

# **An Approach to Artificial Intelligence**

Alan H. Bond, Ph.D.,  
Artificial Intelligence Laboratory, Department of Computer Science,  
Queen Mary College, University of London,  
Mile End Road, London E1, United Kingdom

February 1979 <sup>1</sup>.

## **Abstract**

An approach to artificial intelligence is presented.

Starting from basic research criteria of comparability of work, operational definitions and mathematical theorisability, the paper develops the four main components of the approach:

- (i) Work should investigate complete behaviour patterns using real robots and environments
- (ii) Performance considerations particularly real time constraints should be integral to the concept of computation
- (iii) It is essential to consider the development of the system by autonomous learning and by teaching.
- (iv) Parallel computer architecture should underlie the system. We are particularly interested in array processors.

---

<sup>1</sup>Written in 1978, originally manually typed on A4, then put in troff, then Scribe, finally retypeset in latex in 2001!

# 1 Introduction

In trying to tackle the basic research questions of artificial intelligence for myself, I have, like many before me I suspect, become much exhausted and confused by the body of existing work on the subject.

The history of mainstream A.I. is well known (see (Nilsson 74) for example for a treatment). It consists of a relatively small body of work (about 5000 papers or so), by a few hundred professional A.I. workers, aided by students. There is a small number of basic ideas and a small number of research centres.

The MIT approach includes actors, pattern directed systems, frames, they've been very interested in the blocks world. Carnegie-Mellon have developed production systems and been more interested in psychological modelling. Stanford and SRI have been more eclectic but their work includes associative languages, goal directed vision, mechanised assembly and of course the SHAKEY and STRIPS robot work. Thus the Stanford work with exception of the SHAKEY Project does not involve a single clear approach and method, unlike the MIT and Carnegie-Mellon approaches. This is not to decry their excellent work. Outside the U.S.A. there are a few separate approaches. (Incidentally "schools of thought" seems inappropriate).

In the U.K. for example, there are a number of small A.I. centres most of which try to follow American approaches. This is a very dangerous pastime since one is at a distance from the unwritten experimental basis, one has to take on trust the findings of the American group because one usually cannot verify them for oneself.

There are small groups with their own indigenously original approaches, notably Alexander's group at Brunel University who work on networks of RAM's and Kowalski's group at Imperial College, London who work on predicate logic as a complete and effi-

cient approach to artificial intelligence, (Aleksander 77)(Kowalski 79).

In the U.S. there are a number of groups with their own methods and approaches, for example Arbib's group at Amherst (Arbib 72) (Metzler 77), the Logic of Computers group at Michigan, and Uhr's group at Wisconsin (Uhr 74). In fact, if we look around we find a lot of interesting work away from the mainstream, in the world of the underfunded amateur. This work includes unfashionable subjects such as brain modelling, cellular machines, learning systems etc.

I don't think we should separate artificial intelligence from the rest of computer science. A lot of work that I've found most informative has not been in mainstream A.I.

Theorists of computation tend to think of A.I. as a special case of what they are attempting to do. They believe I think that a lot of the answers to A.I. questions will emerge from the future correct theory of computation in general. (Actually their problems will be solved by the future correct theory of things in general (see Teufelsdröckh (forthcoming))). I do wish them well the theorists of computation and look forward to sharing their new insights. It is possible that artificial intelligence could provide a route for answering the questions of theory of computation. It could produce its own notion of computation namely that involved in intelligent processes and proceed to develop formal methods for its systematisation.

It is possible that such a notion of computation could include the present notions but somewhat rearranged. The relationship of a theory of intelligence to a theory of computation may have similarities with the relationship of formal logic to mathematics in which each simultaneously provides an application of and a theory of the other.

Then there is mainstream computing or computer systems. This is the empirical workshop of computer science and has yielded many important ideas and techniques. A.I. has also developed its own computer system techniques but the two never quite corre-

spond. Computer system developers occasionally indulge in a little "soft" A.I. In spite of many fine words, there is not a lot of real exchange of ideas and techniques between these groups often doing apparently very similar things. I think this is mainly a result of communication problems, but also of the general state of confusion in both areas.

Examples of A.I. techniques that could have but haven't been tried in practical computer systems include

Planner

Conniver

Winston Data Structures

Frames

Conceptual Nets

Techniques that have been tried a bit include

Learning

The Perceptron

Examples of computer system techniques that have not been used in psychology have been given by Minsky (Minsky and Papert 72). I would say that most have not been used conceptually in A.I. and to his list I would add broadcast networks, array processors, ROM, page, and bus.

Now being a theoretical physicist by extraction, I have found it very difficult to understand a subject unless I know its goals, its assumptions, its methods and its finding. What are the methods of the main approaches of A.I. and what are the findings? How can I satisfy myself that I understand them and that I agree with them?

In the next section I suggest some basic criteria that are used in mature disciplines.

A tendency to a natural science position will probably be found in my writing. I probably tend to view intelligence as a phenomenon to be investigated by empirical means. Some researchers I have talked to of more mathematical extraction tend to have a more purist, generalist instinct and to regard intelligence perhaps as a mathematical system, with all the metamorphosis this implies.

In the rest of the paper, I explain my current state of thinking for an approach to artificial intelligence which could address itself to the basic research questions of a theory of intelligence. The purpose of the paper is to express this to other contemporaries, and also to provide a framework for further elaboration. I have tried to explain my sources both technical and imaginative, but I have included some ideas without a great deal of justification. It is intended to be an ideas paper and its contribution I hope is to have wrestled with a wide range of ideas and attempted to evaluate, select and put them into order to form the basis of a practical research program.

## **1.1 Basic Research Criteria**

### **Comparability**

Research results must be comparable with each other, so that differences can be noted. Further this comparison should

- (i) involve a large set of cases on which the two mechanisms to be compared are tried, not just one or two well tried examples.
- (ii) the choice of set of cases should be justified by its relevance and significance to clearly stated goals in a comprehensive approach to all aspects of intelligence.

By approach to intelligence, I mean something like that denoted by the document you now have in your hand; no, not that one, this one.

- (iii) the results of the comparison should be expressed quantitatively. Lord Rayleigh once said that you know very little of a phenomenon until you can put numbers to it.
- (iv) any results should be duplicated and verified by independent research workers.

Why should work be comparable? Well presumably we are all talking about the same thing and proposing methods and answers to questions. If a new method is to supersede or improve upon another then the two methods must be compared. If two results are not comparable how can one judge their relationship?

The number of cases on which a single mechanism is tried is usually quite small but the number of cases in common between two separate pieces of work is often very small. Let's take a moment to look at existing work.

- (a) problem solving No two general problem solvers have been tried on the same problem. The two classic integration programs were compared (Moses 67). Theorem provers have their favourite common cases although the number is quite small.
- (b) perception The work on line finding is quite well compared, this is probably because line finders are not too difficult to code. The work on perfect line drawings didn't use much common material even in the same research centre. Work on real images can perhaps be temporarily excused until interchange of data is easier.
- (c) language processing comparability is not bad in this very large area; the same sentences have been analysed syntactically and semantically by different groups.

There are dangers to concentrating on one or two standard cases. The later GPS work was a courageous attempt to break out into a large set of cases, and showed how difficult this can be. The monkey and bananas puzzle has been attempted in GPS, theorem proving and in STRIPS. The danger is that the small number of cases are assumed "typical" in some sense and that an analysis of these should give insight into a large set of cases. In the problem solving examples mentioned we have no such assurance. The problems chosen could be untypical, even if typical problems do exist. The methodology which lead to these cases is unexpressed in the literature, usually the problems are produced out of a hat (the hat of Zeus?) with a brief wave of the hand.

In the case of GPS, theorem proving with state variables and STRIPS, the fact that they can all do the monkey and bananas problem is not surprising since the mechanism of working back from the goal is probably essentially identical in these otherwise dissimilar systems. The essence is that the goal is represented by a set of features, operators have preconditions and results which are sets of features, and a single matching chain is found from start to finish by standard search.

It is this problem with typicality which leads me to write (i) and (ii) above. The sets of cases or tasks that have been tried are sometimes arranged in a programme of increasing severity. For example the MIT line drawing work gradually added holes then shadows; after all such line drawings could be done, it was not clear which way to go next however.

The line drawing example illustrates another part of (ii) namely that it was not clear how the line drawing work fitted in with and informed other A.I. work. Having got a line drawing analysed a number of crucial questions remained

- (a) how do we obtain drawings from real images anyway?
- (b) is a line drawing a good intermediate representation?

- (c) how is the knowledge in a line drawing program related to other knowledge?
- (d) is the process of complete analysis of a line drawing without any other context a reasonable fragment of an intelligent system?

Waltz's thesis (Waltz 72) is a model of research. His work is clearly and comprehensively explained and his filtering method extracted from its particular environment.

I think approaches and methodologies should be clearly stated in advance rather than as a post hoc rationalisation to pad out one's latest paper. There was I believe a stated policy at MIT some years ago of a return to empiricism - special systems for relatively narrow areas were to be developed to as high a degree of performance as possible. I suppose the symbolic integration work was one example and the scene analysis of line drawings of polyhedra another.

The Newell Simon school also have number of position papers, as well as a tome (Newell and Simon 72).

Hunt(75), Arbib(72) and Uhr(73) have all expressed their general positions.

Coming to (iii), since not much comparison is going on, there is a fortiori not a lot of quantitative comparison.

Areas which have tried quantitative measures are

- (a) theorem proving one or two comparisons of different strategies on the same theorems have been done (Reboh et al 72), and this is very laudable. Unfortunately, the theorems considered have been very simple and non-typical so that the findings do not convincingly generalise to all theorems.

[The set of all theorems to me is a dangerous one, it is quite likely that the computations involved in intelligence form a performance hierarchy and that some large



classes of theorem are very unlikely to arise in intelligent behaviour]

- (b) heuristic search inequalities have been established by Pohl (Pohl 70) and some quantitative work done by Nilsson et al (Nilsson 69)
- (c) game playing search we have the classic analysis of alpha-beta and other quantitative work by Slagle and Dixon, (69), Knuth (74) and others.
- (d) Semantic nets. Some attempts at mathematisation, for example Barnden (75).

Independent verification (iv) is I'm sure done a lot but unsystematically. Many theorem provers capable of the same strategies and proof methods have been written, many simple heuristic search programs.

Anyone getting into vision work naturally duplicates the work of others. All this informal work however simply verifies that the technique works. It never leads to a broad study of the technique or to a verification of the implicit conclusion that the technique is better than other variants.

Comparability is also a major difficulty in other if not all areas of computer science, in theory of computation and in computer systems.

Questions "Where are the papers giving systematic comparisons of  
Prolog and Planner?  
Planner and Conniver?  
Planner and GPS?"

and yet how many of us stand up in class and spout "we now know that  
x.....?"

## **Operational criteria**

The definitions and usages of concepts and research ideas should be relative to an overall intelligent system; as I will explain below I prefer the notion of an autonomous surviving robot. Words that need defining properly so that one has an operational criterion for determining whether they apply to a proposed mechanism or not, include

"intelligence", "concept", "learn", "understanding", "perception", "problem", "solution", "semantics", "knowledge" etc.

## **Mathematical theorisability**

True research results are those data, concepts and insights which transfer from one project to another. "What do I know now as a result of your work that I didn't know before?" assumes a body of knowledge which applies to a wider range of activities than the specific work.

We should try to obtain quantitative results and general theorems.

Any research approach should have a mathematical theory which describes the knowledge within which additional new insights can be gained by mathematical analysis.

Thus, it is part of my approach that these three basic research criteria should be taken seriously in practice. Have you ever been to a conference or even a seminar that addressed these criteria?

## **1.2 Levels of work**

I tend to think of research work being done at several levels which intercommunicate, but with an overall inductive flow i.e. from practice to theory. The people who do the research all have their own unique individual style but presumably gravitate to a

research style, method or posture which has a place in the (international) community of researchers and performs a useful function in the overall effort.

### **1. Ad hoc programming**

The strictly ad hoc programming or hacking of robot control and the ad hoc design of computer systems have their place, but really only indirectly as expedients for getting something working, so that other less ad hoc parts can be studied.

One example which springs to mind, and is perhaps on the borderline with the next level was in the SHAKEY computer system for which a large amount of support code was hacked.

