THE INTELLIGENT HAND: AN EXPERIMENTAL APPROACH TO HUMAN OBJECT RECOGNITION AND IMPLICATIONS FOR ROBOTIC DESIGN

Susan J. Lederman
Queen's University at Kingston
Depts. of Psychology and Computing & Information Science
Kingston, Ontario
Canada K7L 3N6

Abstract

The scientific study of biological systems offers an approach to the development of sensor-based robots that is complementary to the more formal analytic methods currently favoured by roboticists. I initially propose several general lessons from the biological field. Next, I consider a specific example selected from the work of Lederman & Klatzky, which focuses on human haptic object processing. An empirical base and recent theoretical developments from our research program on this topic are described. The human haptic system is an information-processing system that combines inputs from sensors in skin, muscles, tendons, and joints with motor capabilities to extract different object properties. A general model of human haptic object identification, which emphasises how object exploration is controlled, is presented. The model describes major architectural elements, including representations of haptically accessible object properties and exploratory procedures (EPs), which are dedicated movement patterns specialized to extract particular properties. These architectural units are related in processing-specific ways. The resulting architecture is treated as a system of constraints, which guide the exploration of an object during the course of identification. Empirical support for the model is also examined. To conclude, I show how this scientifically-based approach might be applied to developing strategies for active manual robotic exploration of unstructured environments.

1 Outline

In this talk, I treat the concept of "intelligence" as covering a broad domain that includes sensing and perceiving, thinking, and acting on the environment. In Section 2, I argue that there are a number of general lessons offered by the scientific study of intelligent biological organisms for sensor-based robotic design. In Section 3, I provide an example that focuses on the problem of human object perception and recognition. However, rather than follow the mainstream route by investigating vision, Roberta Klatzky and I have selected the "haptic" system for study. We formally define this as an information-processing system that uses inputs from receptors in skin, muscles, tendons and joints to perceive the concrete world and to guide actions within it. I present some of our experimental and empirically-based theoretical work on human haptic object processing, with particular emphasis on the nature and role of active manual exploration. In Section 4, I suggest how this research programme may be modified and extended to guide the development of high-level manual exploration strategies for robots equipped with a haptic perceptual system. Section 5 provides a more general summary of how knowledge of biological systems may contribute to the field of robotics.

the parameters are determined by principles of physics. As 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 sciences are based on careful, systematic and repeatable observation. In addition, when we actually systematically control or manipulate the parameters under investigation, we are 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 pertains to living or artificial systems. It provides formal methods (experimental and statistical) with which to systematically and rigorously test the validity of one's hypotheses, based on empirical results. We see such issues as being critical as well to the successful development and implementation of intelligent sensor-based (tele)robotic systems. The rigorous principles and methodologies of the experimental method expose some of the weaknesses and limitations of current robotic practice. We have argued that the scientific approach offers roboticists a powerful set of general tools with which to complement their formal analytic methods (see Lederman & Pawluk, 1992, for a mini-tutorial on the scientific method and its applications to robotics).

Let us turn now to one example involving the scientific study of intelligent biological systems, specifically the human haptic system, and how it processes and represents objects.

3 The human haptic system and object processing[2]

3.1 Background

Several years ago, we (Klatzky et al, 1985) demonstrated that humans are remarkably skilled at recognizing common objects (e.g., hammer) using only touch. We asked blindfolded observers to identify a set of 100 common objects as quickly and as accurately as possible. Subjects' accuracy approached 100%, while the majority of objects were identified within only 2-3 seconds. This was surprising at the time since others had suggested that the human sense of touch is incapable of such high-level information processing (e.g., Walk & Pick, 1981).

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 our next experiment (Lederman & Klatzky, 1987; Expt. 1), we asked subjects to perform a haptic "match-to-sample" task; on each trial, subjects were initially presented with a "standard" object followed by a set of 3 serially presented multidimensional "comparison" objects. Although all four objects in any set varied along many different object dimensions (e.g., texture, shape, etc., as underscored in Figure 1), subjects were instructed to attend to a single dimension, such as texture. They were to select the comparison object that best matched the standard object on the dimension named. Over the entire experiment, we used different unfamiliar custom-designed object sets for each of the dimension-matching instructions, some of which are shown (underscored) in Figure 1.

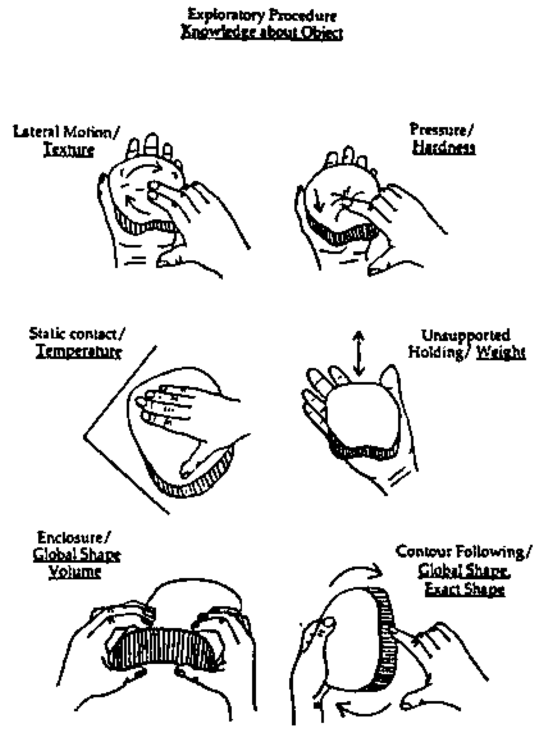

Figure 1. Exploratory Procedures and associated object properties (Lederman, 1991; adapted from Lederman & Klatzky, 1987).

