

RealTime Tracking Meets Online Grasp Planning

Danica Kragić,

Computer Vision and Active Perception
Numerical Analysis and Computer Science
Royal Institute of Technology
SE-100 44 Stockholm, Sweden [†]

Andrew T. Miller and Peter K. Allen

Department of Computer Science
Columbia University
New York, NY 10027 *

Abstract

This paper describes a synergistic integration of a grasping simulator and a real-time visual tracking system, that work in concert to 1) find an object's pose, 2) plan grasps and movement trajectories, and 3) visually monitor task execution. Starting with a CAD model of an object to be grasped, the system can find the object's pose through vision which then synchronizes the state of the robot workcell with an online, model-based grasp planning and visualization system we have developed called GraspIt. GraspIt can then plan a stable grasp for the object, and direct the robotic hand system to perform the grasp. It can also generate trajectories for the movement of the grasped object, which are used by the visual control system to monitor the task and compare the actual grasp and trajectory with the planned ones. We present experimental results using typical grasping tasks.

1 Introduction

Visual control of robotic tasks is a major goal of current robotics research. The advent of real-time vision systems has led to the creation of systems that use cameras to control robotic tasks such as grasping [13]. This paper describes a system that extends this paradigm in a number of ways. Most importantly, we have been able to integrate real-time vision with online, model-based simulation, to create a grasp planning, control and monitoring system that can visually determine an object's pose, plan a stable grasp, execute the grasp, and monitor the task for errors and correctness (see Fig. 1). The vision system is respon-

sible for providing accurate and fast estimates of an object's pose, while the grasping simulator uses these estimates to plan 1) an effective and stable grasp of the object and 2) a collision free trajectory in the workspace for transporting the grasped object to a new position. The simulator output is then used to execute the planned grasp and trajectory, which is monitored by the vision system.

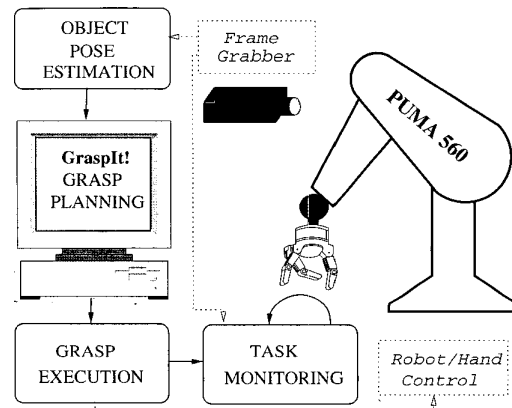


Figure 1: Block diagram of the system.

By merging visual control with grasp planning, we have created a system that can extend the idea of autonomous robot control. Using a calibrated vision system, we only need a CAD model of the object to be grasped to both track the object and determine its pose in real-time. The grasping simulator, previously developed by Miller and Allen [9, 10], is able to compute a stable grasp with a quality measure in real-time using the same CAD model. The grasp stability analysis also includes material properties of the object, and the simulator contains a library of robots and hands to choose from, giving it added flexibility for use in many different robot workcells. Below, we describe the vision and grasp simulation modules, and then discuss the integration of these modules on real grasping tasks using a Barrett hand.

*This work was supported in part by an ONR/DARPA MURI award ONR N00014-95-1-0601 and NSF grant CDA-96-25374. We would also like to thank Professor Gerd Hirzinger and Dr. Max Fischer from the German Aerospace Center (DLR) for providing us with models of their robotic hand.

[†]This research has been sponsored by the Swedish Foundation for Strategic Research through the Centre for Autonomous Systems. The funding is gratefully acknowledged.

2 Pose Estimation and Tracking

Closed-loop control of a robot where vision is used in the feedback loop is commonly referred to as *visual servoing*, [7]. Vision based tracking techniques incorporated in visual servoing systems usually facilitate model based or feature based approaches. A model based approach relies on a CAD model of the object to be tracked [3, 12] while feature based techniques use distinct image features: corners, lines, regions [6]. There are two categories of visual servoing: *position based servoing* and *image based servoing*, [7]. In our work, we use both approaches: position based approach is used for driving the arm and the hand to a pose generated by a simulator. Later, image based servo control is used to visually servo a grasped object to the desired pose.

