

Nonholonomic camera space manipulation using cameras mounted on a mobile base

Bill Goodwine, Michael Seelinger, Steven B. Skaar and Qun Ma

Aerospace and Mechanical Engineering, University of Notre Dame, Notre Dame, Indiana, U.S.A.

ABSTRACT

The body of work called "Camera Space Manipulation" is an effective and proven method of robotic control. Essentially, this technique identifies and refines the input-output relationship of the plant using estimation methods and drives the plant open-loop to its target state. Three-dimensional "success" of the desired motion, i.e., the end effector of the manipulator engages a target at a particular location with a particular orientation, is guaranteed when there is camera space success in two cameras which are adequately separated. Very accurate, sub-pixel positioning of a robotic end effector is possible using this method. To date, however, most efforts in this area have primarily considered holonomic systems. This work addresses the problem of nonholonomic camera space manipulation by considering the problem of a nonholonomic robot with two cameras and a holonomic manipulator on board the nonholonomic platform. While perhaps not as common in robotics, such a combination of holonomic and nonholonomic degrees of freedom are ubiquitous in industry: fork lifts and earth moving equipment are common examples of a nonholonomic system with an on-board holonomic actuator. The nonholonomic nature of the system makes the automation problem more difficult due to a variety of reasons; in particular, the target location is not fixed in the image planes, as it is for holonomic systems (since the cameras are attached to a moving platform), and there is a fundamental "path dependent" nature of nonholonomic kinematics. This work focuses on the sensor space or camera-space-based control laws necessary for effectively implementing an autonomous system of this type.

Keywords: camera space manipulation, robotic vision

1. INTRODUCTION

Camera space manipulation ("CSM") is a proven method for using a vision-based control strategy to precisely control robotic systems. Until now, CSM has been applied primarily to holonomic robotic systems. This paper presents some theoretical aspects of the extension of CSM to general nonholonomic systems and presents simulation and experimental results related to a specific nonholonomic CSM implementation. We restrict our attention in this paper to a hybrid robotic system which possesses both holonomic and nonholonomic features, specifically, a holonomic robotic arm which is attached to a mobile nonholonomic robotic platform. The task is for the end effector of the arm to engage a particular target location in space, usually located on the surface of an object.

The primary difficulties with extending CSM to nonholonomic systems are two-fold. First, unlike its usual implementation,¹⁻³ the class of nonholonomic systems we consider here has the cameras mounted on-board a mobile robot. In this scenario, a desired target location for the end effector does not remain fixed in camera space. This is in contrast to the holonomic case in which the cameras are fixed relative to the desired target point. Second, nonholonomic systems in general are more difficult from a control theoretic point of view than holonomic systems, and this characteristic carries over to the CSM context as well.

This work provides a novel approach to enabling autonomous mobile robot control. While autonomy in general is a desirable characteristic for mobile robots, there are many applications in which nearly complete autonomous control is necessary. One such application is rover-based exploration of locations where constant contact with a human operator is not practical, or time delays in communication are prohibitive, *e.g.*, extraterrestrial surface exploration. The ultimate goal would be the capability for a user to specify a "target location" (probably via a video interface) for an end effector or instrument to which the robot would then autonomously navigate. The results presented in this

Other author information: (Send correspondence to Bill Goodwine)

Bill Goodwine: Email: jgoodwin@nd.edu

Steven B. Skaar.: Email: Steven.B.Skaar.1@nd.edu

Michael Seelinger: Email: mseeling@nd.edu

Qun Ma: Email: qma1@nd.edu

paper are directed toward the aspect of this problem of achieving very precise engagement between the end effector and desired target location.

This paper is organized as follows. Section 2 reviews the theoretical bases for CSM and contrasts the method with other vision-based control strategies, particularly visual servoing and calibration methods and further discusses the complications associated with extending CSM to nonholonomic systems. Section 3 presents the details of the CSM-based control strategy for a nonholonomic robotic system. Section 4 presents simulation results for this system and presents preliminary experimental results of a real implementation of nonholonomic CSM. Finally, Section 5, in addition to drawing conclusions from the presented material, discusses, in detail, planned future developments of this research project.

2. BACKGROUND

Although the present development entails a significant extension — to systems with cameras that move on a mobile, nonholonomic base — it is based upon the method of camera-space manipulation (CSM). CSM is predicated on the idea that the desired three-dimensional spatial relationship between the manipulated body and the stationary workpiece, or receiving body, can be represented as a corresponding geometrical relationship between the two bodies as they appear in a two-dimensional image plane or camera space. In this way, it differs from the calibration approach that typically is applied to employ vision to control robots.