We videotaped and subsequently analyzed subjects' hand movements during each trial. Our results indicated that manual exploration is very systematic; subjects performed highly stereotypical movement patterns that we have called "exploratory procedures" ("EP"s). Further, they chose to execute particular EPs in association with specific dimension-matching instructions, the ones most relevant for the current talk being shown in Figure 1. Thus, a Lateral Motion EP (tangential movements on a surface) was typically

[2]Section 3 is based on material from Klatzky & Lederman (1991 and in press).

2 General lessons for designing sensor-based robots[1]

2.1 Areas of application

It is possible to describe a continuum along which robotic systems can be placed. At one end, we would find those that clearly attempt to reproduce natural living systems. At the other end, we would find those that equally blindly reject the anthropomorphic approach. Yet there is an alternate approach that roboticists may adopt for deriving potentially valuable information from the scientific study of biological systems. According to this approach (e.g., Lederman & Pawluk, 1992; Lederman et al, 1992), one gains new conceptualizations, scientific methodologies, and specific empirical results about how living systems deal with problems that roboticists have yet to solve. Anatomical, biomechanical, neural, and behavioural constraints on information processing are all relevant areas of concern.

The most likely applications of a biological approach are not to be found in highly structured environments (e.g., industrial automation), which may be precisely controlled or modified - under such circumstances, there may be little benefit from copying human processing. It is rather in robotic environments over which the human has little or no control; examples would include those requiring underwater repair and recovery, service and maintenance of the space station, disposal of radioactive waste, exploration of unknown planets, and microrobotic surgery. For operation within such highly unstructured environments, roboticists may benefit from learning how biological systems accomplish complex sensory, cognitive and motor tasks in flexible, efficient ways.

2.2 Overlapping problem domains

Scientists who, study biological systems have addressed many of the same problems that roboticists now face. Consider the following examples: sensor performance, sensor fusion, selection of primitives for scene segmentation and object recognition, object representations, active exploration vs. passive perception, motor control and planning for reaching, grasping and manipulating objects. Both groups need to address hardware considerations to understand how such constraints affect the way information is processed and represented, and how this in turn affects system performance.

[1]Section 2 is based on material from Lederman & Pawluk (1992).

2.2 Living organisms are functioning, multi-level integrated systems

It is important to recognize that biological organisms are complete, multi-level, integrated systems that actually work, despite the complexity of the many problems they must handle (see examples above). As such, living systems clearly demonstrate the complexity of the task facing the roboticist. Both biological scientists and roboticists have found it simpler to treat the different sensory modalities as independently operating units; however, the most recent work with biological systems clearly demonstrates the ultimate fallacy and limits of this approach.

2.3 Designing human-machine interfaces for teleoperation

Initial predictions about the relatively rapid creation of highly flexible, sensor-based autonomous robots have proved overly optimistic. As a result, attention has turned to the design of 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 central nervous system, it is no longer necessary to build an artificial one - as obvious by now, no mean feat! Those of us who study human systems, however, are quick to point out that with this approach, it becomes critical for the roboticist to learn about how our own human sensory systems process information, and about the constraints under which these operate. Such capabilities and limitations must be understood to achieve an effective interface with any teleoperated system. Since we are unable to present all information from the remote workspace, what information should be presented to the human operator? and what are the most effective ways to display it? Until now, such considerations have been ignored or noted too late for appropriate modification.

2.4 What the scientific method can contribute to robotics

In robotics and engineering, it is most common to model a system analytically using differential equations, where performed for the texture-matching instructions; Pressure (applied normal forces or torques about an object axis) was usually selected for hardness: Static Contact (simple static 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 dynamic hefting) was used for weight matching; Enclosure (finger molding to the object envelope) was selected most often to extract both volumetric and global (coarse) shape; Contour Following (edge following) was used most in conjunction with both global shape and exact (fine) shape.

3.2 The macrostructure of human haptic object identification

There are a number of computational models in the field of cognitive science that 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 to the computational model of reading proposed by Just and Carpenter (1980; 1987) and outlined in Figure 2.

READING: FROM EYE FIXATIONS TO COMPREHENSION	HAPTICS: FROM OBJECT EXPLORATION TO IDENTIFICATION
INPUT TEXT	INPUT OBJECT
Move Eyes to Next Word	Move Hand to Next Object Region
	Execute Exploratory Procedure
Local Interpretation of Word: Access word meaning (parallel lookup) Assign sentence function to word	Local Interpretation of Region Compute Object Property (parallel over accessible properties)
Build Representation of Text Meaning	Build Representation of Object (material and geometric properties)
Go back to start (view a new word) or exit ... -to alter previous knowledge representation -to act	Go back to start (move to a region) or exit ... -to discriminate -to classify -to name -to use a tool

Figure 2. Analogous stages between haptic object recognition and the reading model of Just and Carpenter (1980; 1987)

They separate text comprehension into a number of stages: moving the eye to the next location, fixation, local (e.g., lexical) analysis, and creating a global representation. Within a fixation, these stages are performed as completely as possible; over successive fixations, they subsequently recur. The representation is continually updated with the inputs from each new fixation.

In our own model, a period of manual exploration corresponds to a period of eye fixation. During the manual period, what we have called the "selection-extraction loop" takes place. An EP (or EPs) is selected and performed at some area on the object, on the basis of activation flow. 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 of the local and global processing is performed within an exploratory period as " possible; both recur during subsequent exploratory periods and object regions. Eventually, the system recognizes an object or selects the next EP for execution.

