

Department of Computer & Information Science

Technical Reports (CIS)

University of Pennsylvania

Year 1991

Observing A Moving Agent

Ruzena Bajcsy
University of Pennsylvania

Tarek Sobh
University of Pennsylvania

This paper is posted at ScholarlyCommons.
http://repository.upenn.edu/cis_reports/392

Observing A Moving Agent

**MS-CIS-91-01
GRASP LAB 247**

**Ruzena Bajcsy
Tarek Sobh**

**Department of Computer and Information Science
School of Engineering and Applied Science
University of Pennsylvania
Philadelphia, PA 19104-6389**

January 1991

Observing a Moving Agent

Ruzena Bajcsy and Tarek Sobh

GRASP Laboratory

Department of Computer and Information Science

School of Engineering and Applied Science

University of Pennsylvania, Philadelphia, PA 19104

Abstract

We address the problem of observing a moving agent. In particular, we propose a system for observing a manipulation process, where a robot hand manipulates an object. A discrete event dynamic system (DEDS) frame work is developed for the hand/object interaction over time and a stabilizing observer is constructed. Low-level modules are developed for recognizing the "events" that causes state transitions within the dynamic manipulation system. The work examines closely the possibilities for errors, mistakes and uncertainties in the manipulation system, observer construction process and event identification mechanisms. The system utilizes different tracking techniques in order to observe the task in an *active, adaptive and goal-directed* manner.

Keywords : Computer Vision, Control Theory, Decision Making, Discrete Event Dynamic Systems, Finite State Automata, Image Processing, Manufacturing Systems, Motion Estimation, Robotics.

Acknowledgements : This research was supported in part by Air Force AFOSR Grants 88-0244, 88-0296; Army/DAAL Grant 03-89-C-0031PRI; NSF Grants CISE/CDA 88-22719, IRI 89-06770; DARPA Grant N0014-88-0630 and DuPont Corporation.

Contents

1	Introduction	4
2	Discrete Event Dynamic Systems	6
2.1	What is a discrete event dynamic system ?	7
2.2	Modeling	7
2.3	Stability	10
2.3.1	Pre-Stability	10
2.3.2	Stability	11
2.3.3	Pre-Stabilizability	12
2.3.4	Stabilizability	13
2.4	Observability	13
2.4.1	Requirements	14
2.4.2	State Observability	14
2.5	Output Feedback Stabilizability	16
2.5.1	Requirements	16
2.5.2	Strong Output Stabilizability	17
3	Modeling and Observer Construction	18
3.1	Building the Model	19
3.1.1	A Grasping Task	20
3.1.2	A Screwing Task	23
3.2	Developing the Observer	24
4	Event Identification	26
4.1	Image Motion of the Hand	26
4.2	Recovering 3-D events	31
4.3	The Controllable Events	32
5	Sensor Uncertainties	33
5.1	Image Formation Errors	35
5.2	Calibration and Modeling Uncertainties	36

6 Image Processing Uncertainties	38
6.1 Edge Extraction Uncertainties	39
6.2 Computing 2-D Motion Uncertainty	40
7 Refining Image Motion	43
8 Recovering World Events	47
8.1 A 3-D Recovery Algorithm	48
8.1.1 Two-Frame Algorithm	49
8.1.2 Complexity Analysis	50
8.1.3 Observations	51
8.1.4 Multi-Frame Algorithm	51
8.2 Other Algorithms	54
8.3 Recovering 3-D Uncertainties	55
9 Refining World Events	57
10 Navigating the Observer Automaton	58
11 Results	60
12 Discussion	61
13 Extensions and Future Research	62
References	63

1 Introduction

The problem of observing a moving agent was addressed in the literature extensively. It was discussed in the work addressing tracking of targets and, determination of the optic flow [3,9,10,19,22,42], recovering 3-D motion parameters of different kinds of surfaces [14,16,27,33,34,40,41], and also in the context of other problems [2,6,7,8,37]. In this work we try to establish a framework for the general problem of observation, which may be applied to different kinds of visual tasks. We establish “intelligent” high-level control mechanisms for the observer in order to achieve an efficient approach to visually recognizing different processes within a specific dynamic system.

We concentrate on the problem of observing a manipulation process in order to illustrate the ideas and motive behind our framework. The process of observing a robot hand manipulating an object is very crucial for many robotic and manufacturing tasks. It is important to know in an automated manufacturing environment whether the robot hand is doing the correct sequence of operations on an object (or more than one object). It might be a fact that the workspace of the robotic manipulator cannot be accessed by humans, as in the case of some space applications or some areas within a nuclear plant, for example. In that case, having another robot “look” at the process is a very good option. Thus, the observation process can be thought of as a stage in a closed-loop fully automated system where there are robots who perform the required manipulation task and some other robots who observe them and correct their actions when something goes wrong. Typical manipulation processes include grasping, pushing, pulling, lifting, squeezing, screwing and unscrewing. Visual information from the observing robots can be the only kind of feedback, or it can be supplemented by other kinds, like tactile sensing. In this paper, we address the problem of observing a single hand manipulating a single object and “knowing” what is the hand doing, no feedback will be supplied to the manipulating robot to correct its actions.

To be able to observe how a hand manipulates an object, we must be able to identify how the hand moves and how the hand/object physical relationship evolves over time. An obvious way of doing this would be to identify the motion vectors as seen by the observer. In other words, identify the two-dimensional vectors in the observer’s camera plane and use these as a cue to know how the objects under consideration moves in the three-dimensional

space. The problems of recovering the image flow vectors (the two-dimensional motion vectors in the camera plane), and identifying the scene structure and motion have been key problems in computer vision. Many techniques have been developed for estimating the image flow [3,9,15,19,22], and to recover the three-dimensional world structure and motion [14,33,34,36,39,40,41]. Those techniques are not problem-oriented, they are not restricted to a particular problem domain, as is the case with our observer construction problem.

Trying to use the above techniques directly to solve our observer problem will not be efficient. In fact, possibly not feasible to perform in a practical way using the current technology, as the complexity of the manipulation process increases. Due to the fact that we probably know *a-priori* some information about the allowable (or useful) manipulation processes and the geometry of the robotic hand, posing the problem as a structure-from-motion vision procedure is a very naive way of modeling the observer system. It should also be noted that the observer will have to be an *active* one to be able to interact with the manipulation environment in such a way as to be able to “see” at all times. The idea of an active observer was discussed in the literature [2,6], and it was shown that an active observer can solve basic vision problems in a much more efficient way than a passive one.

We use a discrete event dynamic system as a high-level structuring technique to model the manipulation system. Our formulation uses the knowledge about the system and the different actions in order to solve the observer problem in an efficient, stable and practical way. The model incorporates different hand/object relationships and the possible errors in the manipulation actions. It also uses different tracking mechanisms so that the observer can keep track of the workspace of the manipulating robot. A frame work is developed for the hand/object interaction over time and a stabilizing observer is constructed. Low-level modules are developed for recognizing the “events” that causes state transitions within the dynamic manipulation system. The process uses a coarse quantization of the manipulation actions in order to attain an active, adaptive and goal-directed sensing mechanism.

The work examines closely the possibilities for errors, mistakes and uncertainties in the manipulation system, observer construction process and event identification mechanisms. We divide the problem into six major *levels* for developing uncertainty models in the observation process. The *sensor* level models deals with the problems in mapping 3-D features to pixel coordinates and the errors incurred in that process. We identify these uncertainties and

suggest a framework for modeling them. The next level is the *extraction strategy* level, in which we develop models for the possibility of errors in the low-level image processing modules used for identifying features that are to be used in computing the 2-D evolution of the scene under consideration and computing the image flow. In the third level, we utilize the geometric and mechanical properties of the hand and/or object to reject unrealistic estimates for 2-D movements that might have been obtained from the first two levels.

After having obtained 2-D models for the evolution of the hand/object relationship, we transform the 2-D uncertainty models into 3-D uncertainty models for the structure and motion of the entire scene. The fourth level uses the equations that govern the 2-D to 3-D relationship to perform the conversion. The fifth level rejects the improbable 3-D uncertainty models for motion and structure estimates by using the existing information about the geometric and mechanical properties of the moving components in the scene. The sixth and highest level is the DEDS formulation with uncertainties, in which state transitions and event identification is asserted according to the 3-D models of uncertainty that were developed in the previous levels.

We describe the automaton model of a discrete event dynamic system in the next section and then proceed to formulate our framework for the manipulation process and the observer construction. Then we develop efficient low-level event-identification mechanisms for determining different manipulation movements in the system and for moving the observer. Next, the uncertainty levels are described in details. Some results from testing the system is enclosed and future extensions to the system are discussed.

2 Discrete Event Dynamic Systems

In this section we present an overview for the development of a theory for discrete event dynamic systems (DEDS). Dynamic systems are usually modeled by finite state automata with partially observable events together with a mechanism for enabling and disabling a subset of state transitions [26,28,30,31]. We describe a recently developed framework for analyzing and controlling discrete event dynamic systems [28]. We propose that this model is a suitable framework for many vision and robotics tasks, in particular, we use the model as a high-level structuring technique for our system to observe a robot hand manipulating

an object. The approach used in this framework is a state space approach that focuses on controllability issues for DEDS. We consider the issues of stability, observability and stabilizability by output feedback within this framework.

2.1 What is a discrete event dynamic system ?

Discrete event dynamic systems (DEDS) are dynamic systems (typically asynchronous) in which state transitions are triggered by the occurrence of discrete events in the system. Many existing dynamic system have a DEDS structure, manufacturing systems and communication systems are just two of them. The state space approach in representing and analyzing such systems will probably lead to more applications that might be incorporated into the framework of DEDS. It will be assumed in the development of the state space approach of analyzing DEDS that some of the events in the system are *controllable*, i.e, can be enabled or disabled. The goal of controlling DEDS is to “guide” the behaviour of the system in a way that we consider “desirable”. It is further assumed that we are able to observe only a subset of the event, i.e, we can only *see* some of the events that are occurring in the system and not all. In some cases we will be forced to make decisions regarding the state of the system and how to control a DEDS based upon our observations only.

In the next subsection we will discuss the finite state model of a DEDS. This representation of a DEDS will be used the following subsections. This model will be a simple non-deterministic finite-space automaton. Graphical representations for DEDS automatons will be used as examples to explain the different definitions and ideas to be presented. The notions of stability for a DEDS will be introduced and discussed. We then focus on the questions of observability and state reconstruction from intermittent observations of the event trajectory. Further, we address the problem of stabilization by output feedback.

2.2 Modeling

The discrete event dynamic systems under consideration can always be modeled by a non-deterministic finite-state automata with partially observable and controllable events. In particular, one can make the distinction between classical automata theory [18,23,25,32] and our representation of DEDS in terms of the state transitions. In classical automata the events are inputs to the system, whereas in DEDS the events are assumed to be generated internally

by the system and the inputs to the system are the control signals that can enable or disable *some* of these events. We can represent our DEDS as the following quadruple :

$$G = (X, \Sigma, U, \Gamma)$$

where X is the finite set of states, Σ is the finite set of possible events, U is the set of admissible control inputs consisting of a specified collection of subsets of Σ , corresponding to the choices of sets of controllable events that can be enabled and $\Gamma \subseteq \Sigma$ is the set of observable events. Some functions can also be defined on our DEDS as follows :

- $d : X \rightarrow 2^\Sigma$
- $e : X \rightarrow 2^\Sigma$
- $f : X \times \Sigma \rightarrow 2^X$

where d is a set-valued function that specifies the set of possible events defined at each state, e is a set-valued function that specifies the set of events that cannot be disabled at each state, and f is the set-valued function that specifies state transitions from a state under different events. An output process can be formalized simply : whenever an event in Γ happens we see it, otherwise we don't see anything.

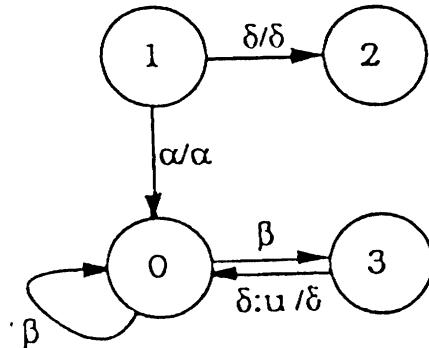


Figure 1 : A Simple DEDS Example

We can visualize the concept of DEDS by an example as in Figure 1, the graphical representation is quite similar to a classical finite automaton. Here, circles denote states, and events are represented by arcs. The first symbol in each arc label denotes the event, while the symbol following “/” denotes the corresponding output (if the event is observable). Finally, we

mark the controllable events by “:u”. Thus, in this example, $X = \{0, 1, 2, 3\}$, $\Sigma = \{\alpha, \beta, \delta\}$, $\Gamma = \{\alpha, \delta\}$, and δ is controllable at state 3 but not at state 1.

Also $d(1) = e(1) = \{\alpha, \delta\}$, $d(3) = \{\delta\}$, $e(3) = \phi$, $f(0, \beta) = \{0, 3\}$ etc. A transition, $x \rightarrow^\sigma y$, consists of a source state, $x \in X$, an event, $\sigma \in d(x)$, and a destination state, $y \in f(x, \sigma)$.

In general, a DEDS automaton A is a nondeterministic finite state automaton, however, if $f(x, \sigma)$ is single valued for each $x \in X$ then A can be termed as a deterministic finite state automaton. A finite string of states, $x = x_0 x_1 \dots x_j$ is termed a path or a state trajectory from x_0 if $x_{i+1} \in f(x_i, d(x_i))$ for all $i = 0 \dots j - 1$. Similarly, a finite string of events $s = \sigma_1 \sigma_2 \dots \sigma_j$ is termed an event trajectory from $x \in X$ if $\sigma_1 \in d(x)$ and $\sigma_{i+1} \in d(f(x, \sigma_1 \sigma_2 \dots \sigma_i))$ for all i , where we extend f to Σ^* via

$$f(x, \sigma_1 \sigma_2 \dots \sigma_i) = f(f(x, \sigma_1 \sigma_2 \dots \sigma_{i-1}), \sigma_i)$$

with $f(x, \epsilon) = x$. In our graphical example (Figure 1), $\alpha\beta\beta\delta$ is an event trajectory.

If we denote a transition labeled by σ by \rightarrow^σ , then we can similarly let \rightarrow^s denote a string of transitions s and \rightarrow^* denote any number of transitions, including no transitions. We can define the range of a state x by

$$R(A, x) = \{y \in X \mid x \rightarrow^* y\}$$

indicating the set of states that can be reached from x , we can also define the range of a subset of states Q in X by

$$R(A, Q) = \bigcup_{x \in Q} R(A, x)$$

An algorithm for computing $R(A, X_0)$ for any $X_0 \subset X$ that runs in $O(n)$ where $n = |X|$ can be easily formalized as follows :

Let $R_0 = Q_0 = X_0$ and iterate

$$R_{k+1} = R_k \cup f(Q_k, \Sigma)$$

$$Q_{k+1} = R_{k+1} \cap \overline{R_k}$$

Terminate when $R_{k+1} = R_k$. Then, $R(A, X_0) = R_k$.

A state $x \in X$ is *alive* if $d(y) \neq \phi$ for all $y \in R(A, x)$. A subset Y of X is termed a *live set* if all $x \in Y$ are alive. A system A is termed *alive* if X is a *live set*.

2.3 Stability

In this section we discuss the notions of stability and the possibility of stabilizing a discrete event dynamic system. In particular, we are going to concentrate on stability notions with respect to the *states* of a DEDS automaton. Assuming that we have identified the set of “good” states, E , that we would like our DEDS to “stay within” or do not stay outside for an infinite time, the problem would reduce to :

- Checking out whether all trajectories from the other states will visit E infinitely often.
- Trying to “guide” the system using the controllable events in a way such that the system will visit the “good” states infinitely often.

We shall start by defining and testing for different notions of stability and then discuss ways to stabilize a system. We shall start by assuming that the DEDS model under consideration is an uncontrolled system with perfect knowledge of the state and event trajectories ($\Sigma \cap \bar{\Gamma} = \phi$), to simplify developing the definitions and examples.

2.3.1 Pre-Stability

To capture the idea of stability, we can suppose that we have already identified a subset of states E in X that returning to E implies being in a position to continue desired behaviour from that point on. We can define the notion of a state in the DEDS being stable with respect to E in two stages. The first stage will be the weaker notion and will be termed pre-stability. We say that $x \in X$ is pre-stable if *all* paths from x can go to E in a finite number of transitions, i.e, no path from x ends up in a cycle that does not go through E .

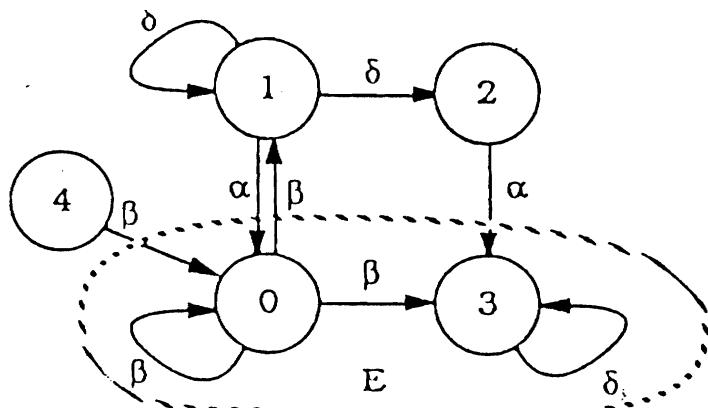


Figure 2 : Stability Example

In Figure 2, states 0, 2, 3, and 4 are pre-stable, since all transitions from them can goto $\{0, 3\}$ in a finite number of transitions. State 1 is not pre-stable since it will stay forever outside E if an infinitely long string of δ 's occurs. A definition of pre-stability can be formalized as follows :

Given a live system A and some $E \subset X$, a state $x \in X$ is pre-stable with respect to E (or E -pre-stable) if for all $x \in \mathcal{X}(A, x)$ such that $|x| \geq n$, there exists $y \in x$ such that $y \in E$. We say that a set of states is E -pre-stable if all its elements are E -pre-stable and a system A is pre-stable if X is E -pre-stable.

The restriction for liveness can be flexible in the sense that if all the dead states are within E , then an automaton might still be E -pre-stable. It follows from the above definition that a state $x \in X$ is E -pre-stable iff $x \in E$ or $f(x, d(x))$ is E -pre-stable. The following algorithm computes the maximal E -pre-stable set X_p within a system :

Let $X_0 = E$ and iterate :

$$X_{k+1} = \{x | f(x, d(x)) \subset X_k\} \cup X_k$$

Terminate when $X_{k+1} = X_k$, then $X_p = X_k$.

In Figure 2, it can be noticed that $X_1 = X_2 = X_p = \{0, 2, 3, 4\}$.