If a scene of an object could not be successfully analysed, the robot would be moved round the object a bit and another attempt made. However this idea was not pointed out as a phenomenon and certainly not as a concept.

### **2. Phenomenology**

The design of special systems for new tasks and the extraction of techniques or phenomena from such experience but without an attempt to generalise or transfer such techniques to other areas. Examples of this are the integration programs SAINT and SIN and the vision programs of Roberts, Guzman, Winston and Waltz. The latter has been generalised into the next level.

### **3. Concept Formation**

When techniques from many special cases are summarised and put together, then we have what one might call a concept. Examples are AND-OR trees, backtracking, Waltz filtering, production systems, pattern directed procedures. It is not clear to me that a concept can be extracted in general from Winston's data structures, but I hope so.

A concept then is normally a control or data idea with a number of existing realisations

that work.

#### **4. The Formalisable Form**

After the concept, which exists most concretely as its set of realisations as programs, each of which may have several rather complex parts, we need to find a simplified elegant version, which is the essence or formalisable form of the concept. One problem is to pitch the level of simplification such that the system will still perform satisfactorily.

Examples are the formalisable forms of heuristic search, GPS and of theorem proving strategies.

#### **5. Mathematical Theory**

This is the cleaned up form that mathematical symbolism and reasoning can be applied to. Hopefully there is actually some content left by this stage.

Examples are alpha-beta pruning (Slagle and Dixon 69), the perceptron (Nilsson 65) heuristic search (Nilsson 69), GPS search (Ernst 69)

Planner (Hewitt 72) is an exceptional example in providing a theory as well as an original computational method.

What I meant to say is that part of my approach is to be clear about the research activity one is undertaking, that the five categories above make some sense to me and that any research work must take its meaning from its place in the overall research effort. Any work should be seen as far as possible in its overall context and attempts made to think through or elaborate its relationships to this context.

If we look at recent conferences and other papers in A.I., I think we see a lot of confusion about the place of individual projects. A lot of projects select tasks which are special purpose and difficult to generalise from, the methods used are unclear technically but rather ineffective and the work falls between categories 1 and 2. There is a lot of work

which claims the discovery of a new concept but falls between 2 and 3 in not establishing the technique as a clear concept.

### 1.3 Psychology

I think of the aim of artificial intelligence as being to explain, describe and produce a theory of intelligence. Just how general one's theory can be is partly a matter of taste but partly unknown at the moment. I imagine there can be a theory of intelligence more general than human intelligence and that human intelligence would be explainable and be an application of the general theory.

It is certainly true that one can give a robot sensory modalities that humans do not have such as ultrasonic detectors, laser vision that detects depth directly, microscope eyes etc. and the type of processing elements and computer architecture can be quite different, but whether this can lead to perceptual and cognitive phenomenon very much different from man is not clear.

I think of artificial intelligence as producing computational models and theories of such, and of psychology being their application to the human case. Having said this, given that artificial intelligence is rather undeveloped, I think that psychology can provide a lot of good concepts and phenomena and that a good approach is to take these seriously and to try to ask what the general significance of each psychological finding is for artificial intelligence.

What psychology has to offer A.I. is experience in conceptualisation and knowledge, descriptions and measurements, of phenomena and effects in the study of the only existing intelligent machine of any sophistication, namely man. What A.I. can offer in return are computational techniques and models, and approaches to theorising about such models.

Some examples of the psychology that I have found interesting are

- (i) theories and models of attention including, sensory and short term store see for example (Broadbent 71). The need for an attention mechanism for robots is not widely recognised.
- (ii) Gibson's theory of vision (Gibson 66) using motion parallax and texture.
- (iii) Bernstein's theory of motor control (Bernstein 66) using a flexible hierarchy
- (iv) and of course Piaget, see for example (Piaget 53), who is the major influence at a more diffuse level.

## 1.4 Applications

My approach to applications is discussed elsewhere. It is one of the attractions of computer science that one can work on imaginatively interesting and intellectually demanding tasks and also be close to their application in industrial and commercial fields. This gives one a social "engagement".

My main approach to applications is to study not a direct industrial analogue of the autonomous surviving robot but to study computer assisted robots in which a man operates a tool or machine but is assisted by a computer. I am interested for example in computer assistance for remote teleoperators in nuclear power stations and in undersea work.

The reasons for this approach are several. The computer assistance problem at the levels required industrially I find more interesting than the problems in industrial robots capable of completely replacing a man. The type of engineering involved in remote operation is also broader and nearer the limits of engineering technique. The assistance of a man

in a very hazardous environment is more socially acceptable to me, I believe that the industrial areas with computer assistance should be developed before or in conjunction with any replacements made by autonomous robots in semi-skilled jobs.

## **1.5 Summary of my approach**

Upon the three research criteria as basis, I put forward four pillars of my approach; each informs and justifies the other. These are choices or assumptions from the totality of possible choices and assumptions. I am saying these are fruitful. The rest of the paper explains, expands and explores these ideas

### **1. Behaviour patterns**

We should work with total behaviours, rather than fragments of systems.

The environment in which the behaviour takes place should be explicitly specified.

A real robot should be used and real environments, as far as possible.

### **2. Performance**

We should make the cost of computation, in space and time, integral to our conception of computation.

In particular time complexity especially real time constraints are essential.

The real time constraint probably only makes sense on a highly parallel architecture.

### **3. Development**

It is essential to consider the generation of the system, which is built up incrementally by information only available at each step and with capabilities only available at each step.

#### **4. Parallelism**

Parallel computer architecture should underlie the system.

I am currently interested in uniform parallelism using array processors, more exactly a small set of strongly interconnected array processors.



## **2 Robot behavior patterns**

### **2.1 Real robots**

The advantage of real robots as opposed to simulated ones is that they provide a harder and truer test of one's methods, a test of which is obviously operational. Of course one can fake robot work.

Robots also provide an extremely rich environment of problems.

We want a real robot and we want it controlled by real computer hardware. The investigation of computer systems, meaning hardware and software, for robot control I see as a fertile and practical way to investigate artificial intelligence. Hans Ernst (70) has written an interesting article trying to summarise robotics from an epistemological point of view.

The robotics approach to artificial intelligence is the study of the design of computer systems for robot control.

### **2.2 Total behavior**

The idea is to work always with the total behaviour of the robot and not just some supposed fragment like vision, problem solving etc. Figure 1 shows a typical experimental robot.

The system we study is a robot in an environment and the environment has a description and is important, in fact we should talk of the robot-environment system. We are interested in complete behaviour in the sense of including perception, decision and action, in the sense of complete meaningful sequences and in the sense of the complete performance of the system for the entire time axis.

I think a lot of Artificial Intelligence work suffers from hidden faults because it works with a presumed subsystem of the total robot control system. It is often the case that output produced by such a subsystem is really unusable in the total system. The demarcations have been drawn without regard to any total plan and further without any responsibility for getting it right. Conversely, often subsystems work very hard and expensively to produce comprehensive information that a robot may never need, or have time to use.

As regards mathematical theory, we hope to be able to formalise the robot-environment system and to define the notion (or notions) of behaviour pattern(s). The robot perhaps computes on the environment or the robot-environment system computes. My idea here is that we restrict the computations considered (the functions computed) to those needed to describe behaviour patterns and that this set of functions would be much smaller than the set of all computable functions. Or perhaps the set of all computable functions is generated in a certain order or hierarchy and this order is a significant fact about intelligence. I am actually going to introduce further restrictions in the next section to make this more plausible. I also hope that the consideration of the computations involved in meaningful behaviour in a reasonable environment might suggest different questions to theorists of computation in general and hence suggest a different notion of computation as such. Computation in general would be the behaviour of computers in environments.

We seek a structure in the set of total behaviour patterns and also descriptive mathematisable concepts which apply to total behaviour patterns. What is the value of the behaviour pattern? What behaviour pattern is produced by this machine in this environment?

The theory must also formalise our idea of **complete behaviour patterns**.

This requirement leads us to study simpler behaviours than is usual in artificial intelligence but to ask nonetheless demanding questions of them.

One classic model for theoretical research is the well-known work on automata and languages which managed to classify automata and the languages they process into categories and to show the relationships between them.

The orientation to sensory motor skills and complete behaviours allows a different factorisation of the total intelligent system into "**horizontal**" rather than "**vertical**" divisions, see Figure 2.

## 2.3 Environments

Environments are neglected in artificial intelligence. Young (76b) discusses methods for simulating a robot's environment, Hendrix (73) and Lowrance and Friedman (77) discuss Hendrix's modelling method which is intended for the robot's internal model of its environment.

Let us briefly consider some real environments, in Figure 3 we show a robot in a rather confusing state-of-the-art environment. Perhaps the best known environment is a real **blocksworld** environment, see Figure 4, following Winograd's simulated blocksworld. A real blocksworld was used for SHAKEY (Figure 5), except it hadn't any arms and so couldn't access all of a blocksworld environment.

Other everyday environments include **outdoor scenes** (Figure 6) and **office environments** (Figure 7). These have been studied for vision (see Ohlander (75) and Tenenbaum and Barrow (76) respectively) but not from the point of view of robot action.

**Industrial environments** (Figure 8) have recently come under scrutiny for robot control. This is the typical "factory floor". This environment is already highly stylised and documented but is still very complex. Colour and texture are probably not good visual cues.

Another environment is the **partial environment of language** , of words and "utterances".

One can try to describe the actual environment of utterances used in a complete environment, presumably including written forms as they occur.

The **infant's environment** (Figure 9) (European version) is something like a playpen. It also has the **breast/bottle environment** (Figure 10) made famous by Melanie Klein, and there is also the initially single human or "**caretaker**" **environment** (Figure 11). The infant's language partial environment is discussed for example by Clark and Clark (77).

## 2.4 Taxonomy of environments

Man's hardware was designed/evolved in a sequence of environments involving water then trees then the plain. Presumably the most recent modifications occurred in the environments of early man, see Figure 12.

One can perhaps define a hierarchy of complexity of environments. Different intellectual functions will be required to survive or equilibrate in different environments.

It is important to be able to imagine different levels and types of intellectual or information processing ability.

Piaget is often said to have been the first psychologist to take the child's mind seriously in its own terms.

## 2.5 Hierarchy of function

Hierarchical organisation is important in motor control, see Bernstein (66), Greene (72) and Kohout (76).

Bronson (65) has proposed a three level structure for the brain, each level has its own response and own motivations.

The hierarchy could be very flexible, allowing processes to alter data at other levels etc. Most hierarchies proposed, like Winograd's natural language system (72) and Fahlman's BUILD system (74) are very hierarchical.

We could perhaps define a hierarchy as a restriction upon a process to only communicate with its not necessarily immediate ancestors, its brothers and its not necessarily immediate offspring, see Figure 13.

## 2.6 Fuzziness

One is lead to having to deal with the apparent "imperfections" in the real world. Noise, inaccuracy and other types of fuzziness and probabilistic effects find their way into your system and have to be integrated into it. I say apparent imperfections because it is conceivable that some intelligent mechanisms will only work in fuzzy environments and not in the smooth limit. Fuzziness allows and indeed forces generalisation of experience. I am thinking here of the fact that one cannot walk without friction and that we would be unable to move without entropy increase.

## 2.7 Continuous functions

Robots also lead one to grapple with continuous functions and real variables in image and speech processing and in motor skills. This need to include continuous descriptions as well as discrete symbolic descriptions also I think gives a better view of intelligence.

The sensory motor basis of human intelligence postulated by Piaget and the motor skill approach to thinking advocated by Bartlett (Bartlett 58) are more difficult to investigate in a purely discrete system.

## 2.8 Motivation

It is likely that the description of a robot control system will require a specification of a **motivational system** or **motivational theory**. That is, "doing what it is told" may not be well defined enough for complex behaviours.

## 2.9 Stability of behavior

Another area which is completely untouched in artificial intelligence is the question of **overall system stability**. This strongly interacts with learning and the development of the system.

## 2.10 Natural decomposition of intelligent robot processes

In studying robots from a performance point of view one is lead to ask different questions from normal artificial intelligence and one is lead to study different functions and phenonema, to analyse the total system into different parts. A lot of these seem to be turning out to be the functions studied by psychologists and biologists in "living robots".

