AD-A131 359 INTERDISCIPLINARY STUDY ON ARTIFICIAL INTELLIGENCE(U) 1/1
MIAMI UNIV CORAL GABLES FLA CENTER FOR THEORETICAL
STUDIES B N KURSUNOGLU JUL 83 ARO-20126.1-AA
UNCLASSIFIED DAAG29-82-K-0198 F/G 6/4 NL

END
FILMED
DTIC

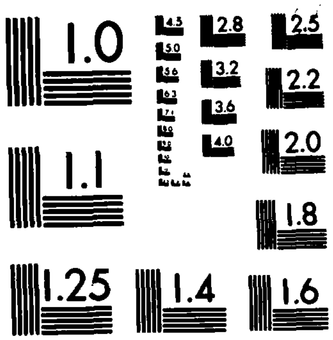

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: 20126.1-~~MA~~ AA
2. GOVT ACCESSION NO.: AD-A131359
3. RECIPIENT'S CATALOG NUMBER:

4. TITLE (and Subtitle): Interdisciplinary Study on Artificial Intelligence "Workshop on Biological Dimensions of Artificial Intelligence"
5. TYPE OF REPORT & PERIOD COVERED: Final: 15 Oct 82 - 14 Oct 83
6. PERFORMING ORG. REPORT NUMBER:

7. AUTHOR(s): Behram N. Kursunoglu
8. CONTRACT OR GRANT NUMBER(s): DAAG29 82 K 0198

9. PERFORMING ORGANIZATION NAME AND ADDRESS: University of Miami, Coral Gables, FL 33124
10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS:

11. CONTROLLING OFFICE NAME AND ADDRESS: U. S. Army Research Office, Post Office Box 12211, Research Triangle Park, NC 27709
12. REPORT DATE: Jul 83
13. NUMBER OF PAGES: 45

14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office):
15. SECURITY CLASS. (of this report): Unclassified
15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

DTIC ELECTE AUG 16 1983 S D

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

D

18. SUPPLEMENTARY NOTES

The view, opinions, and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

artificial intelligence
biology
computerized simulation
brain
neurophysiology
biophysics

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

The interdisciplinary workshop on biological dimensions of artificial intelligence was organized with a very special objective in mind. The objective was to bring together researchers working in a variety of areas directly concerned with intelligence, such as computer modeling of brain processes, experimental neurophysiology, evolutionary programming and adaptability theory, theory modeling and simulation, self-organizing systems, biophysics of information processing, cognitive science, and traditional artificial intelligence. The objective behind this objective was to provide a vehicle for reviewing and analyzing directions of

DD FORM 1473, 1 JAN 73 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

83 08 15 027

ABSTRACT (cont.)

artificial intelligence from the perspective of the full range of scholarly activities relevant to this field. Some of the specifically stated objectives in the original letter of invitation suggested topics such as learning and adaptation, evolutionary algorithms for adaptive pattern recognition and motor control, the comparison of computer and biological organization, knowledge representation and the comparison of biological and computer memory, the potential role of parallelism and the physical limits of computation, and the significance of recent experimental work on biochemical and molecular switching processes inside neurons. At the outset of the metting, it was decided that it would not be desirable to try to arrive at any group conclusion or recommendation. We felt that the most objective and certainly the most useful way to develop a report was for each member to write up an individual recommendation. The body of this report consists of these individual recommendations along with notes on the discussion.

Final Report

INTERDISCIPLINARY STUDY ON ARTIFICIAL INTELLIGENCE

"Workshop on Biological Dimensions of Artificial Intelligence"

held

March 14-25, 1983

at the Center for Theoretical Studies

University of Miami

Coral Gables, Florida 33124

Submitted to

U.S. Army Research Office

Research Triangle Park, North Carolina 27709

by

BEHRAM N. KURSUNOGLU
Principal Investigator

July, 1983

Final Report

to the Army Research Office

INTERDISCIPLINARY STUDY ON ARTIFICIAL INTELLIGENCE

Workshop on Biological Dimensions of Artificial Intelligence

held March 14-23, 1983

The interdisciplinary workshop on biological dimensions of artificial intelligence was organized with a very special objective in mind. The objective was to bring together researchers working in a variety of areas directly concerned with intelligence, such as computer modeling of brain processes, experimental neurophysiology, evolutionary programming and adaptability theory, theory modeling and simulation, self-organizing systems, biophysics of information processing, cognitive science, and traditional artificial intelligence. The objective behind this objective was to provide a vehicle for reviewing and analyzing directions of artificial intelligence from the perspective of the full range of scholarly activities relevant to this field. Some of the specifically stated objectives in the original letter of invitation suggested topics such as learning and adaptation, evolutionary algorithms for adaptive pattern recognition and motor control, the comparison of computer and biological organization, knowledge representation and the comparison of biological and computer memory, the potential role of parallelism and the physical limits of computation, and the significance of recent experimental work on biochemical and molecular switching processes inside neurons.

At the outset of the meeting, we decided that it would not be desirable to try to arrive at any group conclusion or recommendation. We felt that the most objective and certainly the most useful way to develop a report was for each member to write up an individual recommendation. The body of this report consists of these individual recommendations along with notes on the discussion.

At various points in the meeting (after the first round of presentations), each member of the group offered an analysis of the issues under discussion and suggested critical problems which ought to be discussed further or ought to be identified as especially significant research problems. The problems, generally speaking, fell into three categories: biochemical and neuronal mechanisms in the brain, principles of self-organization, and high-level artificial intelligence. Much of the discussion concerned the interrelations of these levels of description, including the use of computer modeling and simulation techniques to study these interrelations.

The bottom-up question concerns the functional significance of the microscopic dynamics of the nervous system. Clearly, it is important to consider the contributions of these dynamics to self-organization and high-level intelligence. This is the first step to harnessing microscopic principles for future technologies. Similarly, principles of self-organization and evolution can only be properly understood against the background of their molecular and physiological substrate, on the one hand, and against the foreground of their contribution to system function, on the other hand. Again, this is the prerequisite for enhancing the contribution of these principles to technology. High-level AI, in large measure, has been pursued in isolation from neurophysiological and evolutionary principles. The extent to which AI ought to contemplate these layers of reality is the question most critical to assessment of current research priorities in this area. A variety of views was expressed, ranging from the view that the computing power of present-day computers is adequate for the problem being addressed by AI to the view that wholly new designs and learning principles can be developed which will both enhance AI and provide valuable insights into those problems which can and cannot be solved. The recommendations on this issue are best left to the individual statements of the participants, though it is probably fair to say the participants in this workshop were strongly inclined to the view that the integration of microscopic, evolutionary, and algorithmic approaches should be identified as an important goal. That the problems in this area are ripe and that the possibility for discovery and technological advance is high are senses which hung in the air at this workshop.

