Multiple Representational Systems - please do not circulate

Satyajit Rao

October 30, 1995

Abstract

Human cognition is distinguished in the *variety* of capabilities that it supports like playing chess, doing arithmetic, planning a trip, learning to ski, or understanding how a device works.*It is the range of these representational and inference capabilities, and the fact that we can do them at all, rather than our questionable competence at doing any particular one that is impressive*.

This view of the essence of human cognition, when taken seriously, suggests a fundamental shift in our efforts to make machines intelligent. Instead of trying to build specialized representations and processes for each task domain (eg modeling the qualitative behavior of liquids, or proving mathematical theorems), we should focus on developing *a Cognitive Alphabet* - a set of representations and processes that can support a wide variety of Cognitive tasks. It suggests that if we are to build human like robots then we should concentrate on the problem of building a *representational platform* that can support a wide variety of cognitive tasks. What shall we choose as the basis set? It seems very likely that humans do not have specialized neural hardware for playing chess, doing arithmetic, or planning a trip, but rather that "cognitive structures for abstract thought are based on earlier evolved structures for objects, locations, motions...". In other words the representations and inference processes involved in perception, action, language, and social intelligence, also form the basis for abstract cognitive tasks like understanding how a toaster works. What implications does this idea have for building a humanoid? and what technical challenges will we face? The main message in this paper is that our focus should be on developing the inference processes of Multiple Representational Systems, and on the interactions between them.

1 What do we want?

Building a robot with human like capabilities has been the holy grail of Artificial Intelligence. After more than 40 years of trying to do this we are still very far from that goal. There is a large gap between the cognitive capabilities of the ant like robots that we can build today and the human like robots that we would like to build. This is paper is an attempt to come up with a program to bridge this gap.

The first issue is that we need to decide on what we really want when we say "human like intelligence" or "cognition". Implicit notions about the capabilities that seem to lie at the crux of human intelligence have directed (and often misdirected) entire research approaches. In this section we briefly review the motivations behind two approaches to achieving cognition, the problems they ran into, and then describe a key problem that we need to solve.

The Symbol Systems Approach

The "Symbol Systems Approach" of classical Artificial Intelligence is based on the separation of reasoning from perception and action. It assumes that knowledge and the manipulation of knowledge, can and should be separated from the particular details of the physical body (the sensors and effectors). The framework in which traditional Artificial Intelligence works is that Perception delivers a symbolic [1] description of the world in terms of a fixed set of predicates like ON(A, B), or LIKES(X, Y). The "Reasoning module" usually contains hand-crafted knowledge about the world expressed in some Knowledge Representation Language (eg if-then rules, frames, scripts, declarative statements in some form of mathematical logic). This knowledge, and the world description, are used by a problem solving mechanism to achieve some goal. The solution is passed to an execution module that performs actual actions in the world to achieve the goal. This is the basic framework.

This approach to achieving cognition arose from a conjunction of two factors - a naive view of human intelligence, and a bias of what the available technology (digital computers and programming languages like LISP) was good at doing - symbol manipulation. The influence of our models of computers on our models of thought, or how the problem was defined to fit the tools is well described by Brooks (1991*a*), so I will briefly describe the other major influence, namely what the early researchers wanted when they said "human like intelligence". Human intelligence was characterized by the ability to solve certain kinds of "difficult problems" - like solving a problem in formal logic (Newell & Simon 1961), proving mathematical theorems, discovering scientific laws from data (Newell, Shaw & Simon 1961), and multicolumn subtraction

[1] There are many different notions of "symbol". In this particular case "symbol" is used in a very narrow sense similar to the tokens and expressions of LISP

(Rosenbloom, Laird, Newell & McCarl 1991). The choice of these problems as representative of human cognition is interesting - it reflects a folk view of intelligence, most of humanity hardly ever engages in these tasks and the few who do are usually the ones regarded as "intelligent", hence (the reasoning went) investigate these "high order intellectual processes" and you have the key to human intelligence. The tool for investigating these "high order intellectual processes" was introspection. For example a student would be given a problem in formal logic and asked to speak his/her thoughts aloud while constructing the proof. These verbalizations would be recorded and then analyzed for clues about the problem solving process (Newell & Simon 1961). After analyzing the problem solving process in the other problems too, it was found that if you strip away the particulars of the problem then you could explain problem solving in these limited class of problems as search in some space of symbols. This fact in conjunction with the bias of the available technology led to the introduction of "symbol manipulation" into a *definition* of intelligence. A physical symbol system was defined as a bunch of symbol structures (tokens, expressions) and some processes to create and modify the structures. The Physical Symbol System Hypothesis states that a physical symbol system has the *necessary* and *sufficient* means for *general intelligent action* (Newell & Simon n.d.). The main problems with the traditional approach are:

- *Building representations for particular tasks instead of representations that could support a variety of tasks.* Consequently we end up with a collection of separate programs that play chess, prove theorems, or do medical diagnosis. There is nothing wrong with this if our intention is to build tools that can assist humans, however we should not fool ourselves into thinking that the collection of these programs says anything about cognition.

