Mapping the mind

Domain specificity in cognition and culture

Edited by

LAWRENCE A. HIRSCHFELD
University of Michigan

SUSAN A. GELMAN
University of Michigan

CAMBRIDGE
UNIVERSITY PRESS

Contents

viii   *Contents*

1 Toward a topography of mind: An introduction to domain specificity

Lawrence A. Hirschfeld and Susan A. Gelman

Over the past decades, a major challenge to a widely accepted view of the human mind has developed across several disciplines. According to a long predominant view, human beings are endowed with a general set of reasoning abilities that they bring to bear on any cognitive task, whatever its specific content. Thus, many have argued, a common set of processes apply to all thought, whether it involves solving mathematical problems, learning natural languages, calculating the meaning of kinship terms, or categorizing disease concepts. In contrast to this view, a growing number of researchers have concluded that many cognitive abilities are specialized to handle specific types of information. In short, much of human cognition is domain-specific.

The notion of domain specificity is not new. Indeed, intriguing (although brief) hints of domain specificity emerge in the epistemologies of Descartes and Kant and in the psychologies of Thorndike, Vygotsky, and de Groot. For example, in *Mind in Society*, Vygotsky argues that

the mind is not a complex network of *general* capabilities such as observation, attention, memory, judgment, and so forth, but a set of specific capabilities, each of which is, to some extent, independent of others and is developed independently. Learning is more than the acquisition of the ability to think; it is the acquisition of many specialized abilities for thinking about a variety of things. Learning does not alter our overall ability to focus attention but rather develops various abilities to focus attention on a variety of things. (1978: 83)

Still, in recent years, increased and detailed attention has turned toward the question of domain specificity. Psychologists with concerns ranging from animal learning to emergent theories of mind and body, cognitivists exploring problem solving and expertise, anthropologists working with color terms and

3

folk taxonomies, psycholinguists investigating auditory perception, and philosophers and others examining reasoning schemata have concluded – often independently – that humans simply could not come to know what they do know in a purely domain-neutral fashion. A major purpose of *Mapping the Mind: Domain Specificity in Cognition and Culture* is to convey the wealth of current research that has resulted from this multidisciplinary exploration.

This introduction will orient readers to a domain specificity perspective. It is divided into three sections. In the first section, we examine work antecedent to the domain perspective drawn from a number of fields. By doing so we hope to give a broad sense of the intellectual traditions from which domain specificity has emerged. Throughout we will highlight the conclusion that the mind is less an all-purpose problem solver than a collection of enduring and independent subsystems designed to perform circumscribed tasks. This common conviction aside, it is important to keep in mind that a domain perspective is not the achievement of a coordinated body of research, unified in a common challenge. Domain researchers have reached some shared conclusions while asking quite diverse questions. In the second section we draw from this multidisciplinary work a common notion of what a domain is. It is important to stress that our intention in this section is to characterize rather than define what a domain is. Finally, in the third section we consider questions that arise by looking at ways in which domain researchers differ in their approaches and conclusions.

It is essential to note, given the diversity of interests and backgrounds of researchers in domain specificity, that conclusions about the nature and scope of the domain specificity approach are not reducible to differences in the traditions from which researchers have engaged the question. Rather, both the major lines of contention *and* commonality evident in these chapters are largely independent of academic discipline or research methodology. We believe that this is one of the most encouraging aspects of domain research, one that provides broad and exciting possibilities for future research directions.

In the introduction we want to raise a number of ideas about domains and issues about their natures. We hope in doing this to motivate the questions that the volume's chapters address: For example, does all domain knowledge reflect the operation of innate devices, or under what conditions might domain-specific knowledge be transferred, or whether initial conceptual organization evolves, is elaborated, or is supplanted in development. In effect, we see our task in this chapter as rendering such questions sensible across disciplines and traditions.

The roots of domain specificity

In this section we review the intellectual antecedents of the contemporary domain perspective. Our goal is twofold. First, we want to indicate the

research and theory that have been crucial to the evolution of a domain approach. Although the authors of some of this work may well not be advocates of a domain-specific perspective, their work has nonetheless been critically important to the development of the approach. Second, we review this work with an eye toward building a characterization, if not a definition, of what a domain is and what a domain is not.

