

SIGART Newsletter

No. 58 June 1976

The SIGART Newsletter is a bimonthly publication of the Special Interest Group on Artificial Intelligence of the Association for Computing Machinery. The Newsletter reports on projects being conducted by the artificial intelligence research community and generally reviews current progress in the state-of-the-art. Correspondents report news from local SIGART Chapters and other AI centers.

Former Chairman: **Les Earnest**
Artificial Intelligence Laboratory
Stanford University Stanford, CA 94305
phone: (415) 497-4971
ARPAnet: LES@SU-AI

Chairman: **Rich Fikes**
Artificial Intelligence Center
SRI Menlo Park, CA 94025
phone: (415) 326-6200 ext. 4620
ARPAnet: FIKES@SRI-AI

Vice chairman: **Dave Waltz**
CSL University of Illinois
Urbana, IL 61801
phone: (217) 333-6071

Secretary/Treasurer: **Don Waterman**
The Rand Corporation
1700 Main Street Santa Monica, CA 90406
phone: (213) 393-0411
ARPAnet: DON@RAND-USD

Newsletter Editor: **Lee Erman**
Assistant Editor: **Mike Rychener**
Computer Science Dept. Carnegie-Mellon University
Pittsburgh, PA 15213
phone: (412) 621-2600 exts. 146 & 141
ARPAnet: SIGART@CMU-10B

The Editors encourage contributions from authors, including letters to the Editor (AI Forum), technical contributions (1 to 6 pages), abstracts (preferably 100-200 words), book reviews, bibliographies of special topics in AI, news items (conferences, meetings, course offerings, personals, etc.), advertisements (new products or classified ads), puzzles, poems, cartoons, etc. **Material may be reproduced from the Newsletter for non-commercial use with credit to the author and SIGART.**

Anyone interested in acting as editor for a special issue of the Newsletter devoted to a particular topic in AI is invited to contact the Editor. Letters to the Editor will be considered as submitted for publication unless they contain a request to the contrary. Technical papers appearing in this issue are unrefereed working papers, and opinions expressed in contributions are to be construed as those of the individual author rather than the official position of SIGART, the ACM, or any organization with which the writer may be affiliated.

You are invited to join and participate actively. To indicate a change of address or to become a member of SIGART, complete the form on the last page of this issue. Direct all business concerning membership and subscriptions to ACM Headquarters, 1133 Avenue of the Americas, NY 10036.

For those with access to the ARPAnet, materials may be submitted by netmail to SIGART@CMU-10B (host 14 decimal). Online access to the currently developing issue is available by connecting to CMU-10B via Telnet and typing "HELP SIGART".

Copy deadline for August issue: July 25.

OFFICERS' MESSAGES

Outgoing Chairman -- Les Earnest

The results of the special SIGART election are:

Chairman: Richard Fikes

Vice-Chairman: David Waltz

Secretary-Treasurer: Donald Waterman

Newsletter Editor: Lee Erman

Their term of office runs from now to May 31, 1977.

This completes the transition from appointed to elected officials that was initiated two years ago and marks a significant step in the maturation of SIGART. On behalf of the SIGART membership, I wish to express thanks to the new officers for their willingness to serve.

Chairman - Rich Fikes

Greetings. As I begin my SIGART Chairmanship, the organization appears to be alive and well. I want to thank Les Earnest for doing a fine job as Chairman, and for helping me and the other new officers get our administration started.

This is the first time (as far as I know) that SIGART has had a secretary-treasurer and vice chairman to assist the chairman and editors. This almost doubles the number of people actively responsible for running the organization, and should noticeably increase our freedom to do more than just "keep the store open". SIGART's primary service to the AI community is the Newsletter, and I see no compelling reason to change that focus. However, we should keep in mind that SIGART has a broad charter and can be called upon to provide many types of services when they are needed. For example, SIGART can sponsor and help organize conferences, workshops, and conference sessions. Currently, plans are being made to cosponsor a conference with SIGPLAN, and three AI sessions have been organized by SIGART for the upcoming National ACM Conference.

The ACM conference will be held October 20-22 in Houston, Texas. Summaries of the AI sessions appear in this issue of the Newsletter. There will also be a SIGART meeting during that conference, probably in conjunction with a dinner. We will announce final plans for the meeting in the next issue of the Newsletter.

The recent election revealed a feature of the SIGART bylaws that many people (including myself) do not like -- namely, that a SIGART election of officers can be held in which there is only one candidate for each of the offices except one (Article V, Paragraph A). We are preparing an amendment to the Bylaws that will help assure, whenever possible, at least two candidates on the ballot for each office.

I am looking forward to an enjoyable year as SIGART Chairman. I and the other officers want to be responsive to the needs of the membership. Our mailboxes are open.

Vice-chairman - David Waltz**Special Issue on Natural Language**

I have volunteered to edit a special issue of the *SIGART Newsletter* to cover the general areas of practical applications of natural language for small task domains, natural language for data-base information retrieval, and data-base systems for AI applications. We expect this would appear in the December issue.

Accordingly, I am soliciting appropriate contributions and suggestions. These can be of various forms, including bibliographic references and notes about works in progress. Submissions should be sent to me (at the address on the inside front cover) or, if accessible, via the ARPAnet to SIGART@CMU-10B. The deadline for these submissions is September 30.

Secretary/Treasurer - Don Waterman

As the new Secretary-Treasurer of SIGART I will do my best to encourage the expansion of SIGART activities and will work toward the development of a budget that accurately reflects the desires and needs of SIGART members. I expect to have a report on the current state of our finances for the next issue of the Newsletter.

Newsletter Editor - Lee Erman

To be Drawn (out) and Quartered?

I would like to propose to the membership that the *SIGART Newsletter* go on to a quarterly schedule, rather than the current bimonthly one. I see several advantages:

1. We would have three months, rather than two, to integrate submitted materials. This would hopefully provide a larger and more uniform (in quality) set to choose from.
2. Since it now takes about five or six weeks from the time we send off the camera-ready copy to be printed until it is in the hands of the readers, there is only about two weeks time left for readers to respond to that issue before we are sending the next one off. Thus, reactions very often have to wait until the second following issue to be printed. A quarterly schedule might actually result in faster turn-around of responses.
3. We could print the same number of pages per year and save money, or, alternatively, print more pages for the same amount. This is because of economies in printing and mailing fewer issues which are larger.

The main disadvantage would be, of course, slightly longer delays (on the average of two weeks) for announcements of time value, e.g., conference announcements and advertisements.

I would very much like to hear reactions from members about this proposal.

REPORTERS' REPORTS

AI at UCLA

J. Pearl and M. Rhodes

Departments of Computer Science and Engineering Systems

University of California at Los Angeles

Sponsored research pursues many interdepartmental interests in AI. A brief description of current research follows.