The field is new, and it will evolve with an interdisciplinary participation by all sciences, especially chemistry, physics, and biology. Most people believe that the ultimate computer may be one whose miniaturized electronic elements are molecules that can even assemble themselves. Thus, we may be talking of a molecular computer. In an old lecture, Richard Feynman stated that, "ordinary machines could build smaller machines that could build still smaller machines, working step-by-step down toward the molecular level." The molecules act as switches, inserted in long, chain-like molecule "wires" that conduct electrons or vibrational pulses (solitons). Solitons, mathematically, are solutions of nonlinear wave equations and are solitary waves that can travel unchanged for long distances. Some people believe that solitons can propagate along chain-like molecules, of which there are many examples. For example, these chain-like molecular wires can conduct, in principle, electrons or vibrational pulses. Thus, electrons or solitons injected into these chains will propagate along the molecule. Combining special purpose molecules with molecular wires can, in theory, produce most of the necessary circuit elements for a computer.

BIOLOGICAL DIMENSIONS OF ARTIFICIAL INTELLIGENCE*

Michael Conrad
Department of Computer Science
Wayne State University
Detroit, MI 48202

* Recommendation to the Army Office of Research based on the Center for Theoretical Studies Conference on Biological Dimensions of Artificial Intelligence, held at the University of Miami, March 14-25, 1983.

Artificial intelligence is sometimes defined as the use of existing computer technology to perform tasks which are usually considered to require human intelligence. This type of work has been heavily supported in the past, with some fruitful practical applications. Continued support is certainly warranted. However, there are a number of dimensions of the subject which have not been adequately cultivated. The main one is the area of computer learning and evolution. Early workers in AI did consider these approaches. But the construction of high-level algorithms tailored for the solution of specific problems appeared to promise more rapid progress. This is indeed the case when the environment in which the program must act is rigidly defined (as in reading universal price codes in a supermarket). It is now painfully evident that many of the most important problems facing AI (such as object recognition and natural language processing) cannot be artificially constrained in this way.

These limitations are at least in part due to the absence of learning capacity and intrinsic adaptability in AI programs. One might seek to remedy this situation by developing further high-level algorithms for learning or by seeking better ways to parameterize large classes of problems. We have taken a third approach which we believe can complement these two. The strategy is to incorporate representations of structural features which enhance the learning and evolutionary capabilities of natural systems.

Here are a series of problems which are critical from this point of view.

1. Development of evolutionary learning algorithms

Evolutionary algorithms have been used over the years for a variety of problems. In general, evolutionary learning is effective only when the system has a gradualism property. By incorporating various redundancies and dynamical features which are present in biological systems, it should be possible to enlarge the class of tasks which can be learned through evolution. Problems which are plausibly addressed (and on which we have concentrated) include adaptive pattern recognition and adaptive motor control.

2. Development of hardware suitable for evolutionary learning

Representing evolution-enhancing features is a computational cost. But the cost can be reduced with appropriate hardware. We are currently experimenting with algorithms which utilize simulated dynamical features. Current technology could be utilized to incorporate these dynamics in an effective manner. Evolutionary algorithms exploit parallelism in a highly efficient way, and the configuring of highly parallel learning systems should be a priority area.

3. Development of new computing primitives using biotechnology

Biological systems are built by evolution for effective evolution and learning. Their primitive switching components, the enzymes, have an intrinsic gradualism property. This is due to the fact that they are tactile pattern recognizers. Their switching function is determined by their three-dimensional shape, which emerges through a dynamic folding processing. A single change in the order of amino acids is thus not unlikely to lead to a slight change in switching function. This type of gradualism property occurs at all levels of biological organization, though, in some cases, the mechanisms are different.

Recent advances in organic chemistry and gene technology now make it feasible to contemplate computer components built out of organic and biological materials. Many of the people working in this field wish to duplicate the silicon chip in more minute form using carbon chemistry. But the real opportunity afforded by these new technologies is to build switching components which incorporate the same type of evolution enhancing features as do the switching elements of biological systems. Harnessing these technologies for computing is further off on the horizon than new silicon designs. But the enormous efforts currently being put into gene technologies make it highly likely that molecular computers will eventually be commercially feasible. I would envisage these as bing speical purpose computing devices which can be adapted by evolutionary mechanisms to fill computing niches which are not effectively filled by present day machines as we know them.

4. Distributed memory systems

The structure of the brain is incompatible with an addressable memory systme of the type used in present day computers. Memory is sometimes identified with the reconstruction of patterns of neural activity under appropriate stimuli. Many models have been constructed. In the model we have been investigating, active reference neurons load all the active neurons they contact. Refiring an appropriate reference neuron can retrieve an arbitrarily complicated pattern of neural activity. The advantage is that it is possible for the brain to compute with arbitrarily complex data structures treated as single entities. This type of global memory system could be implemented with present day fabrication technology.

The study of this and similar memory systems with a view towards developing computers with memory capabilities more similiar to that of the brain is an area which could certainly bear fruit.

5. Brain models

It is now known that cyclic nucleotide mechanisms serve as a link between switching processes at the electrical layer of neural activity and switching processes at the molecular layer. The computing power of the brain is enormous by comparison to that in present day macroscopic computers, both due to the contribution of microphysical processes and to the high-potential parallelism. Our working hypothesis is that this enormous reservoir is captured through the interplay of two mechanisms. One is phylogenetic evolution and ontogenetic learning mechanisms of an evolutionary nature. The second is the reference neuron memory mechanism which allows evolved or learned computing mechanisms to be brought together in new combinations. Other researchers concentrating on different aspects of biological intelligence could justifiably emphasize other features, such as exotic dynamics or the embedding of unorthodox algorithms in brain tissue. The organization of the brain--its high parallelism, its distributed memory system, its unknown timing mechanisms, the relevance of microscopic dynamics, its structural nonprogrammability and amenability to evolution--is clearly quite different than that of our present day, programmable computers. But these magnificent artifacts and the algorithmic languages which have been developed to communicate with them provide the best and possibly the only means of formally modeling the processes of intelligence.

The situation can perhaps be summed up in what I will call the bootstrap conjecture. Present-day electronic computers used as a modeling tool (in cooperation with experiment) provide the most powerful means for studying the biological foundations of computing. Having grasped these foundations, we will be in a better position to build more powerful computing artifacts, or at least to understand the limits of artificial intelligence.

Some AI researchers consider that they are not simply using the computer to duplicate intelligent activities, but that they are, in some sense, studying the mind. This assumption is important, not only on scientific grounds, but on social grounds. Clearly, computers are being given increasing responsibilities in agriculture, economic and social planning, and in military spheres. Where should the computer be used, and at what points should we insist on human judgment and creativity? I would go so far as to say that, in an age of nuclear weapons, this issue assumes a national and, indeed, a world security dimension. It is fundamentally important to understand the similarities and differences between brain and computer. Someday we may exploit some of our new-found knowledge for new useful technologies. But no better service could be performed by such knowledge than to prevent us from misjudging the validity of current technologies.