2.3.2 Stability

The stronger notion of stability corresponds to returning to the set of “good” states E in a finite number of transitions following any excursion outside of E . Thus, given E , we define a state $x \in X$ to be E -stable if all paths go through E in a finite number of transitions and then visit E infinitely often. As an example, in Figure 2, where $E = \{0, 3\}$, only 2 and 3 are stable states. State 1 is not stable since the system can loop at 1 infinitely. State 0 although in E is not stable since the system can make a transition to 1 and then stays there forever, the same applies to state 4. We can use the previously defined notion of pre-stability and define a state to be E -stable if all the states in its reach are E -pre-stable. In Figure 2, 0 and 4 are not E -stable since they can reach 1, which is not E -pre-stable. We can define stability as follows :

Given a live A and $x \in X$, x is E -stable iff $R(A, x)$ is E -pre-stable. A $Q \subset X$ is stable if all $x \in Q$ are stable. A system A is stable if X is a stable set, from which we can conjecture

that A is E -stable iff it is also E -pre-stable.

2.3.3 Pre-Stabilizability

Now, we introduce control and reconsider the stability notions discussed before. We try to “guide” our system or some states of it to behave in a way that we consider desirable. Pre-stabilizability is described as finding a state feedback such that the closed loop system is pre-stable. We can then define pre-stabilizability formally as follows :

Given a live system A and some $E \subset X$, $x \in X$ is pre-stabilizable with respect to E (or E -pre-stabilizable) if there exists a state feedback K such that x is alive and E -pre-stable in A_K . A set of states, Q , is a pre-stabilizable set if there exists a feedback law $K(s)$ (A control pattern) so that every $x \in Q$ is alive and pre-stable in A_K , and A is a pre-stabilizable system if X is a pre-stabilizable set.

As an example, in Figure 3, state 1 is pre-stabilizable since disabling γ pre-stabilizes 1. However, disabling γ at state 2 leaves no other defined events at 2 and “kills” it, so neither state 2 or 3 is pre-stabilizable.

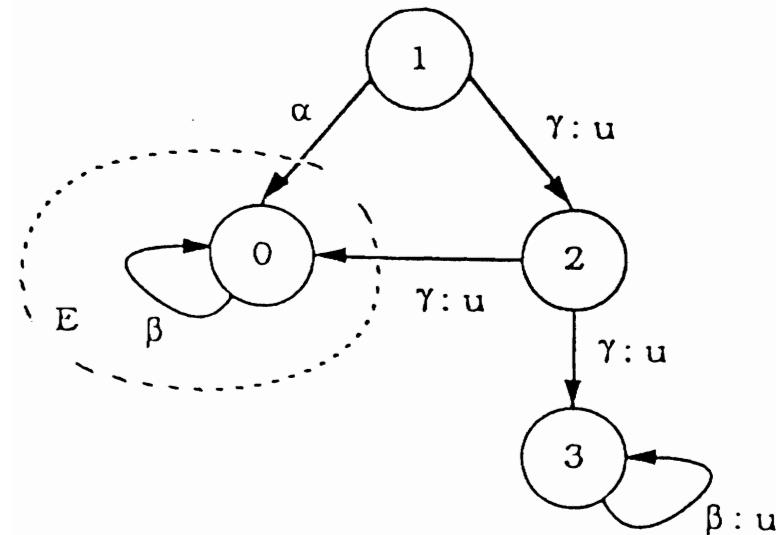


Figure 3 : Pre-Stabilizability Example

2.3.4 Stabilizability

Stabilizability is an extension of pre-stabilizability. Stabilizability is described as finding a state feedback such that the closed loop system is stable. We can then define stabilizability formally as follows :

Given a live system A and some $E \subset X$, $x \in X$ is stabilizable with respect to E (or E -stabilizable) if there exists a state feedback K such that x is alive and E -stable in A_K . A set of states, Q , is a stabilizable set if there exists a feedback law $K(s)$ (a control pattern) so that every $x \in Q$ is alive and stable in A_K , and A is a stabilizable system if X is a stabilizable set.

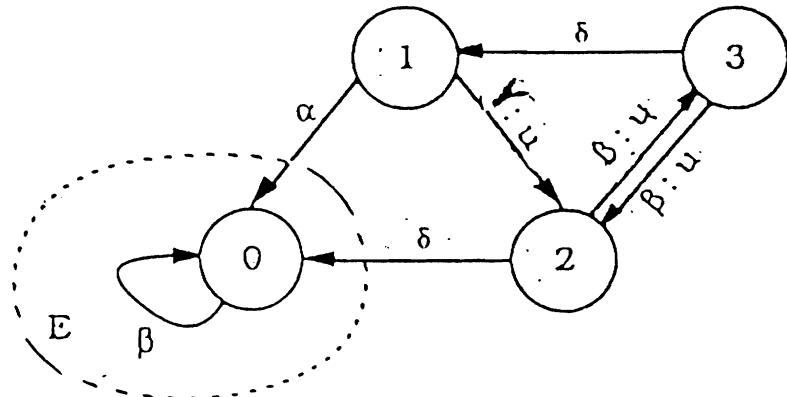


Figure 4 : Stabilizability Example

In Figure 4, disabling β at state 2 is sufficient to make the whole system stable with respect to state 0. Disabling γ at state 1 will help stabilize only state 1, because the system can then continue looping between states 2 and 3. Disabling β at state 3 will not help stabilize or pre-stabilize any state.

2.4 Observability

In this section we address the problem of determining the current state of the system. In particular, we are interested in observing a certain sequence of *observable* events and making a decision regarding the state that the DEDS automaton A might possibly be in. In our definition of observability, we visualize an intermittent observation model, no direct measurements of the state are made, the events we observe are only those that are in $\Gamma \subset \Sigma$, we will not observe events in $\Sigma \cap \bar{\Gamma}$ and will not even know that any of which has occurred. State

ambiguities are allowed to develop (which must happen if $\Sigma \neq \Gamma$) but they are required to be resolvable after a *bounded* interval of events. This notion of observability can be illustrated graphically as in Figure 5.

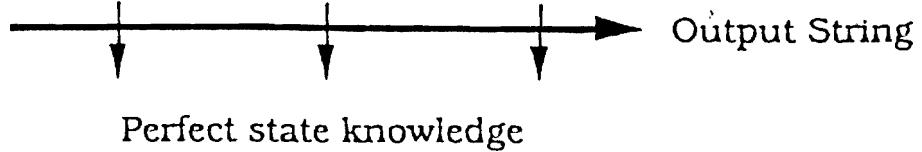


Figure 5 : Notion of Observability

2.4.1 Requirements

In developing the theory and examples we shall concentrate on uncontrolled models of DEDS automata with partial knowledge of the event trajectory. Due to the fact that we are “seeing” only observable events in Γ in our system, it is not desirable to have our automaton generate arbitrarily long sequences of unobservable events in $\Sigma \cap \bar{\Gamma}$. A necessary condition to guarantee this is that the automaton after removing the observable events $A|\bar{\Gamma}$, must not be alive. In fact, it is also essential that every trajectory in $A|\bar{\Gamma}$ is killed in finite time by being forced into a dead state. It can be seen that the condition for a DEDS automaton to be unable to generate arbitrarily long sequences of unobservable events, is that $A|\bar{\Gamma}$ must be *D*-stable, where *D* is the set of states that only have observable events defined (i.e, $D = \{x \in X | d(x) \cap \bar{\Gamma}\}$).

2.4.2 State Observability

As illustrated in Figure 5, a DEDS is termed *observable* if we can use the observation sequence to determine the current state exactly at intermittent points in time separated by a *bounded* number of events. More formally, taking any sufficiently long string, s , that can be generated from any initial state x . For any observable system, we can then find a prefix p of s such that p takes x to a *unique* state y and the length of the remaining suffix is bounded by some integer n_o . Also, for any other string t , from some initial state x' , such that t has the same

output string as p , we require that t takes x' to the same, unique state y .

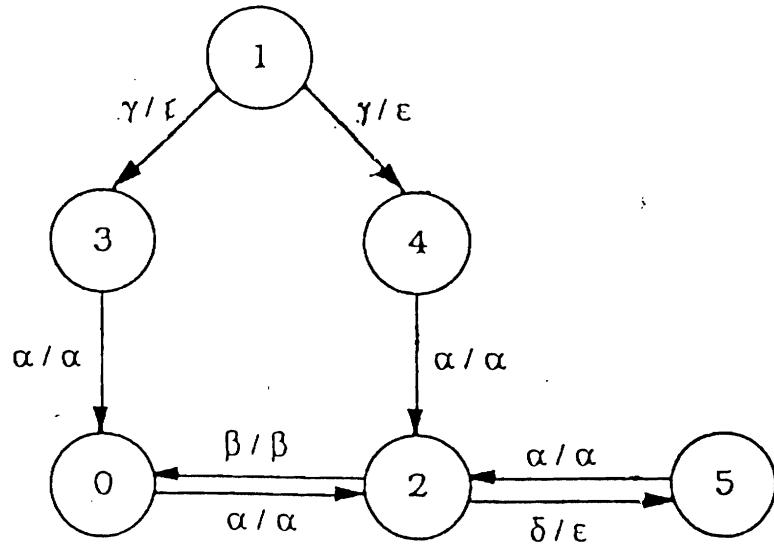


Figure 6.1 : A Simple System

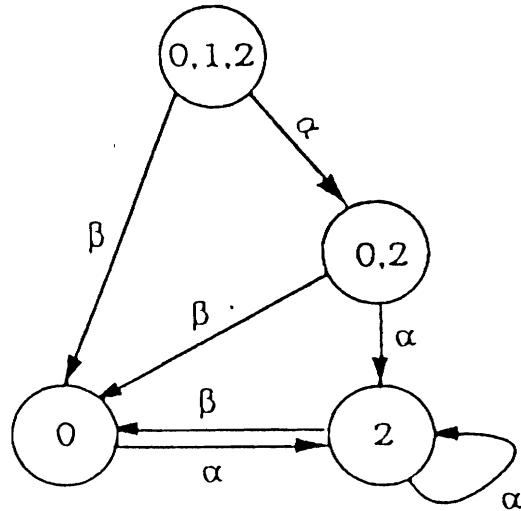


Figure 6.2 : Observer for the System in Figure 6.1

In Figures 6.1 and 6.2 a simple system and its observer are illustrated. It can be seen that the observer will never know when the system will be in states 3, 4 or 5, since the events that take the system to those states are unobservable (δ/ϵ means that $\delta \in \Sigma \cap \bar{\Gamma}$), namely δ and γ . There are two states in the observer which are ambiguous, however, another two states are singleton states, i.e, when our observer reaches them, we'll know the exact state

that the DEDS is currently in. Had it been the case that our observer could, for example, loop forever in ambiguous states, then the DEDS would be unobservable. This leads to the following formal definition of observability that ties it with the notion of stability :

A DEDS automaton A is observable iff E is nonempty and O is E -stable, where O is the observer for A and E is the set of singleton states of O . It can be seen that the observer in Figure 6.2 is stable with respect to the nonempty subset of states $\{0, 2\}$ and thus the DEDS of Figure 6.1 is observable.

2.5 Output Feedback Stabilizability

In this section we combine the ideas discussed in the previous two subsections regarding observability and stability to address the problem of stabilization by dynamic output feedback under partial observations. In this section we concentrate on partially controlled systems with partial knowledge of the event trajectory. In particular, our goal is to develop stabilizing compensators by cascading and a stabilizing state feedback defined on the observer's state space.

2.5.1 Requirements

To attack the problem of output feedback stabilization, it should be noticed that we are actually trying to “manipulate” the system’s observer, in other words, what we have available in a sequence of observable events (the system’s *output*) and we are trying to use this output to control the behaviour of the system using *only* the events that we can control. It is then possible to redefine the problem of output feedback stabilization as the stabilization of the observer by state feedback.

The obvious notion of output E -stabilizability (stabilizability with respect to $E \subset X$) is the existence of a compensator C so that the closed-loop system A_C is E -stable. It is possible that such a stabilizing compensator exists, such that we are sure that the system passes through the subset E infinitely often (E -stable) *but* we never know when the system is in E . A stronger notion of output feedback stabilizability would not only requires that the system passes through subset E infinitely often, but also that we regularly know when the system is in E . In our example and discussion we shall concentrate on this stronger notion of output

stabilizability.

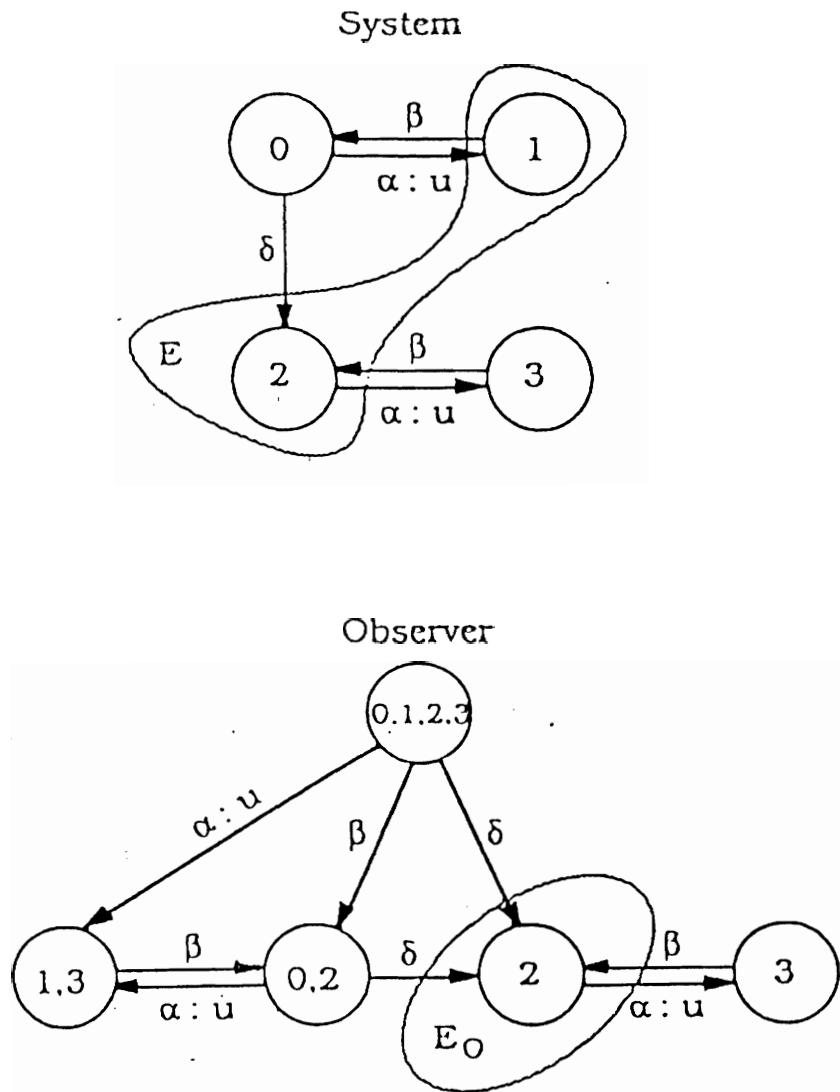


Figure 7 : Example for Output Stabilizability

2.5.2 Strong Output Stabilizability

The basic idea behind strong output stabilizability is that we will know that the system is in state E iff the observer state is a subset of E . The fact that the observer state should be a subset of E instead of having the observer state of interest *includes* states in E is because we want to *guarantee* that our system is within E . Our compensator should then *force* the observer to a state corresponding to a subset of E at intervals of at most a finite integer i

observable transitions. We can then formalize the notion of a strongly output stabilizable system as follows :

A is strongly output E -stabilizable if there exists a state feedback K for the observer O such that O_K is stable with respect to $E_O = \{ \hat{x} \in Z \mid \hat{x} \subset E \}$.

where Z is the set of states of the observer.

As an example, considering the DEDS and its observer in Figure 7, where $E = \{1, 2\}$, we have to check the observer stability (or stabilize the observer) with respect to E_O , because this is the only observer state that is a subset of E . As a start, we do not know which state is our system in (as denoted by the state $\{0, 1, 2, 3\}$), however, using the observer transitions we can see that to achieve E_O -stability for the observer we only need to disable α at the observer state $\{0, 2\}$. It should be noted that all the events are observable in this DEDS automaton.

3 Modeling and Observer Construction

Manipulation actions can be modeled efficiently within a discrete event dynamic system framework. It should be noted that we do not intend to *discretize* the workspace of the manipulating robot hand or the movement of the hand, we are merely using the DEDS model as a high level structuring technique to preserve and make use of the information we know about the way in which each manipulation task should be performed, in addition to the knowledge about the physical limitations of both the observer and manipulating robots. We avoid the excessive use of decision structures and exhaustive searches when observing the 3-D world motion and structure.

A bare-bone approach to solving the observation problem would have been to try and visually reconstruct the full 3-D motion parameters of the robot's hand, which would have more than six degrees of freedom, depending on the number of fingers and/or claws and how they move. The object's motion should also be recovered in 3-D, which is complicated especially if it is a non-rigid body. That process should be done in real time while the task is being performed. A simple way of tracking might be to try and keep a fixed geometric relationship between the observer camera and the hand over time. However, the above formulation is inefficient, not needed and for all practical purposes infeasible to compute in real time. The

limitation of the observer reachability and the extensive computations required to perform the visual processing are motives behind formulating the problem as a hierarchy of task-oriented observation modules that exploits the higher-level knowledge about the existing system, in order to achieve a feasible mechanism of keeping the visual process under supervision.

We do a coarse quantization of the visual manipulation actions which has both continuous and discrete aspects of manipulation dynamics. State transitions within the manipulation domain are asserted according to probabilistic models that determine at different instances of time whether the visual scene under inspection has changed its state within the discrete event dynamic system state space. We next discuss building the manipulation model for two simple tasks, grasping and screwing, then we proceed to develop the observer for these tasks. Formulating the uncertainty models for the state transitions and the inter-state continuous dynamics will be left for the sections that deal with the different uncertainty levels and event identification mechanisms.

3.1 Building the Model

The ultimate goal of the observation mechanism is to be able to know at all (or most) of the time what is the current manipulation process and what is the visual relationship between the hand and the object. It should be noticed that this concept is very similar to the concept of *observability* as defined in the previous section for general DEDS. The fact that the observer will have to move in order to keep track of the manipulation process, makes one think of the *output feedback stabilizability* principle for general DEDS as a model for the tracking technique that has to be performed by the observer's camera.

In real-world applications, many manipulation tasks are performed by robots, including, but not limited to, lifting, pushing, pulling, grasping, squeezing, screwing and unscrewing of machine parts. Modeling all the possible tasks and also the possible order in which they are to be performed is possible to do within a DEDS state model. The different hand/object visual relationships for different tasks can be modeled as the set of states X . Movements of the hand and object, either as 2-D or 3-D motion vectors, and the positions of the hand within the image frame of the observer's camera can be thought of as the events' set Γ that causes state transitions within the manipulation process. Assuming, for the time being, that we have no direct control over the manipulation process itself, we can define the set of admissible control

inputs U as the possible tracking actions that can be performed by the hand holding the camera, which actually can alter the visual configuration of the manipulation process (with respect to the observer's camera). Further, we can define a set of "good" states, where the visual configuration of the manipulation process enables the camera to keep track and to know the movements in the system. Thus, it can be seen that the problem of observing the robot reduces to the problem of forming an output stabilizing observer for the system under consideration, which was discussed in details in the previous section.