Several traditions have converged on a domain perspective. All attempt to solve the central problem of domain specificity, namely, how do humans come to have the wealth of knowledge that they do? These traditions have their roots in the following: (1) Chomsky's theory of natural language grammar; (2) modular approaches to knowledge (particularly vision and auditory speech processing); (3) constraints on induction; (4) philosophical insights into the most intricate knowledge structures created by humans (theories); (5) the learning, memory, and problem solving of our best learners (experts); (6) and the wisdom gained from a comparative perspective (animal, evolutionary, and cross-cultural studies).

Chomsky's theory of natural language grammar

We start with Chomsky's theory of language for two reasons. First, it has special historical interest: Virtually all subsequent domain-specific accounts bear the imprint of Chomsky's arguments about cognitive architecture. Although previous researchers recognized the need for conceiving thought in terms of discrete mental functions, Chomsky elaborated the first modern, sustained, and general account of domain specificity. It would be hard to overestimate the importance that his views have had in forming a broad-ranging domain-specific perspective. Although none of this volume's contributions directly treats natural language grammar, all grapple with issues raised in Chomsky's work.

The second reason for beginning with this theory is the clarity of its claims. Perhaps because it remains controversial, the notion that the language faculty represents a unique mental organ is probably the most widely known domain-specific argument. This attention is well deserved: The study of natural language processing is the arena in which the domain challenge has most continuously and explicitly unfolded. Although not all scholars are convinced that syntax *must* be described in domain-specific terms, the research from which this claim is derived provides an apt and excellent illustration of one domain perspective.

Current Chomskian linguistic theory distinguishes the principles of language structure at the core of the language faculty from language-specific rules derived from these principles. According to this model, (1) understanding a sentence involves assigning it a structural description in terms of abstract categories; (2) operations on sentences necessarily involve interpreting sentences in terms of this abstract phrase structure; (3) this abstract phrase

structure cannot be inferred from surface properties of utterances (such as the linear order of words in the sentence).

For example, consider how a grammatically well-formed question is derived from the following two sentences (the example is drawn from Chomsky, 1980a; see also 1988):

(1)   The man is here. – Is the man here?
       The man will leave. – Will the man leave?

Chomsky suggests that two hypotheses fit these data. The first hypothesis for forming an interrogative from a declarative sentence is the *structure independent hypothesis* (H_1). According to this hypothesis, the speaker processes the sentence from beginning to end, word by word. When the speaker reaches the first occurrence of a class of words, say a verb such as *is* or *will*, he or she transposes this word to the beginning of the sentence. The alternative, *structure dependent* hypothesis (H_2), is the same as the first "but select[s] the first occurrence of *is, will*, etc., following the *first noun phrase of the declarative*" (Chomsky 1980a, emphasis added).

The (first structure independent) hypothesis is less complex in that it relies on superficial features of sequential order rather than requiring speakers to interpret utterances with respect to components of their constituent phrase structure, that is, "the first noun phrase." If the mind prefers "simpler" solutions – that is, is guided by a sensitivity to mental economy – we would expect to find language organized by principles captured with the structure independent hypothesis rather than the more abstract and language-specific structure dependent hypothesis.

The issue is resolved, Chomsky argues, by looking at the different predictions the two hypotheses make for similar sentences and their associated questions. First, on the structure dependent hypothesis the following movements are predicted:

(2)   The man who is here is tall. – Is the man who is here tall?
       The man who is tall will leave. – Will the man who is tall leave?

In contrast, the structure independent hypothesis, in which movements are calculated over surface properties of the sentence (such as word order), predicts a pattern that is not only ungrammatical, but also never encountered:

(3)   Is the man who here is tall?
       Is the man who tall will leave?

The structure dependent claim, accordingly, more adequately captures the linguistic facts.