The interaction of inferences, affects, and intentions in a model of paranoia

The Algorithmic Laboratory of Higher Mental Functions in the Department of Psychiatry conducts research involving computer simulation of human thought processes, in particular in those found in paranoid disorders. We are interested in problems of natural language disorders. We are interested in particular in those of natural language understanding, belief and inference systems, and the planning of patterned action sequences. The project is funded by the National Institute of Mental Health.

Research information: Kenneth Mark Colby, M.D. -- UCLA Department of Psychiatry

Biocybernetic Control in Man-Machine Communication

This neurocybernetics research project has been underway at the Computer Science Department for the last four years. The aim is the incorporation of biological signals, and in particular, of brainwaves in the form of evoked responses into practical man-machine communications. Information from the human brainwave is

directly decoded and interpreted in the context of a man-machine dialogue which forms a closed loop interaction. The research entails the decoding of natural signals and the discovery of their syntax, a task not unlike that of speech recognition. The overall concept is that of a man-robot symbiosis with direct control of robot action for prosthetic and other applications.

Sponsor: National Science Foundation and Advanced Research Projects Agency.

Research information: Jacques J. Vidal -- UCLA Department of Computer Science.

Speech Recognition

The goal of current research in the phonetics Lab, Department of Linguistics, is to learn more about the relationship between the acoustic structure of speech and the associated physiological gestures. Our underlying belief is that more knowledge of this relationship holds the key to decoding the signal from the many degrees of variability between tokens which characterize human behavior. We presently have several operating Fortran programs which convert speech signals (from waveform or spectral structure) into the corresponding vocal tract geometry in various forms. Most of this research is presently concentrated on vowel articulations. Of particular interest is the study of pharyngeal geometry as this is difficult to measure directly, and yet some of the best articulatory data available for a wide variety of languages is in the form of lateral (mid-sagittal view) xrays, which show the tongue in profile. The physical parameter most directly related to the cross-sectional area of the vocal tract as a function of the distance from the lips. Relating these two kinds of information about the positions of the articulators at any instant requires knowledge of the physiological structures which is not presently available.

Research information: D. Lloyd Rice -- UCLA Department of Linguistics.

Robustness in Problem Solving and Natural Language Understanding

Robustness is the ability to respond without program modification to slightly perturbed, or to inexact-specified situations. This quality would seem to be very useful in a number of applications such as assembly line automation, and is moreover characteristic of the way people cope with their environment. This project is exploring the systematic use of robustness and inexactness in natural language semantics and general problem-solving. In particular, we are designing a conversational system to accept and use vague hints about how to run a maze.

Our methods include: careful hierarchical organization; use of procedures to represent knowledge, actions, etc.; fuzzy sets; a general goal of orientation in the system itself; and semantics (i.e., meaning).

Some effort is also being put into broader issues. Evidently, most large-scale public, environmental, and social systems are fuzzy; and we would like them to be robust. The theories of systems and inexactness seem to shed some light here. There are also interesting connections with linguistics and related fields. We are beginning to explore automatic analysis and synthesis of music.

Sponsor: National Science Foundation.

Research information: Joseph A. Goguen -- UCLA Department of Computer Science.

Pattern Recognition and Machine Pattern Analysis

Research in this program is on structural and statistical models in pattern recognition with applications in image processing, computer graphics, and human-computer interface using text-command-interpretation.

A regular decomposition technique for condensing array data like that obtained from scanning pictures is under investigation. Work includes three-level algorithms for creating a tree to retain minimal amounts of array elements while preserving the maximum amount of structure apparent in the original picture.

Programs were developed to use the regular decomposition technique in land-use planning. Test and evaluation of pattern recognition effectiveness requires standard data sets; this research contributed to the development of such sets ("artificial data bases").

Programs for more efficiently using IMLAC computer hardware features for picture display continue to be developed. An interactive language ("SKETCH") was written permitting natural commands for modifying a human face displayed on the IMLAC. Further work includes combining a variety of pattern recognition utility programs in a multi-computer system (The Brain-Computer Interface Laboratory, UCLA).

Un-sponsored research investigated aspects of modeling electrocardiogram data. Another un-sponsored effort is the establishing of a standard glossary of terms in pattern recognition.

Sponsor: Air Force Office of Scientific Research.

Research information: Allen Klinger -- UCLA Department of Computer Science.

Machine Representation of Knowledge

The goal of the study is to understand the tradeoffs between various ways of organizing information in a computer and to devise means by which shifts in data structures could be accomplished automatically. Our main effort focuses on the problem of coining new primitives in terms of which problems can be solved more efficiently. The feasibility of discovering new primitives is being investigated in two problem domains: learning heuristic functions for ordered search and approximating joint probability distributions.

Sponsor: National Science Foundation.

Research information: Judea Pearl -- UCLA Engineering Systems Department.

Computer Aids to Decision-Making

Research in this area concerns the problem of eliciting subjective knowledge from persons and reducing it to a formal structure facilitating machine manipulations.

While most decision-aiding systems employ a large domain-dependent knowledge-base (e.g., medical or legal library), we explore the idea that in many cases it is more advantageous to let the user carry most of the information and only map into the machine that section of knowledge which the user perceives as relevant to his immediate problem. In this mode the machine acts as a sophisticated, friendly 'sounding board'; it does not provide information of its own, but it assists the user in structuring and searching his own knowledge, and provides advice on alternative courses of action.

A domain-independent system which interacts with the user in pseudo-natural English and maps his answers into a decision-analytic problem-structure has been developed in the past year. Work is now underway to introduce domain-specific features into the problem structure (e.g., medical diagnosis), and incorporate non-probabilistic (e.g., fuzzy knowledge) formalisms to capture inexact relations.

Sponsor: National Science Foundation.

Research information: Judea Pearl.

Analysis of Computational Gains From Inexact Information Processing

The fact that human beings organize their knowledge in an inexact manner and are able to perform very complex tasks with such

apparent ease has led many computer scientists to conjecture that a substantial amount of computation could be gained if computer systems, too, (especially in AI applications) were allowed to operate imprecisely. The aim of our research program is to understand and quantify the tradeoffs between complexity and laxity of computations, and to devise means of exploiting impreciseness to achieve more effective schemes of representing knowledge.

In the past year a technique has been identified which relates the complexity of inexact computation to Shannon's entropy function. It was very successful in providing universal tradeoff curves for a number of computational problems such as data retrieval, pattern classification and sorting. It is now awaiting analysis in the area of language recognition and theorem proving. Another task to be undertaken shortly involves a comparative evaluation of various schemes for representing partial knowledge (e.g., statistical, fuzzy logic, etc.) from a computational efficiency viewpoint.