References

M. Conrad (1983). Microscopic-macroscopic interface in biological information processing. Technical report CSC-83-003, Department of Computer Science, Wayne State University. (To appear in the Proceedings of the Orbis Scientiae Conference in honor of P.A.M. Dirac's eightieth year.) This paper provides a recent review of brain theory and evolutionary programming.

M. Conrad (1983). Adaptability: The Significance of Variability From Molecule to Ecosystem. Plenum Press, New York. This book includes a review of evolution enhancing features, compares computers and biological systems, and analyzes conditions under which computing and modeling either enhance or detract from the adaptability of an organization.

COMMENTS AND SUGGESTIONS ON ARTIFICIAL INTELLIGENCE AFTER THE CORAL GABLES WORKSHOP ON AI

Erich Harth
Department of Physics
Syracuse University
Syracuse, New York 13210

After the 10-day exposure to a broad range of topics dealing with AI, I offer the following comments and suggestions:

1. Our understanding of what are intelligent functions of the brain, and by what mechanisms they are carried out, is still at a rudimentary stage. Most neural mechanisms so far explored deal with undirectional processing exemplified by various types of reflexes, or the undirectional mappings observed in both sensory and motor systems. It has become clear that perception and other higher functions involve elaborate feedback mechanisms.

In the visual and auditory systems, strong corticofugal pathways are able to modify sensory stimuli. Thus, input and output, stimulus and response become integrated into one dynamic unit. Another departure from 'classical' neurodynamics must be sought in the strong coupling between the micro and macroscopic scales. No other physical mechanism shows such a continuous range of coupled systems. The separability of macrodynamics and microworld, which is fundamental in every man-made machine, no longer holds for the nervous system. Future studies on natural intelligence must take into account these novel aspects of the nervous system.

2. Computer models and computer simulation studies will continue to be the most useful in studying the synaptic properties of neural systems, since the strong non-linearities and the self-referent pathways make it impossible to make theoretical predictions of neural behavior. We must also be aware, of course, that such simulations are limited by the enormity of the state space to be explored, and by neural dynamics which may well involve the equivalent of strange attractors with resulting chaotic dynamics. In interpreting the results of such studies, increased attention must be paid to bringing together the dynamic and cognitive aspects of neural activity. This research must, therefore, draw on the best from many disciplines, including all aspects of neuroscience,

psychology and computer science.

3. The development of AI has always had a strong footing in the study of natural intelligence. I expect this relationship to continue. Thus, the widening of our understanding of the more sophisticated brain functions will certainly provide new ideas for AI in the future. It is less clear to what extent the particular neural _mechanisms_ accounting for these functions will be translatable into analogous hardware. In this connection, it will be imperative to pay close attention to the relative advantages and disadvantges of biological and engineering solutions to a particular problem. An uncritical _translation_ of a biological mechanism into an engineering analogue may be the most inefficient approach to a problem. I believe that more research should be devoted to this fundamental question of making the best use of the peculiarities and the versatility of components in the design of any system.

4. Finally, some remarks on the long-range outlook for the use of AI in solving human problems. Questions of ethics must ultimately reside in human judgment and cannot be relegated to any mechanism, however sophisticated. Thus, the anticipated superior reasoning power of AI must be backed by only the highest motivation, and appropriate safeguards against misuse must be taken. Given these precautions, the fundamentally amoral character of AI can become its greatest advantage. It may save us from bumbling into disaster.

POSITION PAPER ON COMPUTATIONAL MODELS OF INTELLIGENCE

John H. Holland
The University of Michigan

In approaching computational models of intelligence that involve learning, there are certain difficulties that, it seems to me, must be avoided if we are to make progress in genuinely flexible systems. The first of these arises when the approach requires global consistency of the knowledge base. The computational requirements for maintaining global consistency as experience accumulates are enormous, and such an approach is also unrealistic in the sense that there is little evidence that humans maintain consistency across large ranges of knowledge. (Indeed, there is substantial evidence to the contrary). The second difficulty, closely allied to the first, arises when prediction and plans require detailed proofs in some logical calculus. Again, the computational burdens are enormous, and it is almost impossible to handle ill-defined problems under such restrictions. It is noteworthy that humans rarely if ever produce detailed proofs, even in rigorous studies. Approaches to ill-defined problems rely on techniques such as models, analogies, and metaphor, using conventions and symmetries to bridge gaps. If the problem is complex, the modus operandi would appear to be a hierarchical approach, with progressive insertion of levels of detail until the overall structure is adequate or convincing relative to the ends for which it was constructed. The third kind of difficulty attends natural language approaches to intelligence that proceed in terms of symbols and syntax given a priori. Much can be learned in this way, but it is difficult to provide such systems with procedures for developing analogies, and the like, on the basis of experience. Without consideration of the ontogeny of symbols under a learning process, the problem is much like an attempt to construct a taxonomy of animal species with no consideration of ontogeny and phylogeny. It is possible, but of limited value in understanding the progressive adaptation of structure.

On the positive side, it seems to me important that much more research be done at the level of provision of learning and control procedures for flexible robots. This effort should be in addition to, not to the exclusion of, current research. It is critical that the robot be immersed in an environment (real or simulated) complex enough that the information supplied by the input interface faces the system with the perpetual novelty typical of real situation. The robot's effectors should be adequate for sophisticated behavioral sequences (motion, manipulation, redirection of vision, and the like), and attainment of some of the goals in the environment should require model-based lookahead and linked, contingent action-sequences. The central objective of the learning procedures should be the induction of goal-relevant models of the environment -- models retaining a wide range of hypotheses and analogies, subject to confirmation, and readily usable for generating plausible plans in new situations. For the reasons given above, the developing model, and the procedures attendant upon it, should be at a pre-linguistic level in the sense that the system requires no dictionary of terms, natural-langauge-like syntax, or high-level interpreter to generate its behavior.

It is noteworthy that, even at pre-linquistic levels, one can study the emergence of models, planning, understanding (in the sense of relevant performance), the effects of training, and most of the other issues currently of interest in Artificial Intelligence. In addition, as the internal model develops, one has the possibility of encouraging and studying the origin and development of symbol-oriented behavior (a la conditioning and gesture languages).

In pursuing the development of a cognitive system of this kind, there are several criteria that I think should serve as guideposts for the effort:

1. Building blocks. Knowledge structures should be constructed from a well-defined set of "building blocks" or elements. The elements may be rich and varied, but all allowable ways of combining elements to yield more complex structures should be clearly defined. This criterion serves two purposes. It is a powerful deterrent to ad hoc constructions. More importantly, structures constructed under this criterion can be compared in terms of the elements they hold in common. Such comparisons make possible the discovery of similarities and analogies, transfer of information from one situation to another, and other activities vital to flexible problem solving. The rich conbinatorial possibilities are essential for operation in complex environments.

