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 itinvolves solving mathematical problems, learning natural languages, calculaling 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, judgement, 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 lhan 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 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.

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?

(2) 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* (H1). 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 (H2), 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:

(1) 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:

(2) *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 indepence 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 H2, never H1, on the first relevant occasions.... We cannot, it seems, explain the preference for H2 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 H1, 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.

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. huon 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 procesors that take information from the input systems, in a format aapropriate for the central processors, and use this information to mediate higher functions, such as the fixation of beliefs.

Domains, a characterization

Is it possible, given this survey, to extract what researchers mean when they talk of domain specsficity ? In spite of the wealth of research. curiously we lack an explicit and well-articulated account of what a domain is. It is easier to think of examples of a domain than to give a definition of one. Physical entities and processes, substances, living kinds, numbers, artifacts. mental states, social types, and supernatural phenomena are all candidate domains. Which of these, if any. arc cognitive domains is a question not for general science but for anthropology and psychology: The supernatural may be a domain for human minds though it is not one for science; the cognitive domain of physical phenomena may be quite different from that of physics.

Are there features common to all domains? Our review, though brief, allows us to identify areas of accord, often implicit, in much domain-relevant research. The following we take to be a fairly uncontroversial characterization:

A domain is a body of knowledge that identities and interprets a class of phenomena assumed to share certain properties and to be of a distinct and general type. A domain functions as a stable response to a set of recurring and complex problems faced by the organism. This response involves diffieult-to-access perceptual, encoding. retrieval, and inferential processes dedicated to that solution.

Let us consider each part of this characterization in turn.

Domains as guides to partitioning the world. Most accounts converge on a view that domains function conceptually to identify phenomena belonging to a single general kind, even when these phenomena fall under several concepts. For example, living kinds can be classified in a number of different ways, ranging from foodstuffs to zoo animals. The psychological correlates of competing classifications and their internal structures have significant effects on the way many common categories for living things are sorted, recalled, and recognized (Rosch, Mervis. Gray. Johnson, & Boyes-Braem, 1976). Yet in spite of these competing ways of classifying living things, some beliefs about living kinds are typically early emerging, consistent, and effortlessly acquired. Domain competency facilitates this by focusing attention on a specific domain rather than general knowledge (Chi et al.. 1989).

Domains as explanatory frames. Most researchers would also accept that a domain competence systematically links recognized kinds to restricted classes of properties. Thus a cognitive domain is a class of phenomena that share among themselves, but not with other kinds, a number of relevant properties. Though virtually all domains seem to make reference to causal or otherwise model-derived connections, there is considerable variation .scross domains in how flexible these connections are.

It is not necessary, for example, that all and only members of a given domain share a property. It is recognised property of humans – recognized through the agency of a naive psychology - thai human beings behave in association with their beliefs, whereas artifacts do not. Still, we have little trouble accepting that there are humans who do not have beliefs (e.g., people in deep comas) or humans who do not recognize that others have beliefs (e.g., autistic individuals) or artifacts that are most sensibly dealt with as if they had beliefs (e.g., a chess playing computer). In other domains (say, auditory processing of speech), there may be substantially less flexibility in the degree to which domain properties (say, categorical perception of speech) span domain members.

Domains as functional and widely distributed devices. It is generally accepted among domain researchers that domain competencies are a restricted set of the cognitive skills the organism may develop. Domains of knowledge represent widely shared adaptations targeting recurring problems that an organism faces. Cases where the adaptive aspect of the domain has been challenged, as with the language faculty, have been much debated recently (Pinker & Bloom, 1990). Domains are also generally seen as highly (though not universally) shared among members of a species, not idiosyncratic solutions to individual problems.

Even if a domain skill is unevenly distributed within a population, it must be a solution to a repeatedly encountered problem. To the extent that chess is a domain, the development of perceptual strategies for analyzing chess positions arises *because* chess masters frequently encounter chess problems. Nonmasters do not have a less developed domain of chess; they lack such a domain skill. This relationship between frequency of encounter and domain skill, however, is complex. Some domain skills may appear to be closely tied to differences in the learning environment even if the *underlying domain competence* does *not* depend on environmental conditions.