Much research is directed toward the use of calibration methods.⁴⁻⁷ Such methods entail two separate calibration events which are equally critical for maneuver success: 1. calibration of cameras relative to a fixed, physical "world" coordinate system⁸⁻¹⁰; and 2. calibration of the robot kinematics relative to the physical coordinate system.¹¹⁻¹⁴ Thus the imaging and the manipulation steps are separated. The cameras are used as a kind of measuring device to locate the physical coordinates of the workpiece. After this location, as a separate activity, the internal coordinates of the robot are driven to a terminus based upon an inverse-kinematics calculation.

As early as the mid 1970's (e.g. Tani et al.¹⁵) it was noted that there is no need to separate the imaging and manipulation steps. Organisms provide an example of hand-eye coordination which is certainly based upon pursuing maneuvers in terms of "sensor-space" success, without reference to any absolute coordinate system. In "vision feedback" or "visual servoing," errors as perceived in the sensor reference frame itself — disparities between the actual and desired sensor-frame locations of key features — are used in a control law in order to drive the key features toward a zero-image-error state. Some knowledge of the partial sensitivity of image response to increments in each joint rotation is required, but terminal success tends to be robust to errors in the matrix Jacobians which are for visual servoing.

In Corke¹⁶ and Sanderson and Weiss,¹⁷ diverse forms of implementation of this idea are put forth. The strategy of driving image-plane errors toward zero using inputs as determined through matrix Jacobians continues to be pursued vigorously.^{4,12,13,15,17-22} Each of the various forms of visual servoing has advantages or disadvantages^{16,23} depending upon which of the various and diverse possible applications is addressed.

The alternative of "Camera-Space Manipulation," introduced in 1986 and subsequently developed,^{1-3,7,14,24-29} is highly effective in situations where, during maneuver execution, the participating, uncalibrated, wall- or tripod-mounted cameras can reliably be counted upon to remain stationary between the time of visual access to the workpiece (prior to any blockage) and the time of maneuver culmination, and where the workpiece also remains stationary. CSM is especially useful where maneuver closure itself can cause significant obscuration to cameras of the visual error prior to maneuver culmination. Also, certain practical realities associated with artificial vision — coarseness of the discretized image, as well as computational intensiveness of image-feature identification/location — tend to favor camera-space manipulation with its application of new visual data to input-output estimation rather than to simple feedback.

CSM is particularly recommended in cases where a holonomic robot is used, and where cameras are stationary. Holonomic robots guarantee a continuous, locally monotonic relationship between internal-joint rotation and the camera-space response of the manipulated body; and stationary cameras guarantee time invariance of this relationship. Nevertheless, even for the hybrid robot of the present discussion, where a holonomic arm is placed onto a nonholonomic base, and where the CSM cameras move as a rigid body with that base, the premise still holds: If each of two adequately separated cameras realize the camera-space objectives which are associated with the required three-dimensional relationship of the bodies of interest then the desired terminal pose has been realized. This is

true whether the camera–space relationships are achieved only by moving the manipulated body in each image (as is the CSM norm to date), or whether they are achieved, as discussed herein, by achieving a mix of camera–space positionings — positioning of the manipulated body via exercising the holonomic, on–board degrees of freedom, and positioning of the target body by exercising the mobile base. The present paper discusses a formulation as well as simulation and experimental results pertaining to this prospectively very useful combination.

3. NONHOLONOMIC CAMERA SPACE MANIPULATION

This section presents the details of CSM as applied to the hybrid holonomic/nonholonomic systems under consideration. For a holonomic systems, let $\Theta = (\theta_1, \dots, \theta_m)$ be the joint variables for the system. Let $f(\Theta)$ represent the nominal *a priori* kinematic model of the manipulator and let $g(x, y, z, \mathbf{D})$ represent the model for the mapping between physical space (x, y, z) and camera space, $X = (x_c, y_c)$, *i.e.*,

$$\begin{pmatrix} x_c \\ y_c \end{pmatrix} = g(x, y, z, \mathbf{D})$$

where, for the particular camera, lens and frame grabber combination, the vector of parameters, \mathbf{D} could include as few as six elements as required to capture the position and orientation of the camera in relation to manipulator space.

As a fundamentally open loop method, CSM involves computing a reference trajectory, $\Theta_r(t)$ based upon the kinematic model, $f(\Theta)$ and camera model, $g(x, y, z, \mathbf{D})$, from the current state through to the desired maneuver terminus. This trajectory is initially passed to the joint level controller for the robot, and then is periodically updated. The trajectory is updated when newly processed visual information becomes available at various discrete times, t_i during the evolution of the trajectory. Of course, it is important, at those points in time where transitions in the reference joint trajectories occur due to new information based reference trajectory “upgrades,” that these transitions are adequately smooth.