Sponsor: National Science Foundation.

Research information: Judea Pearl.

AI at Yale

Walter J. Stutzman

Yale University Department of Computer Science
New Haven, CT 06520

Our AI group, directed by Roger Schank, is working on story understanding. During the past year, we developed the SAM system, which uses a particular instantiation of the frame concept called "scripts" to understand stories in a few domains. The system can do the following tasks with 5-10 line English stories about restaurants and trips: summary/paraphrase, question-answering and translation to Chinese, Dutch, Russian and Spanish. Our current work is directed at organizing intentionality in the following way: themes give rise to goals, which in turn utilize instrumental plans or scripts. We are developing the Plan Applier Mechanism which will combine with SAM to extend the capability of our text understanding system. One member of the group is using plans to generate stories.

Abstracts of publications by members of our group (11 graduate students, 2 research associates in addition to Schank and Bob Abelson from the Psychology Department) appear in the Abstracts section of this issue.

BINDINGS -- People Moving

Now in Boston (Apt. 5N, 9 Hawthorne Pl., Boston, Mass. 02144):

Ted Shortliffe (from SUMEX-AIM), doing a medical internship.

As of Autumn, 1976, at *Institut für Informatik I* (Universität Karlsruhe, D-75 Karlsruhe, West Germany):

Laurent Sillosy (from Université de Paris).

Now at BBN (Cambridge, MA 02138):

Andee (Ann D.) Rubin (From MIT-AI).

Richard Rubinstein (From UC-IRVINE).

AI FORUM

How to Write a Program Counts Too

From: David Alan Bourne

May 8, 1976

Dept. of Computer Science University of Vermont
Burlington, VT

After reading Drew McDermott's article,¹ I found myself basically in agreement and have personally experienced the "**Only A Preliminary Version of the Program was Actually Implemented" syndrome.

However, I was a little dismayed by his suggestion to call data structures names like: "G0073". There is no danger of a name like that being misinterpreted, because it most likely will not be interpreted at all. In my opinion it is a programmer's duty to put as much excitement into his code as possible. In fact, the more daring the names, the more honor-bound the programmer will be in living up to those names.

In closing, which would you rather read, despite a little wishful thinking?

DO WHILE THOUGHTS=VALID

...

OD

PANIC: OUTPUT 'I AM SUFFERING FROM MENTAL ILLNESS.'

or

DO WHILE G0072=TRUE

...

OD

ERR1:OUTPUT *SYSTEM ERROR: G0072 FAILED.'

Turing's Test

From: Robert K. Lindsay

May 27, 1976

University of Michigan Ann Arbor, MI 48109

In the AI Forum section of Newsletter 57 (April 1976), there are three letters related to Turing's Test. The first, from Harry M. Murphy, Jr. inquires about the famous confusion involving ELIZA and (I believe) a BBN executive; Murphy asks if such a confusion really constitutes a valid Turing test. The second, from Dennis J. McLeod, argues that simply because a person interprets responses from, say, ELIZA, as "constituting understanding", it does not follow that ELIZA really understood, since its behavior may be (and is) merely predicated on a clever trick. The third, from Yorick Wilks, takes issue with McLeod's argument, asserting that behavioral tests are the only means we have of deciding if something understands, and we have no basis for rejecting apparent understanding as not real merely because we know it to be based on simple mechanisms. The McLeod and Wilks letters were reprinted from CACM because, I presume, they were considered relevant to Murphy's question.² I found this an interesting interchange, particularly because Wilks is one of many who have claimed, on the basis of the incident alluded to by Murphy, that the Turing Test has been passed. (Wilks makes the claim in his book *Grammar, Meaning, and the Machine Understanding of Language*, page 6, but I do not wish to single him out since I have heard and read other well-known computer scientists in the same error). Wilks' present reply to McLeod, while ostensibly an argument on a more general level (the necessity and

1. "Artificial Intelligence vs. Natural Stupidity", *SIGART Newsletter*, 57, pp. 4-9.

2. Editor's note: This presumption is, in fact, not correct -- it was pure chance that the McLeod and Wilks letters were in the same issue as Murphy's; they were reprinted because they were thought to be of particular interest to SIGART readers.

sufficiency of operational definitions of understanding), in fact perpetuates the error by failing to address McLeod's valid objections. Since the confusion is still abroad, I would like to address the question of Turing's Test here.

Even if we grant that behavioral tests are the only reasonable means we have for deciding whether a machine understands (or thinks, or is conscious, or what have you), the trivial test of fooling an unsuspecting, credulous human is not a sufficient behavioral test, nor is it even a difficult one. Since human behavior, including verbal behavior, has an incredible range, any response, even the null response, is consistent with the hypothesis that the communicant is a human, even a highly intelligent human. If, as I drive down the road with the sun in my eyes, I mistake a tree for a computer scientist known to walk this way on occasion, does this constitute evidence that a tree is a sentient being, even if we assume computer scientists are? Of course not. Does a person's mistaking of ELIZA for a human when he had a priori assumed he was talking to a human mean that ELIZA understands? Of course not. The suggestion is patently absurd.

The fact that there are many people in the world who cannot play chess, read, write, do arithmetic correctly, or even count to ten is ample demonstration that computers excel certain humans in certain respects, but this should be of little interest to computer scientists. Presumably we are attempting to mimic man at his finest, or at least at his middling competency, in order to understand better the full range of the human mind and to increase the world's problem solving power; we are not, I take it, merely playing parlor games.

It is to Turing's credit that he realized all of this, and devised his "imitation game" (now widely called the 'Turing Test') to overcome such obvious problems. The original description (see "Computing Machinery and Intelligence, reprinted in Feigenbaum and Feldman, *Computers and Thought*, page 11) is not a complete description, but it is clear that the proposed test had the following ingredients: (1) There were to be two input channels to the interrogator, one connected to a computer, the other to a person; (2) the computer was to attempt to convince the interrogator that it was human; (3) the person was to help the interrogator by being as human-like as possible; and (4) the interrogator was aware of (1), (2), and (3) and was attempting to guess which channel was human-connected, which machine-connected. The machine passes the test if a sufficient number of interrogators fail to identify it sufficiently often.

The Turing Test is a stringent one, more or less so depending upon the length of time permitted, the number of trials, and the intelligence and sophistication of the interrogators and their allies, all of which would need to be specified. But given the basic framework and even loading the other factors in favor of the machine, I can say without an extensive literature search that no machine has passed this test, and am confident that none will in the near or middle future. To my knowledge no one has yet had the hardihood even to PUT any machine to this test.

