Faculty of Arts and Sciences
University of Pennsylvania
December 1979
A Supplement to Almanac

FAS reports

Cognitive Science: Creating New Academic Links

Few programs could better demonstrate the spirit of the University of Pennsylvania's "One University" concept than its new Cognitive Science Group. This program draws from four different departments, computer and information science, linguistics, philosophy and psychology. It also formally links two schools within the University, the Faculty of Arts and Sciences and the School of Engineering and Applied Science.

Continued on next page

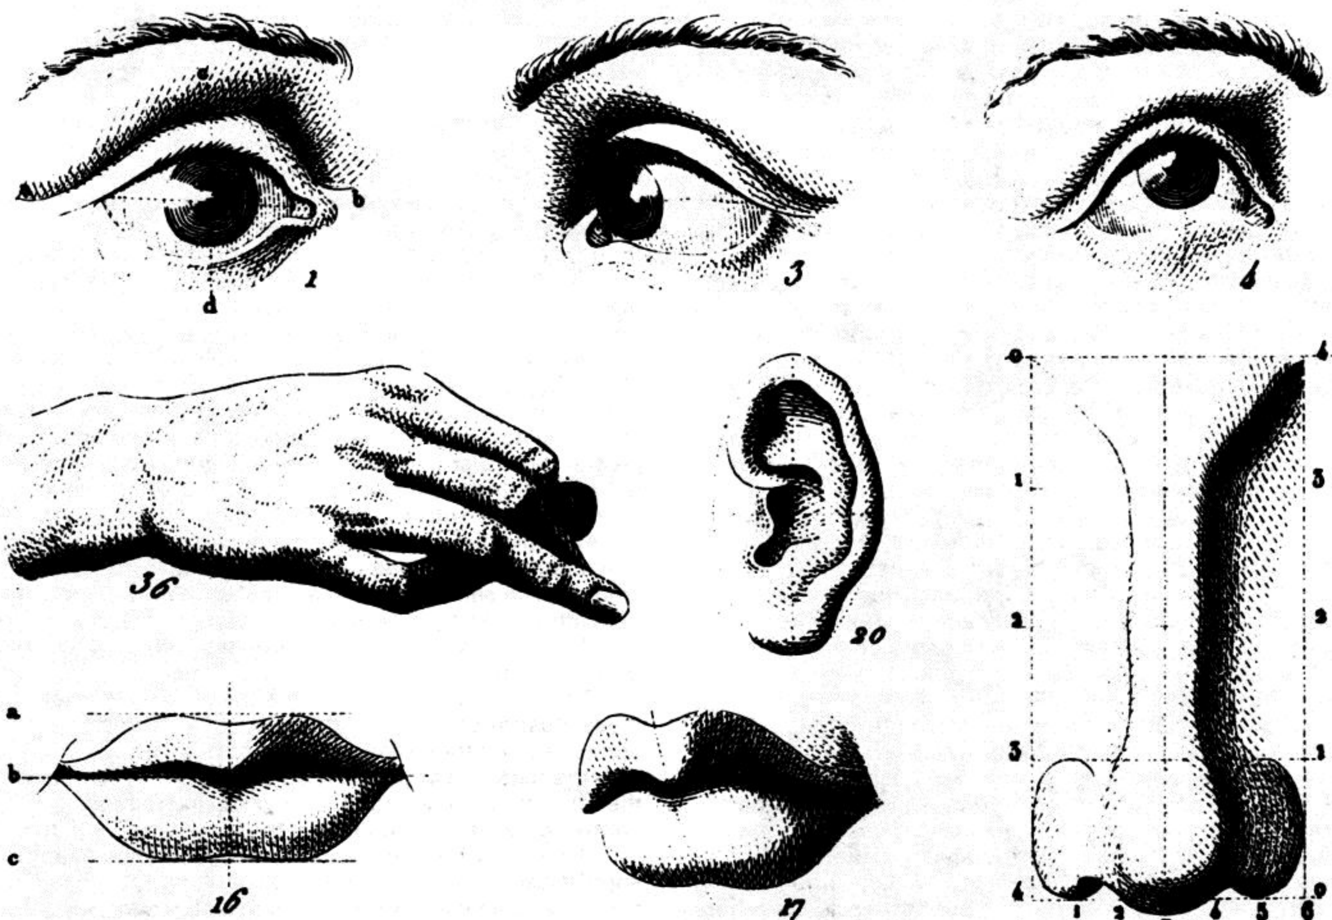

What research interests could a psychologist have in common with a linguist, or a philosopher with a computer scientist? When the issues are intelligence, language and cognition, the objectives of all four disciplines are very similar. For example, one major area of concern in cognitive science is mental representation, or how linguistic and perceptual forms might be represented and manipulated. For a psychologist, the interest is in how people perform certain linguistic and perceptual tasks. A computer scientist, on the other hand, is concerned with the problem of computer representation and how to construct computer systems that carry out similar linguistic and perceptual tasks. Ultimately, a computer scientist hopes to structure a machine that will exhibit to some degree the language and communicative behavior of humans.