- *It evades knowledge acquisition* - It does not address the basic issue of how the knowledge (eg rules of chess, or behavior of liquids) is *acquired* in the first place. The practice of isolating domains and *hand-coding* facts about them is clearly evading the issue and one of the factors responsible for the performance of these systems not scaling.

- *The abuse of Abstraction* - Brooks (1991*b*) points out that coming up with descriptions like ON(A, B) from a scene (both the choice of that particular feature, and the detection of that feature) is the essence of intelligence and the hard part of the problem. He says

  > Under the current scheme the abstraction is done by the researchers leaving little for the AI programs to do but search.

- *Weak Inference* In spite of the emphasis on problem solving, the systems exhibited very weak inference capabilities, they could hardly infer anything beyond what they were told. For example if the system is trying to build a tower of blocks by putting B on C (described as ON(B,C)) and then A on B (described as ON(A,B)) and each block happens to be skewed a little to the right because of some error in positioning. The tower collapses and the description of the history (ON(A,B), ON(B,C)) contains no information about what went wrong, it is in fact identical to descriptions of towers that did not collapse. The tendency of the knowledge engineer is to make the symbolic primitives more elaborate, for example to incorporate degrees of ON-ness in the description. This is just a temporary fix and does not address the source of the problem. The weak inference stems from two sources. The first is that only one kind of representation is being used - the symbolic representation used in the previous example can only make certain types of information explicit (a qualitative representation of ON-ness) and will break when you attempt to use it for tasks for which it it is not suited. The second source of weak inference is the artificial separation of perception from reasoning. Prematurely throwing away the perceptual representation of the image without knowing what's important in it for some task is a mistake. Certain kinds of reasoning (about why the tower collapsed) may be best performed using "perceptual" representations and processes.

- *All Knowledge may not have to be represented explicitly* - A robot can behave in a manner such that a human might attribute some knowledge and goals to it, but that doesn't mean that the robot should be built with an explicit representation of that knowledge or goals.

The Behavior Based Approach

The other approach to achieving human like intelligence is the Behavior based approach pioneered by Brooks (1991*b*)(Brooks 1991*a*), (Brooks 1990) as a reaction to the problems of the Symbol Systems Approach. The main motivation of the behavior based approach is that the ability to perceive and act evolved long before (and took much longer to evolve) than the abilities of playing chess, or proving theorems, hence it is worthwhile to investigate simple perception action loops because they might give you some insight into more complicated forms of intelligent behavior.

In the behavior based approach one attempts to build robots that perform simple tasks in complex environments. The hope is that it is easier to start with a complex environment and some simple tasks and then gradually start increasing the complexity of the task, rather than the other way around. The chief insights of behavior based A.I are:

- Embodiment and Situatedness - Investigate intelligence by working in the real world with real physical robots that directly sense the world. This prevents you from dealing with artifactual problems in simulated environments, and also exploit constraints in the real world.

- Exploiting the complexity of the environment - There is a way to design robots that can lead to complex behavior with non-centralized representations of the world. The complexity of the behavior is a product of the dynamics of the interaction of the robot with a complex environment rather than some complicated internal representation of the world. The trick to designing robots that exploit the complexity of the environment is by designing a collection of behaviors that interact with each other via the environment. The behaviors themselves are very simple independent perception action loops tightly coupled to the environment.

- The world is it's own best model - The world is sensed continuously instead of maintaining and updating some complicated internal model.

With these insights the behavior based approach has been successful in developing a Subsumption Architecture methodology for incrementally building insect like intelligence. It may also be the beginning of a satisfactory explanation of natural intelligence at the level of insects, However, The behavior based approach does not address the aspect of human intelligence that makes it qualitatively different from insect intelligence. The following is my bias about what this aspect is.