By definition, a necessary requirement of a visual servo system is continuous tracking. Our tracking system is based on the ideas proposed in [3]. The general idea and the tracking system operation are presented in Fig. 2.

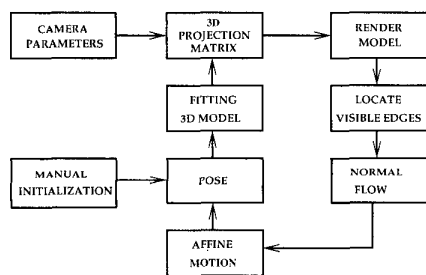


Figure 2: The tracking system.

Pose Estimation

To estimate the position and orientation of the object, $\Omega(\mathbf{R}, \mathbf{t})$, we use the approach proposed by Dementhon [2]. This step is followed by an extension of Lowe's [8] nonlinear approach proposed in [5]. The pose estimation step is called POSE in Fig. 2. The initialization of the tracking system is at this stage done manually. A number of corresponding points on the wire frame model and the image are chosen by the user, see Fig. 3(a) and Fig. 3(b).

Normal Displacement and Affine Motion

Using estimated pose, the model of the object is projected onto the image plane. For each visible edge, tracking nodes are assigned at regular intervals in image coordinates along the edge direction. After that, a search is performed for the maximum discontinuity in the intensity gradient along the normal direction to the edge. The edge normal is approximated with four

directions: 0, 45, 90, and 135 degrees. This way we obtain a displacement vector, $d_i^{\perp}(x_i, y_i) = [\Delta x_i \Delta y_i]^T$, representing the normal displacement field of visible edges. A 2D affine transformation is expressed:

$$\begin{bmatrix} x_i^{t+1} \\ y_i^{t+1} \end{bmatrix} = \begin{bmatrix} x & y & 0 & 0 & 1 & 0 \\ 0 & 0 & x & y & 0 & 1 \end{bmatrix} \Theta = \mathbf{A}(x_i, y_i) \Theta \quad (1)$$

where $\Theta = (a_1, a_2, a_3, a_4, T_x, T_y)^T$ represents the parameters of the affine model. There is a linear relationship between two consecutive images with respect to Θ :

$$\begin{aligned} d_i(x_i, y_i) &= \mathbf{A}(x_i, y_i) \Theta' & (2) \\ &= \mathbf{A}((x_i, y_i)) [\Theta - (1, 0, 0, 1, 0, 0)^T] & (3) \end{aligned}$$

From the previous equation follows:

$$d_i^{\perp} = \mathbf{n}_i^T \mathbf{d}(P_i) = \mathbf{n}_i^T \mathbf{A}(P_i) \Theta' \quad (4)$$

where \mathbf{n}_i is a unit vector orthogonal to the edge at a point P_i . From Eq. 4 we can estimate the parameters of the affine model, $\hat{\Theta}'$ using a M-estimator ρ :

$$\hat{\Theta}' = \arg \min_{\Theta'} \sum_i \rho(d_i^{\perp} - \mathbf{n}_i^T \mathbf{A}(P_i) \Theta') \quad (5)$$

Using the estimate of points in the image at time $t+1$, the POSE step is performed in order to obtain the pose of the object in the camera coordinate system, $\Omega(\mathbf{R}, \mathbf{t})_{init}$.

3D Tracking

Using the estimated affine parameters $\hat{\Theta}'$ and positions of edge nodes at time t , we are able to compute their positions at time $t+1$ from Eq. 1. As already mentioned, the affine model does not completely account for the 3D motion and perspective effects that occur during the object motion. Therefore, the pose space, $\Omega(\mathbf{R}, \mathbf{t})$, should be searched for a best fit given the image data. As proposed in [3], the projection of the object model is fitted to the spatial intensity gradients in the image, using $\Omega(\mathbf{R}, \mathbf{t})_{init}$ as the initialization:

$$\hat{\Omega}(\mathbf{R}, \mathbf{t}) = \arg \min_{\Omega} \left[- \sum_{C_{\Omega}} \|\nabla \mathbf{G}(t+1)\| \right] \quad (6)$$

where $\nabla \mathbf{G}(t+1)$ is the intensity gradient along the projected model edges and C_{Ω} are the visible edges of the model for given a pose Ω .

