


LABORATORY FOR
COMPUTER SCIENCE



MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY

MIT/LCS/TR-563

**A CONSTRUCTIVIST
APPROACH TO ARTIFICIAL
INTELLIGENCE REEXAMINED**

Robert Matthew Ramstad

January 1993

545 TECHNOLOGY SQUARE, CAMBRIDGE, MASSACHUSETTS 02139

This blank page was inserted to preserve pagination.

A Constructivist Approach to Artificial Intelligence

Reexamined

by

Robert Matthew Ramstad

Submitted to the Dept. of Electrical Eng. and Computer Science
on July 8, 1992, in partial fulfillment of the
requirements for the degrees of
Master of Science
and
Bachelor of Science

Abstract

“Made-Up Minds: A Constructivist Approach to Artificial Intelligence”, a Ph.D. thesis by Gary Drescher (MIT, Computer Science, September 1989) and a book published by MIT Press (1991) describe a learning system which controls a simulated robot and gathers information about causes and effects for various actions within the software simulated world the robot inhabits. Beliefs about causality in this world are constructed through a learning process driven by the continuous updating of statistics. Each belief, or *schema*, held by the system has an associated reliability factor, and the system is able to iteratively improve both the reliability and scope of its knowledge base by revising and strengthening previously held beliefs on the basis of new statistically significant information. At any point in time, the amount of knowledge acquired by the system can be determined by direct examination of the schema structures.

Unfortunately, Drescher’s system is computation- and hardware-intensive. This report documents the reimplementaion of this learning system from the ideas in the thesis and book alone, using Common LISP and a UNIX workstation. Execution of the reimplementaion code indicates that Drescher’s results are implementation independent and directly attributable to the ideas in his published works. Results not discussed in Drescher’s works were also discovered.

Thesis Supervisor: Ronald L. Rivest

Title: Professor, Dept. of Electrical Eng. and Computer Science

Thesis Supervisor: Bruce A. Foster

Title: Principal Software Engineer, Digital Equipment Corporation

Keywords

schema mechanism, constructivism, learning by experience, learning by statistics, knowledge acquisition, artificial intelligence

Important Note

This technical report is not the complete thesis submitted to MIT. Rather, it represents the original document less appendices. Pages 101 through 288 are not included, and references to these pages are made in the table of contents and elsewhere in the document.


The complete thesis in PostScript form and other related files (LISP source code, etc.) can be acquired through anonymous FTP to [theory.lcs.mit.edu](ftp://theory.lcs.mit.edu). After providing a valid email address as the password, do “cd pub”, “cd ramstad” and “get README”. The README file explains the contents of the various archives.

Please contact the author at ramstad@theory.lcs.mit.edu if this procedure does not work.

Acknowledgments

I'd like to acknowledge the help of my thesis supervisors, Professor Ronald Rivest, MIT, and Bruce Foster, Digital Equipment Corporation, for their assistance, cooperation and patience throughout all phases of this thesis. The assistance of the VAX LISP group at Digital Equipment Corporation was invaluable in the early going, particularly the help of Paul Anagnostopoulos (macros), Jeff Piazza (dark secrets of LISP compilers), Walter van Roggen (code optimization) and Richard Wells (general Common LISP). Professor Pattie Maes and the MIT Media Lab Music and Cognition group have been a great bunch of people to work with over the last year. I'd like to thank Kevin O'Toole, the VI-A program staff and Marilyn Pierce for assistance throughout with various bits of thesis paperwork and administrative red tape. Robyn Kozierok, MIT Media Lab and AI Lab, assisted me greatly with an unbiased and frank opinion of a draft of this thesis. My fraternity, Beta Theta Pi, has been and continues to be a very important influence on my life, and the brothers of the MIT chapter have been instrumental in helping me keep things in focus over the years. Allison Nelson helped me through some of the worst times when I had been working on the thesis for a while and I was frustrated with the slow progress of the work. Last, but not least, I'd like to thank my mom Sheryl and sister Kristina for all their support — “Why don't you just get a job, Bob?”. All of these people were instrumental in keeping me sane on the long road to the completion of this work.

I would also like to thank my favorite musical artists: David Bowie, Yes, King Crimson, The Who, Peter Murphy / Bauhaus, acoustic Neil Young, Talking Heads, Shriekback and especially The Grateful Dead. Their work has enriched my life and probably improved my code as well. (I usually listen to music when I program, it seems to keep me better focused on the task at hand.)




Contents

1	Introduction	13
1.1	Motivation for work	15
1.2	Sources	16
1.3	The microworld	17
1.4	The schema mechanism	22
1.4.1	Schemas	23
1.4.2	Constructing new schemas via marginal attribution	24
1.4.3	Goal-directed actions	29
1.4.4	Synthetic items	30
1.4.5	Control	31
1.4.6	Summary of schema mechanism	33
1.5	Performance	33
2	Implementation	35
2.1	Differences between this work and Drescher's	35
2.1.1	Goal-directed actions	35
2.1.2	Value	36
2.1.3	Control	37
2.1.4	Overriding conditions	38
2.1.5	Microworld	38
2.1.6	Schema mechanism	39
2.1.7	Reconciling different sources	40
2.1.8	Piagetian influence	40

2.2	Design decisions	41
2.3	Description of reimplementaion code	42
2.3.1	Microworld	42
2.3.2	Schema mechanism	43
3	Results	59
3.1	Test runs	59
3.1.1	Output	60
3.1.2	Analysis program	61
3.2	Drescher's results and schema categories	61
3.2.1	Initial schemas	62
3.2.2	Grasping schemas	62
3.2.3	Coarse visual field shifting schemas	62
3.2.4	Visual field shift limit schemas	64
3.2.5	Foveal region shift schemas	64
3.2.6	Detail shift schemas	65
3.2.7	Visual network schemas	66
3.2.8	Hand movement network schemas	67
3.2.9	Negative consequence schemas	68
3.2.10	Hand to body schemas	69
3.2.11	Seeing hand movements via coarse visual items	69
3.2.12	Seeing hand movements via detailed visual items	69
3.2.13	Further results of the CM2 implementation	71
3.3	Results confirming Drescher's work	71
3.4	Further results of the reimplementaion	72
3.4.1	Shifting the gaze to see the body	78
3.4.2	Using the body as a visual position reference	78
3.4.3	Hand network constraint schemas	79
3.4.4	Visual network constraint schemas	79
3.4.5	Coarse visual shift constraint schemas	79

3.4.6	Detailed visual shift constraint schemas	80
3.4.7	Coarse to detailed visual shift constraint schemas	80
3.4.8	Detailed to coarse visual shift constraint schemas	80
3.4.9	Detailed to coarse visual shift schemas	80
3.4.10	Seeing objects in different visual regions	81
3.4.11	Hand to body constraint schemas	81
3.4.12	Hand movement against object	82
3.4.13	Backward hand movement against body	82
3.4.14	Hand movement from coarse visual to foveal region schemas	82
3.4.15	Support data from test runs	83
4	Analysis, Discussion and Conclusions	89
4.1	Test run evidence confirms Drescher's results	89
4.2	Analysis of all results	94
4.3	Performance	97
4.4	Future directions	98
A	Source code	101
A.1	Microworld simulator	102
A.2	Schema mechanism	131
A.3	Support for fixnum math	231
A.4	Support for float math	233
A.5	Result analysis	235
B	Output	269
B.1	Partial output from sample run	270
B.2	Sample output from analysis program	281
	Bibliography	289



List of Figures

1-1	The hand can move within a 3x3 area in the area in front of the body.	18
1-2	The visual field is 5x5 and can assume nine different orientations, for a total visual region of 7x7.	18
1-3	Five foveal regions in the center of the visual field (front, back, right, left and center) provide detailed visual information.	19
1-4	The two objects in the microworld circulate in a clockwise direction among a series of four contiguous <i>home positions</i>	21
1-5	A sample microworld situation. The hand and the left object are in view, while the right object is out of view.	22
1-6	A basic schema.	24
1-7	A bare schema for the <i>grasp</i> action.	25
1-8	The extended result of the bare schema <i>/grasp/</i> detects that <i>/grasp/hgr</i> should be spun off.	25
1-9	The extended context of the schema <i>/grasp/hgr</i> discovers that <i>tactl</i> improves its reliability and therefore <i>tactl/grasp/hgr</i> is spun off. . . .	26
1-10	Two schemas which predict two items separately can not chain to a schema which has both items in its context: a schema with a conjunctive result is required.	28
2-1	The initial state of the microworld.	44

3-1 Shifting the gaze often causes an object which is within the visual field to appear at a new coarse visual item location. For example, the left hand object in this figure is moved from *vf04* to *vf14* when the glance is shifted to the left. 63

3-2 The detailed visual appearance of an object often shifts from one foveal region to another when a gaze action is taken. The hand in this figure is shifted from the rear foveal region to the center foveal region when the gaze is shifted to the rear. 65

3-3 Shifting the hand results in the hand moving to a new position. 68

3-4 Shifting the hand when it is visible often causes a transition to a new coarse visual item (and possibly to a new foveal region). 70

3-5 Through a large series of schemas, the reimplementation code learns that there are three different objects which each can appear in specific regions of the microworld. 81

List of Tables

1.1	The primitive actions	19
1.2	The primitive items	20
1.3	Rules for determining synthetic item state.	32
2.1	Simple datatypes defined and used by the schema mechanism code. .	45
2.2	Rules for using counters to detect statistically relevant occurrences. .	48
2.3	Compound datatypes defined and used by the schema mechanism code.	49
2.4	Rules for extended context statistics.	51
2.5	Rules for extended result statistics.	53
2.6	Notes on some of the functions in schema.lisp	57
3.1	Further results of Drescher's CM2 implementation, part one.	72
3.2	Further results of Drescher's CM2 implementation, part two.	73
3.3	Statistics for Drescher's classifications, test runs 1-7.	74
3.4	Statistics for Drescher's classifications, test runs 8-14.	75
3.5	Statistics for Drescher's classifications, test runs 15-20.	76
3.6	Statistic summary for Drescher's classifications.	77
3.7	Statistics for new classifications, test runs 1-7.	84
3.8	Statistics for new classifications, test runs 8-14.	85
3.9	Statistics for new classifications, test runs 15-20.	86
3.10	Statistic summary for new classifications.	87

Chapter 1

Introduction

“Made-Up Minds: A Constructivist Approach to Artificial Intelligence”, a Ph.D. thesis by Gary Drescher (MIT, Computer Science, September 1989) [Dre89] and a book published by MIT Press [Dre91], describe a learning system which gathers information about causes and effects while controlling a simulated robot which exists within a *microworld* (a software simulated world). The system he proposes, the *schema mechanism*, is novel in a number of ways. It falls firmly in the category of systems which learn from experience — it learns without any outside assistance. Claims have been made that these systems are “crucial to achieve successful behavior in complex, dynamic, unpredictable environments” [Mae92, p. 5] and as such they are particularly interesting systems to study. This system constructs beliefs about causality in the microworld through a learning process which is driven by the continuous updating of statistics. Each of these beliefs has an associated reliability factor, and the system is able to iteratively improve both the reliability and scope of its knowledge base by revising and strengthening previously held beliefs on the basis of new statistically significant information. The use of reliability factors and the iterative nature of improvement in the knowledge base are different from the methods found in many other learning systems, where the focus is usually on discovering *facts* about causality which are 100% reliable by exhaustive analysis of all possibilities within the problem domain itself. Also, the schema mechanism manipulates, stores and modifies schemas which represent beliefs and are easily understood by humans without using math-

ematical analysis tools. While many connectionist systems also have the ability to iteratively improve the reliability and scope of their knowledge bases, the derivation of the rules represented by the final configuration of the system is difficult. On the other hand, the final configuration of a schema mechanism run can be examined by a person and the highly reliable schemas can be analyzed directly to determine the amount of knowledge acquired by the system.

The schema mechanism provides control to a software simulated robot with body and hand which lives inside a microworld. Given a variety of possible actions and a vector of sensor data, the schema mechanism both attempts to reach states of high *value* (a quantitative measure of the desirability of a given state, based both upon the state itself and the range of states easily reachable from that state) and engages in behavior designed to improve its knowledge of the microworld. The schema mechanism is able to engage in a form of planning by constructing and maintaining *schemas* — structures which can be used to predict the result of taking a specified action in the current situation. The collection of schemas and certain other specialized structures comprises the knowledge the system has at each time step.

The results in the original thesis are startling. The original CM2 implementation of the schema mechanism, with virtually no initial knowledge about the microworld, manages to construct many reliable schemas. These *rules* include: how to grasp an object, how to move the hand from one position to an adjacent one, how to move the glance orientation from one position to an adjacent one, how to move an object closer to the center of the visual field so as to see its details, and then, through the construction of *goal-directed actions*, how to move the hand to any position, how to move the glance orientation to any position, how to move any object so as to see its details, and so on. In fact, the system appears to gain some idea of the concept of *objects* (through the construction of *synthetic items*), as well as some *intermodal* coordination (i.e. schemas which relate one sense to another, for example, schemas which indicate that moving the hand results in seeing the hand at the new location). The CM2 implementation was extremely successful in acquiring a large body of knowledge about the microworld. It is also notable that the progression of learning

exhibited by the program closely matches the progression postulated by Piagetian child development theory, where the concept of “schemata” was first analyzed.

1.1 Motivation for work

The CM2 implementation is relatively time-inefficient, and is also computation and hardware-intensive, utilizing over four thousand processors on a CM2 Connection Machine. A better implementation, in addition to dealing with these problems, might actually be able to learn more — the results of both [Dre91] and [Dre89] are clearly constrained by available memory.

Another purpose served by a new implementation is that of verification of the ideas in the original thesis. It is unclear if the results of the CM2 implementation are solely due to the design as detailed in [Dre91], or perhaps partially due to various specific unknown aspects of the implementation. In other words, the ideas in [Dre91] and [Dre89] may not be sufficient to account for all of the results generated by the implementation. Obviously, in any system of great complexity, seemingly minor implementation decisions may have unforeseen effects on the execution of the program. A new implementation can give concrete evidence that any system built as specified in the original thesis is capable of achieving comparable results.

This thesis documents the reimplementing of this learning system from the ideas in [Dre91] and [Dre89]. Where these two sources conflict, [Dre91] takes precedence. Where the description of the schema mechanism was not sufficiently detailed, the author of the original thesis was consulted. This document is organized as follows:

- **Chapter 1:** Motivation for work, overview of the schema mechanism.
- **Chapter 2:** Modifications and additions to the ideas in [Dre91] and [Dre89] and documentation of the new implementation.
- **Chapter 3:** Results from multiple executions of the new implementation.

- **Chapter 4:** Further analysis of the results described in chapter 3, comparison of the results found by each implementation. Suggestions for future experimentation with the schema mechanism.

Perhaps the most important goal of this project, however, is to encourage other researchers to experiment with the concepts embodied in the original thesis. To this end, the new implementation is written in Common LISP and executes on a general-purpose UNIX hardware platform and was designed with efficiency in mind. It is hoped that making the system available to researchers in a form which can easily be understood, executed and modified will assist the utilization of these ideas in future research.

1.2 Sources

Gary Drescher has published many different works on the schema mechanism, notably a recent book from MIT Press titled “Made-Up Minds” [Dre91] and his Ph.D. thesis [Dre89]. Every attempt has been made to be consistent with [Dre91], as it is the most recent major work. However, the reimplementing effort started before [Dre91] was published, and therefore [Dre89] was also heavily used. Both sources are useful for interpreting the other and therefore both are very valuable. Gary Drescher was also gracious enough to answer many questions via private electronic mail and telephone, which also assisted in constructing what hopefully is an accurate view of the schema mechanism. I freely borrowed (and condensed) from each of these sources as necessary while writing this overview — it is based primarily on the material in [Dre89] but agrees with [Dre91] in all major respects. If a deeper explanation is desired, I suggest reviewing the original sources, particularly [Dre91, section 1.1 (general), section 6.1 (microworld) and chapters 3 and 4 (schema mechanism)] and [Dre89, section 1.2 (general), section 3.1 (microworld) and section 3.2 (schema mechanism)]. The overview presented here definitely emphasizes those parts of Drescher’s schema mechanism which were implemented in detail, and glosses over parts of his system which were not implemented. The following chapter contains a detailed discussion of

the differences between this account, the accounts in [Dre91] and [Dre89] and what was actually implemented.

1.3 The microworld

The *microworld* is a separate system which is intended to be a primitive model of the real world. The area of the microworld is modeled as a 2 dimensional 7x7 grid — all vision is from a *birds-eye* viewpoint. Objects can be placed anywhere within the 7x7 area, but only one object can exist at any single coordinate position — objects are uniform in size, and cannot rotate.

The microworld is inhabited by a simulated robot with body, single hand and visual system which can perform actions in the microworld. The initially supplied *primitive actions* allow shifting of the visual region, movement of the hand, and grasping and ungrasping of objects. Each of these actions may or may not change the state of the microworld — in particular, actions which would take the hand or glance orientation beyond the allowable range have no effect. The hand can move within a 3x3 area near the body (see figure 1-1). Similarly, there are nine glance orientations which allow viewing of any particular 5x5 area within the 7x7 microworld (see figure 1-2). The primitive actions are designed to correspond roughly with the actions available to a stationary infant, and are further described in table 1.1 [Dre89, adapted from table 3.1, p. 66].

The simulated robot receives feedback on the current state of the microworld through roughly one hundred and forty *primitive items* — items which can be on or off (binary) and are directly related to conditions in the microworld. These primitive items include indications of hand position, indications of glance orientation, coarse visual information (an object is present within the visual region of the robot), detailed visual information (when an object is near the center of the visual region, see figure 1-3), tactile indications of the presence of an object (when an object is adjacent to the hand or body), detailed tactile information (when an object is to the left of the hand), detailed taste information (when an object is in front of the body), and indications as




Figure 1-1: The hand can move within a 3x3 area in the area in front of the body.




