Information-Processing Systems in Nature

Aaron Sloman

School of Computer Science
The University of Birmingham
Birmingham, B15 2TT, UK
`A.Sloman@cs.bham.ac.uk`
WWW home page: `http://www.cs.bham.ac.uk/˜axs/`

Abstract. This paper is a sequel to my invited contribution to PPSN2000. It attempts to identify and analyse a collection of issues implicitly taken for granted in the earlier paper and in a great deal of literature which assumes that biological organisms do information processing. Normally it is assumed that we all understand intuitively what it means for something to be an information-processor, whether natural or artificial. I attempt to offer the beginning of an analysis which attempts to justify many of the ordinary ways of talking about information in organisms – some of which attract critical comments from those who are sceptical about attempts to talk about computation and representations in organisms. In the long run I hope to show that that sort of scepticism is misguided.

(Last revised April 22, 2004)

1 The Problem

The claim that organisms process information, use representations and perform computations is taken by many (including me, 30 years ago [12]) to be intuitively obvious.[1] Unfortunately it is not really clear what is meant by 'information' in this context, and most people who talk about such things do not attempt to analyse it. (E.g. it is often assumed to be clear what is meant by saying that DNA includes information.) Recently the use of 'information-processing' terminology, and in particular talk of organisms as performing computations and using representations has been strongly challenged (often with great passion), for instance in discussions at a recent workshop[2] and elsewhere. It seems that there are different research communities, some of which regard such claims as fairly clear and generally correct, while others reject them as based on over-generalising notions like 'computation', 'representation' and 'information processing'. For instance a well known biologist suggests in [10] that only computers, not brains, process information. I'll attempt to show that all the views arguing that objects of a certain type do not process information (or do not compute, or do not use representations) are based on overly narrow conceptions of 'information' (or computation, or representation). The critiques are spurious because they construe the key terms as narrow technical terms, in such a way as to prevent them being applicable to animals. E.g., if we define 'information-processing' in terms of bit manipulation and show that brains are not bit-manipulators that proves that animals do not process information – but only in *that* sense of 'information'.

This sort of argument is invalid if the key concepts are 'cluster concepts': they involve an *indeterminate mixture* of rather loosely defined intuitive concepts and precisely but narrowly defined technical concepts. The narrow technical notion of 'information' developed by Shannon is a purely syntactic notion related to numbers of distinct messages of a given length that can be transmitted on a channel, whereas when people say that organisms acquire, process and use information, the notion of 'information' used usually refers to a much older, less precise, more intuitive, *semantic* notion of information — involving reference, meaning, truth and falsity, implication, consistency, etc.

Likewise arguments that brains do not do computation often depend on a narrowly defined notion of computation, based on notions like 'Turing machine' or 'recursive function', or possibly something like a von Neumann computer architecture in which bit-patterns are manipulated. Proponents forget that there is a more general, much older, notion of 'computation' used to describe activities of humans, and in that older sense computations were performed by brains long before machines were designed to replace humans in certain computing tasks. This is a special case of a general point: systems that are very different at one level of description may be very similar at another level, and that can be true of brains and computers, for instance if both support astronomically large numbers of state-transitions because they have large numbers of independently changeable components. As argued in [20], identifying those high level similarities can enhance our understanding of both systems.

---

[1] Try giving the quoted phrase "information process" together with any of the following words to a search engine: biology, brains, neural, animals, humans, social, psychiatry, learning, perception, emotion, deficit — each produces many thousands of hits.

[2] http://www.ecs.soton.ac.uk/~rid/wgw02/home.html

Finally, claims that natural intelligent systems do not use, or need, representations, e.g. in work inspired by [3], often assume that the word ‘representation’ refers to specific kinds of symbolic structures manipulated by AI programs, because so much work in AI, e.g. [8, 9], has been concerned with a rather limited class of representations which are relatively easy to specify and manipulate using current computer programming languages. Such claims ignore the fact that even in the AI community there have been investigations exploring a wider range of forms of representation, some discussed in [2, 7, 1], others familiar in discussions of connectionism, where it is common to distinguish local and distributed representations, for instance.

