

The Intelligent Hand: An Experimental Approach to Human Object Recognition and Implications for Robotics and AI

Susan J. Lederman and Roberta L. Klatzky

■ The information in this article was originally presented as a keynote invited talk by Susan Lederman at the Thirteenth International Joint Conference on Artificial Intelligence in Chambéry, France; it is based primarily on a joint research program that we conducted. We explain how the scientific study of biological systems offers a complementary approach to the more formal analytic methods favored by roboticists; such study is also relevant to a number of classical problems addressed by the AI field. We offer an example of the scientific approach that is based on a selection of our experiments and empirically driven theoretical work on human haptic (tactual) object processing; the nature and role of active manual exploration is of particular concern. We further suggest how this program with humans can be modified and extended to guide the development of high-level manual exploration strategies for robots equipped with a haptic perceptual system.

We are experimental psychologists interested in the nature of intelligence. Consider the range of work that is being carried out on artificially and naturally intelligent systems and allow us to describe its domain in the broadest sense by including not just thinking but sensing and perceiving, thinking, and motor actions on the environment. We argue that the scientific study of biological systems offers an approach to the development of sensor-based robots that's complementary to the more formal analytic methods favored by roboticists.

First, we comment on why those who

design sensor-based robotic systems might want to consider lessons offered from the scientific study of intelligent living systems. We are certainly not the first to show how such scientific results might relate to intelligent systems. We need only think of how David Marr used biological vision to guide his computational modeling of the processes and underlying object representations for solving different machine vision problems.

We offer you a specific example of how biological work might be applied to robotics that is based on selected results from our own research program. In colloquial terms, the example focuses on how we humans perceive and recognize objects with only our hands. In scientific terms, the problem is known as *haptic object processing*. As formally defined, the *haptic system* is an information-processing system that uses input from receptors in skin, muscles, tendons, and joints to perceive objects and their layout and to guide actions within it.

We introduce you to a selection of our experiments and empirically driven theoretical work on human haptic object processing, with particular attention given to the nature and role of active manual exploration. Finally, we suggest how this research program with humans can be modified and extended to guide the development of high-level manual exploration strategies for robots equipped with a haptic perceptual system. Although this article focuses on robotic sensing and

This article is adapted from an invited talk published in the *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, and used by permission of the International Joint Conferences on Artificial Intelligence, Inc. Copies of this and other IJCAI proceedings are available from Morgan Kaufmann Publishers, Inc. 340 Pine Street, 6th Floor, San Francisco, CA 94104 USA. All changes to the original paper appearing in this article are copyright © 1994, AAAI

exploration, it applies to AI in general, with special relevance for areas such as intelligent interfaces, fine motion planning, object recognition, and knowledge representation.

General Lessons for Robotics and AI

Let's begin by posing the following question: In what ways might the scientific study of living systems help roboticists? We provide the general context behind the question first; then, we suggest five specific lesson areas.

It's possible to describe a continuum along which we can place robotic systems. At one end would be those that clearly try to reproduce natural living systems. At the other end would be those that equally blindly reject the anthropomorphic approach. We propose a different strategy that roboticists might adopt for gleaning potentially valuable information from the scientific study of biological systems. It offers new conceptualizations, new methodologies—scientific ones—and specific experimental findings about how living systems deal with problems that roboticists have yet to solve. This approach requires studying various constraints on biological information processing, which, in turn, involves anatomical, biomechanical, neurophysiological, and behavioral considerations.

Here are the five specific lesson areas:

Lesson 1—Performing in Unstructured Environments: We don't expect the most likely applications of a biological approach to be found in highly structured environments that can precisely be controlled or modified, such as in industrial automation; here, we might find little benefit from studying human processing.

It can be found rather in those robotic environments that we humans have little or no control over—for example, underwater repair and recovery, the service and maintenance of the space station, the disposal of radioactive waste, the exploration of unknown planets, and the performance of microrobotic surgery. To operate within these kinds of highly unstructured environments, roboticists might benefit from learning about how biological systems accomplish complex sensory, cognitive, and motor tasks in flexible and efficient ways.

Lesson 2—Overlapping Problem Domains: Biological scientists have addressed many of the same problems that roboticists now face. For example, we, too, have had to concern ourselves with evaluating the sensing capabilities of many different types of sensors. We, too, have been concerned with the

combining of information from multiple sensors, feature selection for scene segmentation and object recognition, and the nature of object representations. We, too, have come to recognize that perceptual systems must actively explore their environments. We also need to understand the basis of the motor-control and planning processes that allow certain living organisms to reach, grasp, and manipulate their environments so effectively. Finally, both biological scientists and roboticists must address hardware considerations as well to understand how such constraints affect the way information is processed and represented and, thus, the system's performance.

Lesson 3—Living Organisms Are Functioning, Multilevel, Integrated Systems: It's important to recognize that biological organisms are complete multilevel, integrated systems that actually work, despite the enormous complexity of the many problems they must handle. Living systems clearly highlight the complexity of the task facing the roboticist.

Lesson 4: Designing Human-Machine Interfaces for Teleoperation: Initial predictions concerning the relatively rapid creation of highly flexible, sensor-based autonomous robots have proved to be overly optimistic. So many roboticists have now turned their energies to designing teleoperated robots, which retain the human operator in the control loop. The rationale is that it is possible to short circuit the design process by taking advantage of our own considerable sensory, cognitive, and motor competencies. With an intact human nervous system, it is no longer necessary to build an artificial one, which all of you appreciate is no small task!

However, those of us who study human systems are quick to point out that for the teleoperational approach to succeed, it becomes critical for the roboticist to learn about how our own human subsystems process information and about the constraints under which they operate. Such capabilities and limitations must be understood to achieve an effective interface with any teleoperated system. We know of too many costly systems that, from an engineering perspective, have been judged to involve state-of-the-art design, yet they can't be controlled effectively by the human operator! In these cases, the human system was never acknowledged to be an integral component of the design; in other cases, this oversight was recognized but too late to do much about it.

Lesson 5—The Scientific Method Offers a New Approach: In robotics and engineering,

We propose a different strategy that roboticists might adopt for gleaning potentially valuable information from the scientific study of biological systems.

The scientific method provides an alternate and complementary approach to the design of sensor-based robotic systems.

it is most common to model a system analytically using differential equations. The parameters are determined by principles of physics. However, because the system is typically too complex to model without resorting to approximations, the parameter set is reduced arbitrarily by the investigator.

