (3)$$

are both zero at t if and only if (1) p is on the rim along $\Xi(t)$, and (2) the projection of $T_p(S)$ coincides with an image row.

Intuitively, the functions C_1, C_2 of Theorem 1 formalize two conditions that must be simultaneously satisfied when $T_p(S)$ projects to a row in the image: (1) The intersection of $T_p(S)$ with the image along $\Xi(t)$ is parallel to the image rows, and (2) the viewing direction, $\Xi(t)$, is contained in $T_p(S)$. $C_1(t)$ and $C_2(t)$ are jointly minimized if and only if both $\mathbf{R}(t)$ and $\Xi(t)$ belong to the plane spanned by $\Xi(0)$ and \mathbf{Q} , i.e., $T_p(S)$.

While Theorem 1 suggests that a simultaneous minimization of both $C_1(t)$ and $C_2(t)$ leads to the correct viewpoint it does not show *how* such a joint minimization can be achieved. The following lemma shows that it is possible to reach a joint minimum of the two functions (Figure 3):

Lemma 1 Let z, z' be the joint coordinates of axes z_2 and z_3 of the RRR unit, respectively. The function $C_1^2(z, z') + C_2^2(z, z')$ is locally minimized at (z_0, z'_0) if and only if C_1 and C_2 are both zero at (z_0, z'_0) .

Lemma 1 shows that since the only minima of the functions C_1, C_2 are global, a local minimization approach is sufficient to reach one of the joint zeros of C_1, C_2 .

Object Re-Orientation Strategy

Step 1: Let q_1, \dots, q_n be the projections of a collection of fixed 3D points, p_1, \dots, p_n , which can be tracked across frames and are rigidly attached to the viewed object.

Step 2: Establishing an affine reference frame:

- a. Rotate axis z_1 while tracking the projection of p_1, \dots, p_n .
- b. Construct the affine basis vectors, compute the projection matrix as well as the affine coordinates of p_1, \dots, p_n using Eq. 1 [24, 25].

Step 3: Initiating an object re-orientation so that the tangent plane defined by the contour tangent $[u \ v]^T$ along $\Xi(\mathbf{t}_0)$ becomes aligned with an image row:

- a. Let $q_1(t_0), \dots, q_n(t_0)$ be the projections of the tracked points along $\Xi(\mathbf{t}_0)$.
- b. Compute the projection matrix, $\Pi(t_0)$, corresponding to $\Xi(\mathbf{t}_0)$ from $q_1(t_0), \dots, q_n(t_0)$ and their affine coordinates, computed in Step 2.
- c. Let $\mathbf{Q} = u\mathbf{R}(t_0) + v\mathbf{C}(t_0)$.

Step 4: Compute the projection matrix corresponding to $\Xi(\mathbf{t})$.

Step 5: Compute the functions $C_1(t), C_2(t)$ along $\Xi(\mathbf{t})$ from $\mathbf{Q}, \Xi(\mathbf{t}_0), \mathbf{R}(\mathbf{t}_0), \mathbf{C}(\mathbf{t}_0), \Xi(\mathbf{t}), \mathbf{R}(\mathbf{t})$.

Step 6: Control manipulator axes z_2, z_3 in order to decrease the values of C_1, C_2 .

Step 7: Repeat Steps 4-6 until C_1, C_2 are both zero.

Step 8: Rotate axis z_1 until the current projections of p_1, \dots, p_n are identical to those at the initial viewpoint upto a rotation and/or translation in the image plane.

Execution of Step 5 of the Object Re-Orientation Strategy requires controlling a robot manipulator's joints in order to decrease the values of C_1, C_2 . In order to perform this step we use an adaptive controller which continuously updates a local estimate of the *image Jacobian*, relating changes in the control vector $[C_1 \ C_2]^T$ to changes in the manipulator's configuration. Together with Lemma 1, the use of an adaptive controller ensures that convergence of the system can be achieved without any prior estimates of the image Jacobian which, in general, depends on the affine basis points, the configuration of the robot, as well as the relative position of the camera and the manipulator's joint axes. See [20] for details on the controller.

5 Experimental Results

In order to demonstrate the effectiveness of the Object Re-Orientation Strategy we implemented the strategy on our real-time visual servoing system. The system consists of a PUMA760 robot, an 8 processor SUN SPARCserver 2000 multiprocessing system used for vision processing and for running the adaptive controller, and a SPARCserver 330 running RCCL for low-level robot control. A Sony TR-700 camcorder with controllable zoom and at a position approximately 3m from the end-effector provided visual input to the system. A Matlab implementation of the adaptive control loop runs at about 1Hz.

In our experiments, we manually aligned the camera's columns with Joint Axis 4 of the PUMA robot and used Joint Axes 5 and 6 to evaluate our object re-orientation strategy for tangential viewpoint control. No other camera or hand-eye calibration information was used. To evaluate

our viewpoint control strategy we performed three experiments, briefly outlined below. In all of these experiments we simplified point tracking by mounting 5 lights on the robot's end-effector and using a distributed correlation-based point tracker for tracking them at a rate of 30Hz.