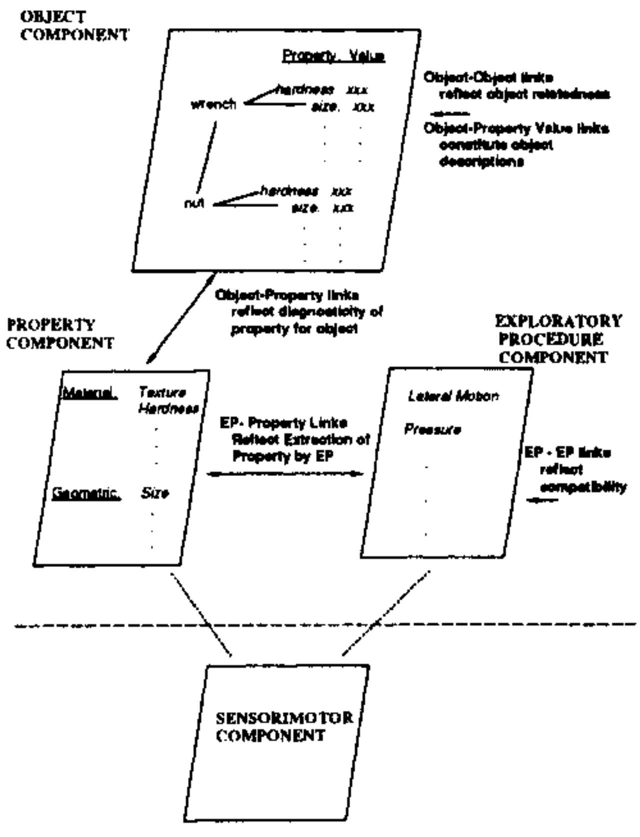

Figure 3. Model of the macrostructure of haptic object identification (from KJatzky & Lederman, 1991 and in press).

Figure 3 presents our model of the macrostructure of the haptic object-identification system, including the different data representations and the links we presume exist among them. An object component represents particular objects (e.g., wrench) and their specific property values. A property component represents the attributes along which an object may potentially vary (texture, hardness, etc.), rather than specific property values. More recently, we have divided properties into two major groups: "material" and "geometric". A material property is defined as having a factor affecting the response of a given material to imposed stimuli and constraints, independent of the shape and size of a given sample, for example, texture (Rosenthal & Asimow, 1971). A geometric (sometimes called technological) property is one that relates to the geometry of a particular material sample (e.g., shape, size). "Hybrid" properties directly reflect both geometry and material (e.g., mass). An EP component represents an exploratory procedure. The underlying neural mechanisms control EP execution and process property information arising from inputs to the sensory receptors. The latter sensorimotor component is not something we have modeled in our work; however, there is much relevant research that addresses both the neurophysiology and psychophysics of the somatosensory system.

Links between object and property components reflect the relative strengths of a given property for a particular object (e.g., texture is important for sandpaper), while links between object components are primarily intended to reflect the hierarchical classification relations documented in the cognitive literature (e.g., Rosch, 1978); for example, wrenches and nuts are linked because of their common tool-related function.

Links between EP and property components represent the precision of information about a property extracted by a particular EP. We have empirical data (Lederman & Klatzky, 1987; Expt. 2) that address this issue. These were obtained from a variant of the match-to-sample experiment discussed above. On any trial, subjects were now constrained to perform a single designated EP in conjunction with a named property, with all possible EP/property combinations performed over the experiment. Both accuracy and response times were measured. With these data, we were able to compare the relative precision with which each EP could extract a designated property. The results are shown in Table Ia in the form of an EP-property weight matrix. 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. An entry of "1" indicated sufficient, but not optimal performance. A "2" indicated that performance with the given EP was optimal and sufficient, although it was not necessary. A "3" indicated that the given EP was necessary as well as optimal. Note that, in general, those EPs that were executed spontaneously in the initial match-to-sample experiment tend to produce optimal performance in the constrained version of the experiment.

Table Ia: EP-to-property weights (from Klatzky & Lederman, in press; adapted from Lederman & Klatzky, 1990a).

	tex	hard	temp	wt	vol	global shape	exact shape
LM	2	1	1	0	0	0	0
PR	1	2	1	0	0	0	0
SC	1	0	2	0	1	1	0
UH	0	1	1	2	1	1	0
EN	1	1	1	1	2	2	0
CF	1	1	1	1	1	1	3

Table Ib: Breadth of sufficiency and average duration* (s) for each EP

	Breadth of Sufficiency	Duration (s)
Lateral Motion (LM)	3	3.46
Pressure (PR)	3	2.24
Static Contact (SC)	4	0.06
Unsupported Holding (UH)	5	2.12
Enclosure (EN)	6	1.81
Contour Following *(CF)*	7	11.20

Durations from Lederman & Klatzky (1987; Expt. 1)

Table Ib provides us with additional important information. By summing the number of non-zero cells across a row in Table Ia, we can represent the relative breadth of sufficiency of each EP; thus, Lateral Motion and Pressure each provide sufficient information about several different properties, whereas Enclosure and Contour Following provide coarse information about most object properties considered in this study. However, the breadth of property information provided by Contour Following must be weighed against its relatively slow execution time, also shown in Table Ib.