The scientific method provides an alternate and complementary approach to the design of sensor-based robotic systems. All science is based on careful, systematic, and repeatable observation. When we systematically control or manipulate the parameters under investigation, we are specifically said to be using the experimental method.

The scientific study of any system provides a coherent framework within which to study a given problem, whether this problem pertains to living or artificial systems. It provides formal methods for data gathering, some merely observational and some experimental; it also provides techniques for statistical analysis. With these methods and techniques, we can systematically and rigorously test the validity of our hypotheses. Our conclusions must be based on empirical data that we have collected according to appropriate scientific procedures.

We see such issues as being critical as well to the successful development and implementation of intelligent sensor-based autonomous and teleoperated robotic systems. The rigorous principles and methodologies of the experimental method expose some of the weaknesses and limitations of current robotics practice.

We have argued that the scientific approach offers roboticists a powerful set of general tools with which to complement their formal analytic methods. For additional information, refer to Lederman and Pawluk (1992) and Lederman, Klatzky, and Pawluk (1992); in the former article, you can find a mini-tutorial on the scientific method and its applications to robotics.

The Human Haptic System and Object Processing

Let's turn now to one example from our work that involves the scientific study of an intelligent biological system—the human haptic system—and how it processes and represents objects in the mind.

Background

Several years ago, we and Tory Metzger (Klatzky, Lederman, and Metzger 1985) experimentally demonstrated that humans are

remarkably skilled at recognizing common objects (for example, a hammer) using only our sense of touch. We asked blindfolded subjects to identify a set of 100 common objects as quickly and as accurately as they could. Their accuracy approached 100 percent, and the majority of objects were identified within just 2 to 3 seconds. This result was really surprising because at that time, many were suggesting that human touch is incapable of such high-level information processing. We began to suspect that how people actively and manually explore such multidimensional objects might play a critical role in uncovering, as well as eventually explaining, the substantial information-processing capacities we had demonstrated.

In the next experiment (Lederman and Klatzky 1987), we asked subjects to perform a haptic match-to-sample task; on each trial, subjects were initially presented with a standard object followed by three comparison objects. Although the four objects in any one set varied along many different dimensions (for example, texture, compliance, shape), subjects were told to attend to only one dimension, such as texture, as shown in figure 1. Subjects had to select the comparison object (marked with an asterisk) that best matched the standard on the dimension named. Over the entire experiment, we produced a number of different custom-designed, multidimensional object sets. Each set was designed specifically to be used with a particular property-matching instruction. Over the entire experiment, the instruction for any single trial was selected equally often from the following property set: texture, hardness, thermal properties, weight, volume, envelope shape, and precise (exact) shape.

We videotaped and subsequently analyzed subjects' hand movements during each trial. Our results indicated that manual exploration is systematic; subjects performed highly stereotypical movement patterns that we call *exploratory procedures* (EPs). They also chose to execute a particular EP in association with a specific property-matching instruction. Here we show you typical versions of the six EPs most relevant to this discussion. We describe their invariant properties and the property-matching instruction with which that EP was most commonly associated (figure 2).

The *Lateral Motion EP* involves back-and-forth tangential movements over a surface; it was typically performed with the texture-matching instructions. A *Pressure EP* involves applying normal forces or torques about an object axis; it was usually selected for hard-

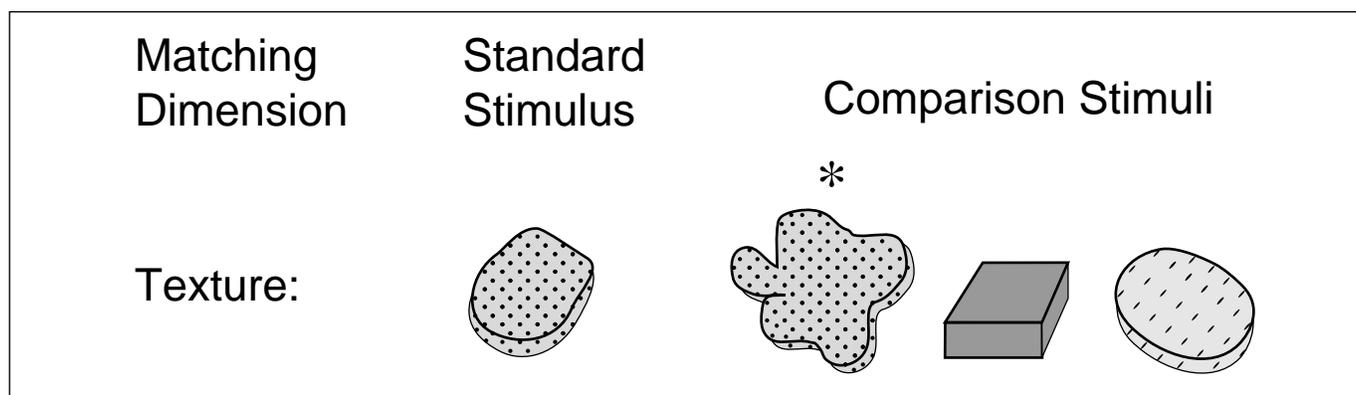


Figure 1. Sample Set of Multidimensional Objects, Consisting of One Standard and Three Comparison Objects, Used with Texture-Matching Instructions in a Match-to-Sample Task.

The asterisk indicates the comparison object that most closely matches the standard in terms of texture.