In vision, we are lead to study:

- (1) The attention function and models of attention.
- (2) Visual registration and orientation.
- (3) Form Recognition function.
- (4) Camera motion and choice of view.
- (5) Multisensory confirmation of perception.
- (6) The role of vision in problem solving and vice versa.
- (7) Low level vision as a separate mechanism, this evokes D. Marr's ideas (Marr (75)).

In the database, we have found separate studies in:

- (1) The spatial map of environmental layout as a distinct separate system.
- (2) Temporal event (episodic) memory.

In natural language we ask:

- (1) What are the uses or functions of language, this evokes Halliday's study (73) of language functions and their acquisition. It also asks whether language uses are speech acts (performatives).
- (2) What do we want to talk to a robot about?
- (3) The use of speech, is it simply equivalent to a teletype channel?
- (4) Nonverbal communication and the perception of other robots/ humans, the sharing of event perceptions.

## 2.11 Functional models and implementation models

I give finally three diagrams (Figures 14, 15 and 16) giving very rough structures of the total robot control system using

- (1) Mainstream A.I.
- (2) Cognitive psychology.
- (3) Neurosciences.

1 and 2 one might call **functional models** in that they do not indicate performance or implementation.

3 I call an **implementation model** since it is related to actual spatial arrangements of neurones.

## 2.12 Computer system design

**Computer systems for robotics.** are made up of all levels of design down to the sensors. It is quite likely that new computer architectures will be designed for robotics since the information processing requirements of robots are severe. We have very high data rates not only for vision but also for other sensors of which there may be many hundred @ 12-bits wide.

The storage requirements for the system are similarly large. A single TV image can occupy anything from 10K (100 x 100) bytes to 1000K (1000 x 1000) bytes. The environmental model or map may also require a large amount of store. A computer system is presumably made up of a number of processing elements (or "PEs") connected by channels of various widths and speeds. Typical processing elements are arithmetic and



logic units (ALUs), instruction decoders and control registers (CUs), RAMs, ROMs, disc units, associative stores, etc. It is quite possible that special new types of processing element will be designed for the robot application. Non electronic computing mechanisms may exist on a robot in the form of mechanical linkage systems and optical systems.

With the advent of microelectronics in LSI form one can contemplate the construction of special computer architectures; at the very least one can have a network of microcomputers. The major scarce resource in present day computer systems seems to be channel capacity or more exactly the ability to get two data items "into contact" i.e., into the same processing element, sufficiently rapidly. This resource is now more scarce than memory size or processing capacity. A computer system is "distributed" in space because **"data is distributed"**. This certainly applies to a robot; it has spatially distributed sensors and it has a large memory of data and programs, some of which will be spatially located on backing stores.

As we shall enlarge upon later, we have been drawn to "highly parallel" computer organizations as being appropriate for high data rates and real time response. The term highly parallel is taken from Murtha's classic article (66) and is intended to mean an organization with a large number of interconnected PEs all working on a single problem. I distinguish this from a looser network of PEs where each PE has its own task which has a meaning independent of the total network. The highly parallel structures currently available are associative stores and processors and array processors.

Running on this distributed computer will be software. There will probably be a system level - a real time system for handling interrupts, tasks, messages, etc. It may also run language interpreters or robot programs in compiled form. These constitute the robot software proper with visual, planning, language and learning computations.

The design of suitable high level languages for such a system may be a fruitful research

strategy.

When it comes to mathematical theory, the formal computational model must be based upon the real computer system. Very little is known of such systems or such theorising, see Hewitt (77) and Hoare (78). A particular study may of course take a particular model and explore its properties.

## 3 Performance

### 3.1 A notion of performance

By performance I mean the cost or efficiency of computations and the time taken. It is important to build the time taken into the notion of computation right from the beginning and I therefore thought of defining a **performance** as a computation carried out in a given time. Thus two identical computations carried out in different times would be different performances, and two different computations carried out in the same time might be considered together whereas the identical computations with different times might be put in quite different categories.

It is very important to be working in terms of performances rather than functions because intelligence must surely involve efficiency and speed of computation. McCarthy and Hayes (68) drew attention to what they called the heuristic and epistemological aspects of representations. It seems that we do know a number of epistemologically adequate representations such as 1st order predicate calculus, w-order predicate calculus and the lambda calculus. A main problem is to use the information rapidly, otherwise we could use the "British Museum algorithm" i.e., exhaustive search, represent everything in predicate calculus and we would have an intelligent robot. I am arguing that efficiency is necessary, not that it is sufficient.

Unless the robot can respond rapidly enough it will not be intelligent. One possible formal approach then to intelligence would be the study of computational complexity. By introducing the additional constraint of studying only robot functions in their natural hierarchical order of complexity, the natural order being derived from an ordering on complexity of environment and performances in the environment, one may be able to make more headway than theorists studying complexity of computation in general.

Next, I really want to be able to produce a theory that will be relevant to psychology. Here again the main lack in artificial intelligence and the main reason why the results of artificial intelligence are not very relevant to psychological modelling is that psychological models model the human mind, which is a highly optimised performance system. Timing has enormous importance, and indeed if we believe that this is a result of evolution the human body presumably is just able to survive when pushed flat out in the most testing environments of the last few million years.

### **3.2 Performance and "normal" AI**

Again, some artificial intelligence results are a result of efficiency requirements, usually imposed by the limited computer resources available to the individual artificial intelligence worker. This has been a rather haphazard constraint, and is regarded more as a nuisance than a welcome discipline. Artificial intelligence workers have not reported such performance data in their published papers very much, Ph.D. students are not required to give performance data in their theses. This is not entirely their fault, it is very hard to compare the performance of two present day computers and so any performance data would not generalise very accurately to other machines, however it would not be too far out, it is unlikely for example that more than a factor of ten discrepancy would be involved between the most disparate mainframes. However, hardware is not the only problem, most artificial intelligence software is written on top of other artificial intelligence software. In the simple case, we have a special A.I. language like Interlisp or Pop-2, but quite often an extra level of interpretation is used. This has lead to a lot of welcome representational and language innovation but it does mean that the program runs as a data structure which is interpreted by a program which is itself a data structure being interpreted by another program.

In these cases, some vague figures like factors of one hundred are estimated for the effects of multiple interpretation. Lisp interpreted compared to Lisp compiled is often about a factor of one hundred. In a comparison between the high level language form and the handcoded assembly form, in the case of A.I., a factor of ten is about right, for example Scott's chess program gave this.

Putting all this together, we could compare the performance of A.I. programs or mechanisms by recoding them all in a standard language, but nobody has the resources to do this. The closest attempt to this was possibly the work of Michael Rychener (76) at Carnegie-Mellon. More importantly, no-one has the inclination as performance is not deemed important. Even supposing that this were done the results would be for present day computer architecture. There are two points here. First, the instruction sets are evolved by market pressures from payroll programs, "scientific" computing and now time-sharing efficiency. Second, to be relevant to psychology, we should surely be using a highly parallel machine. Parallelism of course alters greatly any hierarchy of complexity found for serial computation.

Again for psychological modelling, one could simulate a parallel machine on a serial one and then optimise A.I. mechanisms on this simulated parallel machine. It is quite likely that on present day machines, the cost of running programs on a simulation of a parallel machine would be so high that the mechanisms investigated would have to be limited to rather simple ones.

As most A.I. is done in university environments on highly time-shared machines, the main limitation facing the individual worker has been memory space rather than processor time. Even on paged systems, the memory space requirement is the main one, being now the size of the working set. This constraint actually works against the development of speed optimised A.I. mechanisms. The best empirical environment is a dedicated machine, presumably a mini or micro-computer.

### 3.3 Real time

I came to the conclusion that a real time requirement was what I needed. This means that the environment is time dependent and has autonomous processes which run independently of the robot. The behaviour patterns of the robot involve interacting with time dependent objects and hence each step in the robot's behaviour pattern must be computed sufficiently rapidly for it to be available when needed. Since one can buffer, this does not mean that each new step has to be computed in the same amount of time.

There is some theoretical work on real time complexity in which case a real time computation is one that takes time  $n$  to process an input message of length  $n$ , producing the  $n$ th character of the output result after time  $n$ . Are the functions I am talking about real time computable in this sense, and if so is Yamada's and others' work helpful to us? I do not yet know.

Comparing our notion with a control strategy as in control theory, about which I know nothing, the robot makes a step in each state, the notion of a strategy viz. a function which computes for each state which action to perform, may not be very helpful really. The real time computation requirement is that the strategy function can be computed fast enough so it can make the action decision in the time available.

There is a further convolution which is that a perfectly good strategy may not be real time computable whereas some other not so good strategy might be. I don't think control theorists worry about the time for the controller to compute the next step, they probably have other problems.

### **3.4 Hardware and the unit of time cost**

Because I want the real time theory to be applicable to artificially intelligent machines that could be built, and because I seek security in the concrete, I thought I'd base my measures of computational cost and time shamelessly upon hardware or more precisely upon a conceptual hardware close to real hardware (this gives ultimate security!).

There are various types of hardware. Clocked automata which make one transition per unit time. Our unit of time for computers is between 1/10 and 1 usec.

Serial computer processors make one instruction per unit time. Since an instruction usually makes one memory access and since this is the time limiting part of the instruction we can make one memory access as the unit of performance. A table look up form of an automaton, or a RAM form also gives the same natural unit.

In a later section, we introduce parallel hardware and will extend this definition at that time. Because our measure of complexity is based upon real hardware (and software), the optimisation we are interested in is the same as that studied for general computer systems. This connection between systems viz. integrated hardware, communication, time-sharing, file system and language design and artificial intelligence has always been officially ignored but always occurred in practice. Computer systems designers unfortunately are not interested in robots as users but in humans. Humans have local intelligent processors called brains which make them very different.

### **3.5 Hierarchies of real time response and the theory of iteration**

I realise that one can always define a hierarchy of anything, thereby introducing a fruitless complexity which appears to carry meaning but actually usually gets in the way of clear thinking.

I would like to be able to work with hierarchical hardware that is a set of computers in a line, or a tree, see Figure 17.

The first computer might give local simple control, the next could give a slower more complex and better response and so on. In this way we get a hierarchy of real time responses of increasing time delay and increasing complexity.

Such a configuration seems to occur in applications due to economics, we have a local micro, then a mini, then a mainframe. The data transition time allows the local one superiority over the higher ones for quick simple responses.

Hierarchy of function may be related to a hierarchy of real time response.

Note that there is no artificial restriction on connections, the raw inputs can go to any level and raw outputs can similarly ensue, see Figure 17.

However this will probably have little effect. The upper levels could produce corrections to the basic response and one could regard the overall response of the system as an iteration, which approaches a "correct" value as corrections are added in.

$$t = t_{i0} + t_{i1} + t_{i2} + t_{i3} + \dots$$

Iterative approximation has been defined and studied by Tzichritzis (69).

### 3.6 Performance models

These resemble functional models in consisting of black box models whose input-output relations are known but not their specific implementation, but with one important addition. Each box, each link between boxes and each interface has a specified performance. The exact type of performance measure and specification may depend upon the particular model but I mean a measure of the time taken and also other resources like space.



The time may be specified as a response time to a predictable request or a real time response to an unpredictable request. The performance could be specified by a constraint on timings and likewise a constraint on the usage of other resources.

## 4 Development

### 4.1 The developmental constraint

The third pillar of my establishment is that I am committed to a developmental approach to intelligence. This means that we consider sequences of behaviour patterns, the first behaviour pattern being the initial state of the system and the others successive states as the system learns. The most fundamental consequence of all this is that the form that the system takes, presumably it is the software only which changes, must be such that it could have been learnt and built up incrementally only using the information available and useable each step. The starting state of the system need not be completely *tabula rasa*, it could probably be primed with "instinctive" "reflex" control. Around this initial core and the description of the environment, we must have a generative system which generates each state of the system in turn by modifications to the old. I think this should give a fairly strong constraint on the class of systems involved because

- (a) the system must satisfy well formedness constraints at each point, e.g. survival, and utilisation of data base so far.
- (b) the information upon which modifications are based must enter and be processed by the system so far.

An example of structure occurring because of a developmental constraint has been recently described by Minsky (77), who postulated that the short term memory structure used for passing data between processes would have to remain rooted in the structures in the initial form of the system, but would elaborate upon this root.

The distinction between behaving and developing is unclear and perhaps we shouldn't try to make it. For example, a robot may explore its environment so that it then knows

it and can then move around more easily; this can be regarded as a learning experience or as a perceptual experience.

