1-1-1988

Volume 1, Number 1 - CogSci News (Spring 1988)

Lehigh University Cognitive Science Program

CogSci News

Cognitive Science Program, Lehigh University, Bethlehem, PA 18015

Volume 1, Number 1
Spring 1988

Editorial Policy

This Newsletter is published twice each year, in fall and spring issues, by the Cognitive Science Program at Lehigh University. Its purpose is to inform faculty and students about the interdisciplinary and rapidly growing field of cognitive science, as well as to report the activities of Lehigh's Program.

The Newsletter is distributed free of charge in the United States and Canada to academic programs and individuals interested in cognitive science. Anyone who would like to be added to the mailing list should notify the Editor.

The Editorial Staff welcomes readers' comments and *solicits materials* dealing with cognitive science. We are especially pleased to consider course syllabi, book reviews, short essays, brief descriptions of scholarship and research in progress, and original art work (e.g.,cartoons, line-drawings, computer-generated images).

Address all submissions, comments, and subscription requests to: The Editor, Cognitive Science Newsletter, Lehigh University, Bethlehem, PA 18015. Electronic mail can be sent via Bitnet to JBG1@LEHIGH.

LEHIGH UNIVERSITY OFFERS NEW B.A. IN COGNITIVE SCIENCE

Cognitive science is a recent arrival on college campuses. According to the national Cognitive Science Society, only 30 or so institutions in the United States currently offer an undergraduate degree in this exciting new field. Lehigh University is proud to be among this innovative group. Two years of curricular planning culminated in May, 1987, when the university's faculty approved by unanimous vote a new bachelor of arts degree in cognitive science.

Cognitive science is defined as the study of how knowledge is acquired, represented, processed, and transmitted. The central aim is to detail the functional organization of so-called "intelligent systems," whether human, animal, or machine. The underlying premises are that cognition takes place in terms of mental representations of the world and that the activity we usually call "thinking" consists of various operations performed on these internal representations. The disciplines most commonly participating in this joint enterprise are psychology, computer science, philosophy, linguistics, anthropology, and neuroscience.

At Lehigh, the new major and minor are offered through the College of Arts and Sciences. The program's wide-ranging interdisciplinary character prepares students for a variety of careers and/or graduate programs. Because the course requirements draw from several of the College's B.A. "distribution areas," it is relatively easy for students to complete double majors, and this is encouraged.

The B.A. in Cognitive Science requires a minimum of 47 semester hours—40 within the major itself and at least 7 in collateral areas. Specifically required courses in the major begin with a two-semester introductory sequence (CogS 101 & 102), which students take in their sophomore year. The major's core consists of four intermediate level courses, one from each of the most central disciplines—cognitive psychology, artificial intelligence, philosophy, and linguistics. The final integration of coursework occurs in the required senior seminar, in which students focus on a topic of their choice from a branch of cognitive science.

In addition to the specifically required courses, majors complete six major electives. Students fulfill this requirement by choosing three topical areas (from among six offered) and completing two courses in each.

Collateral requirements include introduction to structured programming and data structures (either CSc 11 & 15 or CSc 17) and at least the first semester of calculus. Additional course work in mathematics is strongly recommended, as are introductory psychology (Psyc 1 or 11), introductory biology and laboratory (Bio 21/22), and human evolution (Anth 12).

Students may also minor in Cognitive Science. The minor requires 19 credit hours—the introductory sequence plus the four disciplinary core courses.

Course Descriptions

Introductory Sequence (7 hrs)

CogS 101. Introduction to Cognitive Science I (3) The conceptual underpinnings of cognitive science, its history, and how its constituent disciplines converge on the analysis of intelligent systems.

CogS 102. Introduction to Cognitive Science II (4) The mathematical tools most widely used in cognitive science.

(continued on page 2)

New CogSci Major (cont.)

Disciplinary Core Courses (12 hrs)

CSc 230. Elementary Artificial Intelligence Applications (3) How computers combine elementary operations to do complex jobs. How computers play chess, compose music, simulate psychiatrists, produce medical diagnoses. No previous knowledge of computers required.

Phil 250. The Minds of Men and Robots (3) Is the nature of thinking illuminated by what computers can do? Is the brain just a complex computer? Could a robot feel pain? Be angry? Recent work in artificial intelligence, psychology, and philosophy.

Psyc 117. Cognitive Psychology (3) Information processing by human beings: attention, memory, language, and thought processes.

CogS 140. Introduction to Descriptive Linguistics (3) Survey of descriptive linguistics. Topics include the relationship between language and mind, formal properties of language, language and society, and how languages change over time.