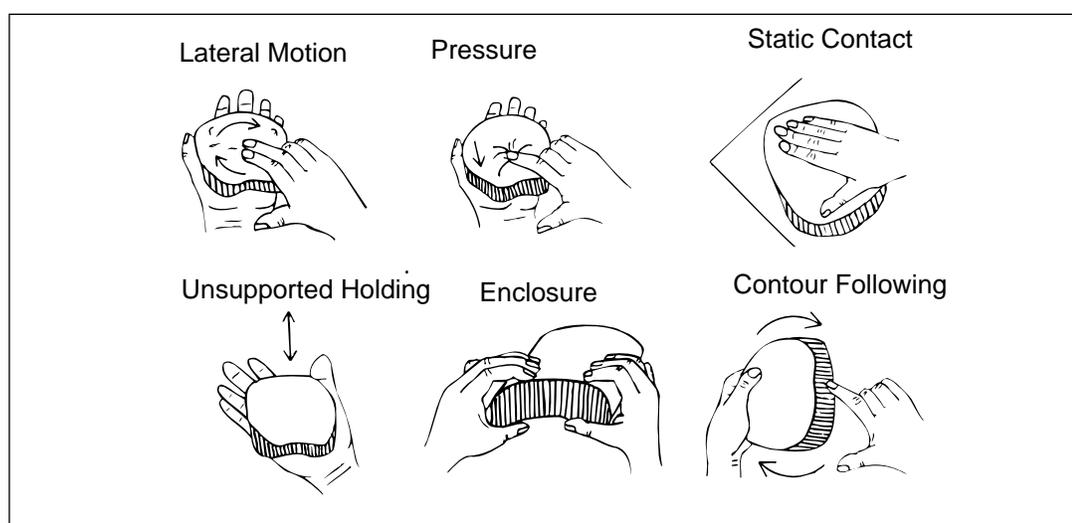


Figure 2. Primary Exploratory Procedure Classes Showing Typical Manual Activity Patterns (revised with permission from Lederman and Klatzky [1987]).

ness. *Static Contact*, or simple contact between an object surface and the skin, was associated with thermal matching. *Unsupported Holding*—lifting the object away from a supporting surface, usually in the form of a dynamic heft—was used for weight matching. *Enclosure*, or molding the fingers and palm to the object contours, was selected to get both volumetric and envelope-shape information. Finally, *Contour Following*, or edge following, was used most often in conjunction with both envelope-shape and exact-shape instructions.

The Macrostructure of Human Haptic Object Recognition

Now let's turn to our proposed macrostructure for the human haptic object-identifica-

tion system. A number of computational models in the field of cognitive science have successfully dealt with broad and complex domains of human information processing.

Our own general approach to haptic object processing can be appreciated first by analogy with the computational model of visual reading proposed by Just and Carpenter (1980), which includes many processes, from low-level eye fixations to high-level text comprehension.

Figure 3a shows a bare-bones outline of the various stages into which Just and Carpenter separate visual text comprehension: Following text input, the eye is moved to the next location, it fixates on a local region, a local analysis of the word is performed (for example, a *lexical analysis*), and a global representa-

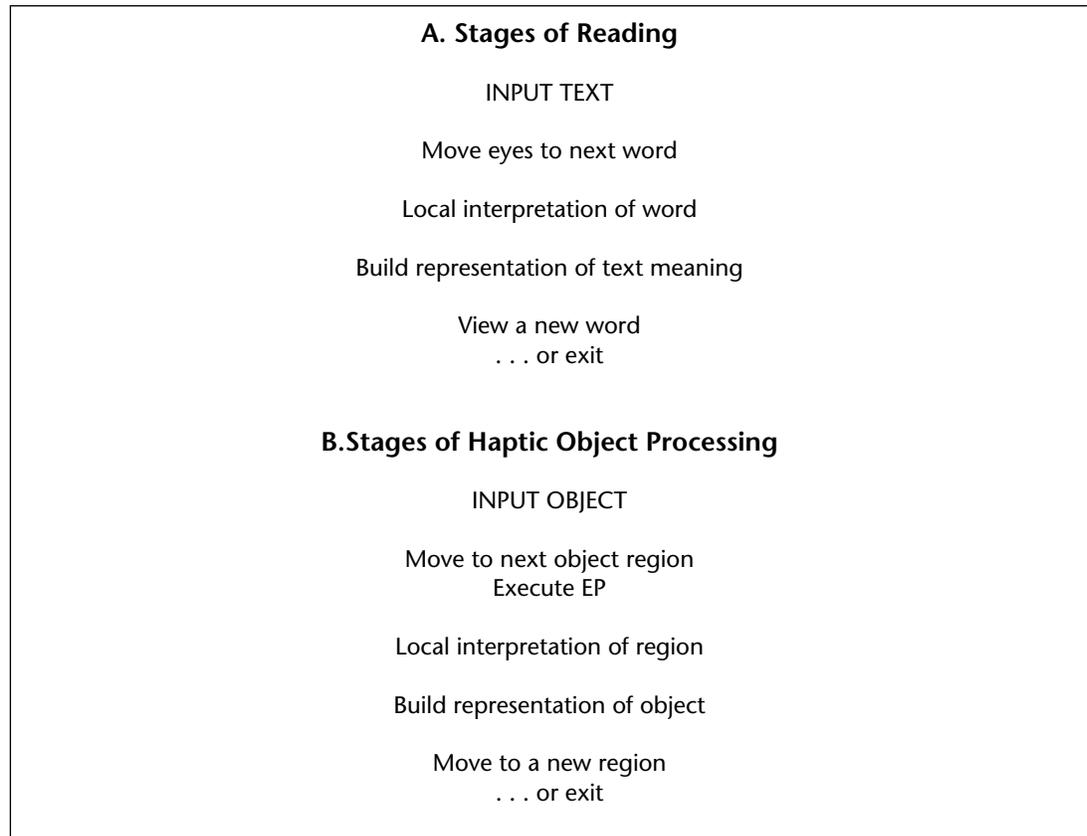


Figure 3. Stages of Reading and of Haptic Object Processing.

A. Summary stages of a model of reading by Just and Carpenter (1980): from eye fixations to comprehension (revised with permission from Klatzky and Lederman [1993]).

B. Parallels with a model of haptic object processing proposed by Klatzky and Lederman (1993).

tion is created. Within a fixation, these stages are performed as completely as possible; over successive fixations, they subsequently recur. The representation is continually updated with the input from each new fixation.

Our model of haptic object processing also includes a number of stages, from manual exploration to object identification. These stages are shown in figure 3b.

A period of manual exploration corresponds to a period of eye fixation. During the manual period, what we call the *selection-extraction loop* takes place. One or more EPs are selected and performed on some area on the object. The resulting data are used to interpret a local object region, which, in turn, is used to build a global object representation for comparison with stored categorical representations. As much as possible of the local and global processing is performed within an exploratory period. Both types of processing recur during subsequent exploratory periods and object regions. Eventually, the system recognizes an object or selects another EP for execution.

Figure 4 presents a simplified version of our model of the macrostructure of the haptic object-identification system. It includes the different data representations and the links we presume exist among them.

The *object component* represents particular objects (for example, a wrench) and their specific property values. The *property component* represents the attributes along which an object might potentially vary (texture, hardness, and so on) rather than specific property values. The *EP component* represents exploratory procedures.