In [16], I offered a general overview of varieties of representations. In this paper, I’ll offer a different overview that makes some basic assumptions more explicit and presents a methodology for clarifying ideas within the ‘design-based’ approach, using thought experiments to explicate various useful distinctions (a method first used, as far as I know, by Galileo). Warning: though useful, the distinctions are not always *sharp*.

2 Varieties of types of information content

The key idea is the notion of ‘information about something’. This includes a variety of different types of information content, which can involve *reference* (to a particular physical object, animal, place, etc.), *an individual state of affairs* (e.g. a fact about the size, shape or other property of some thing, or a relation between several things, such as grasping, containing, touching, being between), a *generalisation* (e.g. unsupported objects fall, things with a certain shape and colour taste bad, big things cannot get through small gaps, etc.) and various other kinds of factual information, including information about what exists, what is possible, what is necessarily the case, and so on.

Besides variation in types of information content there are different ways in which an information-user may relate to the same information content. Some organisms can have only a relation something like *believing* or *not believing* that there is a snake in the grass. In humans, and presumably some other animals, it is also possible to *intend* to put a snake in the grass, to *wonder whether* there is a snake in the grass, or to *attempt to find out whether* there is. The last two require the ability to consider a question without knowing or having an opinion about the answer. This ability to disconnect information from truth is a major aspect of intelligence, and apparently it is one of the many things whose importance evolution ‘discovered’ about information long ago.[3]

The previous examples all concern *factual* information (which includes information about how things might be, even if they are not). There is also *control* information, which is about what to do, which can be very specific imperatives (e.g. run, jump, turn left, chew, freeze, pick that rock up, etc.), conditional imperatives (e.g. if hungry find food), general preferences (e.g. prefer escaping danger to satisfying hunger) or abstract values (e.g. certain types of things are good and should be sought when the opportunity arises or bad and should be prevented when the opportunity arises).

---

[3] Confusion is sometimes caused by the fact that in ordinary language the word ‘information’ implies correctness e.g. ‘Where can I get information about train times?’ But that implication is not universal, e.g. ‘Why did you give me incorrect information?’ I use ‘information’ in the neutral sense, allowing false information, which is more general and more useful for a scientific investigation of information-processing.

Some information-users can handle only a fixed set of possible information contents, whereas others, especially humans, can combine information contents to construct more complex information contents, and that ‘compositionality’ is one of the important functions provided by various kinds of languages (formal and informal, verbal and pictorial, including programming languages). Of course, that requires mechanisms that operate appropriately on those information-bearers, which most organisms lack.

This is not intended as a complete survey of types of information content, merely an indication of some of the variety. Most of the distinctions admit intermediate cases.

3 Information bearers – implicit and explicit

Besides surveying types of information *content* we also need to survey types of information *expression*. Information can be expressed, or stored, or transmitted, or applied in many ways: and although humans often use words, phrases, and sentences as information bearers, not all information has to be expressed that way. But it must be expressed in *some way*. For information to be acquired, transformed, combined, used, or even considered as a possibility (e.g. wondering whether a snake is hiding in the long grass, or deciding whether to climb that tree to get the fruit that is out of reach), there must be something that expresses or encodes the information: an *information-bearer*.

In very simple organisms the only means of expression is by something actively doing something. For instance, as far as I know, a house-fly can have information about an approaching object only while certain sensors and sub-systems activated by sensors are active and exerting some influence on the rest of the system. There is, I suspect, no part of the fly that is able to encode the fact that there *was* an approaching object, or that there *might be* later. There is no enduring ‘symbol’ for *approaching object* that the fly can store and re-use in different contexts, including combining it with other symbols in a formalism with compositional semantics.

We call information that is encoded only in some active state of the system which has particular causal functions in relation to the rest of the system, *implicit* information. In contrast information is *explicit* when it is encoded in some enduring object which may be decoupled from both sensory activation and particular application, but can be accessed, copied, combined with other things, and used when necessary.[4] The use of explicit information-bearers is crucial to the use of compositional semantics.

Even if house-flies are not limited to implicit information, and can use explicit (storeable, directly manipulable) information, it is likely that the simplest organisms have that limitation. The main point here is not a particular empirical claim but a conceptual distinction. For the distinction enables us to ask a question: how and why did evolution develop mechanisms supporting explicit, re-usable, re-combinable information bearers, and which organisms use them and which don’t? Are both kinds used in a new-born human infant? In a six week-old foetus? What intermediate cases are there?