To understand this strategy, consider at an instant, t_i , for the manipulator to be moving toward the goal and for a visual sample to be acquired. Since the image analysis will require some time to process, the arm will continue to proceed toward the maneuver goal using the “current” joint level trajectory plan, $\Theta_r(t)$ calculated prior to t_i . This information will include the camera space coordinates of the end effector acquired at discrete times prior to t_i , *i.e.*, $\{t_{i-1}, t_{i-2}, \dots\}$ and $\{X(t_1), X(t_2), \dots, X(t_{i-1})\}$, along with the simultaneously measured joint locations, $\{\Theta(t_1), \Theta(t_2), \dots, \Theta(t_{i-1})\}$. If the sample acquired at t_i becomes processed before the end of the maneuver, then this new information will be used to update $\Theta_r(t)$ beginning at some future time.

The basis on which the reference trajectories are made and updated is estimation of the input/output relationship where Θ is the input and X is the output, using the “locally” (in joint and camera spaces, as well as time) most relevant information available. Being model–based, the estimation procedure should make use of the geometric structure of the robot and camera models. The strategy that has been adopted to date has been based upon a straight–forward combination of the nominal forward kinematics for the robot, and the pinhole, or perspective, camera model. The pinhole camera model is given by

$$x_{c_i} = f \frac{X_i}{Z_i} \quad (1)$$

$$y_{c_i} = f \frac{Y_i}{Z_i} \quad (2)$$

where f is the focal length of the cameras, the (X_i, Y_i, Z_i) is the location of interest measured from a reference frame located at the focal point of camera i , where the Z –axis is parallel to the orientation of the camera. Substituting the one into the other gives a relationship directly between Θ and (x_c, y_c) ,

$$X = (x_c, y_c) = g(f(\Theta), \mathbf{D}) = h(\Theta, \mathbf{D}).$$

In principle, the six element vector \mathbf{D} could be determined by minimizing over all \mathbf{D}

$$\Gamma(\mathbf{D}) = \sum_i (h(\Theta, \mathbf{D}) - X(t_i))^T W_i (h(\Theta, \mathbf{D}) - X(t_i)),$$

where each W_i is a 2×2 symmetric matrix adjusted to give higher “weights” to those samples acquired nearest to the terminus of interest. However, numerical stability problems associated with the combination of the pinhole camera model and the typically confined span or extent of the applied samples have necessitated, in practice, a modification of this strategy, which is the introduction of an intermediate form of estimation as described in Skaar, *et al.*³ Furthermore, good estimates require an initial diversity of samples, which may require an initial “preplanned trajectory” which is a “sufficiently rich” motion not directly pertinent to the current maneuver.

Due to these problems of numerical stability, an alternative form of estimation model, $h(\Theta, \mathbf{D})$ of the camera space kinematics, denoted $h^*(\Theta, \mathbf{C})$ has been used in practice.³ Estimation of the six elements, \mathbf{C} of $h^*(\Theta, \mathbf{C})$ is the same as above in that it is based upon local samples by minimizing over all \mathbf{C}

$$\Gamma^*(\mathbf{C}) = \sum_i (h^*(\Theta, \mathbf{C}) - X^*(t_i))^T W_i (h^*(\Theta, \mathbf{C}) - X^*(t_i)).$$

The difference here is that the camera space samples are replaced with modified, or “flattened” samples, $X^*(t_i)$ in such a way as to produce the identical effect at the final time as that which would be realized with direct use of the pinhole camera model, h .

The holonomic degrees of freedom for the systems under consideration are controlled exactly in this manner, which we will refer to as “pure CSM.” The nonholonomic degrees of freedom have been controlled in variety of ways, which are detailed in the Section 4 describing both the computer simulation results and physical experimental results.

4. IMPLEMENTATION

We have verified the efficacy of certain aspects of the control strategy discussed in Section 3 using numerical simulations and implementing the strategy on a real physical robotic system. Again, we emphasize that this paper represents “work in progress,” and the complete evolution of our work to pure nonholonomic CSM is the subject of ongoing research. The current simulation results utilize a combination of CSM and certain aspects of visual servoing; whereas, the experimental results are closer to the true nature of CSM.