Links between EP components represent the extent to which the EPs may be co-executed, each extracting the property(ies) for which an EP is optimal/sufficient. For example, Lateral Motion and Pressure may be executed simultaneously, thus providing information about texture and hardness. We have developed a set of visible kinematic and dynamic parameters that formally differentiate the EPs; these parameters were derived from an extensive body of hand-movement data, based on a large number of common and custom-designed multidimensional objects tested over a wide range of experimental conditions. Values of four parameters were observed to occur reliably for a given EP across the different circumstances just described. Each parameter may be described as capturing some constraint inherent in an EP when it must be performed to extract certain types of information. The parameters and their stereotypic values are shown in Table 2a: these are Movement (static or dynamic), Direction (force applied normal or tangential to the surface), Region (of the object contacted by the end effector, i.e., surface, edge, or both), and Workspace constraint (supporting surface required or not). It is assumed 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 manner of exploration. The information in Table 2a allows us to determine whether any two EPs are compatible or not. Clearly, two EPs having identical parameter values would be compatible; however, they could not be differentiated. Still, it is possible to achieve compatibility by selecting some form of exploration that satisfies the constraints inherent in both EPs. For example, if one EP must be executed along the edges, whereas another must be applied to both edges and interior surfaces, in satisfying the second less restrictive constraint the first more restrictive constraint is simultaneously satisfied. Hence, the two EPs can be considered to be compatible.

Table 2a: Values of EPs on four parameters (from Klatzky and Lederman, 1991 and in press)

	Movem'nt	Direct'n	Region	Workspace Constraint?
SC	Static	Normal	Surface	No
PR	Dynam.	Normal	Surface	No
LM	Dynam.	Tang.	Surface	No
EN	Static	Normal	Surf & Edges	No
CF	Dynam.	Tang.	Edges	No
UH	Static	Normal	Surf & Edges	Yes

Table 2b: Compatibility relations between EPs (+ means compatible; - means incompatible), (from Klatzky and Lederman, 1991 and in press)

	PR	LM	EN	CF	UN
SC	+		+		+
PR		+	+	-	+
LM			-	+	
EN				-	+
CF					

We have expressed such compatibilities and incompatibilities in the form of an EP-EP weight matrix (Table 2b). A " + " represents compatibility between two EPs (e.g., Lateral Motion and Pressure); a "-" represents an incompatibility (e.g., Static Contact and Lateral Motion, because it is not possible to resolve the mismatch between two Movement parameter values).

3.3 Selecting an EP and the selection-extraction loop

In keeping with the interactive activation perspective (e.g., McClelland & Rumelhart, 1981), we treat haptic object identification as a parallel interactive process, with sequential constraints imposed by EP execution. Figure 4 shows how the process proceeds in a sequence of selection-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, the object is said to be recognized.

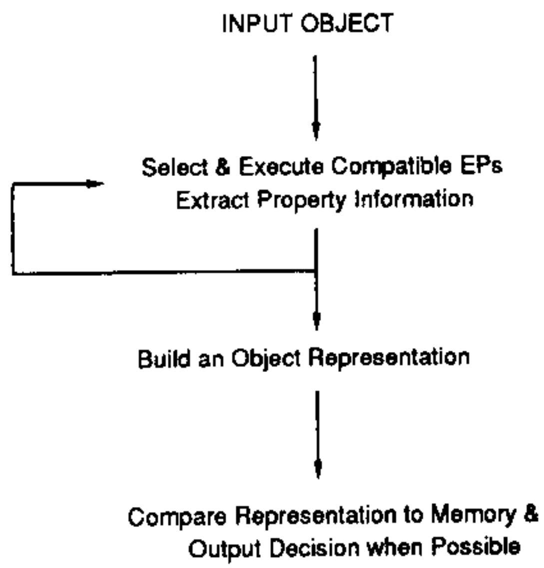

Figure 4. The selection-extraction loop (from Klatzky & Lcdcrman, 1991 and in press).

3.3.1 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 may 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 effect fine shape discrimination. This may generally discourage the use of Contour Following. In contrast, Enclosure might be generally favoured inasmuch as it is relatively fast and provides coarse information about many different properties (broad EP sufficiency).

In principle, these constraints and biases can be represented by the weights between different components in our system. For example, associations between specific objects can be represented by associative weightings within the object component, while expectations concerning the diagnostic value of various properties can be represented by connections between the property and object components. Hand movement precision and breadth of sufficiency are represented by connections between property and EP components (Table lb); EP compatibility (Table 2b) is determined by constraints inherent in the EP parameters (Table 2a). Finally, intrinsic biases, such as duration of execution, may be represented as item-specific bias terms.

3.3.2 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 progressively relaxed until some elements are maximally activated.

A system with symmetric weights and asynchronous updating minimizes a cost function over the set of constraints (weights), eventually selecting an optimum state or activation pattern over the associated units (Hopfield, 1982). In our case, constraint satisfaction serves as a method for selecting the next EP in a sequence during manual exploration. The weights are theoretically and/or 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 the compatibilities between EPs, we implemented the weights in Tables la and 2b as a constraint satisfaction system. The nodes represented the EPs and properties. This is equivalent to examining a single generic object.

Table 3a shows the activation level of each EP (as well as the time to relax and goodness level at the point of relaxation) after constraint satisfaction. Each property was clamped to represent an externally set property goal. In this first output table, the complete weight matrix of Table la was used, thus including positive weights on all EPs that were at least sufficient for extracting a given property. This might represent a condition in which the observer wishes to extract as much information as possible initially. Note that no matter which EP is clamped, the maximum activation level always occurs for an Enclosure, which is not only broadly sufficient but also compatible with other EPs. The next most active element is Unsupported Holding, which is compatible with Enclosure. Thus, this pair of EPs could be selected and executed within the same loop.

