Robust Visual Servoing based on Relative Orientation

Camillo J. Taylor James P. Ostrowski Sang-Hack Jung
General Robotics and Active Sensory Perception (GRASP) Laboratory
University of Pennsylvania, 3401 Walnut Street, Philadelphia, PA 19104-6228
E-mail: {cjtaylor, jpo, sangj}@grip.cis.upenn.edu

Abstract

In this paper the problem of controlling the spatial position and orientation of a robotic platform based on the image data obtained from a video camera mounted on that platform is considered. More specifically, we propose control laws that generate translational and angular velocities that will cause the robot to achieve and maintain a fixed position and orientation with respect to a set of feature points in the scene.

The proposed control schemes make use of well established techniques for computing estimates for the relative orientation of two camera positions from a set of feature correspondences. An important advantage of these control schemes is that it is possible to demonstrate analytically that they are globally convergent even in the presence of large calibration errors in both the intrinsic parameters of the camera and in the extrinsic parameters which relate the frame of reference of the camera to the body frame of the robot platform which is being controlled. Furthermore no a'priori knowledge about the structure of the scene is assumed.

1 Introduction

The use of camera-based techniques to control robotic systems has seen a significant rise in popularity recently. This field, known as *visual servoing*, has been aided by the faster speeds and lower costs of modern microprocessors, coupled with the general availability of high quality cameras. The wealth of information available in image data and the flexibility of the sensor combine to make it an attractive option for a control input if the right algorithms can be developed.

Papanikolopoulos, Nelson, and Khosla [17] have developed a visual tracking system that can track full 3-D motions, by utilizing a Jacobian-based adaptive controller that estimates the depth parameters online in order to track objects. More recent work by Papanikolopoulos has extended these results using optical flow to estimate depth parameters for use in uncalibrated environments [16]. Espiau, Rives, et al. [4, 18, 19] proposed approaches that also utilize the image Jacobian, by developing what they call the interaction screw to describe the relationship between robot motion and image feature motion. They have implemented these systems, for example, in controlling mobile robots performing servoing based on landmarks of known geometries. These techniques, however, require estimation of the depth of features in the scene.

Other researchers have used strictly pose-based control laws, which require explicit calculation of the robot pose with respect to some known coordinate frame. For example, Wilson et al. [21] generate pose estimates using an extended Kalman filter, and use these estimates to drive the robot based on error measurements in the pose space.

Hager, Chang and Morse [7] Hollinghurst and Cipolla [11] and Horaud, Dornaika and Espiau [12] have all investigated the issues involved in controlling robotic manipulators based on the image data acquired with uncalibrated or coarsely calibrated stereo systems. This work is important in the current context, since it established that image-based techniques can have a proven insensitivity to calibration parameters [8].

In this paper the problem of controlling the position and orientation of a robotic platform based on the image data obtained from a single video camera mounted on that platform is considered. More specifically, we propose control laws that generate translational and angular velocities that will cause the robot to achieve and maintain a fixed position and orientation with respect to a set of features in the scene. We make the assumption that the features that are tracked in the image are derived from rigid objects, such as walls, tables, and other fixtures. Thus, they can move with respect to the robot, but not with respect to each other.

We then formulate the problem as one of maintaining a fixed pose relative to a given set of image features, such as points and lines, as viewed by the onboard camera.

The techniques described in this paper would be applicable to a wide range of positioning tasks. They could be employed on mobile platforms such as blimps, helicopters, underwater vehicles, space based robots or terrestrial vehicles to guide the robot to a desired position and orientation with respect to one or more target objects. The schemes could also be used for so-called "eye in hand" servoing applications to position the end effector of a robot arm with respect to a work piece based on the image data obtained from a camera mounted on the robot's gripper.

The proposed control schemes make use of well established techniques for computing estimates for the relative orientation of two camera positions from a set of feature correspondences. As in other visual servoing schemes, the goal position of the platform is actually specified indirectly in terms of the image measurements that would be obtained if the robot were in the desired pose. The system generates appropriate angular and translational velocity commands by estimating the relative orientation between the cameras current position and the target pose from the current and desired image measurements.

The use of relative orientation for pose control has also been proposed by Basri, Rivlin and Shimshoni [1] Malis, Chaumette and Boudet [14, 15], Soatto and Perona [20] and Deguchi [2]. The main contributions of this paper are a theoretical analysis and experimental results which prove that these schemes are convergent even in the presence of large calibration errors in *both* the intrinsic and extrinsic parameters of the camera system.