A schematic representation of the system is illustrated in Figure 1. The variables (x, y, ϕ) represent the position and orientation of the cart on the plane (ground) and the variables $(\theta_1, \theta_2, \theta_3)$ represent the angular position of the first and second wheels and the one degree of freedom arm attached to the cart, respectively. In camera space, the variables (x_{co1}, y_{co1}) and (x_{co2}, y_{co2}) represent the camera space coordinates of the object in cameras one and two, respectively, and (x_{ce1}, y_{ce1}) and (x_{ce2}, y_{ce2}) represent the camera space coordinates of the end effector in the two respective camera spaces. The variables (X_2, Y_2, Z_2) in the figure are coordinates of a point with respect to the coordinate system located at the focal point of camera two, which is used in the mapping from physical space to camera space (recall Equation 1). The radius of the wheels is denoted by r and the distance between the wheels is denoted by b . Point A is the tip of the end effector and point B is the desired target location. The upper left corner of Figure 1 is a representation of camera space two, illustrating the camera space coordinate system, points A and B , and the trace of point A through camera space.

The system is subjected to the usual nonholonomic kinematic constraints for a system of this type, namely,

$$\begin{aligned} \dot{x} &= r \cos \phi \left(\frac{\dot{\theta}_1 + \dot{\theta}_2}{2} \right) \\ \dot{y} &= r \sin \phi \left(\frac{\dot{\theta}_1 + \dot{\theta}_2}{2} \right) \\ \dot{\phi} &= -\frac{r}{b} (\dot{\theta}_1 - \dot{\theta}_2), \end{aligned}$$

and the full set of dynamic equations of motion are given by

$$\begin{aligned} \left(\frac{mr^2}{4} + \frac{Jr^2}{4b^2} + mr^2 + I_1 + \frac{I_2r^2}{2b^2} \right) \ddot{\theta}_1 + \left(\frac{mr^2}{4} - \frac{Jr^2}{4b^2} - \frac{I_2r^2}{2b^2} \right) \ddot{\theta}_2 &= \tau_1 \\ \left(\frac{mr^2}{4} - \frac{Jr^2}{4b^2} - \frac{I_2r^2}{2b^2} \right) \ddot{\theta}_1 + \left(\frac{mr^2}{4} + \frac{Jr^2}{4b^2} + mr^2 + I_1 - \frac{I_2r^2}{2b^2} \right) \ddot{\theta}_2 &= \tau_2, \end{aligned}$$

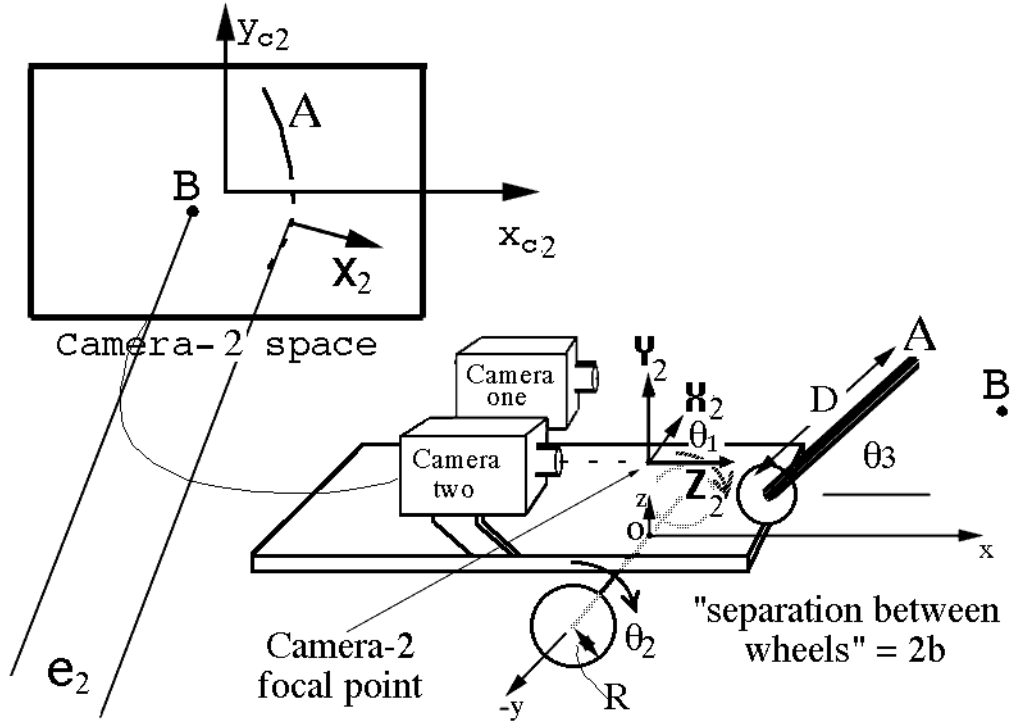


Figure 1. Nonholonomic camera space manipulation system.

where τ_1 and τ_2 represent the torques imparted to the wheels from the motors, m is the mass of the cart, I_1 is the moment of inertia of each wheel about its axis of rotation, I_2 is the moment of inertia of each wheel about the vertical axis and J is the rotational moment of inertia of the cart about the vertical axis.

