Artificial Intelligence-A Personal View

D. Marr
The Artificial Intelligence Laboratory,
Massachusetts Institute of Technology,
545, Technology Square, Cambridge, MA 02139, U.S.A.

Recommended by B. Meltzer

ABSTRACT

The goal of Artificial Intelligence is to identify and solve tractable information processing problems. In so doing, two types of theory arise. Here, they are labelled Types 1 and 2, and their characteristics are outlined. This discussion creates a more than usually rigorous perspective of the subject, from which past work and future prospects are briefly reviewed. Artificial Intelligence is the study of complex information processing problems that often have their roots in some aspect of biological information processing. The goal of the subject is to identify interesting and solvable information processing problems, and solve them. The solution to an information processing problem divides naturally into two parts. In the first, the underlying nature of a particular computation is characterized, and its basis in the physical world is understood. One can think of this part as an abstract formulation of what is being computed and why, and I shall refer to it as the "theory" of a computation. The second part consists of particular algorithms for implementing a computation, and so it specifies how. The choice of algorithm usually depends upon the hardware in which the process is to run, and there may be many algorithms that implement the same computation. The theory of a com- putation, on the other hand, depends only on the nature of the problem to which it is a solution. Jardine and Sibson [6] decomposed the subject of cluster analysis in precisely this way, using the term "method" to denote what I call the theory of a computation. To make the distinction clear, let us take the case of Fourier analysis. The (computational) theory of the Fourier transform is well understood, and is expressed independently of the particular way in which it is computed. There are however several algorithms for implementing a Fourier transform the Fast Fourier transform (Cooley and Tukey [3]), which is a serial algorithm; and the parallel "spatial" algorithms that are based on the mechanisms of coherent optics. All these algorithms carry out the same computation, and the choice of which one to use depends upon the available hardware. In passing, we also note that the dis- tinction between serial and parallel resides at the algorithm level, and is not a deep property of a computation. Strictly speaking then, a result in Artificial Intelligence consists of the isolation of a particular information processing problem, the formulation of a computational theory for it, the construction of an algorithm that implements it, and a practical demonstration that the algorithm is successful. The important point here, and it is what makes progress possible, is that once a computational theory has been established for a particular problem, it never has to be done again, and in this respect a result in A.I. behaves like a result in mathematics or any of the hard natural sciences. Some judgement has to be applied when deciding whether the computational theory for a problem has been formulated adequately; the state- ment "take the opponent's king" defines the goal of chess, but it is hardly an adequate characterization of the computational problem of doing it.[1]

that which decides whether a result in mathematics amounts to a substantial new theorem, and I do not feel uncomfortable about having to leave the basis of such judgements unspecified.[2] This view of what constitutes a result in A.I. is probably acceptable to most scientists. Chomsky's [2] notion of a "competence" theory for English syntax is precisely what I mean by a computational theory for that problem. Both have the quality of being little concerned with the gory details of algorithms that must be run to express the competence (i.e. to implement the computation). That is not to say that devising suitable algorithms will be easy, but it is to say that before one can devise them, one has to know what exactly it is that they are supposed to be doing, and this information is captured by the computational theory. When a problem decomposes in this way, I shall refer to it as having a Type 1 theory. The fly in the ointment is that while many problems of biological information processing have a Type 1 theory, there is no reason why they should all have. This can happen when a problem is solved by the simultaneous action of a con- siderable number of processes, whose interaction is its own simplest description, One computational theory that in principle can solve chess is exhaustive search. The real interest lies however in formulating the pieces of computation that we apply to the game. One presumably wants a computational theory that has a rather general application, together with a demonstration that it happens to be applicable to some class of games of chess, and evidence that we play games in this class. and I shall refer to such a situation as a Type 2 theory.[3] One promising candidate for a Type 2 theory is the problem of predicting how a protein will fold. A large number of

[1] The kind of judgement that is needed seems to be rather similar to One computational theory that in principle can solve chess is exhaustive search. The real interest lies however in formulating the pieces of computation that we apply to the game. One presumably wants a computational theory that has a rather general application, together with a demonstration that it happens to be applicable to some class of games of chess, and evidence that we play games in this class.