---

[4] This has nothing to do with the distinction some psychologists make between ‘implicit’ meaning ‘unconscious’ and ‘explicit’ meaning conscious. Perhaps ‘evanescent’ and ‘re-usable’ might have been better terms. But no words of ordinary language accurately make the distinction. NB: as in all biological distinctions there are probably many interesting intermediate cases used by evolution in making the distinction. That does not make it a matter of *degree*.

The implicit/explicit distinction made here is merely one of a variety of distinctions between ways in which information can be expressed. Other distinctions include whether the information-bearers can vary continuously or discretely, whether they have Fregean or analogical modes of composition ([11]), whether they are context-dependent or not, whether they allow macros, whether they allow quotation, etc.

4 Information bearers – physical and virtual

A particularly important distinction is between physical and virtual information-bearers. An English sentence like 'Bears catch fish' is not a physical object, but an abstract entity made of words that can be expressed in spoken or written utterances or in unuttered thoughts or memories. When words are encoded in letters, the letters can be expressed in many forms, written symbols in a multitude of fonts, semaphore or morse code signals, bit patterns in computers, and whatever human brains happen to use for storing information about how words are spelt. Words composed of letters illustrate the notion of abstract entities that are *multiply realisable*. This is typical of components of virtual machines, which are abstract machines manipulating abstract entities, such as numbers, words, images, plans, proofs, programs, etc. (Throughout this paper, by 'abstract machine' I mean '*running* abstract machine', e.g. the Linux virtual machine on which I am typing this document, not just a mathematical specification.)

Although information is a very abstract notion (somewhat like 'energy', insofar as there are many different forms in which energy can exist, with different physical requirements and different uses) the processing of information always depends on the existence of a concrete instance of some kind of physical machine. In other words, all information-processing is *physically implemented*. However, the relationships between the structure, content and manipulation of information and the underlying physical states and processes can vary widely and can be very indirect. To understand this we need to use the notion of *virtual machine* which is normally used in connection with computing systems, but can be generalised (as in [17]) to a wide range of systems in which abstract states and processes occur, including social and economic systems as well as psychological states and processes.

One advantage of using information-bearers in virtual machines is that virtual machine entities are not subject to the same constraints as physical ones. E.g. two virtual entities can contain each other, which is not possible for physical entities. Also lazily-evaluated virtual entities can be infinite. (It is arguable that the philosopher Kant invented approximately this idea over two hundred years ago.) One enormous advantage of the use of virtual machine structures as information-bearers is that, in a suitably designed implementation, complex virtual structures can be constructed and changed far more rapidly than physical structures. The need for this is particularly obvious in connection with visual perception, since physical motion of the perceiver, including things like rotation of the head, or going round an opaque object, or emerging from a cave, can cause very rapid changes in the visible portions of the environment, requiring rapid changes in the perceptual information content of a perceiver of the environment, even if the environment is unchanging. The changing patterns of 2-D illumination are easily captured in changing states of retinal activity. But changing percepts of trees, animals, people, clouds, walls, doors, cars, and text on a display, are far harder to implement.

Another example is the ability to rapidly consider several possible action sequences, or several complex possible explanations of some perceived unexpected state of affairs, e.g. footprints in the sand, or the absence of offspring from the place where they were hidden. Reading and understanding (or mis-understanding) this text is another example.

Although for us this notion of virtual machine is a recent invention, arising out of developments in computer science and software engineering, it seems that evolution 'discovered' its importance long ago by developing brain mechanisms that can be used for storing and manipulating multiple types of information, as shown most dramatically by the ability of physically very similar human brains to do different things in different cultures and different historical epochs, for instance learning different languages, social customs and explanatory theories. Even identical twins' brains can store wildly different collections of facts, e.g. if they grow up in different cultures.[5]

Things that happen in virtual machines can be causes, or effects, of other things, including physical things.[6] E.g., processes in a spelling checker can cause spelling mistakes to be detected and corrected, which can cause physical changes in a file-store and on a computer display. Decisions in a flight control system can save lives.

