Proceedings of the 1996 IEEE
International Conference on Robotics and Automation
Minneapolis, Minnesota - April 1996

Transparent Grippers for Robot Vision *

Anton Nikolaev and Shree K. Nayar
Department of Computer Science
Columbia University, New York, NY 10027

Abstract

While existing grippers execute manipulation tasks, they occlude parts of the objects they grasp as well as parts of the workspace from vision sensors. We present the concept of a transparent gripper that enables vision sensors to image an object without occlusion while it is being manipulated. The physics of refraction, total internal reflection, lens effects, dispersion and transmittance are analyzed to determine the geometry and material properties of a practical transparent gripper. Algorithms are presented that compensate for image shifts caused by refraction. The experiments demonstrate the proposed gripper to be an effective solution to a variety of problems, including 3D object model generation.

1 Seeing While Grasping

A robot generally interacts with the physical world by means of a gripper mounted on its end-effector. The level of task complexity that the robot can handle is determined to a large extent by the gripper it is equipped with. For this reason, grippers and the process of grasping have been intensely studied in recent years [10][9][6]. This has led to the design and fabrication of a panoply of grippers that range from simple parallel-jaw ones to sophisticated anthropomorphic hands.

While the gripper determines the mechanical dexterity of a robot, sensory feedback determines the level of intelligence a robot can achieve. Vision serves as a powerful component of such a feedback system. It provides a richness of information that can enable a robot to handle uncertainties inherent to a task and react to varying environments.

As all grippers available today are made of materials that are opaque, a robot vision system has no way of observing an object in its entirety, while it is being manipulated or even simply held by the gripper. We attempt to bridge the gap between these two seemingly unrelated areas, namely, gripper design and visual sensing. We seek to develop a gripper that would enable a robot to "see" while it "grasps." Our solution is a transparent gripper that allows a vision sensor to see through to the object it holds.

A transparent gripper has, among others, two pertinent applications. First, it allows vision sensors to monitor an assembly or manipulation task in progress. This improves the perceptual capabilities of the robot system, and hence allows it to recover gracefully from errors.

The second application involves the task of automatic generation of 3D object models. In both machine vision as well as in computer graphics, there is a tangible need for efficient and robust way of recovering the 3D geometry of an object from a sample. Current approaches typically involve special purpose scanning devices [5]. Given that the shape of an object can be fairly complex, scanning it in a finite number of poses does not always suffice. In order to ensure complete shape recovery, the object has to be repositioned or regrasped during the process. This leads to cumbersome registration and calibration procedures required to fuse multiple depth maps.

A transparent gripper provides us with a flexible and fully automated way of recovering complete object geometry. The gripper can grasp an object just once and present it to a depth measurement system in any number of poses. Since end-effector transformations between consecutive poses are precisely known, no additional calibration is needed to construct a complete model from the multiple depth maps.

2 Physics of Transparent Grippers

Transparent materials exhibit a variety of optical phenomena that make the design of transparent grippers an interesting research problem. Here, we review the pertinent physical effects and discuss the considerations they raise in the development of a practical gripper.

The optical effect of primary interest to us is the

*This research was conducted at the Center for Research on Intelligent Systems, Department of Computer Science, Columbia University. It was supported in part by DOD/ONR MURI Grant N00014-95-1-0601 and in part by the NSF National Young Investigator Award.

refraction of light rays as they travel from one medium to another [2]. Refraction causes a light ray to bend at the interface between the two media by an amount determined by Snell's law:

$$\frac{\sin\alpha_1}{\sin\alpha_2} = \frac{c_1}{c_2} = n_{21}. \quad (1)$$

Here, α_1 and α_2 are the angles between the normal to the interface and the rays in the first and the second media, respectively, and c_1 and c_2 are the speeds of light in the two media. The ratio n_{21} is the *refractive index* of the second medium with respect to the first.

Physics handbooks [7][4] provide refractive indices of various materials measured at a given light wavelength with respect to a reference medium - typically air - whose refractive index is taken to be 1. Since light propagates slower in denser media, refractive indices of solids and liquids have values greater than 1. The table below gives indices of some common materials for the wavelength of 5893 Å at 20°C [4]:

Material	Refractive Index
Water	1.33-1.34
Ice	1.31
Glass	1.51-1.67
Celluloid	1.49-1.50
Benzine	1.50
Glycerine	1.47

If it were possible to build a completely invisible gripper, we would only need to concern ourselves with its grasping abilities. Unfortunately, a gripper made from a solid material cannot be absolutely invisible in air because its relative refractive index will be greater than 1. However, from the table above, we see that the ranges of refractive indices for liquids and plastics overlap. Hence, in principle, it is possible to select a combination of a plastic and a liquid with similar indices. Then, if we could immerse the plastic end-effector and the camera into a tank filled with the liquid, the grasped object would appear as if it were floating in space!