Now let's look at the connections between various parts of our proposed macrostructure, which are represented by the arrows in figure 4. Links between objects within the object component (labeled *object relatedness*) are mainly intended to reflect the hierarchical classification relations shown in the cognitive literature. For example, wrenches and pliers are linked because of their common tool-related functions. Links between object and property values (labeled *object descriptions*)

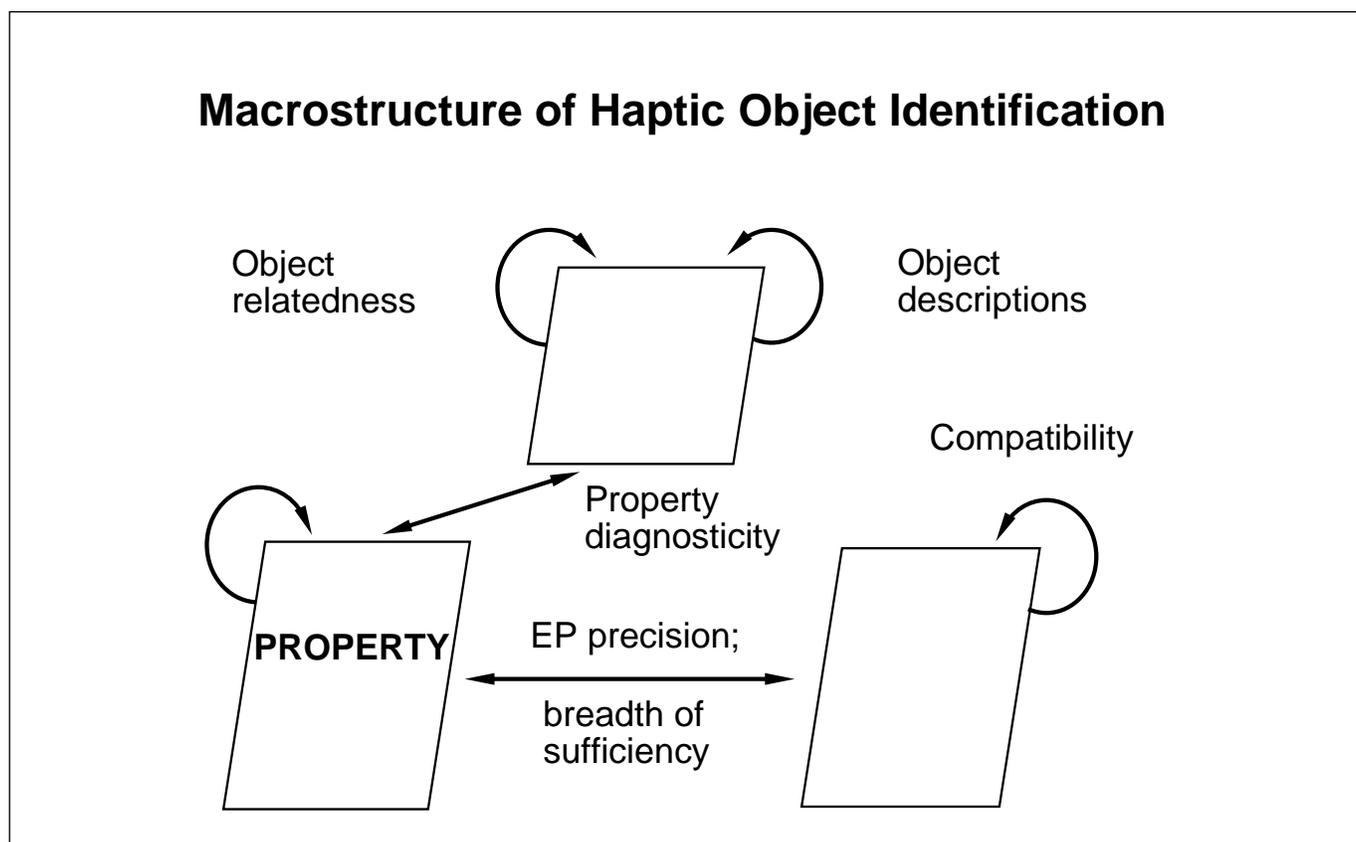


Figure 4. Proposed Macrostructure of Haptic Object Identification (revised with permission from Klatzky and Lederman [1993]).

reflect the relative strengths of a given property for a particular object (as texture is important for sandpaper). Links between object and property components (labeled *property diagnosticity*) represent the relative importance of each property for a given object class. Links between EP and property components represent the relative precision of information about a property extracted by a particular EP and the breadth of EP sufficiency, which we define shortly.

We have empirical data that address both of these factors. They were obtained using a variant of the match-to-sample experiment we discussed earlier. (Later, we suggest how these two experiments, which involve property extraction under free versus constrained exploration, can be applied to the robotic domain.) Let's have a closer look at the variant of the first experiment. On any trial, subjects were now only allowed to perform one of the six EPs shown in figure 2 to extract a single named dimension (for example, weight matching with the Lateral Motion EP). Over the entire experiment, all possible EP-prop-

erty combinations were performed. We measured both accuracy and response time and used these data to compare the relative precision with which each EP could extract a designated property.

The results are shown in table 1 in the form of an EP-property weight matrix, with EPs varying down the rows and property-matching instructions varying across the columns. The entries are based on relative accuracy and speed. A cell entry of 0 indicates that subjects could not perform the property-matching task above chance level with the EP shown (as was the case, of course, with the example in which weight was to be obtained using Lateral Motion). An entry of 1 indicates sufficient but not optimal performance (such as when Lateral Motion is used to extract hardness). A 2 indicates performance with this EP was sufficient and optimal but not necessary (such as when Lateral Motion is used to extract texture). A 3 indicates that the given EP was necessary as well as optimal (as in the case when Contour Following was used to extract precise shape details). In general,

	tex	hard	temp	weight	volume	global shape	exact shape
Lateral Motion	2	1	1	0	0	0	0
Pressure	1	2	1	0	0	0	0
.
.
.
Contour Following	1	1	1	1	1	1	3

Table 1. EP-Property Weightings.

those EPs that were executed spontaneously in the free exploration match-to-sample experiment tended to produce optimal performance in the constrained exploration version of the experiment. To summarize, these data let us determine the relative precision with which each EP can extract a particular property.

