

Cognition and Ecological Psychology

Held at Trinity College, January 23, 1982

Funding for the conference was provided by the Norman Spear Grant of the University Research Institute of Connecticut, with further assistance from Trinity College, The International Society of Ecological Psychology and Adelphi University.

James J. Gibson developed his ecological approach from the point of view of perception, primarily visual perception. Nevertheless he realized that the essence of his ideas extended panoramically in every direction of psychology, as far as the eye could see. He did not often venture far from perception himself, though. Eleanor Gibson, through her work on development, learning, reading, and attention has touched on more cognitive areas, but the Gibsons' ecological psychology has not penetrated far into cognitive psychology. The reason is not hard to understand. The ecological approach emphasizes the roles of real environments and real bodies of active perceivers in controlling perceiving and acting. The more control an analysis assigns to the body and to the environment, the less need for "cognitive" explanatory terms. On the other hand, students of memory, concepts, representation, imagery, speech, and language usually take their subject matter to be mental, and internal to animals. If they are sympathetic to Gibson's ecological work, they assume it applies only to perceiving (a view Gibson himself sometimes voiced) or that he was not really denying "mental" processes as much as emphasizing others. If critics are not sympathetic, they take him to have denied the obvious, and judge him foolhardy. Nevertheless, there has been a persistent refusal on the part of some people from both sides to accept separate but equal status for cognitive and ecological psychology.

Some ecologically oriented people believe they can approach cognitive topics, and some cognitive scientists recognize that environmental constraints are important. Therefore this conference was convened to find out what issues would emerge from a diverse group explicitly discussing ecological psychology and cognition.

Leading representatives of research in perceiving, acting and a number of areas in cognitive psychology were invited to participate in panel discussions. People in art, computer science, theoretical biology, and linguistics were also included. Participants were either actively developing ecological approaches to their subject matter, were trying to discover more about how ecological work fit in with their own, or were merely curious to learn more. None were antagonistic to ecological work. The discussions were meant to be cooperative and constructive. All participants had at least a rudimentary knowledge of the ideas of James Gibson.

Each panel was given 1-1/2 hours. Participants were invited as resource people who could address topics in their areas of expertise of the topics chosen. Two people in each group (Eleanor Gibson and John Bransford in the first, Robert Shaw and Lynn Cooper in the second, and Michael Studdert-Kennedy and Scott Kelso in the third) were given 10 minutes to describe issues of interest to them as a way of establishing topics for the remainder of the panel. The rest of the time was for general group discussion. It was expected that the degree of participation by any one person would depend on the particular turns taken by the conversations.

Group - Memory and Concept Formation

<u>Participant</u>	<u>Affiliation</u>	<u>Specialty Area</u>
Eleanor J. Gibson	Cornell University	Perceptual development, reading
John Bransford	Vanderbilt Univ.	Cognition and education
Jeffery Franks	Vanderbilt Univ.	Cognition
Peter Machamer	Univ. of Pittsburgh	History and philosophy of science
Michael Turvey	Univ. of Connecticut Haskins Laboratories	Organization of action and perception

Memory and Concept Formation

Eleanor Gibson opened the discussion by describing how one might approach memory, concept formation, and other cognitive topics by generalizing from nine characteristics of perceiving, with a hopeful eye to promoting experimentation. "I have a feeling that it's awfully easy, when you're trying to apply a new approach to a new field, to do an awful lot of talking and not get that talk into the sort of shape that you can take to the laboratory," she warned.

She maintained that, in general, memory should be treated functionally. It should not be reified or treated as a faculty with causal or explanatory efficacy. It should not be treated as if composed of elements, or traces, as copies of previous experiences.

Turning to perceiving, she described its function as extracting information about events in the world, the layout of the world, and what these afford for action. Memory, then, could be a function for acquiring knowledge about these things and retrieving it as it helps for behavior, both performatory and exploratory. The characteristics she listed were:

(1) Continuity. Perceiving does not start and stop. It just keeps going on. Since she had in mind James' description of the stream of thought anyway, pointing out that remembering and thinking can be continuous is not new, but it certainly bears repeating.

(2) Perceiving is relational and the information for it is relational. She took this to be a point against elementarism, and equally from James on cognition.

(3) In perceiving, invariants that underlie change are extracted. Can this be applied? Seemingly, concepts could be regarded as the "cognitive fruition of the detection of invariants."

(4) We are motivated to seek order and invariance in perceiving. Are thinking and remembering similarly motivated?

(5) Perceptual learning and development are marked by increasing economy. Remembering, too, seems marked by increasing economy.

(6) What we perceive are the affordances of events and layout. Do we selectively remember affordances?

(7) Perceiving is not generally modality specific. Is remembering modality specific? She said she would like to see research started to turn around the current trend of treating memory as divided into auditory, visual, tactual and so forth. Late in the subsequent period, Peter Machamer used the example of learning to taste wines as an example of increasing modal specificity. Texture, taste, and smell are just one hodge-podge to a novice. These qualities are differentiated in perceptual learning, and therefore more modally specific. Gibson liked that, with the qualifications that there be some utility involved and that it be recognized that this is very different from positing entities such as visual memories.

