COGNITIVE SCIENCE 14, 169–178 (1990)

Safe Takeoffs—Soft Landings

DOUGLAS L. MEDIN, WOO-KYOUNG AHN, JEFFREY BETTGER, JUDY FLORIAN, ROBERT GOLDSTONE, MARY LASSALINE, ARTHUR MARKMAN, JOSHUA RUBINSTEIN, AND EDWARD WISNIEWSKI

University of Illinois at Champaign-Urbana

I. INTRODUCTION

The richness of life's possibilities creates a paradox. It poses fundamental limitations on the ability of an organism to reason and act intelligently, as information in its environment is often ambiguous and consistent with an infinite number of interpretations. At the same time, interpretation of this information is required for intelligent action. To further complicate the situation, many problems that are solvable in principle are intractable in practice. As an example, the number of winning sequences of moves in a chess game is finite. Yet determining these sequences from all possible sequences is impractical for the chess player, as there are more possible sequences than atoms in the universe. Cognitive scientists are certainly familiar with underdetermination and computational complexity in association with work on language learnability theory (e.g., Gold, 1967; Wexler & Cullicover, 1980), but complexity problems extend to virtually every area of cognition from perception to concept learning to decision making.

If organisms cannot, and therefore do not, examine all possibilities in some cognitive task, they must be "prepared" or biased to learn some things rather than others, to draw some plausible inferences rather than others, and in general to favor some possibilities at the expense of others. As evidenced by the present set of exemplary papers by Brown, Gelman, Keil, Markman, Newport, and Spelke, psychologists are increasingly drawing on analyses of constraints or "guiding principles" as a framework for their research. All of the present authors are interested in development, and it is no accident that developmental psychologists have been leaders in recognizing the importance of constraints. Without guiding principles of some form it simply is unclear how the learning process would ever get off the ground.

The purpose of this comment is to underline some of the important points and issues growing out of developmental work on constraints as reflected in the contributed papers. We first discuss a couple of distinctions and definitions associated with research on constraints, then shift attention to the per-

Correspondence and requests for reprints should be sent to Douglas L. Medin, University of Michigan, Perry Bldg., 330 Packard, Ann Arbor, MI 48104.

spective provided by a focus on development, and end with some speculations concerning relationships among constraints within and across domains.

II. DISTINCTIONS AND DEFINITIONS

Frank Keil's discussion brings out a variety of dimensions along which constraints might differ. He focuses on the contrast between domain-specific and domain-general constraints, and innate versus acquired constraints. Whether or not the other contributors can be neatly classified into this two-by-two matrix is less important than the observation that taking a stance on these issues is a step toward a more explicit and complete analysis of constraints. We turn to two other contrasts mentioned by Keil because we believe they also merit emphasis.

A. Hard Versus Soft Constraints

Some researchers have assumed that constraints are necessarily hard-wired (innate) and permit no flexibility. These constraints are hard in the sense that they directly limit variability. For example, linguistic universals provide one candidate set of hard constraints.

In contrast to this view, we interpret a constraint or guiding principle as any factor that favors some possibilities over others. The papers in this issue manifest a wide range of positions with respect to constraints. Markman's study showing a shift from 37% taxonomic choices to 63% taxonomic choices when a novel label was introduced suggests that the taxonomic constraint is a "soft" constraint that nonetheless may exert an important influence on children's learning of word meanings. At the other end of the continuum, Spelke's cohesion, boundedness and spatiotemporal continuity constraints associated with object perception may well be universal. One reason for expecting less variability in notions about objects is that these assumptions are basically correct. To examine the flexibility of assumptions about objects, one would need to create an environment where these constraints are consistently and systematically violated.

It must be noted that the position presented in the papers in this volume is not unanimously accepted in the developmental psychology community. A different perspective is offered by Nelson (1988), who maintains that constraints connote restriction, specificity, and invariance, a "closing down of choice." According to Nelson, if some aspect of cognitive development (in this case, word learning) is supported by a preference or bias rather than driven by a consistent, universal restriction, depiction as a constraint is misleading. Rather, lexical development is better characterized as a process of social convergence on conventional word meanings, in which the child is guided by the adult. Under social convergence, the child is given a greater role in pruning the search space of meaning. During the naming explosion at about 16–18 months of age, it is the child who initiates the request for the name, she suggests. In this way, the child has already determined the entity to be named, and needs only to have the label supplied by the adult.