The optimization method uses a discrete hierarchical search with respect to the pose parameters. However, the correct discrete step of the pose parameters




Figure 3: The key components of the visual tracking system : left (a): the model with points used in POSIT to determine the pose, middle (b): pose estimation (the model is overlaid in black), and right (c): normal flow estimation (the lines represent the direction and not the magnitude).

is crucial in order to find the right minimum value. The affine motion model is particularly sensitive to the rotational motions of the object and large changes in rotation usually result in a loss of tracking. In addition to the original algorithm, we have implemented an adaptive search step determination based on the object's 3D velocity. The basic idea is to dynamically change the size of the search step instead of keeping it constant.

3 Grasping Simulator

GraspIt is a real-time, interactive simulator that allows the user to manipulate a model of an articulated robotic hand, perform grasps with the hand, and visualize stability analysis results on the fly. GraspIt has facilities for modeling a robotic arm and workcell to allow the user to examine the reachability constraints for a given grasping task as well. By consistently using accurate geometric models of all elements of the task, we ensure that grasps planned in the simulator can actually be executed in the real world. In this section we describe how a user interacts with the system and briefly describe some of the components that make up GraspIt (summarized in Fig. 4). Refer to reviews by Bicchi and Kumar [1] and Shimoga [11] for further information regarding grasping analysis.

First the user must decide which robotic hand and which object should be used in this grasping task. Our library of hands currently includes the Barrett Hand, the DLR hand (shown in Fig. 4), the NASA Robonaut hand, the Rutgers hand, and a simple parallel jaw gripper. While our experiments in this paper focus on the Barrett hand, since it is available to us in the lab, they could be replicated with any of the other hands. Furthermore, by using a simple configuration file that specifies a hand's kinematics and geometry,

we have kept the system flexible enough to import new hand designs with ease. An object model can be made from a CAD file, with additional information specifying the surface material and center of gravity. The visual tracking system uses a similar wire frame model representation, which can be derived from the CAD model (see figure 3(a)).

Once the hand and object have been chosen, the simulator constructs the models and allows the user to manipulate the hand model. If a connection is established with the object tracking module, the simulator sets the initial pose of the object in relation to the robot. At this point the user begins to plan the grasp by translating and rotating the wrist within the constraints of the Puma arm. Real-time collision detection prevents the hand model from penetrating any other world elements. After positioning the wrist, the user may manipulate the finger degrees of freedom individually, or perform an auto-grasp where the joints are closed at fixed velocities until contact occurs, providing a faster way to perform an initial grasp.

Each time a contact occurs between a link model and an object model, the system determines the contact area and places friction cones at the points of contact. The size of these cones is determined by the two contacting surface materials and a static Coulomb friction model.

The system will then analyze the grasp's stability on the fly, as each new contact occurs. Using a convex hull operation, the simulator determines the 6-dimensional space of forces and torques that can be applied by the grasp. This is also known as the grasp wrench space (GWS), and if the space includes the origin, then grasp is considered stable because it can resist any outside disturbance force by scaling the forces at the contact points. Several metrics have been proposed to quantify the grasp quality, and we currently




Figure 4: (Left) The internal components of GraspIt! and their functions. (Right) The DLR hand grasping a flask. The volumes in the small windows are projections of the grasp wrench space for this grasp.

compute two of them, although others could easily be added. The first is a worst case measure of the grasp's weakest point, and is defined by the distance from the origin to the closest facet on the boundary of the grasp wrench space. Indicators within the scene also show the direction of the worst case disturbance force and torque, and aid the user in finding this weakest point. The second measure is defined by the volume of the GWS and serves as more of an average case quality measure.

In addition to the numeric measures, the user is also presented with various 3-dimensional projections of the GWS to allow better visualization of a grasp's strong and weak points. This information allows the user to incrementally improve the stability of the grasp, and once satisfied, the user may chose to execute the grasp with an actual robotic hand.