Note that the control laws presented here differ from previous work in that they do not require the estimation of an image Jacobian or interaction screw, used to relate the robot's body velocity to changes in image measurements. Since this image Jacobian depends upon the positions of the features with respect to the camera, control schemes which rely on this approach must either estimate the Jacobian online or assume that this matrix is nominally fixed at some known value. No such estimates are required in this method.

In Section 2 the relative orientation problem is briefly discussed. Section 3 presents an analysis of the pose control problem, proposes control schemes for solving this task and provides conditions under which these strategies are known to be globally convergent. The experimental results that have been obtained with these methods are presented in Section 4 while Section 5 presents some of the conclusions that have been drawn from this research and proposes topics for further investigation.

2 Relative Orientation

The problem of computing the relative position of two cameras from a set of point correspondences in two images has been well studied in both the computer vision and photogrammetry literatures. Most of the approaches to solving this problem proceed by exploiting the epipolar constraint, which relates the position of the projection of a point feature in one image to its projection in the second image.

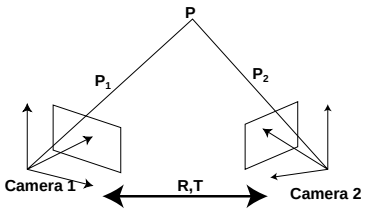

Figure 1: The basic geometry of the relative orientation problem

The basic imaging geometry of the relative orientation problem is shown in Figure 1. Let $\mathbf{P}_1$ denote the coordinates of a point $\mathbf{P}$ in the scene with respect to a coordinate frame of reference centered at the focus of projection of the first camera. The coordinates of that same point $\mathbf{P}$ with respect to a coordinate frame associated with the second camera are given by the expression $\mathbf{P}_2 = R^T(\mathbf{P}_1 - \mathbf{T})$ where R and $\mathbf{T}$ denote the relative orientation and position of the second frame with respect to the first.

Assuming a perspective projection model, the relationship between the coordinates of a point with respect to the camera frame $\mathbf{P} = (X, Y, Z)^T$ and the projective coordinates of the image of that point on the focal plane, $\mathbf{p} = (u, v, 1)^T$ can be written as follows:

$$\mathbf{p} \propto A\mathbf{P}, \tag{1}$$

where A is an invertible upper triangular matrix that represents the *intrinsic parameters* of the camera.

In our application, camera frames 1 and 2 refer to two different positions of the same camera and we will assume that the intrinsic parameters of the camera are fixed, hence A is a constant matrix. This means

that the projective coordinates of the projection of the point into the two images, $\mathbf{p}_1$ and $\mathbf{p}_2$, are given by the following equations:

$$\mathbf{p}_1 \propto A\mathbf{P}_1 \quad (2)$$
$$\mathbf{p}_2 \propto A\mathbf{P}_2 = AR^T(\mathbf{P}_1 - \mathbf{T}) \quad (3)$$

From these two equations one can derive the *epipolar constraint* which is given by

$$\mathbf{p}_1^T A^{-T} J(\mathbf{T}) R A^{-1} \mathbf{p}_2 = 0, \quad (4)$$

where

$$J(\mathbf{T}) = \begin{pmatrix} 0 & -T_z & T_y \\ T_z & 0 & -T_x \\ -T_y & T_x & 0 \end{pmatrix}. \quad (5)$$

The matrix $F = A^{-T} J(\mathbf{T}) R A^{-1}$ is termed the *fundamental matrix* [5]. A number of algorithms have been proposed for estimating this fundamental matrix from a set of point correspondences [13, 9, 10].

If the intrinsic parameter matrix of the camera A is known it is possible to compute the so-called *essential matrix*, E, from the fundamental matrix, F, as follows:

$$E = A^T F A \propto J(\mathbf{T}) R \quad (6)$$

This essential matrix can be decomposed into its constituent parts to obtain estimates for the rotation and translation that link the two camera positions with the understanding that the translation vector $\mathbf{T}$ can only be recovered up to a scale factor [5, 13].