Michael Studdert-Kennedy and John Rickards each commented on Gibson's point during the third session, Speech and Language. Studdert-Kennedy reminded the group that speech is tightly bound up with audition, despite lip reading. He too cautioned against overemphasizing modality independence. Rickards, using the wine tasting example, said that differentiation is only half the task. It must be followed by reintegration so that the differentiation may be usefully applied to wines themselves rather than their parts alone.

(8) Perception is not inference or enrichment. However, she expressed doubt about memory. Certainly one can make inferences based on what one remembers, but is this necessary in a theory of how one remembers?

(9) In the laboratory, she strives for ecological validity of experiments. Should not one do the same for memory? Some memory experiments, for example memory for stories, do seem to be moving in this direction.

John Bransford also chose to look at memory and cognition in terms of functional relations between people and their environments, but he concentrated on acting, rather than perceiving, as his focus. In the second edition of The Sciences of the Artificial, Herbert Simon had suggested that long term memory acted as a second environment, as an extension of the environment in which thought processes take place. Bransford thought it would be interesting to ask what would happen if one pursued this idea of memory as an environment, that one could still be immersed in. He noted three parallels between the real environment and this memory environment.

(1) Both are to a large extent, constructed. It seems straightforward to think of the memory environment as constructed, but it must also be noted that the "real" environment is filled with artifacts of human design, meant to allow various human activities to be performed. Thus it too is largely constructed.

(2) Both may or may not be understood by people. That is people could see or even image a baseball field, but understand little about what they see.

(3) Both permit learning and discovery and thus can be continually modified. For example, the significance of a warning track around a baseball field may be discovered either by actually playing baseball or by imagining acting in that environment and noting what happens if you chase a long fly ball. If one merely tells a novice what a warning track is, the person will usually claim to know what it is as a result. But having the person think through the activities involved allows a person to reason effectively about more possibilities. Later in the session Bransford proposed that people try experiments to get people to represent something in memory, then see if they could discover something new.

Perhaps, then, this ability to imagine activities that can be carried out is one function of memory as an environment in which we operate. In current literature, determining what you know and whether or not it is adequate is called "metacognitive assessment."

Machamer raised a number of questions for consideration. What is the relation between the two environments? Where does the novice have to begin? What determines what gets carried over from the real world? What about false memories? Would Gibson and Bransford be happier with procedural than network models? Gibson said "yes" if that meant something dynamic. Bransford said "yes" if it was understood that "procedure" meant an activity in an environment, like a cook's acting to carry out a recipe in a properly equipped kitchen. Then Machamer asked what this would look like in a way that would not quickly get back to networks.

Turvey argued for examining some very different examples that pointed to what he took to be more fundamental questions underlying memory. He agreed with Gibson that memory should not be reified and asked "What explanatory concepts are needed to explain memory? But he went further in arguing that "memory" and "concepts" are cover terms used in explanations as substitutes for, rather than components of, understanding. Turvey believes that examples of more fundamental problems are behavioral reproducibility and stability. Memory is invoked in a host of disparate situations, often just as an assertion that some knowledge or behavioral ability persists. Turvey's current strategy is to find instances of stability where the possibility of finding natural law explanations is promising. Such examples would then provide a standard of adequacy against which to measure accounts of other stabilities in psychology. The task of repetitively swinging a hammer at a comfortable, natural frequency is one such task. Do people remember the frequency? Or do they return to the same frequency on different occasions because the same laws and conditions hold? What is the information for the repeatable performance? The shape of a flock of flying birds provides another case. The V-shape is apparently dictated by a minimum expenditure

of energy principle. If one has such knowledge of the dynamics of the bird situation, then to say anything about the birds' memories of what shape to fly in, would be gratuitous. Discovering principles like these is surely hard work and will take a very long time. Turvey concluded that patience will therefore be an important virtue. He associated the use of "memory" and "concepts" in explanations with impatience.

Discussion

Agreeing somewhat with Turvey, Bransford noted that to obtain the best reproducibility of much behavior, one should not put experimental subjects in a remembering mode, but should design the demand characteristics of the situation to bring out the stable behavior, to put a good deal of control in the environment.

Machamer mentioned that the panelists had now touched on a wide and wild variety of examples, which surely did not all warrant being treated as the same topic memory. Turvey agreed, arguing that this is precisely the point. If phenomena are carefully analyzed on their own terms, new taxonomies will surely be necessary.

Jeff Franks, recalling his 14 years of research collaboration with John Bransford, much of it on memory experiments, asserted that it was the design of the experiment, the method, that was most interesting to them. The important questions were: How does one design an experiment that works? How does an experimenter set up new situations for a subject that the subject can use? In keeping with this attitude, Franks suggested that the methods section was the heart of a paper, and perhaps the only section that warranted publishing. An extreme of the attitude he was exploring would be to go "whole-hog" functional, to think about memory as a kind of experimental design promoting certain activities such that a modification of knowing occurs. Information content per se, seems of little consequence to him. A contentless theory of memory would address questions about what people do. In order to remember something about the present in the future, an appropriate activity that modifies the person must be performed now. The problem of an experimenter's procedure is "How can I get people to perform the appropriate act?" The answer to that question is what would be in a methods section.