Figure 1-2: The visual field is 5x5 and can assume nine different orientations, for a total visual region of 7x7.

- **handf, handb, handr, handl**: These actions move the hand incrementally forward, backward, right or left.
- **eyef, eyeb, eyer, eyel**: These actions shift the glance orientation incrementally forward, backward, right or left.
- **grasp**: This action closes the hand, grasping any movable object which is immediately to the left of the hand (near its “fingers”) unless the hand was already closed. Once closed or grasping an object, the hand remains in that state for three time units, unless explicitly opened in the interim. Moving the hand moves any grasped object.
- **ungrasp**: This action opens the hand, releasing any object that had been grasped.

Table 1.1: The primitive actions

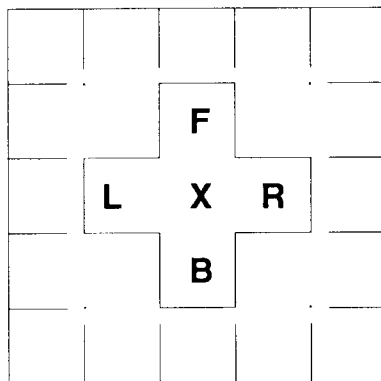


Figure 1-3: Five foveal regions in the center of the visual field (front, back, right, left and center) provide detailed visual information.

to whether or not the hand is closed and (possibly) grasping an object. The primitive items are further explained in table 1.2 [Dre89, adapted from table 3.2, p. 67].

There are two objects in the world, each of which occasionally (at an average of every 200 time units) moves between a series of four contiguous *home positions* in a clockwise direction — see figure 1-4. The right object is out of the range of the hand and therefore cannot be grasped. The left object can, of course, be grasped and moved about. Both objects are often seen by the simulated robot due to their proximity to the body.

- **hp00, . . . ,hp22**: Haptic-proprioceptive (hand position) items, one for each possible hand position, the hand is confined to a 3x3 area (see figure 1-1). Position (0,0) is in the lower left corner of the range; in figure 1-5, the hand appears at position (2,1) which corresponds to item hp21. In figure 1-1 the hand appears at hp10.
- **vp00, . . . ,vp22**: Visual-proprioceptive (visual position) items, one for each possible glance orientation. Coordinate designates the position of the center of the 5x5 visual field, using same conventions as for hand position; in both figure 1-5 and figure 1-2, the glance is oriented at vp01.
- **vf00, . . . ,vf44**: Coarse visual-field items, one for each of 25 glance-relative coordinate positions. Position (0,0) is in the lower left corner of the current visual field; in figure 1-5, the body appears at vf30 while the hand appears at vf42.
- **fovf00, . . . ,fovf33; fovb00-33; fovl00-33; fovr00-33; fovx00-33**: Visual details corresponding to each of five foveal regions: front, back, left, right and center. See figure 1-3. Each has sixteen arbitrary details. In figure 1-5 the left object is in the front foveal region.
- **tactf, tactb, tactr, tactl**: Coarse tactile items, one for each side of the hand: front, back, right, left.
- **bodyf, bodyb, bodyr, bodyl**: Coarse tactile items, one for each side of the body: front, back, right, left.
- **text0, . . . ,text3**: Detailed tactile items, denoting arbitrary textural details of an object touching the left edge (i.e. “fingers”) of the hand.
- **taste0, . . . ,taste3**: Detailed taste items, designating arbitrary surface details of an object touching the front edge (i.e. “mouth”) of the body/head.
- **hcl**: Hand closed.
- **hgr**: Hand closed and grasping something.

Table 1.2: The primitive items




Figure 1-4: The two objects in the microworld circulate in a clockwise direction among a series of four contiguous *home positions*.

The microworld uses three different coordinate systems, microspace, body and glance relative. In each case, the X axis (first position in the coordinate pair) runs left to right while the Y axis (second position) runs bottom to top. This is traditional first quadrant Cartesian coordinate notation [SESA86, p. 92]. Microspace relative coordinates reference the 7x7 world directly where the lower left hand corner is position (0,0) and the lower right hand corner is position (6,0). Body relative coordinates are often used when referring to the center 3x3 area in the microworld. The lower left corner of this area is microspace coordinate (2,2) which is defined as body relative coordinate (0,0). (Translation from microworld to body relative coordinates is accomplished by subtracting 2 from each microworld coordinate; similarly, body relative to microworld coordinate translation is accomplished by adding 2 to each body relative coordinate.) Glance relative coordinates are used when referring to the 5x5 area centered around the current glance orientation. The center of the 5x5 glance relative area is defined as glance relative coordinate (2,2) with the lower left corner of this area defined as glance relative coordinate (0,0).




Figure 1-5: A sample microworld situation. The hand and the left object are in view, while the right object is out of view.

In figure 1-5 the glance is oriented at body relative visual position (0,1) and the hand is at body relative hand position (2,1). The body is visible via the coarse visual field items at glance relative position (3,0), the hand at (4,2), and the left object at (2,3). The right object is not visible. The detailed visual information for the left object is present in the front foveal items. If the hand were in body relative hand position (1,2) adjacent to the left object, it could grasp it. The right hand object is currently out of reach. See table 1.2 for more examples using the various coordinate systems.

It is important to note that the names given to each primitive action and item are for purposes of human readability only. The microworld system provides a series of ten functions corresponding to the ten primitive actions and a series of functions which return the status of each of the primitive items. The schema mechanism begins with absolutely *no* knowledge about which actions and items are related to one another.

1.4 The schema mechanism

The schema mechanism, by utilizing the simulated robot in the microworld, attempts to acquire knowledge about its domain through analysis of its experiences. In the

appropriate situations, the system can create three different types of structures to embody acquired knowledge: *schemas*, *synthetic items* and *goal-directed actions*.

1.4.1 Schemas

Each schema expresses a specific belief about causality in the microworld and is defined by a context, action and result. The context defines the microworld pre-conditions under which the schema can be activated. If a schema is activated, its corresponding microworld action is executed. The result indicates those elements of the microworld state which should change when the schema is *activated* (context satisfied and action executed) — in some sense, indicating what the *effects* of the action are when performed in the given context. Essentially, each schema expresses the context-dependent results of a given action.

The context and result can be single items or conjunctions of items, or be empty. For each included item, the context and result indicate if it is positively or negatively included. A context is considered *satisfied* if each included item matches the current state of the microworld — if an item is positively included, it must be on in the current microworld state, and similarly for negative inclusion and off. A schema is considered *applicable* if its context is satisfied and no overriding conditions exist. (Overriding conditions are detected by a schema's *extended context*, see section 2.1.4 for more details.) If a schema is applicable and its action is taken, the schema has been *activated*. The result *obtains* if each item included in the result matches the state of the microworld after taking the action — if a schema is activated and its result obtains, the schema is said to *succeed*. Note that both primitive and synthetic items (discussed later) can be included in a context or result — however, synthetic items can also be in an *unknown* state, which for purposes of satisfying a context or achieving a result does not match positively or negatively included items.

A schema is notated in the form *context/action/result*, where negated items are indicated by a - and conjunctions of items are constructed by placing $\&$ signs between the items (by convention, the $\&$ can be omitted in the case of items with single letter names). For example, the schema in figure 1-6 is $p\text{-}qr/a/xy$.




Figure 1-6: A basic schema.

A schema is not a *rule* which indicates that the action should be performed when the context is satisfied; the schema just indicates what would happen *if* the action was performed. Note also that the results indicated by a schema are by no means guaranteed — a reliability measure, which indicates how often the result obtains when the schema is activated, is kept by each schema. Schemas may exist with arbitrarily low reliability — as a particular result does not necessarily follow with regularity, schemas cannot be thought of as rules. The notion of a rule also usually includes the notion that a given action should not be performed unless the preconditions are met. In this learning system, each action can be performed at any time — each schema merely asserts what happens when the action is performed when all context conditions are satisfied. The context therefore should *not* be considered a prerequisite for the performance of the action. It is also possible that items not included in the result will change state — the result is not necessarily exhaustive. Finally, a particular schema says absolutely nothing about what might happen if its action is taken when its context is not satisfied.

1.4.2 Constructing new schemas via marginal attribution

The system begins with a set of ten *bare schemas*, one for each primitive action. A bare schema has an empty context (one with no items), and therefore can be activated at any time. A bare schema also has an empty result, and therefore does not make any prediction whatsoever as to changes in the microworld state due to taking the indicated action (see figure 1-7).

As the system executes, a technique known as *marginal attribution* is used to discover statistically important context and result information. This information is then




Figure 1-7: A bare schema for the *grasp* action.




Figure 1-8: The extended result of the bare schema */grasp/* detects that */grasp/hgr* should be spun off.

used to fine-tune existing schemas by creating modified versions of them. Marginal attribution succeeds in greatly reducing the combinatorial problem of discovering reliable schemas from an extremely large search space without prior knowledge of the problem domain.

Result spinoffs

Many different results may occur from the execution of a given action. For every *bare* schema, this facility tries to find result transitions which occur more often with a particular action than without it. For example, my hand ends up closed and grasping an object much more often when the grasp action is taken than with any other action. Results discovered in this fashion are eligible to be included in a *result spinoff* — a new schema identical to its parent, but with the relevant result item included (see figure 1-8). The marginal attribution process can only create result spinoffs from bare schemas.

Specifically, each bare schema has an *extended result* — a structure for holding result correlation information. The extended result for each schema keeps correlation information for each item (primitive or synthetic). The positive-transition correlation is the ratio of the number of occurrences of the item turning on when the schema's action has been taken to the number of occurrences of the item turning on when the




Figure 1-9: The extended context of the schema */grasp/hgr* discovers that *tactl* improves its reliability and therefore *tactl/grasp/hgr* is spun off.

schema's action has not been taken. Similarly, the negative-transition correlation is the ratio of the number of occurrences of the item turning off when the schema's action has been taken to the number of occurrences of the item turning off when the schema's action has not been taken. Note that an item is considered to have *turned on* precisely when the item was off prior to the action and on after the action was performed and similarly for turning off. The correlation statistics are continuously updated by the schema mechanism and weighted towards more recent data. When one of these schemas has a sufficiently high correlation with a particular item, the schema mechanism creates an appropriate result spinoff — a schema with the item positively included in the result if the positive-transition correlation is high, or a schema with the item negatively included in the result if the negative-transition correlation is high. These simple statistics are very good at discovering arbitrarily rare results of actions, especially when the statistics of the non-activated schemas are only updated for *unexplained* transitions. A transition is considered explained if the item in question was included in the result of an activated schema with high reliability (above an arbitrary threshold) and it did, in fact, end up in the predicted state.

Context spinoffs

For schemas which have non-negligible results, the marginal attribution attempts to discover conditions under which the schema obtains its result more reliably. To extend the example, my hand ends up closed and grasping something *much* more often if I can feel an object touching the left edge of my hand before I close my hand with the grasp action. This information is used to create *context spinoffs* — duplicates of the parent schema, but with a new item added to its context (see figure 1-9).

Schemas with non-empty results have an *extended context*. For each item, this structure keeps a ratio of the number of occurrences of the schema succeeding (i.e. its result obtaining) when activated with the item on to the number of occurrences of the schema succeeding when activated with the item off. If the state of a particular item before activation of a given schema does not affect its probability of success, this ratio will stay close to one. However, if having the item on increases the probability of success, the ratio will increase over time. Similarly, if having the item off increases the probability of success, the ratio will decrease. If one of these schemas has a significantly high or low ratio for a particular item, the schema mechanism creates the appropriate context spinoff — a schema with the item positively included in the context if the ratio is high, one with the item negatively included if the ratio is low.

There is an embellishment to the process of identifying context spinoffs. When a context spinoff occurs, the parent schema resets all correlation data in its extended context, and keeps an indication of which item (positively or negatively included) was added to its spinoff child. In the future, when updating the extended context data for the parent schema, if that item is on (if positively included in the spinoff) or off (if negatively included in the spinoff) the trial is ignored and the extended context data is not modified. This embellishment means that the parent schema has correlation data only for those trials where there is no more specific child schema, and it encourages the development of spinoff schemas from more specific schemas rather than general schemas.

Redundancy is also reduced by a further embellishment. If at a particular moment in time, a schema has multiple candidates for a context spinoff, the item which is on least frequently is the one chosen for a context spinoff. The system keeps a *generality* statistic for each item which is merely its rate of being on rather than off — it is this statistic which is used when deciding between multiple spinoff possibilities. This embellishment discourages the development of unnecessary conjunctions when a single specific item suffices [Dre89, p. 104].

parent schema, but with a new item added to its context (see figure 1-9).




Figure 1-10: Two schemas which predict two items separately can not chain to a schema which has both items in its context: a schema with a conjunctive result is required.

Conjunctive contexts and results

The context can be iteratively modified through a series of context spinoffs to include more and more conjuncts in the context. For a variety of reasons, but primarily to avoid combinatorial explosion, a similar process for result conjunctions is undesirable [Dre89, pp. 105-6]. The marginal attribution process therefore requires that result spinoffs occur only from bare schemas, and only one relevant detail can be detected and used as the result for the spinoff schema. However, conjunctive results are necessary if schemas should be able to *chain* to a schema with a conjunctive context. (The goal-directed action facility in particular depends greatly on the detection of chains of reliable schemas where each schema has a result which satisfies the context of the next schema in the chain. See figure 1-10.) This problem is solved by adding a slot to the extended result of each bare schema for each of the conjunctions of items which appear as the context of a highly reliable schema. Statistics are kept for these in the same fashion as those kept for single items, and if one of these conjunctions is often turned on as the result of taking a given action, a result spinoff occurs which includes the entire conjunction in the result. Effectively, this process is able to produce schemas with conjunctive results precisely when such schemas are necessary for chaining.

Summary of marginal attribution

Schemas created by the marginal attribution process are designed to either encapsulate some newly discovered piece of knowledge about causality in the microworld (result spinoff) or to improve upon the reliability of a previous schema (context spinoff). By continuously creating new versions of previous schemas, the system iteratively improves both the reliability and the scope of its knowledge base. It is interesting to note that once created, a schema is never removed from the system. Rather, it may be supplanted by one or more spinoff schemas which are more *useful* due to higher reliability and greater specificity.

1.4.3 Goal-directed actions

Schemas of arbitrarily high reliability can be thought of as *rules* in that if the context is satisfied, taking the indicated action reliably produces the given result. Therefore, once a number of reliable schemas have been produced, it becomes fairly simple to reach a given goal through planning. Over time, the system becomes able to chain various schemas together to produce a variety of desired results. For any desired result, the mechanism can create a *goal-directed action*, an action which is designed to produce the given result. These new abstract actions give the mechanism an ability to bring about a desired result through a number of intermediate actions, and to treat this process as if it were a single discrete action.

In [Dre89], a goal-directed action is created whenever a particular item or conjunction is highly *accessible* — when, at each clock tick, there is usually some chain of reliable schemas which starts with an activatable schema and ends in a schema with the item or conjunction positively included in the result. [Dre91] creates a goal-directed action whenever a result spinoff has a unique result. The reimplementations use the method from [Dre89] as it reduces the proliferation of many actions early on, but the method in [Dre91] is simpler, less compute-intensive and seems more cognitively realistic.

When a goal-directed action is created, a bare schema is constructed which has the new goal-directed action and an empty context and result. The marginal attribution algorithm will then attempt to build reliable schemas which utilize the goal-directed action and encapsulate knowledge about the goal-directed action. (For more details about goal-directed actions see [Dre89, section 3.4.2].)

[Dre91] and [Dre89] use *composite action* where I have chosen to use the term *goal-directed action*. An informal discussion group concluded that *composite* is a word overloaded with meaning — in particular, it suggests the treatment of a specific series of actions as a single action as in the mathematical operation of composition where one constructs a new function by defining it as the result of the sequential use of two separate named functions [SESA86, p. 134]. It was therefore proposed to use the term *goal-directed action* instead, which is more precise in meaning, as a goal-directed action will activate whichever series of schemas will most likely achieve the desired goal state — it does not activate the same series of schemas each time it is executed.

1.4.4 Synthetic items

There are certain concepts that the primitive items are unable to express, for example, that a particular object is present at a particular location while the glance orientation is such that the object is out of view. The schema mechanism contains a facility for building *synthetic items* — items which, when on, indicate that a particular unreliable schema, if activated, would succeed. Suppose a schema */move glance orientation to vp01/fov02* is very reliable if the left hand object in the microworld is in the correct position (see figure 1-5, if the glance orientation is at *vp01* and the left hand object is in the indicated position, it is in the forward foveal region, and could turn *fov02* on). However, this object spontaneously moves between four different positions and so, on average, is only in the correct position about one-fourth of the time. Notably, this schema, if activated and successful, will continue to be very reliable for some period of time (equal to the duration that the object remains in that position), even though on average it is normally not very reliable. To discover such situations, the schema mechanism keeps a *local consistency* statistic which indicates how often the

schema succeeds when its last activation was successful. If a schema is unreliable but highly locally consistent, the mechanism constructs a synthetic item — an item which, when on, indicates that the schema (its *host schema*), if activated, would succeed. Effectively, such an item, when on, predicts what the result of activating the host schema would be. For a variety of reasons (see [Dre91, section 4.2.3]), synthetic items are fundamentally very different from primitive items and express concepts which are inexpressible through any conventional combination of primitive items.