[2] New algorithms for implementing a known computational theory may subsequently be devised without throwing substantial new light upon the theory, just as Winograd's [31] Very Fast Fourier Transform shed no new light on the nature of Fourier analysis.

[3] The underlying point here is that there is often a natural modularity in physics (e.g. under normal conditions, electrical interactions are independent of gravitational interactions), but some

influences act on a large polypeptide chain as it flaps and flails in a medium. At each moment only a few of the possible interactions will be important but the importance of those few is decisive. Attempts to construct a simplified theory must ignore some interactions; but if most interactions are crucial at some stage during the folding, a simplified theory will prove inadequate. Interestingly, the most promising studies of protein folding are currently those that take a brute n force approach, setting up a rather detailed model of the amino acids, the geometry associated with their sequence, hydrophobic interactions with the circumambient fluid, random thermal perturbations etc., and letting the whole set of processes run until a stable configuration is achieved (Levitt and Warshel [8]).

The principal difficulty in A.I. is that one can never be quite sure whether a problem has a Type I theory. If one is found, well and good; but failure to find J one does not mean that it does not exist. Most A.I. programs have hitherto e amounted to Type 2 theories, and the danger with such theories is that they can - bury crucial decisions, that in the end provide the key to the correct Type 1 decomposition of the problem, beneath the mound of small administrative decisions i that are inevitable whenever a concrete program is designed. This phenomenon 5 makes research in A.I. difficult to pursue and difficult to judge. If one shows that a given information processing problem is solved by a particular, neatly circumscribed computational theory, then that is a secure result. If on the other hand one produces a large and clumsy set of processes that solves a problem, one cannot always be sure that there isn't a simple underlying computational theory for one or more related problems whose formulation has somehow been lost in the fog. With any candidate for a Type 2 theory, much greater importance is attached to the performance of the program. Since its only possible virtue might be that it works, it is interesting only if it does. Often, a piece of A.I. research has resulted in a large program without much of a theory, which commits it to a Type 2 result, - but that program either performs too poorly to be impressive or (worse still) has not even been implemented. Such pieces of research have to be judged very harshly, because their lasting contribution is negligible.

Thus we see that as A.I. pursues its study of information processing problems, two types of solution are liable to emerge. In one, there is a clean underlying theory in the traditional sense. Examples of this from vision are Horn's [5] method for obtaining shape from shading, the notion of the primal sketch as a representation of the intensity changes and local geometry of an image (Marr [10]),Ullman's [26] method for detecting light sources, Binford's [1] generalized cylinder

representation, on which Marr and Nishihara's [12] theory of the internal representation and manipulation of 3-D structures was based, a recent theory of stereo

processes involve several at the same time and with roughly equal importance, like protein-folding. Thus the Type 1-Type 2 distinction is not a pure dichotomy, and there is a spectrum of possibilities between them.

vision (Marr [9], Marr and Poggio [13][4] and Poggio and Reichardt's [17] analysis of the visual orienting behaviour of the housefly. One characteristic of these results is that they often lie at a relatively low level in the overall canvas of intellectual functions, a level often dismissed with contempt by those who purport to study "higher, more central" problems of intelligence. Our reply to such criticism is that low-level problems probably do represent the easier kind, but that is precisely the reason for studying them first. When we have solved a few more, the questions that arise in studying the deeper ones will be clearer to us. But even relatively clean Type 1 theories such as these involve Type 2 theories as well. For example, Marr and Nishihara's 3-D representation theory asserts that the deep underlying structure is based on a distributed, object-centred coordinate system that can be thought of as a stick figure, and that this representation is explicitly manipulated during the analysis of an image. Such a theory would be little more than speculation unless it could also be shown that such a description may be computed from an image and can be manipulated in the required way. To do so involves several intermediate theories, for some of which there is hope of eventual Type 1 status, but others look intractably of Type 2. For example, a Type 1 theory now exists for part of the problem of determining the appropriate local coordinate system from the contours formed in an object's image (Marr [11]), but it may be impossible to derive a Type 1 theory for the basic grouping processes that operate on the primal sketch to help separate figure from ground. The figure- ground "problem" may not be a single problem, being instead a mixture of several subproblems which combine to achieve figural separation, just as the dif- ferent molecular interactions combine to cause a protein to fold. There is in fact no reason why a solution to the figure-ground problem should be derivable from a single underlying theory. The reason is that it needs to contain a procedural representation of many facts about images that derive ultimately via evolution from the cohesion and continuity of matter in the physical world. Many kinds of knowledge and different techniques are involved; one just has to sort them out one by one. As each is added the performance of the whole improves, and the complexity of the images that can be handled increases. We have already seen that to search for a Type 2 theory for a problem may be dangerous if in fact it has a Type 1 theory. This danger is most acute in pre- mature assaults on a high-level problem, for which few or none of the concepts that underlie its eventual Type 1 theory have yet been developed, and the con- sequence is a complete failure to formulate correctly the problems that are in fact involved. But it is equally important to realize that the opposite danger exists lower down. For example, in our current theory of visual processing, the notion of the primal sketch seems respectable enough, but one might have doubts about the aesthetics of the grouping processes that decode it. There