We have a natural division of the system into the hard i.e. unchanging part and soft i.e. modifiable part. This gives my **artificial intelligence definition of "hardware" and "software"**, which is implementation independent!

There are those who say that we cannot describe learning until we have first described behaviour. I am pointing out that the learning constraint can be very far reaching. The representations obtained from a study of behaviour only are very unlikely to be suitable for understanding the development of the system or the relationship between different systems with different behaviours. So the "find the behaviour first" view is not necessarily prudent.

I also think of the developmental approach as being an alternative to what one might call the complex programming approach. The latter is to find a complicated program that produces the desired performance. Surely we should aim for the minimal complexity for the job and mistrust complex solutions. The various A.I. mechanisms that have been investigated have not benefited from the additional complexities added to boost their performance. These have merely obscured the performance and prevented critical evaluation of the basic mechanism.

I realise that one cannot know whether one's latest complexity is a hack or a new principle, but this is part of the research discovery at the end of which, upon analysis, the inexplicable complexities can be identified as hacks and removed. By complex programming the researcher puts lots of unquantifiable and unjustifiable information into the system. He cannot explain how this information could come to be there.

A researcher designing a system and following my approach would have to explain the development of his system and explain how each piece of knowledge was generated. It

is quite likely that without some such discipline on the overall well-formedness of a knowledge system, the system could become very ineffective as a result of adding more "knowledge".

## 4.2 Hierarchy of control and development

Following Bronson's (65) model, each level of a hierarchical system would have its own adaptation and its environment would involve other levels.

## 4.3 Learning Sequences

Having defined certain environments and certain behaviour patterns, we can now define standard classes of **learning (or developmental) sequence**. This is an ordered set of behaviour patterns in a given environment.

Examples are

- (i) learning how to find a light (Mott 78). Starting from a system which only makes random movements perhaps with a back-off on touch reflex, we progress to the behaviour pattern of moving towards a light. A more complex version of this sequence has been investigated by Mott in which the system also has to empirically connect the light with food and food with the absence of pain.
- (ii) learning to grasp on sight (Woodward). This sequence has been observed by Woodward and consists of several stages involving (a) holding the head still, visually fixating an object,
  - (b) grasping an object on oral contact
  - (c) grasping an object in the same visual field as the hand
  - (d) grasping an object on visual contact when hand is out of view.

Note that object permanence is not necessary for this sequence.

- (iii) learning to use language. (Halliday 75) sequences of systems have been described by Halliday in which utterances of single words and eventually strings of two or three words are related to their uses in the totally functioning organism.
- (iv) learning to interact (Bruner 78)
  - (a) shared experience and looking
  - (b) dialogue rhythm
  - (c) naming
- (v) deixis (Clark 78) The I-you distinction precedes and informs the formation of other deictic distinctions such as here-there, this-that, come-go and bring-take.
- (vi) seriation (Young 76a) Richard Young has attempted to model the behaviour of children trying to place a pile of blocks in order of size. A progression of strategies seems to correlate with the accretion of rules.

#### **4.4 Stability and Attention Criteria**

As a system learns, a number of types of instability can occur. The usual phenomenon observed is an exponential blowup of stored information, see for example the Case Institute game player (Citrenbaum (72) ). It would also seem that any kind of rote memory approach meets the same fate. What seems to happen is that the programmer devises a data or other computer structure to be as useful as possible. The system then builds structures to increase its performance, but then reaches the limits of the representational ability of the convention used. The more recent schemes of graphs and frames used by Winston have not been tested on complete environments, but it is quite possible that they will exhibit a saturation phenomenon at a fairly low level of intelligence. One should

distinguish between limitations in the knowledge it is logically possible to represent and the limitations in performance of a system due to too much data being required.

The PURR-PUSS system, (Andreae (78) ) being based on a rote memory of sensory fragments will probably also hit a saturation effect.

The "cell assembly" learning system of Cunningham (74) exhibited the problem of adding large numbers of similar new assemblies without any obvious way to select the most useful ones. It also exhibited another type of instability to do with its attention control. After a while, the system got into an attention loop in which it ignored all except a certain type of input and its outputs stimulated these types only. As a result of learning, the system concentrated on one stimulus and shut out any new experiences. This attention control is quite crucial, most animals and also humans have an innate system of behaviour upon birth which not only ensures survival but controls attention patterns around which are built more advanced processes.

Most animals and humans are born with a visual system that innately attempts to follow information-giving information. The eyes track moving objects and nystagmus keeps objects focussed and held in the field of view so that a higher level system can learn from this.

The innate low level vision used to "prime" the higher level one has been shown in chimpanzees to be mediated by a separate part of the brain, emphasizing its separate function. It is interesting that such a priming system is in this case the "old" system from an earlier evolutionary stage of development. A role of teacher and child caretaker is in providing extended attention patterns and in preventing attention loops.

## 4.5 The Development of Mathematics

The axiomatic development of mathematics is closer to what A.I. workers tend to do, than is human development.

Mathematics can be formally developed in several different ways; geometry can for example be developed from points or bodies. The development of geometry by adding axioms would start with projective geometry, the development of algebra from semigroups etc. It is difficult to specify performance criteria for a mathematical system. One should really define development in terms of behaviours instead of abstract axioms.

The abstract development of mathematics has no simple relation to the development of mathematical ability in humans.

The sensory motor knowledge acquired in infancy, if axiomatised would be highly sophisticated, but of course such skills are not "known" in the same conscious way as is required for symbolic manipulation.

To comprehend the notion of set, operation etc. a human will probably need advanced sensory motor knowledge. In order to discover proofs we may need to use imagination developed nonsymbolically.

Because of this to me representational convolution in the development of symbolic manipulative intelligence, I find it a difficult area and not a good candidate for the study of development. By contrast, human sensory motor development does not suffer from such a problem. Further, more is known about it.

## 4.6 The Representational Regression

When we take development seriously, we need to specify an initial system from which intelligent systems develop. But at what level or standard do we take the initial system? What initial concepts and structure are to be initially specified or programmed by the experimenter? Can each experimenter define an initial system suitable for his own particular purpose? One problem is that some apparently simple structures are very powerful and one might build in structures which directly contain within themselves the other structures that it is proposed to develop. Is there a fixed order of development of ideas or are there very many? As we are up against the limits of our imagination, we might try to follow known developmental sequences.

But where to start? In theory we could start anywhere so long as we can specify the system at that point. This has been attempted by child modellers such as Young (76a) but in general most of us experience what I call the representational regression, that is we are driven back to the simplest conceivable behaviours. Some people go back to simple animals, I usually go back to the behaviour of the human neonate.

This system is not so simple. Its I/O (or physiology) is very complex. Munn (65) has listed 75 different reflexes that the neonate has each with a specified input pattern and a specified muscular output pattern. In addition there are coordinations i.e. two different reflexes activating or coexisting with each other and there are sequences which are initially present. We don't have to use exactly the same system, just one that is conceptually similar.

There are two further problems, first that the innate learning mechanisms may be complex and second that growth and maturation effects may obscure the true nature of the system.



## 4.7 Infancy

We are lead then to study infancy; that is, the first two years of human life. Not enough is known of what happens to give a proper specification. The amount of data now available on infancy is colossal. Piaget originally postulated six stages on sensorimotor intelligence in the first eighteen months. There is a development from the initial state of the system by differentiation and refinement of actions and coordination; that is, synchronisation and linking of actions. There is development in problem solving strategy. Later in infancy, we see the development of language, first babbling then single words and finally two and three word utterances. There is a development in overall organization of the system, in attention patterns, and in motivational refinement.

Sensory motor intelligence then is seen as the basis of other intelligence. There are many important questions on the representations that are developed and the overall pattern of behaviour of the system developed. These structures form the basis out of which more advanced thinking develops.

Language is learnt in the context of an existing non-linguistic sensory motor intelligence. Even if the semantic system is developed as an additional separate system, it is difficult to see how it can operate in the behaviour of the system without understanding its control relationship to the sensory motor system. In my fanciful way I tend to think of a bilateral model with a separate semantic system in the language hemisphere and a sensorimotor system in the other.

These two systems share tasks, each doing the things it is good at. The sensorimotor system provides functions like imagery and simulation which are used for inference in the semantic system. The semantic system can control and evoke sensorimotor processes from semantic tokens.

## 4.8 The Brain and Evolution

There is another, two level, type of development corresponding to the phylogenetic and ontogenetic distinction. The phylogenetic development or evolution of the system corresponds to hardware development. One can then compare different stages of hardware development by allowing software or ontogenetic development to occur in specified environments, and noting the results.

Evolutional sequences have been described by several workers for example Olds ( ) and Scott (75). Olds defined the sequence of add-ons as the reticular formation, the cerebellum, the hippocampus, the neocortex and bilateralisation of the neocortex. These were to be correlated with attention control, fine movement control, short term memory, long term memory and language function. There seems to be a sequence of the "chips", used in uniform structure models of the cerebellum, hippocampus and neocortex respectively. Bronson's model (65) could be formed by an add-on evolution of layers.

## 4.9 Growth

The known learning sequences in infancy have the complicating factor of maturation. By maturation one means temporal change independent of experience of the external environment. At the anatomical level, maturation can be divided into growth and differentiation, growth being the replication of units without modification. See for example Balinski (75). The growth in physical size of the child means that the environment is changing and this will have to be taken into account. The growth of the nervous system is also quite large in infancy. Some attempt at representing growth by cellular models has been made (Arbib (72), Baer and Martinez (74) ), but no attempt has been made to integrate such models with cellular models of the behaviour of the nervous system in carrying out the computations involved in behaviour patterns. The differentiation of cell

function by chemical gradients etc. could probably be integrated into the nervous system performance, however simple growth i.e. the creation of more new identical cells without differentiation, produces phenomena not easy to interpret in neural behaviour models.

Examples of developmental changes produced mainly by maturation are

- (i) posture and walking development (see Figure 18)
- (ii) short term memory size
- (iii) the onset of language learning.

## **4.10 Sensorimotor learning in the computer literature**

By concentrating on sensorimotor learning, or more exactly infancy, we concentrate on certain environments, certain behaviour patterns and certain learning sequences. The task is nevertheless formidable. Learning in general is meaninglessly too general for me.

The various learning programs in the computer literature have scratched the surface on the issues Jones (71), Cunningham and Gray (74), Doran (68), Wakeman (70), Sussman (73), Becker (73), Uhr and Kochen (69), Winston (70) and (78). Different learning programs have been tried on different environments, but from our point of view we are interested in what other work can tell us about sensorimotor learning. We can do this by applying known representations and learning methods to the sensorimotor case and also by studying learning in other environments for ideas about what to expect in the sensorimotor case.

## 4.11 Theoretical treatment of learning

Having developed, by an empirical process, the concepts such as environment, behaviour pattern, learning sequence, specific computational models or classes of machine and learning principles or methods, we hope to be able to extract a well defined and simple enough form of system that can be theorised about mathematically. The questions we wish to ask or properties we wish to prove depend upon what our experience suggests to be the case.

How does an initial system develop under certain learning principles?

Are certain sets of learning principles equivalent to each other?

Are systems with certain properties unreachable under given constraints?

Are certain developmental systems likely to exponentially blow up?

What constraints do stability and survivability make on the system states reachable?

## 5 Parallelism

Our fourth insistence is upon parallel hardware and computational models derived from it. This requirement means that performance measures are calculated for systems with parallel hardware, and the empirical work is done with parallel hardware. Of course we cannot survive without some use of simulation on serial machines but this is to be eschewed.

Thus behaviour patterns are to be programmed on the parallel machine, theories of behaviour patterns produced, theories of parallel mechanisms studied and, development on parallel machines studied and theorised about. As remarked in the previous section, the machine itself remains constant during (ontogenetic) development.

The statement of parallelism means very little really unless we can say more about the sort of parallel architecture involved. I am currently very interested in a general array processor model as the class of parallel hardwares. My understanding of these systems is not well enough advanced to insist on this, however. I prefer them to the very general schemes being currently put forward as "distributed computing systems", because they should be easier to program, easier to prove things about and lead one more to an "intelligent memory" style of programming.

One is lead to parallelism out of performance needs. The general computer system these days consists of a set of processing elements of some kind which may be processors, memories of various types or other automata. The set is connected according to some scheme or other - bus, general graph, broadcast net, etc. Having got to this point one's mind is so blown with the vision one has created, that one is lucky not to remain in a state of perpetual awe.