Although Nelson might not endorse our reading of her position, we believe her view diverges from the contributed papers primarily in terms of matters of emphasis. First of all, she uses a much more narrow definition of constraints than the contributors. Second, she appears to differ from Markman in stressing the social rather than the linguistic context of language learning. It is important to identify the roles of children and adults in language learning, but as Markman points out in her discussion of Quine's rabbit example, the fact that the situation is a rich social context does not, by itself, address the learnability problem.

Developmental psychologists are not the only ones who have studied mechanisms that aid learning in a complex world. Researchers who simulate learning routinely build constraints into their computer models, sometimes even without explicitly intending to. For example, researchers often preselect the features that describe the training examples used by inductive machine learning systems. This preselection gives the system an inductive bias (Mitchell, 1982) because it limits the possible features that the system will consider when learning from examples.

Other work in machine learning has discussed two kinds of learning biases that result from the representation language and description language used by a learning program (Utgoff, 1986). The representation language delimits the range of things in the world that the program is capable of representing and thus what it can learn, in principle. In this way, it constrains the system in a "hard" sense. The description language delimits the range of things that the system can currently represent. It therefore provides a bias for those things that the system can learn about an object. In brief, the representation language represents a set of hard constraints whereas the description language provides soft constraints.

One can also see the importance of constraints in connectionist networks. Many systems have certain structural constraints on their architectures and these in turn affect what the system can learn and represent. This point is especially clear if one contrasts simple networks consisting of a layer of input units that are connected directly to a layer of output units with networks that contain intermediate layers of "hidden" units. It is well known that simple networks can learn associations between input vectors and output vectors only if the input vectors are linearly independent. However, linear independence does not hold for many "interesting" cognitive problems such as learning "exclusive or" (Hinton, 1987).

Multi-layered networks can overcome the limitations of single-layered systems. From a constraints perspective, however, this may be a mixed blessing. Constraints are needed to overcome computational complexity problems, and many otherwise powerful computational procedures will not "scale up" to more complex problems. The trick is to provide a system with the "right constraints." For a cognitive psychologist the constraints in a network (or program) should ideally be just the ones revealed as characterizing human performance. Indeed, the constraints perspective has become the framework of choice for much of the recent research in the area of computational vision (e.g., Cavanagh & Leclerc, 1989; Poggio, Torre, & Koch, 1985).

B. Structure, Process, and Outcome Constraints

As Keil points out, structure and process are so interwoven that it is difficult or impossible to disentangle them, even conceptually. Still one can make a few generalizations about factors that tend to be correlated with structure versus process. Processing constraints usually are described in a domain-general manner, with observed performances emerging from an interaction of processing with domain-specific knowledge structures. Consider Newport's paper suggesting that limited working memory may paradoxically constrain induction in a way that favors language acquisition. Here, a processing constraint on working memory capacity affects domain-specific knowledge structures in the domain of language.

Marler's (1964) auditory template hypothesis for song learning in birds provides an example of a structurally-oriented constraint. The basic idea is that birds are born with a template that may be modified, within limits, through experience. The template is biased toward songs of a bird's own species, and in the absence of appropriate auditory input the bird may not develop its normal song, or may develop only fragments of it (Marler, 1970). The auditory input must be roughly in the right form to influence the template (although there appear to be differences both across species and perhaps even among individual birds [Marler & Sherman, 1983]).

Outcome or product constraints seem to be distinct from either structure or process constraints. Both structure and processing constraints restrict the range of possibilities to be considered and therefore directly address computational complexity. Outcome constraints seem naturally to apply to a set of possibilities that has already been generated, as in a chess-playing program that evaluates a large number of candidate moves. In some cases, outcome constraints may represent the theorist's "shorthand" for some yet-to-be specified structural or processing constraints. The contributors to this special issue primarily focus on structure and process.

III. WHY DEVELOPMENTAL WORK ON CONSTRAINTS IS IMPORTANT

Although computational complexity is a central factor motivating an interest in constraints, it certainly is not the only one. The generality of constraints, the interaction of different constraints, and their emergence over time all demand better understanding. Consider the contrast between the challenge offered by a game like chess and the challenge of development. Grandmaster chess players base their expertise on thousands and thousands of games where their favorite variations can be tested and retested. After every game the pieces are reset to common beginning positions. Development is continuous, with no opportunity to "reset the pieces" and systematically explore alternative organizational principles. Initial commitments have far-reaching implications, and an important challenge in the study of development is to trace the course and consequences of these commitments.