Turvey asked for comments from cognitive scientists in the audience. Although the conversation had ranged widely, no one had defended the approaches and models of cognitive science. Later in the discussion, Pat Hayes obliged, explaining "I am not the devil's advocate, I'm the devil... I came here to find out what 'ecological' means." To set the stage, Hayes described the cognitive science strategy in order to encourage a parallel reply. He described flow charts, a target of criticism several times, as not explanations, but as a stage on a path to an ultimate goal. To ecological psychologists, he asked "What is your ultimate goal? What will count as success?"

"What counts as success for us, he said, is a complete account in computational terms... of how the mind works, in information processing terms." It is not known exactly what these terms are going to be, but they should be related to sensing and behavior. Then an account of how the computations are implemented in neuronal hardware is needed. The faith is that the problem can be factored this way, in the way that computer action is factored.

Gibson replied that what she wanted to know was why people do what they do, how they learn what they learn. As a psychologist, she said she certainly had no

interest in the neuronal hardware. Functionally speaking, she said she was a behaviorist.

Machamer proposed that such functional programs intended to link terms referring to the environment and to behavior, with nothing in between as an "internal" term.

Turvey's version of an ecological approach was that it is a program directed toward considerably more careful analyses of the phenomena of interest so as to remove unnecessary mental concepts. One of the most pertinent shortcomings of cognitive science to him is its inability, apparently in principle, to explain the origins of their phenomena. JJ Gibson's attack on the third dimension is a good example of a reanalysis of long standing ideas. The ecological approach cannot identify its ultimate types of explanation with precision, but it points in a direction. It says, "continue to look at the relationship of organism and environment with a critical eye, being open to the increasing number of new variables, new observables, that will creep in. Don't buy into the old observables, and what you will begin to see is that the scientist need not have to attribute to the brain as many special machineries as science now seems to have to... I take the ecological approach to require a lot of patience. There is a tremendous amount of groundwork to be done."

When people were asked for their final remarks of the session, Bransford noted that it was the problem of learning, more than memory, that seemed most important. Gibson concurred, describing the past 15 years as a time when learning, as an enterprise, was dropped in favor of memory. "It's time we got back to learning."

Turvey closed, saying he would like to see JJ Gibson's way of treating time, as an intricate coupling of persistence and change, in nested episodes, taken more seriously in future analyses of memory. What are the constraints of large episodes on the smaller ones nested within?

Imagery and Representation

<u>Participant</u>	<u>Affiliation</u>	<u>Specialty Area</u>
Robert Shaw	Univ. of Connecticut	Perception of growing faces; models of growth
Lynn Cooper	Univ. of Pittsburgh	Cognitive skill learning
Margaret Hagen	Boston University	Perception of pictures, psychology of art
Pat Hayes	Univ. of Rochester	Philosophy, computer science, artificial intelligence
Mari Riess Jones	Ohio State Univ.	Perception of music, rhythm
Nathan Knobler	Philadelphia College of Art	Painter, sculptor, art theorist
Joe Lappin	Vanderbilt Univ.	Perception of motion, mathematical models

Representation and Imagery

By one count, the seven panelists charged with discussing representation and imagery produced eight different issues that concerned them. On the face of it there was little convergence. They were primarily "laying their cards on the table." However, there are some overarching categories that can be used to organize the points without too much distortion. For example, everyone concentrated on representation. No one selected imagery.

There were two main questions that the panelists were concerned with.

- (1) What is a representation?
- (2) How (and by whom) do we want to say that representations are used?

What is a Representation? The statements by Shaw and Knobler touched on the first question. Shaw defined a representation as an object that stands for another object, as a surrogate. Interestingly enough, he took this to be functionally established and therefore implicated the second question. That is, it is the use of one object by some user as a representation that makes the representing object a representation. This is an act of representing. His primary point was that the deeper relation underlying both representation and perception was specification. To be a representation the surrogate must specify that which it represents to someone. It may do so symbolically; that is, by virtue of arbitrary rules known to the user (including associations) or figurally; that is, by virtue of some lawful relation (based on causality) or function (biological), but specify it must. Seen this way, representation depends on successful specification. Shaw realized after the conference that he could have

been more clear had he explained that what he meant by specialization is best modeled by ostensive specification, successful acts of pointing. Construed broadly, pointing is what he meant by direct specification and representation becomes a form of indirect specification. For instance, some aspect of a scene may be "pointed out" to someone by inspecting a picture or a symbol rather than the real scene. The question Shaw hoped to discuss with other participants was whether or not (or how) one might develop theories of direct specification. If so, then when would representations be required in an explanatory role in theories at perception, memory, concept formation, problem solving, etc? He finished by warning against a fallacy he called the First Order Isorphism Fallacy (F.O.I.F.). This is to treat the surrogate function of an object as an explanation for the representation relation rather than as a phenomenon to be explained by clarifying the underlying relation of specification. The ambiguity of whether or not a particular concept was being given descriptive or explanatory status in a discussion was pointed out several times during the conference.