A platform that supports rich representation and inference

The most striking aspect of human intelligence is the astonishingly broad range of things that we can represent and infer, for example; having concepts like "solar system", or understanding how a bicycle pump works, or proving theorems in mathematics. One truly appreciates how peculiar it is that we can do these things when one observes that the representational capabilities of most animals are very tightly coupled to the environmental niche in which they evolved. Homo Sapiens on the other hand can represent things like "atom" or "striped elephant" - things that it may never have perceived, or may not even exist. The key point to remember is that *it is not our questionable expertise at any one of the tasks mentioned above, but that we can do them at all and that we can do so many of them that characterizes human intelligence*. Therefore, a fundamental question in Artificial Intelligence which we must answer if we are to build robots with human like intelligence is what is the source of this representational power? *How can we engineer powerful representational and inference capabilities in a robot?* It cannot be stressed enough that the idea is not to build specialized representations for individual task domains, but rather to build a representational platform that will allow the robot to learn to play chess, or answer a question about how toasters work without a human having to handcode any representations particular to game playing or qualitative physics. This, then is what we want when we say "human like intelligence".

The Problem

An important clue to the origin of representational and inference capabilities in humans lies in our evolutionary history. It seems unlikely that in the course of evolution we suddenly developed specialized neural substrates devoted to playing chess, or doing arithmetic. From an evolutionary standpoint it seems much more likely that such abilities are purely accidental byproducts of more basic previously evolved capabilities like representing space, manipulating objects, natural language syntax, and a representation of social relations (Pinker n.d.) (Bickerton 1990) - representations and processes that evolved as a response to the pressures of being embodied, and being in the company of other similarly embodied humans. The implication of this simple hypothesis is that you may be solving an integral (for example [2]) by using visuospatial operations on spatial representations of some symbols, imputing anthropomorphic roles to the components of the expression, and using inheritance properties of linguistic categories). Whether you believe this particular example or not is not the point. The point is that some such explanation of our abstract cognitive abilities in terms of representations and processes ment for more mundane activities *must* exist. Clearly, ants, frogs, and monkeys have visual, motor, and communication systems too. Yet we know that even the smartest of them cannot even begin to grasp the concept of an atom, a concept that is well within the reach of an adult human. We need a deep understanding of what is special about our representational systems and their interactions that lets us do things that other species cannot. **If we somehow understood what it is about bodily representations and processes and their interactions that gives rise to the "accidental" byproducts of being able to do calculus or think about atoms, then we would be well on our way to understanding the source of human cognition.**

The outline of the program that I intend to pursue to gain such an understanding is as follows.

A Program

- Choose some set of representational systems (say visuospatial, and syntactic). Each representational system involves making a guess about a representation, and a set of associated processes that distinguish that representational system. In other words, make explicit guesses about the machinery that allows one to support the "accidental" byproducts.
- Make a guess about the rules of interaction between the chosen representational systems.
- Show how the previous two choices:
  1. specify the dynamics of the development of cognitive structures of individual representational systems.
  2. make certain things "obvious" which were previously beyond the grasp of any individual representational system.

The following sections discuss the program in more detail.

[2] I choose this example because doing mathematics is considered to be one of the abstract cognitive feats of the human brain

2 Multiple Representational Systems

A *Representational System* is an idealization that couples a representation (for example, a topological representation of space) and a set of procedures or inference mechanism for manipulating and extracting specific information from that particular kind of representation (foe example, ways of extracting different spatial relations). That is, a Representational system is a tuple $\langle R\ G \rangle$, where R is a particular kind of representation, G is a set of procedures that act *only* on R. Representational systems are distinct because the procedures of one cannot act on the representations of another. For instance, it makes no sense to apply the rules for combining phonemes to a visuospatial representation.

The particular set of representational systems that we choose for our robot is very important because at the very outset this may bring whole classes of concepts within its powers and preclude it from ever grasping others.

There is strong evidence from lesion and split-brain studies for the autonomy/modularity of different Representational Systems in humans [3].

Visuospatial $\langle R_v\ G_v \rangle$: Where
R_v = A Topographic representations of space.
G_v = Processes for extracting different spatial relations from R_v, or synthesizing them in R_v.

Motor-Somatosensory $\langle R_m\ G_m \rangle$: Where
R_m = Motor Programs and Somatosensory representations for body posture, movement, and manipulation of various objects.
G_m = Processes for modifying and combining motor programs.

Syntactic $\langle R_s\ G_s \rangle$: Where
R_s = Words, sentences, parse trees.
G_s = Rules for parsing, and synthesis of sentences.