Major Electives (18 hrs)

Two courses each from any three of the following groups.

Artificial Intelligence & Expert Systems:

CSc 262. Programming Languages (3)
CSc 327. Artificial Intelligence Applications (3)
CSc 365. Natural Language Understanding (3, p.req: CSc 262)
CSc 368. Artificial Intelligence Programming (3, p.req: CSc 262)

Formal Models:

Phil 14. Foundations of Logic (3)
Phil 214. Logical Theory (3)
CSc 261. Discrete Structures (3, p.req: Math 21 & CSc 11)
CSc 265. Automata and Formal Grammars (3, p.req: CSc 261)

Philosophy:

Phil 139. Contemporary Philosophy (3)
Phil 220. Knowledge and Justification (3)
Phil 251. Action, Free Will, and Fate (3)

Cognitive Psychology:

Psyc 307. Seminar in Cognition (3, p.req: Psyc 117)
Psyc 320. Psycholinguistics (3)
Psyc 351. Cognitive Development in Childhood (3, p.req: Psyc 117)

Sociocultural Influences on Cognition:

SPsy 135. Human Communication (3)
SPsy 307. Attitudes, Attributions, and Actions (3)
Anth 376. Mind, Self, and Culture (3)

Neuroscience:

Psyc 176. Introduction to Cognitive Neuroscience (3)
Psyc 177. Introduction to Physiological Psychology (3)
Psyc 373. Sensation and Perception (3, p.req: Psyc 176)
Psyc 375. Neuroanatomy of Behavior (3, p.req: Psyc 177)

Senior Seminar (3 hrs)

CogS 301. Senior Seminar in Cognitive Science (3) Integration of the material from cognitive science using topics chosen by the students.

For further information concerning this new major, contact Prof. Edwin Kay, Director, Cognitive Science Program, Lehigh University (215-758-3623).

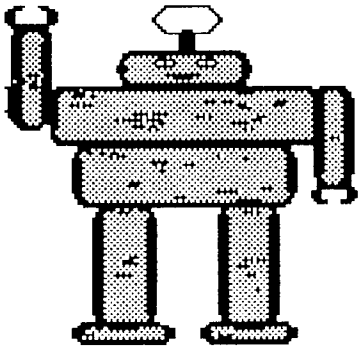

A.I. Active Area at Lehigh

Artificial Intelligence (AI) has recently emerged as a substantial discipline at Lehigh. Several departments are engaged in AI projects in a variety of areas, both in terms of practical applications and scientific study of cognition. The Advanced Technology for Large Steel Structures Engineering Research Center (ATLSS), the Center for Design and Manufacturing Innovation (CDMI), and the Cognitive Science Program all use AI techniques and approaches in their work.

Signaling a commitment to support this growing field, the Department of Computer Science and Electrical Engineering, in its 1987 planning document, identified AI as a major research thrust area. This designation encompasses the activities of most of the Computer Science Division as well as several members of the E.E. Division.

Currently, the department's research focuses on intelligent design and control in industrial applications and on natural language processing. Courses offered by the Computer Science Division cover the full spectrum of Artificial Intelligence—AI languages, object-oriented programming, intelligent robotics, cognitive science, natural language processing, and knowledge-based (expert) systems.

Research and teaching activities are enhanced by the department's new AI laboratory, which features five Sun workstations, two IBM RTs, two Tektronix 4404s, two AT&T 3b2s, and one AT&T 3b15, as well as several AT-compatible microcomputers. A wide range of AI languages and shells are available. There is close cooperation between the AI Lab and others on campus, notably the Intelligent Systems Lab (CSEE/CDMI), the ATLSS Research Center, and the Semiconductor Research Corporation Center for Packaging Research.

Glenn Blank
Computer Science & Elec. Engin.
Lehigh University

BOOK REVIEWS

Lehrer, Adrienne (1983) Wine and Conversation. Bloomington: Indiana University Press. 240+xii Pp.

This is one of the most stimulating and enjoyable books on language I have read in several years. The subject is how people talk about wines and whether, or in what sense, such talk is communicative. More generally, the book is about the conditions that may or may not foster a consensus in the ways people apply words to objects in the world. Its findings and arguments should interest not only a wide academic audience, but also anyone who has been perplexed by the seemingly opaque expressions people use when talking about wines.