## 5.1 Process models in AI

Generally connected hardware has similarities with the general process models used in A.I. that is Conniver and more recently Actors, and in operating system design. The general process model is the main paradigm in computer system design and in A.I. Presumably as the concept of process is clarified it will be differentiated into different types of concept, likewise data scopes, contexts and connection schemes.

Process models in A.I. arose not from real time performance requirements but from the flexible organization of knowledge so that knowledge could be brought to bear at the correct places, so that knowledge was divided into meaningful units and so that knowledge was active or procedural.

Recently there has been influence of parallel architecture on A.I., coming through vision and speech work, see Hanson and Riseman's book (78a), which includes work by Hanson and Riseman (78b) and Uhr (78), and also work on the Hearsay speech project Lesser and Erman (79).

## 5.2 A class of architectures

The class of computational models I am currently interested in is a class of array processors, which I call generalised array machines, it is defined more fully in Bond (80). Basically one can define (infinite) two dimensional arrays of identical processes.

The elements of the array are on the usual 4 neighbourhood NSEW grid and a process has I/O channels to its neighbours. An alphabet of communication characters is defined for each such process array. One can define any number of process arrays, within an array the processes are identical but different arrays can have different processes. One can now define connections between arrays by I/O channels, the connections can be of

three types only 1:1, all:one and one:all (Figure 19).

The latter is a broadcast channel and an array can also broadcast to itself. The second of the three connection types is a competitive retrieval i.e. a process output with an associated weight and the largest or random choice among the equal largest is communicated. The processes are represented in a fairly conventional program form rather than using state transition tables or graphs, however the intention is that they are finite register machines. One uses an Algol type of language with variables whose values are limited to finite ranges.

We define processes in two stages, corresponding to hardware and software. In the first, all processes are synchronous and a set of hardware functions or subprocesses is defined. In the second, these hardware processes can be further combined asynchronously using wait, delay and other synchronisation scheduling primitives, which are defined in terms of the synchronous hardware processes. A number of other derived language constructions have also been added. The idea of the synchronous hardware level is to enable timing measures to be defined, allowing the real time response to be calculated. We have programmed a number of robot control functions such as simple vision and problem solving mechanisms in this system. Most single mechanisms can be done on one process array with a broadcast channel.

This type of programming leads to new programming constructions

- (i) active images - an image is used for spatial representation, and is a process array.
- (ii) activity or attention masks, these are used to define access sets, retrieval sets and processing sets i.e. to divide array temporarily into active or inactive zones, or accessible or inaccessible zones etc.

### 5.3 Sources for our array design

There are a number of influences that have lead me to array processors. The general class of such architectures is described comprehensively in the recent book by Thurber (76) however the corresponding software is a relatively unexplored area. First there is the work on cellular computers, which have no broadcast channel, only neighbour connections. The classic results showed that such machines are very general; they are computation universal and also allow one to specify questions about patterns and their propagation and about the constructibility of machines. The formal work in these directions in the sixties by Wagner (64), Arbib (69) and A.R. Smith III (69) is summarised in A.R. Smith's thesis (69). An attractive feature of such machines is their provable real time properties; the real time recognition of context free languages, generation of primes, addition and multiplication etc. Beyer's work (69) also showed that a lot (in fact he conjectures all) of topological properties of patterns can be computed in a linear function of the time taken for a signal to cross the pattern and return.

Perhaps not surprisingly a number of speed up results are known, by increasing the numbers of states in the automaton. We are really interested in high level descriptions of the component automata, so we must translate such results into increased numbers of variables and control points. There is an interesting technique used in cellular machine work (Beyer (69)), this is layers or laminae of subautomata so that the resulting machine is the cross product of its subautomata. This introduces a useful way of specifying several different functions going on in parallel at every point within the cellular array.

The second influence is that the array discipline can be used at the logical design level, the neuronal level, the architectural level and the high level language level. The relationship between the various levels has not been clarified. A lot of results on the simulation of one array by another either of different neighbourhood connectivity or different di-



mensionality have been given by Smith (69) (and by Wagner (64) in the Mealy case), however we are interested in the correspondence between a high level array description and an equivalent low level logical design. General constructions have been given for a cellular automaton which simulates a given Turing machine and for a modular logical design for a given arbitrary switching function, (see Minnick et al. (66) for example). We are interested in high level language descriptions, where each cell is a process, and in the corresponding micro- implementation, see Figure 20.

Third, a currently reasonable hypothesis about the physiology of the brain is that it is uniformly structured (Szentagothai (72)), it is made up of two dimensional surfaces of identical modules or "chips". Together with our second point above this means that if we work with high level array descriptions of mental functions, we are in the right ball park-performance and description wise- for psychological and physiological theories. Arbib (72) has pursued a uniform layered model of the brain and together with coworkers has produced a few models of brain processes, described at the level of arrays of neurones (Dev (75), Boylls and Arbib (75), Didday (70)).

As mentioned in a previous section, cellular automata have also proved useful in biological modelling in morphogenesis and in cell reproduction. This may in the long haul provide the biological substrate for our brain models and even a unification of biological, physiological and behavioural description. For a survey of this work see Rosen (72) and Waddington (72).

The other main source is the work on associative and array processing hardware (Thurber (76)). In this the set of identical processing elements with or without neighbourhood connections receives broadcast instructions or messages. This is the more practical area and a number of practical programs have been written for ILIAC IV, STARAN, CLIP, DAP etc. and a number of techniques and ideas developed for the utilisation of such architectures.

Algorithms have been given for graph searching, template matching, radar tracking, sorting, matrix calculations and differential equation solution, for a review see Parhami (73) for example. In my own laboratory, Marks (79) has written algorithms for region finding and polyhedral scene analysis for DAP. Data structures for sorting and matrix work have been described by Stone (71) and the great simplification in graphical data structures shown by Stillman et al (71). To such work should be added low level image processing, for example on the CLIP4 architecture Duff (78), and work on digital filtering and signal processing on digital signal computers, see Salazar (77).

## 5.4 The intelligent memory view

Array processors can be viewed in a number of ways, but one way is as generalisations of associative memory. In the latter, only a limited amount of processing is possible in the store whereas in the former any general program can be executed in store; in each and in parallel.

This view I call the **"intelligent memory" view** and it is this image that differentiates our model imaginatively from other process models. Processes are used like memory, they are used in large numbers with identical processes, all work is done by constantly remembering data, control is achieved by attention effects in memory.

Of course in these machines we still have totally synchronised uniform parallelism, each processor not only executes the same program, it must be at the same step in that program, as all the others as well. There is of course an activity bit in each processor which inhibits and facilitates its execution of the step, allowing branching programs to be written.

Our generalised array processor can be mapped fairly naturally onto an existing array processor like the ICL DAP.

Returning to the question of the definition of a performance measure, we would try to use a "conceptual hardware" approach. So the real time response measure depends upon not simply the number of memory accesses but on the movement of data into contiguity. In the purely cellular case, this is the main measure i.e. the time for a signal to cross the array. In the case of a single broadcast array, we can broadcast one data item to all the array in one step and we can extract one data item from the array in one step. When we have several 1:1 interconnected arrays, one data item at a time can be transmitted down each individual connection. The transmission rate need not be the same as the memory access time, but since the data must be accessed before and stored after transmission, the total transmission time includes two memory accesses, and could possibly be simply taken equal to two.

## 6 Interdependencies

We give a figure (Figure 21) indicating the strong interdependencies between our four policies and we give in Table I what each policy tells us about the other three.

### The dependencies among the four bases

Table x/y: what does y tell us about x?

**Behaviour Patterns** Performance- The notion of behaviour pattern particularly in environments with autonomous processes depends upon the speed of response.

Development - Behaviour Patterns are seen as part of a learning sequence and simple patterns gain a justified significance as essential precursors of more complex behaviours.

Parallelism - The theory of behaviour patterns will use a parallel model as its model of the behaving system.

**Performance** Behaviour Patterns - Our notion of performance uses sensorymotor behaviour patterns as its focus.

Development - We are only interested in those systems which could have developed.

Parallelism - We use parallel models for the systems studied, array models have some known good real time properties.

**Development** Behaviour Patterns - We look at sensory motor development initially and semantic development afterwards.

Performance - The developing system must perform adequately at each step of development (not just the first and last).

Parallelism - We may possibly keep the hardware constant and develop the software,

the parallel model avoids some exponential blow up problems in development and maintains real time computability.

**Parallelism** Behaviour Patterns - We focus on computations involved in sensory motor intelligence.

Performance - We design the system so as to have known real time response, if possible.

Development - We are only interested in systems that could have developed.

## 7 Summary of goals

I have suggested an approach to A.I. research, divided it into four interdependent components and within each component made a number of suggestions for research problems and goals. An attempt at exhaustive listing of such goals would probably not terminate, but in any case it is hoped that the reader will form his own goals.

To roughly summarise a way forward

- (i) the definition of intelligent functions as components of robot control.
- (ii) the design and implementation of control systems for robots.
- (iii) the study of learning and training mechanisms and representations.
- (iv) the study of parallel algorithms for intelligent functions.
- (v) the design of programming language constructions for highly parallel machines.

## 8 Possible criticisms

### **Is what I am saying obvious?**

It may take a genius to point out the obvious. Being obvious may mean that everybody knew what I am saying before they read the paper. It may also mean that everyone knew it implicitly and that it now has an explicit form. In the latter case at least I shall have done the difficult job of expression. However, I think many indeed most AI workers will actually disagree with me. Working with real robots is generally thought to be impractical. A lot of people do not regard sensorimotor intelligence as necessarily a basis for other forms of intelligence. My answer to this is that I do not know at all clearly the relationship of sensorimotor intelligence to other forms, in what sense it may or may not form a basis, but it is a very rich area that can provide a basis for intelligent functions and whose information processing representations can at least be studied. Workers opting to study higher forms of intelligence I think will have a very hard time finding suitable representations, without sensorimotor ones.

You only have to look at a few AI conferences and journals to see that most AI workers follow a completely different mainstream way. They want to develop representations based on formal logic or conceptual nets etc.

### **Is it contained in what others have said?**

I would not be writing this paper if there were other such treatments of these ideas. I would like to see more such treatments. Brain modellers such as Arbib (72), Sutro and Kilmer (69) and Albus (79) have written perhaps the closest treatments, but I differ from them considerably. I am software oriented, I'd like to work with high level languages for array machines not neuronal models. I advocate the use of real robots, and my main examples are in human sensory motor intelligence.

**Is it impractical?**

No, we're doing it in the AI laboratory at Queen Mary College, albeit in a modest way. It is possible to build a robot that will work for not much money, but it takes a long time and quite a lot of aggravation. Parallel machines are now on the market, we have access to a 64 x 64 ICL DAP array processor, although only as a batch service, but in a couple of years cheap useable array processors will be bought by individual AI labs.

**Is it inconsistent?**

No.

**Is the type of parallelism advocated too limited?**

Array processors in practice provide a very rich medium for representation and I believe they provide a welcome discipline. However it is quite possible that new forms of parallelism will be required when we have some clearer insights.

**Is my approach too vague?**

It is rather vague and confused, but it is the best I can do. Having written it, one can then proceed to elaborate points of it into clearer forms. We are moving forward on several fronts and expect to be able to produce elaborations of various parts into more specific studies and models.



## 9 Summary of the paper

In spite of a superfluity of fragility, this paper, which has not been produced easily, has I think made the following points.

### The Introduction

I introduced the need for a coherent conscious **approach** to artificial intelligence.

I discussed an inductive **scientific methodology** and argued for **mathematical theorisability**.

I said we should take our relationship to **psychology** more seriously.

### Behaviour Patterns

I talked about complete **behaviour patterns**.

The use of **real robots** providing a rich environment of problems.

The need to deal with **probabilistic** data.

I emphasized the notion of **environment** and described some.

A.I. as the design of **computer systems for robotics**, the data rates, data base and real time response requirements.

The **hierarchical factorization** of the system.

I noted the new issues of the need for a **motivational theory** and the problem of overall **system stability**.

I gave a list of **robot functions** in vision, databases and natural language, which lead to new research lines.

I gave diagrams of a total robot system based on **state of the art** mainstream A.I.,

cognitive psychology and neurosciences.