These data can tell us even more about EPs. By summing the number of nonzero cells across a row in this table, we can also represent the relative breadth of sufficiency of each EP, as shown in table 2. Look at the left column of numbers. Lateral Motion and Pressure (at the top) provide sufficient information only about a few different properties. However, Enclosure and Contour Following (down at the bottom) offer coarse information about most object properties considered in this study. Enclosure and Contour Following are clearly the most broadly sufficient of our EP set. However, the breadth of property information provided by Contour Following must be weighed against its relatively slow execution time, shown in seconds, in the second numeric column. You can see that we've learned a lot about the nature of the links between EPs and properties from these experimental data.

	Breadth of Sufficiency	Duration (in seconds)
Lateral Motion	3	3.5
Pressure	3	2.2
Static Contact	4	0.1
Unsupported Holding	5	2.1
Enclosure	6	1.8
Contour Following	7	11.2

Table 2. Generality and Average Duration for Each EP.

Finally, we want to consider the links between EPs (labeled *compatibility*). Such connections represent the extent to which designated EPs can be coexecuted, each EP extracting the properties for which it is at least sufficient. For example, Lateral Motion and Pressure can be executed together, which, in turn, provide information about both texture and hardness.

We have developed a set of visible kinematic and dynamic parameters that formally differentiate EPs. We derived them from an extensive body of hand-movement data, which are based on a large number of multi-dimensional objects, both common and custom designed. The objects were tested over a broad range of experimental conditions. What we found is that values of four parameters occurred consistently for each EP across these different testing conditions.

We can describe each parameter as capturing some constraint inherent in an EP when it must be performed to extract certain types of information. The four parameters and their possible values are shown in table 3. The parameters include movement (either static or dynamic), the direction of force applied (tangential or normal to the surface), the region of the object contacted by the end effector (interior surfaces, edges, or both), and a workspace constraint (whether a supporting surface was required or not). We assume that compatibility between a pair of EPs only exists to the extent that the constraints inherent in their parameter values can be satisfied simultaneously through some type of exploration.

The pattern of parameter values for each EP allowed us to determine if any two EPs are compatible. Clearly, a pair of EPs with identical parameter values are compatible, but this information isn't informative because EPs can't be differentiated. However, it's still possible to achieve EP compatibility: With certain nonidentical parameter sets, people apparently select some form of exploration that satisfies the constraints inherent in both EPs.

Parameter	Parameter Values
Movement	Static/Dynamic
Direction	Tangential/Normal
Region	Surfaces/Edges/ Surfaces+Edges
Workspace Constraint?	Yes/No

Table 3. EP Parameterization.

	Pressure	Lateral Motion	Enclosure	Contour Following	Unsupported Holding
Static Contact	+	-	+	-	+
Pressure		+	+	-	+
Lateral Motion			-	+	-
Enclosure				-	+
Contour Following					-

Table 4. EP Compatibility.

The rules used to determine EP compatibility are published elsewhere (Klatzky and Lederman 1993); table 4 shows the results of applying these rules in the form of an EP-EP compatibility weight matrix. Our EPs are shown along the rows and columns: A + represents compatibility between two EPs (for example, between Lateral Motion and Pressure); a - represents an incompatibility (for example, between Lateral Motion and Static Contact).

A little later we show you how we used these empirically derived EP-EP compatibility associations in developing a computational model of the EP-selection process, which we turn to now. In keeping with the interactive activation perspective, we treat haptic object identification as a parallel interactive process, with sequential constraints imposed by the execution of exploratory procedures.

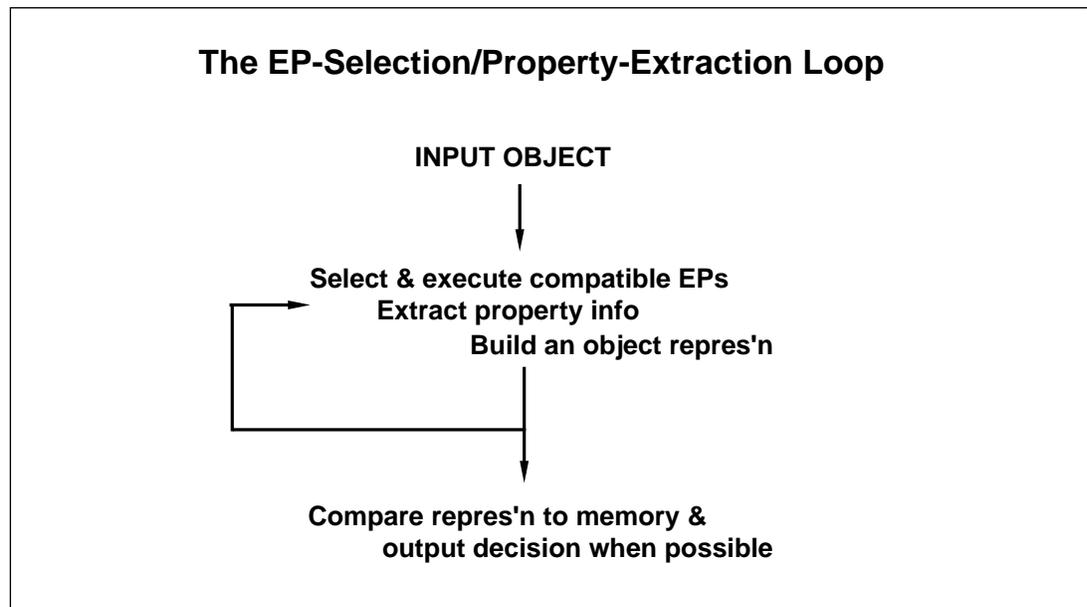
Selecting EP and the Selection-Extraction Loop

Figure 5 shows how the process proceeds in a sequence of EP-selection–property-extraction loops. During each step, an EP is selected and executed along with any other compatible EPs. In this way, information about associated properties is extracted; the precision of the information is determined by the weights on the links between EPs and properties. Over a

sequence of these loops, an object representation is built up and used as a probe to match against stored object representations. When a match criterion is satisfied, the search process is terminated, and the object is said to be recognized.

Constraints on EP Selection The primary goal during the selection-extraction loop is to choose an EP for execution under a number of competing constraints (for example, the need to know as much about the object as quickly as possible or the need to learn about a desired object property). There can also be inherent biases that govern the use of certain EPs; for example, Contour Following is relatively time consuming and, for humans, also fails to provide sufficiently precise contour information to achieve fine shape discrimination. This fault might generally discourage the use of Contour Following. In fact, an Enclosure might be favored more generally because like Contour Following, it, too, provides coarse information about many different properties, but in contrast, it's relatively fast.