Since the Turing Test is so difficult, it has been suggested that we frame simpler tests to serve as easier benchmarks. While there is merit in this suggestion there is none, I feel, to the suggestion that we "redefine" Turing's test to be one of these impoverished versions. Let it stand with its present name, even if it is a goal never achieved. Changing the name will only add to the confusion illustrated by the recent Forum contributions.

It is surprising and, yes, irksome to find widespread misunderstanding of the Turing Test, one of the basic pieces of AI wisdom and lore.

COMPUTER POWER AND HUMAN REASON

Comments by Kuipers, McCarthy, and Weizenbaum

Joseph Weizenbaum, *Computer Power and Human Reason*, W. H. Freeman Co., San Francisco, 1975.

REACTIONS TO WEIZENBAUM'S BOOK

From: Benjamin Kuipers

April 24, 76

MIT AI Lab Cambridge, Mass 02139 BEN@MIT-AI

*"There are more things in Heaven and Earth, Horatio,
than are dreamt of in your philosophy."*

-- Hamlet, Act I, Scene 5.

I had some strong reactions to Joe Weizenbaum's book, *Computer Power and Human Reason*. The book mentions some important concerns which are obscured by harsh and sometimes shrill accusations against the Artificial Intelligence research community. On the whole, it seems to me that the personal attacks distract and mislead the reader from more valuable abstract points. I strongly recommend Samuel Florman's article "In Praise of Technology" in the November, 1975, issue of *Harper's Magazine* to see a different opinion about the role of technology in modern society.

Some of the points below restate concerns which seem to have motivated Weizenbaum to write his book. Others are my own reactions to issues which he raises. In either case, I see ideas like these as being quite current in the AI community, so I was quite puzzled by Weizenbaum's vehement attacks on us for not sharing them.

1. It is important for a scientist to realize that the descriptive methods of his field capture only one aspect of the phenomena he studies.

2. Point 1, notwithstanding, it is a matter of personal faith whether there are aspects of the world which cannot be fully described by some scientific (i.e., empirical) method. It is clear, of course, that many important aspects of the world are beyond our current scientific methods.

3. A scientist should recognize the difference between descriptive and prescriptive statements. Descriptive statements can be based on scientific investigation; prescriptive statements are based on values. A belief that value judgments are trivial can lead the unwise to believe that prescriptive conclusions follow directly from descriptive data.

4. JW says "The very asking of the question, 'What does a judge (or a psychiatrist) know that we cannot tell a computer?' is a monstrous obscenity." (p. 226) On the contrary, it is a fantastically interesting and important question, deserving the attention of serious thinkers. The question is essentially, "What is the difference between wisdom and knowledge?" To declare the asking of such a question obscene is anti-intellectualism at its most blatant. What actually seems to worry JW, however, is not the question, but the potential for a foolish answer.

5. Assuming that we find it possible to build an intelligent computer, there will inevitably be an enormous cultural gulf between it and humans. Social scientists can say a great deal about the amount of common culture which is required between a professional and a client in many cases, such as a psychiatrist or a judge. This could make the use of a computer in one of these roles inappropriate as a technical judgment, rather than as a moral judgment.

6. It seems exceedingly unlikely that the very difficult problems of intelligence can be solved by "hackers" without a deep theory. The primary goal of AI is to develop the computational techniques which will allow such a theory to be formulated precisely. This often requires intimate acquaintance with deep

structure of knowledge. The papers in this session cover all ends of natural language research: parsing, generating, and memory and inference. The advances to be reported on here indicate that the possibility of natural communication with machines may be closer than had been anticipated.

Machine Understanding of Human Intentionality
Robert Wilensky, Yale.

Inferential Question-Answering in a Textual Data Base
Robert F. Simmons, Univ. of Texas.

Prosodic Generation Research
Arvin Levine, Stanford Univ.

A Design for a Wait-and-See Parser for English
Mitchell Marcus, MIT.

FUTURE CONFERENCES

Canadian Society for Computational Studies of Intelligence
CSCSI/SCEIO Summer Conference 1976

Vancouver, British Columbia, August 25-27, 1976

The conference will be held at the University of British Columbia and will feature submitted papers, tutorial talks, and informal sessions. The conference fee of \$40.00 (\$15.00 for students) will enable participants to attend the sessions and to receive the pre-printed conference proceedings.

If you expect to attend, please notify the general chairman, including the number of people in your group and whether you intend to submit an abstract. A subsequent mailing will include details and a booking form for the accommodation available and a conference registration form.

Papers are requested from any of the following areas:
Natural Language Understanding (Text & Speech)
Heuristic Problem Solving and Game Playing
Automatic Programming and Debugging Computer Perception
Psychological Aspects of AI Automatic Theorem-Proving
Knowledge-Based Learning Systems
Representation of Knowledge Applications of AI Robots
Social Consequences of AI

For further details on conference registration and accommodation write to the General Chairman.

<i>General Chairman</i>	<i>Program Chairman</i>
<u>Richard S. Rosenberg</u>	<u>Alan K. Mackworth</u>
Department of Computer Science	
University of British Columbia	
Vancouver, B.C., V6T 1W5, Canada	

SCENE SEGMENTATION BY CLUSTER DETECTION IN COLOR SPACES

B. J. Schacter, L. S. Davis and A. Rosenfeld
Computer Science Center¹
University of Maryland
College Park, MD 20742

Abstract. Ohlander [1] has shown that a variety of scenes can be segmented into meaningful parts by histogramming the values of various point or local properties of the scene; extracting the region whose points gave rise to that peak; and repeating the process for the remainder of the scene. A generalization of this histogram analysis approach is to map the points of the scene into a multi-dimensional feature space, and to look for clusters in this space (a histogram is a mapping into a one-dimensional feature space, in which clusters are peaks). This note illustrates how one of Ohlander's scenes, a house, can be reasonably segmented by mapping it into a three-dimensional color space.

Ohlander [1] has demonstrated that a variety of scenes can be segmented into meaningful parts by a histogram analysis process which proceeds essentially as follows:

- 1) Histograms are constructed of the values, over the scene, of various local or point properties
- 2) A histogram that has a sharply defined peak is selected
- 3) The region whose points gave rise to that peak is extracted
- 4) The process is repeated for the remainder of the scene

Histogram peak detection is one-dimensional cluster detection; the peak defines a region in a one-dimensional feature space which is densely populated in comparison with the neighboring regions. A generalization of Ohlander's approach would be to map the scene points into a multi-dimensional feature space and look for clusters in that space. This should yield more reliable segmentations of the scene, since the histograms are projections of the multi-dimensional space onto individual axes, and there may exist regions which give rise to distinctive clusters in the space but whose projections on each axis overlap those of other regions.