Alas, in most real applications, air is the medium of preference. Therefore, during the design of a transparent gripper, two properties need to be chosen with care: (a) the material of the gripper, and (b) its shape. To understand the tradeoffs involved, we describe the main optical artifacts caused by refraction.

2.1 Total Internal Reflection

According to Snell's law (1), in the case of light traveling from a medium of higher refractive index to one of lower index, the angle of refraction α_2 approaches $\frac{\pi}{2}$ at the *critical angle* of incidence given by

$$\alpha_{1crit} = \arcsin\left(\frac{1}{n_{12}}\right). \quad (2)$$

For larger incidence angles, Snell's law is not valid, and the interface reflects incoming light as a mirror. This can be observed in some of the vertical faces of the cube in Figure 1.

If a light ray attempts to emerge from within the gripper at an angle greater than the critical angle, the ray would be reflected back (once or even several times) before leaving the gripper. From the perspective of gripper design, total internal reflection is probably the most harmful factor as it can render parts of a gripper effectively opaque from certain viewing angles. From (2), we see that in order to have a large critical angle and minimize the chances of internal reflection, we need to keep the refractive index of the gripper as low as possible. Plastic grippers would generally have a higher critical angle than those made of glass, and thus would perform better. This is fortuitous since the brittleness of glass makes it undesirable from mechanical considerations.

2.2 Lens Effect

When the surface of the gripper is not planar, the surface normal varies over the interface. Then, if a parallel set of rays strikes the gripper, each ray hits the surface at a different angle. The refracted rays would either converge or diverge as they enter (or leave) the gripper, depending on whether the surface is convex or concave. In the case of a magnifying lens, this effect serves a useful purpose. In the context of transparent grippers, however, it could cause the image of a grasped object to be distorted and maybe even blurred, as in the case of the sphere in Figure 1. Further, the convergence of light rays can create large invisible volumes behind the gripper, as can be seen in the case of the hemisphere in Figure 1, whose periphery does not reveal the background pattern.

The lens effect is less pronounced for thinner sheets of refracting media. For instance, light bulbs are round but do not exhibit strong lens effects. Alternatively, a rarer material refracts less and hence would have less severe lens properties. It appears that we could alleviate refraction by using materials with lower refractive indices and by making the curved gripper thinner. However, when building a gripper, one has to take into account its mechanical properties. A thinner gripper would require the use of a stronger (more refractive) material. Likewise, the use of a less refractive (weaker) material would force us to make the gripper thicker, and hence more refracting. Therefore, we have a tradeoff between the gripper's optical and mechanical properties.

Grippers with flat faces will not encounter problems associated with the lens effect. However, such shapes tend to suffer more from total internal reflections. We see that the design of a transparent gripper involves

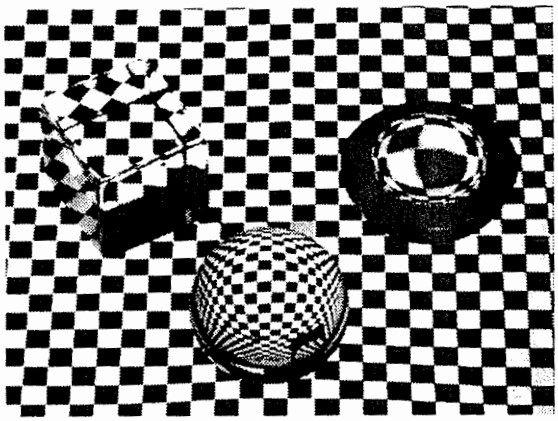

Figure 1: Transparent objects that illustrate total internal reflection (the cube), severe distortion (the sphere), and the problem of invisible volumes (the hemisphere).

several tradeoffs, making it a challenging problem.

2.3 Dispersion

If an incoming ray has components with different wavelengths (as in the case of white light), each wavelength is refracted by a different amount, causing the original ray to be spatially decomposed into a spectrum of monochromatic rays. This effect produces the spectacular colors of a rainbow when sunlight passes through rain drops in the air.