4 Experimental Evaluation

To test the system we set up some simple grasping tasks using two different objects. The task involved determining the object's pose from vision, finding an appropriate and stable grasp for the robotic hand given the object's configuration, executing the grasp, and monitoring the grasped object's trajectory as it moved.

The dextrous robot hand used in this example is the Barret Hand. The details about the hand can be found in [9]. The hand is attached to a Puma560 arm which operates in a simple workcell. The vision system uses a standard CCD camera (Sony XC-77) with a focal length of 25mm and is calibrated with respect to the robot workspace. The camera is mounted on a tripod and views the robot and workcell from a 2m distance. Both the robot controller and the framegrabber are connected to a Sparc20 workstation which is linked via

a TCP socket connection to an SGI Indigo 2 running the GraspIt software.

The process of planning and executing a grasping task is as follows (refer also to the block diagram presented in Fig. 1):

1. After an image of the scene is obtained, interactive pose estimation of the object is performed as presented in Section 2.
2. A pose estimate is sent to the **GraspIt** grasp planning software which assists a user to generate the most suitable grasp.
3. Through **GraspIt** the actual arm and hand are controlled in order to perform the final grasp.
4. Once the object is grasped, the monitoring task begins. Image based visual servoing is used to place the object into the final desired pose. Through **GraspIt**, a few trajectory points (object poses) are manually generated. Each generated pose is used as an intermediate desired pose for the visual servoing routine.

Our two example grasps (Fig. 5) demonstrate steps 1-3 of the system (pose estimation, grasp planning and grasp execution). After the first image of the scene is acquired, a manual initialization step is performed as explained in Section 2, and the estimated pose of the object is sent to GraspIt, which aids the user in selecting an appropriate grasp. After a satisfactory grasp is planned, it is executed by controlling both the robot arm and the robot hand.

The first example shows a successfully planned and executed stable grasp of an L-shaped object. The grasp planned in the second example is not stable but requires precise estimation of the object pose so that the thumb can hook through the handle of the mug




Figure 5: [Row 1] Left: The vision system determines the pose of the L-shaped object within the robot workspace. Middle: A stable grasp planned within GraspIt. Right: The completed grasp. [Row 2] While the grasp planned in this example is not stable, it demonstrates the accuracy of the system. The camera is 2m away from the workspace, yet the thumb was able to enter the handle opening which is only 1.5 cm wider than it.

and it serves to demonstrate the accuracy of the system.

Once the object is grasped the monitoring task begins, and this is done using image based visual servoing to move the object to a desired position. While many existing servoing systems, [7], define the goal pose using a “teach-by-showing” method where the object is physically moved to the destination and feature extraction is performed, our system allows the user to specify the desired end position (as well as way points to define the trajectory) in the simulator and using knowledge of the camera parameters, it is straightforward to estimate the image positions of a number of the control points used to construct the error vector.

According to [7], a simple proportional control is used to guide the robot:

$$\mathbf{T}_G = g(\hat{\mathbf{J}}^T \mathbf{W} \hat{\mathbf{J}})^{-1} \hat{\mathbf{J}}^T \mathbf{W} (\mathbf{f} - \mathbf{f}_d) = g \hat{\mathbf{J}}^\dagger (\mathbf{f}_d - \mathbf{f}) \quad (7)$$

where \mathbf{T}_G is velocity screw of the robot, \mathbf{f}_d , \mathbf{f} are desired and current image features positions, $\hat{\mathbf{J}}$ is a model of image Jacobian, \mathbf{J} which is used in the control expression and \mathbf{W} is a symmetric positive matrix of rank 6 used to select some preferred points out of k available ones.

In this experiment we chose a highly stable grasp around the side of the block, (see Fig. 6). The flat palm of the Barrett hand allows for a strong planar contact on one side of the object and the planar finger surfaces form line contacts when closed around the edges of the opposite face. Looking at the projection

of the GWS into force space, shown in the middle of Fig. 6, we see that this grasp can apply strong forces along the horizontal plane but is weaker in the vertical direction since it relies purely on tangential friction forces. Having planned a stable grasp, we executed it from the simulator as before. Then we moved the simulated object to a desired destination pose, and sent this pose to the visual servoing system. Using the camera parameters, the corners of the object in its destination pose were projected to the image, and the servoing system successfully moved the object to that position as seen in Fig. 7.