In this note we use a simple clustering technique to segment one of Ohlander's scenes, a house, by detecting clusters in a three-dimensional color space. Clustering in color space is the standard method of segmenting and classifying regions in multispectral remote sensor imagery. It has been applied to color scene analysis by Ito [2] and by Tenenbaum, *et al.* [3].

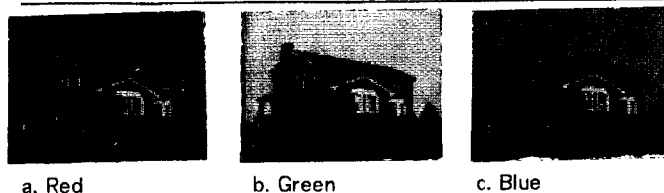


Figure 1. Primary color components of the house scene.

1. The Support of the National Science Foundation under Grant MCS-72-03610 is gratefully acknowledged, as is the help of Shelly Rowe in preparing this paper.

Figure 1 shows the red, green, and blue component images of Ohlander's house scene. The points in this scene were mapped into (red, green, blue) space; i.e., each point maps into the triple of numbers consisting of its values in the red, green, and blue images. Figure 2 shows the projections of this space on to the $(x,y) = (\text{red}, \text{green})$, $(\text{red}, \text{blue})$, and $(\text{green}, \text{blue})$ planes. In these projections, the number of scene points mapping into a given position on the plane is displayed as a gray level (high values are darker). It can be seen that the clusters are rather differently grouped in these three projections, and that in the projections on the axes (i.e., the red, blue and green histograms), the clusters cannot be as well separated as they are in a higher-dimensional space.

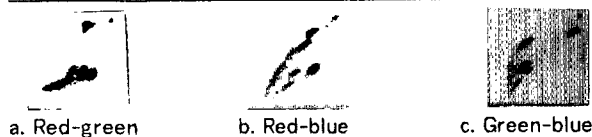


Figure 2. Projections of clusters for house pictures.

The three-dimensional "histogram" can be efficiently stored as a binary tree by using the triples of color values as keys [4], and the number of scene points having each value triple as the stored information. To construct the tree, we examine each scene point and search the tree for its key; if the key is found, the point count at that tree node is incremented by 1, and if not, a node with that key is added to the tree with its count set to 1. To detect clusters, we first find the tree nodes whose point counts are above some threshold t , and we regard such nodes as belonging to the same cluster if they lie within some distance d of each other. (In the experiment reported here, we use $t = 4$ and $d = 1$). The ranges of values, along each dimension, of the points in a cluster define a rectangular parallelepiped bounding the cluster, and any node lying in this parallelepiped can be added to the cluster.

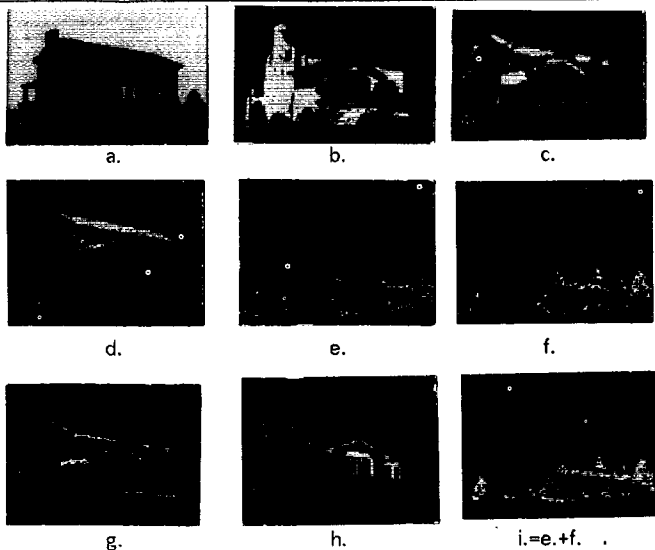


Figure 3. Regions corresponding to the cluster in color space.

Eight clusters were extracted in this way from the (red, green, blue) data for the house scene. The scene points belonging to each of these clusters are shown in Figure 3. The resulting regions can be interpreted as follows:

- a) Sky, including reflections
- b) Brick (sunlit)
- c) Brick (in shadow)
- d) Roof tile
- e,f) Trees and bushes (sunlit and in shadow?)
- g) White trim (sunlit)
- h) White trim (in shadow)

The regions are somewhat noisy, and do not always correspond perfectly to semantically significant objects. However, they represent a finer degree of segmentation than could be expected to be obtained, at the first level of subdivision, by analysis of one-dimensional histograms. Those results suggest that multi-dimensional clustering techniques for scene segmentation deserve further investigation.

REFERENCES

1. R. B. Ohlander, Analysis of natural scenes, Ph.D. Dissertation, Computer Science Department, Carnegie-Mellon University, Pittsburgh, PA, April 1975.
2. T. Ito, Color picture processing by computer, Proc. 4th Intl. Joint Conf. on Artificial Intelligence, Tbilisi, USSR, Sept. 1975, 635-642.
3. J. M. Tenenbaum and S. Weyl, A region-analysis subsystem for interactive scene analysis, *ibid.*, 682-687.
4. J. L. Bentley, Multidimensional binary search trees used for associative searching, *Comm. ACM* 18, Sept. 1975, 509-516.

A SELECTED BIBLIOGRAPHY ON COMPUTER VISION

Martin D. Levine

Dept. of Electrical Engineering¹

McGill University Montreal, Quebec, Canada

Although a few books have been written on the topic of image processing, these have dealt in the main exclusively with the low level aspects of the subject. To a certain extent this is due to the fact that a degree of stability of content has been reached in this field. However, no book is available which covers the more advanced topics in an organized and wide-ranging fashion. The major reason for this deficiency can be attributed to the volatile and unstable nature of research in this area at present. From the research point of view this is a healthy situation but this makes it very difficult to create an advanced graduate course on the subject or to obtain an overview of the field. To overcome this difficulty, this paper presents a selected bibliography which can be used as the basis for such an advanced course on computer vision. The course as it was given in the Department of Electrical Engineering at McGill was modular in design. Students taking the course would already have had a full semester course on image processing. Each of the thirteen sets of articles was maintained in a separate container which the student presented himself for a protracted discussion with the instructor. Since multiple copies of each module were not maintained, the prerequisite structure shown in Figure 1 was employed to allow several students to work on the material at the same time.