The crucial question, Chomsky observes, is how children come to know that structure dependence governs such operations but structure independence does not. It is not, he contends, that the language learner accepts the first hypothesis

and then is forced to reject it on the basis of data such as (2). No child is taught the relevant facts. Children make many errors in language learning, but none such as (3), prior to appropriate training or evidence. A person might go through much or all of his life without ever having been exposed to relevant evidence, but he will nevertheless unerringly employ H_2, never H_1, on the first relevant occasions. ... We cannot, it seems, explain the preference for H_2 on grounds of communicative efficiency or the like. Nor do there appear to be relevant analogies of other than the most superficial and uninformative sort in other cognitive domains. If humans were differently designed, they would acquire a grammar that incorporates H_1, and would be none the worse for that. (Chomsky, 1980a: 40)

Chomsky concludes that the mind is *modular* – "consisting of separate systems [i.e., the language faculty, visual system, facial recognition module, etc.] with their own properties" (Chomsky, 1988: 161). The modular claim has three components: First, the principles that determine the properties of the language faculty are unlike the principles that determine the properties of other domains of thought. Second, these principles reflect our unique biological endowment. Third, these peculiar properties of language cannot be attributed to the operation of a general learning mechanism. Linguistic principles such as structure dependence cannot be inferred from the general language environment alone. Yet children's language development is guided by these principles.

As we observed above, this claim is not uncontroversial. For example, a number of researchers have suggested that the young child's task of inferring the structural properties of language is made easier because adults simplify the language that learners are presented with (Snow, 1972; Furrow & Nelson, 1984, 1986). Cross-cultural work, however, indicates that such simplifications are not a universal feature of the language learning environment (Ochs & Schieffelin, 1984; Pye, 1986). Other studies find that properties of child-directed speech do not correlate with the ease of language learning (e.g., Gleitman, Newport, & Gleitman, 1984; Hoff-Ginsberg & Shatz, 1982). Nonetheless, language acquisition appears to be quite stable and regular across diverse cultural and linguistic environments (Slobin, 1985). The conclusion Chomsky and others have reached is that the child has an innate capacity to learn languages, thus filtering "the input data through an emerging system of rules of grammar" (Gleitman, 1986: 7).

Other evidence lends support to Chomsky's theory. For instance, language learning appears to be stable and regular across significant variation in language *learners* as well as language *learning environments*. Curtiss (1982) has shown that severe disturbances in cognitive capacity do not necessarily result in disrupted language capacity (see also Cromer, 1988). Language development continues to unfold in the typical, predictable sequence for learners who are blind (Landau & Gleitman, 1985) and so have very different sensory experience from sighted children, and for those who are deaf and acquiring language in a different sensory modality (see studies of sign language, such

as ASL; Klima & Bellugi, 1979; Newport & Meier, 1985; Petitto, 1988). Even deaf children who, in their first few years of life, have had little exposure to spoken language and no exposure to sign language, invent "words" and two- or three-word "sentences" (Goldin-Meadow, 1982). These results do not imply that the environment has no effect. For example, delaying exposure to language until later in life can have consequences ranging from moderate to severe (Newport, 1991; Curtiss, 1977). Nonetheless, it is striking that learners manage to construct language systems across a wide array of circumstances.

Modular approaches to cognition

As we observed, Chomsky and others maintain that these findings provide compelling evidence for the claim that the mind is modular, comprising a number of distinct (though interacting) systems (the language faculty, the visual system, a module for facial recognition), each of which is characterized by its own structural principles (1980b, 1988). Clearly this claim is related to the notion that thought is domain-specific, the idea that many cognitive abilities are specialized to handle specific types of information.

Chomsky, however, has also suggested that the mind is modular in a somewhat different way, giving rise to a set of proposals about cognitive architecture stressing the organization and contribution of each of the system's subcomponents rather than the system's overall characteristics. Thus, in other more technical writings, Chomsky has described "modules of grammar" (e.g., the lexicon, syntax, bounding theory, government theory, case theory, etc.) (1988: 135). Here the notion of modularity appears to be tied to specific subcomponents or subsystems of the language faculty rather than to the modular uniqueness of the language faculty itself. The grammar, in the traditional sense, is located at the intersection of these distinct modules.