It should be noted that a DEDS representation for a manipulation task is by no means unique, in fact, the degree of efficiency depends on the person who builds the model for the task, testing the optimality of a manipulation models is an issue that is to be addressed in the future. Automating the process of building a model is another issue that will have to be addressed later. As the observer identifies the current state of a manipulation task in a non ambiguous manner, it can then start using a practical and efficient way to determine the next state within a predefined set, and consequently perform necessary tracking actions to stabilize the observation process with respect to the set of good states. That is, the current state of the system tells the observer what to *look for* in the next step.

3.1.1 A Grasping Task

We present a simple model for a grasping task. The model is that of a gripper approaching an object and grasping it. The task domain was chosen for simplifying the idea of building a model for a manipulation task. It is obvious that more complicated models for grasping or other tasks can be built. The example shown here is for illustration purposes.

As shown in Figure 8, the model represents a view of the hand at state 1, with no object in sight, at state 2, the object starts to appear, at state 3, the object is in the claws of the gripper and at state 4, the claws of the gripper close on the object. The view as presented in the figure is a frontal view with respect to the camera image plane, however, the hand can assume any 3-D orientation as so long as the claws of the gripper are within sight of the observer, for example, in the case of grasping an object resting on a tilted planar surface. This demonstrates the continuous dynamics aspects of the system. In other words, different orientations for the approaching hand are allowable and observable. State changes occur only when the object appear in sight or when the hand encloses it. The frontal upright view

is used to facilitate drawing the automaton only.

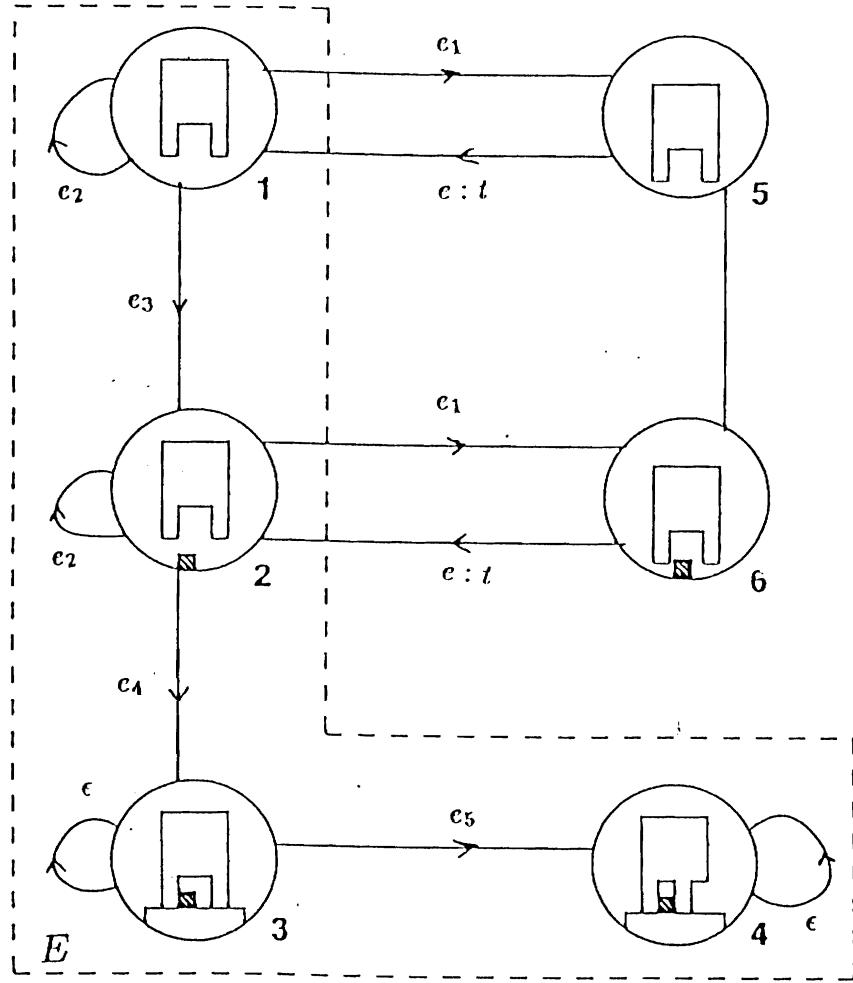


Figure 8 : A Model for a Grasping Task

It should be noted that these states can be considered as the set of good states E , since these states are the expected different visual configurations of a hand and object within a grasping task. States 5 and 6 represent instability in the system as they describe the situation where the hand is not centered with respect to the camera imaging plane, in other words, the hand and/or object are not in a good visual position with respect to the observer as they tend to escape the camera view. These states are considered as “bad” states as the system will go into a non-visual state unless we correct the viewing position. The set $X = \{1, 2, 3, 4, 5, 6\}$ is the finite set of states, the set $E = \{1, 2, 3, 4\}$ is the set of “good” states.

The events are defined as motion vectors or motion vector probability distributions, as will

be described later, that causes state transitions and as the appearance of the object into the viewed scene. The transition from state 1 to state 2 is caused by the appearance of the object. The transition from state 2 to state 3 is caused by the event that the hand has enclosed the object, while the transition from state 3 to state 4 is caused by the inward movement of the gripper claws. The transition from the set $\{1, 2\}$ to the set $\{5, 6\}$ is caused by movement of the hand as it escapes the camera view or by the increase in depth between the camera and the viewed scene, that is, the hand moving far away from the camera. The self loops are caused by either the stationarity of the scene with respect to the viewer or by the continuous movement of the hand as it changes orientation but without tending to escape a good viewing position of the observer. In the next section we discuss different techniques to identify the events. The controllable events denoted by “ $: t$ ” are the tracking actions required by the hand holding the camera to compensate for the observed motion. Tracking techniques will later be addressed in detail. All the events in this automaton are observable and thus the system can be represented by the triple $G = (X, \Sigma, T)$, where X is the finite set of states, Σ is the finite set of possible events and T is the set of admissible tracking actions or controllable events.

It should be mentioned that this model of a grasping task could be extended to allow for error detection and recovery. Also search states could be added in order to “look” for the hand if it is nowhere in sight. The purpose of constructing the system is to develop an observer for the automaton which will enable us to determine the current state of the system at intermittent points in time and further more, enable us to use the sequence of events and control to “guide” the observer into the set of good states E and thus stabilize the observation process. Disabling the tracking events will obviously make the system neither stable or pre-stable with respect to the set $E = \{1, 2, 3, 4\}$, however, it should be noted that the subset $\{3, 4\}$ is already stable with respect to E regardless of the tracking actions, that is, once the system is in state 3 or 4, it will remain in E (as defined by our formulation of the model). The whole system is stabilizable w.r.t. E , enabling the tracking events will cause all the paths from any state to go through E in a finite number of transitions and then will visit E infinitely often.

3.1.2 A Screwing Task

The next model we present is one for a simple screwing task. The task is that of a gripper screwing an object (a nail for example). It is assumed that the claws of the gripper already encloses the nail and that contact is maintained throughout the process, the rotation is allowed to be either clockwise or anticlockwise.

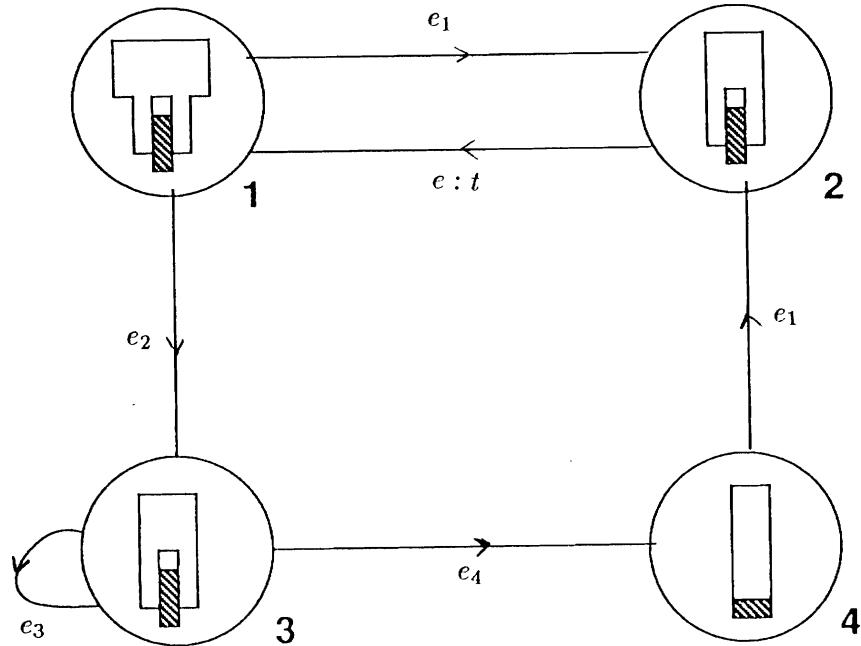


Figure 9 : A Model for a Screwing Task

As shown in Figure 9, the model represents a frontal view of the hand at state 1, with the object between the claws, the hand starts to rotate at state 2 and 3 with some view of the claws and the object still in sight and the claws are occluded at state 4 which represents a side view of the gripper. This specific visual representation was chosen because of the fact that transitions between states 1 and 3 and the self loop at 3 cannot be compensated by a tracking action due to the physical limitations of the tracking arm, in other words, the observing robot might not be able to do 360 degrees rotations around the manipulating hand, especially if the workspaces of both robots do not intersect and both are fixed, non-mobile robots. As mentioned before, the frontal upright view with respect to the camera imaging plane in state one was chosen only to facilitate drawing the automaton. The hand can assume any 3-D orientation as so long as the claws in states 1, 2 and 3 are within sight of the observer, for example, in the case of screwing a nail into a tilted wall.

As shown by our model, the automaton tends to keep the frontal view of the hand as long as possible (as far as the observer robot can rotate), after that the observer will just have to sit idle until rotation of the hand is trackable again. If one define the stable visual state as state 1, then obviously the system cannot be made stable with respect to that state, however, one can think of a screwing action on the whole as a stable set, since the robot hand is always within sight of the observer and it does not tend to escape the viewing field. In that case the set of “good” states E is the same as the set $X = \{1, 2, 3, 4\}$, the finite set of states. The goal of the observer in that case would basically be trying to keep a frontal view as long as it can.

The event e_1 can be defined as rotations that the observer robot can track and keep a frontal position of the hand, while e_2 is the one that makes the observable robot reach its “limit” position where it cannot rotate around the hand in the same direction any longer. The rotations e_3 are the untrackable rotations, which lie beyond the reachable workspace of the observable robot. The event e_4 can be defined as the event that causes the visual scene to be a side view of the gripper.

3.2 Developing the Observer

In order to know the current state of the manipulation process we need to observe the sequence of events occurring in the system and make decisions regarding the state of the automaton, state ambiguities are allowed to occur, however, they are required to be resolvable after a *bounded* interval of events. An observer, as defined in the previous section, have to be constructed according to the visual system for which we developed a DEDS model. The goal will be to make the system a strongly output stabilizable one and/or construct an observer to satisfy specific task-oriented visual requirements that the user may specify depending on the nature of the process. It should be noticed that events can be asserted with a specific probability as will be described in the sections to come and thus state transitions can be made according to pre-specified thresholds that compliments each state definition. In the case of developing ambiguities in determining current and future states, the history of evolution of past event probabilities can be used to navigate backwards in the observer automaton till a

strong match is perceived, a fail state is reached or the initial ambiguity is asserted.

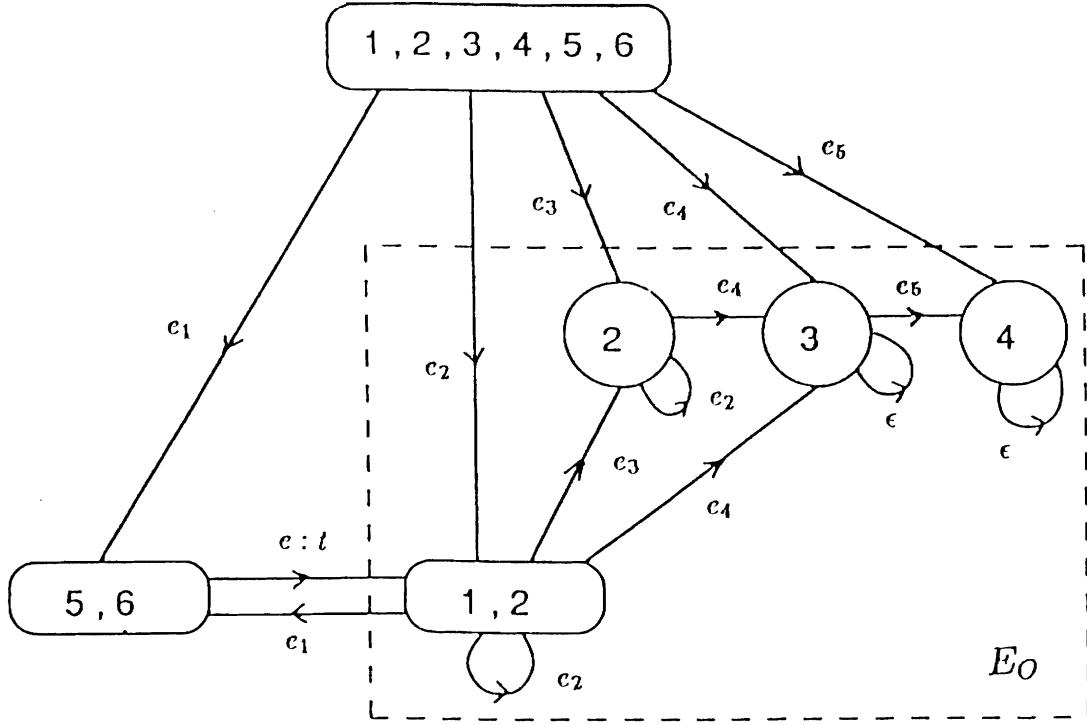


Figure 10 : Observer for the Grasping System

As an example, for the model of the grasping task, an observer can be formed for the system as shown in Figure 10. It can be easily seen that the system can be made stable with respect to the set E_O as defined in the previous section. At the start, the state of the system is totally ambiguous, however, the observer can be “guided” to the set E_O consisting of all the subsets of the good states E as defined on the visual system model. It can be seen that by enabling the tracking event from the state $(5, 6)$ to the state $(1, 2)$, all the system can be made stable with respect to E_O and thus the system is strongly output stabilizable. The singleton states represent the instances in time where the observer will be able to determine without ambiguity the current state of the system.

In the next sections we shall elaborate on defining the different events in the visual manipulation system and discuss different techniques for event and state identification. We shall also introduce a framework for computing the uncertainty in determining the observable visual events in the system and a method by which the uncertainty distribution in the system can be used to efficiently keep track of the different observer states and to navigate in the

observer automaton.

4 Event Identification

In this section we discuss different techniques for calculating the “events” that causes state transitions within the model that we discussed in the previous section. We introduce the concept of uncertainty in recovering the visual actions of the manipulation process and formulate a way of using the uncertainty in the system in an efficient recovery mechanism. Using the formulation in the previous section, it can be shown, from the examples used in modeling the manipulation process, that the events that causes state transitions are either primitives like specific 3-D movements of the manipulating hand and/or events like “there is an object now in view”, “the hand has enclosed the object” and so on. The events that are supposed to be identified and recovered at different states of the observer automaton are highly dependent on the *current* state in the observation process. Thus the observer tends to “look” at specific actions at different instances of time.

We next discuss techniques to be used in identifying the 3-D motion of the manipulation hand and/or the object, which are events that are always important to recover in order to enable the observer to navigate in the automaton. The process is started by identifying the manipulating hand and the object (if it exists) within the observer’s viewing window. We then proceed to develop an algorithm for detecting the two-dimensional motion vectors of the hand on the observer’s camera plane. Overall motion estimation and different tracking strategies are then developed in order to be able to stabilize the observer in the most efficient way.

4.1 Image Motion of the Hand

In order to be able to identify how the manipulating hand is moving within a grasping task, we use the image motion to estimate the hand movement. This task can be accomplished by either feature tracking or by computing the full optic flow. Feature tracking seems to be a good option for determining the hand motion, especially since the same hand will probably be used throughout the manipulation process, and if the system is to be ported to another manufacturing environment, then the interface that tracks specific features can be changed

while maintaining modularity. On the other hand, determining the full optic flow seems to be essential for computing the object motion, as we might not know in advance any shape or material information about the objects to be manipulated.

Many techniques were developed to estimate the optic flow (the 2-D image motion vectors) [3,9,15,19,22,42], we propose an algorithm for calculating the image flow and then we discuss a simpler version of the same algorithm for real time detection of the 2-D motion vectors. As a start, we can use a simple two-dimensional segmentation scheme in order to identify the hand and the manipulated object within the camera view. The input image is thresholded, and all the “objects” within an image are identified. An object is simply characterized by a region with a space of at least one pixel surrounding it from every where, thus regions with holes can be easily recognized using this technique. An edge tracer can be used for this purpose. We can assume that the largest object in the figure is the hand and the second largest object is the manipulated object, or we can make our decision built on the knowledge we have regarding the geometry of the hand and/or the object. As mentioned before, specific features can be identified, for example, the corners, or have a piece of paper with specific features stuck on the hand.

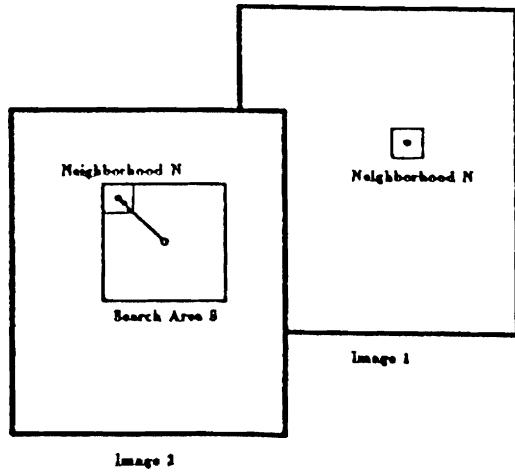


Figure 11 : Identifying the SSD Optical Flow

The image flow detection technique we use is based on the sum-of-squared-differences optic flow. We consider two images, 1 and 2 as shown in Figure 11. For every pixel (x, y) in image 1 we consider a pixel area N surrounding it and search a neighboring area S to seek

a corresponding area in image 2 such that the sum of squared differences in the pixel gray levels is minimal as follows :

$$SSD(\dot{x}, \dot{y}) = \min_{\dot{x}, \dot{y} \in S} \sum_{\Delta x, \Delta y \in N} [E(x + \Delta x, y + \Delta y) - \dot{E}(\dot{x} + \Delta x, \dot{y} + \Delta y)]^2$$

The image flow vector of pixel (x, y) then points from the center of N in the first image to the center of the best match in the second image. The search area S should be restricted for practicality measures. In the case of multiple best matches, we can use the one which implies minimum motion, as a heuristic favoring small movements. It should be noted that the accuracy of direction and magnitude of the optic flow determination depends on the sizes of the neighborhoods N and S .