I distinguished between **functional models** and **implementation models**.

## **Performance**

I didn't say much here, but tried to discuss:-

A notion of **performance** rather than function.

The difficulty of comparing the performance of A.I. programs.

The connection with psychology.

A performance measure based on **memory access time** in conceptual hardware.

**Hierarchy** in real time response.

The notion of a **performance model**.

## **Development**

The **developmental constraint**.

A definition of **hardware and software**.

Well formedness of knowledge.

**Developmental sequences**, some examples.

**Stability and attention**.

Development of mathematics not a suitable area, human development better.

The **representational regression** to

The study of **infancy**.

Remarks on the brain, evolution and maturation.

Theoretical questions about development.

## **Parallelism**

The insistence on a **highly parallel model of computation**.

Description of a **generalised array machine**.

Two level specification into **synchronous and asynchronous** levels.

Some **parallel programming constructs**.

The **sources** of influence are cellular automata, use of array processor descriptions at several levels which could **map** into each other, the uniform anatomical structure of the mammalian brain and hence the possibility of relevance of performance measures to psychology and physiology, associative processors and existing highly parallel computers.

The **"intelligent memory" view**.

Performance measures for array architectures.

## **Interdependencies**

I noted and tried to elaborate the **strong interdependency** amongst the four policies.

## **Summary of goals**

I made brief remarks about the research goals suggested by the approach and a possible **way forward**.

## **Possible criticisms**

I listed possible criticisms that might be levelled at what I have said and made some initial **ripostes**.

## Coda

Having presented a rather wholesale criticism of "AI", I hope this paper will meet the approval of the referee. It is offered in a friendly and creative spirit. The author spent some formative years at Oxford and believes in a community of scholars, communicating in patient, self effacing honesty. (Except oneself of course in whom the ego rides nightly rabid and terrible).

"Something about 'purity of bodily fluids' - we're still working on this one sir....."

from Dr. Strangelove.

## References

- Albus, J., "A model of the brain for robot control", Byte 4 (1979)
- Aleksander, I., "Pattern recognition with networks of memory elements", in Batchelor, B.G. (ed), "Pattern recognition ideas in practice", Plenum, New York (1977)
- Andrae, J.H., "Thinking with the teachable machine", Academic Press, New York (1978)
- Arbib, M.A., "Automata theory in the context of theoretical embryology", in Rosen, R. (ed), "Foundations of mathematical biology", 3, Academic Press, New York (1972) 193-282
- Arbib, M.A., "Theories of abstract automata", Prentice-Hall, New Jersey (1969)
- Arbib, M.A., "The metaphorical brain", John Wiley, New York (1972)
- Baer, R.M. and Martinez H.M., "Automata and biology", Ann.Rev.Biophys. and Bioeng. 3 (1974) 255-291
- Balinski, B.I., "An introduction to embryology", 4th edition, W.B.Saunders, Philadelphia (1975)
- Barnden, J.A., "An approach to semantic net theory", Ph.D. thesis, University of Oxford (1975)
- Bartlett, F., "Thinking", Cambridge University Press (1958)
- Becker, J.D., "An information-processing model of intermediate-level cognition", Ph.D. thesis, Stanford University (1973)
- Bernstein, N.V., "The coordination and regulation of movements", Pergamon, Oxford (1966)
- Beyer, W.T., "Recognition of topological invariants by iterative arrays", Ph.D. thesis

M.I.T. (1969)

Bond, A.H., "A highly parallel language for the real time control of robots", Technical Report (1980)

Boylls, C.C., "A theory of cerebellar function with applications to locomotion", COINS Technical Report 75C-6 (1975) University of Amherst, Massachusetts

Broadbent, D.E., "Decision and stress", Academic Press, London (1971)

Bronson, G., "The hierarchical organisation of the central nervous system: implication for learning processes and critical periods in early development", Beh.Sci. 10 (1965) 7-25

Bruner, J., "Learning how to do things with words", in Bruner and Garton 62-84 (1978)

Bruner, J. and Garton, A., "Human growth and development", Oxford University Press (1978)

Citrenbaum, R.L., "Strategic pattern generation: a solution technique for a class of games", Pattern Recognition, 4 (1972) 317-329

Clark, E.V., "From gesture to word: on the natural history of deixis in language acquisition", in Bruner and Garton 85-120 (1978)

Clark, H.H. and Clark, E.V., "Psychology and language", Harcourt Brace and Jovanovich, New York (1977)

Cunningham, M.A. and Gray, H.J., "Design and test of a cognitive model", Int.J.Man-Machine Studies 6 (1974) 49-107

Dev, P., "Perception of depth surfaces in random dot stereograms: a neural model", Int.J.Man-Machine Studies 7 (1975) 511-528

Didday, R.L., "The simulation and modelling of distributed information processing in the frog visual system", Ph.D. thesis, Stanford University (1970)

Doran, J.E., "Experiments with a pleasure-seeking automaton" in Michie, D., "Machine Intelligence", 3 (1968) 195-216

Duff, M.J.B., "Review of the CLIP image processing system", National Computer Conference (1978) 1055-1060

Ernst, G.W., "Sufficient conditions for the success of GPS", JACM 16 (1969) 517-533

Ernst, H.A., "Computer controlled robots", IBM Research Report RC2781 (1970)

Fahlman, S., "A planning system for robot construction tasks", Artificial Intelligence 5 (1974) 1-49

Gibson, J.J., "The senses considered as perceptual systems", Houghton Mifflin, Boston (1966)

Greene, P.H., "Problems of organisation of motor systems", in Rosen, R. and Snell, F.M. (eds), "Progress in theoretical biology", 2, Academic Press, New York (1972) 303-338

Halliday, M.A.K., "Explorations in the functions of language", Edward Arnold, London (1973)

Halliday, M.A.K., "Learning how to mean", Edward Arnold, London (1975)

Hanson, A.R. and Riseman, E.M., "Visions: a computer system for interpreting scenes", in Hanson, A.R. and Riseman, E.M. (1978) 303-334

Hanson, A.R. and Riseman, E.M., "Computer vision systems", Academic Press, New York (1978)

Hendrix, G.G., "Modelling simultaneous actions and continuous processes", Artificial Intelligence 4 (1973) 145-180

Hewitt, C., "Description and theoretical analysis (using schemata) of planner", Ph.D. Thesis, M.I.T., AI-TR-258 (1972)

Hewitt, C., "Viewing control structures as patterns of passing messages", Artificial Intelligence 8 (1977) 323-364

Hoare, C.A.R., "Communicating sequential processes", CACM 21 (1978) 666-677

Hunt, E.B., "Artificial Intelligence", Academic Press, New York (1975)

Jones, T.L., "A computer model of simple forms of learning", Ph.D.Thesis, M.I.T. (1971)

Kowalski, R., "Logic for problem solving", Elsevier, New York (1979)

Kohout, L.J., "Representation of functional hierarchies of movement in the brain", Int.J.Man-Machine Studies 8 (1976) 699-709

Knuth, D.E., "An analysis of alpha-beta pruning", Technical report STAN-CS-74-441, Stanford University (1974)

Lesser, V.R. and Erman, L.D., "An experiment in distributed interpretation", Technical Report CMU-CS-79-120, Carnegie-Mellon University (1979)

Lowrance, J.D. and Friedman, D.P., "Hendrix's model", Int.J.Man-Machine Studies 9 (1977) 537

Lindenmayer, A. and Rosenberg, G. (eds), "Automata, languages and development", North Holland, Amsterdam (1976)

Marks, P.C., "Low level vision using an array processor", Computer Graphics and Image Processing, 14 (1980) 281-292

Marr, D., "Early processing of visual information", AI memo AI-M-340, M.I.T. (1975)

McCarthy, J. and Hayes, P., "Some philosophical problems from the standpoint of artificial intelligence", in Meltzer, B. and Michie, D., "Machine Intelligence" 4 (1968) 463-502

Metzler, J.(ed), "Systems neuroscience", Academic Press, New York (1977)



Minnick, R.C. et al, "Cellular arrays for logic and storage", Technical report AFCRL-66-613, Stanford Research Institute (1966)

Minsky, M., "Plain talk about neurodevelopmental epistemology", AI Report 430, M.I.T (1977)

Minsky, M. and Papert, S., "Artificial intelligence progress report", AI Report 252, M.I.T. (1972)

Moses, J., "Symbolic integration", Ph.D. Thesis, MAC-TR-47, M.I.T. (1967)

Mott, D.H. and Bond, A.H., "The investigation of learning in a mobile robot", Technical Report, Queen Mary College (1978)

Munn, N.L., "The evolution and growth of human behaviour", Harrap, New York (1965)

Murtha, J.C., "Highly parallel information processing systems", Advances in Computers 7 (1966) 1-116

Newell, A. and Simon, H.A., "Human problem solving", Prentice Hall, Englewood Cliffs, New Jersey (1972)

Nilsson, N.J., "Searching problem-solving and game-playing trees for minimal cost solutions", in Morrell, A.J.H. (ed) "Information processing 68", North Holland, Amsterdam (1969)

Nilsson, N.J., "Learning Machines", McGraw-Hill, New York (1965)

Nilsson, N.J., "Artificial Intelligence" in Rosenfeld, J.L. (ed), "Information Processing 1974", North Holland, Amsterdam (1974) 778-801

Ohlander, R., "Analysis of natural images", Ph.D. Thesis, Carnegie-Mellon University (1975)

Olds, J., "Phylogeny of behaviour", in "Brain mechanisms and the control of behaviour"

(1975) 253-271

Parhami, B. "Associative memories and processors: an overview and selected bibliography", Proc. IEEE 61 (1973) 722-730

Piaget, J., "The origin of intelligence in the child", Routledge, London (1953)

Pohl, I., "First results on the effect of error in heuristic search", in Meltzer, B. and Michie, D. (eds), "Machine Intelligence", 5 (1970) Edinburgh University Press 219-236

Reboh, R. et al, "A study of automatic theorem-proving programs" Technical Note 75, Stanford Research Institute (1972)

Rosen, R. (ed), "Foundations of mathematical biology", Academic Press, New York (1972)

Rychener, M.D., "Production systems as a programming language for artificial intelligence applications" Ph.D. Thesis, Carnegie-Mellon University (1976)

Scott, P.D., "The cerebral neocortex as a programmable machine", Ph.D. Thesis, University of Sussex (1975)

Slagle, J. and Dixon, J., "Experiments with some programs that search game trees", JACM 16 (1969) 189-207

Smith, A.R., "Cellular automata theory" Technical Report SEL-70-016, Stanford University (1969)

Salazar, A.C., "Digital signal computers and processors", IEEE Press, New York (1977)

Stillman, N.J., et al, "Associative processing of line drawings", Spring Joint Computer Conf (1971) 557-567

Stone, H.S., "Parallel processing with the perfect shuffle", IEEE Trans. Computers C20 (1971) 153-161

Sutro, I.L. and Kilmer, W.L., "Assembly of computers to command and control a robot", Spring Joint Computer Conf. 34 (1969) 113-137

Sussman, G.J., "A computational model of skill acquisition" Ph.D.Thesis, M.I.T. (1973)

Szentagothai, J., "The basic neuronal circuit of the neocortex", in Petsche, H. and Brazier, M.A.B. (eds), "Synchronisation of EEG activity in epilepsies", Springer-Verlag, Vienna (1972) 1-45

Tenenbaum, J.M. and Barrow, H.G., "Experiments in interpretation guided segmentation", Technical Note 123, Stanford Research Institute (1976)

Thurber, K.J., "Large scale computer architecture", Hayden, New Jersey (1976)

Tsetlin, M.L., "Automaton theory and modelling of biological systems", Academic Press, New York (1973)

Tzichritzis, D. "Measures on computable sets", Technical Report 8, University of Toronto (1969)

Uhr, L., "Pattern recognition, learning and thought", Prentice Hall, New Jersey (1973)

Uhr, L., "A model of form perception and scene description", Technical Report 231, University of Wisconsin (1974)

Uhr, L., "'Recognition cones' and some test results: the imminent arrival of well structured parallel-serial computers; positions and positions on positions", in Hanson, A.R. and Riseman, E.M. (1978) 363-378

Uhr, L. and Kochen, M., "Mikrokosm and robots", Int. Joint Conf. Artificial Intelligence (1969) 541-555