Nathan Knobler is an artist and theorist of aesthetics who has taught drawing for nearly 30 years. He felt that it was important to remind psychologists what artists must do to make their representations. Since psychologists often seem to have in mind analogies to art, frequently pictures, the artist's products could be taken as the most literal sense of representation. Thus he too was touching on Question (1). In discussing what an artistic representation is, he meant to be underscoring how many aspects of this, if you will, real representation are left behind when the analogies are made and, by implication, to question the success of the analogy.

Knobler's central point was that every representation is incomplete. Consequently, no representation is accurate; indeed there is no standard of adequacy for a representation. Further, making a representation involves a complete series of decisions, not all of which are made in advance. Some of the decisions are about how the artist will move his or her body, others are about color, and materials, quite apart from resulting spatial structure. An art work develops. Some may be heavily constrained by prior intent and level of skill, but not to the last detail. Thus it is not a mere transfer of an internal or external "image" onto the canvas. The artist is not doing what a computer graphics system does when a coded system of points is converted to a display.

Using Representation

Cooper spoke more to the usefulness of representations in her research. She has been studying people learning to solve physics problems and has been struck by how much reorganization of the problem domain she has observed. Learning, changing performance, struck her as a more attractive problem than the study of any particular performance. She echoed John Bransford's focus on this. She has found it useful to say that representational entities used by the problem solvers were what was changing. Whereas Turvey had argued that concepts such as representation and memory were to be resisted because they were default concepts that could not play explanatory roles, Cooper advocated impatience (in the sense of contrast to Turvey's counsel of patience). For her, a psychological theory surely had to attribute something to people to account for learning, changing performance in the face of unchanging problem demands, and Turvey's position seemed to have no place in it for discussing changes in the person. Representations provide the "something" to attribute, although the best form of representations is certainly a question. Thus the main question she presented to the panel was "what kinds of representations shall we use"? Boxes and flow diagrams do not get us very far. Representations surely have a dynamic

quality to them. How should theorists think of them? Thus she arrived at the need to evaluate different representations against one another rather than the need to evaluate the idea or the role of representation.

Mari Jones picked up the "empty organism" problem. She studies people's knowledge of musical structure and rhythm and knows quite well that the level of skill and attentional abilities of a person must be conceptualized. Although she does not attribute "representations" to people and does not use the term, much of her work has been devoted to finding an adequate scientific representation of aspects of the knower and the known (music, rhythm) which are commensurate. This she does by attributing rhythmic structures in transformations and invariants, to persons or animals and their world so that she can capture ideas of synchrony and meshing between them and their world. Time has served as her neutral ground for characterizing both the animal and environment sides of their interaction.

Hayes, from the standpoint of a computer scientist, suggested that people should not be misled by their words. Terms such as "mental image," "memory," "inference," "representation" are understood as technical terms in computer science, just as "group" is a technical term in mathematics. Some kind of logical manipulation or computation must be going on in the head and, thus, some representational medium is required. He intimated that Artificial Intelligence is too modest to commit many of the deep fallacies that arise from prematurely claiming that a theory is explanatory. To him, AI surely puts off many problems, such as a principled semantics, but this is prudence not failure. He expressed wonder that some psychologists doing ecological work seemed to find a conflict with AI. He could understand that more attention needed to be given to the content of representations, as opposed to their form, and that this seemed to be the thrust of ecologically oriented research, but that was no conflict. The logical issues related to the formal aspects of representations remain. Some of his own research deals with the nature of liquids at the level of what he calls naive physics. How might one capture the distinction between a puddle and a pool, for example? A system of representation that can address this distinction allows his programs to run better and thus is readily justifiable on practical grounds.

Discussion

Following on Shaw's remarks about FOIF's, Lappin raised a question about the possibility of theories of direct specification or representation. "How can one have a representation without a representer?" He explained that his fundamental question here was that of how coherence or coordinated activity can arise from a collective. In this form, the question resembled one Howard Pattee began pressing at least 15 years ago: "How does simplicity of biological function arise from physical complexity?" Theorists have been tempted to attribute the responsibility for coherent organization to devices that have been specialized for that purpose, e.g. feature detectors in perception. But these beg the question, as Lappin noted.

Hayes distinguished Lappin's question from Shaw's earlier definition of a representation in terms of a representer. Knobler expressed concern that too much be made of representers because an observer of a work of art may construe the work in many ways that were unintended and unimagined by the original artist. Knobler took this to show that a representation depends on the representer in a very limited way. Earlier, he added to Lappin's question by asserting that examples of coherence arising from collectives abound in art. Any image produced

by a contemporary artist, for instance, is a result of art history.

Lappin closed the conversation by listing a series of 4 examples that could be graded from more to less feasible in terms of constructing a theory of direct specification. At the most concrete end of the spectrum, and hence most feasible, was perceiving the constraints of media and materials used by an artist as discussed by Knobler. Second, was perceiving events and encounters by virtue of flow field information. Third, and much more difficult to envision a thoroughgoing theory of, was music. Fourth, and most abstract, was Cooper's example of solving physics problems.