Virtual machines can contain entities used as information-bearers. For instance, some of the behaviour of a robot could be controlled by collections of condition-action rules expressed in list structures containing symbols in a virtual machine of the sort used by Lisp and other AI programming languages. Some of the robot's factual information about the environment derived from sensory data may use similar list structures. Of course the list-processing virtual machine will be implemented in physical mechanisms, but the structures expressing the rules and factual beliefs for the robot are not physical entities. They cannot be observed by opening up the machine and examining it using devices employed by physicists and chemists, for instance.

This means that 'The physical symbol system' hypothesis of Newell and Simon [9] is incorrect: some symbol systems are not physical, but exist in virtual machines. Instead we assume 'The physically implemented symbol system' hypothesis, since all working virtual machines must be implemented in some physical mechanism. (This is probably what Newell and Simon intended to say.)

So far we have said many things about information acquired, derived, stored, or used in natural systems, especially humans — including things about how it can vary in content, how it can be implicit or explicit, how it can use compositional semantics or not, how it can use physical information-bearers or information-bearers in running virtual machines. We have also begun to hint at the variety of ways it can be used by organisms, and suggested that evolution produced information-users that vary in their sophistication. But we have not said what information *is*.

---

[5] The ability of one machine to contain hugely diverse collections of information is connected with the ability to have very large numbers of different states, and state sequences. It seems that in brains, as in computers, this depends on having large numbers of independently switchable components, as suggested in [20].

[6] Some philosophers find this hard to believe! This is discussed in more detail in [22] and [21].

5 Can we define ‘information?

Like many deep theoretical concepts in the sciences, our concept of ‘information’, or semantic content, cannot be *explicitly* defined in terms of more basic notions, but can be *implicitly*[7] defined at least partially, by specifying facts about information, such as the many facts we have begun to summarise in previous sections – just as predicates and function symbols in a mathematical theory are implicitly defined by the sets of axioms and inference rules in which they occur (which constrain the possible models: [13]) .

Implicit definitions (sometimes analysed as ‘meaning-postulates’ following [4]) constrain the interpretation of a concept by *using* it rather than by equating it to something previously understood. Unlike the mathematical case the concepts implicitly defined in scientific theories are only *partially* defined because the concept may be indefinitely extended or modified by enriching the theory. E.g. theories about information and information processing can develop over time – as also happened to the concept of ‘energy’ in physics, which at the time of Newton was restricted to kinetic energy, gravitational potential energy and elastic energy, and did not include chemical energy or mass-energy. For more on conceptual changes in science see [5].

This paper presents reminders regarding well-known and less well-known facts which constitute an implicit, though partial, definition of ‘information’, and related concepts such as ‘processing’ and ‘using’ information. Whether such concepts are useful will depend on whether the theory incorporating those facts turns out to have predictive and explanatory power, and to be usable in solving practical problems. Whether any reader regards the analysis as fitting his or her understanding of ‘information’ is less important: people are not good at inspecting their own concepts.

6 Evolution as AI engineer

Conjecture: *All organisms are information-users*. Being an information-user involves being more than just a physical object. All physical objects, including very complex things like tornadoes, and planets, and much simpler things, like a grain of sand blown in the wind and a marble rolling down a helter-skelter, produce behaviour which is the result of physical forces acting on and in the object, whereas organisms have an additional feature: they have a store of energy which they are able to deploy through mechanisms that they can control, and they have sensors which provide information that is used to determine how and when to deploy that energy in order to meet some need of the organism, which could be reproduction, survival, repair, prevention of damage, or achieving some sub-goal derived from other needs.

Even the simplest organism which does little more than move through a fluid in the direction of maximum increasing density of nutrients, or contracts its membrane to draw away from contact with something noxious, fits this description, though that does not imply that it has internal states in which there are thoughts, percepts and decisions of sorts that we could express in language. Its information-bearers are only the evanescent states of its sensors and control mechanisms. But even if it is very close to being a limiting case, such an organism fits the description of an information-user. All of these

---

[7] This has nothing to do with the previous contrast between implicit and explicit expressions of information.

capabilities require organisms not to be simply passively driven by the resultants of external forces: they are sometimes able to use an internal store of energy to resist an external force, or to initiate a process that no external force will initiate, e.g. to move to where there is a large source of additional energy.