2. Categorization. Categorization is the central device by which a cognitive system combats the perpetual novelty of the environment. Knowledge structures must be able to represent classes of environmental states that can be treated as equivalent for the purpose of achieving certain goals. Accordingly, there must be simple inductive prodecures to generate and test elements and combinations of elements that can represent categories.

3. Parallelism. In order to avoid postulating a distinct category representationfor each individual entity in the environment (e.g., a "red Saab with a flat tire" mode), it must be natural to represent categories more implicitly by the simultaneous activity of an array of elements. The capacity to evoke simultaneously a set of knowledge structures enables the cognitive system to deal with complex, novel situations.

4. Synchronic and Diachonic Relations. Knowledge structures must incorporate two basic types of relational information. Synchronic relations hold between alternative descriptions of environmental states, encompassing generalizations and associations of categories, default hierarchies, and other declarative knowledge. Diachronic relations represent temporal

transitions between representations of environmental states, making possible predictions of the consequences of actions and other inferences based on procedural knowledge.

5. Gracefulness. It should be simple to insert new, tentative inductions into the cognitive system without radically disrupting existing knowledge structures or useful, well-established behaviors. To this end, the elements of knowledge structures should act as a network of interacting and competing hypotheses. The competition replaces requirements for global consistency with a process of progressive, local (situation-contingent) confirmation.

6. Inductive Efficiency. Elementary considerations of limitations in processing capacity and storage rule out, for realistic cognitive systems, exhaustive or totally random procedures for constructing and modifying knowledge structures. Out of the enormous universe of potential knowledge structures, the inductive procedures must generate only a small subset of plausible structures. And of this subset, the system must retain only a yet smaller subset of actually useful structures. Plausible structures can be generated by both "bottom-up" (stimulus-guided) and "top-down" (category-guided) procedures. Examples are covariation detection (bottom-up procedures that identify regular co-occurences in the environment), and transfer of information (top-down procedures that exploit common elements in categories and knowledge structures).

7. Confirmation and Prediction-based Evaluation. The cognitive system must possess automatic procedures for evaluating the outcome of the competition between the tentatively held, plausible structures. As the system accumulates experience, the relative reliability and usefulness of structures must be readily adjustable, and the outcome must reflect itself directly in the competition. Structures must be modified, stored, or discarded on this basis. Such "feedback" may be covert;

for example, the persistent appearance of "atypical" exemplars of a category may trigger modification of the category. It is important that the system be able to assign credit to structures that act in the early, "stage-setting" portions of goal-attaining sequence (cf., an early move that makes possible a later triple jump in checkers). It is equally important that structures, via diachronic relations, have associated expectations or predictions that can be confirmed or disconfirmed on the basis of experience. This makes it possible for the system to act directly in reducing uncertainties in its models. It is also important that the system be able to assign credit to very general categories that label or "type" a situation, usefully focussing attention by "supporting" more detailed activities in the cognitive structure.

AN APPROACH TO AI BASED ON THE STUDY OF PRINCIPLES AND MECHANISMS UNDERLYING NATURAL INTELLIGENCE

Position Paper

Interdisciplinary Workshop on AI

Roberto R. Kampfner
Department of Computer Science
Wayne State University
Detroit, Michigan 48202

AN APPROACH TO AI BASED ON THE STUDY OF PRINCIPLES AND MECHANISMS UNDERLYING NATURAL INTELLIGENCE

Two closely related, fundamental issues in AI research are the precise characterization of what can be considered intelligent behavior and our lack of sufficient knowledge of the principles and mechanisms underlying natural intelligence. The clarification of these issues is, indeed, a long-term enterprise requiring the coordination of efforts and the integration of results from various domains of knowledge, from Physics, Biophysics and Biology, to Psychology, the Social Sciences and Philosophy. I think, however, that a realistic approach to AI should include the clarification of these issues as an important goal. The pursuit of such a goal would unify AI research and direct it towards the accomplishment of both short- and long-term objectives in an effective manner. Short-term objectives would tackle specific problems and problem areas. The difference with current approaches would be a greater emphasis in the application of principles and mechanisms of proven value in natural systems to problem domains in which they are known to be effective. Long-term objectives, on the other hand, would contemplate the gradual incorporation and integration of solutions to higher-level probelms in a consistent and effective manner. Moreover, moving in this direction would help elucidate the scope and limitation of AI, as well as to identify its most useful domains of applications.

The spirit of the proposed approach can be illustrated by work already done involving the study of adaptive pattern recognition and adaptive motor control processes.[1] In this investigation, principles and mechanisms postulated in the evolutionary selection circuits model[2] were analyzed using computational modeling and simulation. This study, however, represents just the beginning of a promising series of experiments inspired in biological information processing principles.[3]

In a highly speculative way, but one which I personally find extremely suggestive, I shall propose the study of pattern recognition, in the broad sense, as a fundamental mode of information processing underlying natural intelligence. As a means of abstracting environmental features that could be associated with either "harmful" or "beneficial" effects, pattern recognition capabilities seem to have evolved as a resource of organisms for coping with environmental disturbances. For organisms, especially the simplest ones, it is easier to "remember" classes, or patterns, than individual environmental signals. Indeed, it is far more economical. Ascending the ladder of biological complexity, the ability of organisms to classify and generalize features of the enviornment seems to become correspondingly more complex. Organisms become able to cope with a complex hierarchy of superimposed dichotomies of harmful and beneficial environmental features. At these levels of complexity, an organism has to decide on which of these dichotomies are more critical for survival at a given time, depending on its state at the time. At the level of human intelligence, the pattern recognition capability is still ubiquitous. At this level, however, one can also talk about abstract patterns, that is, patterns not necessarily related to the immediate perception of environmental features. In this sense, one can associate pattern recognition with processes of abstraction and generalization, for example. Patterns thus can be extracted (or recognized) from representation of past and present events, and projected into models of present and future situations.

Of course, the concept of pattern recognition has been applied, directly or indirectly, to the explanation of countless phenomena. It has been applied even to the explanation of interactions at the molecular level, as in the case of the affinity between enzyme and substrate. The idea here is, however, to hint to modes of information processing underlying natural intelligence which may be worth studying in more detail. The suggestion is to study pattern recognition processes not only in the context of perception, but also in connection with higher cognitive processes. The use of a more general concept of pattern recognition also might be useful, perhaps with more than a purely instrumental value, in the study of theoretical issues like the symbol-matter problem.[4]

To conclude, I would like also to contrast a fragmented approach to AI, in which solutions are developed for specific problems with the one proposed here. Under the former, our reservoir of solutions to specific problems (such as the understanding of subsets of a natural language, speech syntheses, the automation of certain types of problem-solving activities, or the accumulation and use of expert knowledge in specific areas) would certainly grow significantly. Under the latter, I can envisage a better understanding of principles and mechanisms underlying natural intelligence as a source of tools and techniques that could be fruitfully applied to many areas of AI. However, it would also provide a unified framework and, hence, the means of capitalizing on the knowledge acquired in each specific area of investigation.