The simulation program also models the dynamics of the DC motors in the usual way,

$$\begin{aligned} V_1 &= R_1 i_1 + I_1 \dot{i}_1 + K_m \dot{\theta}_1 \\ V_2 &= R_2 i_2 + I_2 \dot{i}_2 + K_m \dot{\theta}_2, \end{aligned}$$

where V_i is the voltage input, R_i is the resistance, I_i is the inductance, i_i is the current and K_m is the motor constant.

4.1. Computer Simulation

In its current form, the simulation uses a “servoing” approach to drive the cart to the final point of engagement. Since the arm has one degree of freedom, the path of the end effector through each camera space will be a one dimensional trace, as illustrated in Figures 1 and 2. The error signal for each camera has been considered in a variety of ways, but typically is taken to be the shortest distance between the location of the target point, denoted by B and any point on the trace of the end effector through camera space, as illustrated by δ in Figure 2. Another error signal considered was simply the horizontal distance between the object and end effector. We note that this is a visual servoing approach; however, it is one which requires a CSM understanding of the on-board camera space kinematics. Furthermore, the approach involves the use of two separate cameras which act to provide separate camera space requirements, in this case the requirements for successful collocation of points A (the end effector) and B (the target location) in physical space.

While the use of two separate two dimensional camera spaces to culminate a three degree of freedom motion involves objective redundancy in “pure CSM,” and this redundancy would ordinarily call for the algebraic form of

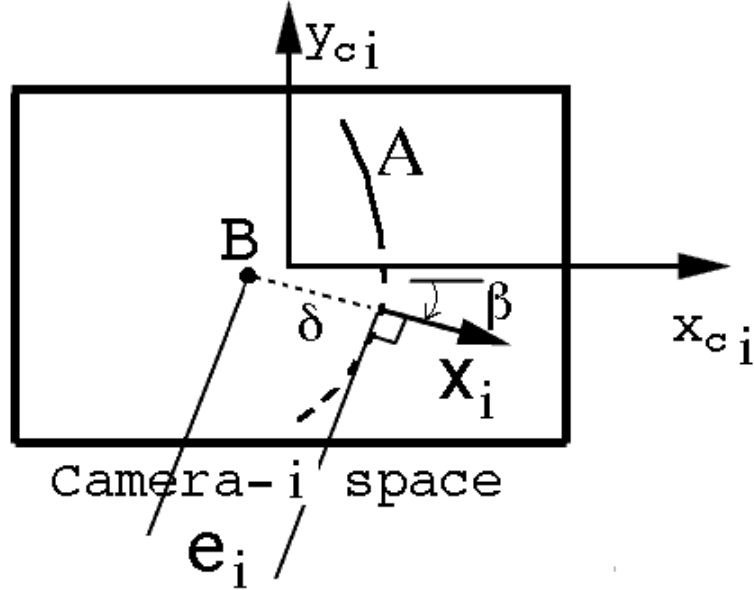


Figure 2. Camera space error measure.

camera space kinematics, we use the differential form of the kinematics (as is necessary for a nonholonomic system) only for a portion of the process, which is facilitated by control of just two degrees of freedom, namely, camera space positioning of the target point B to *any* location on the camera space trace of the end effector tip. Hence, there is no need to deal with the redundancy insofar as the nonholonomic degrees of freedom are concerned. Note that in this context, it is not possible to bring the base to the point where the point B lies on the camera space trace of the tip of the end effector in *both* cameras without this being the correct terminal base location insofar as there exists a θ_3 which in fact produces physical coincidence of the two points of interest.

The problem in visual servoing of image error obscuration near the critical time of closure is avoided due to the CSM aspect of the control scheme. After the camera space trace of the tip of the end effector is determined in each camera space, the arm itself can be kept “out of the way” until the base has reached the above objectives. Control of the holonomic degree of freedom, θ_3 can be accomplished in the usual CSM way: identification and estimation-based open-loop pursuit of camera space objectives separately in each camera, which are compatible with successful realization of the physical maneuver. That a single on-board degree of freedom is adequate to accomplish this is ensured if the target point B is located on the trace of the tip of the end effector in each image plan.

One reason to consider the use of the visual servoing paradigm for wheel control rather than “pure CSM” is the fact that, depending upon surface conditions, wheel slip may prevent the existence of a continuous reliable camera space response to internal rotation that characterized holonomic systems. Inherent in the servoing approach is the fact that the ultimate precision of motion will be limited to something on the order of pixel resolution. Even though, by using visual servoing, some prior calibration is required in order to attain the needed matrix Jacobian, the system should be robust, which is a well known characteristic of visual servoing control strategies. Additionally, it is likely that, ultimately, refinement of the matrix Jacobian using recent input/output experience could be accomplished autonomously on-line to improve convergence performance.