When the displacement between the two camera positions is zero the fundamental matrix degenerates to zero. However, in this case the relationship between image coordinates in the first image and their correspondents in the second image takes the form of a collineation. That is $\mathbf{p}_2 \propto H_\infty \mathbf{p}_1$ where $H_\infty = ARA^{-1}$. Faugeras refers to the matrix H_∞ as the *collineation of the plane at infinity* [6].

This suggests a two stage process for determining the relative orientation of the two camera positions. In the first stage the algorithm tests whether the points in the two images are related by a collineation. This is done by selecting a subset of four image points in the goal image whose coordinates form a projective basis of the real projective plane RP^2 and computing the homography that relates these points to their correspondents in the current image. It is then possible to test how well this collineation predicts the location of the other image features by evaluating the function

$$\lambda(R,T) = \sum_i \left\| \frac{\mathbf{p}_2^i}{\|\mathbf{p}_2^i\|} - \frac{H\mathbf{p}_1^i}{\|H\mathbf{p}_1^i\|} \right\|^2, \quad (7)$$

which measures the disparity between the normalized coordinates of corresponding points. If this residual is sufficiently small then this indicates that the translational displacement between the two cameras is small, and the algorithm should return the computed collineation matrix, H_∞. Otherwise the fundamental matrix must be estimated using any of the previously mentioned techniques.

3 Pose Control

3.1 Controlling Pose

The objective of the proposed visual servoing schemes is to drive the disparity between the robot's current position and desired configuration to zero. As shown in Figure 1 this disparity is characterized by two components, the rotational disparity, $R \in SO(3)$, and the translational disparity, $\mathbf{T} \in \Re^3$. Similarly, the *body velocity* of the platform, that is, its instantaneous velocity with respect to its current pose, can be characterized in terms of its angular velocity $\omega \in \Re^3$ and its translational velocity $\mathbf{v} \in \Re^3$.

For the purposes of this discussion, we will consider the problems of regulating the translational and rotational displacements to zero separately. Let us first set $\omega = 0$ and consider control laws of the form $\mathbf{v} = \alpha(R, \mathbf{T}) L \mathbf{T}$, where $\alpha(R, \mathbf{T})$ is a scalar-valued function of pose, and L is a constant, positive definite matrix which is used to scale the control law.

Next, define a Lyapunov function $\mathcal{L}_T(\mathbf{T}) = \|\mathbf{T}\|^2$, which we can use to test for asymptotic convergence to $\mathbf{T} = 0$ (this is the goal state, since our control is targeted towards zeroing the relative pose). Differentiating this Lyapunov function yields:

$$\frac{d}{dt}\mathcal{L}_T(\mathbf{T}) = -2\mathbf{T}^T\mathbf{v} = -2\alpha(R, \mathbf{T})\mathbf{T}^T L \mathbf{T} \quad (8)$$

From this expression we conclude that for L positive definite, if $\alpha(R,T) > 0 \ \forall \mathbf{T} \neq 0$ then $\mathcal{L}_T(\mathbf{T})$ will decrease monotonically over time and, hence, the translational error will be regulated to zero.

The rotational displacement of the platform can be handled in a similar manner. That is, let $\mathbf{v} = 0$ and consider control laws of the form $\omega = \beta(R) M \mu$ where $\beta(R)$ is a scalar-valued function, M is a constant matrix, and μ is computed from $J(\mu) = R - R^T$. We can construct a Lyapunov function of the form $\mathcal{L}_R(R) = \mathrm{tr}((R-I)^T(R-I)) = 2(3 - \mathrm{tr}(R))$ which gives a measure of the magnitude of the rotational displacement. Differentiating this Lyapunov function

yields:

$$\begin{aligned}\frac{d}{dt}\mathcal{L}_R(R) &= -2\mathrm{tr}(\dot{R}) \\ &= -2\mathrm{tr}(J(\omega)^T R) \\ &= -\mu^T\omega = -\beta(R)\mu^T M\mu \qquad (9)\end{aligned}$$

From which we conclude that this control law will regulate the rotational error to zero if the matrix M is positive definite and $\beta(R) > 0\ \forall R \neq I$.

To regulate the error in both orientation and translation to zero, the control laws described above could be applied sequentially. That is, the orientation control law would be invoked to reduce the rotational displacement to an acceptable error and then the translation control law would be employed to drive the platform to the desired position. In practice , however, we have found that both of these control laws can be applied simultaneously.