Philosophers (especially those concerned with mathematical logic) look at language acquisition from yet a third point of view. In a mathematical framework they work to describe the nature and potential of the human mind. Working in the abstract, philosophers develop mathematical models that may some day provide a step by step analysis of how people learn. As more information is gained about human cognition, the more precise and helpful a philosopher's model becomes.

Of concern to a linguist in cognitive science is how the decisions made about the order and structure of language affect meaning. Much of a linguist's work is looking at word patterns and how sentences are put together to make up conversable language. In addition to apparent meaning in human discourse, linguists look at what lies behind a person's words. They study the implied message that might inspire such a question as, "What did you mean by that remark?"

Pennsylvania's group of linguists, philosophers and psychologists from FAS and computer scientists from SEAS are part of a nationwide program supported by the Sloan Foundation. The Sloan project involves twelve universities and encourages inter-university cooperation as much as it emphasizes working across departmental and school lines. The co-directors of Penn's program, Lila Gleitman (Psychology, FAS) and Aravind Joshi (Computer and Information Science, SEAS), for example, have both done work at other institutions. Gleitman spent a semester at MIT and participated in two workshops at the University of California at Irvine. Joshi was involved in two workshops at Stanford University and one at Irvine. Others in the program, such as Bonnie Webber (Computer and Information Science), Scott Weinstein (Philosophy), Ruzena Bajcsy (Computer and Information Science) and Rochel Gelman (Psyhology) have participated in Sloan workshops at Amherst College, the University of Texas at Austin, Stanford University, the University of California at Irvine and at San Diego. Researchers from schools both formally and informally associated with the Sloan program have reciprocated and sent scholars to participate in the Penn program.

The nature and size of the Sloan project varies from institution to institution. The common thread throughout has been to use language as the tool with which to study cognition. Language, however, is only one possible focus of study. Recently it has become clear to those in Penn's Sloan group that the study of cognition must also focus on perception. As more issues are introduced, additional disciplines and schools may become involved. By bridging departmental and school barriers, the cognitive science group hopes to arrive at a more complete picture of one of the most important of all human concerns, the nature of intelligence.

Talking to the Animals

Questions concerning the nature of language and language learning are asked of animals as well as people. Bees, for example, communicate to each other by sound or movement when honey is found. Several species of birds signal when there is danger. David Premack, Professor of Psychology, is studying chimpanzees to learn more about the nature of intelligence and language. Premack has been working with chimps for thirty years and has succeeded in teaching a chimp to use a vocabulary roughly equivalent to that of a two-year old. This chimp is the now-famous Sarah, who lives in Penn's new laboratory facility at Honeybrook Farms, twenty miles west of Philadelphia. The point for Premack, is not that an ape can learn to use human language. What does concern him are questions like the following: How did the ape learn? How does the animal's cognitive behavior compare to other populations? What do these similarities and differences tell us about human intelligence? Premack put it this way, "The theoretical objective is intelligence and we propose to move in that direction by comparing cognitive performance among several populations—normal children of different ages, retarded children of different ages and different levels of retardation, and chimpanzees of varying ages. We will also compare populations on a number of different bases, such as language, perceptual development, numerical skills and social attribution."

Social attribution is very important to Premack's work. It deals with the states of mind one attributes to another. On a very basic level, social attribution refers to a recognition of the fact that one object has two sides and thus can elicit different interpretations. If a person looks at the convex side of a half-sphere, for example, he will see it differently from a person looking at the concave side of the object. On a more advanced level, social attribution includes the study of how one individual anticipates another individual's wishes or thoughts. Premack is interested in whether or not these skills are exclusive to man. Sarah has shown that at least some of them are not.