From (1), the relation between dispersion at a given wavelength λ and the angle of incidence α_1 is

$$\frac{\partial \alpha_2}{\partial \lambda} = -\frac{\sin \alpha_1}{n_{21}(\lambda)\sqrt{n_{21}^2(\lambda) - \sin^2 \alpha_1}} \frac{dn_{21}}{d\lambda}. \qquad (3)$$

As illustrated in Figure 2, the amount of dispersion increases with the angle of incidence. As a result, clearer images of the grasped object would be visible where the gripper surface is viewed head-on by the image sensor. Generally, dispersion is not a severe problem, since it is always possible to use color filters and treat each wavelength component in isolation. Nevertheless, materials that exhibit less light dispersion are obviously preferred. Many plastics have almost constant refractive indices within the visible spectrum [3]. Glass, on the other hand, is known to produce strong dispersion.

2.4 Energy Transport

At a smooth interface between two media, an incident light ray generally splits into two rays: the refracted ray discussed above and a reflected ray. Exactly how the incident energy is divided between the two depends on the relative refractive index of the media as

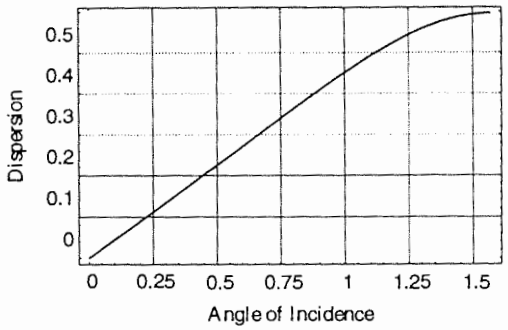

Figure 2: Absolute dispersion increases with the angle of incidence ($n_{21} = 1.5$). This implies that clearer images of the object would be seen for smaller angles of incidence.

well as the polarization of the incoming light. For linearly polarized light, the fraction of the incident energy that is reflected and refracted (*transmitted*), is given by Fresnel's equations [2]:

$$r_\perp = \left[\frac{\sin(\alpha_1 - \alpha_2)}{\sin(\alpha_1 + \alpha_2)}\right]^2, \; t_\perp = 1 - r_\perp, \qquad (4)$$

$$r_\parallel = \left[\frac{\tan(\alpha_1 - \alpha_2)}{\tan(\alpha_1 + \alpha_2)}\right]^2, \; t_\parallel = 1 - r_\parallel. \qquad (5)$$

Here, $r_\perp$ and $t_\perp$ are the reflected and refracted energies for incident light that is polarized perpendicular to the plane of incidence, and $r_\parallel$ and $t_\parallel$ are the energies for incident light with parallel polarization. In the case of unpolarized (for example, ambient) illumination, the fraction of the transmitted energy is simply $(t_\perp + t_\parallel)/2$. For a refractive index of $n_{21} = 1.5$ (glass to air), when the incident light is normal to the surface, 96% of the incident energy is transmitted through and only 4% is lost in reflection. Figure 3 shows a plot of the transmitted energy as a fraction of incident energy, plotted as a function of incidence angle. Finally, not all the transmitted energy passes through the gripper. A part of it is absorbed by the medium of the gripper.

2.5 Specularities

We would like the gripper to have low reflectivity in order to avoid specular reflections that cause undesirable highlights. Since reflectivity increases with the refractive index, materials with lower refractive indices are preferred (again, plastic over glass). Further, the amount of reflected light can be cut significantly by coating the transparent surfaces. In the case of vision sensors that use structured illumination, such as light stripe range finders, incident light may be polarized using a filter. Since specularly reflected light tends to preserve the incident polarization, a cross-polarized

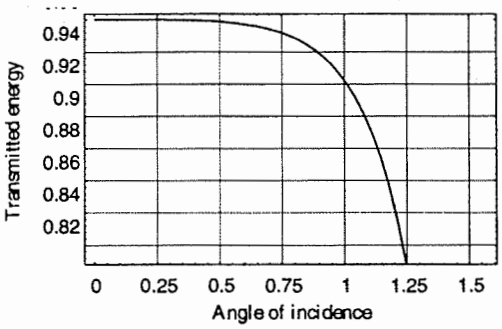

Figure 3: Percentage transmitted energy of unpolarized incident light plotted as a function of incidence angle ($n_{21} = 1.5$).

filter at the sensor end would help block out a large fraction of the reflected light [8].