There are three basic problems with this simple approach, one is that the sum of squared differences will be near zero for all directions wherever the graylevel is relatively uniform, the second is that it suffers from the so-called “aperture problem” even if there is a significant graylevel variation. To illustrate this point, consider a vertical edge moving to the right by one pixel distance, and suppose the N window size is 3×3 pixels and the S window size is 5×5 pixels, the squared-differences at an edge point reaches its maximum for three directions as indicated by the vectors (in pixel displacements); $(1, 0)$, $(1, -1)$ and $(1, 1)$. Figures 12.1 and 12.2 illustrates the aperture problem. The third problem is that the scheme will only determine the displacement to pixel accuracy.

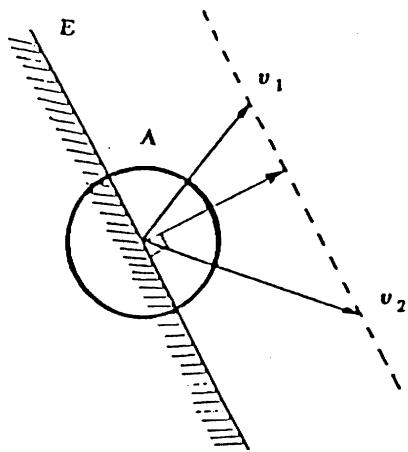


Figure 12.1 : The Aperture Problem

The direction of motion of edge E cannot be determined by viewing E through the aperture A

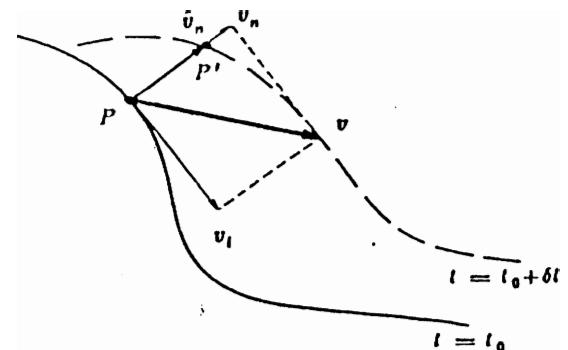


Figure 12.2 : Normal Flow Estimation

We solve the first problem by estimating the motion only at the hand or object pixels (as determined by the two-dimensional segmentation scheme) where the intensity changes significantly. The Sobel edge detector is applied to the first image to estimate the edge magnitude $M(x, y)$ and direction $D(x, y)$ for every pixel :

$$M(x, y) \approx \sqrt{E_x^2 + E_y^2}$$

$$D(x, y) \approx \tan^{-1} \left(\frac{E_x}{E_y} \right)$$

where E_x and E_y are the partial derivatives of the first image with respect to x and y , respectively. The edge direction and magnitude is discretized depending on the size of the windows N and S . The motion is then estimated at only the pixels where the gradient magnitude exceeds the input threshold value. Motion ambiguity due to the aperture problem can be solved by estimating only the *normal* flow vector. It is well known that the motion along the direction of intensity gradient only can be recovered. Then we evaluate the SSD functions at only those locations that lie on the gradient directions and choose the one corresponding to the minimal SSD, if more than one minimal SSD exist we can choose the one corresponding to the smallest movement, as described above. The full flow vector can then be estimated by using the following equation which relates the normal flow vector \vec{v}_n to the full flow vector \vec{v} .

$$\vec{v}_n = \vec{v} \cdot \vec{n}$$

This method works under the assumption that the hand image motion is locally constant. Solving the over-determined linear system will result in a solution for the full flow. The least square error of the system can help us to decide whether the assumption is a reasonably valid one for determining the event that caused the transition in the DEDS. On the other hand, full flow determination can be performed for small clusters of points in the image and a number of full flow estimates is then used for 3-D recovery.

To obtain sub-pixel accuracy, we can fit a one-dimensional curve along the direction of the gradient for all the SSD values obtained. A polynomial of the degree of the number of points used along the gradient can be used to obtain the best precision. However, for an S window of size 7×7 pixels or less and an N window of size 3×3 or so, a quadratic function can be used for efficiency and to avoid optimizational instabilities for higher order polynomials.

Subpixel accuracy using a quadratic function is shown in Figure 13. The subpixel optimum can be obtained by finding the minimum of the function used and using the displacement at which it occurred as the image flow estimate. To avoid probable discontinuities in the SSD values, the image could be smoothed first using a gaussian with a small variance.

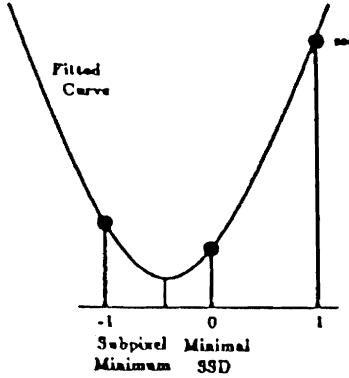


Figure 13 : Subpixel Accuracy for Optical Flow

A simpler version of the above algorithm can be implemented in real-time using a multi-resolution approach [42]. We can restrict the window size of N to 3×3 and that of S to 5×5 , and perform the algorithm on different levels of the gaussian image pyramid. A gaussian pyramid is constructed by the successive applications of gaussian low-pass filtering and decimation by half. The pyramid processor, PVM-1 is capable of producing complete gaussian pyramid from a 256 by 256 image in one video frame ($\frac{1}{30}$ of a second). Maxvideo boards can be used for the simultaneous estimation of image flow at all the levels of the pyramid for all the pixels. Image flow of 1 pixel at the second level would correspond to 2 pixels in the original image, 1 pixel displacement at the third level would correspond to 4 pixels in the original image, and so on. The level with the smallest least square fitting error of the normal flow can be chosen to get the full flow and the motion vector is scaled accordingly. This method is crude in the sense that it only allow image flow values of 1,2,4 or 8 pixel displacement at each pixel, but it can be used for detecting fast movements of the hand.

By either using a flow recovery algorithm or a feature identification and tracking algorithm,

we end up having a set of values for 2-D displacements of a number of pixels. The problem is how can we model the uncertainty in those 2-D estimates, which are to be used later for 3-D parameter recovery. For example, if the estimate is - for a specific 3-D feature - that pixel (x_i, y_j) has moved to pixel (x_m, y_n) , then the problem reduces to finding space probability distributions for the four indices. The sensor acquisition procedure (grabbing images) and uncertainty in image processing mechanisms for determining features are factors that should be taken into consideration when we compute the uncertainty in the optic flow. In sections 5, 6 and 7 we discuss these problems in details.

4.2 Recovering 3-D events

One can model an arbitrary 3-D motion in terms of stationary-scene/moving-viewer as shown in Figure 14. The optical flow at the image plane can be related to the 3-D world as indicated by the following pair of equations for each point (x, y) in the image plane [27] :

$$v_x = \left\{ x \frac{V_Z}{Z} - \frac{V_X}{Z} \right\} + \left[xy\Omega_X - (1 + x^2) \Omega_Y + y\Omega_Z \right]$$

$$v_y = \left\{ y \frac{V_Z}{Z} - \frac{V_Y}{Z} \right\} + \left[(1 + y^2) \Omega_X - xy\Omega_Y - x\Omega_Z \right]$$

where v_x and v_y are the image velocity at image location (x, y) , (V_X, V_Y, V_Z) and $(\Omega_X, \Omega_Y, \Omega_Z)$ are the translational and rotational velocity vectors of the observer, and Z is the unknown distance from the camera to the object.

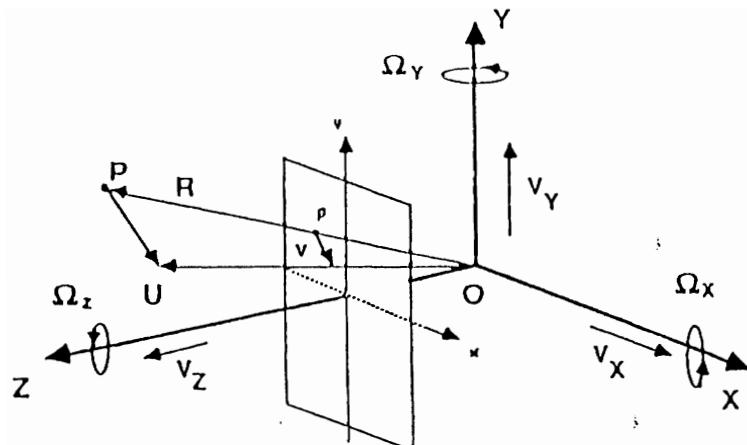


Figure 14 : 3-D Formulation for Stationary-Scene/Moving-Viewer

In this system of equations, the only knowns are the 2-D vectors v_x and v_y , if we use the formulation with uncertainty then basically the 2-D vectors are random variables with a known probability distribution. In case that the real 3-D relationships between feature points (on the hand) are known, then recovering the absolute depth is a simple process, The equations can then be formalized, in case that the 3-D features lie on a planar surface, as follows :

$$v_x = (1 - px - qy) \left(x \frac{V_Z}{Z_o} - \frac{V_X}{Z_o} \right) + \left[xy\Omega_X - (1 + x^2) \Omega_Y + y\Omega_Z \right]$$

$$v_y = (1 - px - qy) \left(y \frac{V_Z}{Z_o} - \frac{V_Y}{Z_o} \right) + \left[(1 + y^2) \Omega_X - xy\Omega_Y - x\Omega_Z \right]$$

where Z_o is the absolute depth, p and q are the planar surface orientations. It should be noticed that the resulting system of equations is nonlinear, however, it has some linear properties. The rotational part, for example, is totally linear. In section 8 we discuss different methods for solving the system of equations and thus recovering the 3-D parameters in real time with and without uncertainty formulation.

A part of the events definition, as mentioned before, is the recognition of the existence of an object, for example. In other words, identifying objects in the visual scene and not only recovering 3-D motion. Orientation of the object relative to the observer's camera and its shape can always be asserted by a simple 2-D segmentation strategy as mentioned in the discussion about computing the 2-D motion vectors. A data base of different shapes and orientations for different sized objects with the associated state that they may be manipulated in may be used and updated by the system. Correlation-based matching techniques can be used to compare 2-D object representations, while moment computations are used to scale, shift and re-orient the shapes to be correlated. New objects can still be recognized and stored in this data base to facilitate future accesses.

4.3 The Controllable Events

The only kind of control inputs that can be supplied to the observer robot are the tracking actions. Depending on the nature of the manipulation process, the observer has to keep track of the hand and object within the camera image plane in such a way so as to be able to observe the process. The intelligent tracking control is supplied by the DEDS formulation. Simple-minded tracking ideas, like keeping fixed 3-D relation between the camera and the

manipulating agent are not to be used in our system. The manipulation action might be a simple one that does not require complex tracking, such as screwing and unscrewing, however, more complex events, where the hand may occlude the manipulation process, or when the hand starts moving away from the observer, might suggest the need for complex tracking mechanisms, including translations and rotations of the observing robot hand on which the camera is mounted.

A subset of the three-dimensional motion and structure parameters would have to be calculated using two or more frames [14,36,39,41]. The size of the subset will depend on the *expected* kind of 3-D motion, as the current state of the DEDS system will specify. Our system needs to track the object while using all the six degrees of freedom of the observer robot in order to position the observer at the best feasible position at different states of the automaton. Using rotations only to follow the end effector of the manipulating robot is not sufficient for the stabilizing observer.

Two kinds of tracking mechanisms can be used, in the first kind, the two images on which the motion estimation algorithms will be used, will be taken while the camera is stationary and then the camera will move and the process will be repeated after the camera stops. The observer movement will be a “jerky” one. Another scheme can be used where the camera can grab images while the robot arm holding it is moving, in this case one should compensate for the moving arm before calculating the image flow of the hand and/or object. Thus, the problem reduces to finding the image flow due to the camera movement using the stationary-scene/moving-viewer 3-D formulation. In the absence of translations, for example, we can compensate for the rotational part in a very fast and efficient way. Compensation will have to be performed before using the structure and motion recovery algorithms. Velocity control for moving the observer’s camera can be used to match the moving agent’s speed.

5 Sensor Uncertainties

In this section and the next two sections we develop and discuss modeling the uncertainties in the recovered 2-D displacement vectors. As mentioned in the section describing techniques for recovering the image flow, the uncertainty in the recovered values results from sensor uncertainties and noise and from the image processing techniques used to extract and track

features. When dealing with measurements of any sort, it is always the case that the measurements are accompanied by some error. Mistakes also occur, where mistakes are not large errors but failures of a system component or more. A description of errors, mistakes and modeling them can be found in [4,5].

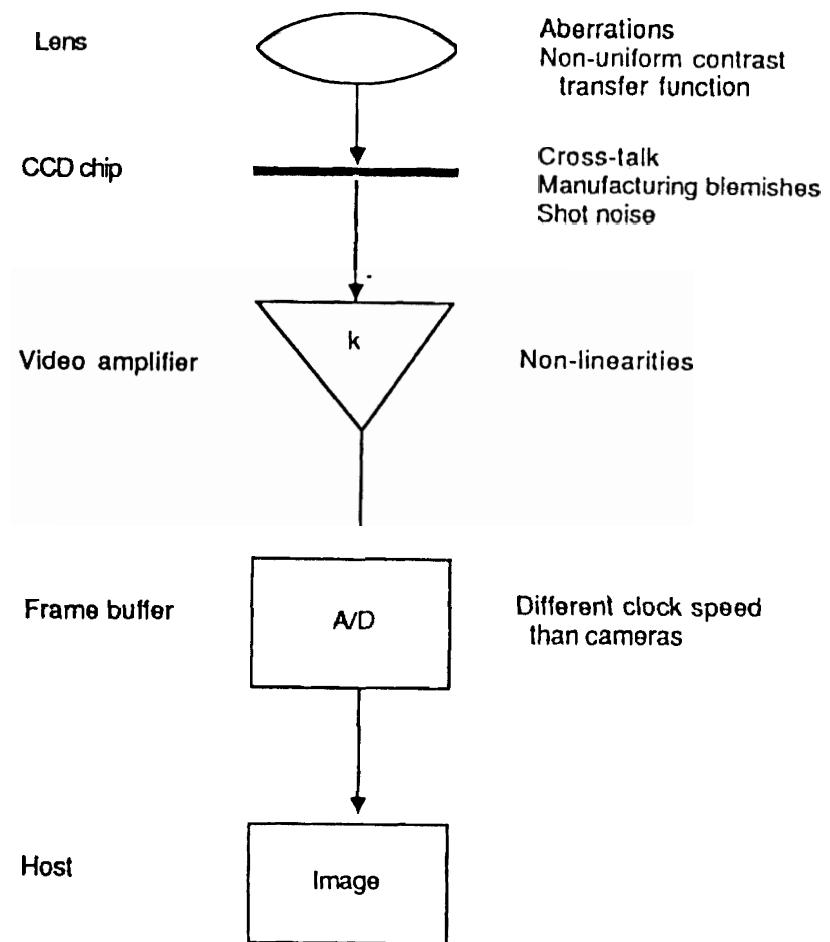


Figure 15 : Image formation.

In this section we discuss errors in image formation. The observer robot uses a camera to grab and register images of the manipulation system, so we need to know errors in mapping

from the 3-D world features to the 2-D domain which we use in forming 3-D hypothesis about the task under supervision. The accuracy, precision and modeling uncertainty of the camera as our sensor is an important issue and the first step towards forming a full uncertainty model for recovering the 3-D events in the observer automaton.

In Figure 15 (redrawn from [5]), a model of the image formation process is illustrated, which lists some salient features of each component. As a lot of the image processing algorithms compute derivatives of the intensity function, noise in the image will be amplified and propagated throughout the observation process. The goal of this treatment is to find a distribution for the uncertainty of mapping a specific 3-D feature into a specific pixel value. In other words, if the feature 2-D position was discovered to be (i, j) , then the goal is to find a 2-D distribution for i and j , assuming that there is no uncertainty in the technique used to extract the 2-D feature, the technique's uncertainty will be discussed in the next section. The end product of modeling the sensor uncertainty is to be able to say a statement like : “The 3-D feature F is located in the 2-D pixel position (i, j) with probability p_1 or located in the 2-D pixel position $(i, j + 1)$ with probability p_2 or given that the registered location is (l, m) , such that $p_1 + p_2 + \dots + p_n = 1$, and \mathcal{A} error in the 2-D feature recovery mechanism.”

5.1 Image Formation Errors

The errors in the image formation process are basically of two different kinds. The first type is a spatial error, the other type is a temporal error. The spatial error due to the noise characteristics of a CCD transducer can be due to many reasons, among which are dark signatures and illumination signatures. The technique to be used is to take a large number of images, we can denote the image intensity function as a 3-D function $I(u, v, t)$, with spatial arguments u and v and temporal argument t . The sample mean of the image intensities over N time samples can be denoted by $\bar{I}(u, v)$.

$$\bar{I}(u, v) = \frac{1}{N} \sum_{t=1}^N I(u, v, t)$$

The spatial variance in a 5×5 neighborhood of the means is computed by:

$$s^2(u, v) = \sum_{i=-2}^2 \sum_{j=-2}^2 (\bar{I}(u + i, v + j) - \bar{I}(u, v))^2$$

The dark signature of the camera can be determined by computing $\bar{I}(u, v)$ of each pixel with the lens cap on. It will be found that a small number of pixels will have non-zero mean and non-zero variance. The specific pixel locations are blemished and should be registered. The uniform illumination is computed by placing a nylon diffuser over the lens and computing the mean and variance. It will be noticed that due to digitizing the CCD array into a pixel array of different size, and the difference in sample rates between the digitizer and camera, the border of the image will have different mean and variance from the interior of the image. Some “stuck” pixels at the location of the blemished pixels will also be noted. The contrast transfer function will also be noted to vary at different distances from the center of the lens. Temporal noise characteristics can also be identified by taking a number of experiments and notice the time dependency of the pixels intensity function. In our treatment and for our modeling purposes we concentrate on the spatial distribution of noise and its effect on finding the 2-D uncertainty in recovering a 3-D feature location in the pixel array.

5.2 Calibration and Modeling Uncertainties

Methods to compute the translation and rotation of the camera with respect to its coordinates, as well as the camera parameters, such as the focal length, radial distortion coefficients, scale factor and the image origin, have been developed and discussed in the literature [8,21,37]. In this section we use a static camera calibration technique to model the uncertainty in 3-D to 2-D feature locations. In particular we use the sequence of steps used to transform from 3-D world coordinates to computer pixel coordinates in order to recover the pixel uncertainties, due to the sensor noise characteristics described previously.

As shown in Figure 16, the sequence of steps is used for a coplanar set of points in order to obtain the rotation and translation matrices, in addition to the camera parameters. The input to the system are two sets of coordinates, (X_f, Y_f) , which are the computer 2-D pixel image coordinates in frame memory and (x_w, y_w, z_w) , which are the 3-D world coordinates of a set of coplanar points impressed on a piece of paper with known inter-point distances. A discussion of the exact mathematical formulation of the inter-step computations to find all the parameters can be found in [8].