[4]The notion of cooperative computation, or relaxation labeling (Zucker[331), is a notion at the algorithm level. It suggests a way of implementing certain computations, but does not address th problem of what should be implemented, which seems to be the real issue for vision no less than elsewhere.

are many of them, their details are somewhat messy; and seemingly arbitrary preferences occur (e.g. for vertical or horizontal organizations). A clear example of a Type 2 theory is our assertion that texture-vision discriminations rest on these grouping processes and first-order discriminations applied to the information held in primal sketch of the image (Marr [10]). As such, it is less attractive than Julesz's [7] clean (Type 1) theory that textured regions are discriminable only if there is a difference in the first or second-order statistics of their intensity arrays. But as Julesz himself found, there exist patterns with different second-order statistics that are nevertheless indiscriminable; and one can in fact view our own work as attempting to define precisely what characteristics of the second-order statistical structure cause dis- criminability (see Schatz, 1977, in preparation). This inevitably forces us to relinquish the beauty of Julesz's concise theory, but I feel that one should not be too distressed by the need at this level of investigation to explore rather messy and untidy details. We already know that separate modules must exist for computing other aspects of visual information-motion, stereoscopy, fluorescence, color and there is no reason why they should all be based on a single theory. Indeed one would a priori expect the opposite; as evolution progressed, new modules came into existence that could cope with yet more aspects of the data, and as a result kept the animal alive in ever more widely ranging circumstances. The only important constraint is that the system as a whole should be roughly modular, so that new facilities can be added easily. So, especially at the more peripheral stages of sensory information processing, and perhaps also more centrally, one should not necessarily give up if one fails to find a Type 1 theory there may not be one. More importantly even if there were, there would be no reason why that theory should bear much relation to the theory of more central phenomena. In vision for example, the theory that says 3-D representations are based on stick-figure coordinate systems and shows how to man ipulate them, is independent of the theory of the primal sketch, or for that matter of most other stages en route from the image to that representation. In particular, it is especially dangerous to suppose that an approximate theory of a p peripheral process has any significance for higher level operations. For example, because Julesz's second-order statistics idea is so clean and so neatly fits much data, one m ight be tempted to ask whether the idea of second order interactions is in some way central to higher processes. In doing so one should bear in mind that the true explanation of visual texture discrimination may be quite different in nature even if the theory is very often a correct predictor of performance. The reason for drawing this point out at such length is that it bears upon another issue, namely the type of theory that the grammar of natural language might have. The purpose of human language is presumably to transform a datastructure that is not inherently one-dimensional into one-dimensional form for transmission as a sequential utterance, thereafter to be retranslated into some rough copy of the original in the head of the listener. Viewed in this light, it becomes entirely possible that there may exist no Type 1 theory of English syntax of the type that trans- formational grammar attempts to define-that its constraints resemble wired-in conventions about useful ways of executing this tedious but vital operation, rather than deep principles about the nature of intelligence. An abstract theory of syntax may be an illusion, approximating what really happens only in the sense that Julesz's second order statistics theory approximates the behaviour of the set of processes that implement texture vision and which, in the final analysis, are all the theory that there is. In other words, the grammar of natural language may have a theory of Type 2 rather than of Type 1. Even if a biological information processing problem has only a Type 2 theory, it may still be possible to infer more from a solution to it than the solution itself. This comes about because at some point in the implementation of a set of processes, design constraints attached to the machine in which they will run start to affect the structure of the implementation. This observation adds a different perspective to the two types of research carried out by linguists and by members of the artificial intelligence community. If the theory of syntax is really of Type 2, then any impor- tant implications about the CNS are likely to come from details of the way in which its constituent processes are implemented, and these are often explorable only by implementing them.