The ability to develop and understand mathematics may be rooted in some fairly specific cognitive mechanisms, which human beings are innately endowed with. But if so, many cultures do not require that people use this ability. Nor is it occasioned by every environment. Mathematics does not spontaneously arise irrespective of social context, but seems to require a richer and more sustained sequence of experience and instruction in order to flourish than, say, basic grammatical knowledge, color perception or appreciation of living kinds. (Atran, 1988: 8)

Such mathematical skills, involving a formal language of mathematics, are distinct from other, universally emerging arithmetic competencies such as those principles underlying counting and cardinal enumeration (R. Gelman & Brenneman, this volume; R. Gelman & Gallistel, 1978).

Domains as dedicated mechanisms. Domain-specific processing is typically seen as independent of will and accessible to consciousness only with difficulty (if at alD. This property is easily evident in domains such as color perception or phonetic interpretation, cases in which some innate mechanism is difficult to doubt. But it is also apparent when an innate mechanism seems less likely, as, for instance, in what some consider a marginal case of domain specificity, like chess. Chess masters differ from chess novices in their visual-perception processing of chess information, not in logical-deductive thinking or memory processes (Chase & Simon, 1973). Accordingly, domain operations generally involve focused, constrained, and involuntary perceptual, conceptual, or inferential processes.

What domains are not. Domains can be characterized by what they are not as well as what they are. In this regard it is useful to contrast domains w sometimes supplied by the fact that the effects of domain-specific knowledge organization occur in the absence of lexical indices (as, e.g., in the case of focal color categories [Heider, 1972] or covert folk biological taxa [Berlin, 1974]). In contrast to *scripts* and *schema,* domain structures involve expectations about the model-derived connections between domain elements. In a restaurant script, payment follows service by a convention that the script itself cannot account for. Another way in which conceptual interconnectedness is achieved is via *prototypes.* According to a widely held position, categories are not built around defining features but around central members or prototypes. Because prototypes are described as consisting of collections of correlated attributes, they contribute to underlying conceptual interconnectedness. The particular category structure found in prototypes, however, applies to a wide range of phenomena; hence, domain competencies cross-cut them. Thus, prototypes could apply to any domain. (Conversely, they may not truly characterize any domain. Lakoff, 1987, suggests that prototypes are effects, reflecting other kinds of conceptual structure.) Finally, *analogical transfers* contribute to conceptual interconnectedness and mental economy. Analogical transfer may in fact represent a means for integrating domain knowledge across domains (Brown, 1989; Hirschfeld, this volume). But unlike domain structures, transfers can be both idiosyncratic and not functional.

In addition to these ways of producing conceptual interconnectedness that are clearly not domainlike, there are several means for accomplishing conceptual interconnectedness that domains may not be. These include the notions of *category* and *motoric competencies.* Our characterization of domain stressed the expectation that there is some specific "formal" or "syntactic" property of the mental representations pertaining to a given domain that accounts for their distinctive cognitive role. Unfortunately, this does not exclude arrays that we would hesitate to call domains. For example, single concepts have distinctive semantic and formal properties, yet we would not call single concepts domains. Similarly, pairs of concepts like male or female, key and lock, or hot and cold are too narrow to be domains although they fit most of our criteria. Note also that our characterization does not exclude strictly motoric competencies such as riding a bicycle, which many may not find a compelling example of domain-specific skills.

Domain differences

Clearly, a principled way of defining what a domain is continues to elude us. Many of the chapters in this volume put forward solutions to this question. And other chapters take issue with these proffered solutions. For this reason, in the present introduction we have tried to *characterize* rather than *define* domains. We have highlighted several qualities of domains, particularly their functional and semantic features. In the rest of this chapter we explore four suggestions for more closely defining domains and their consequences. These are: domains (1) as innate mechanisms, (2) as distinct ways of acquiring knowledge, (3) as reflections of specific relations between the world and our knowledge of it, (4) and as the product of a distinct research orientation.