Thematic $\langle R_t\ G_t \rangle$: Where
R_t = Representations of social and affective relations - likes, dislikes, emotions, social relations
G_t = Rules for inferring elements of R_t from structures in R_t.

The following description of what Representational Systems do and how they interact is distilled from observations from Neuroscience, and Developmental Psychology. They should be treated like the rules describing a dynamical system, the claim is that the net effect of these few rules is to support the rich representation and inference capabilities

[3] Localized trauma to the brain can selectively knock out one representational system while leaving the others intact

that we seek.

What do they do?

1. **Each representational system interprets the world according to its representations only.** This is the most basic property of Representational Systems. Consequently different Representational Systems simultaneously represent an event in very different ways, and also apply very different kinds of inference mechanisms, this is the chief virtue of having multiple Representational Systems. To give an example - suppose our robot with Multiple Representational Systems is looking at an event where one block collides with another. The Visuospatial system may represent the sizes, shapes, motions, and spatial relations involved in that event (depending on what exactly was being attended to). The Socio-Thematic system simultaneously represents the event in a very different space, this system cares about whether the objects are alive or not, it may assign anthropomorphic roles of the doer and receiver to the blocks. The Motor-Somatosensory system only cares about the motor programs necessary to reach out and handle these blocks. Depending on the exact task at hand whether it is to decide if one of the blocks is going to fall over the table or if one of the blocks is going to break, the inference systems of very different Representational Systems might come into play. The inference capabilities of the Visuospatial system may be more competent to answer the first question whereas for the second question the knowledge about the brittleness of that particular object may be extracted from some linguistic representation of knowledge, Or the experience and inference properties of many Representational Systems may be tapped in the process of answering these questions.

2. **Representational systems detect regularities or "suspicious coincidences" within their representations.** A visuospatial representational system might discover that unsupported objects fall towards the earth. A motor representational system might simultaneously discover that adopting a particular hand configuration (opening the hand) while holding an object might result in the disappearance of the sensation of holding the object (the object fell). These regularities constitute a model of the environment, they are used to predict events before they happen. Notice that the model of the workings of the environment is distributed across different representational systems. They detect regularities in different aspects of an event. It is interesting to note that that in very simple organisms like Aplysia the mechanisms of habituation and sensitization serve the same purpose of detecting regularities, except that the kinds of regularities detected and the behaviors they drive are hardwired into the neural circuitry. The evolution of the cortex, brings the detection of regularities to the level of representational systems. One of the many existing hypothesis of the function of the cortex (Barlow 1985), views it as a sophisticated coincidence detector. The regularities that we have been discussing are *within* the representations of a representational system, but what about

detection of suspicious coincidences *across* representations of different representational systems? This is accomplished by a temporal associative memory.

3. **Each representational systems carves up the world into different categories based on its representations.** The visuospatial system represents aspects like size, positions, and motions of objects, so it might develop a distinction between objects that move by themselves and objects that don't. The motor system deals with motor programs, so it might distinguish between objects based on the motor programs used to deal with them, for example between objects that can be grasped in one hand and those that require both (because you need very different motor programs to deal with them). The Syntactic system classifies words like ran, fell, and jumped into the same category because they perform the same syntactic function in sentences. Within a Representational System the categories themselves may not be represented explicitly as categories but arise as the global effect of the application of a local associative memory rule.

4. **Representational systems have a synthetic capability.** The defaults or patterns that representational systems acquire during their interaction with the environment can be used to synthesize representations for hypothetical situations. From this viewpoint at a certain level there is no difference between "perception" and "imagination" - they both share the same substrate and the same inference processes apply to both.

How do they interact?

Representational systems interact with each other directly and also via the environment. The direct mode of interaction is only through a temporal associative memory mechanism, they do not have direct access to each others representations and processes. The indirect means of interaction is via the environment.