Table 3. Activation level of each EP after constraint satisfaction when each property has been activated externally. Results arc given for two sets of weights. Also shown are the time to relax (in multiples of 52 updates) and the goodness level at the point of relaxation, (from Klatzky & Lederman, 1991 and in press)

Activated Property	Activated EP: LM	PR	CF	EN	UH	SC	Time to Relax	Goodness Level
A. FULL WEIGHT MATRIX								
Texture	.45	.58	.48	.63*	.56	.56	8	9.1
Hardness	.40	.60	.47	.63*	.59	.54	7	9.1
Exact Sh.	.36	.54	.58	.60*	.55	.53	9	8.6
Global Sh.	.33	.55	.46	.66*	.59	.57	9	9.0
Size	.33	.55	.46	.66*	.59	.57	9	9.0
Weight	.33	.55	.47	.64*	.62	.54	9	8.8
Temp.	.39	.58	.46	.63*	.59	.59	8	9.2
B. OPTIMAL EP-PROPERTY WEIGHTS ONLY								
Texture	.53*	-.34	.50	-.50	-.46	-.46	12	4.4
Hardness	-.36	.52*	-.51	.50	.46	.46	12	4.4
Exact Sh.	.39	-.38	.62*	-.50	-.46	-.46	11	4.7
Global Sh.	-.39	.37	-.50	.59*	.45	.45	11	4.5
Size	-.39	.37	-.50	.59*	.45	.45	11	4.5
Weight	-.40	.38	-.50	.50	.56*	.46	11	4.4
Temp.	-.40	.37	-.50	.50	.46	.56*	11	4.4

*highest activation for the given property

Table 3b models a different situation, which might occur following the initial stage of coarse exploration -- if more precise information is required than can be extracted by an EP that is merely sufficient for that property, the optimal EP would be required. To model this we used a second set of weights in which only optimal EP-property weights in the matrix were included. Note the very different activation levels - here we see that clamping a particular property results in the optimal EP being selected.

3.4 Behavioural support for a constraint satisfaction approach

3.4.1 Human experiments on the 2-stage EP sequence

The results of our modeling are supported by evidence of a 2-stage exploratory sequence, which was adopted by subjects during an object classification task (Lederman & Klatzky, 1990b). On each trial, subjects were initially asked a yes/no question of the type: "Is this X further a Y?" (e.g., "Is this wood-working implement further a piece of sandpaper?"). An object was then placed in the subjects hands for exploration; on half the trials, the object was in fact an exemplar of both X and Y classes named in the question, while on the remaining trials, an object from the same X but different Y class was presented (e.g. file). The most diagnostic property for each object class named in the questions had been determined previously in a separate experiment (Lederman & Klatzky, 1990b).

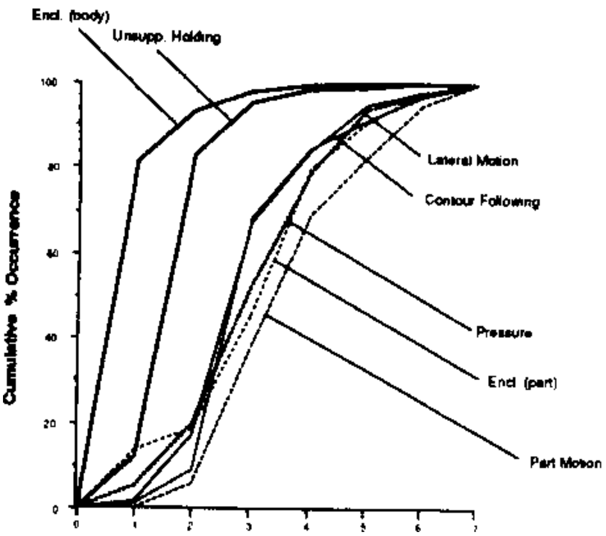

Figure 5. Cumulative percentage of occurrence of each EP as a function of position in the exploratory sequence (from Lederman & Klatzky, 1990b). The solid lines indicate the grasp/lift combination. Static Contact is not included because it occurred very infrequently.

The hand movements from each trial were analyzed as a sequence of EPs. This analysis indicated two separate stages of manual exploration, as evident in Figure 5. Each function depicts the cumulative percentage of EP occurrence as a function of serial position in the EP sequence. The two solid dark lines indicate that the first two EPs in the sequence were an Enclosure followed by an Unsupported Holding, which together comprise what we refer to as a grasp/lift routine. Both of these are relatively broadly sufficient and would presumably provide considerable coarse information about any object. The remaining EP functions occurred after this initial exploratory sequence, the particular EP being predicted by the property that was known to be most important for object identification. For example, since texture was most diagnostic for deciding whether or not the wood-working tool in the subject's hands was a piece of sandpaper, we predicted that Lateral Motion would be selected following execution of a grasp/lift routine, since it is optimal for extracting texture. Such predictions were confirmed. The second stage of exploration was more specifically directed toward extracting further precise information about the critical property. In short, these data on manual exploration during object identification support our approach to exploratory control as a constraint satisfaction process.

3.4.2 Experiments on the selection-extraction loop and property extraction

In addition to determining how EPs are selected, our model also addresses how EP selection affects the precision with which information about an object can be extracted. We assume that the strength of relations between EPs and properties (Table Ia) should predict the extent to which an object's properties may be perceived and learned using a particular form of manual exploration. This assumption has several implications. Learning about a property should be fastest when exploration involves an optimal EP, because less sampling is necessary to obtain a given amount of information. When only a single optimal EP is executed, incidental information about other properties will be restricted to those for which that EP is sufficient. Selecting compatible EPs allows for co-execution, thus making available information about all properties for which either is sufficient. Presumably, selecting one EP is assumed to prevent the co-occurrence of any other incompatible EPs also selected; rather, these must be performed in sequence. If the eliminated EPs are necessary for a given property, then no learning about that property will occur. Clearly, these predictions highlight the "gatekeeper" role that EPs play during object perception.