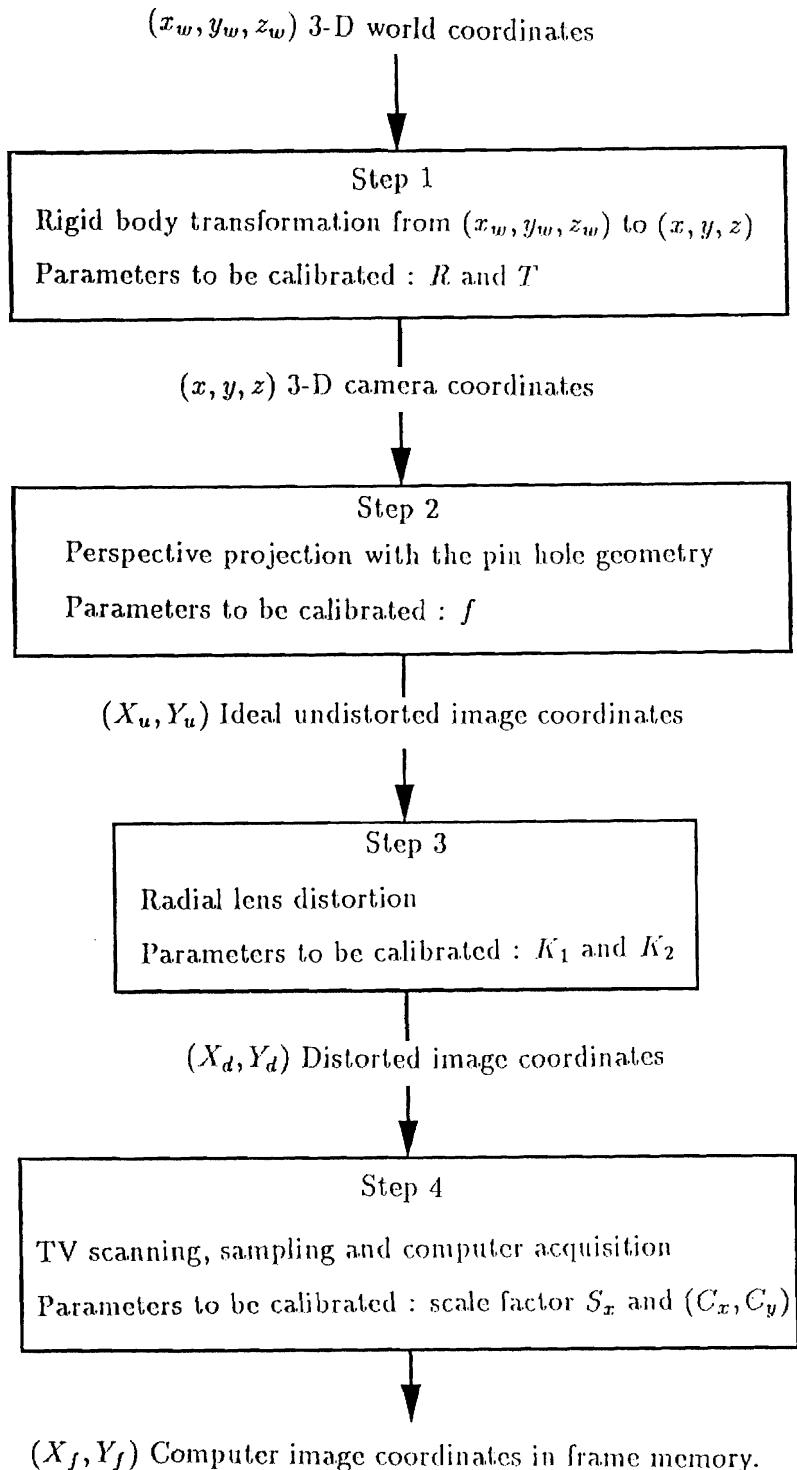


Figure 16 : The Four-Steps Transformation from 3D World Coordinates to Computer Image Coordinates

Our approach is to treat the whole camera system as a black box and make input/output measurements and develop a model of its parametric behaviour. The next step is to utilize the recovered camera parameters and the number of 3-D points which we created in order to formulate a distribution of the 2-D uncertainty. The points used in calibration and later in recovering the distribution can be the actual features on the robot hand that are to be tracked and thus providing a similar experimental environment to the one that the observer will operate in.

The strategy used to find the 2-D uncertainty in the features 2-D representation is to utilize the recovered camera parameters and the 3-D world coordinates (x_w, y_w, z_w) of the known set of points and compute the corresponding pixel coordinates, for points distributed throughout the image plane a number of times, find the actual feature pixel coordinates and construct 2-D histograms for the displacements from the recovered coordinates for the experiments performed. The number of the experiments giving a certain displacement error would be the z axis of this histogram, while the x and y axis are the displacement error. Different histograms can be used for different 2-D pixel positions distributed throughout the image plane. The three dimensional histogram functions are then normalized such that the volume under the histogram is equal to 1 unit volume and the resulting normalized function is used as the distribution of pixel displacement error, thus modeling the sensor uncertainty. The black box approach is thus used to model errors in a statistical sense.

6 Image Processing Uncertainties

In this section we describe a technique by which developing uncertainties due to the image processing strategy can be modeled. In addition, we end the discussion by combining both the sensor uncertainties developed in the previous section and the models developed in this section to generate distribution models for the uncertainty in estimating the 2-D motion vectors. These models are to be used for determining the full uncertainty in recovering the 3-D events that causes state transitions between states of the observer automaton.

We start by identifying some basic measures and ideas that are used frequently to recognize the behaviour of basic image processing algorithms and then proceed to describe the technique we use in order to compute the error model in locating certain features from their 2-D

representation in the pixel array. We concentrate on modeling the error incurred in extracting edges, as edge extraction is a very popular mechanism that is used for both identifying feature points on the manipulating hand and also for computing 2-D contours of the object under supervision. When we discussed flow recovery techniques before, it was discussed in details that the optic flow recovery algorithm using local matching works well for the intensity boundaries and not for the inside regions.

6.1 Edge Extraction Uncertainties

Edge extraction strategies and methods to evaluate their performance qualitatively and quantitatively have been presented and discussed in the literature [11,13,24,29]. There are many types of edges, ideal, ramp and noisy edges as shown in Figure 17 are only three of them. Different curvatures in the edges also constitute another dimension to be taken into consideration when it comes to asserting the types of edges that exist.

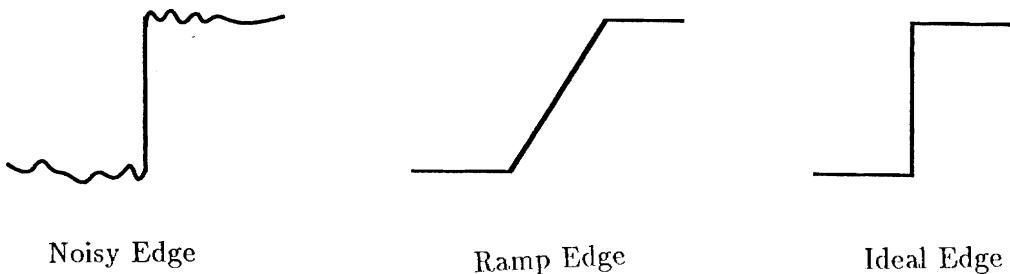


Figure 17

The goal of developing the error models for edge extraction is to be able to say a statement like : “Given that the 2-D feature recovered using the edge recovery S is in pixel position (x, y) , then there is a probability that the feature was originally at pixel position $(x+1, y)$ with probability p_1 or etc. due to the noise in the pixel image, such that $p_1 + p_2 + \dots + p_n = 1$.” The problem is to find the probabilities.

It should be obvious that there may be different types of noises and also different levels of those types that might vary at different locations in the sensor image plane. This adds to

the different models that we might have to construct. Our approach is to use ideal, that is, synthesized edges of different types, locations and also orientations in image frames then corrupt them with different kinds and levels of noises. We know the ideal edge points from the ideal image, for which we shall use the edge detector that is to be used in the observer experiment. The corrupted images will then be operated upon by the detector and the edge points located. The edge points will differ from the ideal image edge points. The problem reduces to finding corresponding edge points in corrupted and ideal images then finding the error along a large number of edge points. A 2-D histogram is then constructed for the number of points with specific displacement errors from the ideal point. The volume of the histogram is then normalized to be equal to 1, the resulting 3-D function is the 2-D probability density function of the error of displacements.

In Figure 18, an ideal box is drawn, then corrupted with an additive gaussian noise with σ equal to 3, 10, 20, 30 and 50 respectively and then the edges computed as shown. In the box there are four different kinds of ideal edges (different orientations with the object inside or outside of the background). The correspondence between edge points in the corrupted and ideal is established by choosing the point with the *minimal* distance from the ideal edge point, *such that* it does not correspond to another ideal edge point. The histogram is constructed for each edge and then normalized. For practicality measures, the process can be repeated for orientations differing by 15° and the set of distributions preserved. Whenever the observer automaton deals with a specific edge while extracting features, the corresponding distribution is referenced.

6.2 Computing 2-D Motion Uncertainty

In this section we describe how to combine sensor and strategy error models to compute models for the recovered image flow values. To simplify the idea, let's assume that we have recovered a specific feature point (x_1, y_1) in an image grabbed at time instant t and the corresponding point (x_2, y_2) at time $t + 1$. The problem is to figure out the distribution of v_x . As an example, to explain the procedure, let's assume that from the 3-D sensor distribution we have have computed the marginal density function of the x coordinate of x_1 in the point:

$$f_X(x) = \int_R f_{X,Y}(x,y)dy$$

where R is all the possible y values within the sensor uncertainty model.

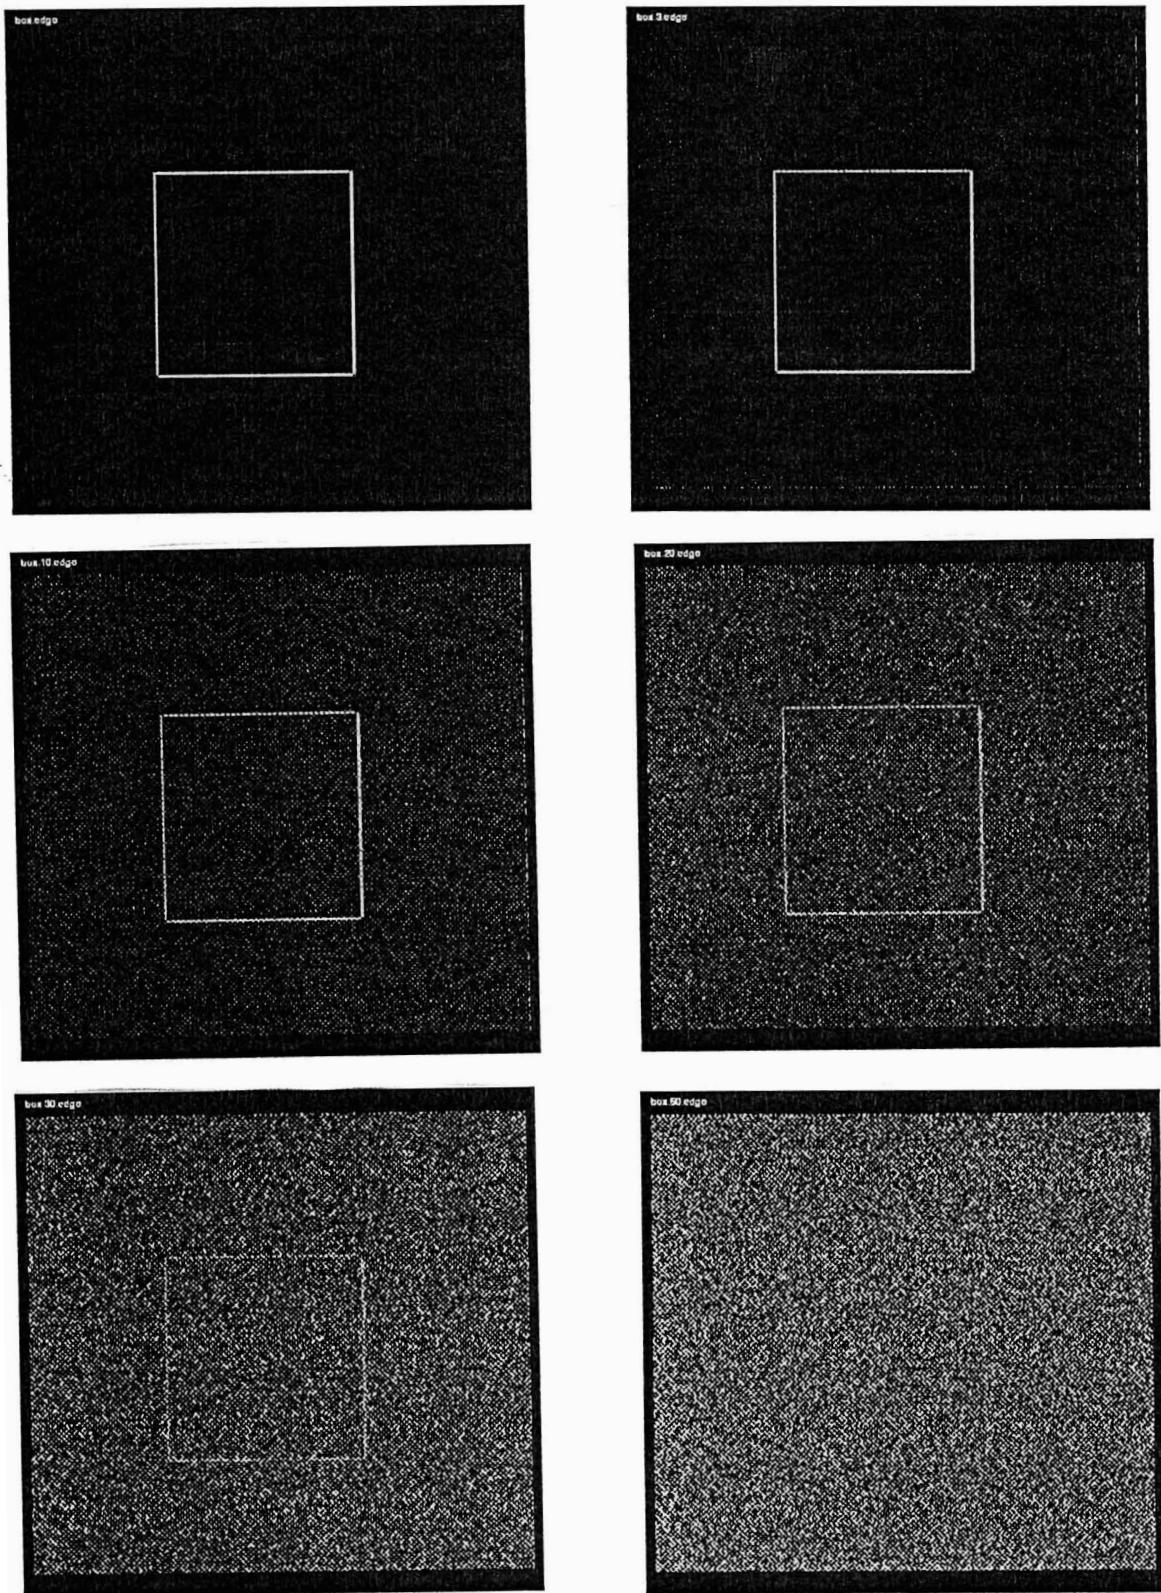


Figure 18 : Edge Detection Results for Different Noise Levels

The same process is applied for the strategy distribution and another function is recovered. To simplify things, let's assume that both distributions look like the distribution in Figure 19, that is, there is an equal probability equal to $\frac{1}{3}$ that the x coordinate is the same, or shifted one position to the left or the right. Combining both distributions in a filtering-through process would produce the distribution shown in Figure 20, which is the error probability density function of having the 3-D feature x 2-D coordinate in the recovered image 2-D x position. Further more, assume that x_2 distribution is the same.

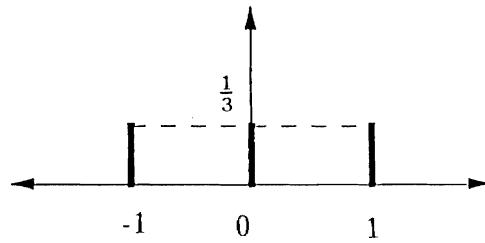


Figure 19 : Distribution of the x -coordinate

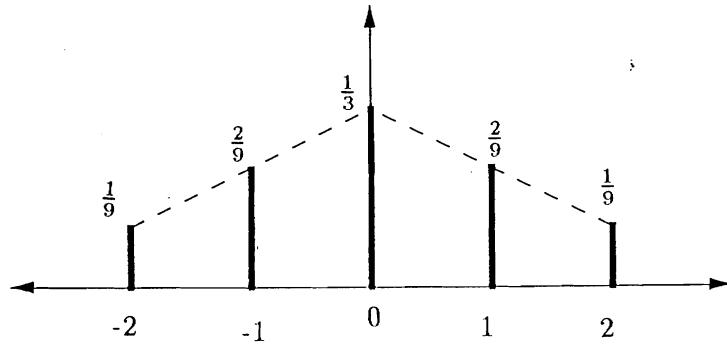


Figure 20 : Combined Sensor and Strategy Distribution

The problem reduces to finding the distribution of the optic flow x component, using these two combined distributions. As an example, if $x_1 = 10$ and $x_2 = 22$, then all probability statements can be easily computed, a set of some of these probability statement is shown :

$$P(v_x = 8) = P((x_1 = 12) \wedge (x_2 = 20)) = \frac{1}{9} \times \frac{1}{9} = \frac{1}{81}$$

$$P(v_x = 9) = P(((x_1 = 12) \wedge (x_2 = 21)) \vee ((x_1 = 11) \wedge (x_2 = 20))) = (\frac{1}{9} \times \frac{2}{9}) + (\frac{2}{9} \times \frac{1}{9}) = \frac{4}{81}$$

$$P(v_x = 10 | x_1 = 10) = \frac{P(x_1 = 10 \wedge x_2 = 20)}{P(x_1 = 10)} = \frac{\frac{3}{9} \times \frac{1}{9}}{\frac{3}{9}} = \frac{1}{9}$$

Consequently, all distributions and expected values can be computed from the combination of the sensor level and strategy level uncertainty formulation. Those flow models are then passed to the higher levels for 3-D recovery. In the next section we discuss a method for refining the measured 2-D motion vectors and we then proceed to formulate the 3-D modeling of events as defined by the observer automaton.

7 Refining Image Motion

In this section we describe a method to refine the recovered 2-D motion vectors on the image plane. Having obtained from the sensor and extraction strategy uncertainty levels distribution estimates for the image flow of the different features, we now try to eliminate the unrealistic ones. We concentrate on the flow estimates for the motion of the manipulating hand and develop a technique that is to be used during the observation process as a means to reject faulty estimates. Faulty estimates can result from noise, errors or mistakes in the sensor acquisition process, manipulation or visual problems like occlusion, modeling the uncertainties in the previous two levels may still leave room for such anomalies.

We assume that the features to be tracked on the hand lie on a planar surface or that segmenting the hand as a polyhedra object into planar surfaces is simple, although the modification would be very simple to allow for arbitrary 3-D positions of the feature distribution. Since we know a-priori some information about the mechanical capabilities and limitations and geometric properties of the hand, also about the rate of visual sampling for the observer, since we actually control that, we might be able to assert some limits on some of the visual parameters in our system.