The contrast is further extended when considering the generalizability of each experience. Although chess skill may be correlated with more general spatial reasoning abilities, it probably is a relatively isolated skill. (We say probably because it may be a mistake to prejudge the issue.) Grandmasters probably would not be likely to show any more or any less ability in gymnastics or second-language learning by virtue of their chess playing skill. Although Chomsky (1965, 1975) and others have argued that there may be a language organ or module, there is increasing evidence that development is coordinated between domains in intricate ways. We will later suggest ways in which separate domains share constraints and are interdependent.

In brief, there are several reasons why developmental work on constraints is of special importance. In addition to taking both the notion of a domain and the notion of a constraint seriously, developmental work forces us to consider the temporal unfolding and interactive character of multiple constraints. Also, ongoing work provides a rich source of ideas and challenges with respect to the treatment of constraints.

A. Domains and Unfolding

Developmental research leads naturally to questions about the unfolding of constraints and the organization of knowledge within and across domains. For example, one might speculate that the later a constraint exerts its influence, the more specific its influence will tend to be. The idea is that if an early constraint gets things headed in the right direction, then only minor adjustments may be needed later on. Another factor that might favor less specific constraints is that they may be borrowed or propagated across domains. For example, there is evidence suggesting that the recognition of causality is present very early in infants (e.g., Leslie & Keeble, 1987; Michotte, 1963).

B. Cooperation and Borrowing of Constraints Across Domains

Constraints may not always operate in isolation. A rudiment of one function may represent a refinement of another. A nice example of the cooperation of constraints comes from studies of the feeding behavior of herring gull chicks (Hailman, 1967). The newborn chick has little coordination and its pecks are poorly aimed. The red spot on the adult gull's lower mandible acts as a sign stimulus that triggers or "releases" pecks by the chick. The chick cannot recognize food, but by aiming at the bill of the adult and missing it strikes food and rapidly learns to recognize it. These feeding experiences also allow the chick to learn to recognize its parents. With practice, the chick comes to differentiate a food begging peck from its feeding peck. The fixed action pattern of pecking at the red spot enables both parental recognition and the development of feeding behaviors. This example illustrates Frank Keil's point that one cannot meaningfully parse development into innate and acquired components.

In considering human cognitive development, Spelke speculates that early constraints on object perception provide a bootstrap for learning how properties are associated with objects. Organizing scenes into objects is originally based upon constraints such as "Assume objects move as wholes" and "Assume objects move independently of each other." While these constraints do not include "Assume objects have the same color or texture throughout," this further constraint is discoverable once the objects have been parsed on the basis of the early constraints. In extension, such organizational constraints make possible the acquisition of still further properties that are associated with objects. Similarly, Gelman postulates a constraint that functions as a conceptual "skeleton," dividing objects into animates and inanimates and thereby allowing a child to make extended inferences about other properties of animals that may be nonvisible. This skeleton provides the conceptual framework upon which to organize properties that are correlated with the animate and inanimate distinctions. Brown demonstrates that a sensitivity to children's theories and assumptions concerning fundamental processes about causality can be used to produce learning and transfer performance that far surpasses conventional studies uninformed by a knowledge of guiding principles.

Cooperation may also function across domains in the same manner as across constraints. The language model relies on, and is relied on by, other domains. Markman's mutual exclusivity constraint, whereby ambiguous new labels are associated with unfamiliar as opposed to familiar objects, can only be applied once a visual field has been parsed into objects. This parsing of the world into objects, Spelke argues, will be governed by its own unique constraints. The language module, in turn, aids other cognitive processes. The mutual exclusivity constraint, for example, affects inferences drawn concerning properties. If a child already has a label for an object, a new label associated with the object will often be treated as referring to a *property* of the object. In this manner, language constraints may shift attention from the whole object to components of the object, thus playing an instrumental role in the development of induction.