It is not clear whether these two notions of modularity are to be distinguished, and if so how to interpret the relationship between them. One possibility is that modules are nested, that is, the language faculty is a separate module that in turn consists of distinct component operations or modules. Another interpretation – supported indirectly by the fact that Chomsky speaks of the language faculty *as* a module to nonlinguists but speaks of the language faculty *as consisting of* modules to linguists – is that the mind is, strictly speaking, modular with respect only to these second-level component modules. The language faculty itself would accordingly be a more vaguely defined construct resulting from the operation of these modules, but one that in itself is not modular in the sense of being defined in terms of a distinct set of principles.[1]

Modular accounts of other cognitive competencies more often resemble the second modular interpretation of Chomsky's position than the first. Thus, for example, although the visual and auditory systems are often compared

with the language faculty as contrasting modules (Chomsky, 1988, 1980b; Fodor, 1983), detailed accounts of these systems typically analyze their structure in terms of a set of component modular operations, each of which accounts for only part of the overall system's functional output. Thus, descriptions of such systems adhere to what Marr (1982) called the *principle of modular design*, "the idea that a large computation [such as vision] can be split up and implemented as a collection of parts that are as nearly independent of one another as the overall task allows" (p. 102).

Marr's own theory of vision is a clearly elaborated example of this sort of modular explanation. The theory's principal goal is to understand how it is that we see stable and identifiable images in spite of great variation and "noise" in the input. For example, although we perceive colors, shapes, and sizes as constant, the stimulus information available to the visual system is not sufficiently constrained to permit us to infer constancy without additional interpretation. Areas of unequal shading (which makes some areas of a single color appear darker than others), the possibility of object movement (which makes the same object appear smaller or larger depending on whether it is moving toward or away from the viewer), or partial occlusion (which obscures large parts of objects that are nonetheless perceived as a single whole) mean that visual information alone often underdetermines our perception of color, shape, and size constancy.

To explain such judgments, Marr puts forward a computational theory of vision that analyzes the perception of shape, size, and motion into representations constructed from a set of specific algorithms. These algorithms transform representations by means of *modular* devices that detect edges, apparent motion, surface texture, and the like. Vision, the process of *seeing*, involves the coordination of these atomic visual modules into a coherent whole.

Other modular devices seem to control auditory processing. A considerable body of research emerging from the Haskins Laboratory under Alvin Liberman provides a computational theory of auditory processing. Central to this work is the demonstration that the phonetic analysis of speech involves mechanisms different from those that affect the perceptual analysis of auditory nonspeech (Mattingly, Liberman, Syrdal, & Halwes, 1971; Liberman & Mattingly, 1989).

Drawing on this empirical work in vision and speech processing, in *Modularity of Mind*, Fodor (1983) offers the first general discussion of the implications of modularity for a wide set of domains. Fodor lists a number of candidate modules, including color perception, analysis of shape, analysis of three-dimensional spatial relations, recognition of faces, and recognition of voices.

Fodor's model involves a functional cognitive taxonomy that distinguishes between input systems or modules (which produce knowledge about the world, such as edge detectors) and transducers (which compile information *from* the

world, such as perceptual organs). Input systems, in turn, are distinguished from central processors that take information from the input systems, in a format appropriate for the central processors, and use this information to mediate higher functions, such as the fixation of belief.

Thus, according to Fodor's modular view, knowledge of the different aspects of the world is mentally represented in distinct formats. Perception accordingly involves not only interpretation, but an interpretation that is constrained by the format under which particular world knowledge is represented. In other words, input systems are not simply conduits for perceptual encodings of information; they are mental modules that "deliver representations that are most naturally interpreted as characterizing the arrangement of *things in the world*. Input analyzers are thus inference-performing systems" (1983: 42).

Modular views of cognition represent a major challenge to predominant, domain-general approaches found in psychology, linguistics, philosophy, and anthropology. As such, they have important implications for any domain-specific perspective. Yet modular and domain-specific approaches also contrast in significant ways. The principal difference is the former's emphasis on specificity in functional cognitive architecture and the latter's focus on specialization for specific types of knowledge. In the following three sections we consider the direct intellectual antecedents to domain specificity. The work we examine, rather than focusing on cognitive architecture (a modular issue), is concerned with the mental activities that operate on that architecture. We turn first to the issue of constraints on representations.