References

Kampfner, R. and M. Conrad, 1983. "Computational modeling of evolutionary learning processes in the brain." Bulletin of Mathematical Biology, (to appear).

Conrad, M., 1976. "Complementary molecular models of learning and memory." BioSystems 8, 119-138.

Conrad, M., 1983. "Microscopic-macroscopic interface in biological information processing." Technical report CSC-83-003, Department of Computer Science, Wayne State University. (To appear in the proceedings of the Orbis Scientiae Conference in honor of P.A.M. Dirac's eightieth year.)

Pattee, H., 1982. "Cell Psychology: An Evolutionary Approach to the Symbol-Matter Problem." Cognition and Brain Theory, 5(4), p. 325-341.

FROM: Harry Klopf

SUBJECT: Recommendations for Research in Artificial Intelligence (Prepared in conjunction with the AI Workshop at the University of Miami, 14-25, March, 1983)

TO: Dr. Herman Robl
U.S. Army Research Office

1. First, I want to note that the following opinions are my own and do not represent the official position of the Air Force.

2. Regarding the questions you raised at the outset of the workshop, I will state the questions as I heard them and then attempt brief answers.

Is it productive at this time to undertake computer simulations of brain functions as a means of furthering AI research?

I would say, definitely "yes." I view this as a high payoff area at this point. Such simulations will, I believe, be mutually beneficial for our understanding of natural intelligence and artificial intelligence.

Is it time to undertake hardware realizations of brain-like circuity?

I would say, in general, "no." Brain theory has not advanced far enough and computer simulations remain a very efficient research vehicle. This answer has to be qualified though because there may already be exceptions where hardware realizations could be instructive.

Is there any hope for self-organizing systems or will we always have to program the knowledge base in advance?

I would say that for truly powerful AI systems, for such applications as image and speech understanding, our only hope is self-organizing systems. The immense knowledge bases that will be required will, I believe, constitute an insurmountable obstacle to programming approaches. We must automate the process of knowledge acquisition (learning), and that is what self-organizing systems are all about.

3. Because I do recommend a neurobiologically oriented approach to AI, let me say a little bit about how I think such an approach differs from most current AI research. Below, I have listed what I believe are some fundamental or essential characteristics of natural intelligence. I have put the characteristics in two categories, in terms of whether they arise explicitly or implicitly in the case of natural intelligence:

NATURAL INTELLIGENCE

EXPLICIT CHARACTERISTICS	IMPLICIT CHARACTERISTICS
Actions	Symbols
Goals	Rules
Control of inputs	Control of outputs
Shaping of behaviors	Searching for behaviors

Before I comment on this table, some explanations are required. First of all, regarding the distinction between control of inputs and control of outputs, this is fundamental and has not generally been appreciated. I recommend Powers (1973) for his discussion of this point. One other explanatory comment: In the case of natural intelligence, animals, in general, don't search for behaviors; instead, they behave and let the environmental feedback shape their behavior. The next time a similar situation arises, their behavior is more appropriate. This illustrates the kind of distinction being noted here between searching and shaping.

It should also be noted that all of the characteristics on the right in the table above tend to become explicit to some degree in the most advanced forms of natural intelligence, namely homo sapiens. However, I would argue that, even when this happens, it is the characteristics on the left that are fundamental to natural intelligence. I consider this point important because I believe that, to a considerable degree, AI research has focused on the characteristics on the right in an explicit fashion and has then tended to address the characteristics on the left only implicitly, if at all. The result, I think, is that AI systems differ markedly in character from natural intelligence. This would be

OK if AI research were yielding robust system designs. However, I think AI research is not and may not in the future with current approaches. I believe the way to obtain more robust system designs is to employ a neurobiologically oriented approach, an approach that explicitly addresses the characteristics I have noted on the left in the above table.

One final comment: I feel concerned that the above discussion is altogether too brief and that it will precipitate misunderstandings. Permit me to recommend Klopf (1982) for a fuller discussion of some of these issues.

4. In concluding, I want to express my appreciation to you and the Army Research Office for supporting the AI Workshop at the University of Miami. Every participant in the workshop, I felt, contributed significantly and it was evident that all were genuinely interested in and seriously investigating the problems we were discussing. I found the week to be most valuable and productive.

References

Klopf, A. Harry (1982) The Hedonistic Neuron: A Theory of Memory, Learning and Intelligence, Hemisphere Publishing Corporation, New York.

Powers, William T. (1973) Behavior: The Control of Perception, Aldine Publishing Company, Chicago.

POSITION PAPER of HUGO M. MARTINEZ

for the

Interdisciplinary Study on Artificial Intelligence, March 14-25, 1983

1. Overview of the Various Disciplines

I see mainstream AI as emphasizing natural language and expert systems research. These areas are undoubtedly indispensable for the eventual design of systems which can exhibit nontrivial intelligent behavior. Particularly intriguing as a sub-area of expert systems is the new concern with meta-level architecture. This concept captures, to my way of thinking, the essential aspect of intelligence: the ability to reason about one's own capabilities and behavior. What does not seem to be of general interest, however, because of the nearly exclusive use of high-level languages like LISP, is the combined software/hardware problem and how it may be that radically new computational architecture may be required to achieve significant progress in the chosen areas.

On the other hand, I interpret the non-mainstream approach as being principally concerned with the software/hardware problem in an attempt to capture the structure-function relationship exhibited by natural nervous systems. The need to pay attention to parallelism for speed and how it is that concurrent systems (networks of asynchronous processors) can be sensibly organized and programmed to yield intelligent behavior are among the pressing issues which I think are being addressed by this sector of research.

2. My Emphasis

Coming from the side of trying to understand how natural systems are organized and how they perpetuate their design in both a conservative and exploratory manner, I am pursuing two projects. One aims at encoding a hierarchical organization for concurrent systems in a manner reminiscent of how a multicellular organism develops from a single cell (processor) by a series of binary divisions. The other seeks to generate feature detectors in adaptive classifiers by means of genetic algorithms in the sense of Holland and Reitman (1978) and of Smith (1980), and also by means of a formalism analogous to that of quantum mechanics.