Are all domains innate modules? One proposal for defining domains is to restrict candidate domains to sets of computationally relevant concepts that are the product of innate mechanisms or innately guided learning systems. Even if this proposal is adopted, significant issues remain unresolved. For example, Fodor's modular hypothesis sees domains in terms of discrete subsystems, each tied to a specific perceptual channel. Other biologically oriented approaches, including Atran's claims about biological beliefs (this volume) or Leslie's contention about theory of mind (this volume), construe domains in terms of much larger competencies implicating numerous perceptual modalities. It is thus not evident that their positions and Fodor's would pick out the same candidate domains. Chomsky (1986) similarly argues that viewing the language module as an input system is too narrow.

Can domains be denned by their mode of acquisition? Domains may be defined in terms of the specific pattern of learning associated with them. Keil (1990) and Atran and Sperber (1991) argue that candidate domains may be identifiable with respect to their mode of acquisition. Many domain skills are acquired with ease by virtually all members of the species. A broad consensus has emerged that the language faculty is a universally distributed and rapidly acquired cognitive system whose acquisition depends on an innately guided learning device. Other domains seem to be considerably less widely distributed, acquired only with great effort, and appear to be outside the scope of any innately dedicated program of learning.

Chess mastery, for example, is an unevenly distributed cognitive ability that requires determined and often formal training to acquire. Even within a single domain there may be considerable variation. Many aspects of naive biology emerge spontaneously across quite different learning environments (Boster, 1988). Yet, some biological knowledge appears to be sensitive to differences in expertise, such as the degree to which elements within a content area are seen as causally interconnected (Chi et al., 1989). Inagaki (1990a) suggests that familiarity with a domain increases not only the amount of factual knowledge children have but also their conceptual organization of the domain.

To what extent can domains be denned by their relationship to the external world? Some domains appear to carve the world at its joints. When we think of domains, we think of ranges of phenomena comparable to the ranges of facts that are the subject matter of the different sciences. Naive linguistics, biology, psychology, physics, mathematics, cosmology, and so forth, have all been proposed as domains. It is tempting to go from this observation to a view that some causal link exists between the structure of empirical sciences and domain competencies (Carey, 1985). Domains, on this view, would *necessarily* be comparable in size to the subject matter of the different sciences. Is there reason to believe that we are unduly influenced by the way scientific disciplines partition the universe?

Yes. For example, some domains not so much carve nature at its joints as they create such joints. Color terminology and perception, for example, reflect discrete categories of the mind imposed on a continuous natural phenomenon. In processing speech we automatically ignore variations between the representations of what we take to be the same sounds across speakers and within the utterances a single speaker produces. Yet, when attending to the same human voice singing we just as automatically are sensitive to variations in pitch. In short, color and speech perception are less discoveries about nature than interpretations of it. By the same token, a fairly extensive literature documents domainlike competencies that span contrived phenomena not typically associated with a unique science or discipline - ranging from chess to chicken sexing (Biederman & Shiffrar, 1987) to reading x-rays (Lesgold, Rubinson, Feltovich, Glaser- Klopfer, & Wang, 1988) - whereas others encompass phenomena that lack scientific validity or correspondence (for example, race, magic, or the supernatural).

The degree to which a domain is dependent on the world is complex and variable. Some domains (e.g., naive biology and physics) may be less closely linked to our scientific understandings of the world than is sometimes appreciated. Some changes in scientific understanding quite lapidiy become part of commonsense. Thus, whales are widely understood to be mammals, not fish, even by children. Examples of this sort prompted Putnam (1975) to speak of a linguistic division of labor, in which science is seen as determining the meaning of natural kind terms. But changes in the validity of a given scientific or formal description do not always alter our commitment to the correspondence between these systems and formal systematics in biology, is not a new discovery. Ernst Mayr (1949) documented close parallels between nonliterate Papuan folk biology and formal systematics. However, these parallels were interpreted as evidence for the scientific validity of the species concept rather than as evidence of a shared cognitive mechanism for conceptualizing living kinds (Diamond, 1966). In other words, it was assumed that a close correspondence between common-sense and science reflected regularities in the world external to cognition, rather than indicating a set of shared cognitive dispositions. As several of the chapters in this volume suggest, the empirical and structural parallels between folk beliefs and science are indeed informative; but from the perspective of understanding mental domains, these parallels are more pertinent to the cognitive than the other sciences.