Now, this error signal can be utilized in a variety of ways. One approach would be compute the Jacobian that relates changes in wheel velocities to changes in camera space coordinates (a 4×2 matrix), *i.e.*,

$$\begin{pmatrix} \dot{x}_{co1} \\ \dot{y}_{co1} \\ \dot{x}_{co2} \\ \dot{y}_{co2} \end{pmatrix} = J \begin{pmatrix} \dot{\theta}_1 \\ \dot{\theta}_2 \end{pmatrix}.$$

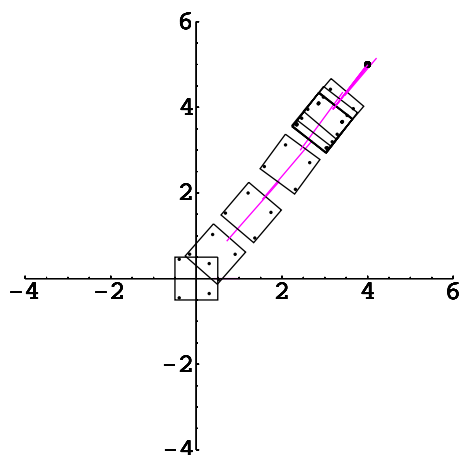


Figure 3. Computer simulation results.

Then, using a pseudo-inverse or other technique that may exploit the particular geometry of this problem, we can reduce the 4×2 Jacobian to an invertible 2×2 matrix.

Using the pseudo-inverse technique as well as even more *ad hoc* techniques, many hundreds of simulations indicate the likelihood of global convergence of the system to the desired end point. A graphical illustration of the evolution of a cart in a typical simulation is shown in Figure 3, which is a top view of the cart. The square is the outline of the body of the cart, the four dots in the rectangle represent the points of contact between the wheels and the ground, the larger dot is the desired target location and the line attached to the body represents the one degree of freedom arm. Future work with the simulation will include investigating experimentally difficult subjects such as wheel slip, uneven terrain, and the evolution to a more “pure CSM” control strategy, which will possibly involve updating an estimated Jacobian and an optimal control path planning strategy.

4.2. Experimental Implementation

Our experimental platform is shown in Figure 4. In the picture, the cart is facing to the right with the one degree of freedom arm extended to the right and is constrained to move in an arc with an axis of rotation parallel to the axes of rotation for the front two drive wheels. The white ovals near the end of the arm are cues which simplify the image processing demands. The two front wheels are the drive wheels, each of which is independently actuated, and the rear wheels are casters. The cart is tethered to a PC which does the necessary image processing and provides control commands.

The arm on this system is controlled via “pure CSM,” which entails estimation and updating the parameters \mathbf{C} . The nonholonomic degrees of freedom directly make use of the CSM based control of the holonomic degree of freedom in the following way. Basic geometry dictates that given an initial orientation for the cart and a desired final position for the end effector, there exists a unique circular arc tangent to the initial position of the cart and going through the target point. Using the camera model and estimated values for the vector \mathbf{C} , an estimated spatial location of the target point can be determined, from which the circular arc trajectory is simple to determine. The cart is controlled by commanding the ratio of wheel velocities which result in the desired radius of curvature. When updated parameters are obtained, a revised \mathbf{C} is obtained, which then can update the radius of curvature for the circular trajectory.

As an aside, we note that as a spatial estimation of the actual location of the target, is actually substantially inaccurate relative to the extreme precision accomplished by the total motion of the cart. However, this is irrelevant since the ultimate goal is *camera space* coincidence (guaranteeing physical coincidence) of the location of the end effector tip and target point, which illustrates one of the fundamental benefits of working in camera space, and with CSM in particular.



Figure 4. Nonholonomic camera space manipulation experimental platform.

Current experiments indicate that at final engagement the end effector engages the target location with usual tolerances indicated in Table 1. A close-up photograph of a typical engagement is illustrated in Figure 5. For scale, the cross section of the arm is one inch by one inch. The desired target point is the center of the inner circle near the end effector tip.

x -direction (normal to target surface):	$\pm 2.0\text{mm}$
y -direction (tangent to target surface):	$\pm 0.5\text{mm}$
z -direction (tangent to target surface):	$\pm 0.5\text{mm}$

Table 1. Nonholonomic camera space manipulation engagement tolerances.