3.2 Control Schemes

In the most straightforward case, the intrinsic parameters of the camera are known and the body frame of the robot platform and the camera frame are coincident or can be made so by a suitable change of coordinates. In this situation, the essential matrix can be extracted from the fundamental matrix and estimates for R and $\mathbf{T}/\|\mathbf{T}\|$ can be computed as described in Section 2.

The estimate for the rotation matrix can be used to regulate the orientation error to zero simply by applying the control law $\omega = \mu$ where $J(\mu) = R - R^T$. This control law trivially satisfies the conditions for convergence set forth in Section 3.1.

Similarly the control law $\mathbf{v} = (\lambda(R,\mathbf{T})/\|\mathbf{T}\|)\mathbf{T}$ with $\lambda(R,\mathbf{T})$ given in Equation 7 can be employed to regulate the translational displacement to zero since the scalar valued function $\lambda(R,\mathbf{T})/\|\mathbf{T}\|$ is guaranteed to be greater than zero for all configurations where $\mathbf{T} > 0$ as long as a sufficient number of feature points in general position are used.

3.3 Control in the presence of extrinsic calibration error

In practise the relationship between the camera frame and the robot's body frame can be difficult to measure accurately. In this case there will be some discrepancy between the frame in which the body velocities, ω and $\mathbf{v}$, are applied and the camera frame. For the purposes of this discussion, we will assume a worst-case scenario, where the displacement between the robot and camera is completely unknown. Let these unknown rotational and translational displacements be denoted by $\tilde{R}$ and $\tilde{\mathbf{T}}$, respectively, which we will call the *extrinsic calibration errors*.

In this case there will be a discrepancy between the commanded body velocity of the camera frame and the actual velocity, since these two velocities are related to each other via the adjoint transformation, which depends on R and $\mathbf{T}$. More specifically, in the case of a purely translational motion where the velocity of the robot frame is $\mathbf{v}$, the body velocity of the camera frame will be $\tilde{R}\mathbf{v}$. Similarly if the angular velocity of the robot frame is ω, then the angular velocity experienced by the camera frame will be $\tilde{R}\omega$.

The equations for the control law must be amended to account for this discrepancy. In the translational case the effective velocity of the camera frame will be given by:

$$\mathbf{v} = (\lambda(R,\mathbf{T})/\|\mathbf{T}\|)\tilde{R}\mathbf{T} \qquad (10)$$

From the analysis in Section 3.1 we can conclude that this control law will still regulate the translational error to zero if the matrix $\tilde{R}$ is positive definite. This will be true if $\tilde{R}$ represents a rotation of less than $\pi/2$ degrees. Notice that this is quite a generous margin of error.

Similarly, when the orientation control law is amended to account for extrinsic calibration error, we obtain an expression of the form:

$$\omega = \tilde{R}\mu \qquad (11)$$

We can again conclude that the orientation error will be regulated to zero as long as $\tilde{R}$ is positive definite. It should also be noted that while we have shown these control laws to be valid even under very large extrinsic calibration errors, in practice it is expected that some information will be known about the calibration, so that the relative calibration error will be quite small.

3.4 Control in the Presence of Intrinsic Parameter Errors

In the previous control laws it was assumed that the intrinsic parameters of the camera were known perfectly. In practise the estimates for these parameters will contain errors. The relationship between the estimate for the intrinsic parameters, A_{est}, and the actual values, A, can be expressed as follows:

$$A = A_{est}\tilde{A} \qquad (12)$$

Where $\tilde{A}$ is an unknown invertible upper triangular matrix.

In this case, we can recover an approximation to the essential matrix, which we denote by $\tilde{E}$, from the fundamental matrix, F:

$$\begin{aligned} \tilde{E} &= A_{est}^T F A_{est} && (13) \\ &= A_{est}^T (A_{est}\tilde{A})^{-T} J(\mathbf{T}) R (A_{est}\tilde{A})^{-1} A_{est} \\ &= \tilde{A}^{-T} J(\mathbf{T}) R \tilde{A}^{-1}. && (14) \end{aligned}$$

From this matrix it is possible to recover an estimate for $\tilde{A}\mathbf{T}$ by observing that

$$(\tilde{A}^{-T} J(\mathbf{T}) R \tilde{A}^{-1})^T (\tilde{A}\mathbf{T}) = 0. \quad (15)$$

If this estimate is used to to generate translational velocity control signals the effective control law in the presence of extrinsic calibration errors will take the form:

$$\mathbf{v} = (\lambda(R, \mathbf{T})/\|\tilde{A}\mathbf{T}\|)\tilde{R}\tilde{A}\mathbf{T}. \quad (16)$$

By the analysis in Section 3.1, we can conclude that this control scheme will be globally convergent if the matrix $\tilde{R}\tilde{A}$ is positive definite. This result implies that it is possible to construct control laws which servo the robotic platform to a desired position even in the presence of large errors in both the intrinsic and extrinsic parameters.