Premack was limited to non-verbal forms of expression in his experiments with the chimps. He used flash cards, each illustrating a different activity, and resolution cards that relate to each activity. Each card answered a different question about the activity, such as what colors are there; does the person look happy or sad; what is he doing; why is he doing it. The resolution card selected helps indicate what ability level coincides with a given stage of physical development. For example, the card showing a person jumping up and down beneath a banana tree was put before three-year old chimps. The same test, with an analagous situation (a person beneath a cookie jar) was given to two groups of three-year old children. One group of children was retarded and the other normal.

The three populations solved the problem presented in these illustrations. To Premack's surprise, he discovered, "the chimps solved for intention." They chose cards responding to *why* the person was jumping up and down. "The children, on the other hand, responded to physical likenesses," by choosing cards that duplicated the colors in the picture.

It is widely believed that intention demands a far more sophisticated level of cognition than mimicry. If this theory is true, one logical interpretation of Premack's findings is that chimps have a heightened level of cognitive development at this stage of their lives. Premack, however, is testing another theory which he feels makes more sense.

> Perhaps intentions are the most primitive of cognitive information, more basic even than physical likeness. The earliest instinct in most animals is self protection. When a predator is preparing to attack, the victim's first awareness is intent, not what the attacker looks like!

Comparative studies of this type have pointed out several other facts of learning peculiar to man. In studies of quantities, for example, where different amounts of liquids and solids are shown to children and chimps, Premack notes that a very young child is quick to understand how much of the substance exists. The chimp, on the other hand, is much less apt to deal with amounts and more likely to respond to the color or length of the quantity shown.

In addition to Gleitman's work on language acquisition in deaf and blind children and Premack's work with chimpanzees, there are two other psychologists involved with the Cognitive Science Group: Rochel Gelman, researching number acquisition in children, and Deborah Kemler, working on perception of the retarded. "The more populations included and the more expressions of intelligence we study," Premack concludes, "the more far-reaching and thorough our understanding of cognition becomes."

Sarah: the chimp who learned to speak.

Communication Between Man and Machine: The Interschool Tie

Computers are no longer a novelty of flashing lights in a science fiction movie. They can now be found in the offices of even the smallest businesses. What is new and not well-appreciated is that computers are not simply number crunchers, but are symbolic processors. It is the computer's capacity for symbol manipulative tasks that makes them such interesting objects for investigation in an interdisciplinary area such as cognitive science.

Aravind Joshi, Professor of Computer and Information Science, and Bonnie Webber, Assistant Professor of Computer and Information Science, are part of a team in the School of Engineering and Applied Science working on the communication between man and machine via natural language. The main problems this team of computer scientists has been grappling with are how to store large bodies of complex data and how to make the data accessible to the public in a convenient manner.

Accessibility involves the speed and method by which a piece of information is found. A simple organizational technique such as a list was all that was required initially. But as information became increasingly complex, such simple structures ceased to be adequate. Complicated structures must now be formed to describe the intricate relationships between all the different pieces of data. As Joshi explains, "It is one thing to establish this complex network of relationships; but, once having done it, it is another major problem to use it most effectively." The key issue is to develop techniques that enable people to get at the information, or, more accurately, to communicate with the machine. This is an area of great interest to psychologists as well as computer scientists because it involves an understanding of how we link words to information in our daily language.

Joshi, Webber and a number of graduate students in computer science are approaching the issue of natural language in man-machine communication by using question and answer systems. These systems will establish an interaction between man and machine similar to that between two people who want to share some information. The goal, Webber says, "is to make the machine a bit more helpful and graceful in its relationship with people, and more able to cope with how people actually speak and behave."

Although the computer's ability to handle large amounts of data seems limitless, there remains the continuing problem of flexibility in the accompanying system of communication. Joshi explains, "Rigidity in such a communication network means that the person interacting with the computer must already know exactly how the information is organized. He therefore has to know how to ask just the right question. This is no longer possible; the information held in a computer is now far too complicated for such comprehensive understanding. Instead, a person must be allowed to formulate the query in terms of his own understanding of the data. It is the machine's job to then figure out what that person's view is and try to match it with its own."