3 An Example Gripper

The optimal gripper would be one whose shape and material properties minimize the undesirable effects described above, namely, total internal reflections, invisible volumes, lens distortion, dispersion, low transmittance, and specular reflections. Here, as an example, we show how our understanding of these effects can be exploited to design a gripper that performs quite well. This gripper is shown in Figure 4 holding an object (a coke can).

We have seen that certain plastics have desirable material properties. We have chosen the poly(methyl methacrylate) moulding powder (commercially available as Lucite [3][1]). Lucite has the following important properties: low refractive index of 1.49 (critical angle 42°), negligible dispersion, high transmittance, and low light absorption.

Our experiments were conducted using the Adept-1 5-DOF robot that comes with a pneumatic parallel-jaw mechanism. To match this setup, we chose two rectangular Lucite plates as the gripper components. The planar faces of the gripper avoid the lens effect, and any plane of light incident on the gripper undergoes only a constant shift. This makes it easy to compensate images for the effects of refractions. The small side faces of the two plates do produce total internal reflections, but these project onto relatively small image areas and hence can be made up for when multiple images of the object are taken. As is evident from Figure 4, the image of the area behind the gripper is shifted due to refraction.

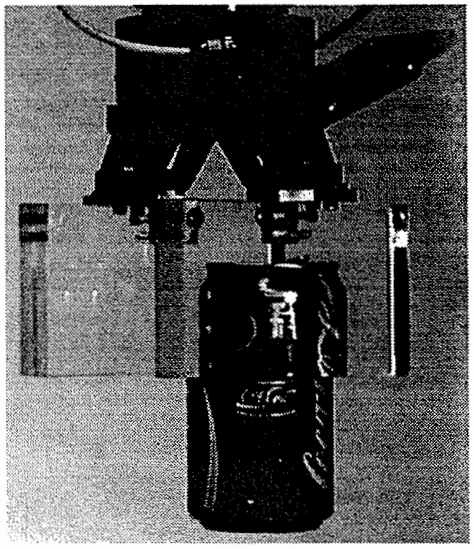

Figure 4: An object being held by a parallel-jaw transparent gripper. The image shift of the object area behind the gripper is caused by refraction.

4 Gripper-Camera Calibration

In order to determine which areas in the image are affected by refractions, we need to know the image projections of the two transparent parallel jaws of the gripper. The position and orientation of the gripper in the world coordinate system are always known since the gripper is attached to the end-effector. In our setup, the camera remains fixed, so in order to find the projection of the gripper onto the image plane we need to calibrate the camera with respect to the world frame.

If we neglect lens distortions in the imaging optics, and use the perspective projection model, the correspondence between the world and image coordinates becomes

$$\widetilde{x} = \frac{m_{11}x + m_{12}y + m_{13}z + m_{14}}{m_{41}x + m_{42}y + m_{43}z + 1}, \quad (6)$$

$$\widetilde{y} = \frac{m_{21}x + m_{22}y + m_{23}z + m_{24}}{m_{41}x + m_{42}y + m_{43}z + 1}, \quad (7)$$

where $\widetilde{x}$ and $\widetilde{y}$ are measured image plane coordinates, x, y, and z are world coordinates, and m_{ij} is the 4x4 projective transformation matrix [11] to be determined. m_{44} is set to 1 in order to get rid of the arbitrary scale factor.

Using 64 calibration points and solving the resulting overdetermined linear system (6)-(7), we were able to calibrate the camera to within 1 pixel accuracy.

Once the calibration is done, it is easy to project the 8 corners of each of the two gripper plates onto the image to find projection polygons corresponding to the

planar faces of the gripper. These polygons determine the regions of refraction in the image. Figure 5 shows the real-time display of the refraction polygons of one of the parallel jaws, overlaid on the image of the gripper.

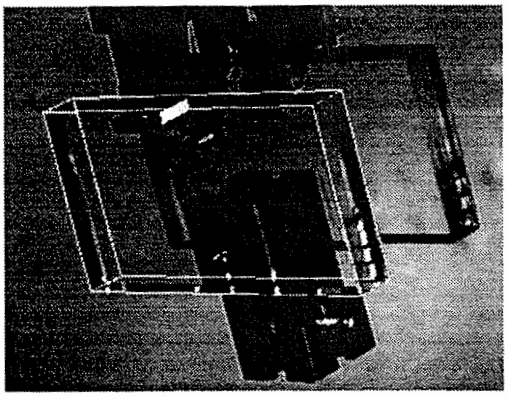

Figure 5: The camera-gripper calibration enables real-time display of the refraction polygons.