Waddington, C.H., "Towards a theoretical biology: four essays", Edinburgh University Press (1972)

Wagner, E.G., "An approach to modular computers", IBM Research Reports RC1107 and RC1414 (1964)

Waterman, D.A., "Generalisation learning techniques for automating the learning of heuristics, Artificial Intelligence 1 (1970) 121-170

Waltz, D.L., "Generating semantic descriptions from drawings of scenes with shadows", Ph.D.Thesis AI-TR-271, M.I.T. (1972)

Winograd, T., "Understanding natural language", Academic Press, New York (1972)

Winston, P.H., "Learning by creating and justifying transfer frames", Artificial Intelligence 10 (1978) 147-172

Winston, P.H., "Learning structured descriptions from examples", Ph.D.Thesis, M.I.T. (1970)

Witten, I.P., "Learning to control sequential and non-sequential environments", Technical Report EES-MMS-CON 75, University of Essex (1975)

Young, R.M., "Seriation by children: a production system approach", Birkhauser-Verlag, Basel (1976a)

Young, R.M., "Design choices for a world modelling system", AISB Conf. (1976) 376-386




Figure 1: Robot

Use of nonstandard hypotheses	Complex planning & cogitation	Complex speech Puns Poetry	3
Object Identification	Combination of learnt units	Simple Deixis	2
Simple vision orientation etc.	Reflexes	Gesticulation Basic Communicative Functions	1