To illustrate the idea behind the approach, consider Figure 21, assume all the curves are 2-D parabolic functions $y = ax^2 + bx + c$, if the set of data points are as shown in the figure, then a least square error fit will produce the function D . However, if we know some upper and lower limits on the values of the coefficients a , b and c then we might be able to construct an upper and lower function parabolas A and C as an enclosing envelope, outside which we can reject all the data points. In that case, we can do a fit for the points that lie inside the

envelope and obtain a more realistic function as shown by the curve *B*.

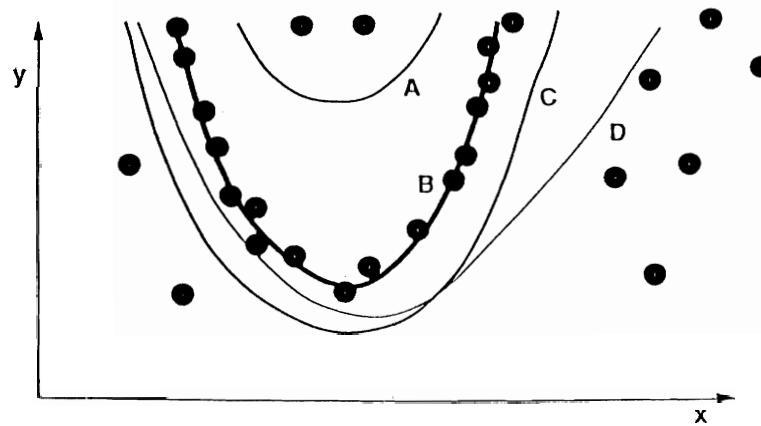


Figure 21 : Fitting Parabolic Curves

The situation for rejecting estimates for the image flow is not much different. We know equations that govern the behaviour of the image flow as a function of the structure and 3-D motion parameters, as follows :

$$v_x = (1 - px - qy) \left(x \frac{V_Z}{Z_o} - \frac{V_X}{Z_o} \right) + \left[xy\Omega_X - (1 + x^2) \Omega_Y + y\Omega_Z \right]$$

$$v_y = (1 - px - qy) \left(y \frac{V_Z}{Z_o} - \frac{V_Y}{Z_o} \right) + \left[(1 + y^2) \Omega_X - xy\Omega_Y - x\Omega_Z \right]$$

Which are second degree functions in *x* and *y* in three dimensions, $v_x = f_1(x, y)$ and $v_y = f_2(x, y)$.

In addition, we know upper and lower limits on the coefficients *p*, *q*, V_X , V_Y , V_Z , Ω_X , Ω_Y , Ω_Z and Z_o , as we know that the mechanical abilities of the robot arm holding the hand will make the relative velocity and distance between the camera impossible to exceed specific values within visual sampling timing period. So the problem reduces to constructing the three dimensional envelopes for v_x and v_y as the worst case estimates for the flow velocity and rejecting any measured values that lie outside that envelope. Figure 22a indicates the maximal v_x that can ever be registered on the CCD array of the camera, the *x* and *y* are in

millimeters and the $x - y$ plane represents the CCD image plane, the depth Z is the maximal v_x in millimeters on the CCD array that can ever be registered. Figure 22b indicates the minimal v_x , it can be noticed that they are symmetric due to the symmetry in the limits of the coefficients.

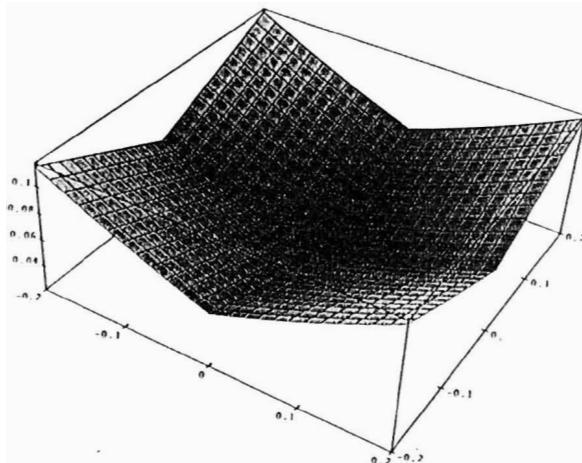


Figure 22.a : Maximal v_x

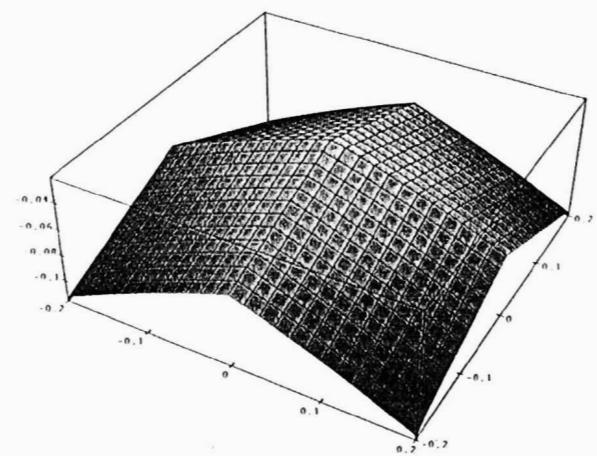


Figure 22.b : Minimal v_x

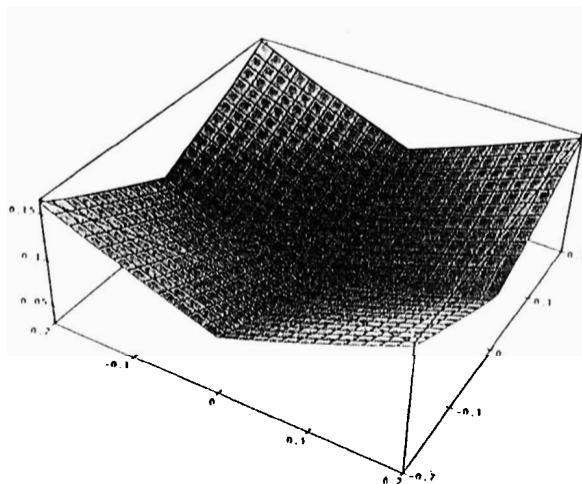


Figure 22.c : Maximal Flow Magnitude

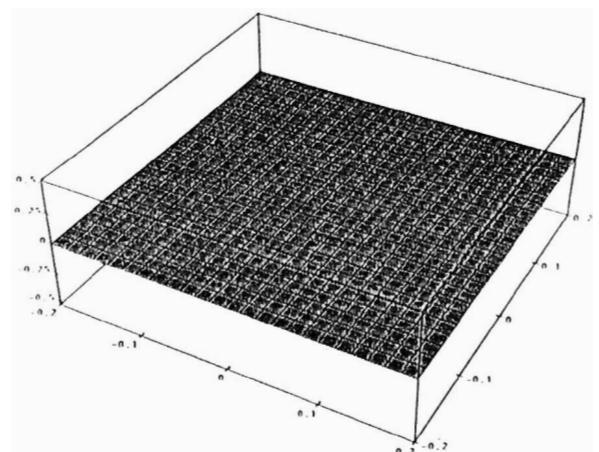


Figure 22.d : Minimal Flow Magnitude

As an example, we write the equation governing the maximum v_x value in the first quadrant of the $x - y$ plane (x^+, y^+).

$$v_{x_{max}} = \left(-\frac{fV_{X_s}}{Z_{o_s}} - f\Omega_{Y_s} \right) + \left(\frac{V_{Z_l}}{Z_{o_s}} + \frac{\max(p_l V_{X_l}, p_s V_{X_s})}{Z_{o_s}} \right) x + \left(\frac{\max(q_l V_{X_l}, q_s V_{X_s})}{Z_{o_s}} + \Omega_{Z_l} \right) y + \left(\frac{\Omega_{X_l}}{f} - \frac{\min(q_l V_{Z_s}, q_s V_{Z_l})}{f Z_{o_s}} \right) xy - \left(\frac{\min(p_l V_{Z_s}, p_s V_{Z_l})}{f Z_{o_s}} + \frac{\Omega_{Y_s}}{f} \right) x^2$$

where the subscripts s and l denote lower and upper limits, respectively. At first sight the problem of determining the maximum value of v_x seems to be a constrained non linear optimization problem, which is true, however, assuming that the upper and lower limits of the coefficients are equal in magnitude and opposite in directions (except for Z_o , which is used only as $Z_{o_s}^+$) makes the input to the *max* and *min* functions in the above equations always equal and thus providing one more degree of freedom in choosing the parameters and making the choice consistent throughout the equation. Thus the problem becomes simply to write eight equations as the above one for each of v_x and v_y , to draw the function in each of the four quadrants for maximum and minimum envelopes. We shall not rewrite the sixteen equations here, but we show the results for v_x in Figures 22a and 22b, Figures 22c and 22d are the maximum and minimum magnitude $m(x, y)$ for the the image flow at any given point, where :

$$m(x, y) = \sqrt{v_x^2 + v_y^2}$$

It should be noted that the maximum absolute possible value of the image flow is minimal at the origin of the camera image plane and increases quadratically as the distance increases from the center.

The above envelopes are then used to reject unrealistic 2-D velocity estimates at different pixel coordinates in the image. As a further note, it should be mentioned that some on-line elimination procedures can be implemented depending on the current positions in the observer automaton, for example, the image flow field tends to assume certain configurations in the image plane depending on the 3-D motion, independent of the object's or the hand's structure, if the motion is only relative rotational velocities, the flow vectors all tend through pass from the same point. In other words, in addition to off-line a-priori estimation of the envelopes and on-line testing of measurements, we can also develop custom rejection techniques for certain observer automata states.

8 Recovering World Events

In this section we describe different techniques for recovering the 3-D events. In particular we utilize the refined 2-D motion distributions that were computed in the previous levels in order to achieve a robust estimation of the three dimensional motion and structure vectors of the scene under observations. We develop some techniques for finding estimates of the required parameters and discuss mathematical formulations that will enable us to determine the 3-D event distributions.

We concentrate in our treatment of the subject on determining the manipulating hand parameters, as the hand configuration is well defined, we also continue using the assumption that the feature points lie on a planar surface. As argued before, the extension to arbitrary configurations is straight forward. The object behaviour can be asserted using similar techniques and/or by observing conveniently located surface patches under similar assumptions. We start by describing a deterministic method to recover 3-D parameters, then we describe other approximate methods and we conclude by discussing some mathematical formulations for using the same techniques for recovering variable distributions of the world events at different observer states.

The problem of recovering scene structure and the camera motion relative to the scene has been one of the key problems in computer vision. Many techniques have been developed for the estimation of structure and motion parameters (Tsai and Huang [36], Weng et al. [41] etc.). A lot of existing algorithms depend on evaluating the motion parameters between two successive frames in a sequence. However, recent research on structure and motion has been directed towards using a large number of frames to exploit the history of parametric evolution for a more accurate estimation and noise reduction (Ullman [39], Grzywacz and Hildreth[14] etc.)

Next, we describe a method for recovering the 3-D motion and orientation of the planar surface (on which lies the hand features) from an evolving image sequence. The algorithm utilizes the image flow velocities in order to recover the 3-D parameters. First, we develop an algorithm which iteratively improves the solution given two successive image frames. The solution space is divided into three subspaces - the translational motion, the rotational motion and the surface slope. The solution of each subspace is updated by using the current solution

of the other two subspaces. The updating process continues until the motion parameters converge, or until no significant improvement is achieved.

Second, we further improve the solution progressively by using a large number of image frames and the ordinary differential equations which describe the evolution of motion and structure over time. Our algorithm uses a weighted average of the expected parameters and the calculated parameters using the 2-frame iterative algorithm as current solution and continues in the same way till the end of the frame sequence. Thus it keeps track of the past history of parametric evolution.

The solution is further improved by exploiting the temporal coherence of 3-D motion. We develop the ordinary differential equations which describe the evolution of motion and structure in terms of the current motion/structure and the measurements (the 2-D motion vectors) in the image plane. As an initial step we assume that the 3-D motion is piecewise uniform in time. The extended Kalman filter is then used to update the solution of the differential equations.

8.1 A 3-D Recovery Algorithm

One can model an arbitrary 3-D motion in terms of stationary-scene/moving-viewer as shown previously in Figure 14. The optical flow at the image plane can be related to the 3-D world as indicated by the following pair of equations (In case of a planar surface), for each point (x, y) in the image plane :

$$v_x = (1 - px - qy) \left(x \frac{V_Z}{Z_o} - \frac{V_X}{Z_o} \right) + \left[xy\Omega_X - (1 + x^2) \Omega_Y + y\Omega_Z \right]$$

$$v_y = (1 - px - qy) \left(y \frac{V_Z}{Z_o} - \frac{V_Y}{Z_o} \right) + \left[(1 + y^2) \Omega_X - xy\Omega_Y - x\Omega_Z \right]$$

where v_x and v_y are the image velocity at image location (x, y) , (V_X, V_Y, V_Z) and $(\Omega_X, \Omega_Y, \Omega_Z)$ are the translational and rotational velocity vectors of the observer, p and q are the planar surface orientations. The situation becomes, for each point, two equations in eight unknowns, namely, the scaled translational velocities V_X/Z_o , V_Y/Z_o and V_Z/Z_o , the rotational velocities Ω_X , Ω_Y and Ω_Z and the orientations p and q . Differential methods could be used to solve those equations by differentiating the flow field and by using approximate methods to find the flow field derivatives. The existing methods for computing the derivatives of the flow

field usually do not produce accurate results. Our algorithm uses a discrete method instead, i.e, the vectors at a number of points in the plane is determined and the problem reduces to solving a system of nonlinear equations.

It should be noticed that the resulting system of equations is nonlinear, however, it has some linear properties. The rotational part, for example, is totally linear, also, for any combination of two spaces among the rotational, translational and slope spaces, the system becomes linear. For the system of equations to be consistent, we need the flow estimates for at least four points, in which case there will be eight equations in eight unknowns.

8.1.1 Two-Frame Algorithm

The algorithm takes as input the estimate of the flow vectors at a number of points ≥ 4 obtained from motion between two images. It iterates updating the solution of each subspace by using the solution of the other two subspaces. Each update involves solving a linear system, thereby it requires to solve three linear systems to complete a single iteration. This process continues until the solution converges, or until no significant improvement is made. The algorithm proceeds as follows :

1. Set $p, q = 0$;
 input the initial estimate for rotation ;
 Solve the linear system for translation;
2. Use the translation and rotation from step 1 ;
 Solve the linear system for the slope ;
3. Set $i=1$;
 While ($i \leq \text{Max. Iterations}$) and (no convergence) Do
 Solve for the rotations using latest estimates of translations, p and q ;
 Solve for the translations using latest estimates of rotations, p and q ;
 Solve for p, q using latest estimates of translations and rotations;
 end While ;

8.1.2 Complexity Analysis

As we mentioned earlier, one should notice in the equations relating the flow velocities with the slope, rotational and translational velocities that they are “quasi-linear”, if one can say so. The equations exhibit some linear properties. This suggests that a purely iterative technique for solving non-linear equations might not be an excellent choice, since, the variables are linearly related in some way. To think of a way of “inverting” the relations might be a good start, although to do that without a framework based on iterating and gravitating towards a solution is not a good idea.

This makes one think of applying a method which converges faster than a purely iterative scheme like Newton’s method. However, the complexity of Newton’s method is determined by the complexity of computing the inverse Jacobian, which is of an order of N^3 , or $N^{2.81}$ multiplications as the lower bound using Strassen’s technique. In our case, since we have at least 8 equations in 8 unknowns, the complexity is of order $8^3 = 512$ multiplications at every iteration, and the method does not make any use of the fact that the set of equations at hand exhibits some linear properties.

The algorithm proposed, on the other hand, makes very good use of the fact that there are some linearity in the equations, by inverting the set of relations for each subspace at every iteration. The complexity at every iteration is of the order of the complexity of computing the pseudo-inverse which is of the order of $(3^3 + 3^3 + 2^3)$ multiplications at each iteration, where the first 3 comes from solving the system for the rotational variables, the second 3 is for the translations, the last 2 is for p and q . This is equal to 62 multiplications at every iteration, which is significantly less than the 512 multiplications in a method like Newton’s for example. It was noticed that the algorithm converged to solution in a very small number of iterations for most experiments we have conducted so far. The maximum number of iterations was 6.

Using the latest solution obtained from the two-frame analysis as the initial condition for the next two-frame problem in the image sequence would further decrease the complexity, as the next set of parameters would, most probably, be close in values to the current parameters, thus the number of iterations needed to converge to the new solution would decrease significantly.

8.1.3 Observations

- The algorithm is not sensitive to the initial condition of the orientation parameters. The plane is simply assumed to be a frontal one at the beginning. The slope parameters evolves with iterations.
- The algorithm is sensitive to input noise just like other existing algorithms, some experiments shows the sensitivity with respect to the change of viewing angle. Similarly, the algorithm performs better for a large number of points that are evenly distributed throughout the planar surface, than it does for clustered, smaller number of image points.
- It is proven that there exists dual solutions for such systems. However, if our method gravitates towards a “fixed point” in the solution space we can find the other explicitly in terms of the first one from the relations given by Waxman and Ullman [40].

8.1.4 Multi-Frame Algorithm

The ordinary differential equations that describe the evolution of motion and structure parameters are used to find the expression for the expected parameter change in terms of the previous parameter estimates. The expected change and the old estimates are then used to predict the current motion and structure parameters.