5 Summary and Conclusion

We have implemented a robotic grasping system that integrates real-time vision with online grasp planning via a simulator to execute stable grasps on objects, and which can also monitor a grasped object’s trajectory for correctness. The system represents a step forward in autonomous control of tasks such as grasping.

Although the system is still in an initial phase, we believe it is an important step in integrating real-time visual control, grasp planning, and online task monitoring. All of the system components are still open research issues in the field of robotics and we will continue to improve the performance of each of them.

The visual tracking component has as its major goals improved robustness and increased speed. The current update rate of 10Hz is far from satisfactory.




Figure 6: (Left) A stable grasp of the L-shaped object. The axes are positioned at the object's center of gravity (Z axis points into the table). (Middle Top) The 6D grasp wrench space projected into torque space by setting the net force to 0. (Z axis points up). (Middle Bottom) The GWS projected into force space by setting the net torque to 0. (Right) The task is planned in GrasPlt by moving the object from its original location (empty wireframe) to its destination point and saving the new object pose.




Figure 7: Snapshots of the servoing near the beginning and end of the task. The object is being tracked and the corners of the destination pose are shown as white dots in the image.

To improve the robustness in the case of significant occlusions and natural background, we will continue to test the two implemented tracking techniques more extensively. In addition, none of the techniques incorporates any kind of velocity or acceleration prediction which can be easily added (e.g. a Kalman filter approach).

Currently, the grasp planning function is only semi-automated; the user needs to place the hand in proximity of the object and an automated finger contact procedure is instituted. We are exploring adding rules to properly synthesize a grasp given an object's pose and environmental constraints [4].

References

- [1] A. Bicchi and V. Kumar. Robotic grasping and contact: A review. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, pages 348–353, 2000.
- [2] D. Dementhon and L. S. Davis. Model-based object pose in 25 lines of code. *International Journal of Computer Vision*, 15:123–141, 1995.
- [3] P. E. Marchand and F. Chaumette. A 2D–3D model-based approach to real-time visual tracking. Technical report, INRIA, March, 2000.
- [4] K. Goldberg, E. Smith, K. Bohringer, and J. Craig. Computing parallel-jaw grip points. In *Proc of the IEEE Int. Conf. on Robotics and Automation*, 1999.
- [5] R. L. C. H. Araujo and C. Brown. A fully projective formulation for Lowe's tracking algorithm. Technical report 641, The University of Rochester, CS Department, Rochester, NY, November, 1996.
- [6] G. Hager and K. Toyama. The XVision system: A general-purpose substrate for portable real-time vision applications. *Computer Vision and Image Understanding*, 1(69):22–37.
- [7] S. Hutchinson, G. D. Hager, and P. I. Corke. A tutorial on visual servo control. *IEEE Transactions on Robotics and Automation*, 12(5):651–670, 1996.
- [8] D. G. Lowe. Robust model-based motion tracking through the integration of search and estimation. *International Journal of Computer Vision*, 8(2):113–122, 1992.
- [9] A. Miller and P. Allen. Examples of 3D grasp quality computations. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, pages 1240–1246, 1999.
- [10] A. Miller and P. Allen. GrasPlt!: A versatile simulator for grasping analysis. In S. Nair, editor, *Proc. of the ASME Dynamic Systems and Control Division*, 2000.
- [11] K. B. Shimoga. Robot grasp synthesis algorithms: A survey. *International Journal of Robotics Research*, 15(3):230–266, June 1996.
- [12] P. Wunsch and G. Hirzinger. Real-time visual tracking of 3D objects with dynamic handling of occlusion. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, pages 2868–2872, 1997.
- [13] B. Yoshimi and P. Allen. Visual control of grasping. In D. Kriegman, G. Hager, and S. Morse, editors, *The Confluence of Vision and Control*, pages 195–209. Springer-Verlag, 1998.