Validating the predicted variation of the control vector. In order to determine whether the variation in the affine functions C_1, C_2 defined in Section 4.1 agree with the predictions of Theorem 1 and Lemma 1, we computed these functions as follows. First Joint Axis 4 of the PUMA was rotated and the 5 feature points on its end-effector were tracked to establish an affine frame of reference. The axis was then rotated back to its original position, where Joint Axis 6 was along the camera's optical ray. The values of C_1 and C_2 were subsequently computed for Joint 5 and 6 values between ± 20 degrees of the initial position, and for different orientations of the vector $[u \ v]^T$ defining the tangent plane in the initial image.

Figure 4 shows that the measured values closely match the predicted variation of C_1, C_2 (Figure 3). Furthermore, Figures 4(a),(b) which correspond to the case where the optimal joint configuration is the initial position, indicate that the initial position lies in the intersection of two zero-sets of C_1, C_2 , exactly as predicted by Theorem 1. In our experiments, we have also found that when the tangent plane orientation at the initial position varies from horizontal, the effect is a "shift" in the location of the valleys of C_1, C_2 .

Validation of the object re-orientation strategy. Figures 4(c)-(f) show results from one execution of the Object Re-Orientation Strategy. We mounted a doll on the PUMA's end-effector and applied the strategy so that the tangent plane at the top of the doll's head in Figure 4(c) becomes aligned with the image rows. Its intersection with the initial image is a diagonal line, indicated by the projection of the white rod that is mounted to the top of the doll's head to serve as a reference line. Figure 4(d) shows the final position of the doll. The projection of the rod is aligned with the image rows and the initial and final images are related by a rotation in the image, as required by the object re-orientation goal (Figure 2(a)).

Figures 4(c),(d) show the path in joint space generated by our viewpoint control strategy, projected onto the contour lines of functions C_1 and C_2 . Zeros of the two functions correspond to the middle "valley" between the contour lines, which is diagonal for C_1 and vertical for C_2 . The paths show the behavior of our adaptive control algorithm in relation to the variation of C_1 and C_2 : The robot first controls the joint axes by following a path that leads it to a zero of C_1 and then adjusts its position in the direction of the zero-set of C_1 so that function C_2 can also be minimized. Note that neither the structure of the zero-sets of C_1 and C_2 nor their positions were given as input to the system, but were "discovered" by the adaptive controller.

Accuracy of object re-orientation. In order to determine the accuracy of our Object Re-Orientation Strategy we mounted an object on the PUMA's end-effector and rotated Joint Axes 5 and 6 so that the tangent plane of a specific point on the object's occluding contour was aligned with

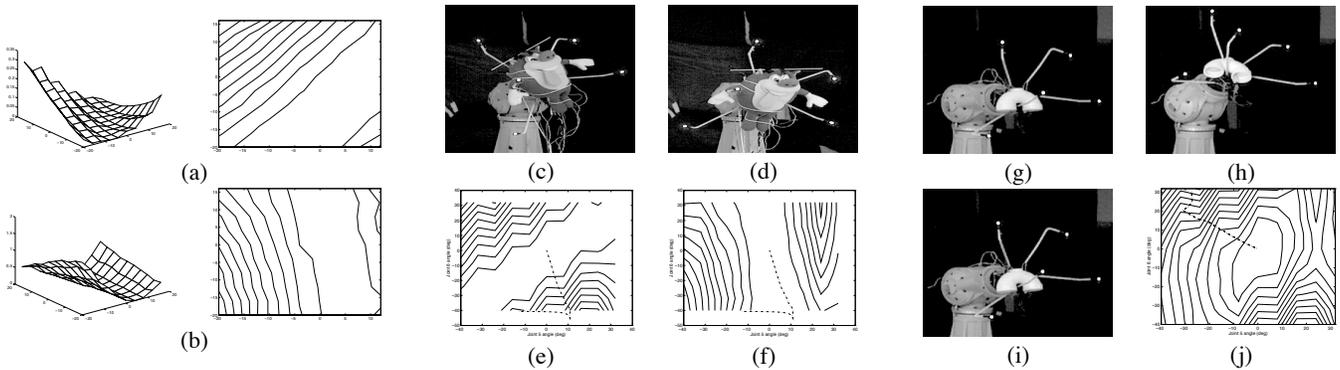


Figure 4: Experimental results. (a),(b) Variation of C_1 , C_2 , respectively, with the PUMA's Joint 5 and 6 angles. $T_p(S)$ is along an image row at the robot's initial position. Joint angles vary between -20 and $+20$ degrees from that position. Contour plots of C_1 , C_2 are shown to the right. (c)-(f) A run of the viewpoint control process. (g)-(j) Determining re-orientation accuracy.

the image rows (Figure 4(g)). We then recorded this position, moved both joints to a random position approximately $\pm\pi/4$ radians from the original position (Figure 4(h)), and executed the Object Re-Orientation Strategy to force the selected point's tangent plane to become horizontal again (Figure 4(i)). The joint trajectory taken to reach the goal position for one of the runs is shown in Figure 4(j), overlaid with the contour plot of $\sqrt{C_1^2 + C_2^2}$. In the approximately 20 experimental runs we performed the accuracy of the re-orientation process was between 2 and 7 degrees measured in joint space coordinates. These errors are partly due to the limited accuracy of point tracking, since the tracked lights are considerably larger than a single pixel, and partly due to the near-zero values of C_1 , C_2 close to the goal position.