Joshi is using a computer system designed by one of his former graduate students, Dr. Samuel J. Kaplan. The system is set up for question and answering tasks to try to learn how best to organize the questioning process. The problem is complicated by the fact that when a question is asked, certain assumptions are made by the person about how the information may have been stored. As an example Joshi says,

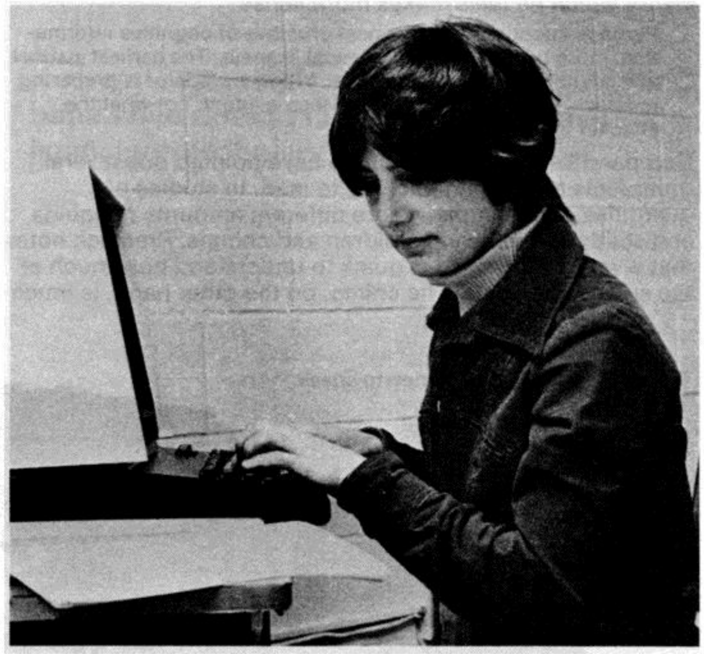

The goal is to make the machine a bit more helpful and graceful in its relationship with people, and more able to cope with how people actually speak and behave.

—Bonnie Webber

"If I ask the computer the question, 'Is there a book written by Lila Gleitman in the library?', most likely I am assuming Gleitman has written a book. I am therefore not only asking the question, 'Is a book written by Gleitman in the library?', but am conveying implicitly the assumption that such a book exists." If that assumption is false, a poorly programmed computer could still give the answer "No," without being aware that the response is due to the fact that Gleitman has not written a book at all. The answer is correct, but misleading. It would allow the questioner to go away still believing that such a book exists, but not in this library. Such behavior has been called "stonewalling." It occurs when apparently correct answers are given to questions whose implied assumptions are false.

What is more desirable is a system that in a sense cooperates with the programmer. For example, for the previous question (if the assumption is false), the computer would respond by saying, "No, Gleitman has not written any books." It may even go further and say, "... but she has written two articles, A and B. Are you interested in them?" Rather than simply answering "No," and encouraging the questioner in this incorrect assumption, the computer indicates which piece of information is wrong and perhaps describes what is wrong with it.

Bonnie Webber's work in natural language communication is concerned with what sociolinguists call "repairs" or "side sequences." Like Joshi, Webber is interested in the general problem of mismatched information. She is studying what occurs when people correct or clarify themselves "online" (mid-sentence). For example, the questioner may realize

FAS reports

that his or her question is misleading, but rather than begin again, he might add a phrase like "if any" or "that is": "Which millionaire students are studying at Penn, ... if any?"; or "What is your mother's address—that is, at work?" The speaker may also remedy confusion by adding examples later in the conversation if he does not feel his explanation was sufficiently clear. Possibly the speaker will re-order his information if later the data is found to be faulty. In all of these cases the questioner has limited or clarified his question to avoid receiving a correct, but inappropriate, answer.

Side sequences, as these mid-conversation repairs are called, present unusual problems for a natural language system between programmer and computer. In human discourse they are accompanied by a change in tone of voice, a facial expression or hand gesture. A computer has a much narrower information band than does a human. It has only what formal, written, or spoken data has been programmed into it. To make language between man and computer truly "natural", a computer must have access to the physical data humans acquire through sight. Norman Badler and Ruzena Bajcsy, computer scientists in SEAS, are both doing research on computer vision. Badler is investigating methods to convert descriptions of human movement onto a human body model, and Bajcsy is constructing computer systems that can interpret two- and three-dimensional images or pictures.

The interaction between researchers in computer vision and human vision is crucial. For computer scientists, a better understanding of how people supplement oral and written information with visual data would lead to more flexible computer vision and communication systems. Given the work of Joshi, Webber, Badler and Bajcsy, computers may soon have the ability to interpret any and all types of human data—written, oral, visual or tactile. A machine could then become a very effective model for psychologists, linguists and philosophers to use to test their own theories about human cognition!