A result of natural selection is that, over generations, information-users develop mechanisms that help to ensure their ability to reproduce, and that sometimes includes mechanisms that extend the abilities of individuals to survive, including mechanisms that enable them to acquire nutrients and other substances they need, to detect and avoid harmful entities, and to repair damage. Some organisms, instead of using only their own energy supplies, find ways to use external sources, for instance plants that release pollen or seeds into the wind. In some cases this uses sensors to determine details (including timing) of a small change, using internal energy, that can benefit from a larger external source of energy.[8]

In addition to the information gained by individual organisms through sensors, there is also information gained by the genome through natural selection: this is, roughly, information about which designs work. This implicitly encodes information about certain aspects of the environment, but not in any form that could be regarded as encoding facts like ‘winds are sometimes strong enough to carry seeds’, or ‘some passing animals have fur that allows seeds with hooks to be transported’.

7 Architectures for information-users

One of the most subtle and most impressive aspects of evolution is production of more and more complex and sophisticated information-users. Increasing sophistication involves many steps including using explicit information-bearers, separating the time of acquisition and the time of use of information, developing capabilities supporting compositional semantics, developing deliberative mechanisms for exploring questions, hypotheses and plans, and developing the ability to represent and reason about states and processes in information-processors. However, these evolutionary steps are invisible, since although fossil and other records directly reveal changes of *physical* properties, such as shape, size, kind of covering, etc., fossils do not directly reveal behaviour, least of all internal behaviour that involves the processing of information in perception, learning, reasoning, and the development of ontologies suited to the environment. Developing good theories about the evolution of virtual machines will inevitably depend on much speculation and very indirect testing. (Likewise developmental psychology).

Evolution of ever more sophisticated information-processing capabilities is to be expected once it is acknowledged that using information can be beneficial to organisms. It is now a common-place of modern biology that evolution (i.e. random genomic changes along with natural selection) produced increasingly complex and varied means of locomotion, of consuming prey, of hiding from or out-running predators, of defend-

---

[8] Thanks to John McCarthy for the point about using external energy and to Kay Hughes for the observation that this can still involve use of internal energy, e.g. when birds that use gravity and up-drafts as external energy sources use internal energy to change their wing configurations appropriately. In many windmills, the wind provides energy for both purposes.

ing territory, of attracting mates.[9] So it should not come as a surprise that evolution produced a wide variety of ways of acquiring, analysing, interpreting, transforming, storing, combining, transmitting, and using information. *How* it happened, is another matter. There are probably more intermediate cases than we have dreamt of.

Some obvious cases of this diversity are variations in types of sensory transducers: capable of detecting chemicals (in air or water), temperature, pressure, friction, sounds, photons, magnetic or electric fields, gravitational fields, and a host of internal states. Other obvious cases involve mechanisms for controlling muscles and other devices that produce motion, chemical signals, growth, repair, assimilation of food, sexual and other displays, and so on. Less obvious is the fact that evolution, sometimes aided by a learning-bootstrapping process, has produced many different kinds of *internal* mechanisms for analysing and interpreting sensory information from similar sensors, and for controlling the use of similar effectors. For example, not all animals that have acute auditory sensors appear to have the internal physical and virtual machines required to analyse acoustic signals at different levels of abstraction concurrently so as to detect the phonemes, words, phrases, sentences, stress-patterns, intonation-contours, and tonal features that form part of the process of understanding language including factual reports, requests, pleas, exhortations, stories and poetry.[10] Similar comments can be made about the use of multiple levels of control in virtual machines for producing linguistic output, whether in speech, hand-writing, semaphor signals, sign-language or text typed at a keyboard, such as this paper.

It was conjectured in [17] (and in older papers, [15, 19]) that some of the developments in sensory and motor information-processing capabilities were related to the evolution of different layers in central processing mechanisms, concerned with learning, goal-formation, decision-making and action control. For example, although many reactive mechanisms can operate with continuously varying sensor information, the use of a planning mechanism that can construct, compare and choose between *sequences* of steps into the future requires *discretisation* of information so that options can be ‘chunked’ for exploration of possible futures. As far as I know there is no way for a search engine directly to explore branching continua of continua. The use of chunks supports the learning of generalisations relating such chunked options, so that distinct types of actions that are possible in distinct situations can be learnt, and the consequences of those actions learnt. This need for chunking, it is conjectured, produced evolutionary pressure for development of specialised perceptual mechanisms that chunked perceptual input at different levels of abstraction. These could have been the earliest precursors for the current mechanisms in humans that analyse acoustic and visual *linguistic* input at different levels of abstraction, as discussed above.