Primitive items get their state directly from the microworld. On the other hand, the schema mechanism itself must maintain and update the state of all synthetic items. The rules the mechanism uses to determine the state of a given synthetic item are summarized in table 1.3. The use of synthetic items effectively allows the schema mechanism to invent new concepts — concepts which are not expressed well by the microworld or cannot be expressed by conjunctions of boolean values at all.

1.4.5 Control

The CM2 implementation cycles between periods where schemas are chosen for activation on the basis of their value, and periods where the system is emphasizing experimentation with recently created schemas [Dre89, section 3.2.2]. The primary goal of the reimplementations of the schema mechanism is to validate the results found in [Dre91] and [Dre89]. There is no analysis of the ability of the CM2 implementation to find and obtain states of high value in either source, rather, the results presented are the structures (schemas, goal-directed actions and synthetic items) which the system built in a reference run. As goal-seeking behavior is not documented in the results of either source, and this reimplementations is an attempt to verify the results, the reimplementations detailed in this thesis does not need to cycle between periods of goal-seeking and experimentation, and therefore doesn't. The reimplementations also does not make use of any notion of *value* (see [Dre89, pp. 78–83] for a discussion of value). Rather, the reimplementations merely selects one of the currently defined actions at random at each time step. As there are bare schemas for each action, and bare schemas are always activatable, each action is always selectable, and so picking

- **Host schema activated:** If the host schema for a synthetic item is activated, the item is turned on if the schema succeeded. If the schema failed, the item is turned off.
- **Host schema overridden:** If the host schema is context overridden, the synthetic item is turned off. For purposes of updating the synthetic item state, a schema is considered overridden if an item is correlated by its extended context at least 75% in the direction opposite the current state of the item. See section 2.1.4 for more details.
- **Context spinoff:** Context spinoff schemas may be created from the host schema in an attempt to improve the reliability of the host schema. These spinoff schemas have the same action and result as the host schema. Whenever a reliable schema is applicable, its parent schema is checked to see if it is a host schema. The fact that a reliable schema with the same action and result is applicable implies that the host schema would succeed if activated, and therefore the synthetic item is turned on. If a reliable schema is applicable, but overridden, its parent schema is not checked.
- **Result predictions:** If a reliable schema which contains a synthetic item in its result is activated, the mechanism assumes that the schema succeeded, and turns the item on (if positively included) or off (if negatively included).
- **Local consistency:** When the mechanism turns a synthetic item on or off for any of the reasons listed here, the item stays in that state for the length of the expected duration for that transition. (The schema mechanism keeps two statistics for each synthetic item: average duration the item stays on once turned on, and a similar statistic for off.) If that period of time ends without the item being turned on or off by the mechanism, the item is placed in the unknown state.

Host schema evidence has the highest priority when determining the state of a synthetic item, as a synthetic item is defined in terms of success or failure of its host schema. Context and result evidence have the next highest priority — if they disagree, the synthetic item is placed in the unknown state. Local consistency evidence has the lowest priority.

Table 1.3: Rules for determining synthetic item state.


from all actions randomly is perfectly acceptable. This also has the nice side effect of ensuring that all actions are exercised roughly equally.

1.4.6 Summary of schema mechanism

Each of these facilities has an important role. While the marginal attribution technique is a powerful and central part of the system, it can only perform induction from what is already known. Goal-directed actions give the system the ability to abstract the details away from a process which is designed to bring about a desired result, while synthetic items allow the system to invent useful and arbitrarily complex concepts. Together, these two abilities enable the system to discover and define concepts and procedures of its own — contributions to the knowledge base which could not be made by marginal attribution.

1.5 Performance

While the marginal attribution algorithm is fairly compute intensive, especially as the number of schemas increases, it is not intractably inefficient. See [Dre89, section 3.3] for a discussion of the architecture of the CM2 implementation (a parallel machine). This implementation completed its reference run in roughly one day of real time. The reimplementaion, somewhat simplified but running on a DECStation 5000/120 with 16 megabytes of memory and using Lucid LISP 4.0, completes a reference run in slightly more than two days of real time. The schema mechanism and microworld, while complicated, are not so compute intensive as to make them cognitively implausible.



Chapter 2

Implementation

2.1 Differences between this work and Drescher's

2.1.1 Goal-directed actions

(Note: the term *goal-directed action* is used in this work wherever Drescher's works would use the term "composite action". See section 1.4.3.)

While Drescher's work discusses goal-directed actions at some length, they do not seem to be vital for accounting for many of the results he found. In particular, his work does not include statistics for how the system behaves or performs when engaging in goal seeking behavior.

The reimplementaion is focused on evaluating the accuracy of the results presented in [Dre91] and [Dre89]. This primary goal, combined with a desire to finish this work in a timely fashion and some technical problems surrounding storage of goal-directed action controller data in reasonable amounts of memory, led to a decision to leave most of the goal-directed action ideas unimplemented.

In particular, the reimplementaion uses the method in [Dre89] for determining if a goal-directed action should be created. When a new goal-directed action should be created, the reimplementaion merely displays a message to this effect. The reimplementaion does *not* create bare schemas for the new goal-directed action, and therefore none of the schemas created by the reimplementaion pertain to goal-directed

actions. The reimplementaion does not create goal-directed action structures, does not update them, and cannot activate them.

In the reimplementaion, only positive items or conjunctions (with items which are positively or negatively included) are eligible to be the goal of a goal-directed action. Having a negated conjunction as a goal would be fairly useless, as it is equivalent to a disjunction of negated items, which is something the system doesn't work with or understand. A goal which is the negation of a single primitive item may be useful in some rare cases, but generally, goal-directed actions which turn a given item on are more useful. The reimplementaion currently supports only positive single items, but could very easily be modified to support goal-directed actions which have negated items as the goal. In fact, some support for this is already in place, but it seemed fairly unimportant to finish given that the goal-directed action execution code is incomplete.

With the limited implementation of goal-directed actions, some other statistics and structures are no longer necessary. In particular, duration and correlation statistics for schemas are no longer kept — each primitive action has a duration equal to one clock tick and schemas are never activated for their result due to a simplified control mechanism (see section 2.1.3), so neither statistic is needed.

However, many of the structures and code required for finishing the implementation of goal-directed actions are in place in the code, and with an inspired solution to the memory problem noted above, the code could be finished fairly easily.

2.1.2 Value

Drescher's work describes a system whereby various items are given delegated and instrumental *value*, and the schema mechanism, when engaging in goal seeking behavior, tries to reach states which have high value. As there is no analysis of the ability of the CM2 implementation to find and obtain states of high value in either source and no goal seeking behavior in the reimplementaion, the notion of value is not necessary for the reimplementaion and therefore omitted. With no notion of

value, there can be no notion of the *cost* of a schema, and therefore this statistic is not maintained either. (See [Dre89, pp. 78–83].)

2.1.3 Control

Drescher's work indicates that the schema mechanism should cycle between periods of goal seeking behavior and periods of behavior designed to generate more knowledge about the world [Dre89, section 3.2.2]. Goal seeking behavior is brought about primarily through the activation of schemas which have goal-directed actions and therefore encourage the bringing about of a desired result through the explicit activation of a series of schemas. The schema mechanism selects a schema for explicit activation when it is applicable and contains a number of desired results.

As mentioned above, however, the reimplementaion does not support the creation or execution of goal-directed actions. Therefore, at each time step, the reimplementaion must only choose between the ten primitive actions which are provided by the microworld. In addition, the reimplementaion *never* needs to activate a specific schema for its result, due to the lack of goal-seeking behavior and no notion of value, the system can activate any applicable schema at each time step. Therefore, the reimplementaion does not select a particular schema for explicit activation. Rather, at each time step, one of the ten primitive actions is selected and executed. As there are bare schemas for each action, and bare schemas are always eligible for activation, each action is always selectable, and so picking from all actions randomly is perfectly acceptable. This also has the nice side effect of ensuring that all actions are exercised roughly equally. All schemas which share the selected action and have a satisfied context are considered *activated*. Drescher's distinction between explicit and implicit activation is unnecessary, as there is no process of selecting a particular schema for explicit activation in this implementation.

2.1.4 Overriding conditions

[Dre89, p. 106] specifies that a schema which is applicable can not be explicitly activated if an overriding context condition occurs. Due to the simplified control system of the reimplementation, this distinction is unnecessary. The control system implemented does not select a schema for explicit activation and therefore the reimplementation does not look for overriding conditions when picking an action to execute.

(Note for future implementors: in a private electronic mail message, Gary Drescher indicated that the CM2 implementation was designed to override a schema when an item is correlated by its extended context at least 50% in the direction opposite the current state of the item. Example: if the extended context of $/a/x$ indicates that the counter for item p is at least halfway to the value which would force creation of a context spinoff with p positively included, and p is negative at the start of this cycle, then $/a/x$ is suppressed — but it is still considered applicable and if action a is taken by another schema, it will be considered implicitly activated for purposes of updating its statistics. On the other hand $/a/x$ cannot be selected for explicit activation, because of the overriding condition.)

2.1.5 Microworld

There are some minor differences in the microworld from that described in [Dre91, section 6.1]. In particular, the two objects are not in precisely the same position, and an arbitrary decision was made to have both objects rotate clockwise among their home positions (this wasn't specified in the original sources).

Hand motions require that the destination square for the hand is empty — if the hand is currently grasping an object, the destination square for the object must be empty also. Note, however, that when the hand moves left, the hand ends up occupying the former position of the grasped object, and vice versa for movements to the right. Again, this was not specified in the original sources.

[Dre91, table 6.2] labels the haptic and visual proprioceptive items as having their origin at position (1,1) while the other items using a 2D representation (such as the

visual field items) have their origin at (0,0). This is confusing, and in fact [Dre89] uses a (0,0) origin for all 2D items. The microworld items are precisely as described in [Dre91, table 6.2] except that all (1,1) based 2D items have been translated to (0,0) based systems, for example, *hp* and *vp* items now range from (0,0) to (2,2) instead of from (1,1) to (3,3).

In the reimplementaion, if the hand is closed, it is automatically opened after three time units pass. This is different than either [Dre91] or [Dre89], and was changed to show that the precise value of the duration is not important.

2.1.6 Schema mechanism

Gary Drescher answered many questions via private electronic mail and phone conversations. Most of the answers served to illuminate material which was already present in his work — these answers were used in writing the explanations of the microworld and schema mechanism in chapter 1.

One detail which was not adequately explained in either [Dre91] or [Dre89] was precisely when, in the order of events, the system is supposed to randomly move objects and open the hand if it's been closed for more than a given clock tick duration. Gary Drescher explained that part of the job of the mechanism is to differentiate between transitions which are caused by the execution of an action and those which are completely external. Therefore, on each cycle, the system selects an action to perform, executes the action, calls the `clock-tick` function (which randomly moves the objects and opens the hand if necessary, as well as incrementing the clock) and then takes the statistics needed by the schema mechanism. The results of calling the `clock-tick` function are considered to be part of the effect of executing the action on this particular occasion — its effects are attributed to the selected action. (As it turns out, the statistics kept by the schema mechanism are fairly immune to this source of noise in the data.)

2.1.7 Reconciling different sources

[Dre91] and [Dre89] contradict one another occasionally. In each case, the reimplementation is based upon [Dre91]. Some examples from the microworld sections of each source:

- in [Dre91] there are 9 possible hand positions in a 3x3 area, while in [Dre89] there are 25 possible hand positions in a 5x5 area
- in [Dre91] there are 9 possible gaze orientations in a 3x3 area, while in [Dre89] there are 25 possible gaze orientations in a 5x5 area
- in [Dre91] when the hand is closed it remains closed for 2 time units, while in [Dre89] it remains closed for 20 time units (note that the reimplementation actually keeps the hand closed for a total of 3 time units, just to show this is a non-critical value, but this value is intentionally similar to that given in [Dre91])

Again, in all cases, the reimplementation is based on the account in [Dre91] except for the creation of goal-directed actions as noted in section 1.4.3 where the reimplementation uses the accessibility method given in [Dre89].

2.1.8 Piagetian influence

As a reimplementation of major portions of Drescher's work and an investigation into the reliability of his results, this work does not directly reference the psychologist Jean Piaget in any major fashion. Piagetian ideas are clearly the driving influence in Drescher's work and the milestones postulated by Piaget are used as a benchmark for the success of the CM2 implementation. For a discussion of the progression postulated by Piaget and some of his theories see [Dre91, chapter 2]. His ideas are the underpinnings of this work and Drescher's works — without his observations and contributions to child psychology, the results found in these works would be much less interesting.

2.2 Design decisions

The reimplementaion work began with a number of important goals in mind. Speed of the code was considered to be quite important, as multiple test runs were desired to give statistical validity to the results. The readability and usability of the code was also a significant goal. Finally, the system had to execute on a general-purpose UNIX platform. These parameters led to a number of key decisions:

- The microworld was implemented and tested separately, to make the system more modular and easier to understand.
- A set of macros were developed to increase the speed of the math computations through LISP declarations without making the code in the rest of the reimplementaion unreadable (see appendix sections A.3 and A.4).
- As the schema mechanism code is pretty complicated, an effort was made to use data abstractions whenever possible to increase the readability and usability of the code. In each case, an effort was also made to ensure that the use of these data abstractions would not drastically reduce the speed of the program.
- When necessary, tradeoffs were made to save memory at the expense of additional computational overhead. The speed of the program may have been reduced, but from informal benchmarking, it was determined that virtual memory paging with the larger structures was similarly slow. The result is code which can run well on machines with fairly modest amounts of memory.

Given the complexity of the schema mechanism, it was inevitable that the reimplementaion code would be computation and time intensive, but the final version of the code is not intractably so. It also, thanks to the constant development of better UNIX platforms, can be used as the basis for future work which is more computation and memory intensive without switching to a different platform.

It is notable that the final version of the code seems to have achieved a reasonable balance between these design goals.

There are some differences, largely omissions, between the reimplementations and the original sources, as noted above in section 2.1. An incomplete implementation of goal-directed actions, no concept of value and the simplified control mechanism are all different than that proposed by Drescher. However, the marginal attribution algorithm and the notion of synthetic items are completely implemented, and these account for the majority of the results reported in [Dre91] and [Dre89].

2.3 Description of reimplementations code

2.3.1 Microworld

The microworld is implemented as a completely separate unit (in its own LISP package) with an interface consisting of 141 functions which can be called to get each primitive item status, 10 functions which are primitive actions (i.e. may change the item status), a `clock-tick` function (occasionally randomizes object position, automatically ungrasps the hand after 3 time units) plus a function which initializes the world and a function which changes the initial random number generator seed.

The microworld supports the definition and initial placement of world objects. Each world object has 16 `t` or `nil` visual characteristics, 4 `t` or `nil` tactile characteristics, and 4 `t` or `nil` taste characteristics. Each object also has a flag indicating if the object is movable or not — only a movable object can be grasped by the hand. Currently, the only immobile objects are the body and hand.

The state of the microworld is kept internally through a handful of special variables. Nothing other than status messages or the values `t` or `nil` are returned by any routine. A primitive item function which returns `t` indicates to the schema mechanism that the item is currently on, while `nil` indicates the item is off.

The primitive actions are implemented as described in section 1.3. Hand and glance motions work if the resulting position is within the designated area. Hand motions further require that the destination square for the hand is empty — if the hand is currently grasping an object, the destination square for the object must be

empty also. Note, however, that when the hand moves left, the hand occupies the former position of the grasped object, and vice versa for movements to the right, so only one vacant square is required for hand movements in either the left or right direction regardless of if the hand is grasping an object or not. Grasping only succeeds in picking up an object if the hand is not currently closed and the object to the left of the hand is movable — whenever the hand is closed, it is automatically opened after 3 calls to `clock-tick`.

The `clock-tick` function always checks to see if the grip should be undone, and at random intervals averaging 200 clock ticks moves all of the non-grasped movable objects to their next home position. It also increments the clock and returns the new value. Because of this use of random numbers (and the modified control system which selects between each of the ten primitive actions at random at each time step), the random number generator start is saved, and new random number generator seeds can be created and saved. This facilitates the ability to repeat a given test run.

The microworld can be initialized (with the current random seed) by calling the `init-microworld` routine, which zeros the clock, opens the hand, places the hand at body relative position (1,1), glance orientation at body relative position (1,1), places the body in the customary microworld position (3,1), and puts two generic objects into the world as well (the initial state of the world is shown in figure 2-1).

2.3.2 Schema mechanism

The schema mechanism code (see appendix section A.2) is nearly four thousand lines long. It was therefore essential to keep the code as organized as possible. Data abstraction was the major strategy used throughout to combat complexity. Macros were utilized to implement the data abstractions in a fashion which would keep the time cost at execution minimal. (The use of macros forced simpler datatypes to be defined earlier in the code than more complex datatypes to avoid compiler warnings.)

In addition, with an ambitious program with many complicated algorithms, it is vital to be able to debug sections of the program separately and together. A series of constants and macros are defined at the start of the file. By changing a given




Figure 2-1: The initial state of the microworld.

constant from `nil` to `t` and recompiling the program, execution of the program will send the indicated detailed output to the output file. When a constant is `nil`, all code relating to output for that particular section of the program is eliminated entirely via the macros. This setup gives a great amount of flexibility, as data can be observed as desired without affecting the speed of the program when no output or minimal output is desired.

To aid in the discussion of the code, *simple* datatypes are defined as datatypes which are made up of LISP datatypes and at most one other simple datatype. On the other hand, *compound* datatypes are made up of LISP datatypes and multiple simple datatypes. Each of the major structures of the schema mechanism has a corresponding compound datatype. Each of the compound datatypes also has an array which stores all the created instances of that type.