Although the computer's ability to handle large amounts of data seems limitless, there remains the continuing problem of flexibility in the accompanying system of communication.
—A. Joshi

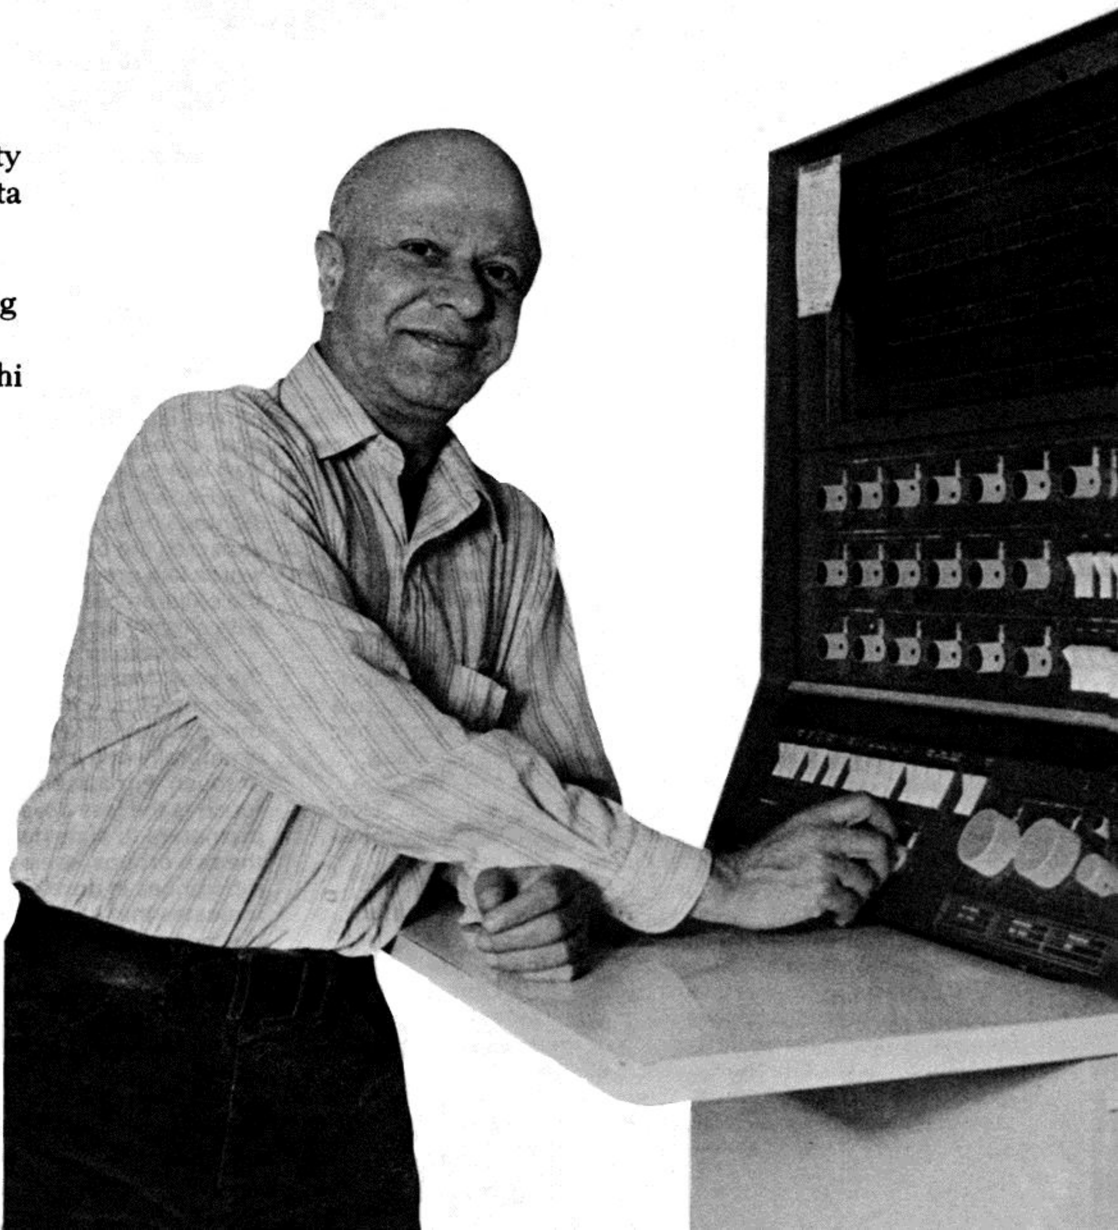

Learning Despite Deprivation

...despite seemingly acute sensory deprivation, a blind child does in fact acquire a normal language vocabulary and structure.