The hierarchical encoding project is an attempt to deal with the problem of how to program concurrent systems. It is designed to take advantage of the hierarchical lay out of multiple processors on a VLSI chip as advocated by Mead and Conway (1980). The idea is for the program at the root of the tree to copy itself into the daughter processors by means of a series of forks (new process creation as used in an operating system like Unix). At each fork, the parent process instructs the daughter process what specific program it should be carrying out, etc. At each node of the tree, a processor may be functioning as a supervisor of its two daughters or as a non-supervisor, depending on the problem the system is designed to be solving. This project is still not documented.

The feature generation project is an attempt to deal with the problem of how to create sensibly "new points of view" from existing ones. The actions possible for a classifier are voted upon by a finite set of internally contained "voters". Given an input, each voter compares it with stored templates, one for each possible action, and votes for that action whose template is "most similar" to the input. On the basis of past experience, the classifier has assigned a "confidence" weight to each voter and chooses the action to be performed according to these weights.

Each vote represents a point of view, and new voters are obtained by performing a binary operation on the templates of two existing ones. The binary operation can be as obtained with a genetic algorithm or by means of matrix operations if one looks upon the templates of a voter as the eigenvectors of an "observable" with the actions as the corresponding eigenvalues. Details in Martinez (1983).

References

Holland, J.H. and Reitman, J. (1978) "Cognitive Systems Based on Adaptive Algorithms", in Pattern-Directed Inference Systems, eds. Waterman and Hayes-Roth, Academic Press, 1978.

Smith, S.F. (1980) "A Learning System Based on Adaptive Algorithms", Doctoral thesis, University of Pittsburgh, 1980.

Mead, C. and Conway, L. (1980) "Introduction to VLSI Systems", Addison-Wesley Publishing Co., 1980.

Martinez, H.M. (1983) "Computer-Aided Studies of Evolutionary Systems", Final Report, NSF Grant #MCS-8018242.

RADA'S POSITION PAPER

What are the contributions which various disciplines may make to AI and what perspective may be particularly worth emphasizing?

1. Contributions of Various Disciplines. Mainstream AI is properly a part of Computer Science. Cognitive psychology and linguistics have many results that relate to natural language processing and that currently link those two fields to the AI subdiscipline of natural language processing. This influence should wane in about ten years as natural language processing is automated. Work in the neurosciences and biopsychology will be most useful to AI when problems with vision and motor control are solved by machines in about 20 years. In 40 years, the push of AI will be for creative, adaptive machines that are based on radically new hardware, and then contributions from molecular biology and evolution theory should be critical.

2. Approach to Intelligence. What are necessary and sufficient conditions for intelligence? A necessary property seems to be an endless iteration of generate and test. But the intelligent system internalizes testing, and this introduces the need for populations of competing components (Holland, 1975). Each component is capable of generating parts or all of itself. These are necessary properties of intelligence, but they are not sufficient. The parameters of the model may easily be set so that the behavior quickly becomes periodic (Rada, 1982). Hierarchies of components must naturally evolve. Continuity in structures must exist (Conrad, 1979; Lenat, 1982).

More progress towards delineating the properties needed is stymied by the lack of a good measure of when a system has intelligence. In an intelligent system each component tries to spread its influence. It spreads its influence by copying parts of itself into other components. At any time t the system should have a subset s of components which A) uses less than half of the available resources of the system, but B) after some time, t+j s and its products use more than half of the system's resources (Rada, 1981). Given some non-zero level of noise in the system, finding initial conditions for the system so that it manifests "intelligence" for a long time is very difficult. With a quantifiable metric, however, results from various experiments can be effectively compared.

References

Conrad, M., 1979. "Bootstrapping on the Adaptive Landscape," Biosystems 11:167-182.

Holland, J., 1975. Adaptation in Natural and Artificial Systems. University of Michigan Press: Ann Arbor.

Lenat, D.. "The Nature of Heuristics," Artificial Intelligence 19, 189-249.

Rada, R., 1981. Evolutionary Structure and Search, Ph. D. Thesis, Department of Computer Science, University of Illinois, Urbana.

Rada, R., 1981. "Evolution and Gradualness" BioSystems 14, p. 214-218.

WALTZ'S POSITION PAPER

Mainline AI (the variety currently done on serial digital computers) has long been centered around the symbolic representation of knowledge. While, in part, this direction has been followed by AI because of expedience, it also seems to me to be the best way to address problems involving thought, judgment, reasoning, natural language, beliefs, planning, knowledge, goals (of the kind people would generally say they pursue), actions, events, modeling of the world (e.g., cause and effect) -- in short, nearly everything we include under the terms "culture" and "knowledge." In this view, people operate at "the knowledge level" [Newell, 1981a] and are able to do this because, via some not-well-understood mechanisms, our brains are effectively "physical symbol systems" [Newell, 1981b]; that is, the net result of the action of our vast numbers of neurons is to allow us to operate as though we were "perfect processors" [Simon, 1965] following abstract symbolic rules given to us through our language and education, induced through interaction with the world, or innately embodied in us.

While there seems to be much truth in AI's view, there are nagging problems that have not been very well-explained at the knowledge level. Examples include perception, sensory-motor interaction with the world, mental imagery, intuition, and some types of learning, particularly adaptation and learning of skills. It is these apparently non-symbolic problems that seem most appropriate for neurally-inspired modeling, and it is here, indeed, that most of the neural-net, automata-based, evolutionary, and adaptive network efforts have been focused. Even the best mainline AI work on vision [Marr and Hildreth, 1979] was based heavily on results in neurophysiology.

Almost without exception, mainline AI has ignored bottom-up learning, and has instead concentrated on producing "instant adults."

Unfortunately, a wide gap exists between the neural-net type research and mainline AI. The gap will remain, I believe, until the neural-net community shows that its programs and devices can perform symbolic operations in an explicit, convenient, and satisfying manner. Soon after birth, infants seem capable of making figure/ground distinctions, and of segmenting units, i.e., items that can be named and composed in larger units, eventually resulting in mental life that centers on the symbolic. I very much doubt that our neural net models will accomplish anything similar unless they are given a great deal of structure, and possibly even "innate knowledge" [Fodor, 1975]. Randomly structured masses of neurons with random weights seem very unlikely to me to ever learn symbol manipulation.

What are the prospects for self-organizing systems? Already, a great deal of recent work in mainline AI has demonstrated promise in learning of scripts (generalized event sequences) [DeJong and Waltz, 1983], learning by analogy [Winston, 1980], learning by discovery [Lenat, 1977], and learning of natural category descriptions through experience [Michalski, Carbonell and Mitchell, 1983]. While other researchers at the workship reported progress in learning (e.g. Holland's robot arm system at Polaroid that learns to touch a spot on a CRT screen, or Barto's system that learns to maneuver in a simple world), most systems described did not seem yet to merit fully the term "self-organizing", and none seem to me to have any near-term practical application value. My guess is that mainline AI will produce useful learning (if not quite self-organizing) systems well before neural net research pays off.

Should neural net-type research be continued and encouraged?