Simple datatypes

Flags are a simple datatype which can be either true or false. Flag-records and flag-arrays are defined to support storage and manipulation of many flags at the same time. In particular, flag-arrays have some specialized functions which support inclusive ORing of two flag-arrays (`flag-array-ior`) and comparing two flag-arrays to see if all the flags which are true in one flag-array are also true in another flag-

- **flag**: Flags can either be true or false.
- **flag-record**: Flag-records contain 24 flags, each of which can be either true or false. They are used only to help define the flag-array datatype.
- **flag-array**: Flag-arrays contain a number of flags equal to the length of the array times 24. They are used to store large amounts of true/false data in a very compact way.
- **state**: States can be on, off or unknown.
- **state-record**: State-records contain 12 states, each of which can be on, off or unknown. They are only used to help define the state-array datatype.
- **state-array**: State-arrays contain a number of states equal to the length of the array times 12. They are used to store large amounts of state data in a very compact way.
- **counter**: Counters are used to keep extended context and extended result correlation data. They have a toggle flag, a value which ranges from 0 to 15 and a positive flag which indicates the sign for the value.
- **counter-record**: Counter-records contain 4 completely separate counters and are used to help define the counter-array datatype.
- **counter-array**: Counter-arrays contain a number of counters equal to the length of the array times 4. They are used to store large numbers of counters in a compact way.
- **rate**: Rates keep track of how often something occurs.
- **weighted-rate**: Weighted-rates keep track of how often something occurs, weighted towards more recent trials.
- **average**: Averages keep a fixnum value which indicates the average of all the values passed to it via average-update since it was first initialized. They are fairly accurate except for rounding error and are used to keep average durations for synthetic items.

Table 2.1: Simple datatypes defined and used by the schema mechanism code.

array (`flag-array-included-p`). These specialized functions help simplify a number of complicated routines. Flags are implemented as 0 (off) or 1 (on) and flag-records are fixnums which hold 24 flags. Flag-arrays are simply arrays of fixnums. This representation allows 10 fixnums to contain 240 flags, a big memory savings over using the LISP symbols `t` or `nil` (where 240 atoms consume as much space as 240 fixnums in Lucid LISP 4.0) and an implementation which is under a lot more control of the programmer as opposed to using bit vectors. In particular, the specialized functions take advantage of fixnum comparison operations which are much faster than the corresponding bit vector operations, as 24 flags can be compared at once rather than one at a time, and array initialization is faster as well.

States can be on, off or unknown. As with flags, state-records and state-arrays are defined to store multiple states and help simplify complex algorithms. Functions for state-arrays include methods to copy the on or off states in a state-array into a flag-array (`state-array-copy-pos/neg-flag`) and generation of a human-readable string for state-arrays which represent conjunctions of items (`state-array-get-print-name`), as well as also supporting inclusive ORing of two state-arrays (`state-array-ior`) and comparing two state-arrays to see if all the states which are on/off in one state-array are also on/off in the other state-array (`state-array-included-p`). State-arrays are used often, especially to represent conjunctions of items and schema contexts. States are implemented as 2 (on), 1 (off) or 0 (unknown). One trick which is used occasionally is a *both* state (represented as 3) which includes both on and off states, is impervious to inclusive ORing when using `state-array-ior` and matches both on and off states when using `state-array-included-p`. Similar to flags, state-records are represented as fixnums which hold 12 states, and state-arrays are arrays of fixnums. A simple representation would have each state as a fixnum, whereas this method uses 1/12th of the memory and again supports quicker algorithms for specialized operations on state arrays.

Rates and weighted-rates are used to keep track of how often something occurs. Weighed-rates are weighted towards more recent occurrences, while standard rates are not. Averages keep the average value of all the values given to them since they

were created. These three datatypes are used occasionally, again largely to simplify sections of code.

The spinoff detection machinery requires a lot of statistics to be kept in the extended context and extended result of each schema. To keep the memory demands reasonable, the counter method detailed in [Dre89, page 109-11] was used. Each counter has a toggle flag, a value which ranges from 0 to 15 and a positive flag which indicates the sign for the value. The toggle flag is used for purposes of alternating between trials which can give positive evidence for a particular correlation and those which can give negative evidence.

The standard way of utilizing counters is to use them to track whenever an event *occurs* and tabulate statistics to see if either a) having a particular item on or off helps the result to obtain (extended context) or b) taking a particular action helps an item to transition from on to off or from off to on (extended result). Counters are explicitly designed to help gather these statistics. The rules for modifying counter values are summarized in table 2.2. An example might help explain these rules. A counter starts at zero, sign positive. On this trial, the event being tabulated occurs. The sign stays the same, the value is incremented to two. On the next trial, the desired transition does not occur, and the value is therefore decremented. As the resulting value cannot be lower than zero, the value is set to zero, and the sign is still positive as zeros are considered positive. On the next trial, the desired transition does not occur again, so the value is decremented again. The value ends up at two but the sign is now negative. Further negative evidence with the desired transition not occurring will continue to increment the value, eventually maximizing the value at 15 with the sign negative. However, a series of positive evidence trials will decrement the value back to zero, switching the sign to positive, and then increment the value, maximizing the value at 15 with the sign positive. Note that this scheme always exerts pressure towards zero due to the fact that the increment (+2) is smaller than the decrement (-3), which protects against maximization of a counter value due to random noise. The toggle is used to either a) alternate between trials with the item on or the item off (extended context) or b) alternate between trials with the schema's action taken

positive	occurred	counter modification
YES	YES	increment
YES	NO	decrement
NO	YES	decrement
NO	NO	increment

- To increment a counter, add two to its value.
- The maximum value is 15.
- To decrement a counter, subtract three from its value.
- If the value for a counter is non-zero and it is decremented, the resulting value can not be lower than zero.
- If the value for a counter is zero and it is decremented, the sign is set negative and the value becomes two (equal to one correlation in the negative direction).
- A value of zero always has the sign set positive.

Table 2.2: Rules for using counters to detect statistically relevant occurrences.

and those with a different action taken (extended result). Alternating between the two types of trials ensures that the counter accurately reflects the significance of the item in question, as it could otherwise be overwhelmed by a series of trials with a particular item on or a series of trials with a given action taken repeatedly.

The simple datatypes are presented here in a distinct order. Flags and states are very simple notions to grasp. Rates, weighted-rates and averages are only slightly more complex. Each of these five datatypes could find uses in other programs. The most complex of all the simple datatypes is the counter datatype, mostly because the extended statistics algorithm which uses counters is difficult to understand and demands particular behavior from the counters.

Compound datatypes

The compound datatypes use a variety of different datatypes, both LISP datatypes and simple datatypes as defined above, to represent structures which the schema

- **schema:** Each schema structure represents a schema which has been discovered by the schema mechanism or is part of the initial repertoire of the mechanism, with all the added data required for various statistics. Schemas are created in an attempt to encapsulate some bit of knowledge about causality in the microworld and as a springboard to further development of more complicated schemas.
- **item:** Each item represents an item which is either primitive and given to the schema mechanism initially, or a synthetic item which is constructed by the schema mechanism itself in response to a schema which is highly locally consistent but unreliable.
- **conj:** Each conj represents a conjunction of items. Conjs are constructed by the schema mechanism when a schema with a conjunctive context becomes highly reliable, and they support the construction of schemas with conjunctive results. No two conjs can exist with the same conjunction of items, the system never creates duplicate conjs.
- **syn-item:** Each syn-item represents a synthetic item, and stores extra information which is not part of the the item datatype. A synthetic item has two associated data structures, an item which stores the standard information for an item, and a syn-item which stores information particular to synthetic items. The state of a synthetic item is not determined by the microworld, rather, it is determined by the schema mechanism itself, and certain specialized statistics are required to support this.
- **action:** Actions are used to connect a human-readable string to each primitive action. (If goal-directed actions were fully implemented, the structure would be similar to that for synthetic items: two data structures defined for each goal-directed action, an action structure which has a string and a function to call to execute the goal-directed action, and a specialized structure which supports everything necessary for the execution of the goal-directed action.)

Table 2.3: Compound datatypes defined and used by the schema mechanism code.

mechanism uses. The most complicated of these is the schema datatype which is used to store blank schemas which the system starts out with as well as all schemas which it creates during execution.

Each schema has a context, an action and a result which define the schema and never change. Schemas are *never* duplicated, each one is unique. A schema claims that if its context is satisfied, taking the indicated action yields a specified result (with a given reliability factor). Each of these three parts are stored within the schema datatype. If the context is empty, the flag context-empty is set true. Otherwise the context for the schema is stored in the context-array, which can represent conjunctions of items as well as single items. To simplify certain parts of the algorithm, if the context is a single item, the context-single flag is set true, and if a conjunction has been created for the context (the schema is very reliable), the context-conj flag is set true and the conjunction number is stashed in the context-item slot for the schema. The result part of the schema is stored somewhat similarly. If the result is empty, the result-empty flag is set true. If the result is a conjunction, the result-conj flag is set true and the result-item slot holds the conjunction number. Otherwise the result-item slot gives the item number for the result, and if the item is negated, the result-negated flag is set true (conjunctions are never negated).

If the action-gd flag for a particular schema is true, the action for the schema is a goal-directed action and the action-item slot contains the index for the goal-directed action. Otherwise, if the action-gd flag is false, the action-item slot contains an index for a standard action. The action slot contains a function which, when executed, performs the action and modifies the state of the microworld. (The action slot is not set for goal-directed actions, due to the incomplete implementation of goal-directed actions, goal-directed actions cannot be selected for execution. Also, this implementation does not select a schema for explicit activation, rather, it chooses between all the available actions at each time step, so this slot is unnecessary. It is provided merely for support for later development of the system.)

Synthetic items can be defined for each schema. For each schema, if the synthetic item flag is true, the reifier slot gives the synthetic item index for the synthetic item

toggle	item	result obtains?	counter modification
ON	ON	YES	event occurred, toggle toggle
OFF	OFF	YES	did not occur, toggle toggle
toggle state = item state		NO	toggle toggle

Table 2.4: Rules for extended context statistics.

created for the schema. First-tick is used to store a clock tick time used in updating the on and off-durations for the reifying synthetic item. If the syn-item flag is false, no synthetic item has been created for this particular schema.

The schema mechanism requires a number of statistics to be kept for each schema. Many of these statistics deal with various situations when a schema is activated (context satisfied and action taken). The reliability of a schema indicates how often the result obtains when the schema is activated (i.e. how often the schema succeeds) and is biased towards more recent trials by using the weighted-rate functions.

If the result of a schema is non-empty, its extended-context looks for items which, when ON or OFF before the action is taken, affect the probability of success. The system keeps track of both positive and negative correlations for each individual item via a series of counters, one per item. When a schema is activated (context satisfied and action taken), the counter for each item in its extended context is updated according to table 2.4. For each item, the system uses the counter toggle to alternate between taking statistics with the item on before the action is taken and with it off. In this scheme, the single counter for each item will go maximally positive if having the item on before the action is taken makes the result occur more reliably, and the counter will go maximally negative if having the item off is relevant. (See table 2.2 for an explanation of precisely how the counter values are updated.) A maximized counter indicates that the schema may be chosen as the parent of a spinoff schema.

(If the schema has a goal-directed action, it also keeps a similar set of correlations in its extended-context-post slot but with respect to the value of items after the goal-directed action is executed. These statistics are used to determine items which need to have their state sustained for the duration of the execution of the goal-directed

action [Dre91, section 4.1.6, pp. 80–81]. As mentioned earlier, goal-directed actions are not currently executable, and therefore updating of the extended-context-post statistics is not yet implemented, but the structure already exists to support future program development.)

On the other hand, if a schema has an empty result, its extended-result slot looks for item transitions which appear to be caused by the activation of the schema. The system keeps separate positive and negative transition correlations for each item as well as positive transition correlations for conjunctions which represent the contexts of reliable schemas. It alternates between trials with and without the activation of each schema, for each trial, extended statistics for the current schema and item are only taken if the state of the counter toggle for that item matches the activation of the schema. The rules for updating the counters for each schema are explained fully in table 2.5. A counter will become maximally positive if its item undergoes a transition more often when the schema is activated than when it is not. (For more detail, see table 2.2 which explains exactly how counters are updated.)

The extended-context and extended-result counters indicate when a given spinoff schema should be created. The extended-context for a given schema will have a maximally positive or maximally negative counter which indicates which item should be positively or negatively included in a context spinoff schema. The extended-result for a given schema will have a maximally positive counter in either the positive transition or the negative transition statistics which indicates which item should be positively or negatively included in a result spinoff schema.

As mentioned before, schemas are never duplicated. A group of three structures, context-children, result-children and result-conj-children, keep the mechanism from accidentally duplicating an existing schema without having to search through all the schemas checking for duplicates whenever a new spinoff schema is proposed. For example, when a context spinoff occurs, the parent schema notes that the item was spun off (positively or negatively) and makes a similar notation in its new child. No schema descended from either schema will ever attempt to spinoff another context

Result transition statistics for an item are only taken when:

- The item was (off for positive transition statistics, on for negative transition statistics) at the end of the last cycle.
- Each schema which was not activated this cycle only updates statistics for items which were not explained by the activation of a reliable schema (i.e. the transition must not have been predicted).
- If the schema was activated, the counter toggle for the item must be true, and if the schema was not activated, the counter toggle for the item must be false. This forces the statistics to alternate between trials when the schema was activated and those when the schema was not activated.

Positive transition statistics:

item	activated	predicted	positive counter modification
OFF \Rightarrow ON	YES	don't care	event occurred, toggle toggle
OFF \Rightarrow ON	NO	NO	did not occur, toggle toggle
OFF \Rightarrow OFF			toggle toggle

Negative transition statistics:

item	activated	predicted	negative counter modification
ON \Rightarrow OFF	YES	don't care	event occurred, toggle toggle
ON \Rightarrow OFF	NO	NO	did not occur, toggle toggle
ON \Rightarrow ON			toggle toggle

Table 2.5: Rules for extended result statistics.

schema with that item added. (The context-children slot also supports deferral of taking extended context statistics to a more specific child schema. See section 1.4.2.)

To save memory, a data slot holds many important status bits in a compressed form. The applicable, overridden and activated flags are used by the toplevel control code to keep track of the status of the various schemas. Marked is a flag used by the update accessibility routines to keep track of which schemas have already been visited by the algorithm. Succeeded-last is a flag which indicates that the schema succeeded the last time it was activated and is used to help keep local consistency data. Lc-consy (short for local-consistency) is a rate used to keep the probability of successful activation given that the previous activation was successful. If lc-consy is high, the lcly-cons (short for locally-consistent) flag is set true, a synthetic item is created for the schema, the reifier slot of the schema is set so as to point to the new synthetic item, and the schema mechanism begins to keep track of the on and off duration for the synthetic item.

For purposes of data abstraction, a series of macros are defined for each of the flags and rates which are compressed into the data slot. These macros make it appear as if each flag and rate were defined as their own slots by following the standard Common LISP naming conventions for structures. The use of these macros also makes the code much more readable.

The item structure is used to store information about each primitive and synthetic item. The print-name string, the syn-item-p flag, and code if primitive or syn-item-index if synthetic define the item and never change. A synthetic item is indicated by having syn-item-p true. A primitive item has its code slot bound to the appropriate microworld function (see `init-item`) which returns the state for the item. On each clock tick, the state as returned by the code function is placed in the current-state slot, with the old value placed in the last-state slot. (For synthetic items, the mechanism itself determines the state of the item at each time step, and likewise the the current state is placed in the current-state slot and the old value moved to the last-state slot.) Generality is the rate of the item being on rather than off and is used when selecting between more than one possible context spinoff schema. When many context spinoffs

are possible for a given schema, the system picks the “most specific” item, which is defined as the one which is *on* less frequently. This is the item with the least generality. Accessibility is the rate of being reachable by a reliable chain of schemas beginning with an applicable schema, highly accessible items have goal-directed actions created with them as the goal. The *gd-created-p* flag indicates if a goal-directed action has in fact been created with this item as a goal.

Conjunctions of items (primitive or synthetic) are stored in *conj* structures (short for conjunction). The items included in a given *conj* are stored in an *item-array* (a *state-array*) with a corresponding pair of *positive/negative-flag-arrays* which have a flag set if that item is included positively/negatively in the *item-array*. An *inclusion-array* indicates which other *conjs* are included by a particular *conj*. Example: the *conj a & b & c* would include the *conjs a & b* and *a & c* in addition to others. The flag arrays and the *inclusion-array* are used to speed up certain algorithms required by the system. Highly accessible conjunctions are eligible to be the goal of a goal-directed action, the data slot for each *conj* keeps appropriate statistics to support this. Finally, the state of each conjunction is computed at each clock tick and placed in the *current-state* slot, while the old value is moved to the *last-state* slot.

Synthetic items are used to designate validity conditions of unreliable schemas which are locally consistent. The unreliable schema which causes creation is called the *host-schema*, while the synthetic item is the schemas “reifier”. Synthetic items are included as items in the *item array* and are treated 100% as if they were primitive items, they can be in the context and result of schemas and conjunctions. However, the state of a synthetic item is determined by the schema mechanism rather than by a call to a *microworld* function. Certain extra data is needed to support the mechanism determining the state directly. This data is stored in a *syn-item* structure. The *host-schema* slot contains the index into the *schema array* for the *host schema*. *Item-index* is the index into the *item array* for the entry for this synthetic item. *Maybe-state* is used by the routines which calculate the state for the synthetic items — it is merely a place to stash an intermediate value before deciding what the current value is. *On-duration* and *off-duration* are the lengths of time the synthetic item tends to stay

on or off once placed in that state respectively. Set-time is the clock tick when the item was last modified, while unknown-time is the clock tick when the item should be automatically set unknown.