—Lila Gleitman

How do alterations and differences in mental capacity affect learning potential? How much impact does one's natural environment have on learning ability? By studying groups of people who have different endowment levels and groups who have been exposed to different environmental conditions, psychologists like Professor Lila Gleitman hope to learn more about the properties of the human mind. As Gleitman explains, "If you take away various parts of the total environment and look at which components of language survive, the basic ingredients in the language learning process begin to emerge."

Towards this end, Gleitman has been involved in two major studies of children who have been deprived of some aspect of the natural environment. One study focuses on deaf children, the other study on blind children.

The studies of deaf children were carried out primarily by Drs. Heidi Feldman and Susan Goldin-Meadow while they were still graduate students in psychology at the University of Pennsylvania. Working with Rochel Gelman (Psychology) and Gleitman, Feldman and Goldin-Meadow selected children who were from families of hearing parents and had lived in normal physical surroundings. They were studied before a formal sign language was taught to them and before they had acquired any training in oral language such as lip reading. The investigators found that certain properties of language can and do develop without speech. Gleitman explains, "Deaf children develop a well-formed, useful means of communication through hand gestures that has many of the same characteristics as spoken language."

The question that next arises is whether such a gesture language actually functions in the same manner as a natural (spoken and written) language, with structure, consistency and organization. On one level they discovered that it did. The deaf children they studied were living within their own families, isolated from each other. Yet, there was a certain consistency apparent; the language each developed was not very different from the others. One obvious example is the expression of hunger by pointing to one's mouth.

The study of blind children is a project recently begun by Barbara Landau (Psychology) and Gleitman. They hope to learn how a blind child makes the connection between a physical object and the word used to represent that object. For example, how does a child relate a round, throwing toy with the word ball? A seeing child has the entire visual environment and the simultaneous visual/oral connection with which to establish such language reference. "It is remarkable," Gleitman says, "that despite seemingly acute sensory deprivation, a blind child does in fact acquire a normal language vocabulary and structure." The child has been able to establish some alternative means of discovering how words refer (such as feeling).

Equally remarkable is a blind child's ability to use the sighted vocabulary of a normal child. Gleitman and Landau expected the blind children to have serious difficulties with words like green, red, look and see. Instead, they found that the children both use and understand these words, although with slightly altered application.

The verb to see does not only mean to literally look with the eyes. More generally, Gleitman explains, "it means to apprehend by some sensory transaction." With a blind child this is a tactile transaction, not a visual one. The visual sense is naturally at a lower level of development than is the sense of touch. For a seeing person, on the other hand, the sense of touch is the less developed sense. Although it is still premature to state definitively, it appears that the language of the two groups, seeing and non-seeing, reflects these differences in development. Gleitman's and Landau's work indicates that a blind child possibly makes the subtle distinction between the words touch and tactile in the same way seeing people distinguish the verbs to see (e.g. "I never saw him.") from the more emphatic, to set eyes upon (e.g. "I never set eyes on him!")

Commenting on her studies, Gleitman notes, "The deaf and blind children have been able to compensate for rather serious natural deprivations. They are amazing proof of the resiliency and flexibility of the human cognitive system."

A Model of Language Acquisition

Some theoretical issues in psychology and computer science merge into areas of interest to philosophers. Behind human memory and learning and behind computer programming is a system of progressively complex assumptions and inferences linked together by logic. Scott Weinstein, Assistant Professor of Philosophy, is using the techniques of mathematical logic to analyze this system that allows us to acquire language, make assumptions and draw conclusions. His final project is to be a formal model that abstracts the mental processes involved with language acquisition.

A child builds a grammar that organizes the data he begins receiving from the day he is born. As Weinstein says, "The child 'learns' the language when it converges on this grammar." The question is *how* does the child develop this grammatical framework and what limits might there be to the number or complexity of languages a child can learn?

Weinstein has set up a model that describes the properties of an abstract learner, "D". The model explains how the learner might acquire an appropriate grammar for any given set of data. It then describes how increasingly difficult grammars are developed that organize more complex data. In the final stages, the learner (D) is said to learn a language (L_I) when he formulates a grammar (G_n) so comprehensive that no additional data can threaten it.