1. **Different Representational systems are connected together by a temporal associative memory mechanism.** If the representation of an event is distributed across different representational systems, then what binds these separate representations as one event? A temporal associative memory mechanism that runs through different representational systems is the answer. Representations that get activated within some time window in different Representational Systems get bound together by this mechanism. So the next time when some of the representations get activated, the other representations associated with that event get activated in other representational systems. This is a way of detecting "suspicious coincidences" across representational systems (previously we described it within a Representational System). One example of the consequence of detecting suspicious coincidences across different representational systems, is the formation of a distributed concept of the notion of an "object". Gary Drescher's Schema mechanism (Drescher 1991) demonstrates this by noticing that every time

the hand moves to some visual location where an object exists, there arises a touch sensation in the somatosensory-motor system, this repeatable association forms the basis of the system's model of its actions and also its conception of an object. The particular mechanism used to implement a temporal associative memory could be K-lines as suggested by Minsky (1985*a*), and/or the phase locking properties of networks of neurons as suggested by Damasio (1988) whereby the simultaneous firings of separate collections of neurons may settle into a fixed pattern of firing that distinguishes it as an event.

2. **Different Representational systems cause representational changes in each other via the environment** Representational systems interact with each other via the environment much like behaviors do. A Representational system may be driven by the environment to perform some operations involving its representations, this in turn may result in some action thus causing a change in the environment which in turn drives representational changes in other systems. Understanding nature often requires the establishment of different layers of complexity. The kinds of phenomenon exhibited by a collection of interconnected neurons (eg population coding) is qualitatively different from the behavior of an individual neuron (eg shunting inhibition), hence the establishment of a different repertoire of concepts for describing them. Brooks showed how the complexity of the interactions of many simple behaviors can gives rise to sophisticated overall behavior (which might give the illusion of complicated underlying control and planning). The same is true of the dynamics of interactions of multiple representational systems. The interactions of different kinds of representations and inference mechanisms will lead to sophisticated problem solving giving the illusion of some powerful general inference mechanisms where none might explicitly exist. For example a statement that is intuitively obvious to anyone who knows the rules of chess is that starting from any square the king can move to any part of the chess board, this inference could well be the result of the interaction of visuospatial, motor, and syntactic representational systems rather than some general problem solver.

How are they connected to behaviors?

Representational Systems modulate behavior. Behaviors are independent of Representational Systems, Initially Representational Systems have no models of behaviors, gradually they learn to modulate them.

Behaviors are the initial agents of representational change (eg correlation)

cognitive development

Piaget described four distinct stages of cognitive development that he claimed all humans must pass through before reaching adult level competence at by the age of 15. "Cognitive development is a coherent process of successive qualitative changes of cognitive structures” (Wadsworth 1984). The capabilities associated with each stage are described in detail however there is only a little discussion of *what* the cognitive structures are that develop and the *mechanism* by which they develop. The basic unit of development is called a “schema”. It can be viewed as as the representation of some regularity in the environment (eg every time I execute some motor command at some position, my hand moves to some other region of space (Drescher 1991)). The processes “assimilation” reinforces old schemas, and the process of “accommodation” uses new-experience to either modify old schemas or construct completely new ones. That in a nutshell is the Piagetian description of the structure and mechanism. Lately

3 Testing the ideas

Let’s review where we are: We started out by adopting the view that distinguishing feature of human cognition is the *variety* of representational and inference capabilities that it displays. We then made two hypothesis about the origin of this representational power - the basis set and grounded primitives hypothesis. We then made a particular choice of this basis set in Multiple Representational Systems and described what they do and how they interact. If the central claim of this paper is true then we should be able to support various “abstract” cognitive tasks on the interaction of the Representational Systems. Two examples of abstract cognitive tasks which we will try to elicit from our representational systems are:

1. Obvious Cognitive Judgments/Common sense reasoning.
2. Understanding simple sentences about spatial events.

The reason for choosing these tasks is because there has been very little success in solving them in the past.

Obvious Cognitive judgments or Common sense reasoning

There are a number of “facts” about the world that most humans after an early age, find obvious. Here are three such statements:

1. If point A is inside a circle and point B outside it, then any path connecting A and B must cross the circumference of the circle somewhere.

2. Large things are inconvenient to lift.
3. If A likes B, then A will seek the company of B.

What is intriguing is that it is precisely these kind of facts - the ones that we regarded as obvious - that Artificial Intelligence programs give up on. The quick and dirty patch of hand coding such statements as axioms has traditionally run into a severe problems - there are *too many* of such facts and no matter how many we hand code - it only marginally delays the time before the theorem prover gives up for want of another fact that is obvious to us. Common sense reasoning is usually the name given to feats of reasoning that are so effortless for humans that the results seem obvious. Marvin Minsky (Minsky 1985*b*) says about it:

> Common sense is not a simple thing. Instead, it is an immense society of hardearned practical ideas—of multitudes of life-learned rules and exceptions, dispositions and tendencies, balances and checks. If common sense is so diverse and intricate , what makes it seem so obvious and natural? This illusion of simplicity comes from losing touch with what happened during infancy, when we formed our first abilities.