One might also expect borrowing of constraints across domains. Shepard (1984) has provided an eloquent case for the idea that human cognition of music derives from our spatial representation system. When constraints are borrowed from an earlier domain, there may be telltale signs that reveal the constraints' origin. If imagery processes are borrowed from visual perception, for example, then we might expect visual constraints to be evident in imagery. Indeed, Finke and his associates have found that many visual constraints are also exhibited when subjects are asked to visually imagine objects. For example, creation of the illusion of apparent fusion requires a longer SOA than does the illusion of apparent explosion, paralleling a bias in the visual system for natural, irreversible transformations (Finke, 1985; Freyd, 1983). Next, Finke & Freyd (1985; also Finke, Freyd, & Shyi, 1986) demonstrated that representational momentum (a shift in visual memory induced by implied rotation) obeys some of the same laws as physical momentum. As a final example, Finke and Kosslyn (1980) obtained parallel changes in resolution with distance from the focal point of both perceived and imaged patterns.

The fact that constraints may cooperate to serve more than one function should not be interpreted as implying a planful "conspiracy." Nature has to work with the tools it is given, and one might imagine that exploitation of one structure for something other than its initial function occurs as well as, or instead of, co-adaptation (e.g., Gould & Lewontin, 1984; Piattelli-Palmarini, 1989).

IV. Soft Landings

The various aspects of constraints discussed by the authors combine to form an intriguing perspective on learning and development. Constraints are perhaps best seen as enablers, with the mosaic of constraints working to ensure a "fine tuning" of learning as soft landings. Constraints not only limit the range of possibilities considered; they also permit the later acquisition of new constraints. The constraints that Spelke discusses are interesting not only because they limit what scene fragments will be combined into objects, but also because they permit the later acquisition of other object-property systematicities. Construing constraints as "permitters" as well as "limiters" suggests that they do more than just weed out possible alternatives. Although seemingly paradoxical, a system with constraints may actually have more expressive power than an unconstrained system, if the constraints provide a "language" for expressing other processes, distinctions, and constraints.

We offer the speculative analogy: The organization of various cognitive domains may parallel the organization of the thermoregulatory system. Recent review articles by Satinoff (1978, 1983) trace the evolution of thinking concerning temperature regulation in mammals. Initially investigators thought that animals had a single internal thermostat that could trigger reflexes or behaviors designed to adjust body temperature. Next the idea that reflexes and behaviors were controlled by separate thermostats became popular. Most recently the preponderance of evidence suggests that there are a number of parallel mechanisms for regulating temperature, each with its own thermostat, that act in concert to maintain body temperature within a narrow range. Furthermore, most of these systems were originally used for other purposes and only over long periods of evolution came to be controlled by temperature detectors. For example, all animals breathe, and panting, which is an excellent means of losing heat, involves borrowing from or taking advantage of the respiratory systems. In Satinoff's words, "Eventually there came to be many such systems, each evolved at separate times and each independent of the others but not destined to remain independent." (Satinoff, 1983, p. 461). We think that the development of human cognition may manifest parallels with the phylogenetic development of thermoregulation.

Perhaps we should have said "*very* speculative analogy" for we have little to support it. We suggest that different cognitive domains may have a variety of parallel constraints that allow the successive refinement and coordination of learning. Medin and Ortony (1989) propose one such parallel in their discussion of psychological essentialism. They argue that learners have a bias for thinking that things that are superficially similar are alike in deeper, more principled ways. Adults are notoriously flexible about what counts as similar (e.g., Tversky, 1977), but young children may not be (Smith, 1989). The inflexibility of young children's perception of similarity may be precisely what allows them to set up categories that will be useful later on. The more principled classification schemes would represent a refinement of the categories given by the perceptual system.

V. SUMMARY

The set of papers in this special issue is indeed provocative. Using an analysis of constraints to trace the course of development may not yet be "the only game in town" but, in our opinion, it provides a fundamental and critically important perspective on cognitive development in particular as well as cognitive science in general. Further research on developmental constraints promises to provide more accurate descriptions of mental representation for computer modelers of human cognition, to stimulate new approaches to learning, and perhaps to propose common origins of cognitive abilities that were previously thought to be unrelated. It is no mean accomplishment to face up to computational complexity and still produce safe takeoffs and soft landings.

REFERENCES

Cavanagh, P., & Leclerc, Y.G. (1989). Shape from shadows. *Journal of Experimental Psychology: Human Perception and Performance, 15*, 3–27.

Chomsky, N. (1967). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.

Chomsky, N. (1975). *Reflections on language*. New York: Random House.

Finke, R.A. (1985). Illusions of apparent visual explosion and fusion. *Bulletin of the Psychonomic Society, 23*, 321–324.

Finke, R.A., & Freyd, J.J. (1985). Transformations of visual memory induced by implied motions of pattern elements. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 780–794.