Constraints

An appeal to constraints begins with the problem of induction. As Rochel Gelman phrased it (1990: 3), "How is it that our young attend to inputs that will support the development of concepts they share with their elders?" She raises two significant difficulties with developing the appropriate concepts: Experience is inadequate in that many of the critical concepts children need to learn never appear; it is "pluripotential" in that it is logically open to many alternative construals. As R. Gelman points out, "the indeterminacy or inadequacy of experience and the pluripotentiality of experience ... are central to current discussions of the acquisition of syntax (Landau & Gleitman, 1985; Wexler & Culicover, 1980), visual perception (Marr, 1982; Ullman, 1980), the nature of concepts (Armstrong, Gleitman, & Gleitman, 1983; Medin & Wattenmaker, 1987), and the learning of word meaning (Macnamara, 1982; Quine, 1960)" (pp. 3–4).

The inadequacy and pluripotentiality of experience are implicit in many accounts of learning, including Quine's treatment of word meaning acquisition (1960) and Peirce's discussion of hypothesis generation in science (1960). For the child to learn word meanings without constraints is akin to an alien

trying to discover the laws of nature by examining the facts listed in the Census Report. Both would be doomed to positing thousands upon thousands of meaningless hypotheses. The child might wonder whether "rabbit" refers to a certain patch of color, or the positioning of a limb; the alien may wonder whether there is a meaningful causal relation between the number of babies born in Cancun and the height of women in Brazil whose names start with "Z." If left unconstrained, induction would yield meaningful knowledge only rarely (if at all), and even then only by chance.

One promising response to the induction puzzle is to suggest that there are constraints on the form development takes. Constraints are restrictions on the kinds of knowledge structures that the learner typically uses (Keil, 1981: 198). With constraints, the induction problem is simplified because the learner need not consider every possible reading of the input. For example, regarding the acquisition of word meanings, Markman (1989) suggests that children first assume that nouns refer to whole objects that are taxonomically related (the *taxonomic* and *whole object* constraints). These constraints would exclude from consideration meanings for *rabbit* such as "white fur" or "things that hop [including pogo sticks and wallabees]." Keil (1981) also proposes domain-specific constraints on number concepts, deductive reasoning, ontological knowledge, and natural language syntax.

All theorists acknowledge the need for constraints of some sort. Even traditional learning theorists propose constraints on learning (e.g., perceptual constraints; contiguity). Disagreement remains as to the importance of constraints, how much focus they deserve, and on how best to characterize their nature (see Behrend, 1990; Nelson, 1988).

Keil points out that constraints could be in the learner or outside the learner. Even focusing just on those in the learner, Keil observes that there are still strong disagreements about whether they are innate or acquired, probabilistic or absolute, regarding process or structure, domain-specific or domain-general, and so forth. Thus, a constraints view need not be domain-specific. For example, in an ingenious argument, Newport (1990) suggests that there are domain-general information-processing constraints on attention that help children acquire language. Other theorists are agnostic as to whether the constraints they propose are domain-specific (e.g., Markman, 1989).

However, in the present context, the suggestion of domain-specific constraints is of particular interest. Indeed, there may be a natural affinity between constraints and domain specificity. If constraints are appealing because they make the induction problem easier, then domain-specific constraints are all the more appealing because they make the induction problem all the more easy (Keil, 1981). The argument is that

it is necessary to grant infants and/or young children domain-specific organizing structures that direct attention to the data that bear on the concepts and facts relevant to

a particular cognitive domain. The thesis is that the mind brings domain-specific organizing principles to bear on the assimilation and structuring of facts and concepts, that learners can narrow the range of possible interpretations of the environment because they have implicit assumptions that guide their search for relevant data. (R. Gelman, 1990)

There appears to be a rich array of such constraints. Spelke (1990) proposes a variety of constraints on object perception that seem to be operating from early infancy. R. Gelman (1990) provides support for constraints on early numeric understanding (specifically, principles for counting) and causal understanding of animate and inanimate movement. She refers to these constraints as "skeletal principles" because they are the framework on which developing knowledge depends and grows. Brown (1990) shows evidence for constraints on interpretations of causal relations, and argues that such constraints guide the kinds of analogical transfer subjects find easy or difficult.