The first three chapters deal with the semantics of wine descriptors as these are defined and used in the writings of authorities. Wines are complex stimuli that differ on a number of dimensions, and wine descriptors follow suit. Lehrer presents a preliminary list of 241 commonly encountered expressions and then arranges these into lexical subsets according to the qualities of wines they describe. The most important finding of this semantic analysis is that almost all wine descriptors involve a measure of affective evaluation, that is, the expressions within each lexical subset imply varying degrees of praise with respect to the quality of wines they specify. A second point is that the vocabulary of wine talk is not a closed set, but rather new expressions are coined as appropriate. Many, if not most, of the expressions describing wines are metaphors.

The empirical heart of the book lies in five chapters reporting a series of experimental studies. Taken collectively, these studies were designed to determine (1) whether people identify and describe the same wines with similar terms, (2) whether the amount of time spent conversing about wines increases the group's referential consensus, and (3) whether subjects who are knowledgeable about wines speak with greater precision than less knowledgeable subjects and thus evidence greater referential consensus. The wine talk of three groups of adult subjects, differing in their expertise, was studied. Experimental tasks ranged from writing descriptions for stimulus wines to checking off appropriate descriptors from extensive lists to describing wines so that another person could identify them to discussing wines until a group consensus on the descriptions was reached to semantic clustering tasks for each subject's wine vocabulary. The most striking finding of these experiments is that even when people agree on the relative meaning of wine terms—the intralinguistic meaning of scalar descriptors—they apply these to stimuli differently, apparently as a function of differences in taste preferences.

Chapters 9 through 11 examine the referential imprecision of wine talk from a communications perspective, focusing on why people talk about wines in the first place. Lehrer characterizes wine talk as critical communication, the function of which is not so much to develop shared and stable relations between words and things as to get another person to notice some aspect of a wine. In this sort of informal communicational setting, referential accuracy and precision are not really required. By contrast, in scientific discourse, referential accuracy is at a premium, and Lehrer briefly traces the history of scientific terminologies from several fields. Chapter 12 provides a more extended case history of the problems scientists must overcome in converting natural language vocabulary (personality descriptors) into a more precise scientific terminology.

The book's major contribution is its demonstration, thoroughly accomplished, that the meaning of words involves both intralinguistic and referential relations. Further, consensus as regards the former does not guarantee consensus in the way people connect words to the world. The degree of referential precision is related to the communicational objectives of the speech context. Most conversations do not require referential agreement, whereas scientific and legal discourses do. These conclusions have far-reaching implications. For example, studies that use identification tasks (including communication accuracy tests) as the means of determining "shared meanings" are inadequate because they fail to recognize that meaning consists of both extralinguistic and intralinguistic relations. As Lehrer's analysis of wine talk shows, people may agree on the relative meaning of words and yet disagree on instances.

(Reprinted from Language and Society, 1985, 14:429.)

John B. Gatewood
Social Relations Department
Lehigh University

Haugeland, John (1985) Artificial Intelligence: The Very Idea. Cambridge, Mass: A Bradford Book, MIT Press. 287 Pp.

"The real issue," John Haugeland writes, "has nothing to do with advanced technologies (or corporate specialties), but with deep theoretical assumptions. According to a central tradition in Western philosophy, thinking (intellection) essentially is rational manipulation of mental symbols (viz., ideas)" (p.4). An examination of the promise and the problems of this assumption in the context of contemporary AI is the topic of this clearly written and accessible book.

The historically oriented first chapter illuminates some of the crucial steps (taken by Copernicus, Galileo, Hobbes, and Descartes) which led to the idea that ratiocination just is computation. This Hobbesian formula was made possible by the Copernican divorce between the way the world seemed and the way his mathematics described it, by the way in which Galileo used geometry (so that a line might represent a time, a velocity, or an acceleration just as well as something spatial), and by the Cartesian doubt that suggested

(continued on page 4)

Book Reviews (continued)

Descartes's beloved physics might be just a dream or delusion—without being altered in its intrinsic character one whit! The idea that in thinking the mind might be just engaging in computation over elements intrinsic to itself—a syntactic engine—is the result.

Haugeland shows us that this suggestion poses two fundamental problems. One he calls the Mystery of Original Meaning. Suppose we have elements, formulae composed of them, and a syntactic engine which transforms some of them into others. So far, none of this means anything. Intelligent beings such as ourselves, however, operate not just with marks or noises, but with meaningful elements. The Mystery is where meaning comes from. How does a syntactic engine get to be a semantic engine?