5. CONCLUSIONS

This paper presents the initial extension of camera space manipulation to general nonholonomic systems. While still work in progress, both computer simulation and experimental results illustrate the efficacy of using CSM for nonholonomic problems which results in millimeter-scale precision control of the final engagement of the end effector with a desired target point.

Current work includes completely extending the model based estimation aspect of CSM to the nonholonomic degrees of freedom for the system. Also, more demanding scenarios must be investigated, including investigating the effect of irregularities of the surface upon which the robot rolls, including perhaps very wavy terrain which would possibly sometimes occlude the view of the target point. Additionally, one important extension of this work would be to investigate the effects of wheel slip, both in the servoing context as well as a more pure CSM based approach. One important application of this research is rover-type vehicles, which often operate on sandy surfaces which may

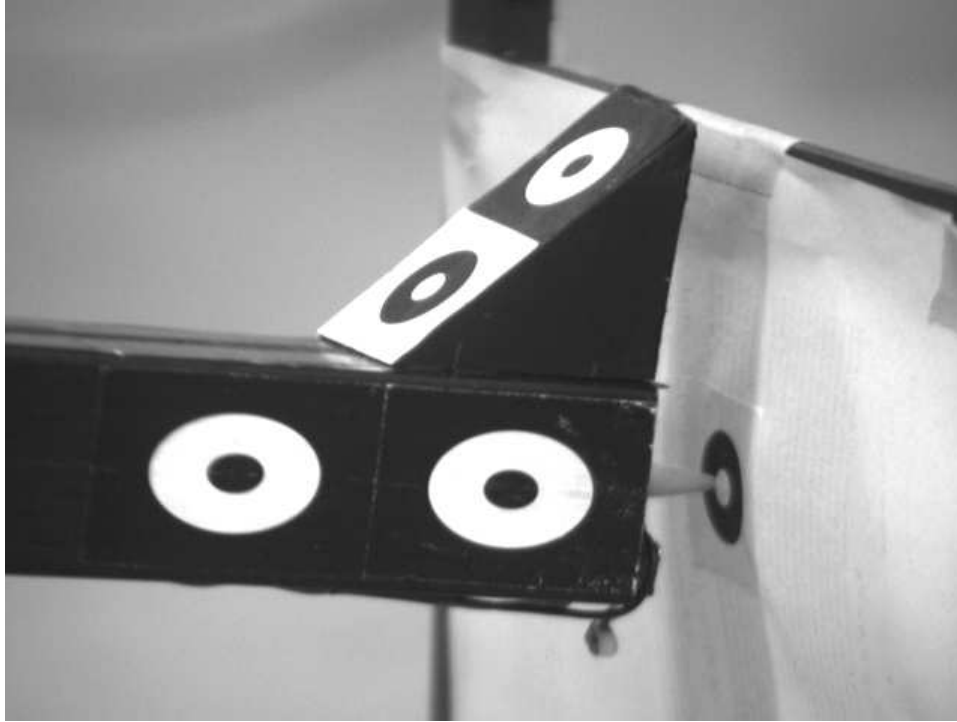


Figure 5. End effector engaging target point.

often yield substantial and irregular wheel slip. Finally, other issues such as efficient algorithms or methods to avoid obstacles that may arise along the nominal path of the robot must be considered before our system is ready for “real world” implementation.

Acknowledgments

The authors gratefully acknowledge the Jet Propulsion Lab in Pasadena, California, for providing an SBIR which provided support for hardware and Michael Seelinger, and the National Science Foundation, NSF: EEC-9700702.

REFERENCES

1. S. B. Skaar, W. H. Brockman, and R. Hanson, “Camera space manipulation,” *International Journal of Robotics Research* **6**(4), pp. 20–32, 1987.
2. S. B. Skaar and E. Gonzalez-Galvan, “Versatile and precise manipulation,” in *Teleoperation and Robotics in Space*, S. B. Skaar and C. F. Ruoff, eds., pp. 241–279, AIAA, Washington, D.C., 1994.
3. S. B. Skaar, U. A. Korde, and W. Z. Chen, “Application of a precision enhancing measure in 3-d rigid-body positioning using camera-space manipulation,” *International Journal of Robotics Research* **16**(2), pp. 240–257, 1997.
4. G. V. Puskorius and L. A. Feldkamp, “Global calibration of a robot/vision system,” in *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 190–195, 1987.
5. R. Y. Tsai, “A versatile camera calibration technique for high accuracy 3d machine vision methodology using off-the-shelf tv cameras and lenses,” *IEEE Journal of Robotics and Automation* **RA-3**(4), pp. 323–344, 1987.
6. J. Weng, P. Cohen, and M. Herniou, “Camera calibration with distortion models and accuracy evaluation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **14**(10), pp. 965–980, 1992.
7. H. Zhuang and Z. S. Roth, *Camera-Aided Robot Calibration*, CRC Press, Boca Raton, 1996.
8. S. Maybank and O. D. Faugeras, “A theory of self-calibration of a moving camera,” *International Journal of Computer Vision* **8**(2), pp. 123–151, 1990.