We have investigated the role of EPs in limiting accessibility to object properties in a series of experiments (Klatzky et al., 1989; Reed et al, 1990; Lederman et al, in press). The tasks generally required subjects to learn to classify sets of multidimensional objects into groups according to different classification rules. For these tasks, we designed sets of objects that varied factorially in a number of properties, such as texture, hardness, shape and size. At least one of these properties was used to divide objects into categories.

Take an example of a 1-property rule:, the roughest objects were all in category "A", the intermediate roughness objects all in category "B", and the smoothest all in category "C". Over a set of trials, subjects were required to indicate to which category the objects in a set belonged. Response time for object classification was measured as time from contact to verbal response.

Across one experiment, the categorization tasks demanded that subjects extract information about one or more diagnostic properties of each presented object. We expected response time to decrease as the sets were repeatedly presented, because subjects would learn to eliminate those EPs that were not optimal for learning about the diagnostic property. And indeed, we note in Figure 6 that response time decreased to an asymptotic value for three different sets of objects regardless of the number of diagnostic object properties used to classify objects in a given set (1, 2, or 3).

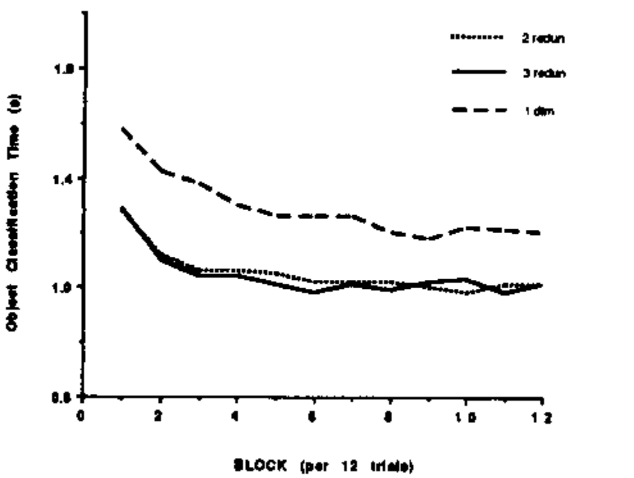

Figure 6. Mean response time for classifying each object based on 1, 2-redundant, and 3-redundant property classification rules (adapted from Klatzky et al, 1989). The 1- and 2-redundant functions were both produced by averaging data from several different property-classification conditions.

One important reason for the reduction in response time becomes evident when we examine the corresponding EPs performed over successive time periods in the 1-property case (Figure 7). We observe that subjects chose to streamline their manual exploratory activity over time. Thus, EPs that were optimal for extracting the diagnostic property (texture, shape, or hardness) continued to occur frequently and across blocks; in contrast, the other EPs scored in the study occurred less often initially and subsequently declined. Similar results were obtained for the 2- and 3-property categorization conditions discussed next.

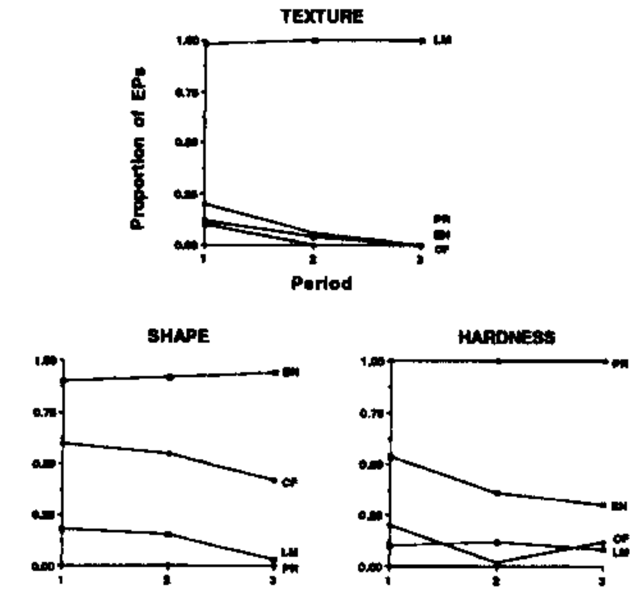

Figure 7. Proportion of EP occurrence (Lateral Motion, Pressure, Contour Following, and Enclosure) over sequential periods for three different 1-property classification conditions (adapted from Klatzky et al. 1989). LM= Lateral Motion; PR= Pressure; EN= Enclosure; CF=Contour Following.

We also predicted that when more than one property redundantly defines the categories (e.g., all "A"s are both very rough AND very hard, all "B"s are of intermediate roughness AND intermediate hardness, and all "C"s are both very smooth AND very soft), categorization times should be faster than with a 1-property classification rule. This is evident in Figure 6 by the fact that the 2- and 3-property curves lie below the 1-property curve, indicating a "redundancy gain". We argue that this reflects the savings resulting from the fact that the EPs that were optimal for extracting two properties were compatible (e.g., Lateral Motion and Pressure for redundant texture and hardness). That there is no additional savings from adding a third redundant dimension (shape) to redundant texture and hardness specifically reflects incompatibility of Contour Following with either Lateral Motion or Pressure.