"Common sense knowledge" is a term that has come to encompass all knowledge that seems obvious to any human but has eluded machines. I suggest that common sense knowledge has eluded machines because firstly we have ignored the role of embodiment, perception, and action as a basis for such knowledge, and secondly that such knowledge is a product of the interaction of multiple representational systems.

Natural Language Semantics

Linguists are fond of pointing to language as "the jewel in the crown of cognition". They have focused primarily on syntax, it has proven more tractable than semantics. Understanding a sentence like "The ball fell into the bucket" requires many different kinds of expertise - Knowledge about syntax, knowledge about the thematic structure of the sentence, visuospatial knowledge about balls, buckets, their shapes, sizes, and motion, and how a ball could fall into a bucket. It has long been acknowledged that natural language semantics must somehow be based on perception and action, I suggest that we organize this grounding in terms of multiple representational systems. Representational power of language arises from the interaction between the expressive properties of generative grammar and the representational capabilities of other representational systems.

What next?

The task now is to first close the loop. We need to:

- Build a few of the Representational Systems (say Visuospatial and Syntactic) after making specific choices about the representations and processes.
- Implement the kinds of interaction described above. The list of what Representational Systems do and how they interact as described in the previous section may seem very terse, for instance there is no explicit specification of how the visuospatial Representational System interacts with the Syntactic system. This omission is intentional, the rules of interaction mentioned are powerful in themselves and we should try them before doing something more complicated.
- Choose some small set of behaviors (eg Tracking humans, and trying to engage them in conversation). The issue of how Representational Systems influence behaviors is very important and has not been described in this paper yet.
- Turn on the system and watch the interaction of the Representational Systems lead to cognitive development. At this preliminary stage in our understanding of dynamics of interactions of Representational Systems, we will have to learn by trial and error, i.e small changes in the workings of the Representational Systems can lead to large changes in the cognitive structures that develop.

This is a lot of work, but I have confidence that it is a viable path to achieving human like intelligence. My proposed thesis " A Visuospatial Representational System" is a first step in this direction.

References

Barlow, H. B. (1985), Cerebral Cortex as a Model Builder, *in* D. Rose & V. G. Dobson, eds, 'Models of the Visual Cortex', John Wiley and Sons Ltd, chapter 4.

Bickerton, D. (1990), *Language and Species*, first edn, The University of Chicago Press.

Brooks, R. A. (1990), 'Elephants Don't Play Chess', *Robotics and Autonomous Systems* **6**, pp3–15.

Brooks, R. A. (1991*a*), Intelligence Without Reason, A. I. Memo 1293, M. I. T. Artificial Intelligence Laboratory.

Brooks, R. A. (1991*b*), 'Intelligence without Representation', *Artificial Intelligence* **47**, pp139–159.

Damasio, A. R. (1988), The Brain Binds Entities and Events by Multiregional Activation from Convergence zones, *in* '??', ??

Drescher, G. L. (1991), *Made up Minds - A Constructivist Approach to Artificial Intelligence*, MIT Press.

Minsky, M. L. (1985*a*), *The Society of Mind*, first edn, Simon and Schuster Inc, chapter K-Lines: A Theory of Memory, p. 82.

Minsky, M. L. (1985*b*), *The Society of Mind*, first edn, Simon and Schuster Inc, chapter Common Sense, p. 22.

Newell, A. & Simon, H. (1961), GPS: A program that simulates Human Thought, *in* Feigenbaum & Feldman, eds, 'Computers and Thought'.

Newell, A. & Simon, H. A. (n.d.), Computer Science as Empirical Inquiry: Symbols and Search, *in* J. Haugeland, ed., 'Mind Design'.

Newell, A., Shaw, J. & Simon, H. (1961), Empirical Exlorations with the Logic Theory Machine: A Case Study in Heuristics, *in* Feigenbaum & Feldman, eds, 'Computers and Thought'.

Pinker, S. (n.d.), The Evolution of Cognition, Notes for an IAP 92 talk at M.I.T.

Rosenbloom, Laird, Newell & McCarl (1991), 'A preliminary analysis of the Soar architecture as a basis for general intelligence', *Artificial Intelligence* **47**, pp289–325.

Wadsworth, B. J. (1984), *Piaget's Theory of Cognitive and Affective Development*, third edn, Longman.