A similar argument for robustness in the face of intrinsic calibration errors can be made for orientation control when the matrix H_∞ is available. As was noted earlier this matrix cannot, in general, be recovered solely from point correspondences. It can, however, be estimated if the translational displacement between the two camera frames is known to be zero or if additional information such as correspondences between vanishing points in the two images can be recovered.

When this matrix *is* available we can recover an estimate for the rotational displacement between the two camera frames by noting that the H_∞ matrix takes the form of a rotation matrix under a similarity transform. Since similarity transformations preserve eigenvalues it follows that after an appropriate scale factor is applied, the eigenvalues of H_∞ should be 1, $e^{i\theta}$ and $e^{-i\theta}$, and its trace should be $1 + 2\cos\theta$ where, $R = \exp(J(\theta\mu))$ and $\|\mu\| = 1$. Furthermore the eigenvector of H_∞ corresponding to the unit eigenvalue will be $\eta = A\mu$. From this, we can compute $\tilde{A}\mu$ by premultiplying η by A_{est}^{-1}.

From these estimates, we can apply a control law of the form:

$$\omega = (\theta/\|\tilde{A}\mu\|)\tilde{R}\tilde{A}\mu. \quad (17)$$

This control law will again regulate the orientation error to zero as long as $\tilde{R}\tilde{A}$ is positive definite.

Note that these convergence results are *global* in nature, that is, they do not depend on the robots starting or ending positions. This distinguishes the presented analysis from local stability results egs [3].

4 Experimental Results

Experiments were carried out both in simulation and on a Nomadics XR4000 mobile robot platform. The simulation experiments modeled the action of a robot with all 6 degrees of freedom acting under the control schemes outlined in the previous section.

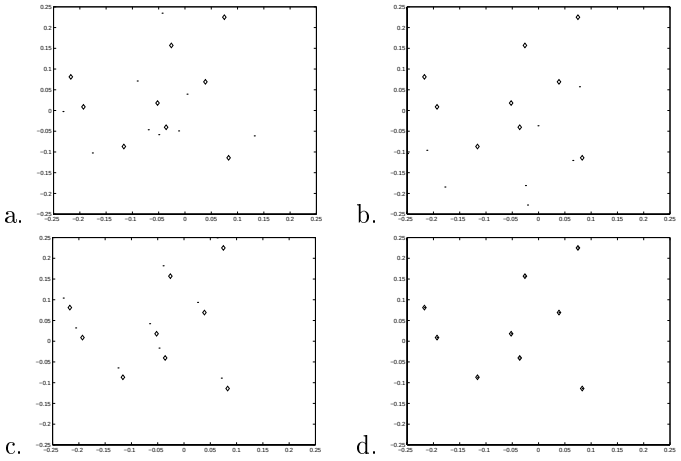

Figure 2: This figure demonstrates how the control law drives the feature points to the desired values over time even in the presence of large calibration errors. The dots denote the current positions of the projections of the feature points in the image while the diamonds represent the goal configuration.

In the simulation experiments the robot was started in an arbitrary position and instructed to servo the image measurements to a specified configuration. Figures 2a thru 2d show how the image of the feature points change as the robot moves towards its goal. The diamonds represent the goal configurations of the points while the dots denote the current projections of those features. These simulations verified that the control laws regulated the platform to the desired position and orientation on every occasion even in the presence of significant errors in the intrinsic and extrinsic calibration parameters.

During these experiments the rotation and translation control laws were employed concurrently so the robot translated and rotated at the same time. The fact that the system was able to converge perfectly

in every trial provides some evidence that the translational control law is robust in the face of disturbances, since the angular velocity of the platform actually causes some undesired translation of the camera frame due to extrinsic calibration errors.