At time instant t , the planar surface equation is described by

$$Z = pX + qY + Z_o$$

To compute the change in the structure parameters during the time interval dt , we differentiate the above equation to get

$$\frac{dZ}{dt} = p \frac{dX}{dt} + X \frac{dp}{dt} + q \frac{dY}{dt} + Y \frac{dq}{dt} + \frac{dZ_o}{dt}$$

The time derivatives of (X, Y, Z) in the above expression are given by the three components of the vector $-(\mathbf{V} + \boldsymbol{\Omega} \times \mathbf{R})$ that represent the relative motion of the object with respect to the camera. Substituting these components for the derivatives and the expression $pX + qY + Z_o$ for Z we can get the exact differentials for the slopes and Z_o as

$$\begin{aligned}
dZ_o &= Z_o [(\Omega_Y + V_X)p - (\Omega_X - V_Y)q - V_Z] dt \\
dp &= [p(\Omega_Y p - \Omega_X q) + (\Omega_Y + \Omega_Z q)] dt \\
dq &= [q(\Omega_Y p - \Omega_X q) - (\Omega_X + \Omega_Z p)] dt
\end{aligned}$$

Using the above relations, we can compute the new structure parameters at time $t + dt$ as

$$\dot{p} = p + dp, \dot{q} = q + dq \text{ and } Z_o = Z_o + dZ_o$$

Thus the slope parameters evolve at time $t + dt$ as follows :

$$\begin{bmatrix} \dot{p} \\ \dot{q} \end{bmatrix} = \begin{bmatrix} p \\ q \end{bmatrix} + \begin{bmatrix} \Omega_Y p - \Omega_X q & \Omega_Z & \Omega_Y \\ -\Omega_Z & \Omega_Y p - \Omega_X q & -\Omega_X \end{bmatrix} \begin{bmatrix} p \\ q \\ 1 \end{bmatrix} dt$$

The new translational velocity V at time $t + dt$ can be found in the absence of accelerations from

$$\dot{\mathbf{V}} = \mathbf{V} + \mathbf{V} \times \boldsymbol{\Omega} dt$$

Dividing $\dot{\mathbf{V}}$ by \dot{Z}_o we get the new expected scaled translational velocity components at time $t + dt$ as follows :

$$\begin{bmatrix} \dot{V}_X \\ \dot{V}_Y \\ \dot{V}_Z \end{bmatrix} = \begin{bmatrix} V_X \\ V_Y \\ V_Z \end{bmatrix} + \begin{bmatrix} -s & \Omega_Z & \Omega_Y \\ -\Omega_Z & -s & \Omega_X \\ \Omega_Y & -\Omega_X & -s \end{bmatrix} \begin{bmatrix} V_X \\ V_Y \\ V_Z \end{bmatrix} dt,$$

where s is expressed as follows :

$$s = (\Omega_Y + V_X)p - (\Omega_X - V_Y)q - V_Z$$

The expected rotational parameters at time $t + dt$ remain equal to their values at time t since

$$\dot{\boldsymbol{\Omega}} = \boldsymbol{\Omega} + \boldsymbol{\Omega} \times \boldsymbol{\Omega} dt = \boldsymbol{\Omega}$$

and thus

$$(\dot{\Omega}_X, \dot{\Omega}_Y, \dot{\Omega}_Z) = (\Omega_X, \Omega_Y, \Omega_Z)$$

Our first multi-frame algorithm uses a weighted average of the expected parameters at time $t + dt$ from the above equations and the calculated parameters using the two-frame iterative

algorithm as the solution at time $t + dt$, and continues in the same way until the end of the frame sequence. Thus it keeps track of the past history of parametric evolution. We further develop the first multi-frame algorithm to exploit the temporal coherence of 3-D motion. We develop the ordinary differential equations which describe the evolution of motion and structure in terms of the current motion/structure and the two-dimensional flow vectors in the image plane. We assume that the 3-D motion is piecewise uniform in time, i.e, $\vec{\Omega} = \vec{V} = \mathbf{0}$. We then use the equations expressing the time derivative of the slope derived above and the fact that the derivative of the rotational velocities is zero and develop the following expressions for the scaled translational velocities and the depth Z_o :

$$\frac{dV'_X}{dt} = -V'_X \frac{1}{Z_o} \frac{dZ_o}{dt}, \quad \frac{dV'_Y}{dt} = -V'_Y \frac{1}{Z_o} \frac{dZ_o}{dt} \quad \text{and} \quad \frac{dV'_Z}{dt} = -V'_Z \frac{1}{Z_o} \frac{dZ_o}{dt}$$

$$\frac{1}{Z_o} \frac{dZ_o}{dt} = -V'_Z - pv_x - qv_y$$

The extended Kalman filter is then used to update the solution of the differential equations. Where the state vector can be written as :

$$X = [V'_X \ V'_Y \ V'_Z \ \Omega_X \ \Omega_Y \ \Omega_Z \ p \ q]$$

and the measurement vector is expressed as :

$$Z = [v_x \ v_y \ \frac{\delta v_x}{\delta x} \ \frac{\delta v_y}{\delta x} \ \frac{\delta v_x}{\delta y} \ \frac{\delta v_y}{\delta y} \ \frac{\delta v_x}{\delta t} \ \frac{\delta v_y}{\delta t}]$$

The behaviour of the two-frame algorithm and the multi-frame algorithm can be conceptualized as a control system as shown in Figures 23a and 23b. Parallel implementations could be designed for the system, thus solving for the structure - motion parameters for each surface separately. In fact, solving the linear system at each iteration could also be parallelized. Extra processing is needed to segment the polyhedra-like hand into separate planar surfaces.

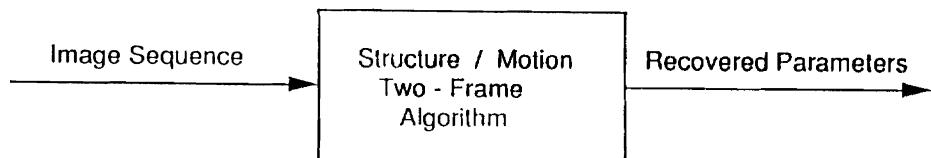


Figure 23.a : Two - Frame Algorithm

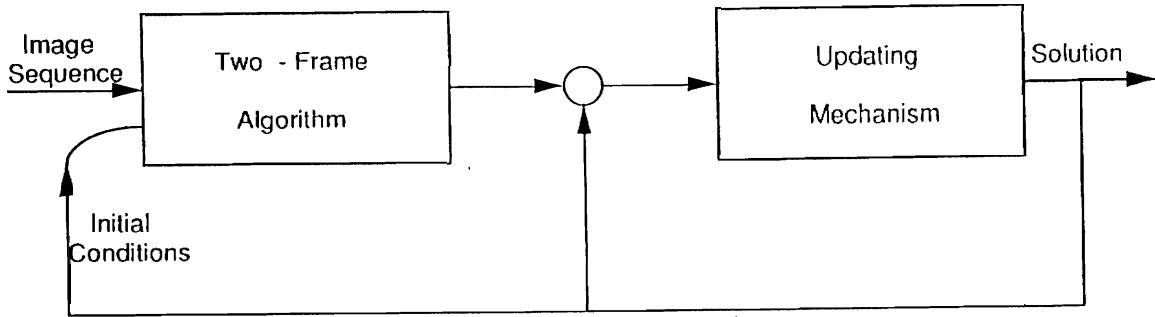


Figure 23.b: Multi - Frame Algorithm

8.2 Other Algorithms

There are other non-iterative techniques for recovering the 3-D parameters resulting from 2-D motion between two frames. The methods that will be mentioned here rely on specific assumption regarding the hand's geometry and/or world manipulating actions. Assuming that the actual relations between feature points that lie on the hand plane is well defined than a closed form solution for the structure parameters and depth can be estimated by using a method like the one described by Fischler and Bolles [12]. The motion parameters can then be easily recovered by solving a linear system in six parameters.

It should be noticed that we try to use alternative methods in order to make the system of equations “as linear as possible”, the motive behind that is the fact that linear systems can be solved in a pseudo-real time framework for a relatively small number of feature points and in addition a closed form solution always results. Another idea is to assume that the surface of the manipulating hand is frontal at the time of capturing the frame to be processed with the previous one, thus p and q are equal to zero, and the problem reduces to solving a linear system in six parameters for the motion parameters, while the depth is easily computed by knowing the 3-D distance between any two feature points, thus Z_o is equal to :

$$l^2 = \left[\left(\frac{x_2}{f - px_2 - qy_2} - \frac{x_1}{f - px_1 - qy_1} \right)^2 + \left(\frac{y_2}{f - px_2 - qy_2} - \frac{y_1}{f - px_1 - qy_1} \right)^2 + \left(\frac{1}{1 - \frac{px_2}{f} - \frac{qy_2}{f}} - \frac{1}{1 - \frac{px_1}{f} - \frac{qy_1}{f}} \right)^2 \right]$$

where f is the focal length of the lens, l is the real 3-D distance between two feature points on the hand and (x_1, y_1) and (x_2, y_2) are the CCD coordinates of the two image points.

The assumption here being that the observer always locates itself to a position in which

the hand is frontal with respect to the camera image plane, and that manipulating movements while the camera is moving and during computations is negligible. Other formulations may attempt to find pseudo-close form solution of the non-linear second order system and other assumptions, like the absence of rotational and/or translational motion reduces the complexity significantly.

8.3 Recovering 3-D Uncertainties

Having discussed methods for computing the three dimensional motion vectors and structure parameters between two image frames, we now use the same formulations described earlier for 3-D recovery but using 2-D error distributions as estimates for motion and/or feature coordinates in order to compute 3-D uncertainty distributions for the real world motion vectors and structure instead of single values for the world events.

As an example to illustrate the idea, let's assume that we have a linear system of equations as follows :

$$x + 3y = z_1$$

$$2x + y = z_2$$

The solution of this system is very easily obtained as

$$x = \frac{3}{5}z_2 - \frac{1}{5}z_1$$

$$y = \frac{2}{5}z_1 - \frac{1}{5}z_2$$

That is, a linear combination of the right hand side parameters. If the parameters z_1 and z_2 were random variables of known probability distributions instead of constants, then the problem becomes slightly harder, which is, to find the linear combination of those random variables as another random variable. The obvious way of doing this would be to use convolutions and the formula :

$$P_{X_1+X_2}(y) = \sum_R P_{X_1, X_2}(x, y - x)$$

for the sum of two random variables X_1, X_2 for any real number y and/or the formula for linear combinations over the region R , which is for all x such that $P_{X_1, X_2}(x, y - x) > 0$.

Using the moment generating function or the characteristic function seems also to be a very attractive alternative. The moment generating function M of a linear combination of random variables, for example X_1, X_2 can be written as :

$$M_{aX_1+bX_2+c}(t) = \exp^{ct} (M_{X_1}(at)M_{X_2}(bt))$$

for independent random variables X_1, X_2 . That is, the problem of solving linear systems on the form $Ax = b$, where b is a vector of random variables, may be reduced to finding closed form solutions for x in terms of the random parameters (using any elimination technique) and then manipulating the results and finding different expectations using moment generating or characteristic functions.

The 2-D to 3-D conversion problem, as discussed in details earlier, is a non linear system on the form $F(x) = b$ where b is the vector of 2-D random variables obtained from the previous levels. An approach to solving this system might be to try and approximate mathematically the problem to finding the roots using an iterative technique which calculates the Jacobian at every iteration and use Newton's method iterative formula for an over-determined system at the n th step as follows :

$$X_{n+1} = \left(J^T(X_n)J(X_n) \right)^{-1} J^T(X_n)F(X_n)$$

where J is the Jacobian of the system, however the Jacobian and F will contain positional and motion random variable nonlinear combinations at every iteration and we may have to use the following formulae for product and quotient of random variables :

$$f_{XY}(y) = \int_R \frac{1}{|x|} f_{X,Y} \left(x, \frac{y}{x} \right) dx$$

$$f_{\frac{X}{Y}}(y) = \int_R |x| f_{X,Y}(yx, x) dx$$

Obviously, such elaborate computations at every stage of descending towards a solution for the non-linear combinations of random variables is very hard and expensive to compute in pseudo-real time, if not impossible.

The solutions we suggest to this problem of finding the random variable solution for the 3-D parameters utilize the techniques we described in the previous two subsections. Using

either the two-frame iterative technique or the closed form algorithms, it should be noticed that the problem reduces to either solving multi-linear systems or a single one. In that case, using elimination and characteristic functions for computing the required expectations and/or distributions is straight forward, as all the systems become linear or pseudo-linear. In the iterative two-frame algorithm expectations can be used to avoid multiplication of random variable estimates for the structure and translational parameters when solving for the rotational random variable error parameters. Also, the same can be used for the positional parameters on the CCD camera array.

Thus, we have suggested algorithms for the quick estimation of the 3-D uncertainties in the structure and motion of the manipulation system. The next step would be to refine these estimates and use them for asserting the world events with uncertainty modeling and compensation. This will be described in the following two sections.

9 Refining World Events

In this section we describe techniques for eliminating and refining the 3-D models of manipulation under observation, whose recovery was discussed in the previous sections. In particular, we discuss a strategy to reject improbable events that might have been computed due to noise and uncertainties that were not compensated for in the distribution formulation, also because of unsmooth visual artifacts. We employ both existing knowledge about the mechanical properties of the manipulation and also knowledge from the current state of the observer automaton.

We concentrate our treatment of the subject on the three dimensional behaviour of the hand that is used in manipulation. The hand is assumed to be a well defined entity, and as we mentioned before, changing the hand and/or its characteristics can be modeled by simply plugging in a module that describes the new characteristics, the *same* hand is used throughout the entire manipulation activities.

Knowing the joint limits of the manipulating robot will enable us to reject improbable recovered 3-D motion vectors, that could not have occurred in the real 3-D world. As an example, assuming that we use a gripper with two "claws" having only one degree of freedom, then, obviously, any recovered 3-D rotational velocities for the claws should be rejected. Unreal-

istic slope estimations should also be rejected, knowing the robotic reachability of the end effector, with respect to the viewer.

The current position in the observer automata will allow refining the recovered 3-D event distributions, as it might well be the case that impossible manipulation actions at a specific manipulation stage are recovered. It is impossible, for example, due to the visual sampling rate, that the hand is in an upright position holding a nail in the center of the image plane at a time step, then having it disappear or hold another object at a dramatically distant 3-D position in the next time step, unless, of course a manipulation or viewer system failure has happened. In that case, some designated fail state should be accessed, discarding the recovered parameters. Limits on V_X , V_Y , V_Z , Ω_X , Ω_Y , Ω_Z and Z are asserted for every observer subset of states, and used for refining the recovered 3-D world events.

10 Navigating the Observer Automaton

At this point in the hierarchy of recovery and uncertainty levels, we have established methods and algorithms for recovering the refined three dimensional velocity and structure of the scene under observation. In addition, we computed the distribution of the uncertainty in the numerical values of the parameters in real-time. For example, the computed value for the translational velocity V_X might be a random variable lying between two values V_1 and V_2 with a known probability distribution \mathcal{F} . The same applies for all the other parameters for the different components in the scene.

The problem now is how to make use of these distribution values in order to be able to navigate in the observer automaton as defined in section 2 and demonstrated by examples in section 3. In other words, having built the DEDS automaton model of the visual system and its observer, we have a set of events that are defined as ranges on the visual scene parameters that cause state transitions between the automaton states. For example, there might be two different events branching from a state in some screwing task observer automaton and causing state transitions to two other states, and a self loop caused by the continuous dynamics within a coarse quantization of a DEDS state, as follows :

$$S_1 \xrightarrow{e_1: \Omega_1 < \Omega_Y < \Omega_2} S_2$$

$$\begin{array}{ccc}
 & e_2: -\Omega_1 < \Omega_Y < \Omega_1 & \\
 S_1 & \xrightarrow{\hspace{2cm}} & S_1 \\
 & e_3: -\Omega_2 < \Omega_Y < -\Omega_1 & \\
 S_1 & \xrightarrow{\hspace{2cm}} & S_3
 \end{array}$$

In addition to other limits on the other scene parameters. That is, if Ω_Y occurs within a specific range, then the corresponding state transition should be asserted according to the above set of event description.

The problem then reduces to computing the corresponding areas under the refined distribution curves obtained from the hierarchy levels. In the case of the presence of more than a single parameter in the transition event description, then the corresponding area under *each* parameter curve should be computed and multiplied for each parameter in the event definition. The goal is to find the probability of the occurrence of each event. In the above example, the goal would be to find the probability of e_1, e_2 and e_3 .

An obvious way of using those probability values is to establish some threshold values and assert transitions according to those thresholds. For example, if for any event in the set (e_1, e_2 and e_3), the computed probability of the range is > 0.85 , then the corresponding state transition should be asserted. It should be noted that those threshold values are highly task and state-dependent, appropriate values for the thresholds can be determined by performing many experiments for different task descriptions. The thresholds can also be updated adaptively according to the current manipulation patterns under observation. Many problems may arise after having obtained the above probabilities at the current automaton state. It might be the case that none of the obtained probability values exceeds the set threshold value and/or all values are very low. In that case, there is a good chance that we are at either the wrong automata state, or that a gross error has occurred in manipulation or some system failure.

The remedy to such problems can be implemented through time proximity, that is, wait for a while (which is to be preset) till a strong probability value is registered and/or *reback* in the automaton model for the observer till a high enough probability value is asserted, a fail state is reached or the initial ambiguity is asserted. The rebacking strategy can be implemented using a stack-like structure associated with each state that has already been traversed. A

stack of the latest computed probability values sorted in descending order as an index to the corresponding event. As soon as a forward traversal is performed, the top value should be popped. Rebacking can be done by using the top of the stack value and do the corresponding transition and compute the new probabilities for the events. For states that have not been visited at all, new stacks and computations should be performed.

Having established techniques for navigating the observer, the model description is now completed. The formulation uses uncertainties to assert current states of the manipulation system and attempts to recover from mistakes and errors. The model uses different intermediate levels for computing uncertainties, from the sensor level to the observer automaton level. Next, we discuss some results and discuss our approach. Then, we suggest ideas for extensions and future research.

11 Results

A substantial portion of the proposed system is already implemented and tested. Experiments were performed to observe the robot hand. The Lord experimental gripper is used as the manipulating hand. Different views of the gripper are shown in Figures 24.a to 24.c. Feature tracking is performed for some dots on the gripper in real time, using the Maxvideo system. Approximate algorithms to allow linearizing the optical system are used as described in section 8.2. A static look-and-move strategy was then used for tracking the hand features. The visual tracking system works in real time and a position control vector is supplied to the observer manipulator.

The 2-D uncertainty levels were tested. Edge detection with uncertainty is performed using different noise levels as shown in section 6, the enclosing “envelopes” were determined for the mechanical system and plotted in 3-D in section 7, the rejection algorithms are completed. A grasping task using the Lord gripper, as seen by the observer, is shown in Figure 25.a to 25.d. The sequence is defined by our model, and the visual states correspond to the gripper movement as it approaches an object and then grasps it.

The image flow algorithm described in section 4.1 is tested on the image of the gripper. The 2-D flow vectors resulting from the detection algorithm when applied to diagonal movements of the gripper’s image are shown in Figures 26.a and 26.b. The motion was upwards to the

left and downwards to the right. It can be seen that the resulting optic flow vectors are consistent with the actual motion. The image gaussian pyramid of the gripper is shown in Figure 27, the pyramid is formed by successive applications of gaussian low-pass filtering and decimation by half, five levels of the pyramid are shown. Simple segmentation and edge tracing are shown in Figure 28 and 29, as applied to the hand. Thus, event identification for the motion of the hand is computed. Tracking mechanisms are demonstrated and shown to work in real-time to follow the hand, uncertainty levels are also developed.

12 Discussion

We have proposed a new approach to solving the problem of observing a moving agent. In particular, we described a system for observing a manipulation process. Our approach uses the formulation of discrete event dynamic systems as a high-level model for the framework of evolution of the hand/object relationship over time. The proposed system utilizes the a-priori knowledge about the domain of the manipulation actions in order to achieve efficiency and practicality.

We started by describing the automaton model of a discrete event dynamic system then proceeded to formulate frameworks for the manipulation processes, and the observer construction. We developed efficient low-level event-identification mechanisms for determining different manipulation movements in the system and for moving the observer. Next, we defined and constructed six different levels for converting the raw 2-D image data into meaningful 3-D descriptions of the world events. The formulation includes computing uncertainty models resulting from errors in the 2-D and 3-D recovery mechanisms. The formulation allows the observer to navigate in real time with a stable behaviour through the automaton state space and thus assert world events efficiently.

The approach used can be considered as a frame work for a variety of visual tasks, as it lends itself to be a practical and feasible solution that uses existing information in a robust and modular fashion. The work examines closely the possibilities for errors, mistakes and uncertainties in the manipulation system, observer construction process and event identification mechanisms. Ambiguities are allowed to develop and are resolved after finite time, recovery mechanisms are devised too. Theoretical and experimental aspects of the work supports

adopting the framework as a new kind of basis for performing many task-oriented recognition, inspection and observation of visual phenomena. In the next section we examine extension ideas and future research opportunities for which the formulation can be considered as the backbone.