The second problem is called the Paradox of Mechanical Reason. Supposing a mechanism which manipulates the tokens according to syntactic rules, there seem to be two possibilities. Either the manipulator pays attention to what the tokens mean, or it doesn't. But if it does, it apparently can't be entirely mechanical—since meanings (whatever they might be) don't seem to exert physical force. (Homunculus problems lurk here.) On the other hand, if it doesn't pay attention to meaning, its manipulations seem not to be instances of reasoning—since reasoning depends on meaning. So "if a process is mechanical, it can't reason; if it reasons, it can't be mechanical" (p.39).

I have dwelt on these introductory matters because Haugeland sees AI as the most promising tack to date for resolving the paradox and unveiling the mystery. And the rest of the book is an exploration of the contributions that work in AI has made to these deep theoretical questions. Chapter 2 is devoted to a careful examination of the idea of an automatic formal system. Chapter 3 is on semantics, setting out the constraints on interpretation, defining the notion of an interpreted formal system, and reaching a climax in the Formalist's Motto: "If you take care of the syntax, the semantics will take care of itself" (p.106).

The fourth chapter is devoted to computer architecture and discusses in an elementary way Turing machines, von Neumann machines, the structure of LISP, and production systems. Those who know computers will find nothing new here, but I found it a useful survey.

The heart of the book is in the final two chapters, called "Real Machines" and "Real People." In the former, one meets the notions of heuristic search, microworlds, stereotypes, and the frame problem. There are discussions of GPS, SHRDLU, mechanical translation, and expert systems. Particularly useful are the sharply drawn portrayals of the difficulties these programs and ideas have run up against and the problems that remain recalcitrant.

In "Real People" there is a consideration of pragmatic sense, an aspect of meaning in ordinary situations which seems very difficult to formalize. Whether mental images play a role in intelligence is examined, making reference to the well-known "mental rotation" experiments of Shepard and Metzler. There is a categorization of feelings into seven kinds, with some speculations about how (and whether) they might be handled by AI. And finally the question of whether a sense of self or ego- involvement is an essential aspect of intelligent behavior is explored.

It is a rich and varied menu that Haugeland sets before us, and it will be an unusual person who will not find something satisfying on his table. There is no discussion of recent work in parallel processing, but apart from this lacuna, it would be hard to find a more competent and clear survey of AI and what it means for the project of trying to understand intelligence.

Norman Melchert
Philosophy Department
Lehigh University

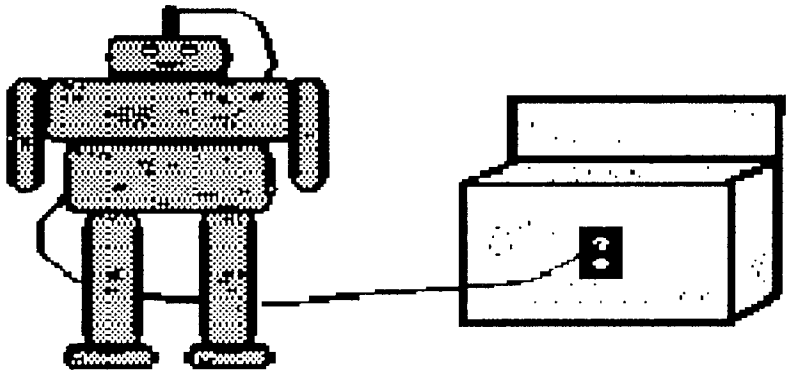

MEETINGS OF INTEREST

Cognitive Science Society

The 10th Annual Conference of the Cognitive Science Society will be held August 17-19, 1988, at McGill University, Montreal, Canada. Topics include problem solving, education and professional applications of cognitive science, language processing, cognitive development, the relationship between cognitive and neural sciences, and recent developments in parallel distributed systems.

Connectionist Models Summer School

The 1988 Connectionist Models Summer School will be held June 17-26, 1988, at Carnegie-Mellon University, Pittsburgh, PA. This special summer school is open only to graduate students. Each session will be led by a major researcher in the field.

Psychonomic Society

The 29th Annual Meeting of the Psychonomic Society will be held November 10-12, 1988, at the Palmer House in Chicago, IL.