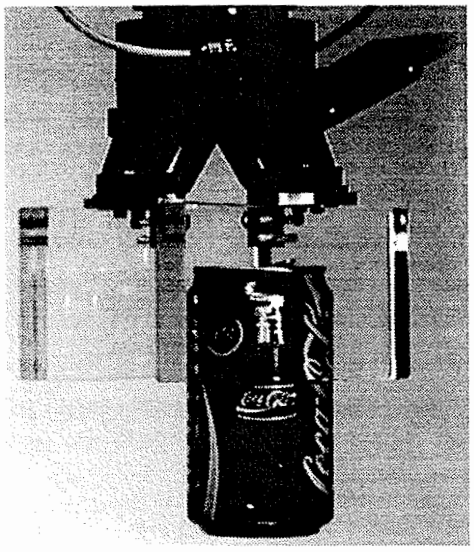

Figure 6: The image in Figure 4 after refraction compensation. The occluding boundary of the object is seen to have shifted back to its correct position.

5 Refraction Compensation

In the absence of total internal reflection, a flat slab of transparent material does not change the direction of light passing through it. Since the media involved in refraction at the entry into the slab and the point of exit are the same, but play reverse roles, the relative refractive indices are reciprocal to each other. Thus, the overall effect of refraction is no more than a parallel shift by the amount Δ, which depends on the refractive index n_{21}, the incidence angle α_1, and the thickness of the refracting slab h:

$$\Delta = \frac{h \sin(\alpha_1 - \alpha 2)}{\cos \alpha_2} = h \left(\sin \alpha_1 - \frac{\cos \alpha_1 \sin \alpha_1}{\sqrt{n_{21}^2 - \sin^2 \alpha_1}} \right). \quad (8)$$

From the above formula, it follows that a light plane with the unit normal $\{\alpha, \beta, \gamma\}$, passing through a flat transparent gripper, will change its equation as

$$\alpha x + \beta y + \gamma z + \delta = 0 \longrightarrow \alpha x + \beta y + \gamma z + (\delta \pm \Delta) = 0. \quad (9)$$

The sign in front of Δ depends on the sign of the dot product of the normals to the light plane and the plane of refraction. Since any straight line can be represented as a pair of plane equations, (9) can be used for both planes and lines.

Using (8) and (9) it is possible to restore an image in the case of orthographic or weak-perspective projection by shifting a portion of the image. Figure 6 shows the result of such a correction applied to the image in Figure 4.

6 3D Object Model Acquisition

In this section, we present our preliminary results on the application of the transparent gripper to the problem of 3D object model acquisition. The transparency of the gripper offers two possible approaches to the measurement of depth: passive or active techniques. Well-known passive depth estimation methods, such as stereo, structure from motion, and depth from defocus, can be applied with slight modifications to account for the refractions. Since it is desirable that the recovered shapes are accurate, we use an active sensor, namely, a light stripe range finder.

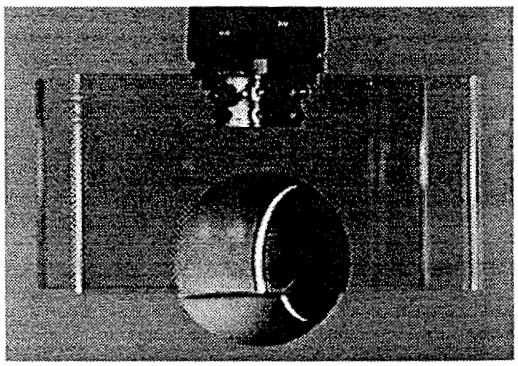

Figure 7: An object is actively scanned past a light stripe by the transparent gripper. The transparency of the gripper permits shape recovery of the entire object.

Only a single fixed stripe is used for recovery, since

the robot has sufficient degrees of freedom to show the object of interest to the light stripe in any desired orientation. In general, this enables the robot to show deep concavities and surface patches that suffer from self-occlusions to the sensor. This flexibility is difficult to get from scanning devices that rotate object about fixed trajectories.