Finke, R.A., Freyd, J.J., & Shyi, G.C.-W. (1986). Implied velocity and acceleration induce transformations of visual memory. *Journal of Experimental Psychology: General, 115,* 175–188.

Finke, R.A., & Kosslyn, S.M. (1980). Mental imagery acuity in the peripheral visual field. *Journal of Experimental Psychology: Human Perception and Performance, 6,* 244–254.

Freyd, J.J. (1983). The mental representation of movement when static stimuli are viewed. *Perception & Psychophysics, 33,* 575–581.

Gold, E.M. (1967). Language identification in the limit. *Information and Control, 10,* 447–474.

Gould, S.J., & Lewontin, R.C. (1984). The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist programme. In E. Sorber (Ed.), *Conceptual issues in evolutionary biology.* Cambridge, MA: Bradford Books/MIT Press.

Hailman, J.P. (1967). How an instinct is learned. *American Scientist,* 98–106.

Hinton, G.E. (1987). Connectionist learning procedures. *Artificial Intelligence.*

Leslie, A.M., & Keeble, S. (1987). Do six-month-olds perceive causality? *Cognition, 25,* 265–288.

Marler, P. (1964). Inheritance and learning in the development of animal vocalization. In M.C. Busnel (Ed.), *Acoustic behavior of animals.* (pp. 228–243). Amsterdam: Elsevier.

Marler, P. (1970). A comparative approach to vocal learning: Song development in white-crowned sparrows. *Journal of Comparative and Physiological Psychology.* (Monograph 71, No. 2, Part 2.)

Marler, P., & Sherman, V. (1983). Soap structure without auditory feedback: Emendations of the auditory template hypothesis. *The Journal of Neuroscience, 3,* 517–531.

Medin, D.L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similiarity and analogical reasoning.* New York/London: Cambridge University Press.

Michotte, A. (1963). *Perception fo causality.* London: Methuen.

Mitchell, T.M. (1982). Generalization as search. *Artificial Intelligence, 118,* 203–226.

Nelson, K. (1988). Constraints on word learning? *Cognitive Development, 3,* 221–246.

Marler, P. (1970). A comparative approach to vocal learning: Song development in white-crowned sparrows. *Journal of Comparative and Physiological Psychology.* (Monograph 71, No. 2, Part 2.)

Marler, P., & Sherman, V. (1983). Soap structure without auditory feedback: Emendations of the auditory template hypothesis. *The Journal of Neuroscience, 3,* 517–531.

Medin, D.L., & Ortony, A. (1989). Psychological essentialism. I S. Vosniadou & A. Ortony (Eds.), *Similiarity and analogical reasoning.* New York/London: Cambridge University Press.

Michotte, A. (1963). *Perception fo causality.* London: Methuen.

Mitchell, T.M. (1982). Generalization as search. *Artificial Intelligence, 118,* 203–226.

Morgan, J.L. (1986). *From simple input to complex grammar.* Cambridge, MA: MIT Press.

Nelson, K. (1988). Constraints on word learning? *Cognitive Development, 3,* 221–246.

Piatelli-Palmarini, M. (1989). Evolution, selection and cognition: From "learning" to parameter setting in biology and in the study of language. *Cognition, 31,* 1–44.

Poggio, T., Torre, V., & Koch, C. (1985). Computational vision and regularization theory. *Nature, 317,* 314–319.

Satinoff, E. (1978). Neural organization and evolution of thermal regulation in mammals. *Science, 201,* 16–22.

Satinoff, E. (1983). A reevaluation of the concept of the homeostatic organization of temperature regulation. *Handbook of behavioral neurobiology, V, 6: Motivation.* 443–474.

Shepard, R.N. (1984). Attention and the metric structure of the stimulus space. *Journal of Mathematical Psychology, 1,* 54–87.

Smith, L.B. (1989). From global similarity to kinds of similarity: The construction of dimensions in development. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning*. Cambridge/New York: Cambridge University Press.

Tversky, A. (1977). Features of similarity. *Psychological Review, 84,* 327–352.

Utgoff, P. (1986). Shift of bias for inductive concept learning. In R.S. Michalski, J.G. Carbonell, & T.M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach,* Vol. 22. (pp. 101–148). Los Altos, CA: Morgan Kaufman Publishers, Inc.

Wexler, K., & Cullicover, P. (1980). *Formal principles of language acquisition*. Cambridge, MA: MIT Press.