The flourishing of a constraints perspective and the wealth of evidence being amassed on domain-specific constraints do not lead to convergence on what is meant by "domain." As Keil points out (1990: 139):

The notion of domain varies considerably across researchers. In some cases, ... the domains cover very broad areas of cognitive competency such as the representation of space or physical objects. In other cases, the domains may be locally circumscribed bodies of expertise. The critical common factor in all cases is that domain specific constraints are predicated on specific sorts of knowledge types and do not blindly constrain any possible input to learning.

Theories

Another sense of domain specificity arises from considering everyday knowledge as falling into folk or commonsense theories. Theories are by nature domain-specific. A theory of biology cannot be applied to the phenomena of physics. Theories make different ontological commitments (biologists appeal to species and DNA; physicists appeal to quarks and masses). They put forth domain-specific causal laws (e.g., gravity does not affect mental states; biological processes such as growth or respiration cannot be applied to force dynamics). So, if human thought is in important ways analogous to scientific theories, then it should be organized separately for distinct domains.

The claim that there are theories is a controversial and substantive one. Theories are not in principle required for getting around the world. It is possible to form biological categories without the benefit of theories, as when pigeons classify birds and trees (Herrnstein, 1979) or humans form groupings of the biological world (Atran, 1990, this volume). It is possible to respond to others' mental states by simply reflecting on one's own (Harris, this volume).

Moreover, at first blush the claim that everyday knowledge is theory laden may seem implausible. If "theory" means "scientific theory," then it is certainly

not the case that everyday knowledge is organized into theories. It is clear that few of us have the detailed, explicit, formal understanding that PhD biologists or physicists have. We rarely, if ever, conduct scientific experiments to test our everyday hypotheses. We often lack conscious awareness of those principles that we do understand implicitly. Consider, for example, how subjects reason about physical laws regarding motion and velocity. Kaiser, Proffitt, and Anderson (1985) found that subjects implicitly know the correct natural trajectory of an object in motion, but perform badly when asked to judge the naturalness of static representations of the same events.

However, a commonsense or folk theory is not the same as a scientific theory. This point has been made by Karmiloff-Smith and Inhelder (1975), Murphy and Medin (1985), Carey (1985), Keil (this volume), Gopnik and Wellman (this volume), and others. Instead, everyday thought may be theory-like in its resistance to counterevidence, ontological commitments, attention to domain-specific causal principles, and coherence of beliefs. We sketch out some examples of these properties in the following sections.

The argument that ordinary knowledge can be likened to commonsense theories is rooted in several distinct strands of research. Karmiloff-Smith and Inhelder (1975), in an important demonstrational study, showed that children construct hypotheses (akin to miniature theories-in-action) that are resistant to counterevidence. They gave children a series of blocks to balance on a fulcrum. Some were symmetrical and thus balanced in the center, but others were asymmetrical, either visibly so (e.g., having an extra, visible weight at one end) or invisibly so (e.g., having a hidden weight at one end). Many of the children at first used trial and error to balance the blocks. They were fairly successful on the task, because they were approaching the task strictly as empiricists. However, as the session continued some children formed the explicit hypothesis that the blocks balanced in the middle, and so started making errors that they did not make previously. Children had particular difficulty with the invisibly asymmetric blocks. Some of the children, after repeated attempts, finally abandoned the blocks that would not balance in the middle, reporting that they were impossible to balance. This demonstration suggests the importance of theoretical beliefs in organizing input.

Another set of demonstrations of theorylike beliefs in ordinary thought emerges from examining semantics and categorization. Murphy and Medin (1985) propose that theories are needed to account for the insufficiencies of similarity as a construct. On a similarity view, word meanings and categories are constructed on the basis of similarity of members to one another, whether these similarities are computed in terms of prototypes, feature lists, or similarity to exemplars (Smith & Medin, 1981). However, Murphy and Medin (1985) note (following Goodman, 1972) the problems with a pure similarity view.

To give an intuitive example, if you see someone jump into a swimming pool with all his clothes on, you may classify him as a "drunk" even though