6 Concluding Remarks

Two limitations of the current implementation are (1) its fairly primitive point tracking system, and (2) the inability to handle cases where the robot manipulator reaches a joint limit. We are planning to incorporate into our system recent methods for affine-based point tracking [27], as well as information about the structure and topology of the space of control vectors in order to improve re-orientation accuracy and avoid the manipulator's joint limits.

References

- [1] Y. Aloimonos, "Purposeful and qualitative active vision," in *Proc. Int. Conf. on Pattern Recognition*, pp. 346–360, 1990.
- [2] D. H. Ballard, "Animate vision," *AIJ*, v. 48, pp. 57–86, 1991.
- [3] D. Wilkes and J. K. Tsotsos, "Active object recognition," in *Proc. CVPR*, pp. 136–141, 1992.
- [4] C. E. Smith and N. P. Papanikolopoulos, "Computation of shape through controlled active exploration," in *Proc. Robotics Automat. Conf.*, pp. 2516–2521, 1994.
- [5] C. B. Madsen and H. I. Christensen, "Localizing un-calibrated, re-active camera motion in an object centered coordinate system," in *Proc. Work. on Visual Behaviors*, pp. 119–123, 1994.
- [6] M. J. Taylor and A. Blake, "Grasping the apparent contour," in *Proc. 3rd ECCV*, pp. 25–34, 1994.
- [7] P. Whaithe and F. P. Ferrie, "Autonomous exploration: Driven by uncertainty," in *Proc. CVPR*, pp. 339–346, 1994.
- [8] K. N. Kutulakos and C. R. Dyer, "Occluding contour detection using affine invariants and purposive viewpoint control," in *Proc. CVPR*, pp. 323–330, 1994.
- [9] K. N. Kutulakos and C. R. Dyer, "Recovering shape by purposive viewpoint adjustment," *IJCV*, v. 12, n. 2, pp. 113–136, 1994.
- [10] K. N. Kutulakos, "Affine surface reconstruction by purposive viewpoint control," in *Proc. 5th ICCV*, 1995.
- [11] K. N. Kutulakos and C. R. Dyer, "Global surface reconstruction by purposive control of observer motion," in *Proc. CVPR*, 1994. pp. 331–338.
- [12] K. Hashimoto, *Visual Servoing*. World Scientific, 1993.
- [13] G. Hager and S. Hutchinson, eds., *ICRA Work. on Visual Servoing: Achievements, Applications, and Open Problems*, 1994.
- [14] N. Hollinghurst and R. Cipolla, "Uncalibrated stereo hand-eye coordination," in *Proc. BMVC*, 1993.
- [15] G. D. Hager, "Calibration-free visual control using projective invariance," in *Proc. 5th ICCV*, 1995.
- [16] P. A. Beardsley, A. Zisserman, and D. W. Murray, "Navigation using affine structure from motion," in *Proc. 3rd ECCV*, pp. 85–96, 1994.
- [17] J. J. Koenderink, *Solid Shape*. MIT Press, 1990.
- [18] R. Vaillant and O. D. Faugeras, "Using extremal boundaries for 3-d object modeling," *T-PAMI*, v. 14, pp. 157–173, 1992.
- [19] K. Hosoda and M. Asada, "Versatile visual servoing without knowledge of true jacobian," in *Proc. IROS Conf.*, 1994.
- [20] M. Jagersand and R. C. Nelson, "Adaptive differential visual feedback for uncalibrated hand-eye coordination and motor control," Tech. Rep. 579, University of Rochester, 1994.
- [21] J. H. Rieger, "Three-dimensional motion from fixed points of a deforming profile curve," *Optics Letters*, v. 11, n. 3, pp. 123–125, 1986.
- [22] J. L. Mundy and A. Zisserman, eds., *Geometric Invariance in Computer Vision*. MIT Press, 1992.
- [23] J. J. Koenderink and A. J. van Doorn, "Affine structure from motion," *J. Opt. Soc. Am.*, vol. A, no. 2, pp. 377–385, 1991.
- [24] D. Weinshall and C. Tomasi, "Linear and incremental acquisition of invariant shape models from image sequences," in *Proc. 4th ICCV*, pp. 675–682, 1993.
- [25] S. M. Seitz and C. R. Dyer, "Complete scene structure from four point correspondences," in *Proc. 5th ICCV*, 1995.
- [26] H. G. Barrow and J. M. Tenenbaum, "Interpreting line drawings as three-dimensional images," *AIJ*, v. 17, pp. 75–116, 1981.
- [27] I. D. Reid and D. W. Murray, "Tracking foveated corner clusters using affine structure," in *Proc. 4th ICCV*, 1993. pp. 76–83.