Does the question asked make a difference? It is possible that the research question posed may have profound impact on the candidate domains we discover. It may not, in short, be possible to distinguish between the research interest that prompted a set of data and what those data can be interpreted as supporting. A case in point may be the relation between a competence associated with a given domain and the knowledge represented in that domain. Keil (1981) makes a similar point regarding approaches to developmental change, when he distinguishes between theories that are knowledge based from those that are competence based. Keil's concern is with constraints, but his observation can be extended to a discussion of how to approach domains. In broad terms, knowledge-based research sees domains as being derived from knowledge structures, whereas competence-based research focuses more on how differing states of knowledge are derived from distinct competencies. The strongest versions of each of these are, respectively, exoertise studies and modular approaches.

The competence-based view lays emphasis on the domain-specific principles of organization that shape the sorts of knowledge to which the organism attends and the kinds of knowledge structures the organism develops (sec Gallistel et al., 1991). A central concern of the competence-based researcher is how attention is guided toward a circumscribed set of observations, not only during development but also during mature processing. Knowledge-based studies, in contrast, focus more on the consequences of holding domain knowledge than on the prirciples that underlie its acquisition (see Shantz, 1989; Chase & Simon, 1973; de

Groot, 1966). The principal question posed in this work is what impact does domain knowledge have on subsequent processing? Given a certain level of knowledge integration, what are the implications for thought?

Accordingly, the two viewpoints could be seen simply as different perspectives on the same phenomenon, in which somewhat distinct emphases are applied. Knowledge- and competence-based approaches would accordingly be slightly different points of view on a single issue, namely, the relationship between knowledge of something and the competence underlying that knowledge. The two orientations, however, can also be seen as different levels of analysis, yielding quite different kinds of generalizations. Keil, for example, argues that the knowledge-based approach does not provide a principled way of "distinguishing what is merely a change in a knowledge base from what is a change in computational and representational machinery" (1981: 204). Knowing that a cognitive skill changes with experience thus does not permit us to decide whether that change is best described in knowledge-or competence-based terms. By the same token, knowing that a knowledge-based hypothesis accounts for a set of observations does not permit us to conclude that a competence-based vantage would construe the observations as coming from a bounded type.

Consider, for example, the problem of deciding whether a change in the way knowledge is represented reflects a change in domain competence. There appears to be a domain of dinosaur knowledge for which some children develop a special expertise. Thus Chi and her associates (1989) have shown that increased knowledge about dinosaurs leads to an enriched network of causally linked beliefs. Generally, the conceptual effects of such causal networks appear to be quite limited, in the sense that they do not readily transfer to closely allied content areas (Inagaki, 1990b). From a competence-based perspective, however, dinosaur knowledge is not so much a domain as a subset of another domain, namely, living things. What are the implications of this?

One implication is that knowledge-based claims may not commit the researcher to a particular opinion on questions like "What is a domain?" "What are domain boundaries?" or "How are domains implicated in the learning process?" Discovering the relationship between domain-specific and other sorts of knowledge is presumably dependent on determining what is and what is not domain knowledge, and by extension what is and what is not a domain. Thus, like the contrast between domain-general and domain-specific perspectives, knowledge-based and competence-based views may differ in terms of the kinds of generalizations that each can enlist, particularly about domains.

Conclusion

In this introduction we have sketched out several distinct intellectual traditions that have contributed to the view that human cognition can be fruitfully considered domain-specific. Clearly, the approach seems to us extremely promising. Still, a series of difficult questions remain: What is a domain? How many domains are there? Do they concern processes, representations, or both? Do we have conscious access to domains? Can domains change over time, and if so, how? Do the changes that occur in childhood differ in important ways from those changes in domain knowledge that take place after maturity? Does conceptual change occur because of or in spite of domains? Do domains "communicate" among themselves, and if so how? We leave these issues for the reader to consider.