The selection of the material naturally represents personal biases and in no way is it intended that the bibliography be

1. The research described in this paper was partially supported by the National Research Council under grant no. A4156.

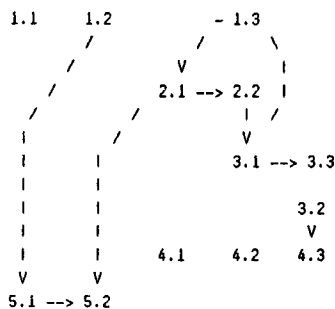


Figure 1. Suggested prerequisite structure. The preferred order is sequential (1.1, 1.2, . . . , 5.1, 5.2) but variations as shown above are equally good. Wherever no arrow leaves or enters a node, it may be considered more or less on its own although it is preferable to have read the modules preceding it in the sequence.

complete in any sense. Rather the intention was to give the students a comprehensive overview of the field. The first chapter is concerned with an introduction to artificial intelligence, a discussion of psychology as it relates to certain aspects of vision, and some important early work. The next four chapters deal with current research and emphasize perception of polyhedra, natural scene analysis, special problems, and knowledge representation, respectively. The detailed bibliography follows below.

Overview

1. Introduction
 - 1.1 Introduction to AI
 - 1.2 Psychology and Vision
 - 1.3 Early Work
2. Perception of Polyhedra
 - 2.1 Polyhedron Recognition - Abstract
 - 2.2 Polyhedron Recognition - Real Data
3. Natural Scene Analysis
 - 3.1 Region Analysis
 - 3.2 Organizing Local Features
 - 3.3 Image Understanding
4. Special Problems
 - 4.1 Object Recognition
 - 4.2 Face Recognition
 - 4.3 Waveform Analysis
5. Knowledge Representation
 - 5.1 Image Descriptions
 - 5.2 Learning

1. INTRODUCTION

1.1 Introduction to AI

- Good, I. J., Human and Machine Intelligence: Comparisons and Contrasts, *Impact of Science on Society*, Vol. XXI, No. 4, 1971, pp. 305-322.
- Minsky, M., Steps Toward Artificial Intelligence, in *Computers and Thought*, E. A. Feigenbaum and J. Feldman, ed., McGraw-Hill, 1963, pp. 406-475.
- Minsky, M. L., Introduction, in *Semantic Information Processing*, M. L. Minsky, editor, the MIT Press, 1968, pp. 1-32.
- Nievergelt, J. and Farrar, J. C., What Machines Can and Cannot Do, *Computing Surveys*, Vol. 4, No. 2, June 1972, pp. 81-95.
- Nilsson, N. J., Artificial Intelligence, *IFIP Congress 1974*, Stockholm, Sweden, Aug. 5-10, 1974.

1.2 Psychology and Vision

- Minsky, M., Papert, S., Research on the Laboratory in Vision, Language, and other Problems of Intelligence, Mass. Institute of Technology, AI Lab, Memo 252, Jan. 1, 1972.
- Newell, A., Remarks on the Relationship Between Artificial Intelligence and Cognitive Psychology, in *Theoretical Approaches to Non-Numerical Problem Solving, Part IV*, R. Banerji and M. D. Mesarovic (eds.), Springer-Verlag, New York, 1970.
- Price, K., A Comparison of Human and Computer Vision Systems, *SIGART Newsletter*, No. 50, Feb. 1975, pp. 5-10.
- Pylyshyn, Z. W., Minds, Machines, and Phenomenology: Some Reflections on Dreyfus', "What Computers Can't Do". The Univ. of Western Ontario, Department of Psychology, Research Bulletin No. 235, June 1972.
- Pylyshyn, Z. W., What the Mind's Eye Tells the Mind's Brain: A Critique of Mental Imagery, *Psychological Bulletin*, Vol. 80, pp. 1-24, 1973.
- Thomas, A. J., Binford, T. O., Information Processing Analysis of Visual Perception, A Review, Stanford Univ., AI Lab, Memo AIM-227, June 1974.

1.3 Early Work

- Evans, T. G., A Program for the Solution of Geometric-Analogy Intelligence Test Questions, in *Semantic Information Processing*, M. L. Minsky editor, The MIT Press, 1968, pp. 211-353.
- Roberts, L. G., Machine Perception of Three-Dimensional Solids, *Symposium on Optical and Electro-Optical Information Processing Technology*, Boston, 1964, pp. 159-197.
- Sutherland, I. E., Sketchpad: A Man-Machine Graphical Communication System, Mass. Institute of Technology, Lincoln Lab, Tech. Report No. 296, Jan. 30, 1963.
- Turner, K. J., Scene Analysis and Object Recognition: A Survey, Univ. of Edinburgh, Dept. of Machine Intelligence and Perception, Experimental Programming Reports: No. 21, July 1971.

2. PERCEPTION OF POLYHEDRA

2.1 Polyhedron recognition - abstract

- Clowes, M. B., On Seeing Things, *Artificial Intelligence*, Vol. 2, No. 1, 1970, pp. 79-116.
- Guzman, A., Computer Recognition of Three-Dimensional Objects in a Visual Scene, Mass. Institute of Technology, Ph.D. Thesis, MAC-TR-59, Dec. 1968.
- Huffman, D. A., Impossible Objects as Nonsense Sentences, *Machine Intelligence*, Vol. 6, ed. B. Meltzer and D. Mitchie, American Elsevier, 1971, pp. 295-323.
- Michie, D., On Not Seeing Things, Dept. of Machine Intelligence and Perception, Univ. of Edinburgh, Experimental Programming Reports: No. 22, September, 1971.
- Waltz, D., Understanding Line Drawings of Scenes with Shadows, in *The Psychology of Computer Vision*, P. H. Winston, ed., McGraw-Hill 1975, pp. 19-91.

2.2 Polyhedron recognition - Real Data

- Ambler, A. P., Barrow, H. G., Brown, C. M., Burstall, R. M., Popplestone, R. J., A Versatile Computer-Controlled Assembly System, *Proc. 3rd Inter. Joint Conf. on AI*, Stanford Univ., Aug. 20-23, 1973, pp. 298-307.
- Duda, R. O., Hart, P. E., Experiments in Scene Analysis, Stanford Research Institute, Tech. Note 20, Jan. 1970.
- Duda, R. O., Some Current Techniques for Scene Analysis, Stanford Research Institute, AI Group Tech. Note 46, Oct. 1970.
- Falk, G., Computer Interpretation of Imperfect Line Data as a Three-Dimensional Scene, Stanford Univ., AI Lab, Memo AIM-132, Aug. 1970.
- Grape, G. R., Model Based (Intermediate Level) Computer Vision, Stanford Univ., Comp. Sci. Dept., AI Memo AIM-201, May 1973.

Shirai, Y., Analyzing Intensity Arrays Using Knowledge About Scenes, in *The Psychology of Computer Vision*, P. A. Winston, ed. McGraw-Hill 1975, pp. 93-113.