Implications of this view

If one accepts this view of A.I. research, one is led to judge its achievements according to rather clear criteria. What information processing problem has been isolated? Has a clean theory been developed for solving it, and if so how good are the arguments that support it? If no clean theory has been given what is the evidence that favors a set-of-processes solution or suggests that no single clean theory exists for it, and how well does the proposed set of mechanisms work? For very advanced problems like story-understanding, current research is often purely exploratory. That is to say, in these areas our knowledge is so poor that we cannot even begin to formulate the appropriate questions, let alone solve them. it is important to realize that this is an inevitable phase of any human endeavor, personally risky (almost surely no exploring pioneer will himself succeed in finding a useful question), but a necessary precursor of eventual success. Most of the history of Al. (now fully 16 years old) has consisted of exploratory studies. Some of the best-known are Slagle's [24] symbolic integration program, Weizenbaum's [30] Eliza program, Evans' [4] analogy program, Raphael's [19] SIR, Quillian's [18] semantic nets and Winograd s [32] Shrdlu. All of these programs have (in retrospect) the property that they are either too simple to be interesting Type 1 theories, or very complex yet perform too poorly to be taken seriously as a Type 2 theory. Perhaps the only really successful Type 2 theory to emerge in the early phase of A.I. was Waltz's 27] program. And yet many things have been learnt from these experiences mostly negative things (the first

20 obvious ideas about how intelligence might work are too simple or wrong) but including several positive things. The MACSYMA algebraic manipulation system (Moses [1 5]) is undeniably successful and useful, and it had its roots in programs l!k Slagle's. The mistakes made in the field lay not in having carried out such 5tudies-they formed an essential part of its development-but consisted mainly in failures of judgement about their value, since it is now clear that few of the early studies themselves formulated any solvable problems. Part of the reason for these internal failures of judgement lay in external pressures for early results from the field, but this is not the place to discuss what in the end are political matters. Yet, I submit, one would err to judge these failures of judgement too harshly. They are merely the inevitable consequence of a necessary enthusiasm, based on a view of the long-term importance of the field that seems to me correct. All impor- tant fields of human endeavor start with a personal commitment based on faith rather than on results. A.I. is just one more example. Only a sour, crabbed and unadventurous spirit will hold it against us.

Current trends

Exploratory studies are important. Many people in the field expect that, deep in the heart of our understanding of intelligence, there will lie at least one and probably several important principles about how to organize and represent knowledge that in some sense captures what is important about the general nature of our intellectual abilities. An optimist might see the glimmer of such principles in programs like those of Sussman and Stallman [25], of Marr and Nishihara [12], in the overall attitudes to central problems set out by Minsky [14], and possibly in some of Schank's [21, 22] work, although I sometimes feel that he failed to draw out the important points. While still somewhat cloudy, the ideas that seem to be emerging (and which owe much to the early exploratory studies) are:

(1) That the "chunks" of reasoning, language, memory, and perception ought to be larger than most recent theories in psychology have allowed (Minsky [14]). They must also be very flexible at least as flexible as Marr and Nishihara's stick-figure 3-D models, and probably more. Straightforward mechanisms that are suggested by the terms "frame" and "terminal", are certainly too inflexible.

(2) That the perception of an event or of an object must include the simultaneous computation of several different descriptions of it, that capture diverse aspects of the use, purpose or circumstances of the event or object.