The pose control scheme was also implemented on the Nomadics XR4000 mobile platform. In this case a simpler form of the relative orientation algorithm was employed which exploited the fact that the motion of the robot was confined to the plane. This relative orientation computation only required four point correspondences rather than the usual eight. The robot was instructed to drive the image measurements to a specified configuration as shown in Figure 3. The marked points on these image indicate the features that the robot used to accomplish its task. The intrinsic parameters of the camera in this example were known only approximately and the body frame of the robot was assumed to be coincident with the camera frame although this was not in fact the case.

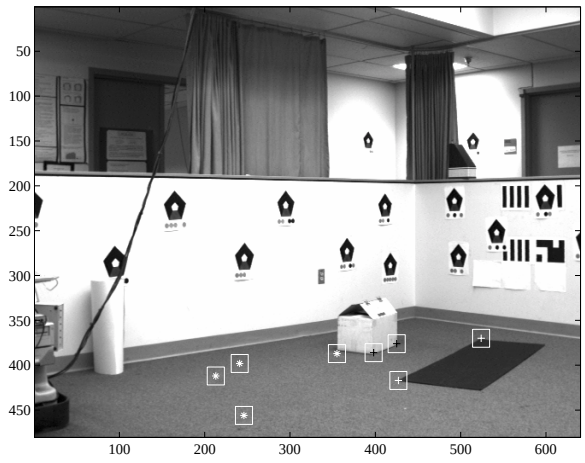

Figure 3: This figure shows the four features that were used for pose control. The crosses denote the positions of the features in the current image while the asterisks denote the target positions of the features. The targets shown on the walls were *not* used by the system.

Experiments were carried out to determine how accurately the pose control system could regulate the robot's position. The first step in these experiments was to place the robot at the target location and acquire an image of the scene from this vantage point. The robot was then moved to an arbitrary starting location and instructed to apply visual servoing until the positions of the tracked features matched their locations in the starting image. The robot achieved the desired pose by first rotating to the required orientation and then translating to the desired position. The robot's position and orientation were independently measured by the odometry system. Over a sequence of 5 trials the average error in the robots position at convergence was 9.584 cm, the average error in orientation was 1.1574 degrees and the average disparity between the measured image features and their target locations was 3.5192 pixels.

A second implementation of our pose control scheme on the mobile robot made use of vanishing points in the image to achieve better control of the robot's orientation. By tracking two parallel lines in the image the robot was able to obtain an estimate for the homography of the plane at infinity and, hence an estimate for its rotation which was used for pose control. Using the vanishing point information improved the accuracy of the pose control system significantly. Over a series of five trials the average error in the robot position at convergence was 7.04 cm, the average error in orientation was 0.5844 degrees and the average disparity between the image features and their target locations was 3.760 pixels. Figure 4a shows how the error in the robots orientation changes over time during the orientation phase of the motion while Figure 4b shows the robots trajectory in the plane as it moved towards the target position which is denoted by a diamond. The discontinuities in these graphs are artifacts of the implementation of the controller which involved discontinuous changes in the robot's velocity profile.

a.
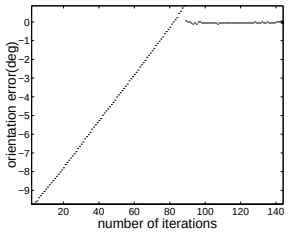

b.
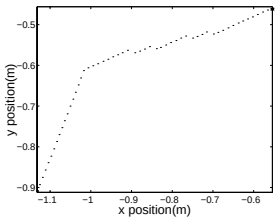

Figure 4: Figure a indicates how the orientation error of the robot decreases over time during the rotation phase of the control when vanishing points in the image are exploited while figure b indicates the robots trajectory towards the target position during the translation phase.

We expect that all of these results could be improved with better feature extraction and tracking algorithms and better control over the robot's velocity.

5 Conclusions and Future Work

This paper describes an approach to controlling the pose of a robotic platform using image measurements obtained from a camera rigidly mounted on that platform. The techniques make use of well-established techniques for computing the relative orientation of two camera positions from feature correspondences between the two images. These techniques do not require any a'priori knowledge of the locations of the features in the scene nor do they attempt to estimate such information online.

The techniques differ from methods based on pose estimation since they do not attempt to measure the location of the platform with respect to any particular target object in the scene. They also differ from other visual servoing techniques since no attempt is made to estimate the image Jacobian that relates control actions to changes in the image measurements.