9. H. Trivedi, "A semi-analytic method for estimating stereo camera geometry from matched points," *Image and Vision Computing* **9**, pp. 227–236, 1991.
10. E. Baumgartner, M. Seelinger, M. Fessler, A. Aldekamp, E. Gonzalez-Galvan, J. D. Yoder, and S. B. Skaar, "Accurate 3-d robotic point positioning using camera-space manipulation," in *24th Midwest Mechanics Conference*, T. J. Rudolph and L. W. Zachary, eds., 1995.
11. D. J. Bennet, D. Geiger, and J. M. Hollerbach, "Autonomous robot calibration for hand-eye coordination," *International Journal of Robotics Research* **10**(5), pp. 550–559, 1991.
12. R. Bernhardt and S. L. Albright, *Robot Calibration*, Chapman and Hall, London, 1993.
13. F. Chaumette, P. Rives, and B. Espiau, "Classification and realization of the different vision based tasks," in *Visual Servoing*, K. Hashimoto, ed., pp. 257–284, World Scientific, 1993.
14. B. W. Mooring, Z. S. Roth, and M. Driels, *Fundamentals of Manipulator Calibration*, John Wiley and Sons, New York, 1991.
15. K. Tani, M. Abe, K. Tanie, and T. Ohno, "High precision manipulator with visual sense," in *Proceedings of ISIR*, pp. 561–568, 1977.
16. P. I. Corke, "Visual control of robot manipulators — a review," in *Visual Servoing*, K. Hashimoto, ed., pp. 1–32, World Scientific, 1993.
17. P. Corke, "Video-rate robot visual servoing," in *Visual Servoing*, K. Hashimoto, ed., pp. 257–284, World Scientific, 1993.
18. J. T. Feddema and O. R. Mitchell, "Vision-guided visual servoing with feature-based trajectory generation," *IEEE Transactions on Robotics and Automation* **5**(5), pp. 691–700, 1989.
19. J. T. Feddema, C. S. G. Lee, and O. R. Mitchell, "Feature-based visual servoing of robotic systems," in *Visual Servoing*, K. Hashimoto, ed., pp. 105–138, World Scientific, 1993.
20. K. Hashimoto, *Visual Servoing*, World Scientific, Singapore, 1993.
21. W. Jang, K. Kim, M. Chung, and Z. Bien, "Concepts of augmented image space and transformed feature space for efficient visual servoing of an "eye-in-hand" robot," *Robotica* **9**, pp. 203–212, 1991.
22. B. Nelson, N. P. Papanikolopoulos, and P. K. Khosla, "Visual servoing for robotic assembly," in *Visual Servoing*, K. Hashimoto, ed., pp. 139–164, World Scientific, 1993.
23. S. Hutchinson, G. Hager, and P. Corke, "A tutorial on visual servo control," *IEEE Transactions on Robotics and Automation* **12**(5), pp. 651–670, 1996.
24. E. Gonzalez-Galvan and S. B. Skaar, "Servoable cameras for three-dimensional positioning with camera-space manipulation," in *Proceedings IASTED Robotics and Manufacturing*, pp. 260–265, 1995.
25. E. J. Gonzalez-Galvan and S. B. Skaar, "Efficient camera-space manipulation using moments," in *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 3407–3412, 1996.
26. E. Gonzalez-Galvan, M. Seeliner, J. D. Yoder, E. Baumgartner, and S. B. Skaar, "Control of construction robots using camera-space manipulation," in *Robotics for Challenging Environments, Proc of RCE II*, L. A. Demsetz, ed., pp. 57–63, 1996.
27. E. J. Gonzalez-Galvan, S. B. Skaar, U. A. Korde, and W. Z. Chen, "Application of a precision enhancing measure in 3-d rigid-body positioning using camera-space manipulation," *International Journal of Robotics Research* **16**(2), pp. 240–257, 1997.
28. U. A. Korde, E. Gonzalez-Galvan, and S. B. Skaar, "Three-dimensional camera-space manipulation using servoable cameras," in *Proceedings Intelligent Robots and Computer Vision*, pp. 658–667, SPIE, 1992.
29. R. K. Miller, D. G. Stewart, H. Brockman, and S. B. Skaar, "A camera space control system for an automated forklift," *IEEE Transactions on Robotics and Automation* **10**, pp. 710–716, October 1994.