(3) That the various descriptions described in (2) include coarse versions as well as fine ones. These coarse descriptions are a vital link in choosing the appropriate overall scenarios demanded by (1), and in establishing correctly the roles played by the objects and actions that caused those scenarios to be chosen.

An example will help to make these points clear. If one reads
(A) The fly buzzed irritatingly on the window-pane.
(B) John picked up the newspaper.

the immediate inference is that John's intentions towards the fly are fundamentally malicious. If he had picked up the telephone, the inference would be less secure. It is generally agreed that an "insect-damaging" scenario is somehow deployed during the reading of these sentences, being suggested in its coarsest form by the fly buzzing irritatingly. Such a scenario will contain a reference to something that can squash an insect on a brittle surface a description which fits a newspaper but not a telephone. We might therefore conclude that when the newspaper is mentioned (or in the case of vision, seen) not only is it described internally as a newspaper, and some rough 3-D description of its shape and axes set up, but it is also described as a light, flexible object with area. Because sentence (B) might have continued "and sat down to read", the newspaper may also be being described as reading-matter; similarly, as a combustible article, and so forth. Since one does not usually know in advance what aspect of an object or action is important, it follows that most of the time, a given object will give rise to several different coarse internal descriptions. Similarly for actions. It may be important to note that the description of fly-swatting or reading or fire-lighting does not have to be attached to the news- paper; merely that a description of the newspaper is available that will match its role in each scenario.

The important thing about Schank's "primitive actions" seems to me not the fact that there happens to be a certain small number of them, nor the idea that every act is expressed solely by reduction to them (which I cannot believe at all), nor even the idea that the scenarios to which they are attached contain all the answers for the present situation (that is where the missing flexibility comes in). The importance of a primitive, coarse catalogue of events and objects lies in the role such coarse descriptions play in the ultimate access and construction of perhaps exquisitely tailored specific scenarios, rather in the way that a general 3-D animal model in Marr and Nishihara's theory can finish up as a very specific Cheshire Cat, after due interaction between the image and information stored in the primitive model. What after sentence (A) existed as little more than a malicious intent towards the innocent fly becomes, with the additional information about the newspaper, a very specific case of fly-squashing.

Marr and Nishihara have labelled the problem of providing multiple-description s for the newspaper its 'reference-window problem". Exactly how it is best done and exactly what descriptions should accompany different words or perceived objects, is not yet known. These insights are the result of exploratory studies, and the problems to which they lead have yet to be precisely formulated, let alone satisfactorily solved. But it is now certain that some problems of this kind do exist and are important; and it seems likely that a fairly respectable theory of them will eventually emerge.

Mimicry versus exploration

Finally, I would like to draw one more distinction that seems to be important when choosing a research problem, or when judging the value of completed work. The problem is that studies-particularly of natural language understanding, problem-solving, or the structure of memory can easily degenerate into the writing of programs that do no more than mimic in an unenlightening way some small aspect of human performance. Weizenbaum [30] now judges his program Eliza to belong to this category, and I have never seen any reason to disagree. More controversially, I would also criticize on the same grounds Newell and Simon's work on production systems, and some of Norman and Rummelhart's [16] work on long term memory.

The reason is this. If one believes that the aim of information-processing studies is to formulate and understand particular information-processing problems, then it is the structure of those problems that is central, not the mechanisms through which they are implemented. Therefore, the first thing to do is to find problems that we can solve well, find out how to solve them, and examine our performance in the light of that understanding. The most fruitful source of such problems is operations that we perform well, fluently (and hence unconsciously) since it is difficult to see how reliability could be achieved if there were no sound underlying method. On the other hand, problem-solving research has tended to concentrate on problems that we understand well intellectually but perform poorly on, like mental arithmetic andcriptarithmatic or on problems like geometry theorem-proving, or games like chess, in which human skills seem to rest on a huge base of knowledge and expertise. I argue that these are exceptionally good grounds for not studying how we carry out such tasks yet. I have no doubt that when we do mental arithmetic we are doing something well, but it is not arithmetic, and we seem far from under- standing even one component of what that something is. Let us therefore con- centrate on the simpler problems first, for there we have some hope of genuine advancement.