The action datatype is simply a way of relating human readable print names with a series of functions which are called to execute a given action. It isn't even a LISP structure, rather, it is a pair of arrays which have the same relative indexing.

The rest of the code

It is notable that even with the complexity of the algorithms required for the schema mechanism that of approximately 3800 lines of code, roughly 1800 lines are used to implement the datatypes referenced above. A solid grip of all of the datatypes is essential before the remainder of the implementation code makes sense. The remaining code is precisely what would be expected based on the description of the algorithms in chapter 1 and the datatypes as described above and should be fairly understandable. Some comments on some of the more unusual and complex functions are in table 2.6.

The main function `run` ties the whole system together. Everything is initialized, and a loop is entered which saves the output from each 50 clock ticks into a different file. For each clock tick, the system first shows the state of the microworld. Applicable schemas are marked, accessibility of items and conjunctions is updated and goal-directed actions are created. (In the current implementation, goal-directed actions aren't fully supported, and so the system merely outputs a message indicating that the given goal-directed action would have been created.) An action is selected at random and executed, and the microworld `clock-tick` function is called. The state of all items and conjunctions is then updated. Schemas which were applicable and shared the executed action are then marked activated. The state of the synthetic items is then updated. Each schema is marked to indicate if its result obtained, and this data is used to update the reliability factor for each schema. Predicted results are noted, and this information is used in updating the extended statistics for each schema. Note that the extended statistics are taken *after* the microworld `clock-tick`

- **schema-update-print-name:** This function takes a schema which has a context, action and result defined as described in the schema datatype code and in this document and sets its print-name slot to a human-readable description of the schema. This function is somewhat sensitive and will *break* or produce unusual results if the flags referenced in the schema datatype code are not set correctly.
- **init-everything:** An unusual thing about this function is the selector argument, which selects between three different microworlds. 0 selects the standard microworld, 1 selects a microworld with only the left hand microworld object and 2 selects a microworld with only the right hand object. This supports more than one type of test run. This function also takes a random state filename as an argument which allows the system to begin from a different random number generator seed for different test runs. The main `run` function and the batch files for the test runs take advantage of these two arguments.
- **update-accessibility:** This function is extremely complicated and requires a lot of convoluted computations. It is also heavily commented. Note that [Dre91] does not use this method to decide when to make a goal-directed action. [Dre91] uses a simpler method which just creates a goal-directed action for each item or conjunction which appears as the result of a schema. See section 1.4.3. The method implemented here finds items which are accessible through a path of .75 reliability forward from any currently applicable schema. The maximum length of a path is 5 schemas and the reliability of a path is the product of the reliabilities of each schema on the path. Also, only positive accessibility is important, as the program is finding items and conjunctions which it wants to make composite actions for, and negative conjunctions don't make sense (they would be equivalent to disjunctions), while negative items aren't very interesting as goals (and also would proliferate much too easily).
- **syn-item-update-state:** The functions which implement the synthetic item state algorithm are fairly complicated, but a good grasp of the datatypes and the material in table 1.3 should serve to illuminate the code.
- **run:** As mentioned above in `init-everything`, this function takes a selector and a random state filename as an argument, which it passes on to the `init-everything` function. It also takes a string prefix which is concatenated with a continuously increasing three digit number to produce a series of output filenames for each fifty clock ticks.

Table 2.6: Notes on some of the functions in `schema.lisp`

function has been called, they are expected to discern between the effects of the action and random effects caused by the clock-tick function. Finally the system may add one new schema, any number of conjunctions and any number of new synthetic items at the end of each clock tick.

Chapter 3

Results

3.1 Test runs

[Dre91] and [Dre89] each have results which are drawn from a single *reference run*. For this work, the reimplemention code was executed twenty times with a different random number seed each time in a desire to get more statistically significant proof of the ability of the schema mechanism to acquire large amounts of knowledge about the microworld. The various test runs also serve to demonstrate that certain categories of knowledge seem to form regardless of the order in which various actions are performed.

Each test run was ran until 10000 clock ticks went by, or the space for schemas (3600 schemas maximum) or conjunctions (200 maximum) was exhausted. Many of the test runs ended before clock tick 10000. However, all made it to at least clock tick 7000 before terminating due to saturation and most made it much farther than that. Two test runs made it to clock tick 10000, eight were terminated when they ran out of conjunction space, and ten ran out of schema space. (See the first few lines in tables 3.3, 3.4 and 3.5 for the last clock tick that data was saved for each test run and the reason for early termination if the test run ended before clock tick 10000.)

The reference runs in [Dre91] and [Dre89] terminate when all of the available memory is exhausted. The CM2 implementations also utilize a number of separate processors in parallel, where each processor designates a schema or goal-directed action. [Dre89, p. 94] indicates that the reference run from this source was limited to

a total of roughly 3600 schemas due to hardware, while [Dre91, p. 105–6] indicates that the reference run from this source was limited to roughly 7400 schemas.

When using a single UNIX workstation, there is the option of using virtual memory and there are no complications involving dividing up structures among physical processors. Given enough time and patience, and a big enough hard disk for memory paging, substantially more than 7400 schemas can be supported by the reimplementation code. The lower number of 3600 schemas was chosen and used due to the desire to perform a series of test runs, especially as the running time for each clock tick in a single processor serial implementation increases linearly with the number of schemas. It was deemed unreasonable to spend more than the forty CPU-days required to execute the twenty test runs.

3.1.1 Output

The output file from each test run lists the schemas, synthetic items and goal-directed actions which were built and the clock tick when they were built. The program also checks through all of the schemas every 50 clock ticks, and outputs the number of each schema which is highly reliable. An abridged output from a reference run is in appendix section B.1. It has the reliability data removed and is only the output from the first 2300 clock ticks. When the reliability data is included and the test run complete, the output from a test run is a formidable amount of information to wade through and understand, as it can run over one megabyte in length. On the other hand, with one line per structure built by the program and a space-saving tradeoff for reliability of schemas (just outputting the schema numbers of those above .9 reliability each 50 clock ticks, as opposed to outputting some continuous reliability measure or outputting the actual reliability numbers for all schemas each 50 clock ticks), this output is basically as minimal as it can possibly be and still summarize most of the important things which happened during a given test run.

3.1.2 Analysis program

An analysis program was written which reads in an output file from a test run and analyzes it, selecting certain schemas, categorizing them and calculating certain statistics for each category. By using this program, all twenty test runs were held up to a consistent amount of scrutiny within a constrained amount of time. Another benefit of using the analysis program was that certain regularities among the various test runs were made obvious, many of which would have probably been obscured if the raw output files had been analyzed by hand. The code for this program is in appendix section A.5, and sample output from the program in in appendix section B.2.

Of course, this approach brings along with it the potential for abuse, as it is possible that the program itself is obscuring schemas which are important by not outputting schemas which do not fit a category and are not highly reliable. As used in this work, the analysis program has been extremely useful and there have been no perceived negative effects. Indeed, the major purpose of the analysis program is to establish a focus on the schemas which belong to various categories or are highly reliable and therefore worthy of further investigation.

All analysis of the output file did not concern structures which contained goal-directed actions or synthetic items, as goal-directed actions weren't fully implemented in the reimplementaion, and the synthetic items mentioned in Drescher's results reference a goal-directed action. (See 3.2.13.)

3.2 Drescher's results and schema categories

In the result sections of [Dre91], Drescher identifies a number of schemas which he feels are notable and display a significant amount of learning on the part of the schema mechanism. The analysis program discussed above in section 3.1.2 identifies schemas which belong to these categories, among other things. The following sections summarize these categories and the results found by the original CM2 implementation. A more detailed treatment of this material can be found in [Dre91, pp. 119–141]. Note that some of the examples below are taken directly from Drescher's

work, and therefore may reference coordinates which only make sense in the CM2 implementation.

3.2.1 Initial schemas

Both the CM2 implementation and the reimplementations start out initially with ten bare schemas, one for each primitive action. These schemas are */handf/*, */handb/*, */handr/*, */handl/*, */eyef/*, */eyeb/*, */eyer/*, */eyel/*, */grasp/* and */ungrasp/*.

3.2.2 Grasping schemas

The first schema built by the CM2 implementation is */grasp/hcl*. This schema asserts that taking the grasp action results in the hand being closed. This schema is reliable, despite an empty context, as the result follows unconditionally from the action. The marginal attribution process builds this schema quickly because the result occurs only when the *grasp* action is taken. A similar schema, */grasp/hgr*, indicates that the grasp action often leads to grasping an object. However, this action is unreliable, and the CM2 implementation eventually builds *tactl/grasp/hgr* which indicates the importance of having an object to the left of the fingers. A schema is classified as a grasping schema by the analysis program if it has either *grasp* or *ungrasp* as its action.

3.2.3 Coarse visual field shifting schemas

A series of schemas built by the CM2 implementation connect the various coarse visual items to one another via the incremental gaze actions. A typical schema is *vf04/eyel/vf14*. The full complement of these schemas indicate that if an object is seen at a particular position, shifting the gaze will result in that object appearing at a specific adjacent position (depending on which direction the glance action was taken in). Note that the two coarse visual items must relate correctly to one another, for example, moving the glance to the left causes a given object to shift to the right. (See figure 3-1.) There are a total of eighty such schemas, twenty for each glance action,

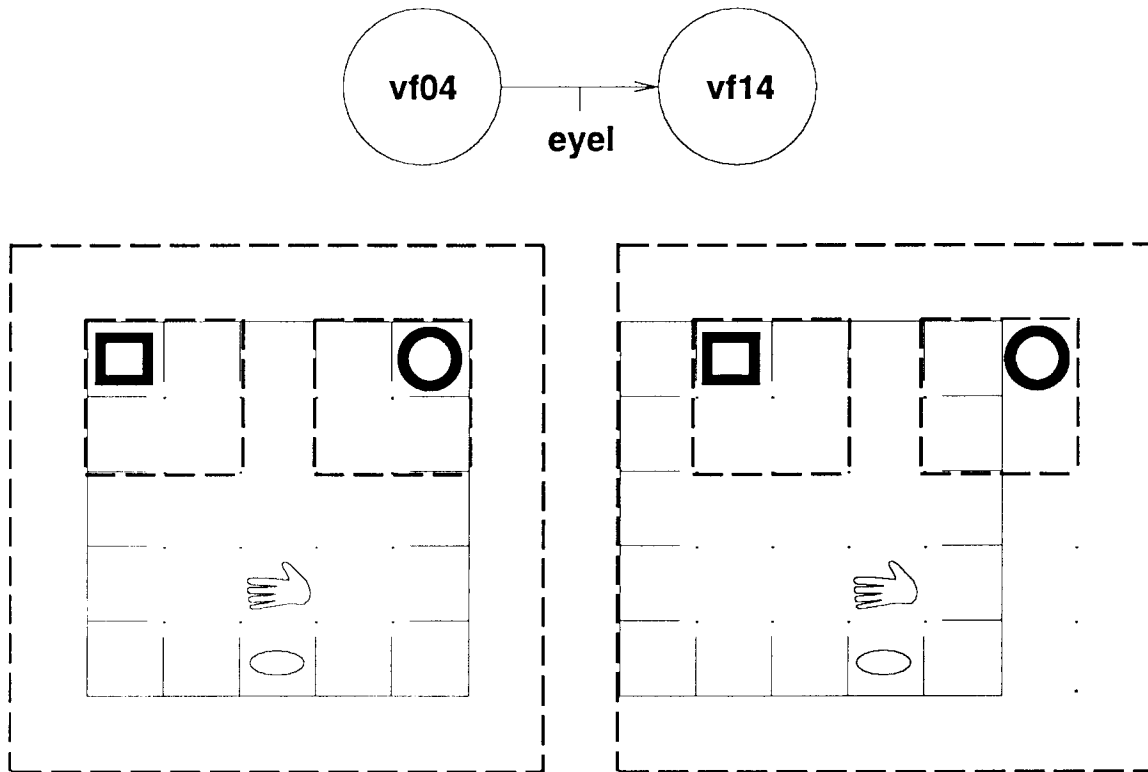


Figure 3-1: Shifting the gaze often causes an object which is within the visual field to appear at a new coarse visual item location. For example, the left hand object in this figure is moved from *vf04* to *vf14* when the glance is shifted to the left.

which together form a network which elaborates the spatial relationships between the various elements of the coarse visual field. The CM2 implementation builds fifty-five of these schemas in its reference run [Dre91, p. 123].

The CM2 implementation also builds twenty-four schemas that concern the special case of moving the image of the body. A typical schema of this type is *vp11/eye/vf30* where the image of the body appears at *vf20* when the glance orientation is *vp11*. (Again, see figure 3-1, the image of the body shifts from *vf20* to *vf30* through the gaze action, but it is just as valid to use *vp11* as an appropriate context.)

The analysis program categorizes any schema with the format *vf??/eye?/vf??* (using the traditional UNIX notation where “?” indicates a position occupied by any single character and “*” is used to represent any number of characters concatenated together, including zero characters) as a coarse visual shift schema provided that the two items and action relate correctly to one another as discussed above, and the two

items are not identical. It categorizes any schema with the format $vp??/eye?/vf??$ as a schema which sees the body via coarse visual items provided that the image of the body when the gaze is oriented at $vp??$ appears at a position which relates correctly to the action and the coarse visual item.

The coarse visual shifting schemas are reliable except when either a) the gaze is already at the extreme orientation for the gaze shifting action or b) the object happens to move as a result of a call to the `clock-tick` function after the action has been taken and before the statistics are updated.

3.2.4 Visual field shift limit schemas

The coarse visual shift schemas referenced above are unreliable when the gaze is already in an extreme orientation. In response, the mechanism builds a few schemas like $-vp01\&vf04/eyel/vf14$ which relate the importance of not having the glance orientation at an extreme position. For example, in figure 3-1, moving the gaze to the left works to shift the object from $vf04$ to $vf14$. However, if the action *eyel* was taken again, none of the coarse visual items would shift at all because the glance is already oriented maximally to the left at $vp01$. The schema $-vp01\&vf04/eyel/vf14$ and others like it are built in response to this inconsistency. There are a very large number of these schemas, the CM2 implementation just begins to learn about them in the reference run [Dre91, p. 123]. The analysis program categorizes a schema as one of this type if it has the form $-vp??\&vf??/eye?/vf??$, the coarse visual items relate correctly to the incremental glance action and are not identical, and the visual position item represents an extreme where taking the glance action would not change the gaze orientation.

3.2.5 Foveal region shift schemas

In a similar fashion to the coarse visual field shifting schemas, the mechanism discovers a number of schemas which deal with the shifting of objects within the detailed foveal regions. $fovb11/eyeb/fovx12$ is a typical schema which is built by the mech-

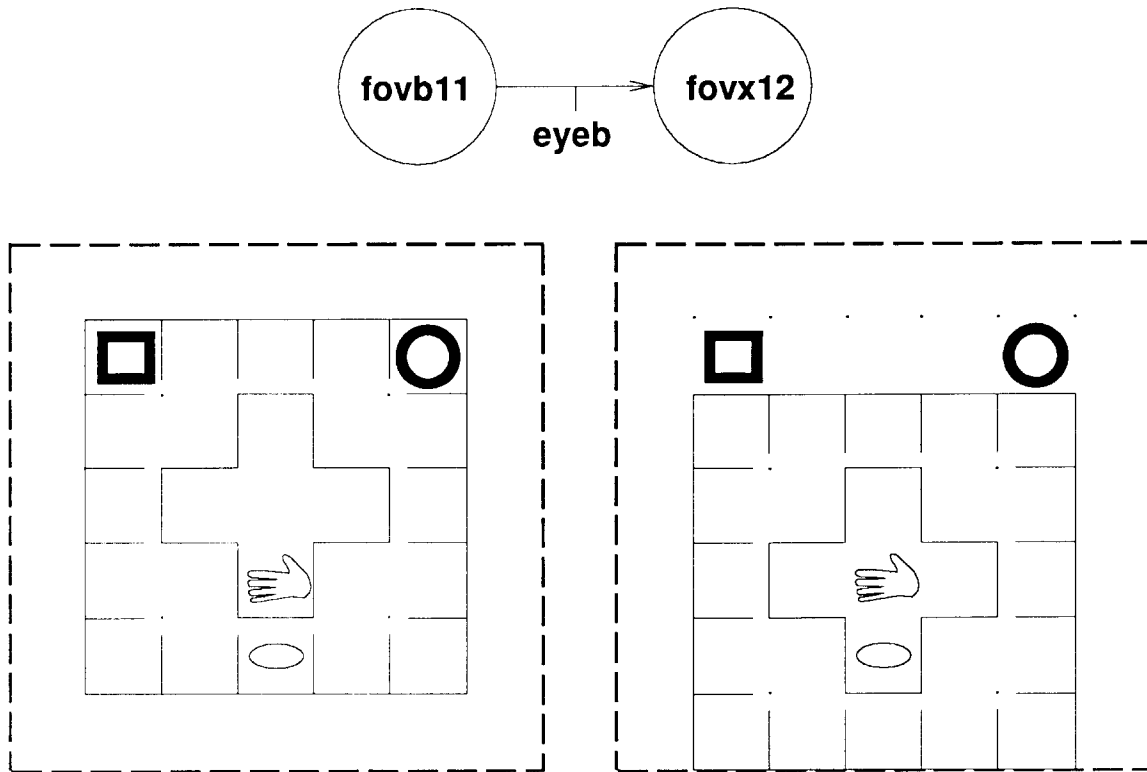


Figure 3-2: The detailed visual appearance of an object often shifts from one foveal region to another when a gaze action is taken. The hand in this figure is shifted from the rear foveal region to the center foveal region when the gaze is shifted to the rear.

anism. Many visual details tend to co-occur between objects, and therefore this schema is just as valid as the more obvious *fovb12/eyeb/fovx12*. Of course, if the mechanism later sees an object which doesn't have both detail 11 and detail 12, the reliability of the first schema may drop, and the second more obvious schema may be built. The analysis program classifies schemas as this type if they have the format *fov??/?/eye?/?/fov???* and that the two different foveal regions relate correctly to one another through the indicated gaze action. (See figure 3-2.)