A particular type of chunking would categorise types of actions, such as grasping, releasing, pushing, throwing, catching, entering, leaving, hitting and many more. Some organisms might need to be able to categorise both their own actions and also the actions of others in terms of these chunks. For example, learning by imitation requires

---

[9] It is also sometimes argued that non-genetic evolution can contribute to such processes, but that debate makes no difference to what follows.REF BJPS 2004.

[10] Chapter 9 of [12] illustrates an example of concurrent multi-level processing in interpreting complex and noisy visual depictions of words.

a perceptual process that abstracts from the precise details of an action performed by another animal, to produce a characterisation that can also be applied to one's own actions, not only accounting for detailed differences in shape, size, dispositions of limbs, etc. but also abstracting from the difference in viewpoint when representing perceived actions of others and actions that might be done by oneself: a still unsolved problem in computer vision, except for 'easy' cases.

Merely discovering that the same neurons fire in the brain for similar actions done by others and by oneself, as reported in connection with 'mirror neurons' explains nothing. In particular, it leaves completely unexplained the very complex perceptual and control processes that can relate the details of actions done by others and by oneself to the same abstract representation of that type of action.

Once an animal has the ability to chunk processes into different sorts of actions, including those that it can perform, it may also be able, using *explicit* symbols, to relate features of objects and situations to the actions that can be performed on those objects or in those situations and to the possible goals that would be achieved thereby, for instance seeing ways of grasping a mug and the different consequences of different sorts of grasps. This seems to be a requirement for the perception of affordances [6, 14, 18].

8 Representing information-processors: self and others

Another conjectured development involved the production of mechanisms for processing information about systems capable of processing information! This includes both the representation of *other* information-processing systems, such as other organisms, and representing *one's own* information-processing. The former includes information about what another agent intends, or can see, or has learnt. The latter includes information about one's own intentions, percepts, strategies, etc. In both cases, there are important differences between describing properties, relations and movements of physical objects, and describing aspects of information-processing in virtual machines. For instance the latter can introduce referential opacity, e.g. 'Fred believes the alligator biting his leg will eat him' can be true even when there is no alligator.

Conjecture: just as evolution produced animals with a useful, though incomplete and partially erroneous ontology for types of physical objects, properties, relations and behaviours, so also did it produce, in a smaller set of species, a useful, though incomplete and partially erroneous ontology for information-processing phenomena. In particular human children seem to be born with a bootstrapping mechanism that very early makes them start treating other humans, e.g. parents, as more than just physical objects. I doubt that this is merely a by-product of some totally general learning mechanism shared with all other animals that learn, though much of it may be shared with other primates. It seems to be the case that, unless checked by a coldly rational culture, such mechanisms readily over-generalise and over-interpret, causing mentality to be attributed almost to anything that moves. This could be more conducive to survival than a less generous ontology: false positives do no harm, whereas false negatives can get you eaten. Likewise, internal perception can include over-simplifications and errors, yet still be useful.

9 Conclusion

The vast majority of work on biologically inspired computational modelling either investigates very simple simulated creatures, or applies general mechanisms, such as neural learning or genetic algorithms, to particular practical problems. I have been trying to recommend a broader vision for researchers which includes trying to understand how evolution could produce complex, multi-functional organisms that in addition to having many physical properties and capabilities also have a wide variety of different concurrently active information-processing capabilities many of which though probably not all (e.g. the ability think about infinite sets) are shared with other animals. Progress in this area will require interdisciplinary investigations including AI, Computer Science, Software engineering, psychology (many kinds), neuroscience, biology, ethology, and social science. Making progress with this grand challenge,[11] is far more important than continuing the factional battles in which researchers try to argue that only the architectures, mechanisms and forms of representations that *they* study are the right ones to use. Likewise the ideas and distinctions presented here are illustrative rather than definitive and there are probably many interesting intermediate cases still to be discovered by us, already discovered by evolution, and possibly used in individual development.