We further predicted that when a single EP is used to explore an object, incidental knowledge about other properties will accrue depending on the weight between each property and that EP. To assess this, we turn now to another experimental approach, which involves what we have called the "withdrawal" paradigm. We presented subjects with sets of objects that were redundantly defined with different 2-redundant property classification rules (texture/hardness; texture/shape; hardness/shape). Subjects were initially told to classify the objects on the basis of a single named dimension (e.g., texture), even though objects varied redundantly on two dimensions (texture and hardness); when performance asymptoted, the object set was switched to a one-dimensional rule (i.e., objects only varied in texture) since values of the second property (hardness) were now held constant. That is, variation on the second property was "withdrawn". We reasoned that if subjects had previously incidentally learned about the second property, then their response times should increase just after it was withdrawn. Figure 7 presents the results. Note a strong withdrawal effect for texture/hardness redundancies, regardless of which dimension was withdrawn. In contrast, withdrawal effects for the other two combinations were very small, and typically not statistically significant. This result was also expected as Contour Following (for shape) is incompatible with both Lateral Motion (texture) and Pressure (hardness) for the set of planar stimulus objects used: shape information was only available at the edges, while texture and hardness were both found in the interior surface areas. In contrast, when texture and shape are available in the same local region (created by using fully 3-dimensional ellipsoids of revolution), we obtained strong withdrawal effects. We predicted this outcome because information about both shape and texture was simultaneously available from the now fully compatible EPs, Contour Following and Lateral Motion.

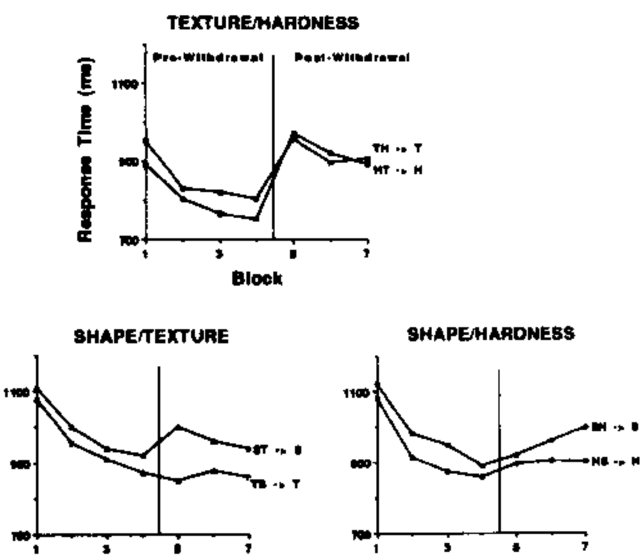

Figure 8. Classification response time as a function of period for three 2-redundant property classification rules. Separate effects of targeting one property while withdrawing the second are shown in each panel. X/Y -> X should be read as properties X and Y are both initially presented together; Property Y is subsequently withdrawn, (from Klatzky et al, 1989)

In the previous set of experiments, a particular property was targeted. Another series of experiments focused on EP selection under conditions in which no property was targeted. We predicted that knowledge about objects would be determined on the basis of which EPs were spontaneously selected. The prediction was tested experimentally by having subjects sort, according to perceived object similarity, the complete set of multidimensional planar shapes used in the immediately preceding series. We found (Klatzky et al, 1987; Summers et al, submitted) that when subjects could use only haptic exploration, they preferred to sort by hardness and texture, their selection apparently reflecting the cost (speed, accuracy) of executing a Contour Following to extract shape information. However, note what occurred when the instructions stressed attending to visual images or to visual cues when they were also provided. Such instructions presumably created a bias toward shape; as would be expected, Contour Following and Enclosure were selected most frequently.

4 Application to robotic exploration[3]

In keeping with Gibson's earlier observations (e.g. 1966), much research with humans and other living organisms has highlighted the importance of active exploration in perceptual activities. Our own work again confirms this general principle with respect to human haptic object processing. In the robotic domain, Bajcsy (1989) may be credited with emphasizing the need for exploration, particularly when information about object properties must be used to interact with unstructured environments. Presumably, such is true whether or not identification is requisite.

Bajcsy and a number of others have recently linked the biological and robotic exploration domains by specifically adoping the concept of an EP as a systematic testing procedure, and by implementing robotic versions of the human EPs described above (e.g., Allen & Michelman, 1990; Bajcsy & Campos, 1992; Sinha, 1992; Stansfield, 1988). However, neither the selection of particular EPs nor the sequence in which these should be performed is intuitively obvious. We advocate adopting the experimental paradigm used by Lederman & Klatzky (1987, Expt. 2; see Section 3.2 of this paper) to develop a more general robotic search solution for EP selection when multidimensionally varying objects are explored. This experimental approach allows the systematic determination of the relative performance characteristics of the set of robotic EPs selected, which could then be used in conjunction with a constraint satisfaction approach to select efficient EP sequences.

[3]Section 4 is based on material from Lederman et al (1992).

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 may also be sufficient for extracting several others. The properties and associated EPs will depend upon the particular robotic end effector and sensing system selected, and could be quite different from what humans use (unless an anthropomorphic design has been deliberately 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 listed in Table la may be applied to any exploring system, along with any others that are specifically relevant to the robotic domain. The experiment described in Section 3.2 can be used as a methodological guide for systematically evaluating the relative performance of each EP in extracting each property; note that the tasks should involve different levels of property discrimination if EP performance is to be fully and properly assessed. Also relevant to this approach is the extent to which robotic EPs may be co-executed, that is, the issue of EP compatibility. To determine the compatibilities and incompatibilities between EPs, the EP descriptions must be specifically defined in terms of robotically relevant constraint parameters. Collectively, the results concerning the relative strengths of the EP-to-property and inter-EP-compatibility associations can be used to rank robotic EPs for use in associated computational models of EP selection.