Of course! (Though I would like to see much more emphasis on structure and programmability, as mentioned above.) No account of human cognition can be satisfactory unless it accounts for the detailed nature of the workings of the brain and its components, down to the neural level (and perhaps below). Moreover, mainline computer architecture has let AI down: it has opted either for (1) keeping the standard "von Neumann" architecture and attempting to increase its performance through methods such as pipelining, caching, cryogenics, instruction pre-fetch, and small levels of parallelism, or (2) constructing parallel vector-type machines, where all processors are more or less synchronized. Neither of these approaches meets mainline AI's needs, and I, for one, hope that neural-net researchers can help us.

References

DeJong, G. and D.L. Waltz, 1983. "Understading Novel Language," International Journal of Computers and Mathematics, to appear.

Fodor, J., 1975. The Language of Thought, Thomas Y. Crowell Co., New York.

Lenat, D.B., 1977. "Automated Theory Formation in Mathematics," Proc. 5th Intl. Joint Conf. on Artificial Intelligence, Cambride, MA, pp. 833-842, August 1977.

Marr, D. and E. Hildreth, 1979. "Theory of Edge Detection," Proc. Royal Soc. London B, 207, pp. 187-217.

Michalski, R., J. Carbonell and T. Mitchell, 1983. Machine Learning: An Artificial Intelligence Approach, Tioga Press, Los Altos, CA.

Newell, A., 1981a. "The Knowledge Level," AI Magazine, 2,2, pp.1-20.

Newell, A., 1981b. "Physical Symbol Systems," Perspectives on Cognitive Science, D.A. Norman (ed.), Lawrence Erlbaum Assoc., Hillsdale, NJ, pp. 37-86.

Simon, H.A., 1969. The Sciences of the Artifical, MIT Press, Cambridge, MA.

Winston, P.H., 1980. "Learning and Reasoning by Analogy," Commun. ACM 23, 12, December.

AI AS AN ENGINEERING DISCIPLINE: THE FLIGHT OF BIRDS ANALOGY

Position Paper

Interdisciplinary Workshop on AI

Bernard P. Ziegler
Department of Computer Science
Wayne State University
Detroit, MI 48202

I shall take the position that AI is concerned with the design of certain kinds of systems, viz., those whose behavior evidences more "intelligence" than commonly accepted as normal for artifacts. From this perspective, AI is a special branch of systems engineering or, if the systems are restricted further, of software engineering. I shall consider the impact of this perspective on the directions in which AI should move from two aspects: the methodological discipline that is characteristic of engineering at its most professional level, and the use of empirical lawful relationships in design. This perspective implies that the AI field should seek more firm design-methodological and bio-scientific foundations.

1. Impact of Methodological Discipline on AI. From an engineering perspective, it follows that an AI project should be subject to the methodologies of systems or software design: objectives should be well-formulated, requirements should be translated into specifications, design should proceed top-down, validation and verification procedures should be prespecified, etc., etc.

Of course, we know that such methodologies are not universally adopted for systems design generally, and in AI specifically. On the one hand, the methodologies are not sufficiently advanced so as to be convenient to use.[1] On the other hand, there has to be a firm base of experimental knowledge in the foundations on which a system is designed to make the methodological discipline at all meaningful. One can't meaningfully lay out objectives, for example, if one hardly knows whether these objectives are realistically achievable.

The latter is probably the case in AI. Its practitioners would, no doubt, argue that the methodological approach is premature. We are only beginning to understand the limits of the possible and how to achieve them. Nevertheless, ambitious plans for fifth generation computers,[2] and other similar commercial plans, indicate that intense development, as opposed to research, of AI systems is being considered seriously.

So it is my position that we ought to start relating to AI projects from a systems and software design perspective. It does not seem to me that having to carry out a design-methodological discipline would overly inhibit progress. Indeed, taken with a grain of salt, it might serve to direct thinking in AI into more fundamental, and lasting considerations, than has been the case heretofore. For example, the single idiosyncratic program would be less of an objective, than the concepts that support a whole design approach to a particular problem area.

2. Relation to Natural Science: the Flight Analogy. This perspective also throws some light on the relation of AI to cognitive and biological science. Traditionally, AI workers have asserted the independence of intelligent behavior from its physical realization.[3] The analogy with flying has been accepted at face value: trying to mimic the flying behavior of birds was not the way to realize artificial flying vehicles. Therefore, the argument goes, mimicking human behavior is not necessarily the way to achieve artificial intelligence.

But if we examine the flight situation carefully, we see that artificial flight was achieved when certain principles about the real world (e.g., thrust, lift, drag, weight) were grasped and manipulated in a form that made the objective feasible. These principles, now formalized in aeronautical engineering, apply both to birds and to airplanes. Both types of flight are incorporated in a unified framework: dynamic models that generate flight behavior can be set to different parameter settings to characterize the different realizations of stable flight. Indeed, this parametric manipulation capability now makes human powered flight feasible through very sophisticated simulation model-based structural design (and the advent of super-light materials, an engineering application of science as well).

It is the analogy, more deeply understood, that I feel should drive the accepted paradigm. Taking AI seriously as an engineering discipline, with engineering objectives, means that principles of intelligent behavior should be sought that apply to both natural and man-made systems. Just as with flight, an artifact that is man-made must still be fashioned from natural material and must obey the rules of the natural world. We do not, at this point, understand the levels at which real world constraints impact intelligent behavior. The assumption that only the level of "symbol system"[4] is significant may be an entirely wrong lead. It would seem much safer research strategy to look at physiological, biological, and other levels as well.

In sum, were we to understand the principles of intelligent behavior, we would have the parametric manipulation capabilities that we have in the case of flight. And, as in that case, simulation models incorporating these principles would be developed to support the design process.[5] This is, no doubt, in the future. But shouldn't we start in that direction today?

References

Ziegler, B.P. and R. Rada, (1938), "Abstraction-Based Methodology: A Framework for Computer Support," J. Information Processing and Management, (in press).

Motta-Okka, M. (1982), Fifth Generation Computing Systems, North Holland.

Haugeland, J. (1981), "The Nature and Plausibility of Cognitivism," reprinted in Mind Design, (ed: J. Haugeland) MIT Press, Cambridge, MA.

Newell, A. and H.A. Simon (1981), "Computer Science as Empirical inquiry: Symbols and Search," reprinted in Mind Design, op. cit.

Ziegler, B.P. (1983), "Theory and Application of Modelling and Simulation: a Software Engineering perspective," Handbook on Software Engineering (ed: C.R. Vick, C.V. Ramamoorthy): Von Nostrand, NY (to appear).