An analysis has been presented which demonstrates that the proposed control schemes are convergent in the face of large calibration errors in both the intrinsic and extrinsic parameters. From this analysis, we are able to derive a precise characterization of the magnitudes of the errors that can be tolerated. Further analysis is required to elucidate the rate of convergence of these schemes and their disturbance rejection properties.

Acknowledgments

This work was supported by the NSF under grant number IRI-9711834, and DoD MURI DAAH04-96-1-0007

References

[1] Ronen Basri, E. Rivlin, and I. Shimshoni. Visual homing: surfing on the epipoles. In *International Conference on Computer Vision*, 1998.

[2] Koichiro Deguchi. Optimal motion control for image-based visual servoing by decoupling translation and rotation. In *Proceedings of the 1998 IEEE/RSJ International Conference on Intelligent Robots and Systems*, page 705, 1998.

[3] B. Espiau. Effect of camera calibration errors on visual servoing in robotics. In *3rd International Symposium on Experimental Robotics*, October 1993.

[4] B. Espiau, F. Chaumette, and P. Rives. A new approach to visual servoing in robotics. *IEEE Transactions on Robotics and Automation*, 8(3):313–326, June 1992.

[5] Olivier Faugeras. *Three-Dimensional Computer Vision*. MIT Press, 1993.

[6] Olivier Faugeras. Stratification of three-dimensional vision: projective, affine and metric representations. *Journal of the Optical Society of America A*, 12(3):465, March 1995.

[7] G. Hager, W. Chang, and A.S. Morse. Robot feedback control based on stereo vision: Towards calibration-free hand-eye coordination. *IEEE Control Systems Magazine*, 15(1):30–39, 1995.

[8] Gregory D. Hager. Calibration-free visual control using projective invariance. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 1009–1015, 1995.

[9] Richard E. Hartley. Estimation of relative camera positions for uncalibrated cameras. In *Second European Conference on Computer Vision*, page 579, May 1992.

[10] Richard I. Hartley. In defense of the eight-point algorithm. *IEEE Trans. Pattern Anal. Machine Intell.*, 19(6):580, June 1997.

[11] Nicholas Hollinghurst and Roberto Cipolla. Uncalibrated stereo hand-eye coordination. *Image and Vision Computing*, 12(3):187–192, April 1994.

[12] R. Horaud, F. Dornaika, and B. Espiau. Visually guided object grasping. *IEEE Trans. on Robotics and Automation*, 14(4):525–533, August 1998.

[13] H.C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293:133–135, September 1981.

[14] Ezio Malis, Francois Chaumette, and Sylvie Boudet. 2d 1/2 visual servoing stability analysis with respect to camera calibration errors. In *Proceedings of the 1998 IEEE/RSJ International Conference on Intelligent Robots and Systems*, page 691, 1998.

[15] Ezio Malis, Francois Chaumette, and Sylvie Boudet. Positioning of a coarse-calibrated camera with respect to an unkown object by 2d 1/2 visual servoing. In *IEEE Int. Conf. on Robotics and Automation*, page 1352, 1998.

[16] N. P. Papanikolopoulos and P. K. Khosla. Eye-in-hand robotic tasks in uncalibrated environments. *IEEE Trans. Robotics and Automation*, 13(6):903–914, December 1997.

[17] N. P. Papanikolopoulos, B. J. Nelson, and P. K. Khosla. Six degree-of-freedom hand/eye visual tracking with uncertain parameters. *IEEE Trans. Robotics and Automation*, 11(5):725–732, October 1995.

[18] R. Pissard-Gibollet and P. Rives. Applying visual servoing techniques to control a mobile hand-eye system. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 166–171, Nagoya, Japan, May 1995.

[19] D. Simon, K. Kapellos, and B. Espiau. Design of control procedures for a free-floating underwater manipulation system. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 2158–2165, Albuquerque, April 1997.

[20] S. Soatto and P. Perona. Structure-independent visual motion control on the essential manifold. In *Proc. of the IFAC Symposium on Robot Control (SYROCO), Capri, Italy*, pages 869–876, Sept 1994.

[21] W. J. Wilson, C. C. Williams Hulls, and Graham S. Bell. Relative end-effector control using Cartesian position based visual servoing. *IEEE Trans. Robotics and Automation*, 12(5):684–696, October 1996.