3. NATURAL SCENE ANALYSIS

3.1 Region analysis

- Barrow, H. G., Ambler, A. P., Burstall, R. M., Some Techniques for Recognizing Structures in Pictures, *Inter. Conf. on Frontiers of Pattern Recognition*, Honolulu, Hawaii, Jan. 1971.
- Barrow, H. G., Popplestone, R. J., Relational Description in Picture Processing, *Machine Intelligence*, Vol. 6, 1971, pp. 377-396.
- Brice, C. R., Fennema, C. L., Scene Analysis Using Regions, Stanford Research Institute, AI Group, Tech. Note 17, April 1970.
- Harlow, C. A., Image Analysis and Graphs, *Computer Graphics and Image Processing*, Vol. 2, 1973, pp. 60-82.
- Harowitz, S. L., Pavlidis, T., Picture Segmentation by a Directed Split-and-Merge Procedure, *2nd Inter. Joint Conf. on Pattern Recognition*, Aug. 13-15, 1974, Copenhagen, Denmark, pp. 424-433.
- Preparata, F. P., Ray, S. R., An Approach to Artificial Nonsymbolic Cognition, *Information Sciences*, Vol. 4, 1972, pp. 65-86.
- Yakimovsky, Y., Scene Analysis Using a Semantic Base for Region Growing, Stanford Univ., AI Lab, Memo AIM-209, June 1973.
- Zucker, S. W., Region Growing: Childhood and Adolescence, Univ. of Maryland, Comp. Sci. Tech. Report TR-370, April 1975.

3.2 Organizing Local Features

- Chandrasekaran, B., and Kanal, L., On Linguistic, Statistical, and Mixed Models for Pattern Recognition, Computer and Information Science Research Center, Tech. Report OSU-CISRC-TR-71-3, The Ohio State Univ., Columbus, Ohio, March 1971.
- Evans, T. G., A Grammar-Controlled Pattern Analyzer, *IFIP 1968*, Amsterdam 1969, pp. 1592-1598.
- Gips, J., A Syntax-Directed Program that Performs a Three-Dimensional Perceptual Task, *Pattern Recognition*, Vol. 6, No. 3/4, Dec. 1974, pp. 189-199.
- Levi, G., Sirovich, F., Structural Descriptions of Fingerprint Images, *Information Sciences*, Vol. 4, No. 4, Oct. 1972, pp. 327-355.
- Miller, W. F. and Shaw, A. C., Linguistic Methods in Picture Processing - A Survey, *FJCC*, pp. 279-290, Dec. 1968.
- Moayer, B., Fu, K. S., A Syntactic Approach to Fingerprint Pattern Recognition, *Pattern Recognition*, Vol. 7, No. 1/2, June 1975, pp. 1-23.
- Ramer, E. U., Computer Edge Extraction from Photograph of Curved Objects, New York Univ., Tech. Report CRL-34, Dec. 1973.
- Tsuji, S., Fujiwara, R., Linguistic Segmentation of Scenes into Regions, *2nd Inter. Joint Conf. on Pattern Recognition*, Aug. 13-15, 1974, Copenhagen, Denmark, pp. 104-108.

4. SPECIAL PROBLEMS

4.1 Object Recognition

- Agin, G. J., Representation and Description of Curved Objects, AI Lab, Memo AIM-173, Stanford Univ., Oct. 1972.
- Agin, G. J., Binford, Thomas O., Computer Description of curved Objects, *3rd Inter. Conf. on Artificial Intelligence*, Stanford Univ., Stanford, CA, Aug. 1973.
- Horn, B., Obtaining Shape from Shading Information, in *The Psychology of Computer Vision*, P. H. Winston, ed., McGraw-Hill, 1975, pp. 115-155.
- Levine, M. D., Computer Determination of Depth Maps, *Computer Graphics and Image Processing*, Vol. 2, 1973, pp. 131-150.
- Nevatia, R., Binford, T. O., Structured Descriptions of Complex Objects, *Proc. 3rd Inter. Joint Conf. on AI*, Stanford, Aug. 20-23, 1973, pp. 641-647.
- Shirai, Y., Recognition of Polyhedrons with a Range Finder, *Pattern Recognition*, Vol. 4, 1972, pp. 243-250.

4.2 Face Recognition

- Fischler, M. A., Elschlager, R. A., The Representation and Matching of Pictorial Structures, Lockheed Palo Alto Research Lab, Information Sciences, Report LMSC-D243781, May 1972.
- Goldstein, A. J., Harmon, L. D., Lesk, A. B., Identification of Human Faces, *Proc. IEEE*, Vol. 59, No. 5, May 1971, pp. 748-760.
- Goldstein, A. J., Harmon, L. D., Lesk, A. B., Man-Machine Interaction in Human-Face Identification, *Bell System Tech. J.*, Vol. 51, No. 2, Feb. 1972, pp. 399-427.
- Kelly, M. D., Visual Identification of People by Computer, Stanford Univ., AI Lab, Memo AIM-130, Aug. 1970.
- Sakai, T., Nagao, M., Konde, T., Computer Analysis and Classification of Photographs of Human Faces, *1st USA-Japan-Computer Conf. Proceedings*, pp. 2-7-1 to 1-7-8, Oct. 1972.

4.3 Waveform Analysis

- Ehrich, R. W., Foith, J. P., Representation of Random Waveforms by Relational Trees, Univ. of Mass. at Amherst, Electrical and Computer Engineering, June 1975.
- Lozano-Perez, T., Parsing Intensity Profiles, Mass. Institute of Technology, AI Lab, AI Memo, 329, May 1975.
- Pavlidis, T., Waveform Segmentation Through Functional Approximation, *IEEE Trans. on Computers*, Vol. C-22, No. 7, July 1973, pp. 689-697.
- Pavlidis, T., Horowitz, S. L., Piecewise Approximation of Plane Curves, *Proc. of the 1st Inter. Joint Conf. on Pattern Recognition*, Washington, D.C., Oct. 1, 1973, pp. 396-405.