3.2.6 Detail shift schemas

The mechanism also discovers the relationship between the coarse visual field items and the detailed foveal regions. In particular, the system builds a number of schemas such as *vf21/eyeb/fovx12*. These schemas relate a particular coarse visual field position with an eye movement which brings the visual details of the object to a particular

foveal region. Note that since *fovb11* is never on unless *vf21* is on, the mechanism will never create *fovb11/eyeb/fovx12* once it has created *vf21/eyeb/fovx12* due to the deferral to a more specific schema discussed in section 1.4.2.

A schema is considered to be a detail shift schema by the analysis program if it has the format *vf??/eye?/fov???* or the format *vf??&fov??/eye?/fov???*. The coarse visual item (and the initial foveal region, if the schema is of the second format) must also relate correctly with the given eye movement and final foveal region.

Figure 3-2 illustrates a number of schemas of this type. *vf21/eyeb/fovx12* indicates that shifting the glance to the rear moves the hand into the center foveal region. *vf20/eyeb/fovb03* is similar, but the body is initially only visible in the coarse visual region, as opposed to the previous example where the details of the hand were already initially apparent in the *fovb* items. Many of the schemas of this category indicate how to shift the gaze so as to make the details of a given object apparent. Finally, a schema like *vf21&fovb11/eyeb/fovx12* may be formed by the mechanism if *vf21/eyeb/fovx12* has been created and the schema mechanism sees some objects which do not have visual detail 11 or 12 moving from *vf21* into the center foveal region. These schemas improve on the reliability of their predecessors by constraining which objects have *fovx12* on when moved to the center foveal region.

3.2.7 Visual network schemas

Incremental shifting of the glance changes the glance orientation in a specific and reliable way. *vp11/eyeb/vp10* is a typical schema built by the system in response to this regularity. Figure 3-2 illustrates an example of a successful activation of this schema. These schemas form a lattice which shows the spatial relationship between each of the visual proprioceptive items.

The system also builds schemas such as *vf20/eyeb/vp10* because of the continuous relative position of the body; the body appears at *vf20* when *vp11* is on. The schema *vf20&vp11/eyeb/vp10*, which explicitly relates the two visual proprioceptive items, is built later by the mechanism. Figure 3-2 shows a situation where both of these schemas are also successfully activated.

The CM2 implementation reference run builds 17 of the 24 $vp^{??}/eye^{?}/vp^{??}$ schemas, and seven pairs of $vf^{??}/eye^{?}/vp^{??}$ and corresponding $vf^{??}\&vp^{??}/eye^{?}/vp^{??}$ schemas. Note that the CM2 results do not include schemas where the vp items did not change. The analysis program recognizes schemas of all three formats as belonging to this category provided that the various items relate correctly to one another and to the gaze action taken, but also accepts schemas where the vp items are identical, provided that the gaze is in an extreme position where taking the indicated gaze action does not change the gaze orientation. This increases the number of possible $vp^{??}/eye^{?}/vp^{??}$ schemas to 36, and also increases the number of possible pairs placed in this category. These schemas are included in this category because they encapsulate valuable information about the limits of the possible gaze orientations.

3.2.8 Hand movement network schemas

Similarly, the mechanism builds a series of schemas which relate the various haptic proprioceptive items to one another. $hp10/handl/hp00$ is a typical schema. The system may build schemas like $taste0/handl/hp00$ instead in situations where the hand is immediately in front of the body before taking the hand action. The $taste0$ item always occurs with $hp10$ as $taste0$ is a detail indicating what the hand “tastes like” when in front of the body, and the system makes an arbitrary decision as to which of the two items is included in the context spinoff. Figure 3-3 illustrates both of these schemas.

[Dre91, p. 127] does not clearly indicate if schemas like $hp00/handl/hp00$, where the hand is at an extreme orientation and doesn’t actually move, are to be considered part of this network. If they are included, there are a total of 36 schemas, 9 for each incremental hand action. The analysis program includes these schemas in the hand network category; all schemas of the format $hp^{??}/hand^{?}/hp^{??}$ are included provided that the hand position items relate correctly to one another and the incremental hand action. If the items are identical, they must indicate an extreme position where the given hand action does not change the position of the hand.

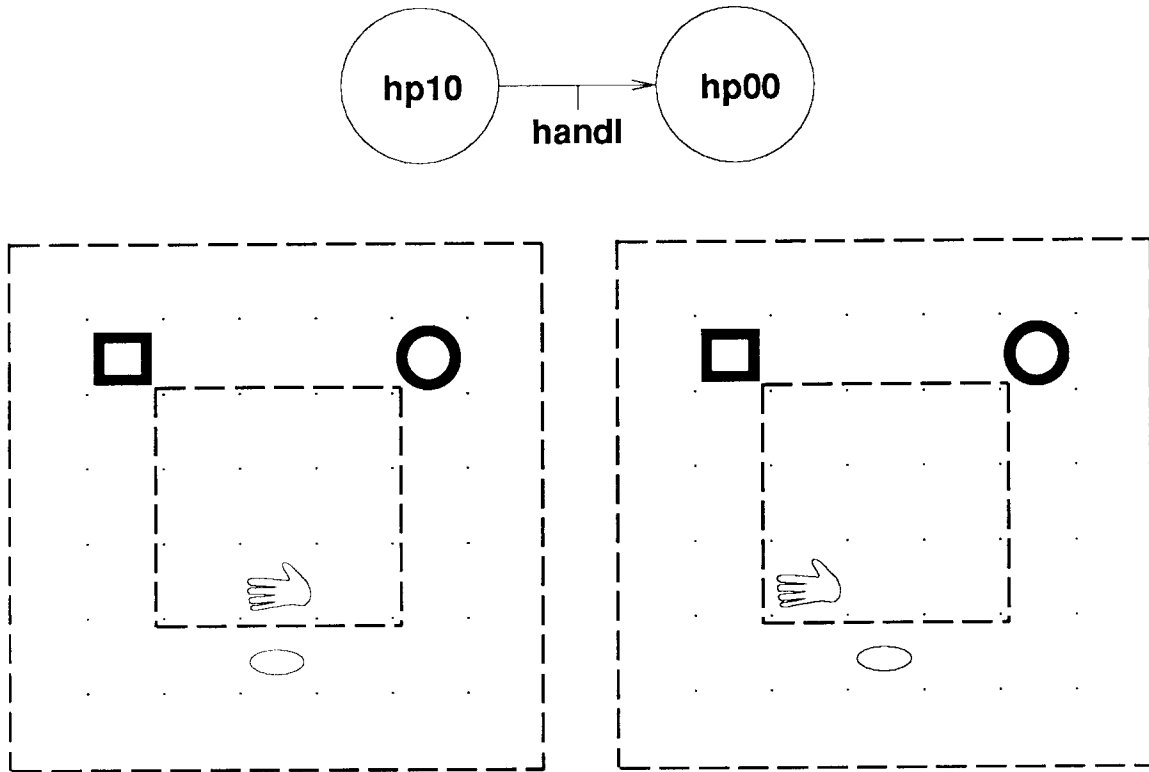


Figure 3-3: Shifting the hand results in the hand moving to a new position.

The CM2 implementation builds all of the schemas in the haptic proprioceptive network in its reference run, with the exception noted earlier where two schemas of the form *taste?/hand?/hp??* are built instead of the corresponding *hp??/hand?/hp??* schema. The analysis program considers these schemas to be part of this classification, as well as *bodyf/hand?/hp??* and *tactb/hand?/hp??*, provided that the hand motion and resulting hp item indicate that the hand was, in fact, immediately in front of the body at *hp10* before the action was taken.

3.2.9 Negative consequence schemas

Taking a glance action or hand motion action changes the current position of the hand or gaze orientation (unless the hand or gaze is already at an extreme position). [Dre91, pp. 126-7] shows */eyel/-vf23*, */eyel/-vp33* and */handb/-hp12* as representative schemas of this type. However, [Dre89, pp. 138-9] shows the schemas *vf23/eyef/-vf23*, *vp12/eyel/-vp12* and *hp34/handb/-hp34*, which is actually a fairly radical difference

in definition. The analysis program accepts a schema as a member of this category if it has a *hand* or *eye* action and has a result which is the same as its context, but negated. In other words, the analysis program is looking for schemas of the type described in [Dre89] for this category.

3.2.10 Hand to body schemas

The schema mechanism also learns that when the hand is near the body, it can be brought immediately in front of the body, which affects the taste and coarse tactile sensations felt by the body and hand. Schemas such as *hp11/handb/taste2*, *hp11/handb/bodyf* and *hp11/handb/tactb* are formed by the CM2 implementation to represent this knowledge. These schemas are similar to those which are in the hand network, just as the hand position in front of the body can be referenced to the hand positions around it through hand motions, the inverse can be done as well. The analysis program accepts any schema with the format *hp??/hand?/(taste? or tactb or bodyf)* as a member of this classification provided that the initial hand position and incremental hand action result in the hand being in front of the body at *hp10*.

3.2.11 Seeing hand movements via coarse visual items

When the hand is moved and is currently visible, the transitions taken by the various coarse visual items lead to the formation of schemas such as *vf21/handf/vf22*. (See figure 3-4.) These schemas are unreliable, as the object appearing at a given coarse visual field position doesn't necessarily have to be the hand. The analysis program considers schemas of the format *vf??/handf/vf??* to be members of this class provided that the two coarse visual items are not identical and they relate to one another and the incremental hand action correctly.

3.2.12 Seeing hand movements via detailed visual items

When the hand starts out in one of the foveal regions, however, the schema mechanism can easily discern between the hand and other objects. A large number of schemas





Figure 3-4: Shifting the hand when it is visible often causes a transition to a new coarse visual item (and possibly to a new foveal region).

such as *fovb22/handf/fovx10* form to relate two foveal regions to one another through a hand movement. Over time, the CM2 implementation adds details which allow the mechanism to discern between the hand and other objects which appear in the various foveal regions to the context. This process eventually culminates in the formation of schemas like *SeeHand@21/handf/SeeHand@22* where *SeeHand@21* is shorthand for a series of visual details which serves to distinguish the hand from all other objects. (Position 21 corresponds to the rear foveal region and the *fovb* visual detail items, while position 22 corresponds to the center foveal region and *fovx* visual detail items.) The analysis program classifies schemas of the format *fov??*/hand*/fov??** in this category provided that the foveal regions referenced by the first item in the context and result relate correctly to one another and the hand action.

3.2.13 Further results of the CM2 implementation

Drescher's CM2 implementation accomplishes a number of other interesting milestones in its reference run, which are summarized in tables 3.1 and 3.2. Again, each example in the table is taken directly from Drescher's work, and therefore may reference coordinates which only make sense in the CM2 implementation. By a notational convention, if an item or conjunction of items appears in the action position for a schema, the schema has a goal-directed action with the goal of turning each of those items on (if positively included) or off (if negatively included).

The reimplemention code is unable to verify any of these further results, as each of them requires executable goal-directed actions in order to form.

3.3 Results confirming Drescher's work

The results of analyzing the twenty test runs for the schemas discovered and categorized by Drescher are shown in tables 3.3, 3.4 and 3.5. This data is further summarized in table 3.6. Discussion of this data and conclusions can be found in chapter 4. Note that these tables only concern the categories detailed in section 3.2.1 through section 3.2.12. In particular, the structures found by Drescher which utilize

- **Touching what is seen:** The mechanism creates a series of schemas such as *fov02/SeeHand@22/tactf* which are reliable and give the mechanism the ability to touch an object which is seen through a series of incremental hand actions.
- **Seeing what is touched:** Similarly, *tactf/SeeHand@32/vf33* forms, which gives the mechanism the ability to shift the gaze so as to bring the hand and the touched object into view.
- **Moving the hand into view:** *vp23/hp23/SeeHand@22* is one of a number of schemas which relate various gaze orientations with the ability to see the details of the hand by shifting it to a given orientation.
- **Shifting the gaze to see the hand:** A network of schemas such as *hp33/vp33/SeeHand@22* relate various hand positions with the ability to see the hand by shifting the gaze to a specific orientation.

Table 3.1: Further results of Drescher’s CM2 implementation, part one.

goal-directed actions and are detailed in tables 3.1 and 3.2 cannot be duplicated by the reimplementaion, as goal-directed actions are not fully supported by the reimplementaion.

3.4 Further results of the reimplementaion

The lack of executable goal-directed actions limits the ability of the reimplementaion to verify all of the results shown in Drescher’s work. However, this does not keep the system from generating many interesting schemas which were not discussed in either [Dre91] or [Dre89].

The analysis program was initially designed to find and categorize only those schemas discussed in section 3.2.1 through section 3.2.12. It also displayed all schemas which were not categorized and had been quite reliable during their life span. By manually examining a number of these lists of schemas, additional categories were added to the analysis program. This process was iterated a few times until categories existed in the analysis program for most of the schemas that were understandable, deemed interesting and occurred in multiple test runs. This process was by no means

- **Persistent positional palpability:** (Try saying that three times fast!) The synthetic item `[/hp23/tactl]`, when on, indicates that there is a palpable object at body relative position 1,3. A series of these synthetic items serve as a map for touchable objects.
- **Persistent positional visibility:** The synthetic item `[/vp21/vf14]` indicates, when on, that there is a visible object at microworld position 3,5. Again, a series of these synthetic items serve as a map for visible objects.
- **Persistent identity details:** The system also generates synthetic items which can identify certain objects, such as `[/hp23/text0]` and `[/hp02/text3]` which denote two different unique objects if *text0* and *text3* do not co-occur between the two objects. Similarly, `[/vp23/fovr10]` and `[/vp23/fovr20]` may form and allow the system to tell the difference between two different objects which can both appear at the same position at different points in time (such as the hand and either of the microworld objects).
- **Inversely indexed persistence:** The schema mechanism, for many synthetic items like `[/hp23/text0]`, also builds `[/text0/hp23]`. To the extent to which *text0* is unique, this serves as an inverse index, where the first structure could be thought as a map indicating what was adjacent to *hp23*, the second structure is a map which says where a specific object is located.
- **Coordinating visible and palpable-object representations:** The creation of various synthetic items culminates in the creation of schemas such as `[/hp23/tactl]/vp23/fovl33` and `[/vp23/fovl03]/hp23/tactl`. The first schema relates that if the world is in the state where an palpable object is at body relative position 1,3, moving the gaze orientation to *vp23* will bring the palpable object into view. The second schema is analogous, but in the other direction, stating that a visible object can be touched.
- **Relational items:** Synthetic items such as `[/vf02/vf10]` relate the position of one object to another, as opposed to the other synthetic items discussed earlier which relate objects to the frame of reference of the body.

Table 3.2: Further results of Drescher’s CM2 implementation, part two.

test run	1	2	3	4	5	6	7
total clock ticks	9850	9350	9050	9350	7000	9250	8500
schema saturated	yes		yes	yes			yes
conj saturated		yes			yes	yes	
0: initial	10	10	10	10	10	10	10
reliable	10	10	10	10	10	10	10
1: grasping	8	7	7	7	2	7	8
reliable	6	6	6	6	2	6	6
2: vf/eye/vf shift	51	40	40	43	30	42	43
reliable	49	36	36	36	30	40	40
3: vf shift limit	16	10	3	7	0	6	5
reliable	6	7	2	5	0	5	4
4: fov/eye/fov shift	282	112	153	161	174	156	318
reliable	164	68	115	78	129	117	220
5: vf/eye/fov	198	131	128	124	98	121	162
reliable	171	114	106	104	75	95	125
6: visual network	63	56	59	46	40	47	43
reliable	49	52	51	40	36	43	34
7: hand network	38	41	36	41	25	36	42
reliable	34	36	34	34	24	31	35
8: $x \Rightarrow -x$	0	0	0	0	0	0	0
reliable	0	0	0	0	0	0	0
9: hand to body	15	16	15	15	12	20	15
reliable	14	15	15	15	12	20	11
10: seeing hand move vf	1	2	1	3	1	2	0
reliable	1	0	1	2	1	1	0
11: seeing hand move fov	103	26	129	36	28	73	39
reliable	79	21	106	20	28	42	30

Table 3.3: Statistics for Drescher's classifications, test runs 1-7.

test run	8	9	10	11	12	13	14
total clock ticks	8950	8350	9500	8250	8900	8550	10000
schema saturated			yes		yes	yes	
conj saturated	yes	yes		yes			
0: initial	10	10	10	10	10	10	10
reliable	10	10	10	10	10	10	10
1: grasping	8	7	7	4	7	7	8
reliable	7	6	5	3	6	6	7
2: vf/eye/vf shift	46	44	44	36	45	48	42
reliable	39	39	41	32	42	43	40
3: vf shift limit	9	6	8	0	7	4	3
reliable	8	5	6	0	6	3	3
4: fov/eye/fov shift	222	310	257	188	242	311	116
reliable	157	216	181	138	161	222	78
5: vf/eye/fov	135	155	208	127	183	182	132
reliable	118	124	173	112	143	160	119
6: visual network	56	42	54	47	62	49	59
reliable	51	32	44	41	51	39	51
7: hand network	36	32	38	33	34	36	40
reliable	28	28	30	29	29	33	38
8: $x \Rightarrow -x$	0	0	0	0	0	0	0
reliable	0	0	0	0	0	0	0
9: hand to body	8	15	20	15	15	16	14
reliable	7	15	18	15	12	14	12
10: seeing hand move vf	0	1	0	3	1	0	1
reliable	0	1	0	3	1	0	0
11: seeing hand move fov	13	32	18	124	44	27	102
reliable	7	26	11	95	38	20	81