13 Extensions and Future Research

The proposed formulation can be extended to accommodate for more manipulation processes. Increasing the number of states and expanding the events set would allow for a variety of manipulating actions. The system can be made more “modular” by constructing a general automaton model of a discrete event dynamic system and defining the states, events and the certainty thresholds for them in an *automatic* way through a learning stage. In other words, different manipulation actions can be performed and “shown” to the observer and then the possible states, events and sequences of operations are automatically embedded in the general dynamic model. Thus, the manual formulation of the DEDS model for the task would not be needed anymore.

More powerful models for the DEDS could be sought, for example, context sensitive grammars, pushdown automata, turing machines and/or μ -recursive functions. The model building process can be thought of as forming a compiler with the object, sensor, task description and learning model as inputs, and the algorithm to follow the observer automaton with uncertainty as the output. Feedback can be supplied to the manipulating system in order to correct its actions, thus closing the vision-manipulation loop. The system could be generalized to an arbitrary number of mobile manipulating robots and mobile observing ones, a scheme would have to be devised to allow for distributed and parallel control of the observation and feedback process in an efficient way and to prevent deadlock and/or starvation problems.

The characteristics of the workspaces of both the manipulating robot and the observer can be utilized in order to avoid problems like collision and occlusion. This might be necessary to explore if both workspaces intersect in a 3-D volume. This can occur in a simple laboratory setup with two fixed manipulators, visualizing the volume of intersection and the holes and voids [1] within each robot reachable workspace will be necessary for planning and

constructing the model and its observer.

Foveal and peripheral vision strategies can be applied to “focus” on a specific aspect of the scene under considerations, according to the present observer state. Pyramid approaches for locating actions can be used. Logarithmic sensors, like cameras whose CCD array resembles the human eye can be utilized as the observer’s visual sensor for shifting attention to the interesting parts of the image.

Parallelizing the whole process by forming simultaneous observers can be explored. This will be necessary in case of multiple observing robots, manipulating robots and/or different kinds of sensors (tactile, range, vision ..etc) so as to allow for modular and efficient planning, “seeing” and recovery mechanisms. Inter-parallelization of different algorithms should be explored too. Overcoming delays in communication links between different observers and between the vision, control and parallelization modules within the same observer module should be addressed, specially if the modules are *physically* distant within the laboratory setup. Overcoming delays when feedback is supplied to the manipulating hand would be necessary.

The idea of DEDS as skeletons for observation under uncertainty can be explored further to allow for various other visual tasks. We discussed observing manipulation as a subset of observing moving agents, however, similar formulation can be described for other tasks, like recognizing stationary objects with optimal observation costs, i.e, minimal motion events. Perturbation analysis [17,35] can be performed for the average task behaviour of frequent visual events within a specified manipulation domain. Disappearing objects and partially occluded objects can also be recognized optimally using the proposed scheme, using time proximity as another dimension for asserting the identity of different targets, that is, allow recognition and/or tracking to be completed within a pre-specified, task-dependent time frame.

References

- [1] T. Alameldin, “Visualization of 3-D Workspaces”, Ph.D. Thesis, Computer and Information Science Department, University of Pennsylvania, *In preparation* (1991).

- [2] J. Aloimonos and A. Bandyopadhyay, "Active Vision". In *Proceedings of the 1st International Conference on Computer Vision*, 1987.
- [3] P. Anandan, "A Unified Perspective on Computational Techniques for the Measurement of Visual Motion". In *Proceedings of the 1st International Conference on Computer Vision*, 1987.
- [4] H. L. Anderson, *GRASP lab. Camera Systems and Their Effects on Algorithms*, Technical Report MS-CIS-88-85 and GRASP lab. TR 161, University of Pennsylvania, 1988.
- [5] R. Bajcsy, E. Krotkov and M. Mintz, *Models of Errors and Mistakes in Machine Perception*, Technical Report MS-CIS-86-26 and GRASP lab. TR 64, University of Pennsylvania, 1986.
- [6] R. Bajcsy, "Active Perception", *Proceedings of the IEEE*, Vol. 76, No. 8, August 1988.
- [7] R. Bajcsy and T. M. Sobe, *A Framework for Observing a Manipulation Process*. Technical Report MS-CIS-90-34 and GRASP Lab. TR 216, University of Pennsylvania, June 1990.
- [8] N. M. Benahmed, *Camera Calibration for Dynamic Environment*. M.S. Thesis, Department of Electrical Engineering, University of Pennsylvania, 1989.
- [9] P. J. Burt, C. Yen, and X. Xu, "Multiresolution Flow-Through Motion Analysis". In *Proceedings of the 1983 IEEE Conference on Computer Vision and Pattern Recognition*.
- [10] P. J. Burt, et al., "Object Tracking with a Moving Camera", *IEEE Workshop on Visual Motion*, March 1989.
- [11] E. S. Deutsch and J. R. Fram, "A Quantitative Study of the Orientation Bias of some Edge Detector Schemes", *IEEE Trans. Comput.*, C-27, No. 3, March 1978.
- [12] M. Fischler and R. C. Bolles, *Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography*, Readings in Computer Vision, Morgan Kaufmann Publishers, 1987.

- [13] J. R. Fram and E. S. Deutsch, "On the Quantitative Evaluation of Edge Detection Schemes and Their Comparison with Human Performance", *IEEE Trans. Comput.*, C-24, No. 6, June 1975.
- [14] N. M. Grzywacz and E. C. Hildreth, *The Incremental Rigidity Scheme for Recovering Structure from Motion: Position vs. Velocity Based Formulations*, MIT A.I. Memo No. 845, October 1985.
- [15] D. J. Heeger, *Models for Motion Perception*. Ph.D. Thesis, Computer and Information Science Department, University of Pennsylvania, September 1987.
- [16] J. Heel, "Dynamic Motion Vision", In *Proceedings of the SPIE Conference on Computer Vision*, November 1989.
- [17] Y. Ho, "Performance Evaluation and Perturbation Analysis of Discrete Event Dynamic Systems", *IEEE Transactions on Automatic Control*, July 1987.
- [18] J. E. Hopcroft and J. D. Ullman, *Introduction to Automata Theory, Languages and Computation*, Addison-Wesley, 1979.
- [19] B. K. P. Horn and B. G. Schunck, "Determining Optical Flow", *Artificial Intelligence*, vol. 17, 1981, pp. 185-203.
- [20] B. K. P. Horn, *Robot Vision*, McGraw-Hill, 1987
- [21] A. Izaguirre, P. Pu and J. Summers, "A New Development in Camera Calibration: Calibrating a Pair of Mobile Cameras", In *Proceedings of the International Conference on Robotics and Automation*, pp. 74-79, 1985.
- [22] D. Keren, S. Peleg and A. Shmuel, *Accurate Hierarchical Estimation of Optic Flow*, TR-89-9, Department of Computer Science, The Hebrew University of Jerusalem, June 1989.
- [23] Z. Kohavi, *Switching and Finite Automata Theory*, McGraw-Hill, 1979.
- [24] E. Krotkov, *Results in Finding Edges and Corners in Images Using the First Directional Derivative*, Technical Report MS-CIS-85-14 and GRASP lab. TR 37, University of Pennsylvania, 1985.

- [25] H. R. Lewis and C. H. Papadimitriou, *Elements of the Theory of Computation*, Prentice-Hall, 1981.
- [26] Y. Li and W. M. Wonham, “Controllability and Observability in the State-Feedback Control of Discrete-Event Systems”, *Proc. 27th Conf. on Decision and Control*, 1988.
- [27] H. C. Longuet-Higgins and K. Prazdny, *The interpretation of a moving Retinal Image*, Proc. Royal Society of London B, 208, 385-397.
- [28] C. M. Özveren, *Analysis and Control of Discrete Event Dynamic Systems : A State Space Approach*, Ph.D. Thesis, Massachusetts Institute of Technology, August 1989.
- [29] T. Peli and D. Malah, “A Study of Edge Detection Algorithms”, *Computer Graphics and Image Processing*, vol. 20, 1982, pp. 1-21.
- [30] P. J. Ramadge and W. M. Wonham, “Supervisory Control of a Class of Discrete Event Processes”, *SIAM Journal of Control and Optimization*, January 1987.
- [31] P. J. Ramadge and W. M. Wonham, “Modular Feedback Logic for Discrete Event Systems”, *SIAM Journal of Control and Optimization*, September 1987.
- [32] G. E. Révész, *Introduction to Formal Languages*, McGraw-Hill, 1985.
- [33] T. M. Sohn and K. Wohm, “Recovery of 3-D Motion and Structure by Temporal Fusion”. In *Proceedings of the 2nd SPIE Conference on Sensor Fusion*, November 1989.
- [34] M. Subbarao and A. M. Waxman, *On The Uniqueness of Image Flow Solutions for Planar Surfaces in Motion*, CAR-TR-113, Center for Automation Research, University of Maryland, April 1985.
- [35] Rajan Suri, “Perturbation Analysis : The State of the Art and Research Issues Explained via the GI/G/1 Queue”, *Proc. of the IEEE*, January 1989.
- [36] R. Y. Tsai and T. S. Huang, “Estimating three-dimensional motion parameters of a rigid planar patch”, *IEEE Transactions on Acoustics, Speech and Signal Processing*, ASSP-29(6), December 1981.
- [37] R. Y. Tsai, “An Efficient and Accurate Camera Calibration Technique for 3-D Machine Vision”, IBM Report.

- [38] S. Ullman, “Analysis of Visual Motion by Biological and Computer Systems”, *IEEE Computer*, August 1981.
- [39] S. Ullman, *Maximizing Rigidity: The incremental recovery of 3-D structure from rigid and rubbery motion*, AI Memo 721, MIT AI lab. 1983.
- [40] A. M. Waxman and S. Ullman, *Surface Structure and 3-D Motion From Image Flow: A Kinematic Analysis*, CAR-TR-24, Center for Automation Research, University of Maryland, October 1983.
- [41] J. Weng, T. S. Huang and N. Ahuja, “3-D Motion Estimation, Understanding and Prediction from Noisy Image Sequences”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-9(3), May 1987.
- [42] K. Wohn and S. R. Maeng, “Real-Time Estimation of 2-D Motion for Object Tracking”, In *Proceeding of the SPIE Conference on Intelligent Robotics*, November 1989.

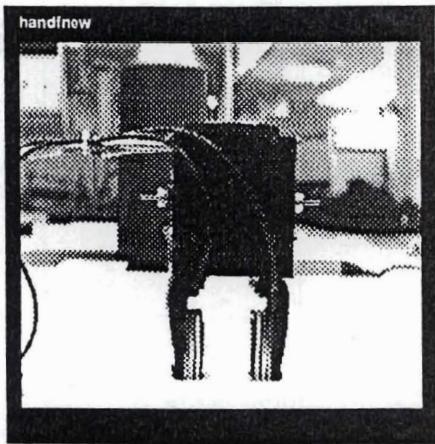


Figure 24.a



Figure 24.b



Figure 24.c

Different Views of the Lord Gripper

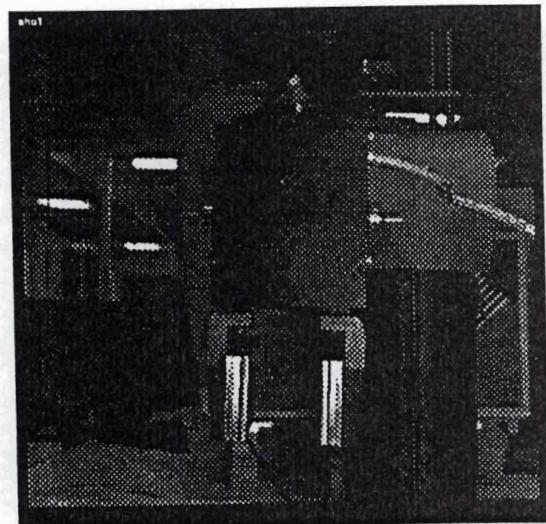


Figure 25.a

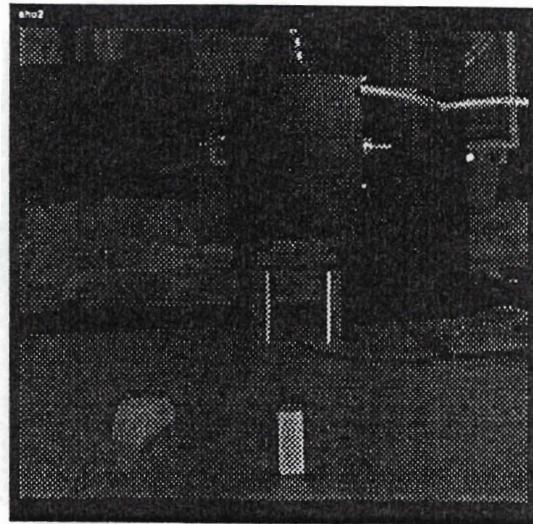


Figure 25.b

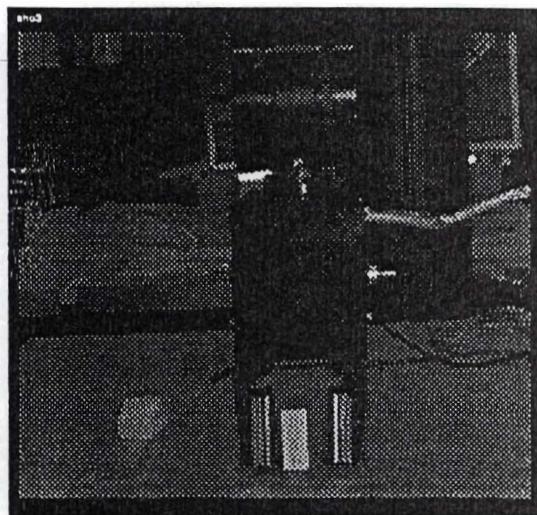
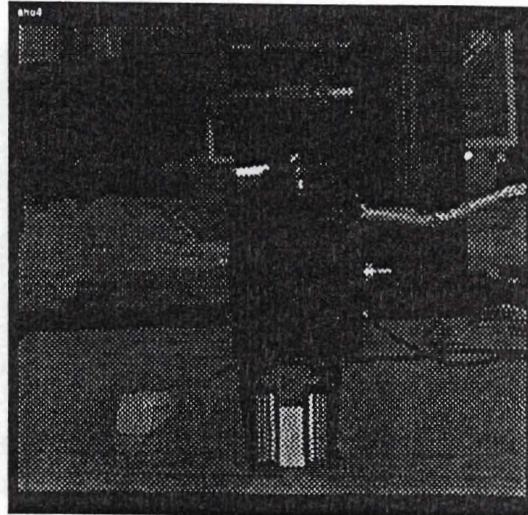


Figure 25.c



The Grasping Task

Figure 25.d

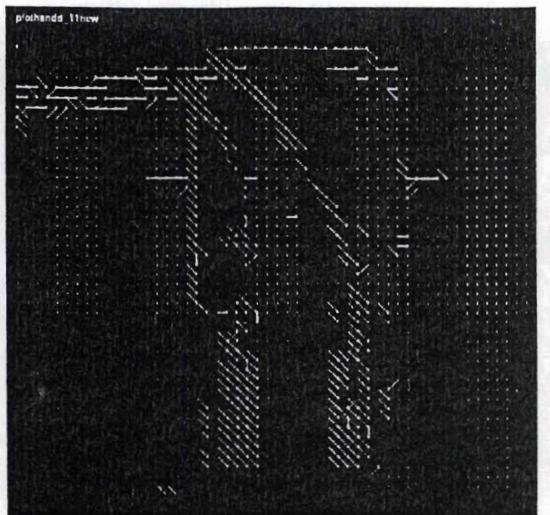


Figure 26.a

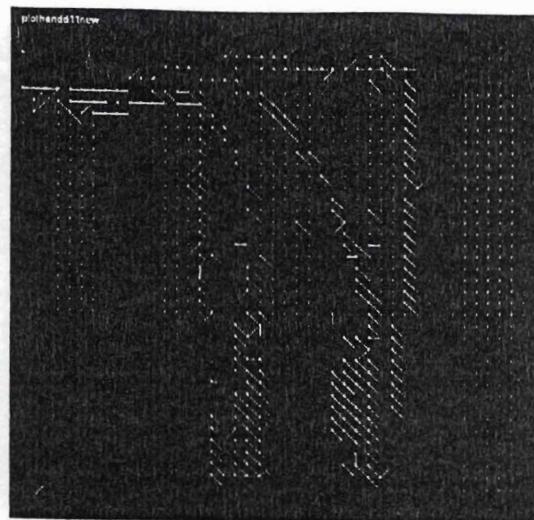


Figure 26.b

Optic Flow Vectors

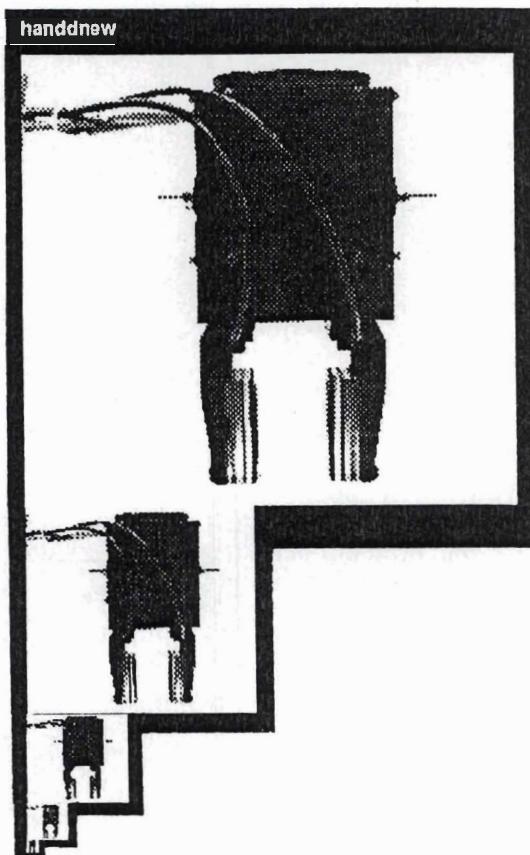


Figure 27 : Gaussian Pyramid of the Hand

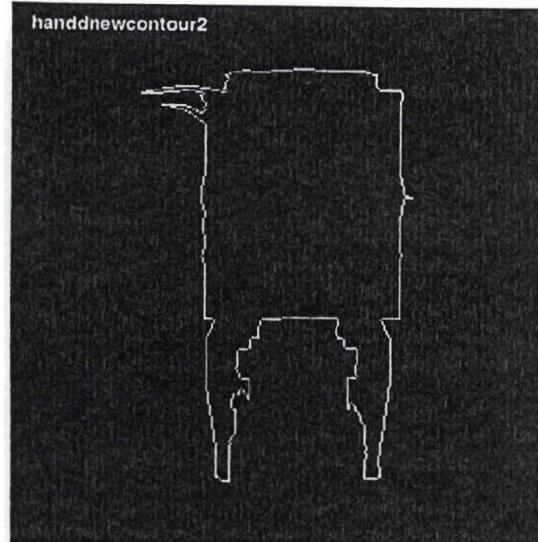


Figure 28 : Results of Edge Tracing



Figure 29 :
Results of 2-D Segmentation to Find the Hand