If one ignores this stricture, one is left in the end with unlikely looking mechanisms whose only recommendation is that they cannot do something we cannot do. Production systems seem to me to fit this description quite well. Even taken on their own terms as mechanisms, they leave a lot to be desired. As a programming language they are poorly designed, and hard to use, and I cannot believe that the human brain could possibly be burdened with such poor implementation decisions at so basic a level.

A parallel may perhaps be drawn between production systems for students of problem-solving, and Fourier analysis for visual neurophysiologists. Simple operations on a spatial frequency representation of an image can mimic several interesting visual phenomena that seem to be exhibited by our visual systems. These include the detection of repetition, certain visual illusions, the notion of separate linearly adding channels, separation of overall shape from fine local detail, and a simple expression of size invariance. The reason why the spatial frequency domain is ignored by image analysts is that it is virtually useless for the main job of vision-building up a description of what is there from the intensity array. The intuition that visual physiologists lack, and which is so important, is for how this may be done. A production system exhibits several interesting ideas- the absence of explicit subroutine calls, a blackboard communication channel, and some notion of a short term memory. But just because production systems display these side-effects (as Fourier analysis "displays" some visual illusions) does not mean that they have anything to do with what is really going on. My own guess would be, for example, that the fact that short-term memory can act as a storage register is probably the least important of its functions. I expect that there are several "intellectual reflexes" that operate on items held there, about which nothing is yet known, and which will eventually be held to be the crucial things about it because they perform central functions like opening up an item's reference window. Studying our performance in close relation to production systems seems to me a waste of time, because it amounts to studying a mechanism not a problem, and can therefore lead to no Type 1 results. The mechanisms that such research is trying to penetrate will be unravelled by studying problems, just as vision research is progressing because it is the problem of vision that is being attacked, not neural visual mechanisms.

A reflexion of the same criticism can be made of Norman and Rummelharts work, where they studied the way information seems to be organized in long term memory. Again, the danger is that questions are not asked in relation to a clear information-processing problem. Instead, they are asked and answers proposed in terms of a mechanism-in this case, it is called an "active structural network" and it is so simple and general as to be devoid of theoretical substance. They may be able to say that such and such an "association" seems to exist, but they cannot say of what the association consists, nor that it has to be so because to solve problem X (which we can solve) you need a memory organized in such-and- such a way; and that if one has it, certain apparent "associations" occur as side- effects. Experimental psychology can do a valuable job in discovering facts that need explaining, including those about long-term memory, and the work of Shepard [23], Rosch [20] and of Warrington [28] (for example) seems to me very successful at this; but like experimental neurophysiology, experimental psychology will not be able to explain those facts unless information-processing research has identified and solved the appropriate problems X.[5] It seems to me that finding such problems X, and solving them. is what A.I. should be trying to do.

[5]In the present state of the art, it seems wisest to concentrate on problems that probably have Type I solutions, rather than on those that are almost certainly of Type 2.

ACKNOWLEDGEMENT

Although I take full responsibility for the purely personal views set out here, any virtues that they may have are due in part to many conversations with Drew McDermott. This report describe5 work done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology.W

Support for the Laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract number NOOO 1 4-75-C-0643.