5 Summary

In closing, I would like to repeat my 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, anthropomorphic design. On a general level, I have argued (Lederman & Pawluk, 1992; Lederman et al, 1992) that such work: a) addresses many of the same problem domains, b) provides an example of, and framework for, designing working, multilevel integrated systems, and c) offers valuable suggestions for presenting robotically extracted information to a human operator in teleoperation. Further, d) the scientific method highlights the value of properly constraining the problem, formulating testable hypotheses, designing rigorous and unbiased experimental tests of the hypotheses, and using statistical techniques for assessing the validity, reliability and generality of the experimental findings. On a more specific level, based on our familiarity with the scientific results of experiments on biological touch, my colleagues and I have further proposed a number of particular suggestions for designing robotic tactile/haptic systems, including the example discussed in Section 4. For biological scientists, who attempt to understand the bases of natural intelligence, and roboticists, who attempt to create such behaviour in machines, serious collaboration may provide a new and potentially valuable approach to robotic design.

Acknowledgments

The research program on human haptics is based on a long-standing collaboration with Dr. Roberta Klatzky (University of California at Santa Barbara). I also wish to thank Dr. Ruzena Bajcsy (Pennsylvania) for numerous fruitful discussions about related biological and robotic issues. My thanks to the other valuable collaborators with whom some of the work reported here was performed: Dr. Cathy Reed (Pennsylvania), Dr. Craig Summers (Laurentian), and Dianne Pawluk (Harvard). Cheryl Wilson helped format this manuscript for publication. The paper was prepared with the financial assistance of the Institute for Robotics and Intelligent Systems, a Canadian federal centre of excellence, and the Information Technology Research Centre of Ontario centre of Excellence.

References

Allen, P.K. & Michelman, P. (1990). Acquisition and interpretation of 3-D sensor data from touch. IEEE Transactions on Robotics and Automation, 6(4), 397-404.

Bajcsy, R. (1989). Active perception and exploratory robotics. In I. Plander (Ed.) Artificial intelligence and information-control systems of robots. New York: North Holland.

Bajcsy,R. & Campos, M. (1992). Active and exploratory perception. CVG1P: Image Understanding, 56,31-40.

Gibson, J.J. (1966). The senses considered as perceptual systems. Boston: Houghton-Mifflin.

Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. Proceedings of the National Academy of Sciences. 79, 2554-2558.

Just, M.A. & Carpenter, P.A. (1980). A theory of reading: From eye fixations to comprehension. Psychological Review, 87, 329-354.

Just, M.A. & Carpenter, P.A. (1987). The psychology of reading and language comprehension. Newton, Mass: Allyn & Bacon.

Klatzky, R.L., & Lederman, S.J. (1991). Toward a computational model of constraint-driven exploration and haptic object identification (Tech. Rep. RPL-TR-9104). Kingston, Ontario, Canada: Queen's University; to appear in Perception (in press).

Klatzky, R.L., Lederman, S.J., & Metzger, V. (1985). Identifying objects by touch: An "expert system". Perception & Psvchophvsics. 37(4). 299-302.

Klatzky, R., Lederman, & Reed, C. (1989). Haptic integration of object properties: Texture, hardness, and planar contour. Journal of Experimental Psychology: Human Perception & Performance. 15 (1), 45-57.

Lederman, S.J. & Klatzky, R.L. (1987). Hand movements: A window into haptic object recognition. Cognitive Psychology, 19(3), 342-368.

Lederman, S.J. & Klatzky, R.L. (1990a). Flexible exploration by human and robotic haptic systems. Proc. Ann. Intern. Conf. IEEE/Engineering in Medicine & Biology Society, 12(5), 1915-1916.

Lederman, S. J. & Klatzky, R.L. (1990b). Haptic classification of common objects: Knowledge-driven exploration. Cognitive Psychology, 22, 421-459.

Lederman,S.J. & Klatzky, R.L. (in press). Extracting object properties by haptic exploration. Acta Psychologica.

Lederman, S.J., Klatzky, R.L., & Pawluk, D. (1992). Lessons from biological touch for robotic haptic sensing. In H. R. Nicholls (Ed.), Advanced tactile sensing for robots. (World Scientific Series in Robotics and Automated Systems, Vol. 5). Singapore: World Scientific Publishing.

Lederman, S.J., Klatzky, R.L, & Reed, C. (in press). Constraints on haptic integration of spatially shared object dimensions. Perception.

Lederman, S.J., & Pawluk, D. (1992). Lessons from biological touch for robotic tactile sensing. In H. R. Nicholls (Ed.), Advanced tactile sensing for robots. (World Scientific Series in Robotics and Automated Systems, Vol. 5). Singapore: World Scientific Publishing.

McClelland, J.L. & & Rumclhart, D.E. (1986). A distributed model of human learning and memory, in J. L. McClelland & D.E. Rumelhart (Eds.), Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2, Cambridge, MA: MIT Press, 170-215.

Reed, C, Lederman, S.J. & Klatzky, R.L. (1990). Haptic integration of planar size with hardness, texture, and planar contour. Canadian Journal of Psychology. 44(4), 522-545.

Rosch,E. (1978). Principles of categorization. In Rosch, E. & Lloyd, B. (Eds.) Cognition and categorization. Hillsdale,NJ: Erlbaum, 27-48.

Stansfield, S. (1988). A robotic perceptual system utilizing passive vision and active touch. International Journal of Robotics Research. 138-161.

Sinha, P.R. (1992). Robotic exploration of surfaces and its application to legged locomotion. Proc. of IEEE International Conf. on Robotics & Automation. 1. 221-226.

Walk, R.D. & Pick, H.L.,Jr. (1981). Intersensoy perception and sensory integration. New York: Plenum.