Figure 2: Horizontal and vertical factorisation




Figure 3: Robot in state of the art environment




Figure 4: Display of Winograd's blocks world






Figure 5: Shakey in its real blocks world




Figure 6: Outdoor scene of Ohlander



Figure 7: Office environment of Tenenbaum





Figure 8: Factory floor environment




Figure 9: Playpen environment





Figure 11: Caretaker environment




Figure 12: Environment of early man






Figure 13: Access restrictions on a possible flexible hierarchy

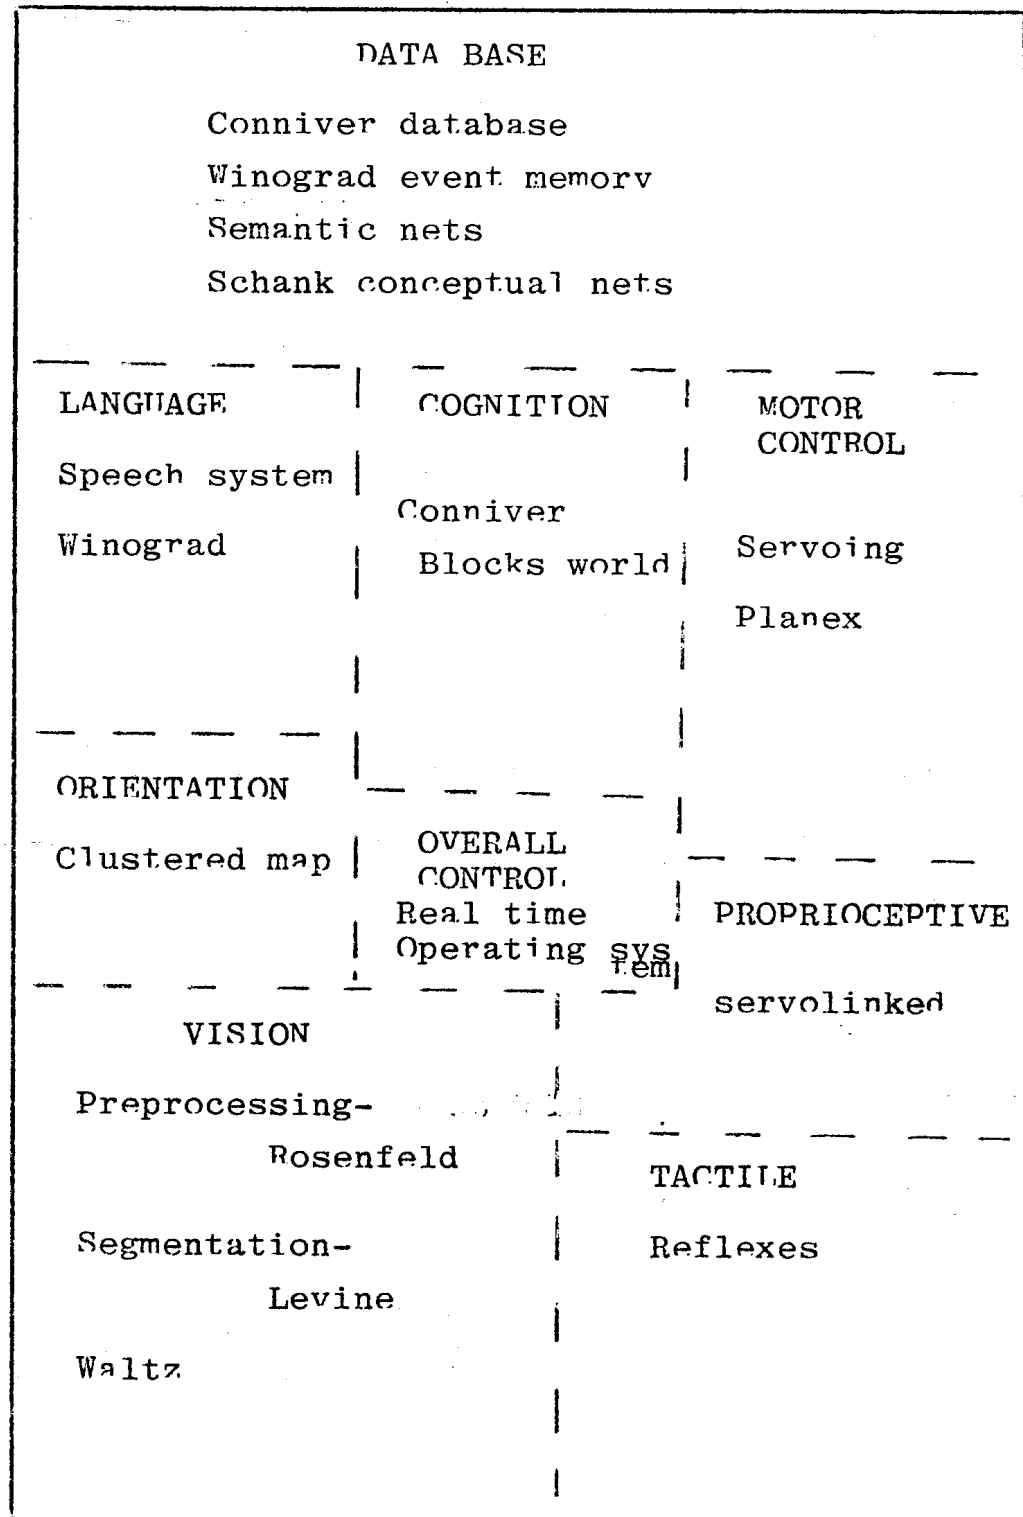


Figure 14: Functional model based on AI

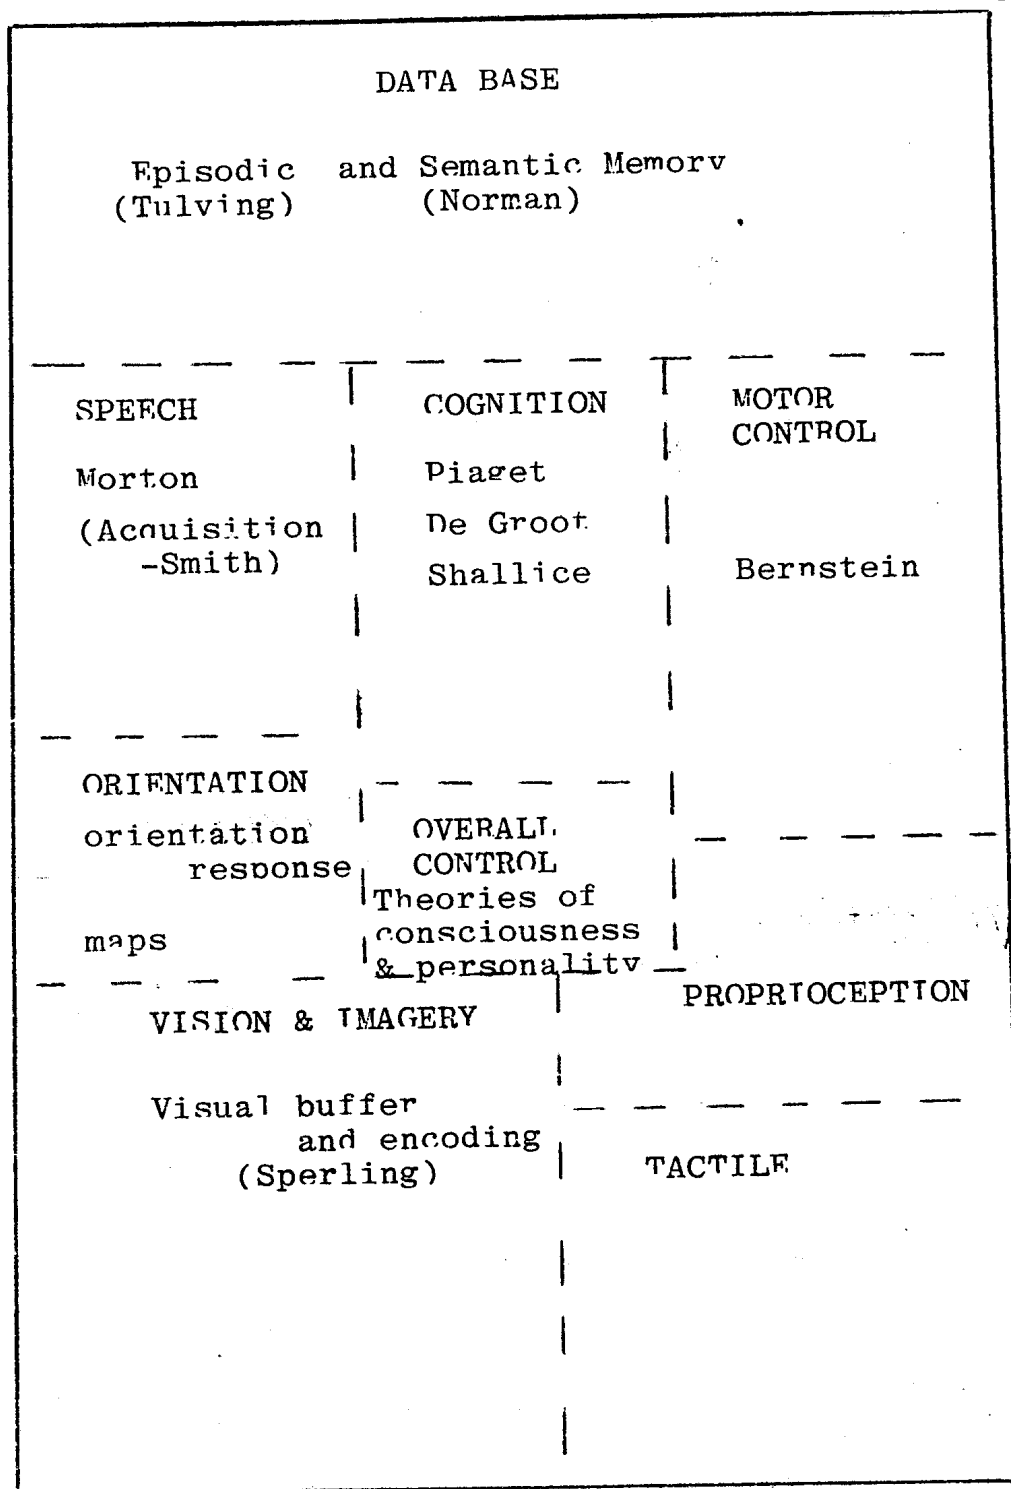


Figure 15: Functional model based on cognitive psychology

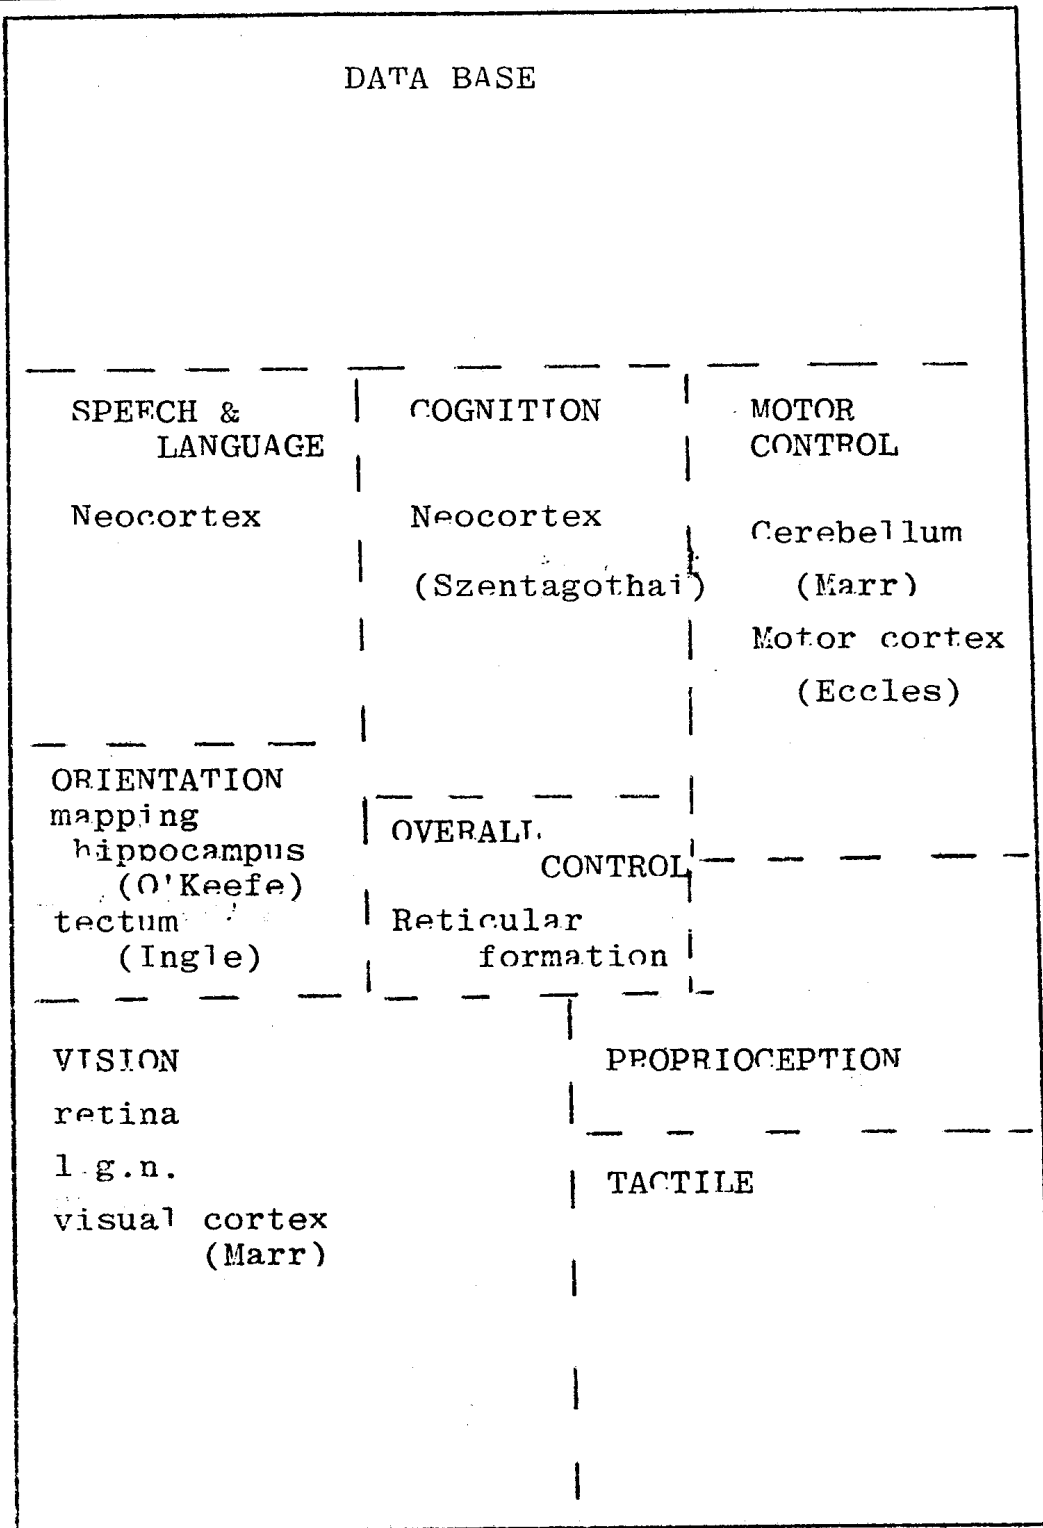


Figure 16: Implementation model based on neurosciences






Figure 17: Hierarchies of computers



Shirley's developmental sequences. Adapted from "The First Two Years, a Study of Twenty-Five Babies, I, Intellectual Development," by M. M. Shirley. In Institute of Child Welfare Monographs, Series No. 8, 1933. Reprinted by permission of University of Minnesota Press.

Figure 18: Development of posture and walking




Figure 19: Interconnection schemes for process arrays




Figure 20: Mapping arrays between levels of detail



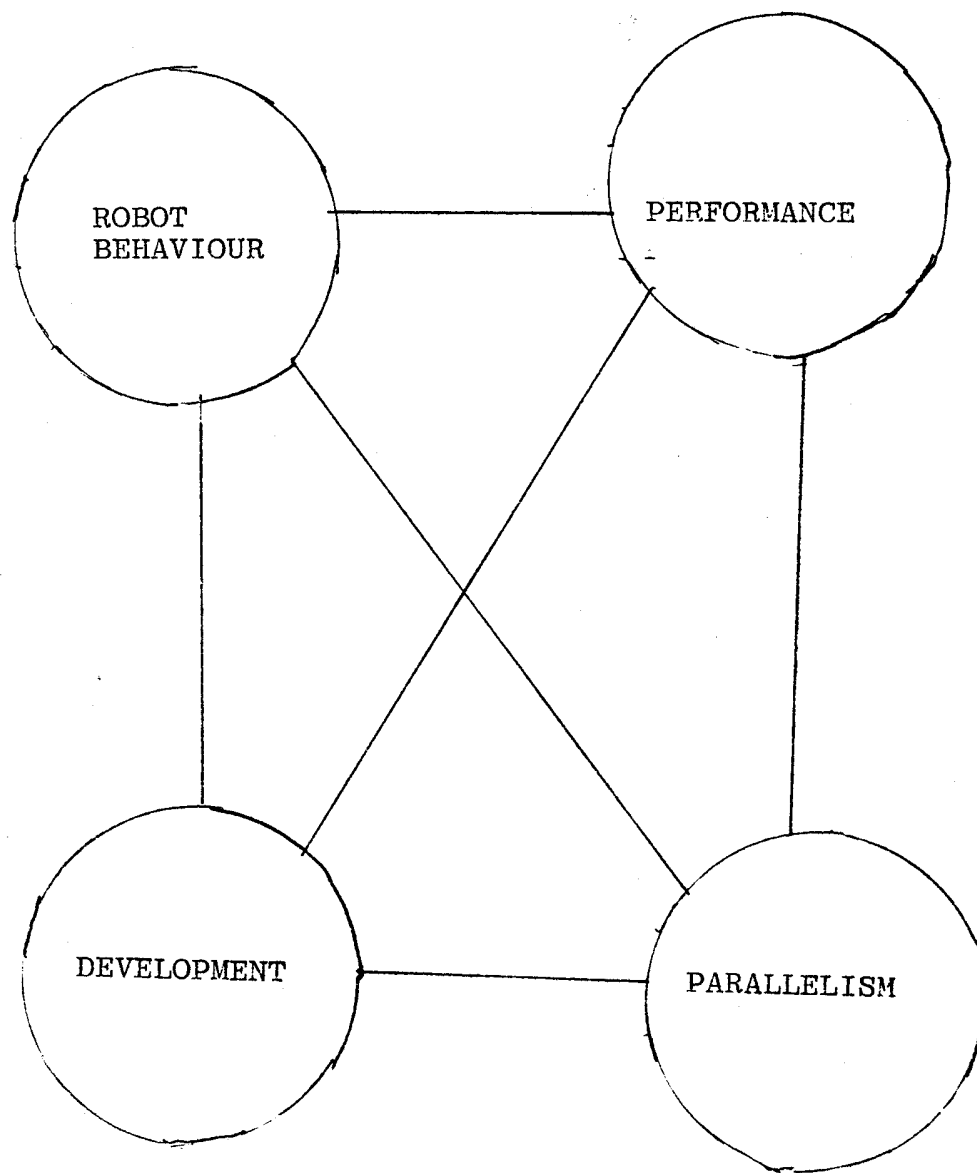


Figure 21: Interdependencies between the four policies of my approach

# Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
1.1	Basic Research Criteria . . . . .	5
1.2	Levels of work . . . . .	10
1.3	Psychology . . . . .	13
1.4	Applications . . . . .	14
1.5	Summary of my approach . . . . .	15
<b>2</b>	<b>Robot behavior patterns</b>	<b>17</b>
2.1	Real robots . . . . .	17
2.2	Total behavior . . . . .	17
2.3	Environments . . . . .	19
2.4	Taxonomy of environments . . . . .	20
2.5	Hierarchy of function . . . . .	21
2.6	Fuzziness . . . . .	21
2.7	Continuous functions . . . . .	22
2.8	Motivation . . . . .	22
2.9	Stability of behavior . . . . .	22
2.10	Natural decomposition of intelligent robot processes . . . . .	22
2.11	Functional models and implementation models . . . . .	24
2.12	Computer system design . . . . .	24
<b>3</b>	<b>Performance</b>	<b>27</b>
3.1	A notion of performance . . . . .	27
3.2	Performance and "normal" AI . . . . .	28
3.3	Real time . . . . .	30
3.4	Hardware and the unit of time cost . . . . .	31
3.5	Hierarchies of real time response and the theory of iteration . . . . .	31
3.6	Performance models . . . . .	32

<b>4</b>	<b>Development</b>	<b>34</b>
4.1	The developmental constraint . . . . .	34
4.2	Hierarchy of control and development . . . . .	36
4.3	Learning Sequences . . . . .	36
4.4	Stability and Attention Criteria . . . . .	37
4.5	The Development of Mathematics . . . . .	39
4.6	The Representational Regression . . . . .	40
4.7	Infancy . . . . .	41
4.8	The Brain and Evolution . . . . .	42
4.9	Growth . . . . .	42
4.10	Sensorimotor learning in the computer literature . . . . .	43
4.11	Theoretical treatment of learning . . . . .	44
<b>5</b>	<b>Parallelism</b>	<b>45</b>
5.1	Process models in AI . . . . .	46
5.2	A class of architectures . . . . .	46
5.3	Sources for our array design . . . . .	48
5.4	The intelligent memory view . . . . .	50
<b>6</b>	<b>Interdependencies</b>	<b>52</b>
<b>7</b>	<b>Summary of goals</b>	<b>54</b>
<b>8</b>	<b>Possible criticisms</b>	<b>55</b>
<b>9</b>	<b>Summary of the paper</b>	<b>57</b>

## List of Figures

1	Robot . . . . .	69
2	Horizontal and vertical factorisation . . . . .	70
3	Robot in state of the art environment . . . . .	71
4	Display of Winograd's blocks world . . . . .	72
5	Shakey in its real blocks world . . . . .	73
6	Outdoor scene of Ohlander . . . . .	74
7	Office environment of Tenenbaum . . . . .	75
8	Factory floor environment . . . . .	76
9	Playpen environment . . . . .	77
10	Breast environment . . . . .	78
11	Caretaker environment . . . . .	79
12	Environment of early man . . . . .	80
13	Access restrictions on a possible flexible hierarchy . . . . .	81
14	Functional model based on AI . . . . .	82
15	Functional model based on cognitive psychology . . . . .	83
16	Implementation model based on neurosciences . . . . .	84
17	Hierarchies of computers . . . . .	85
18	Development of posture and walking . . . . .	86
19	Interconnection schemes for process arrays . . . . .	87
20	Mapping arrays between levels of detail . . . . .	88
21	Interdependencies between the four policies of my approach . . . . .	89