The model indicates limits to a learner's capacity. If, for example, the learner hears sentence A, applies a simple grammar, and continues to apply the same, basic grammar to every sentence it hears, then it is likely that the learner has not acquired the language. He has simply adopted a string of words, void of any unifying structure. Other learners may formulate an appropriate grammar for one language, but given data from any other language, they will continue to apply the grammar that fit the first language. This would indicate that one language is the limit to this learner's ability.

Weinstein hopes to refine his mathematical model of acquisition by making use of the work of linguists and psychologists on human language acquisition. Their studies suggest what elements must be incorporated into his model if it is to represent the processes involved in human language acquisition. They also indicate the degree of importance of some of these elements, such as the visual environment, in the process of learning.

Each additional piece of information further refines Weinsten's model. As it becomes more and more precise, it in turn could be used by psychologists and linguists to learn more about the steps involved in the language acquisition of different groups of people. It may, for example, be used someday in the study of deaf children to better understand what steps in the learning process are lacking because of their physical deficiency.

Insofar as Weinstein's model deals with an abstract learning device, it may be generalized to a number of different populations. The model is questioning the capacity for language acquisition and is thus dealing with some fundamental issues about the range and boundaries of all human learning.

Scott Weinstein.

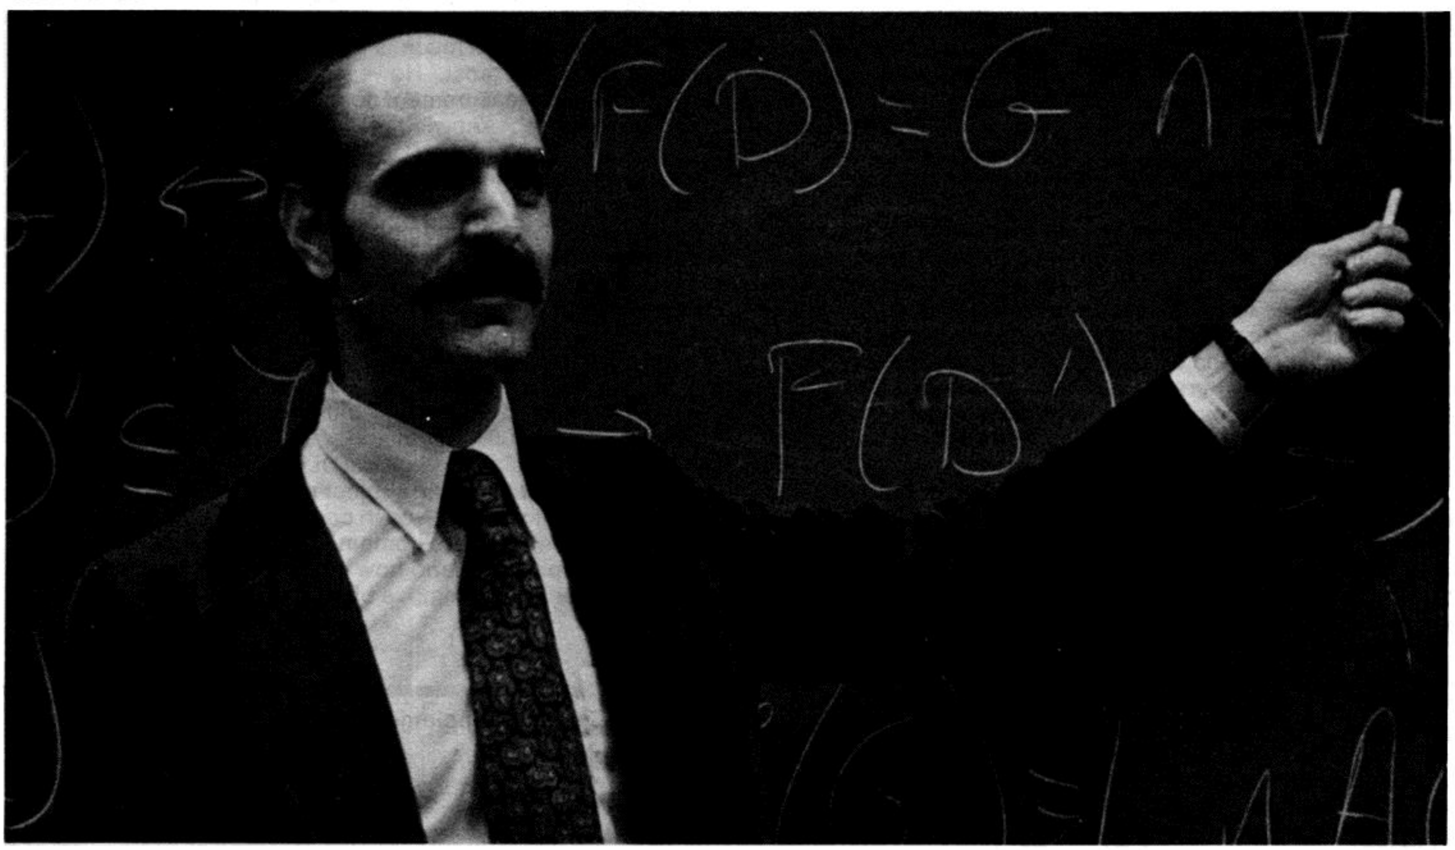

The Structure of Language

"The thing to bear in mind, gentlemen, is not just that Daisy has mastered a rudimentary sign language but that she can link these signs together to express meaningful abstract concepts."

Adding another dimension to the study of language are linguists who, like computer scientists and philosophers, strive to discover and describe the structure of language. Of particular concern to linguists such as Ellen Prince, is "what distinguishes a random list of sentences from a coherent text."

Ellen Prince is now working on a proposed project that will study how word organization, or "packaging," affects meaning. The project involves two other members of the cognitive science group, Lila Gleitman and Bonnie Webber. The group is looking at the ABC captioned, evening news, telecast for the deaf on public television. Through oral and written transcripts, they have found that the captions do not always correlate exactly with the oral text. Prince and her colleagues will be studying the differences and how they affect comprehension. The television program is directed towards an audience of deaf adults of average reading ability, as well as those with as low as a fourth grade reading level, who became deaf before learning English.