Addressing Shaw's interest in theories of similarity and Knobler's claim that a representation could be separated from the process of its generation, Hagan asked if it were possible to make 2 representations that are the same in different ways? Knobler answered that people have been misled by theories of organization in art that try to explain too much. To him, an artist produces a "carrot" which sustains an involvement on the part of a viewer that allows the viewer to create his or her own organization. A work is usually so complex that there are large numbers of possible relations that could be selected by an observer. Careful observers impose organization in a series of hierarchical refinements. Knobler is interested in why certain relations are selected and not others. Perceivers may perceive the same things in different works if they are looking for the same things, e.g. curvilinear or triangular gross organization. The grossest levels may be forced by the artist, but not the more refined ones. Hagan felt this approach still left representations largely equivalent and did not allow for an ordering of differences.

Using the example of perceiving growing faces, Shaw listed different kinds of similarity one might use: geometric, kinematic, and dynamic. In the first there are corresponding structures, in the second there are corresponding motions, and in the third corresponding forces. Then, an issue would be: what units are derived and what are fundamental? Elementarism is when microscopic units are taken to be fundamental and macroscopic units derived. Correspondences would be explained at a first order isomorphism, point for point. The opposite end, dynamic similarity, would use field comparisons, not points. Points would become derived, not fundamental units. If radar guns for highway patrolmen had been developed before the yardstick, velocity could have been a primary measure and distance derived. Whether or not a theory of perceiving looks direct will depend partly on what are chosen as fundamental units.

Hayes added clarifying remarks about his use of the terms "representation" and "representor." "Representor" means "that which represents" to him. How is the person related? One can't fall into Hume's problem. One does not need a perceiver or user for representation. What might be more accurate would be to say that we are our representations.

Cooper repeated her question about what the implications of learning to solve physics problems might be for a theory of representation, knowledge structure. Lappin suggested that such a notion of representation would not be explanatory in Turvey's sense, but is a construction of the problem itself. Perhaps it is an explanation, in a way, of the knowledge someone has.

As the conversation turned more to expertise, Knobler characterized 3 levels of expertise: Where someone (1) learns to copy, (2) learns to do it well, and (3) comes to do what no one has done before. This is where the unique individuality of the creator enters. Cooper agreed emphatically, saying that cognitive

reorganization in learning was what she is trying to get at.

Shaw brought out the role of negative facts (suggested by Bransford) in highlighting the problem of intentionality in perception. One case is Knobler's example from Picasso who said that he knows when he has done something right when it doesn't remind him of what he has done before. Another is noticing that an expected assassin is not in the room. An ecological approach, as well as any other, must address these cases too.

Language and Speech

Michael Studdert-Kennedy	Queens College Haskins Laboratories	Speech perception and production
Scott Kelso	Haskins Laboratories	Organization of action
Randy Diehl	Univ. of Texas	Speech Perception
Howard Pattee	SUNY Binghamton	Physics, Physical foundations of biology, symbols
John Rickards	Univ. of Connecticut	Educational psychology, cognition, memory
Mary Smith	Univ. of Connecticut	Linguistics

Speech and Language

Michael Studdert-Kennedy attributed his interest in ecologically stimulated ideas to the promise he saw in them for attaching the problems of segmentation and invariance (or lack of it) in speech. Historically, it has been difficult to assign invariant acoustic underpinnings to all repetitions of phonetic segments. Therefore, special "devices" in the head were posited to "compute" segments from sound structure.

Studdert-Kennedy grew dissatisfied with the explanatory value of this way of thinking and was encouraged by Michael Turvey to imagine that there still might be ways to derive speech from physics. He cited work by Lindblom which derived phonemes from principles of maximum perceptual contrast and articulatory ease (along with assumptions about how vowels might be embedded) as promising. Try as he might, however, he still cannot see how to eliminate a level of autonomous linguistic organization. For example, in spoonerisms, nouns are substituted for nouns and verbs for verbs. Later he mentioned the error in pluralization when "cows track" is substituted for "cow tracks". The proper phonetic [z] is used to pluralize cow, not the [s] from [tracks]. Even though [cows] (rhymes with "louse") is physically simple, it is not said. An appeal to linguistic rule seems to be the only way to handle such data.

Turning to perceptual learning, as raised by Eleanor Gibson, Studdert-Kennedy mentioned speech data that called for distinguishing several levels. Pat Kuhl has reported a number of phonemic distinctions that infants make across speaker and context (showing that the babies are not bothered by the apparent lack of invariants).

Yet some of the same distinctions, sa - sha, fa - tha, give 3-year olds trouble when they try to produce them for purposive communication. Following Winifred Strange's suggestion that production and perception be treated as a unit, Studdert-Kennedy wondered if the 3-year olds showed the same decrement in perceptual development as speech came under intentional control. Scott Kelso continued the theme introduced by Turvey by advocating a strategy of trying to understand a system's dynamics a first - before invoking concepts like representations or programs - in order to determine how much organization could emerge "for free." This is a strategy he applies to speech as well as other acts and the one Studdert-Kennedy commented on. He used examples of mass-spring systems and gait to illustrate properties of stability and equifinality that emerge from underlying organization without being imposed by an agent, programmer, designer, etc. That is, much control is clearly not computational. There is not necessarily an a priori, stored reference level required to account for equilibration.