In principle, these constraints and biases can be represented by the weights between different components in our haptic object-identification system. For example, connections between specific objects can be represented by associative weightings on the



*Figure 5. The EP-Selection-Property-Extraction Loop
(reprinted with permission from Klatzky and Lederman [1993]).*

object-object links within the object component. Expectations concerning the diagnostic value of various properties can be represented by links between the property and object components. Hand-movement precision and breadth of sufficiency can be represented by connections between property and EP components, as shown in table 1. Recall how we derived the associated weights from the results of our constrained exploration match-to-sample experiment. Finally, EP compatibility can be determined by constraints inherent in the EP parameters, producing the matrix of binary EP-compatibility weights in table 4.

A Constraint-Satisfaction Approach
Collectively, these constraints function to select an EP given a specific object and certain prior expectations. In connectionist terms, the EP-selection process can be treated as a constraint-satisfaction algorithm in which the weights serve as constraints to be relaxed progressively until some elements are maximally activated.

In our case, we have used constraint satisfaction as a method for selecting the next EP in a sequence during manual exploration. The weights are theoretically and empirically derived associations among EPs and properties (and, potentially, objects). As the system progressively relaxes, a stable activation pattern eventually emerges that is used to predict which EP will be executed next in some exploratory situation.

To consider the consequences of the associations between EPs and properties and of the compatibilities between EPs, we implemented the weights shown in the two associated weight matrix tables as a constraint-satisfaction system. The nodes represented EPs and properties. This system is equivalent to examining a single generic object.

Each property was clamped to represent an externally set property goal; for example, look for texture. To represent the situation in which an observer initially wants to extract as much information about an object as possible, we used the full matrix of EP-property weights. We found that no matter which EP was clamped, that is, regardless of the specific property instruction, the maximum activation level always occurred for an Enclosure (or grasp), which is not only broadly sufficient but also compatible with other EPs. The next most active element was Unsupported Holding (or lifting), which is compatible with Enclosure as well as broadly sufficient and relatively quick to execute. Thus, this pair of EPs could be selected and executed within the same loop to provide a great deal of information rapidly.

Now we describe what happened when we modeled a different situation, which might occur after a stage of coarse exploration, that is, when you want more precise information than can be obtained with an EP that is only sufficient for that property. You have to execute the optimal EP. To model this scenario,

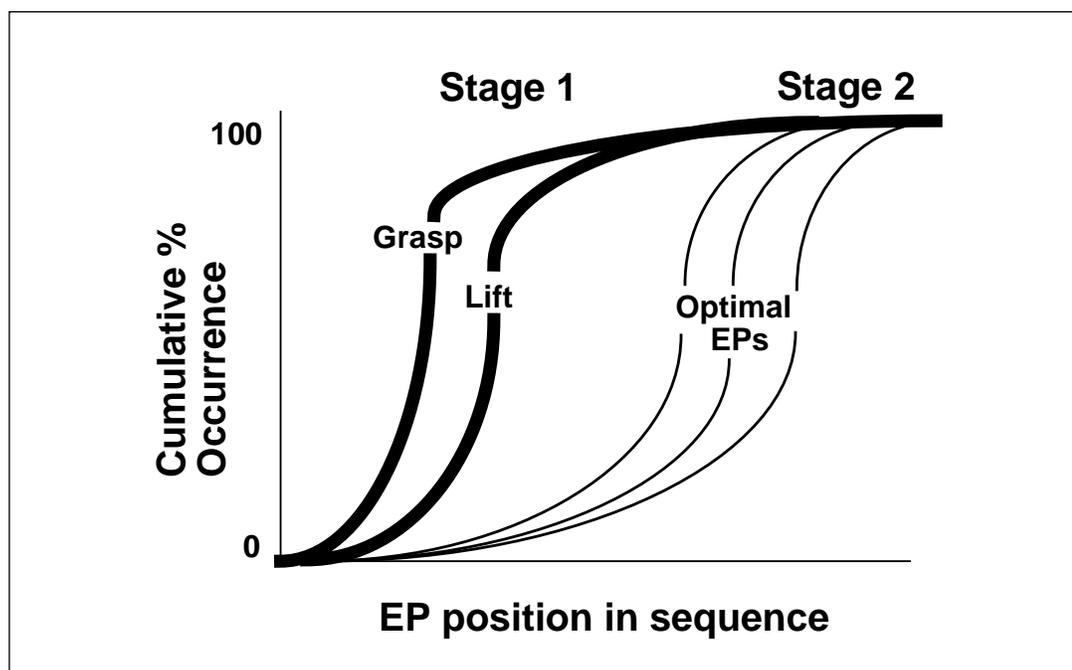


Figure 6. Schematic Representation of the Two-Stage Sequence of Manual Exploration in a Common Haptic Object-Classification Task (revised with permission from Lederman and Klatzky [1990]).

The occurrence of each EP (indicated as the cumulative percent occurrence of all EPs) is plotted as a function of its temporal position in the EP sequence.

we used a second set of weights in which only the optimal EP-property weights in the matrix were included. Here, we found different activation levels; now, clamping a particular property resulted in only the optimal EP being selected, such as Lateral Motion for texture.

Behavioral Support

Let's turn now to some of the behavioral experiments we performed to test certain predictions made by our constraint-satisfaction approach.

Predicting the Two-Stage Sequence of EPs

The results of our modeling are supported by experimental evidence of a two-stage exploratory sequence (Lederman and Klatzky 1990), which was adopted by subjects in a common object-classification task. On each trial, we first asked our subjects a yes-no question such as, Is this abrasive surface further a piece of sandpaper? An object was then placed in the subject's hands to explore. In one-half the trials, the object was actually the sandpaper named in the question. However, in the remaining trials, a different abrasive surface was presented (that is, a metal file). We had already determined the most appropriate diagnostic property for each object class named in the questions in a previous experiment.

We analyzed the hand movements from each trial as a sequence of EPs. Our results indicated two separate stages of manual exploration. The data are schematically presented in figure 6. Each function depicts the cumulative percent of EP occurrence as a function of its serial position in the EP sequence. The two solid, dark lines indicate that the first two EPs in the sequence were Enclosure followed by Unsupported Holding, which together make up what we call a grasp-lift sequence. In keeping with the model's predictions, both EPs are broadly sufficient, compatible, and fast.