REFERENCES

1. Binford, T. 0., Visual perception by computer, IEEE Conf. Systems and Control, Miami (December 1971).
2. Chomsky, A. N., Aspects of the Theory of Syntax (M.I.T. Press, Cambridge, MA., 1965).
3. Cooley, J. M. and Tukey, J. W., An algorithm for the machine computation of complex Fourier series, Math. Comp. 19 (1965) 297-301.
4. Evans, T., A program for the solution of geometric-analogy intelligence test questions, in: M. Minsky, Ed., Semantic hiformation Processing (M.I.T. Press, Cambridge, MA., 1968) 271-353.
5. Horn, B. K. P., Obtaining shape from shading information, in: P. H. Winston, Ed., The Psychology of Computer Vision (McGraw-Hill, NY, 1975)115-155.
6. Jardine, N. and Sibson, R., Mathematical Taxonomy (Wiley, NY, 1971). 7. Julesz, B., Experiments in the visual perception of texture, Scientific American 232 (April 1975) 3443.
8. Levitt, M. and Warshel, A., Computer simulation of protein folding, Nature 253 (1975) 694 698.
9. Marr, D., A note on the computation of binocular disparity in a symbolic, low-level visual processor, M.I.T. A.I. Lab. Memo 327(1974).
10.Marr, D., Early processing of visual information, Phil. Trans. Roy. Soc. B. 275(1976)483-524. 1
11.Marr, D., Analysis of occluding contour, Proc. Roy. Soc. B., in the press.
12.Marr, D. and Nishihara, H. K., Representation and recognition of the spatial organization of three dimensional shapes, Royal Soc., submitted for publication.
13.Marr, D. and Poggio, T., Cooperative computation of stereo disparity, Science 194 (1976) 283-287.
14.Minsky, M., A framework for representing knowledge, in: P. H. Winston, Ed., The Psychology of Computer Vision (McGraw-Hill, NY, 1975) 211-277.
15.Moses, J., MACSYMA-the fifth year, SIGSAM Bull., ACM, 8 (1974)105-110.
16.Norman, D. A. and Rumelhart, D. F., Explorations in cognition (W. H. Freeman and Co., San Francisco, 1974) e.g. The active structural network, 35-64.

17.Poggio, T. and Reichardt, W., Visual control of the orientation behaviour of the fly: towards the underlying neural interactions. Quarterly Reviews of Biophysics 9 (1976) 377–t38.
18.Quillian, M. R., Semantic memory, in: M. Minsky, Ed., Semantic Informatioii Processing (M.I.T. Press, Cambridge, MA, 1968) 227-270.
19.Raphael, B., SIR: semantic information retrieval, in: M. Minsky, Ed., Seniantic Jiformation Processing (M.I.T. Press, Cambridge, MA, 1968) 33-145.
20.Rosch, E., Classifications of real-world objects: origins and representations in cognition, Bulletin de Psychologie, in press.
21.Schank, R. C., Identification of conceptualizations underlying natural language, in: R. C. Schank and K. M. Colby, Eds., Computer Models of Thought and Language (W. H. Freeman, San Francisco, 1973).
22.Schank, R. C., Conceptual Information Processing (North-Holland, Amsterdam, 1975).
23.Shepard, RN., Form, formation, and transformation of internal representations, in: R. Solso, Ed., Information Processing and Cognition: The Loyola Symposium (Lawrence Eribaum Assoc., Hillsdale NJ, 1975) 87-122.
24.Slagle, J. R., A heuristic program that solves symbolic integration problems in freshman calculus, in: E. A. Feigenbaum and J. Feldman, Eds., Computeis and Thought (McGraw-Hill, NY, 1963)191-203.
25.Sussman, G. J. and Staliman, R. M., Heuristic techniques in computer-aided circuit analysis, IEEE Transactions on Circuits and Systems, CAS-22 (1975) 857-865.
26.UlIman, S., On visual detection of light sources, Biol. Cybernetics 21 (1976) 205-212.
27.Waltz, D. L., Understanding line drawings of scenes with shadows, in: P. H. Winston, Ed., The Psychology of Computer Vision (McGraw-Hill, NY, 1975)19-91.
28.Warrington, F. K., The selective impairment of semantic memory, Q. J. Exp. Psychol. 27 (1975) 635-657.
29.Weizenbaum, J., ELIZA-a computer program for the study of natural language communication between man and machine. Commun. ACM 9 (1965) 36-15.
30.Weizenbaum, J., Computer Thought and Human Reason (W. H. Freeman, San Francisco, 1976).
31.Winograd, S., Computing the discrete Fourier transform, Proc. Nat. Acad. Sci. 73K1976) 1005- 1006.
32.Winograd, T., Understanding Natural Language (Academic Press, NY, 1972).
33.Zucker, S. W., Relaxation labelling and the reduction of local ambiguities, University of Maryland Computer Science TR-451 (1976).

Received June 1976; revised version received February 1977