In gait, there are a number of stable organizations that emerge when something like velocity is varied. In a quadruped, shifts from a walk to a trot and from a trot to a gallop each seem to occur at certain critical values of velocity. Thus locomotion shares with language the properties of having categorical segmentation, yet there were no initial units "glued" together. Rather, the units, the gaits, emerged as nonlinearities on the overall system organization.

To apply his strategy to speech, Kelso advocated looking at larger stretches of time and wider ranges of conditions than is usually done to discover hitherto unknown variants. Thus, analogous to the locomotion example, one might have a person speed up or slow down speech or produce it under the bite block conditions as Betty Tuller has done.

Mary Smith, speaking from the point of view of linguistics, asked "What can ecological psychology do for me?" taking her goal as understanding the process of successful communication. Accepting the competence-performance distinction, she assumed a theorist might choose to make the grammar as small as possible, assigning as much organization as possible to the nonlinguistic, performance aspects of a speech system. How much would a child have to differentiate a speech signal in order to learn to speak? What are the right units? What are the relevant dimensions of the physical environment to attend to? How are these put together in a sensible expression (i.e., what relations should be established - ratios, ratios of ratios, etc.)? There is a juggling act involved in going back and forth between acoustics and linguistic units. Finally, how are expressions evaluated? She suggested that relative values are more plausibly used than absolute values. Taken together, these points amounted to a call for vigorous forays into ecological acoustics.

Extending the range of data that challenge ecological acoustics and theories of tuning, she mentioned second language learning, aphasias, and historical language change.

Finally, turning to what might be taken to be data on the role of intentionality, she described a study of children with cerebral palsy who could produce no language-like sounds even though they seemed motivated and bright. However, given access to a source of sound that they could control, they began to teach themselves to imitate the sounds. A therapist pushing their faces around to various speaking positions had no such effect. How does intention get into anyone's picture?

One of the primary questions for any theory of perception is: "What are the proper objects of perception?" Many of the controversies concerning representations rage over this point. James Gibson, for example, took it that sensation based theories of

perception always required that objects other than those of the real world be the direct objects of perception. Both retinal images and mental representations of any sort counted as surrogate direct objects. Gibson maintained, on the contrary, that meaningful properties of an environment could be what animals experience - without intervening or mediating objects. Randy Diehl was the first person at this conference to raise the point. There are nonarbitrary dimensions of language and much of the research in speech is devoted to fleshing this out. However, implicit in some speech research is the idea that configurations or trajectories of a vocal tract are primary objects of perceiving. They seem to be what is specified acoustically. But this seems patently wrong. Why would one want to know anyone else's articulatory states? Meaningful speech or speech acts seem more likely candidates for the natural object of perceiving, or speaking for that matter. Since these involve arbitrary symbol systems somewhere along the line, it is difficult to see how meanings could be acoustically specified. Is there a way out?

Howard Pattee was not about to offer an easy way out. He argued that it is an important strategy in physics to focus on the worst case for the available theories. For languages, he thinks this worst case is the arbitrariness of symbol systems. Working at the molecular level (genetic codes) he has noted several characteristics that might well be universal over symbol systems. The hierarchical level problem is very serious. Theories tend to work best at one level and have difficulty conceptualizing changes of level. Looking at the genetic case, he said that the string of DNA that makes up genes is obviously physical, but its role is not described by physics. The resultant folded structure of protein is critical to its functioning but the gene does not describe the folding. Rather, the conditions for folding are established and the folding proceeds automatically according to physics. Since the folding is not described, Pattee regards it as ineffable. It is also inscrutable. One cannot look at the functioning enzyme, the protein, and go backward to derive the string, because the mapping or code is arbitrary. One does not seem to be able to go back and forth between linguistically functioning entities and physics period. Yet both are essential to a full description of living systems. "Is that bad enough?" Pattee asked in closing.

In later discussion, Kelso questioned Pattee's long standing position that linguistic-like codes and dynamics are complementary in Bohr's sense of complementarity. Kelso maintained that dynamics must come first. Codes must emerge. Pattee agreed only that linguistic activity should be kept to a minimum.

Harking back to the themes of transitions from novice to skilled practitioner, Rickards mentioned a distinction between thinking about something (e.g. wishes) and thinking in terms of it. The expert, to him, achieves the ability to think in @U terms of the object, and it is this transition that we need to understand. Mary Smith countered that the skilled-unskilled distinction did not apply to first language learning, though it might to second language learning. To Studdert-Kennedy, first language learning should not even be considered learning, but rather a kind of embryological development. Smith seemed to agree. In second language learning, people "can't get the prosody right, but they can get the segment inventory down in about two hours. The kid, on the other hand gets the prosody in about six months but can't get the segments for five years." These reverse trends argue against accepting strong parallels between first and second language learning. Studdert-Kennedy continued Smith's point by contrasting the radical discontinuity between prelinguistic infant noises and the babbling of one year olds advocated by Jakobson on one hand and the recent work of Rachel Stark on the other. Studdert-Kennedy argued that a system of notation different from Jakobson's, allowed Stark to see infant phonemes develop as emergents from processes such as eating, sucking, and crying.