5. KNOWLEDGE REPRESENTATION

5.1 Image Descriptions

- Baird, M. L., Kelly, M. D., A Paradigm for Semantic Picture Recognition, *Pattern Recognition*, Vol. 6, No. 1, June 1974, pp. 61-74.
- Barrow, H. G., Tenenbaum, J. M., Representation and Use of Knowledge in Vision, *SIGART Newsletter* No. 52, June 1975, pp. 2-9.
- Firschein, O., Fischler, M. A., Describing and Abstracting Pictorial Structures, *Pattern Recognition*, Vol. 3, 1971, pp. 421-443.
- Firschein, O., Fischler, M. A., A Study in Descriptive Representation of Pictorial Data, *Pattern Recognition*, Vol. 4, No. 4, 1972, pp. 361-377.
- Frijda, N. H., Simulation of Human Long-Term Memory, *Psychological Bulletin*, Vol. 77, No. 1, 1972, pp. 1-31.
- Kunii, T., Weyl, S., Tenenbaum, J. M., A Relational Data Base Schema for Describing Complex Pictures with Color and Texture, *Proc. 2nd Inter. Joint Conf. on Pattern Recognition*, Lyngby-Copenhagen, Denmark, Aug. 1974, pp. 310-316.
- Minsky, M., A Framework for Representing Knowledge, in *The Psychology of Computer Vision*, P. H. Winston, ed., McGraw-Hill, 1975, pp. 211-277.
- Rosenfeld, A., Models, Univ. of Maryland, Comp. Sci. Center, Tech. Report N00014-67-A-0239-0062, Jan. 1973.
- Winograd, T., Representation: Formalisms for Knowledge, Lecture 3 in Five Lectures on Artificial Intelligence, Stanford Univ., AI Lab, Memo AIM-246, Sept. 1974.

5.2 Learning

- Underwood, S. A., Coates, C. L., Visual Learning and Recognition by Computer, The Univ. of Texas at Austin, Department of Electrical Engineering, Tech. Report No. 123, April 6, 1972.
- Winston, P. H., Learning Structural Descriptions from Examples in *The Psychology of Computer Vision*, P. H. Winston, ed., McGraw-Hill, 1975, pp. 157-209.

CHESS and BACKGAMMON

Is Brute Force Backgammon Possible?

David Levner

Computer Science Department

Carnegie-Mellon University Pittsburgh, PA 15213

LEVNER@CMU-10A

In view of the recent interest in backgammon, two mathematical questions can be asked: how many backgammon positions are there, and how hard would it be to solve by brute force?

The first question shall be answered here, but the second is very difficult. Von Neumann and Morgenstern¹ solved all finite zero-sum games of perfect information of bounded length; but backgammon can continue arbitrarily long since it is possible to repeat a position arbitrarily often; worse, backgammon is not a finite game because the players can double the stakes any number of times. No known algorithm computes the optimal strategy, nor has anyone proven that such a strategy exists.

The number of backgammon positions is much easier to compute. The number of times the players have doubled cannot be considered for then there would be infinitely many positions. Only the three states of possession of the doubling cube (it can belong to player one, player two, or both players) are counted.

A backgammon board consists of three areas: white's offboard (which is also black's bar), black's offboard (which is also white's bar), and the twenty-four points of the backgammon board. There are fifteen white men and fifteen black. Men of opposite color can coexist in all areas, but they cannot coexist on any of the 24 points. A backgammon position is determined by seven factors: the player on the move, the state of possession of the doubling cube, the numbers of white men on the bar, or borne off, the numbers of black men on the bar, or borne off, and the occupation of the 24 points by the remaining men.

It is useful to define $B(m,n)$, the number of ways to arrange m men of a single color on n points (or n areas). To calculate $B(m,n)$, $m+1$ cases are distinguished: case i ($0 \leq i \leq m$) corresponds to placing i men on the first point, leaving $m-i$ men to be arranged on $n-1$ points, which can be done in $B(m-i, n-1)$ ways. Summing the $m+1$ cases yields equation (1):

$$(1) \quad B(m,n) = \sum_{i=0}^m B(m-i, n-1)$$

Note that every position is counted once and only once. Table 1 lists the values of $B(m,n)$ for $m=0..5$ and $n=1..6$.

Table 1

m	0	1	2	3	4	5
n						
1	1	1	1	1	1	1
2	1	2	3	4	5	6
3	1	3	6	10	15	21
4	1	4	10	20	35	56
5	1	5	15	35	70	126
6	1	6	21	56	126	252

There are $B(15,3)$ ways to allot the white men among the three areas, and just as many ways for the black, for a total of $B(15,3)^2 = 136^2 = 18496$ possibilities. Each possibility leaves a

1. Morgenstern, O. and Von Neumann, J. *The Theory of Games and Economic Behavior*, Princeton University Press, Third edition, Princeton NJ, 1953.

certain number of men, say w white men and b black men, to be distributed over the 24 points.

To compute the number of occupations of w white men and b black men on 24 points, w cases are distinguished. Case i ($1 \leq i \leq w$) corresponds to placing the w white men on exactly i points. At least one white man must occupy each of the i points, and the remaining $w-i$ men can be arranged on the i points in $B(w-i, i)$ ways. There are

$$\binom{24}{i}$$

sets of i points. Once the white men are in place, $B(b, 24-i)$ arrangements of the b black men on $24-i$ points are possible. Hence case i contains a total of

$$\binom{24}{i} B(w-i, i) B(b, 24-i)$$

occupations of the white and black men. The full number is obtained from the sum of the w cases:

$$(2) \quad oc(w,b) = \sum_{i=1}^w \binom{24}{i} B(w-i, i) B(b, 24-i)$$

Every occupation has been counted exactly once.

The number of backgammon positions (BKG) is calculated by equation (3).

$$(3) \quad BKG = \sum_{\text{for all } \langle p,d,woff,wbar,boff,bbar \rangle} oc(15-woff-wbar, 15-boff-bbar)$$

where p is a player, d is a state of possession of the doubling cube, $woff$ is the number of white men borne off, $wbar$ the number of white men on the bar, and $boff$ and $bbar$ are defined similarly. Note that there are two players and three possible states of the doubling cube. BKG was calculated from these three equations and found to be 1.1×10^{20} .

7th U.S. Computer Chess Championship

At ACM Conference -- Houston, October 19-21, 1976

Entries are being solicited to the Seventh U.S. Computer Chess Championship to be held October 19-21, 1976, in conjunction with the ACM Annual Conference in the Hyatt Regency Hotel, Houston, Texas. A four round Swiss style tournament is planned with the first round on Tuesday, October 19th, beginning at 8 p.m., and the final round on Thursday, October 21st, at 8 p.m. The field will be limited to 12 teams. David Levy, International Master from England, will serve as tournament director.

For further information and for an application form, write to Monroe Newborn, School of Computer Science, McGill University, Montreal, Quebec H3C 3G1, Canada.

ABSTRACTS

Understanding LISP Programs is Improving LISP Programs (in French)
Harold Wertz

U.E.R. Informatique & Linguistique Université de Paris 8

A system designed to understand and improve programs has to combine reasoning about programming with a good deal of knowledge about the programming language used. Our program-improving-system detects and eliminates two classes of errors: *surface-errors*, detectable by local analysis, and *deep-errors*, detectable only by global analysis. We describe some knowledge and some tactics of reasoning needed to improve programs supposed to be written in a subset of LISP: *first-order-LISP*. This subset is nicely suited to describe the know-how of a human programming apprentice. We illustrate these concepts by a hand