Table 3.4: Statistics for Drescher's classifications, test runs 8–14.

test run	15	16	17	18	19	20
total clock ticks	9550	8200	10000	8800	8200	8700
schema saturated	yes				yes	yes
conj saturated		yes		yes		
0: initial	10	10	10	10	10	10
reliable	10	10	10	10	10	10
1: grasping	6	4	7	6	8	8
reliable	5	3	5	5	7	8
2: vf/eye/vf shift	35	40	40	42	36	38
reliable	34	37	38	38	34	38
3: vf shift limit	6	1	6	5	0	7
reliable	5	1	4	4	0	6
4: fov/eye/fov shift	147	79	163	302	223	267
reliable	84	47	118	201	169	211
5: vf/eye/fov	120	121	115	162	140	127
reliable	103	95	104	127	120	116
6: visual network	58	45	59	57	51	54
reliable	54	40	51	48	49	43
7: hand network	42	41	46	37	36	40
reliable	37	34	39	33	31	35
8: $x \Rightarrow -x$	0	0	0	0	0	0
reliable	0	0	0	0	0	0
9: hand to body	15	14	15	15	13	11
reliable	15	11	15	14	12	11
10: seeing hand move vf	0	0	2	4	1	1
reliable	0	0	1	1	1	1
11: seeing hand move fov	24	71	32	86	106	3
reliable	15	52	27	57	69	3

Table 3.5: Statistics for Drescher's classifications, test runs 15–20.

category	min	max	median	mean
0: initial	10	10	10	10.0
reliable	10	10	10	10.0
1: grasping	2	8	7	6.65
reliable	2	8	6	5.55
2: vf/eye/vf shift	30	51	42	41.25
reliable	30	49	38	38.1
3: vf shift limit	0	16	6	5.45
reliable	0	8	4	4.0
4: fov/eye/fov shift	79	318	188	209.15
reliable	47	222	138	143.7
5: vf/eye/fov	98	208	131	143.45
reliable	75	173	116	120.2
6: visual network	40	63	54	52.35
reliable	32	54	44	44.95
7: hand network	25	46	37	37.5
reliable	24	39	33	32.6
8: $x \Rightarrow -x$	0	0	0	0.0
reliable	0	0	0	0.0
9: hand to body	8	20	15	14.7
reliable	7	20	14	13.65
10: seeing hand move vf	0	4	1	1.2
reliable	0	3	1	0.75
11: seeing hand move fov	3	129	36	55.8
reliable	3	106	28	41.35

Table 3.6: Statistic summary for Drescher's classifications.

exhaustive, and it is certain that there are significant schemas built by the system that are not mentioned in this paper.

3.4.1 Shifting the gaze to see the body

These schemas are similar to the ones which form part of the visual network, but instead of indicating what the resulting gaze orientation is likely to be if an object (most likely the body) is present at a given location and the gaze moved, they indicate, from a series of gaze orientations, which gaze movements will bring the body into view, and where it will be seen. A typical schema built by the reimplementation is *vp11/eyel/vf30* (illustrated by figure 3-1). The format recognized by the analysis program is *vp??/eye?/vf??* where the elements of the schema must relate correctly to one another and the object seen by the coarse visual item must be the body for it to be included in this category.

Another related, but separate category, contains schemas which bringing the body into detailed view from a given gaze orientation. A typical schema is *vp11/eyeb/fovb20* (illustrated in figure 3-2). The analysis program uses a similar format and has similar requirements for this category as it does for the prior category.

3.4.2 Using the body as a visual position reference

These schemas do the opposite indexing as the previous category, as when the body is seen in a foveal region and recognized, moving the gaze in a given direction will result in a specific final gaze orientation. A typical schema is *fovb20/eyer/vp20*.

This category also includes schemas with the foveal item negated, as long as the visual position item is also negated, as in the schema *-fovb20/eyef/-vp11*. This schema is actually a very strong statement, as it indicates if *fovb20* is off, moving the gaze forward will *never* result in a gaze orientation of *vp11*. In other words, this schema says that *fovb20* is *always* on whenever the gaze is oriented at *vp10*, as opposed to the less constrained schema *fovb20/eyer/vp20* which merely says that if *fovb20* is on, and the gaze is moved right, the resulting gaze orientation makes *vp20* on. This

second schema, even if highly reliable, admits the possibility of the existence of other schemas with the same action and result and completely different context items. The first schema, if highly reliable, makes a very strong statement about the conditions required to reach the result state.

The analysis program recognizes schemas of either format *fov??/eye?/vp??* or *-fov??/eye?/-vp??* as belonging to this category provided that the items relate correctly to one another and the object seen is, in fact, the body.

3.4.3 Hand network constraint schemas

Distinctly related to the hand network schemas, these schemas indicate which transitions are not possible. A typical schema is *-hp21/handl/-hp11*. Again, these schemas, when reliable, make stronger statements about the microworld than the previously discussed hand network schemas. In particular, this schema says that *hp11* is only reachable through a *handl* action if the hand is at *hp21* before the action is taken. Schemas which belong to this category must satisfy precisely the same requirements as for the hand network schemas, except that both items are negated.

3.4.4 Visual network constraint schemas

A similar set of constraint schemas exist for the visual network. *-vp02/eyeb/-vp01* is a schema which is built by the reimplementation and included by the analysis program in this category. The analysis program looks for schemas which match the format *-vp??/eye?/-vp??* and have the correct relationship between the various components.

3.4.5 Coarse visual shift constraint schemas

-vf??/eye?/-vf?? schemas, if the coarse visual items relate correctly between one another and the gaze action, are placed in this category by the analysis program. These schemas give an alternate perception of the coarse visual field shift network described earlier. *-vf11/eyef/-vf10* is a typical schema built by the reimplementation

which indicates that *vf10* cannot be turned on by the *eyef* action if *vf11* is off before the action is taken.

3.4.6 Detailed visual shift constraint schemas

These schemas give a different description of the foveal region shift network described earlier. If the foveal regions relate to one another through the gaze action correctly, the analysis program puts all schemas of the format *-fov???*/*eye?*/*-fov???* in this category. *-fov01/eyeb/-fov01* is a typical schema, however, the two details need not be identical. (Note: this category appears as category 24 in tables 3.7, 3.8 and 3.9, not in the same order as these categories are presented here.)

3.4.7 Coarse to detailed visual shift constraint schemas

-vf31/eyer/-fovb10 is a typical member of this category. The analysis program looks for schemas which match this format and have a gaze action which moves the object at a given coarse visual coordinate to a foveal region.

3.4.8 Detailed to coarse visual shift constraint schemas

-fovr02/eyef/-vf31 is a typical member of this category. The analysis program looks for schemas which match this format and have a gaze action which moves the object in a given foveal region to a specified coarse visual coordinate.

3.4.9 Detailed to coarse visual shift schemas

Strangely enough, neither [Dre91] nor [Dre89] report any schemas of the format *fov???*/*eye?*/*vf???*, even though the development of these schemas occurs in parallel with the coarse to detailed visual shift schemas which were reported in these sources. Figure 3-1 illustrates a situation where the schema *fovb02/eyel/vf31* has been activated and successful due to the shifting appearance of the hand. The analysis program has the same requirements for membership in this category as for the preceding one, except that the items are not negated.




Figure 3-5: Through a large series of schemas, the reimplemention code learns that there are three different objects which each can appear in specific regions of the microworld.

3.4.10 Seeing objects in different visual regions

The reimplemention builds a number of very reliable schemas which match one of the formats $vp??/eye?/vf??$, $vp??/eye?/fov???$, $-vp??/eye?/-vf??$ or $-vp??/eye?/-fov???$. Examined separately, each schema makes some statement describing where objects appear in the microworld. Taken together, these schemas seem to describe which areas of the microworld each object may appear in. In particular, the left hand and right hand objects are each constrained to be at one of four microworld positions, and the hand is constrained to be in one of nine microworld positions. (The body is, of course, immobile and is only seen in one position.) These three regions where the three movable objects can be seen are effectively separated and defined by the series of schemas defined by the mechanism. (See figure 3-5.) The analysis program accepts schemas which match any of the four formats provided that the components relate to one another correctly.

3.4.11 Hand to body constraint schemas

The hand position in front of the body is only reachable from certain hand positions. $-hp20/handl/-taste1$ is a typical schema which belongs to this category. The analysis

program has the same requirements for these schemas as it does for the hand to body schemas, except that the context and result items must be negated.

3.4.12 Hand movement against object

The reimplementation learns that an object which is felt by the hand cannot be pushed out of the way and indicates this knowledge by building the schemas *tactr/handr/tactr*, *tactl/handl/tactl*, *tactb/handb/tactb* and *tactf/handf/tactf*. The system also builds schemas which have various *text?* items substituted in positions where the *tactl* makes sense, as the former items are detailed versions of the latter coarse item.

These schemas work even when the hand is grasping an object. If the hand is moving left, the object moves with the hand, and the *tactl* or *text?* item is still on after the action is taken. If the hand is moving in any other direction, the object in that direction still blocks the movement of the hand.

The analysis program recognizes each of the four schemas noted above, plus variations where a *text?* item is substituted for a *tactl* item.

3.4.13 Backward hand movement against body

Similarly, the reimplementation learns about situations where the hand is in front of the body and moved backward against it. A series of schemas of the format *(taste? or tactb or bodyf)/handb/(taste? or tactb or bodyf)* are created which represent the various items which are on whenever the hand is in front of the body. Schemas which match this format are placed in this category by the analysis program.

3.4.14 Hand movement from coarse visual to foveal region schemas

[Dre91] and [Dre89] mention discovery of *vf??/hand?/vf??* and *fov??/hand?/fov??* schemas by the CM2 implementation, but they do not mention the construction of *vf??/hand?/fov??* schemas. This category is included for completeness, and contains

those schemas which match this format and have the correct relationships between components.

3.4.15 Support data from test runs

The twenty test runs were also analyzed for the categories detailed in this section. The data gathered is presented in tables 3.7, 3.8 and 3.9. It is summarized in table 3.10. Discussion of this data and conclusions can be found in chapter 4.

test run	1	2	3	4	5	6	7
total clock ticks	9850	9350	9050	9350	7000	9250	8500
12: vp/eye/vf body	17	14	16	18	12	11	12
reliable	16	12	16	15	11	10	11
13: vp/eye/fov body	31	37	27	47	25	24	25
reliable	29	35	20	47	25	20	25
14: fov body/eye/vp	8	6	12	10	5	11	9
reliable	7	4	7	6	5	9	8
15: hp required for hp	10	5	12	5	11	12	6
reliable	10	5	11	5	10	11	6
16: vp required for vp	5	6	6	7	6	7	9
reliable	5	6	6	7	6	7	9
17: vf required for vf	22	26	29	27	25	30	29
reliable	20	22	27	23	21	26	26
18: vf required for fov	71	78	101	84	102	115	91
reliable	61	70	97	68	86	108	77
19: fov required for vf	18	28	35	23	26	12	26
reliable	12	21	28	11	20	9	22
20: fov/eye/vf	77	50	46	53	40	58	61
reliable	54	38	34	33	37	41	43
21: visual regions for objects	53	71	50	45	102	57	51
reliable	43	62	43	31	90	54	41
22: hp required for body	5	4	5	5	5	0	5
reliable	5	4	5	5	5	0	5
23: hand against object	8	7	7	6	1	6	8
reliable	6	6	7	5	1	5	4
24: fov required for fov	64	142	197	136	182	68	75
reliable	50	109	142	95	129	56	60
25: hand against body	20	24	22	15	17	12	22
reliable	12	19	19	10	13	9	12
26: hand movement vf \Leftrightarrow fov	44	25	37	61	8	53	15
reliable	30	20	34	40	3	33	10

Table 3.7: Statistics for new classifications, test runs 1-7.

test run	8	9	10	11	12	13	14
total clock ticks	8950	8350	9500	8250	8900	8550	10000
12: vp/eye/vf body	15	13	15	13	20	11	16
reliable	11	13	11	11	15	9	15
13: vp/eye/fov body	36	25	25	25	47	24	36
reliable	32	21	25	24	45	23	29
14: fov body/eye/vp	6	9	8	9	9	7	8
reliable	4	6	4	5	8	6	6
15: hp required for hp	8	11	8	11	11	7	6
reliable	8	9	8	11	10	7	6
16: vp required for vp	7	6	8	5	5	6	5
reliable	7	6	8	5	5	6	5
17: vf required for vf	23	26	26	24	27	21	29
reliable	22	25	23	24	23	18	23
18: vf required for fov	85	96	65	75	79	71	82
reliable	70	89	56	73	73	54	64
19: fov required for vf	32	16	43	10	23	19	45
reliable	21	12	32	8	17	16	31
20: fov/eye/vf	54	61	43	45	63	63	35
reliable	43	46	30	36	48	45	26
21: visual regions for objects	51	102	66	95	67	70	74
reliable	44	84	52	84	40	49	50
22: hp required for body	9	5	0	5	5	4	5
reliable	9	5	0	5	5	4	5
23: hand against object	8	3	3	2	8	7	7
reliable	5	1	1	2	8	6	7
24: fov required for fov	148	122	135	81	147	104	209
reliable	93	86	78	71	109	70	147
25: hand against body	8	17	15	20	19	10	24
reliable	6	12	8	17	13	3	21
26: hand movement vf \leftrightarrow fov	39	9	29	41	24	16	49
reliable	17	3	22	32	15	13	27


Table 3.8: Statistics for new classifications, test runs 8–14.

test run	15	16	17	18	19	20
total clock ticks	9550	8200	10000	8800	8200	8700
12: vp/eye/vf body	16	9	17	15	14	13
reliable	14	7	15	15	13	11
13: vp/eye/fov body	35	22	31	45	37	4
reliable	34	22	26	45	37	4
14: fov body/eye/vp	14	11	8	8	8	8
reliable	11	10	2	7	7	7
15: hp required for hp	6	6	5	11	6	7
reliable	6	6	5	8	6	7
16: vp required for vp	7	9	6	7	6	7
reliable	7	9	6	7	6	7
17: vf required for vf	34	30	30	30	36	34
reliable	31	28	30	25	34	31
18: vf required for fov	108	124	122	85	98	110
reliable	100	109	99	58	77	100
19: fov required for vf	26	29	38	32	53	40
reliable	20	25	27	22	37	31
20: fov/eye/vf	54	37	38	47	29	49
reliable	41	23	31	33	20	42
21: visual regions for objects	47	62	62	29	41	70
reliable	43	48	46	19	32	59
22: hp required for body	5	5	5	5	5	9
reliable	5	5	5	5	5	9
23: hand against object	7	6	6	3	7	8
reliable	7	6	6	2	7	7
24: fov required for fov	111	223	134	36	207	161
reliable	85	165	95	21	166	131
25: hand against body	24	10	23	18	12	18
reliable	22	9	18	12	11	12
26: hand movement vf \Leftrightarrow fov	17	46	39	63	31	29
reliable	11	42	21	51	25	15

Table 3.9: Statistics for new classifications, test runs 15–20.

category	min	max	median	mean
12: vp/eye/vf body	9	20	14	14.35
reliable	7	16	12	12.55
13: vp/eye/fov body	4	47	27	30.4
reliable	4	47	25	28.4
14: fov body/eye/vp	5	14	8	8.7
reliable	2	11	6	6.45
15: hp required for hp	5	12	7	8.2
reliable	5	11	7	7.75
16: vp required for vp	5	9	6	6.5
reliable	5	9	6	6.5
17: vf required for vf	21	36	27	27.9
reliable	18	34	24	25.1
18: vf required for fov	65	124	85	92.1
reliable	54	109	73	79.45
19: fov required for vf	10	53	26	28.7
reliable	8	37	21	21.1
20: fov/eye/vf	29	77	49	50.15
reliable	20	54	37	37.2
21: visual regions for objects	29	102	62	63.25
reliable	19	90	46	50.7
22: hp required for body	0	9	5	4.8
reliable	0	9	5	4.8
23: hand against object	1	8	7	5.9
reliable	1	8	6	4.95
24: fov required for fov	36	223	135	134.1
reliable	21	166	93	97.9
25: hand against body	8	24	18	17.5
reliable	3	22	12	12.9
26: hand movement vf \Leftrightarrow fov	8	63	31	33.75
reliable	3	51	21	23.2

Table 3.10: Statistic summary for new classifications.



Chapter 4

Analysis, Discussion and Conclusions

4.1 Test run evidence confirms Drescher's results

Drescher, in [Dre91], discusses results which are found by the CM2 implementation which this paper splits into 22 different categories. Of these categories, 10 are unable to be built by the reimplementation code, due to the fact that the reimplementation does not fully implement goal-directed actions. The other 12 categories formed the basis of the analysis program.

Table 3.6 summarizes the results of analyzing all twenty test runs for these 12 schema categories. As a whole, the test runs support the results found by Drescher in 11 of the 12 categories.

All twenty test runs built */grasp/hcl* and further, each built */ungrasp/-hcl*. 19 build */grasp/hgr*, and 18 of these build *tactl/grasp/hgr*. Each of these schemas is mentioned by Drescher, and this is excellent verification of these results.

Beyond these results, 16 test runs have one or more schemas of the form *text?/grasp/hgr* which have the ability to discern between various objects, as the body is immobile and cannot be grasped successfully, and *tactl/grasp/hgr* is therefore not very reliable. 17 test runs build */ungrasp/-hgr*, which is closely related to */grasp/hgr* and */ungrasp/-hcl*.