For each bright (illuminated) point in the image, the world coordinates of the corresponding scene point can be determined by solving a system of three linear equations: the equation of the light plane and the two equations given by the straight line passing through the image point and the focal point of the camera. All these three equations need to be corrected for refraction using (8) and (9) if the scene point lies behind the transparent gripper.

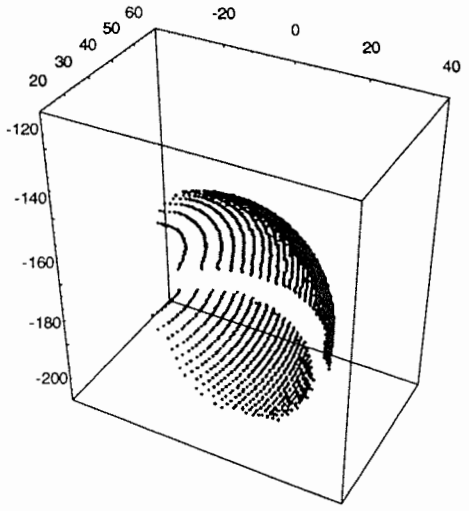

Figure 8: Depth map of the ball in Figure 7 computed without refraction compensation. As expected, refraction causes the depth map of the object region behind the gripper to be offset from the real map.

Figure 7 shows an object being actively scanned past a single light stripe by the transparent gripper. Bright image points that lie inside the refraction polygons in the image are corrected before they are used for depth computation. Figure 8 shows the depth map of the object recovered without refraction compensation and Figure 9 shows the depth map after correction. The computed map is seen to be very accurate.

7 Summary and Conclusions

The results reported here have opened up a number of interesting directions for future work. A natural extension would be the design of transparent grippers with more complex shapes as well as materials that

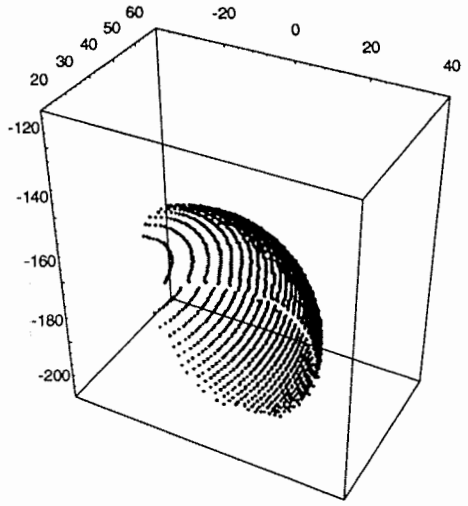

Figure 9: Depth map of the ball computed after refraction compensation.

yield even better optical properties. The vision sensors we have used thus far have been, in a sense, distant observers of the object. With a transparent gripper, one can attach a variety of sensors to its outer surfaces, looking into the manipulation workspace. These could vary from proximity sensors to regular CCD cameras.

References

[1] M. Ash and I.Ash, *Encyclopedia of plastics, polymers, and resins,* Chemical Publishing Co., 1982-1983.

[2] M.Born and E.Wolf, *Principles of Optics,* Pergamon, 1965.

[3] J.A.Brydson *Plastic materials, 6th ed.,* Butterworth-Heinemann, 1995.

[4] W.H.J.Childs *Physical constants, selected for students, 6th ed. rev.,* Methuen, 1951.

[5] B. Curless and M. Levoy, "Better Optical Triangulation through Spacetime Analysis," *Proc. of Intl. Conf. on Computer Vision*, pp. 987-994, June 1995.

[6] M. Cutkosky, "On grasp choice, grasp models, and the design of hands for manufacturing tasks," *IEEE Trans. on Robotics and Automation*, Vol 5, No. 3, June 1989.

[7] W.G. Driscoll, editor, *Handbook of Optics,* McGraw Hill Inc., 1978.

[8] S. Mersch, "Polarized Lighting for Machine Vision Applications," *Proc. of RI/SME Third Annual Applied Machine Vision Conf.*, pp. 40-54, February 1984.

[9] I. Kato, *Mechanical Hands Illustrated,* Hemisphere Publishing Corporation, 1982.

[10] G. Lundstrom, B. Glemme, and B. W. Rooks, *Industrial Robots - Gripper Review,* International Fluidics Services Ltd., 1977.

[11] R.Y.Tsai "An Efficient and Accurate Camera Calibration Technique for 3D Machine Vision," *IEEE Transactions on Robotics and Automation*, p. 364-374, No. 5, 1986.