The remaining EP functions (the thinner lines) occurred after this initial, highly efficient exploratory sequence; the EP that was most often selected on any trial was successfully predicted by the property known to be most diagnostic for identifying the object. For example, in the case of our sandpaper question, we predicted Lateral Motion would be executed after the grasp-lift sequence because it is optimal for extracting texture, which we experimentally showed to be most diagnostic of sandpaper. This second stage of exploration was more specifically directed toward extracting additional precise information about the critical property. In short, these data on manual exploration during object

identification provide experimental support for our approach to exploratory control as a constraint-satisfaction process.

Experiments on the Selection-Extraction Loop and Property Extraction In addition to which EPs are selected, our model addresses how a person's choice of EPs affects the precision of the property information that can be extracted. We assume that the strength of the associations between EPs and properties should predict how much a person can perceive and learn about an object with whatever EPs he or she has selected. This assumption permits us to make several different predictions based on our model that highlight the gatekeeper role played by EPs in limiting the availability of information about object properties for purposes of object perception and recognition (Klatzky, Lederman, and Reed 1989).

The tasks generally required subjects to learn to classify multidimensional objects into groups, according to different classification rules. For these tasks, we designed subsets of planar waferlike objects that varied along four property dimensions: texture, hardness, shape, and size. Each object had one of three texture values (high, intermediate, or low roughness), one of three hardness values (high, intermediate, or low compliance), one of three shapes (one, two, or three lobes), and one of three sizes (small, medium, or large). Here we discuss one experiment that supports one of the predictions based on our assumption that EPs serve as the gatekeepers of property information about objects. (The other experiments are described in the full text of Lederman's talk published in the conference proceedings for the Thirteenth International Joint Conference on Artificial Intelligence.)

The model predicts that when a single EP is used to explore an object, incidental knowledge about other properties accrues depending on the associative weights between each property and that EP. To assess this prediction, we developed an experimental approach that we call the *redundancy withdrawal paradigm*.

We presented subjects with sets of planar objects that were redundantly defined according to different two-redundant-property classification rules. An example of a texture-hardness redundancy classification rule might be, All A's are very rough and very hard; all B's are of intermediate roughness and hardness; all C's are very smooth and very soft. In the experiments, we actually used sets of objects that were classified not only by texture-hard-

ness redundancy rules but also by shape-hardness and shape-texture redundancies. Subjects were initially told to classify the objects on the basis of a single named dimension (for example, texture), even though objects varied redundantly on two dimensions (for example, texture and hardness); when performance was asymptotic, unbeknown to the subject, the object set was switched. This new set could only be classified using a one-property classification rule. To continue our example, in this new set, objects only varied in texture because the value of the second property, hardness, was now constant or, as we say, *withdrawn*. We reasoned that if subjects had previously incidentally learned about the second property, then their response times should increase just after it is withdrawn.

As shown in figure 7, we obtained a strong withdrawal effect for texture-hardness redundancies regardless of which dimension was withdrawn (shown by two separate functions). Response times are plotted as a function of ordered blocks of trials. You can see the large increase in response times for both functions at the vertical line, which indicates the point at which the second dimension was withdrawn. However, the withdrawal effects for the other two redundant combinations (texture-shape, hardness-shape) were negligible. This result was also expected because Contour Following (which is necessary as well as optimal for extracting full shape details) is incompatible with both Lateral Motion (used for texture) and Pressure (used for hardness). With these planar objects, shape is typically available only at the edges, but both texture and hardness are extracted best from interior surface areas. To summarize, this study supports one prediction made by the model: Selecting compatible EPs allows for their coexecution, which makes available information about all properties for which either EP is at least sufficient.

Some years ago, the psychologist J. J. Gibson argued eloquently for the importance of active exploration by biological perceptual systems, and indeed, much of the subsequent work has supported his position—a tribute to his early insights. Our own work clearly confirms and also extends the importance of active exploration with respect to human haptic object processing.

Application to Robotic Exploration

In the robotic domain, Ruzena Bajcsy can be credited with emphasizing the need for active machine exploration, particularly when

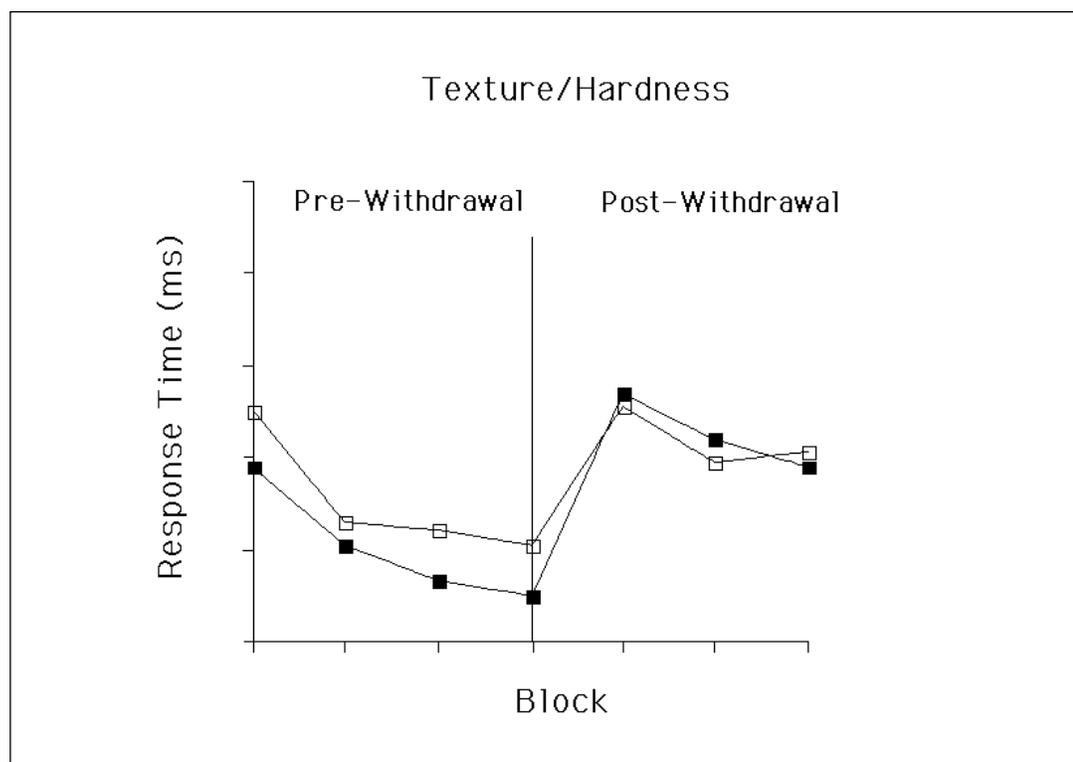


Figure 7. Response Time as a Function of Ordered Blocks of Trials for a Texture-Hardness Redundancy Classification Task.