Diehl steered the conversation back to the more general topic of representation by returning to Turvey's earlier example of analyzing conditions for a flock of birds to make a V formation. While accepting least effort principles and physical conditions, Diehl claimed that these principles do not compel the birds to fall into the V pattern. They must elect to take advantage of these possibilities - and that seems like a representation of an action plan to characterize knowledge of the physical conditions. Turvey agreed that the patterns of turbulence set up by the birds in the air were not enough to fully explain the behavior, but rather than settle for a "plan" or the like in a theory, urged that his strategy be applied recursively. Accordingly, this new problem of how the birds "choose" what to do would be looked at in light of physical constraints and optimality conditions. Hayes was puzzled by Turvey's resistance to theoretical terms that involved representation. Turvey repeated his view that representational concepts are default concepts which, at best, act explicitly as placeholders for explanations and, at worst, masquerade as explanations. He cited temperature regulation as an example of a phenomenon in which the usefulness of a representational concept (set point) diminished with increasing understanding of underlying physics and physiology. He asked what criteria a theorist should use to decide when to bring in representations, particularly when the system of interest controls and harnesses energy.

Hayes granted the wisdom of pursuing supporting physics where possible but warned that important generalizations could be lost if algorithmic components of theories were rejected out of hand. In the case of the gull flocking patterns, a simple rule about where to line up with a neighbor's wing could account for much of what is observed. Details of physics would lose that completely.

Turvey agreed that this was a point not to be missed and a legitimate place for "rules." Principles may be elegantly captured by such an algorithm, but this is a descriptive usage. A rule, used in this way should not be taken to be embodied. It should not be given command or causal status. Natural laws are not written in impositional vocabulary.

"Are you balking at the explicitness of rules?" Hayes wondered. "Yes," answered Turvey, "and also the causal role people give them." Studdert-Kennedy again sympathized with the desire for understanding that penetrated more deeply than a level of "rules." He characterized rules as placeholders, like the work "intention," but he also saw no substitutes. Language is full of too many recalcitrant cases to expect alternative nonrule-like accounts to be successful.

Discussion

Nine identifiable discussions emerged in the period open to full audience participation.

1. Direct Memory. Steve Wilcox opened by championing a direct theory of memory that would forego concepts of storage and representation. He argued that theorists rely too heavily on recall tasks for their models. Pattee and Hayes offered spirited resistance.
2. Pragmatic vs. Explanatory Approaches. Jim Todd noted a contrast in positions represented at the conference which he thought were somewhat incompatible despite their common aims to be ecological. Concerns both with practical goals (engineer-

ing) or natural environments often give a lower priority to theoretical rigor. Approaches such as Turvey's emphasize theoretical rigor but may appear to have limited practical value. Is either more "ecological" than the other? Shaw agreed that science could proceed on two tracks, the applied and the explanatory, but did not see these as contradictory.

3. Imagery and Perception. Turvey asked Cooper to describe her criteria for establishing parallels between imagery and perception. She said she used functional criteria, primarily a variable's having a common effect on perception and imagery tasks. Further, she stressed that she took an action oriented view toward imagery. That is, she preferred to think of it roughly as a "readiness for action" a la Festinger rather than in figural or pictorial terms.

4. Invariants. How do people know when they've found an invariant? In speech, Michael Studdert-Kennedy said, the conventional strategy is to present listeners with a systematic range of synthetic cases in order to determine the acoustic characteristics of some speech property. However, he took this to be a case of the machine acting as an arbiter and contrary to what he understood Shaw to claim in the morning. Shaw had said that to create or build something was not to explain it. Studdert-Kennedy seemed to take this as an attack on synthetic stimuli. Turvey, Robert Remez, and Eleanor Gibson all spoke in defense of using synthetic displays in experiments. Shaw said that the point is to show how the displays are used and is not really whether or not they are synthetic.

5. Memory and Inference. The longest exchange of the session was triggered by Michael Studdert-Kennedy's comment that Jeff Franks seemed to be proposing a Skinnerian operational definition of memory, that memory amounted to no more than the conditions for certain behaviors to occur. Franks felt that that was not quite right because he was not thinking of explaining memory. He was referring to situations where animals seem to know what to do as a result of having performed certain acts. He was emphasizing effective action as the object of his interest and in that respect was comfortable with an "old style" functionalism. Turvey said he wanted the ecological approach to be richer than that.

Shaw proposed to include Franks' concerns within an alternative scheme of classifying types of knowledge first offered by James Gibson. Here, one could think of retrospective, perspective, and prospective knowing and consider them as arranged on a gradient of constraint of degrees of freedom. Looked at retrospectively, the degrees freedom of world events are fully constrained. That is, there are no degrees of freedom. Perspectively, world events are highly, but not completely constrained; and perspective there would be much less constraint on degrees of freedom. Theories of constraint, then, would be what are called for, and could address phenomena from memory to inference. Hayes accepted this formulation with the qualification that one not call them 3 kinds of knowledge but knowledge about 3 (or better, a gradient of) different things.