Understandably, the text has been changed for the sake of simplification. However, the network's simple grammatical changes have altered meaning. For example, when a spoken passive sentence is made active when written, the emphasis, and often the entire context, of the sentence is changed. In the phrase, 'It has been discovered that...', the discovery is being highlighted, whereas 'John Doe discovered that...', emphasizes the discoverer. Similarly, replacing a pronoun by a full noun phrase may result in a shift of importance from activity to actor.

Have in fact the network's changes made comprehension easier? "My own role in this study," Prince says, "among other things, will be to design comprehension tests. The objective is to see if what a deaf person comprehends from the written text is the same as what a hearing person comprehends from the oral."

Prince has been working on the issue of packaging in another, quite different study. In this project she is studying how people convey what they consider to be "shared—or presupposed—knowledge." She offers the example, "If I say, 'Gregorian told me that...', I would be assuming you know who Gregorian is. On the other hand, if I say, 'The Provost of the University told me that...', I would not be basing my statement on such assumed shared knowledge."

In this project, Prince is studying the discourse of physicians. She is analyzing tape recordings of their morning conferences to learn what each assumes to be common knowledge. However, her concern ultimately is not the shared knowledge related to technical or medical facts that all are expected to have learned in medical school. Rather, Prince is studying the assumed shared knowledge that relates to the ethical and moral issues that arise so often in the medical decision-making process.

The first step is to determine exactly what is said. "It is common practice for people, especially in a conference situation," Prince explains, "to be vague and indefinite. Perhaps because so much is at stake, doctors seem to be vague more than most." However, behind doctors' evasive symptom and measurement descriptions lies a base of shared (or presupposed) knowledge. Doctors use the phrases "sort of (tender)" or "(blue)ish" in describing symptoms. In citing measurements, they might say, "His blood pressure was around..."

Another type of hedging Prince found deals with unusual relationships between the speaker and what he says. When a person uses the phrases, "I think...," "As far as I can tell..." or "According to...", he is expressing doubt and a reluctance to take responsibility for the implications of his remarks. Again, as important as what is said, is what the speaker has implicated.

Prince's study is dealing with the unsaid side of language—presupposition and implicature—as much as the apparent goal of a common set of ethical beliefs among physicians. Although the subject matter is quite different from that of others involved in cognitive science research, the issues, Prince feels, are compatible. "When A says X, what does B comprehend, and how did B arrive at the understanding, whether B is a chimp, a blind child, a machine, or a physician?"

FAS Reports: A Supplement to Almanac

Editor .... Jan Brodie
Design.................................... E.R. Landesberg