The vertical line indicates the point at which either texture (solid symbols) or hardness (open symbols) was withdrawn (revised with permission from Klatzky, Lederman, and Reed [1989]).

information about object properties must be used to interact with unstructured environments. Presumably, such is the case whether or not identification is required.

In keeping with this notion, Bajcsy and others have now adopted the concept of an EP as a systematic manual testing procedure and have implemented robotic versions of our EPs. Roboticians next face the problems of EP selection and the choice of appropriate execution sequences for intelligent robotic exploration.

A Scientific Approach We suggest adopting our constrained exploration match-to-sample experimental paradigm to develop a more general robotic search strategy for EP selection when multidimensionally varying objects are explored. This experimental approach permits one to systematically determine the relative performance characteristics of robotic EPs selected, which could then be used in conjunction with a constraint-satisfaction approach to choose efficient EP-execution sequences.

Consider the following scenario. We begin with the view of an EP as a motoric routine that is optimal for extracting one property,

although it might also be sufficient for extracting several others. The properties and associated EPs will depend on the particular robotic end effector and sensing system chosen and could be different from what humans use, unless an anthropomorphic design has deliberately been adopted.

Having selected a set of properties and EPs, one can now experimentally test the relative constraints on EP performance; however, this is defined by the roboticist. The constraints on human EPs can be applied to any exploring system, along with others that are specifically relevant to the robotic domain.

The constrained match-to-sample experiment can be used as a methodological tool for systematically evaluating the relative performance of each EP in extracting each property. Also relevant to this approach is the extent to which robotic EPs can be coexecuted, that is, the issue of EP compatibility.

Collectively, the results concerning the relative strengths of the EP-to-property and EP-EP compatibility associations can be used to rank robotic EPs for use in computational models of EP selection during active exploration of highly unstructured environments.

Summary

In closing, we want to repeat the claim that the scientific study of biological organisms can further the development of current sensor-based robots in many different ways without being constrained by, or limited to, an anthropomorphic design.

General Lessons

On a general level, we argued that such work addresses many of the same problem domains; provides an example of, and framework for, designing working, multilevel, integrated systems; and offers valuable suggestions for presenting robotically extracted information to a human operator for purposes of teleoperation. In addition, the scientific method highlights the value of properly constraining the problem; formulating testable hypotheses; evaluating these hypotheses experimentally with rigorous and unbiased tests; and using statistical techniques for assessing the validity, reliability, and generality of the experimental findings. Such comments are just as relevant to the field of AI.

Substantive Lessons

On a more specific level, we have also made a number of particular suggestions for designing robotic tactile and haptic systems, including the example presented here. All the suggestions are based on the scientific results of experiments on biological touch.

An Interdisciplinary Approach

For biological scientists who attempt to understand the bases of natural intelligence and for roboticists and AI researchers who attempt to create such behavior in machines, serious collaboration might provide a new and potentially fruitful approach to system design.

Acknowledgments

We would like to acknowledge Dianne Pawluk, Cathy Reed, and Tory Metzger, with whom we published the work discussed here. We'd also like to thank Ruzena Bajcsy for the many stimulating interactions we have had with her and her colleagues and students in the GRASP (general robotics and active sensory perception) Lab at the University of Pennsylvania concerning the connections between natural and artificial sensing systems. Preparation of this talk was supported by a grant from the Institute for Robotics and Intelligent Systems Centre of Excellence of Canada.

References

Just, M. A., and Carpenter, P. A. 1980. A Theory of Reading: From Eye Fixations to Comprehension.

Psychological Review 87:329–354.

Klatzky, R. L., and Lederman, S. J. 1993. Toward a Computational Model of Constraint-Driven Exploration and Haptic Object Identification. *Perception* 22:597–621.

Klatzky, R. L.; Lederman, S. J.; and Metzger, V. 1985. Identifying Objects by Touch: An "Expert System." *Perception and Psychophysics* 37(4): 299–302. Klatzky, R. L.; Lederman, S. J.; and Reed, C. 1989. Haptic Integration of Object Properties: Texture, Hardness, and Planar Contour. *Journal of Experimental Psychology: General* 15(10): 45–57.

Lederman, S. J., and Klatzky, R. L. 1990. Haptic Classification of Common Objects: Knowledge-Driven Exploration. *Cognitive Psychology* 22:421–459. Lederman, S. J., and Klatzky, R. L. 1987. Hand Movements: A Window into Haptic Object Recognition. *Cognitive Psychology* 19(3): 342–368.

Lederman, S. J., and Pawluk, D. 1992. Lessons from Biological Touch for Robotic Tactile Sensing. In *Advanced Tactile Sensing for Robots*, ed. H. R. Nicholls, 151–192. Singapore: World Scientific Publishing.

Lederman, S. J.; Klatzky, R. L.; and Pawluk, D. 1992. Lessons from Biological Touch for Robotic Haptic Sensing. In *Advanced Tactile Sensing for Robots*, ed. H. R. Nicholls, 193–219. Singapore: World Scientific Publishing.



Susan J. Lederman is professor of psychology and computing and information science at Queen's University in Kingston, Ontario, Canada. She holds an M.S. and a Ph.D. in experimental psychology from the University of Wisconsin at Madison

and the University of Toronto, respectively. Her interests include human perception, cognition, and motor control, with particular emphasis on the sense of touch. She recently became interested in applying research on biological sensorimotor systems to the design of autonomous sensor-based robots and human-machine haptic interfaces for teleoperation and virtual environments. She and Klatzky have actively collaborated for 14 years.



Roberta L. Klatzky is a professor and head of psychology at Carnegie Mellon University. She received a B.S. in mathematics from the University of Michigan and a Ph.D. in experimental psychology from Stanford University. Her research empha-

sizes cognitive aspects of perception and motor performance. She has particular interests in knowledge representations achieved through touch, locomotion, and vision; interactions between knowledge-based and sensory processes; and action planning.