Wilcox was surprised to hear Hayes argue for a unitary type of knowledge and mentioned that AI people usually like to attribute differences in phenomena to differences in process or mechanism rather than content. Hayes said that Roger Schank might attack a problem that way but that he (Hayes) chose to think of processing as one "massive enchanted loom" of inference making. Hayes chose this occasion to further clarify what he meant by inference, which was merely "mechanical" or formal symbol manipulation. Eleanor Gibson had associated deliberation with her sense of "inference" and Hayes stressed that this was not his meaning. He did not use "inference" to refer to a human act. It is not a natural kind term to him. Turvey stressed that it was precisely Hayes' meaning that he (Turvey) would

dispute. Turvey argued that if one chose different fundamental observables (e.g. David Lee's time-to-contact instead of distance and time), what looked like a series of stepwise calculations from one view could look like an amalgamated unit from another. Even more, Turvey said, he would avoid Hayes' kind of inference in a perceptual theory because it leaves no firm foundation for perceptual knowledge. Perception should be thought of as nonpropositional.

Hayes agreed but noted that something goes on in vision and that beliefs do arise. Obviously, one has beliefs, but they cannot explain perception, Turvey answered. Hayes then moved on to David Marr's work as an example of defensible work that includes his kind of inference making, processing with a semantics, something that preserves truth.

Trying further to clear away all misunderstandings, Hayes described the way his colleagues tend to interpret claims by Gibson-inspired ecological psychologists. Rather than looking at AI work such as Marr's and saying "OK, we're not doing that," they seem to say "that ought not to be done." But then the alternative seems to be that a magic semantic wand sends "chairness" to the eye when looking at a chair. How can this be? Hayes queried. Even if there was plenty of meaning "in the light," something would have to discover it or extract it. That could not be by natural laws. The symbolic relations involved are too arbitrary.

Shaw replied that for Gibson, meaning was not something to be put into a system but that it flows through or bathes a system by virtue of the perceiving/acting loop. Without the acting side to complete the loop, meaning would be left out.

Eleanor Gibson added that her husband's meaning of direct perception had never been a denial of the nervous system and some kind of "processing," but was an assertion that the perceiver did not have to add information.

Hayes then suggested that another reason some people looked askance at ecological theories was an implicit question, "How could you build one of those?" Later, in passing, Hayes asserted that ecological psychologists believed that perceiving was direct and given. Shaw quickly pointed out that these are not the same, that neither he nor any other adherent of direct perception takes "direct" to mean "given" (where "given" refers to presuppositions or assumptions). Rather "direct" is a kind of mapping.

Turvey then explained that direct perception might best be taken as the name of a program of research where the object is to elaborate science in such a way as to sustain the argument. It is a target for science. That is another reason there are "mechanisms" of the familiar type.

"OK, Hayes summarized, that's your placeholder. The rules are our placeholder."

6. Objects of Research. Mace suggested that Gibson and Bransford and Franks differed from information processing theorists by directing some of their research to the phenomena themselves rather than tests of underlying hypotheses. Thus, Gibson seemed to be genuinely wrestling with the question of what perceiving is like (particularly in its exploratory modes) while AI researchers never raise the question. That is, they seem to assume that certain pattern recognition or classification tasks embody perception truly enough to concentrate on developing models that would work in any simple instances. Mace asked Bransford if this characterized his work. Bransford answered that he was still seeking the phenomena that one might later try to formalize. At the risk of being misunderstood, he hastened to

add that he did not randomly select situations to see if learning occurred in them.

7. Skill Learning. A number of people during the day had expressed interest in examples of skill learning. Elliot Saltzman proposed to turn to a different part of the question - the novice. There are some attempts to characterize what an expert has learned, but what is a novice doing?, he asked.

Nathan Knobler said that the novice has a great deal of information that gets in the way. An art teacher raises questions about the novice's theories. Saltzman observed that novices tend not to appreciate natural constraints. They think all possibilities are viable. Smith reminded the group that once again first language learning is outside this domain. It does not look like skill acquisition as much as embryological development. There is relatively little exploration of odd alternatives in first language learning. Children do not make weird errors much. Saltzman agreed, saying that learning to walk was the same. What he had in mind was more like second language learning. Knobler mentioned that novices show a good deal of rigidity. Saltzman added that their joints are literally stiff and that they are reluctant to yield control.

8. Perceptual Learning. Citing his dissertation research, Bill Warren introduced some facts about stair climbing. The energy demands of stair climbing vary with riser height and the height of the person. Shown pictures of stairs in a way that maximizes their chances of seeing the scale in terms of the room they are in, people express preferences for stairs whose riser heights are .26 of their leg length. How are we to think of this?

Bob Remez pointed out that arm lengths vary a good deal more than leg length for the same height and wondered, therefore, what would happen if people were to choose stairs to crawl up?

9. Comparing Types of Search. Steve Braddon completed the discussion period with a question about cognitive search. That is, we often hear talk about searching memory or knowledge for information. How is this similar or different from a search for information in a real environment?