EVENTS

SEMINARS

26 Feb 86
Discussion of Gardner's "The Mind's New Science" (philosophy chapter)
Steven Goldman

19 Mar 86
Discussion of Gardner (A. I. chapter)
Glenn Blank

2 Apr 86
Discussion of Gardner (linguistics chapter)
Herbert Rubenstein

9 Apr 86
Discussion of Gardner (psychology chapter)
Martin Richter

16 Apr 86
Discussion of Gardner (anthropology chapter)
John Gatewood

23 Apr 86
Discussion of Gardner (neuroscience chapter)
George Shortess

21 May 86
Decision processes under conditions of risk
John Gatewood

4 Jun 86
Children's acquisition of concepts
Maureen Callanan

18 Jun 86
Formulae and rhetoric in Homer's writings
Julie Williams

2 Jul 86
Natural language syntax and semantics: A register vector grammar
Glenn Blank

16 Jul 86
Theories of categories and categorization
Barbara Malt

30 Jul 86
Artistic visions of the nervous system
George Shortess

16 Mar 87
The mystery of the Chinese room
Robert Barnes

30 Mar 87
Discussion of Dreyfus & Dreyfus's "The Mind of a Machine"
Norman Melchert & Gordon Beam

13 Apr 87
A course in cognitive neuroscience course - What is it?
George Shortess

27 Apr 87
Inference in expert decision-making
John O'Connor

11 May 87
Can thoughts cause behaviors and feelings?
Lloyd Williams

10 Jun 87
Can thoughts cause behaviors? (continued)
Lloyd Williams

24 Jun 87
Human and computer understanding of metaphors
Glenn Blank

1 Oct 87
A distributed-processing model of memory
Robert Barnes

15 Oct 87
Discussion of Lakoff's "Women, Fire, and Dangerous Things"
Barbara Malt

22 Oct 87
Discussion of Lakoff (continued)
Barbara Malt & Donald Campbell

29 Oct 87
Discussion of Lakoff (continued)
Donald Campbell

19 Nov 87
Audio-tape of Noam Chomsky's Swarthmore talk

25 Feb 88
Background to Rumelhart & McClelland's "Parallel Distributed Processing"
Robert Barnes

10 Mar 88
Discussion of Rumelhart & McClelland
Edwin Kay

24 Mar 88
Discussion of Rumelhart & McClelland (continued)
Edwin Kay

7 Apr 88
Discussion of Rumelhart & McClelland (continued)
Edwin Kay

COLLOQUIA

5 Feb 87
Silicon neurons: The design of "neural-net" computers
Larry Jackel

4 Mar 87
The role of fuzzy logic in common sense reasoning
Lotfi Zadeh

5 Mar 87
The management of uncertainty in expert systems
Lotfi Zadeh

9 Apr 87
"Top-down" versus "bottom-up" approaches to the philosophy of mind
Joseph Margolis

5 Aug 87
Information search and retrieval on CD-ROM: Designing and developing applications based on the user's needs
John O. Cole

(continued on page 6)

Colloquia (continued)

29 Sep 87
Human consciousness and the brain
Daniel Dennett

8 Oct 87
What are mental representations (if there are any)?
Daniel Dennett

9 Oct 87
Some problems of consciousness
Daniel Dennett

11 Nov 87
Evolution and rationality
Stephen Stich

10 Dec 87
Against connectionism as a theory of mind
Jerry Fodor

29 Jan 88
Reconceptualizing the process of cognitive development
Kurt Fischer

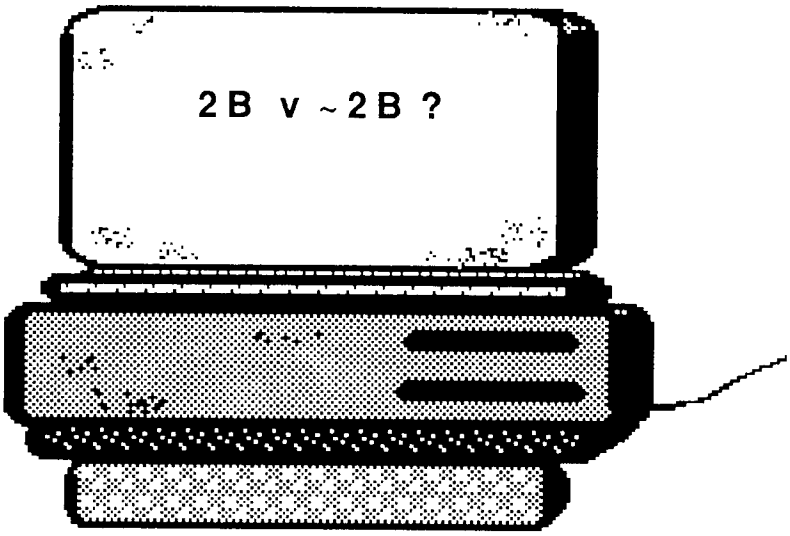

Cognitive Science Program
Lehigh University
Bethlehem, Pennsylvania 18015
U. S. A.

Non-Profit
Organization
U.S. Postage
PAID
Bethlehem, Pa.
Permit #230