There was another interesting result, too, one which is not very intuitive. *-hcl/grasp/hcl* is built in 9 test runs. This schema indicates the prerequisite that the hand cannot be already closed when the grasp action is taken for it to become closed reliably, as very occasionally, if the hand is closed and the grasp action taken, the call to clock tick forces the hand open as the number of consecutive clock ticks with the hand closed is three. This indicates how good the marginal attribution algorithm is, actually, as this schema indicates understanding the necessity of opening the hand first if one wants to close the hand and keep it closed for a while reliably.

All twenty test runs build at least 30 reliable coarse visual item shifting schemas. The coarse visual item shift limit schemas, however, are quite a bit more rare. In three of the twenty test runs, not a single shift limit schema was formed. This is probably due to the fact that, in order to form, the system must both repeatedly try to shift the gaze from a particular extreme position in a particular direction, and the object which is seen via the coarse visual item must be in the same position. In detail, to have *vf11/eyer/vf01* spinoff *-vp20&vf11/eyer/vf01*, the system must believe that creating a child schema with *-vp20* added will make a more reliable version of this schema. The marginal attribution algorithm, to come to this conclusion, must see a series of trials which alternate between having *vp20* on and off with the parent schema activated (i.e. *vf11* on and *eyer* action taken). The series of trials must be at least fifteen in length with at least eight *vp20* off trials where the result obtains and at least seven *vp20* on trials where the result does not obtain in order to have the system decide it is a significant context item in the negative direction. These are a pretty serious group of requirements, which is most likely why not very many of these schemas are built in each test run.

While the CM2 reference run builds a total of 79 coarse visual field schemas (including those that reference a visual position to the position of the body), the test runs, on average, only built 41 such schemas. This discrepancy may be due to a misunderstanding. Neither [Dre91] or [Dre89] specify if schemas like *vf00/eyer/vf00* should be included in this category. The analysis program requires that the two items be different and related to one another correctly via the gaze action in order to be

placed in this category, which eliminates schemas such as the one above. Another possibility is that the alternation between random and goal-directed behavior in the toplevel control loop for the CM2 implementation might serve to bias the system towards discovering more about coarse visual items, as they are some of the first goal-directed actions built, there would be a natural bias towards these types of actions.

Massive numbers of foveal shift schemas are built in the reimplemention test runs. A number of these schemas are unreliable, as the system does not know which details co-occur between objects, and has to discover this. However, most of the test runs also have a number of schemas such as *fovf01&fovf02/eyef/fovx32* where the system is clearly trying to improve the reliability of the result by defining precisely which object is meant by the context. As the body can only appear in the rear foveal region (*fovb* items), it cannot be referred to by this schema, and in fact, as it cannot be shifted into any other foveal region, it can never be the subject of a schema of this type. Interestingly enough, of the three other objects in the world, only two objects satisfy the context when in the front foveal region, and these two objects both have detail 32 on, so this schema is reliable provided that the gaze isn't oriented at an extreme position. This is an improvement on its parent schema, *fovf01/eyef/fovx32*, whose context is satisfied by every object in the microworld but whose result only follows reliably for two objects. This process of iterative additions to the context eventually culminates in schemas with contexts which match only those objects which have the particular visual detail referenced in their results. *fivr00&fivr02/eyer/fovx03* is a good example, as its context can only be satisfied by the hand (as the body cannot appear in the right foveal region) and the hand has visual detail 03. Context items may also be negated, for example, *fovl01&-fovl20/eyel/fovx11* has a satisfied context when either the hand or the right hand object is in the left foveal region. Again, both objects have visual detail 11 on, and so further identification of each object is unnecessary. As noted earlier, however, a good number of intermediate schemas which are unreliable due to incorrect detail relationships must be built in order to build reliable schemas with more specific contexts. All visual shift schemas may also

be unreliable due to the fact that they do not take into account the limits of the gaze orientation, as mentioned earlier.

Every test run also has a number of schemas which concern shifting an object from a coarse visual item position into a foveal region. While most of these schemas in each test run are of the format *vf??/eye?/fov???*, at least one in each test run includes a detailed foveal item added to the context, such as *vf12&fovl00/eyel/fovr12*. The context in this schema, to improve reliability, discerns between the two objects which can appear at *vf12*, the hand and the left object, where the hand has both details 00 and 12 and the left hand object has neither. Through this series of schemas the system is again relating precisely which objects have the various visual detail items on. However, most of these schemas deal with the more common situation where the details of the object are not apparent until the object is shifted into a foveal region. In these situations, the schemas which are most reliable are those which have a visual detail result which is shared among all of the objects which can appear in that foveal region.

The numbers reported for the visual network schemas in table 3.6 don't quite match with those reported in [Dre91], where 17 visual network schemas and 7 pairs of schemas which relate a coarse visual item seeing the body (and visual position, for the other schema in the pair) to a given visual position were described with non-identical visual position items. As mentioned earlier, however, the analysis program accepts schemas where the two visual position items are identical as long as they are at an extreme where the gaze action will not change the gaze orientation, which definitely increases the number of schemas belonging to this category.

To get a better idea of what the numbers in the table for this category represent, the first test run was examined directly. The analysis program found 63 visual network schemas in the first test run. 19 schemas were traditional visual network schemas as reported in [Dre91] such as *vp12/eyer/vp22*, where the two visual position items are different from one another. There were a number of schema pairs as discussed in [Dre91], however, many of these pairs covered material which was already adequately covered by other schemas. Only one schema pair was found which effectively added

a segment of the network which was missing. A total of another 5 schema pairs were found, but they did not add anything new to the network. 12 schemas were the complete set of limit schemas such as *vp00/eyel/vp00*. 7 schemas dealt with seeing the body when at an extreme gaze orientation, such as *vf31/eyel/vp00* where the glance orientation doesn't change. Another 5 schemas combined these two ideas together, adding coarse visual information to limit schemas, such as *vp21&vf10/eyer/vp21*, where the body appears at *vf10* when the gaze is oriented at *vp21*. This brief analysis indicates that the system did, in fact, succeed in acquiring a substantial body of knowledge about the visual position network. A quick analysis of a few of the other test runs again yielded data which supported an estimate of roughly 80% of the 24 possible visual position schemas are built in each test run. Similarly, roughly 80% of the hand position network is constructed in each of the test runs.

The system was very good at building schemas which relate information about how to move the hand to the body. In most cases, the system built over 75% of the 20 possible schemas of the form *hp??/hand?/(taste? or tactb or bodyf)*. (20 schemas are possible as the hand has three taste details and can be moved in four different directions.) In two of the twenty test runs, the entire group of schemas is built.

The test runs rarely built schemas which related the situation where the movement of the hand was seen via the coarse visual items. [Dre91] makes it clear that these schemas were uncommon in its reference run. This, coupled with the fact that the system was able to see the hand move in other modes, leads to the conclusion that the lack of *vf??/hand?/vf??* schemas in the test runs is not notably unusual. In particular, the schema mechanism was usually successful in finding a number of schemas which related the hand shifting between various foveal regions. The formation of these *fov??/hand?/fov??* schemas seemed to be very dependent on the particular test run, however. These schemas will not form if the hand is rarely moved around in the foveal regions. A more likely reason for a lack of these schemas is that the hand, when in a foveal region, was often in an extreme position and the hand action didn't result in a new hand position. When the system was successful in finding these schemas in a given test run, it even built schemas such as *fovx01&fovx03/handl/fovl32*

where the context accurately picks the hand out of all objects which can appear in the center foveal region.

One category was basically not confirmed by the reimplementations, that of negated consequence schemas. In this particular case, [Dre91] and [Dre89] disagreed in what was expected to form a negated consequence schema, and the more constrained format of $x/(eye? \text{ or } hand?)/-x$ shown in [Dre89] was used. Not a single schema of this format was found in any of the test runs. However, examination by hand showed that many schemas of the format given in [Dre91] were formed, in fact, for many of the schemas in the other categories, $/(eye? \text{ or } hand?)/-x$ schemas are required precursors to the formation of the appropriate child schemas.

4.2 Analysis of all results

The division between Drescher's results and those which were discussed in section 3.4 is an artificial one. With the massive amounts of data, and Drescher's interest in spending most of his time explaining and discovering the most complex structures the CM2 implementation built in its reference run, it is likely that the CM2 implementation did build some of each of the schemas first discussed in this work, but that their significance wasn't commented upon. However, these schemas give a much stronger body of proof of the ability of the marginal attribution algorithm to find and codify regularities in the microworld.

With the more detailed results given above, it's clear that the system has a good grasp of how the hand actions work. (Apologies for the awful pun... I just couldn't restrain myself.) In particular, the $-hcl/grasp/hcl$ schema, built in nine of the twenty test runs, shows an unexpectedly pure insight into the mechanics of the microworld. The microworld, if the hand is closed, keeps a count and only allows the hand to remain closed for three time units. The extended context of $/grasp/hcl$ keeps track of each item to see if it helps predict the successful activation of the schema. In particular, it is determined that hcl has the peculiar effect of needing to be off in order for the result of having hcl on occur more reliably. This is precisely because

occasionally, when *hcl* is on and the *grasp* action taken, taking the action doesn't change the situation at all, and the call to clock tick which occurs after each action is taken results in the hand being opened, as it has been closed for three clock ticks. The system learns that the hand ends up closed most reliably when it is open immediately before the grasp action is taken. To amplify this, if the system had a goal of ending up with the hand closed, it would have the knowledge necessary to recommend opening the hand first, and then closing it, which reliably gives the desired result.

By any reasonable measure, the system clearly understands the relationships between its various visual items and the glance actions. In every test run, it generates hundreds of schemas which relate coarse visual items to one another, detailed visual items to one another, coarse visual items to detailed visual items, and vice versa. It also learns how to relate sensed visual positions (through the visual proprioceptive items or from seeing the body at a given coarse visual item position or detailed foveal region) to other visual positions through glance actions. Each of these relationships also explore the abilities and limits of the visual system, through indications of how to view the visual details of an item by shifting the gaze since the foveal area is limited, as well as through schemas which indicate the limits of the gaze orientation.

These relationships are further strengthened by groups of *constraint* schemas. These schemas, first discussed in this work, help to define the various networks discussed in the preceding paragraph by indicating the items required for a given transition to take place. Each test run found a number of constraint schemas corresponding to each of the types mentioned above — $-vf??/eye?/-vf??$, $-fov??/eye?/-fov??$, $-vf??/eye?/-fov??$, $-fov??/eye?/-vf??$ and $-vp??/eye?/-vp??$. These schemas indicate, in a very strong way, which transitions involving visual items are *impossible* in the microworld, given certain situations. $-vf11/eyef/-vf10$ stands in contrast with its relative $vf11/eyef/vf10$, as the first schema, if reliable, indicates that none of the other coarse visual items have anything to do with the state of *vf10* when the *eyef* action is taken; if *vf11* is off, then *vf10* always ends up off. Essentially, when this schema is reliable, it makes a statement about all 25 of the visual items at once, as 24 of the items are apparently unrelated to achieving the result, and only one item is

related to the result. The second schema, when reliable, still admits the possibility of other schemas with different contexts which do not include *vf10* and are also reliable. These constraint schemas are a major development by the schema mechanism; when taken along with all other related schemas, they indicate an understanding which seems to go beyond any sort of simple cause and effect nature.

Similarly, but more simply, the system has a substantial body of knowledge about the relationship between hand position and the hand actions, as well as between hand positions and the position of the body. The system learns both *hp??/hand?/hp??* and constraining *-hp??/hand?/-hp??* schemas, effectively representing a network which relates the various hand positions to one another through the hand actions. Gaps in this network are filled in by relationships between hand positions and tactile feedback from having the hand in front of the body. The system, by exploiting these relationships, is able to move the hand to the body if it is in an adjacent position. It is also able to predict which position the hand will end up in after it is moved when the system can feel the hand in front of the body through the tactile items. It is also able to understand the constraints involved, as the tactile feedback unique to having the hand in front of the body can only be achieved by moving the hand into the correct position; no other situation can satisfy these constraints.

The system also gains the ability to understand that pushing its hand against an object which it can feel is futile; the object does not move. In the case of pushing against the body, the system also learns that the hand does not change position and remains in front of the body, as well as learning precisely which items are on when the hand is in front of the body.

The abilities of the system aren't just relating the senses and actions having to do with a particular faculty to one another. It is able to, and in fact does, relate the visual and hand systems to one another. *vf??/hand?/vf??*, *fov??/hand?/fov??*, *vf??/hand?/fov??* and *fov??/hand?/vf??* schemas all form to indicate the visible result of a hand motion.

The system also learns a lot about objects in the microworld. It can predict where the image of the body will appear in response to a given gaze movement from schemas

such as $vp??/eye?/vf??$ and $vp??/eye?/fov???$. It also can use the visual perception of the body as a way of determining the current gaze orientation. $fov??/eye?/vp??$ schemas can give this reliably when the foveal detail item is one which is included only by the body. In addition, the system effectively creates a map of the microworld with an understanding of the regions in which each object can appear. This map is created over time, and is represented as a series of $vp??/eye?/vf??$, $vp??/eye?/fov???$ and their corresponding constraint schemas which collectively indicate which objects (and with which detail items) appear at each region of the microworld. The body location schemas mentioned earlier are similar and can be thought of as a special case where the object is restricted to a one unit area and therefore does not move.

In summary, the system displays an incredible amount of understanding about the world, creating a large body of knowledge merely by observing what happens when actions are taken. This is especially notable when it is remembered that the system starts out with virtually *no* information about the microworld whatsoever, just a set of blank schemas, one for each action. It does not know which items relate to one another or to a particular action; the names by which each of these is referred to in this paper is merely a series of names which are given to the items and actions by convention and are only human readable. The system itself discerns between items strictly by item number and between actions by action number. Given this impoverished start, its accomplishments are even more notable.

4.3 Performance

Direct examination of the timing data generated when executing the test runs indicates that the length of time required to process each clock tick increases linearly with the number of schemas in the system. This is most likely largely due to the serial implementation of the system, as each schema must be updated in sequence throughout the process. A parallel implementation may be able to achieve substantially better performance if carefully coded. The current implementation finishes a test run to 10000 clock ticks in roughly 2 1/2 days of CPU time on a DECStation

5000/120 with 16 megabytes of memory. However, as has been noted earlier, faster UNIX workstations which run Lucid LISP already exist and will continue to be developed. Indeed, if the current code was ran on the fastest available UNIX machine, it would probably finish in roughly 12 hours, as an estimate. This is one big advantage enjoyed by this implementation — as better hardware and Common LISP software becomes available, this system can take advantage of it without requiring anything but minor changes.

In addition, it is almost certain that a clever algorithm designer could find ways of improving the speed of certain crucial portions of code by combining various steps, building new data structures to speed up certain portions of the algorithm, or perhaps even coming up with a completely different but functionally equivalent algorithm.


4.4 Future directions

This work, while begun as an attempt to verify Drescher's results, is now incredibly open-ended. The possibilities for future work based upon this one are endless, perhaps even more so as the implementation uses Common LISP and future researchers are strongly encouraged to use and modify the code which is included in appendix A for this purpose.

These suggestions are by no means exhaustive, and range approximately from the simple to the complex.

- Run the code on a faster UNIX workstation with more memory, and see if the system learns more in a longer test run.
- Find more schema categories and add them to the analysis program, see which of them show up in most or all of the test runs. See if there is a strong correspondence between what schemas are built in a given category and which are built in another category; it is entirely possible that there are inter-category dependencies, and this hasn't been examined at all.

- Try running the code with fewer primitive items in the microworld, and see how this affects the types and numbers of the schemas built. Do the same with a richer microworld with more primitive items. Try decreasing or increasing the number of possible gaze orientations or hand positions.
- Improve the analysis program so the system can compare test runs to see exactly which schemas were built in every test run. Use this data to further refine the categories and the overall impression of what knowledge the system learns about the microworld. Also use this data to develop a full picture of when the system develops each category.
- Analyze the microworld rigorously, deciding which regularities should be able to be learned and how complex learning each of the regularities should be (i.e. how often is there a counterexample, and how often does the given situation occur?) Give a theory for which categories should exist, and in which order they should be created. See if the mechanism can build all of these categories, and if other categories are discovered, use this to redo the analysis of the microworld.
- Figure out a good way of storing the controller data required for execution of the goal-directed actions. Complete the implementation of goal-directed actions, remove the accessibility tests and instead use the method given in [Dre91] to decide when a new goal-directed action should be built.
- With goal-directed actions complete, make them executable, adding appropriate controls to make sure the system does not favor one set of actions over another. See if the rest of Drescher's results can be verified. If not, why not?
- Analyze the algorithms used in the system and redesign them for greater efficiency.



Bibliography

- [Dre86] Gary Drescher. Genetic ai – translating piaget into lisp. Technical Report A.I. Memo No. 890, Massachusetts Institute of Technology Artificial Intelligence Laboratory, February 1986.
- [Dre89] Gary Drescher. *Made-Up Minds: A Constructivist Approach to Artificial Intelligence*. PhD thesis, Massachusetts Institute of Technology Artificial Intelligence Laboratory, September 1989.
- [Dre91] Gary Drescher. *Made-Up Minds: A Constructivist Approach to Artificial Intelligence*. MIT Press, Cambridge MA, 1991.
- [Mae92] Pattie Maes. Slides from 4.996: Modeling and building autonomous agents. (lecture given on the 19th), March 1992.
- [SESA86] Douglas Smith, Maurice Eggen, and Richard St. Andre. *A Transition to Advanced Mathematics*. Brooks/Cole Publishing, Monterey CA, 2nd edition, 1986.