References

1. M. Anderson, B. Meyer, and P. Olivier, editors. *Diagrammatic Representation and Reasoning*. Springer-Verlag, Berlin, 2001.
2. R.J. Brachman and H.J. Levesque, editors. *Readings in knowledge representation*. Morgan Kaufmann, Los Altos, California, 1985.
3. R. A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.
4. R. Carnap. *Meaning and necessity: a study in semantics and modal logic*. Chicago University Press, Chicago, 1947.
5. L.J. Cohen. *The diversity of meaning*. Methuen & Co Ltd, London, 1962.
6. J.J. Gibson. *The Ecological Approach to Visual Perception*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1986. (originally published in 1979).
7. J. Glasgow, H. Narayanan, and B. Chandrasekaran, editors. *Diagrammatic Reasoning: Computational and Cognitive Perspectives*. MIT Press, Cambridge, Massachusetts, 1995.
8. J. McCarthy and P.J. Hayes. Some philosophical problems from the standpoint of AI. In B. Meltzer and D. Michie, editors, *Machine Intelligence 4*. Edinburgh University Press, Edinburgh, 1969. (Accessible as http://www-formal.stanford.edu/jmc/mcchay69/mcchay69.html).
9. A. Newell. Physical symbol systems. *Cognitive Science*, 4:135–183, 1980.
10. S. Rose. *The Making of Memory*. Bantam Books, Toronto, London, New York, 1993.
11. A. Sloman. Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence. In *Proc 2nd IJCAI*, London, 1971. Reprinted in *Artificial Intelligence*, vol 2, 3-4, pp 209-225, 1971, and in J.M. Nicholas, ed. *Images, Perception, and Knowledge*. Dordrecht-Holland: Reidel. 1977.
12. A. Sloman. *The Computer Revolution in Philosophy*. Harvester Press (and Humanities Press), Hassocks, Sussex, 1978. Online at http://www.cs.bham.ac.uk/research/cogaff/crp.
13. A. Sloman. What enables a machine to understand? In *Proc 9th IJCAI*, pages 995–1001, Los Angeles, 1985.

---

[11] Described further here http://www.cs.bham.ac.uk/research/cogaff/gc/

14. A. Sloman. On designing a visual system (Towards a Gibsonian computational model of vision). *Journal of Experimental and Theoretical AI*, 1(4):289–337, 1989.
15. A. Sloman. The mind as a control system. In C. Hookway and D. Peterson, editors, *Philosophy and the Cognitive Sciences*, pages 69–110. Cambridge University Press, Cambridge, UK, 1993.
16. A. Sloman. Towards a general theory of representations. In D.M.Peterson, editor, *Forms of representation: an interdisciplinary theme for cognitive science*, pages 118–140. Intellect Books, Exeter, U.K., 1996.
17. A. Sloman. Interacting trajectories in design space and niche space: A philosopher speculates about evolution. In *et al.* M.Schoenauer, editor, *Parallel Problem Solving from Nature – PPSN VI*, Lecture Notes in Computer Science, No 1917, pages 3–16, Berlin, 2000. Springer-Verlag.
18. A. Sloman. Evolvable biologically plausible visual architectures. In T. Cootes and C. Taylor, editors, *Proceedings of British Machine Vision Conference*, pages 313–322, Manchester, 2001. BMVA.
19. A. Sloman. How many separately evolved emotional beasties live within us? In R. Trappl, P. Petta, and S. Payr, editors, *Emotions in Humans and Artifacts*, pages 35–114. MIT Press, Cambridge, MA, 2002.
20. A. Sloman. The irrelevance of Turing machines to AI. In M. Scheutz, editor, *Computationalism: New Directions*, pages 87–127. MIT Press, Cambridge, MA, 2002. (Available at http://www.cs.bham.ac.uk/research/cogaff/).
21. A. Sloman and R.L. Chrisley. Virtual machines and consciousness. *Journal of Consciousness Studies*, 10(4-5):113–172, 2003.
22. A. Sloman and M. Scheutz. Tutorial on philosophical foundations: Some key questions. In *Proceedings IJCAI-01*, pages 1–133, Menlo Park, California, 2001. AAAI. http://www.cs.bham.ac.uk/˜axs/ijcai01.