POSITION PAPER

ON

INTERDISCIPLINARY STUDY ON ARTIFICIAL INTELLIGENCE

Howard Pattee
Department of Systems Science
School of Advanced Technology
State University of New York at Binghamton
Binghamton, New York 13901

POSITION PAPER of HOWARD PATTEE

for the

Interdisciplinary Study on Artificial Intelligence, March 14-25, 1983

Present Artificial Intelligence (AI) research is dominated by an explicit and principled disregard for hardware realizations, both physical as well as biological, while simultaneously it embraces a conceptual dependence on current logical and technological computational models (e.g., Newell, 1980; Pylyshyn, 1980). Of course, these logics and technologies are very powerful, and AI research should continue to make the most of them.

The question is whether they are powerful enough. Does this disregard of biological architectures and evolutionary strategies place limitations on what can be achieved by current approaches? Much of the early thinking about computation made direct use of speculations on how the brain works (e.g., Turing, von Neumann). We now know that most of these speculations were erroneous or greatly oversimplified; but, despite this fact, the course of computer development has continued along these early lines and, by now, has become entirely technology-driven with little reference to molecular or cellular processes in nervous systems that display learning and intelligence. The rapid rise in the past decade of the interdisciplinary field called the <u>cognitive sciences</u> is, in many ways, an attempt to correct this lack of interaction of computer models with psychology and the neurosciences, but the fact remains that AI research proceeds largely unaffected by biological fundamentals.

What biological fundamentals offer a complementary correction to current AI trends? These well-established facts of life appear universal:

(1) Linear copolymer molecules (e.g., nucleic acids, proteins) are the basic units of instruction, recognition, and action at all levels of biological function, including the nervous system. These macromolecules are continuously

constructible and modifiable, with enormously high specificities and intrinsic powers of self-assembly into higher functional units. They cannot be usefully compared to present fixed computer memories, architectures, and logics. AI research needs to look more realistically at simulations of these macromolecular computations (e.g., Conrad, 1983).

(2) The cell is the basic coherent, self-sustaining unit of all living systems, and the cell is a highly evolved, exceedingly complex informational and physical system. It is totally unrealistic to imagine a neuron as a simple gating element, however modifiable its assumed logical function may be. AI research needs to explore models of neurons as self-organizing units with realistic cellular powers of adaptation and learning (e.g., Pattee, 1982).

(3) Biological intelligence is a product of gradual evolution over billions of years. Formal symbol systems and computation is only the very latest, most abstract, and most artificial form of representation that life has created. In so far as it is formal, it is inherently meaningless. The essential interpretative functions necessary to provide a semantics is not well understood. AI research needs to consider the evolutionary steps that are prerequisites for the meaningful applications of formal logics and representations.

References

Newell, A. Physical symbol systems. Cognitive Sciences, 1980, 4, p. 135-183.

Pylyshyn, Z. Computation and cognition. Behavioral and Brain Sciences, 1980, 3, p. 111-169.

Conrad, M. Microscopic-macroscopic interface in biological information processing, Dept. Computer Science, Wayne State Univ. Report of work done on NSF grant MCS-82-05423 (Distr. at this meeting).

Pattee, H. Cell psychology: An evolutionary approach to the symbol-matter problem. Cognition and Brain Theory 1982, 5(4), p. 351-341. (Distr. at this meeting).

POSITION PAPER

ON

INTERDISCIPLINARY STUDY ON ARTIFICIAL INTELLIGENCE

Hans J. Bremermann
Department of Mathematics
University of California, Berkeley
Berkeley, California 94720

POSITION PAPER of HANS J. BREMERMANN

for the

Interdisciplinary Study on Artificial Intelligence, March 14-25, 1983

At present, many systems (like cars, planes, plants, etc.) carry power to subsystems (motors, lights, hydraulics, etc.) over wires from switches and controls that are near the operator but remote from the subsystems.

It is possible to replace distant switches by local switches, actuated by messages. The messages can be sent over thin (nonpower-carrying) wires, or by means of a carrier: sound, light, radio frequency, etc. A special case is the nervous system. Some messages travel over fibres as action potential, others in the form of chemical signals (hormones).

For large, complex systems, questions of reliability arise: a local switch may fail to respond, a channel may be noisy, a neuron may die, a receiving unit may malfunction.

The problem is not only to analyze the reliability of a system as built or to make the components as reliable as possible, but to design systems for reliability.

In spite of the success of complex engineering systems in the space program, in spite of the reliability of modern computers, there remains much unexplored territory in this field.

In computers, for example, John von Neumann, in the fifties, worried about how to synthesize reliable computers from unreliable components, and he wrote a famous paper about it (in 1956) which also makes reference to nerve nets. Von Neumann analyzed what could be done by providing multiple copies of switching components (multiplexing).

Winograd and Cowan (1963) showed that multiplexing is an inefficient design for reliability and that ideas from Shannon's analysis of noisy communication channels can be applied to circuit. Subsequently, the phenomenal reliability of solid-state devices probably was responsible for a decline in interest and work in the area of design of reliable computers from unreliable components. It seems that there are many as yet unexplored possibilities.

In a system where there is a human controller (or group controllers), subsystems send messages about their status to the control center. Unless the operator is to be confronted with all the raw data, the control center requires a systems model that can be consulted to sound an alarm or to take automatic action to correct deviation from set values of the state variables.

Control systems have, of course, been studied for decades. However, the mathematical analysis has often been limited to linear systems dynamics and controls and to simple objective functions. The reason is that, mathematically, analysis becomes quickly untractable when these simplifying (and often oversimplifying) assumptions are dropped. It is here that the methodology of artificial intelligence holds great promise. The advent of inexpensive, powerful microprocessors has opened up numerous new possibilities. For example, parallel processing is still in its infancy. It may well be worthwhile to take a fresh look at parallel algorithms, parallel complexity, parallel computer architecture, and cybernetics.

Switching by means of messages and tuned receivers also happens in genetic regulation. Here, the messages are regulatory proteins that bind to specific sequences of the DNA double-helix. These binding proteins can act as repressors or promoters and prevent or promote the transcription of specific structural genes (which, in turn, can switch on or switch off other genes, etc.) In spite of the pioneering work of Jacob and Monod in the sixties, gene regulation has remained poorly understood. Recently, specific details of the λ-switch mechanisms from lysogenic to lytic infection have become known.

There is some older work by Kauffman onstable cycles in randomly constructed gene-regulatory nets. I think that analogous results may be important for nerve nets.

It should also be interesting to investigate evolution processes in genetic nets in view of the new knowledge about transposable genetic elements (which change the network structure as such).

Finally, it seems that cancer is a problem of systems reliability peculiar to living systems with elements (cells) that are capable of proliferation.

The topics outlined seem to hold much promise. I would hope that further discussions could explore not only nerve nets but genetic (molecular) nets, reliability of biological systems as well, and also allow for pursuit of relevant problems in artificial intelligence, algorithms, and theoretical computer science.

in neurophysiology.

END

FILMED

9-83

DTIC