Available online at www.sciencedirect.com
ScienceDirect
Cognitive Systems Research 8 (2007) 131–133
Cognitive Systems RESEARCH
www.elsevier.com/locate/cogsys

Book review

Review of Symmetry, Causality, Mind, Michael Leyton; Cambridge, Massachusetts: MIT Press, 1992, 630 pages

Action editor: Stefan Wermter

Monika Krishan

Department of Psychology, Rutgers University, New Jersey, USA

Available online 11 December 2006

Cognitive scientists often make use of the expression *Representation is Explanation*[1] while referring to the nature of the relationship between the mind and the external world. That cognition consists of "explaining" a particular stimulus appears to agree with our intuitions. But what, in the context of the mind, is Explanation? "Symmetry, Causality, Mind" (SCM) attempts a definition of the term and in doing so shows how it is possible for the mind to make sense of the endlessly diverse stimuli in our environment. This interpretation of mental representation is presented as a general theory of cognition that unifies the various modules of Cognitive Science. The book also provides a quantification of the notions of Explanation and Representation in terms of abstract group theoretic concepts.

How does the mind arrive at a particular representation of a stimulus? Why this representation and none other? Why do we sometimes alternate between perceptions of the same stimulus? What does the stimulus tell us about itself? Why does the mind operate under the assumption that there's more to the stimulus than meets the senses? SCM offers radical solutions to the problems surrounding these issues in a wide spectrum of domains. Chapter 1, "Recovering Process History" provides an overview of the proposed principles of stimulus organization. Chapters 2 and 3 ("Traces" and "Radical Computational Vision") discuss the application of these principles to visual stimuli, while Chapter 7 considers linguistic representations. Chapter 8 extends the treatment of visual cognition to the perception of art. Chapter 9 reflects on the interaction between the quality of the social environment and state of the human mind. Chapters 4, 5 and 6 form the core of SCM. Here Leyton presents his ideas as a computational theory of cognition and shows how the abstract nature of groups allows these ideas to be generalized to almost any domain of study.

One of the key ideas presented in this book is that representation involves giving the stimulus a machine-like description characterized by a set of operations that generate this particular stimulus. A stimulus, therefore, is designated by a series of operations that act to give the stimulus its present form. According to Leyton, every stimulus object is perceived to originate in a state of least distinguishability i.e., one that is maximally homogeneous. The associated set of operations then transform the featureless "original" into something with greater differentiability, much like the movement of an embryo through the various foetal stages, with increasing definition at each stage. Thus the claim is that, to the mind, a stimulus is not a static object, frozen in time but rather that it comes with a past, a past that is inferred from the current form of the stimulus. That is, the mind extracts, from the shape of an object, information about the forces that lead to its being just so. For instance, a soda can with a dent is perceived as something that was first a smooth can which was then acted on by some other object that caused the can to undergo a deformation. Leyton's suggestion is that the mind, in the process of forming a representation of a stimulus object, attempts to uncover the *history* of the object, as it were. Since representation involves tracing the stimulus object's history in terms of the various processes that caused the stimulus to be what it is, the result is a *causal* history of the stimulus. Leyton points out that processes that leave no trace cannot be recovered. Therefore, the perceived shape of an object, in effect, gives it its history.

E-mail address: monika.krishan@gmail.com

[1] This expression was coined by the author of Symmetry, Causality, Mind, Michael Leyton and is now frequently used by the cognitive science community.

1389-0417/$ - see front matter
doi:10.1016/j.cogsys.2006.11.001

So how does the mind accomplish the recovery of stimulus history? Leyton has put together several principles that govern the representational process with emphasis on their generality. These include the History Minimization Principle which states that the representation of a stimulus object is one that is the most concise with respect to the number of operations required to transform the stimulus from its initial state to its current state. In other words, this representation corresponds to the shortest causal history of the stimulus. Consider the example of a perfectly smooth soda can. Although the smoothness of this soda can might be attributable to the application of a sequence of opposite, deforming forces, Leyton's claim is that the mind constructs a much shorter causal history of the soda can, characterized by the absence of any kind of deformation. Multiple representations of an object are said to arise from multiple causal histories. For example, a metal can with its diameter equal to its height may be seen as the result of a rotation of a vertical line in space or as the translation of a circle. The idea of a "minimal" representation has been in existence in one form or another since the beginning of the information processing trend in psychology. However, the minimization referred to in the literature typically involves a reduction in the descriptive complexity of a stimulus item. Leyton's proposal departs from traditional Kolmogorov/Chaitin complexity in at least two ways. The latter does not make any claim about the actual representation generated by the mind, only that it is possible to conceive of a "most-compact" description for a stimulus, given a particular language of description. Further, "description" in the latter refers only to the current state of the stimulus while Leyton's "representation" includes references to previous states of the stimulus that may or may not be available for inspection in the present. For instance, Leyton's mind would represent a rectangle not just as a "4-sided shape with opposite sides of equal length" but a "4-sided shape with opposite sides of equal length resulting from a stretch operation performed on a 4-sided shape with all sides of equal length, i.e., a square".

Complementing the History Minimization Principle, are the Symmetry Principles given in SCM which further determine the recovery of causal history. Humans have always been fascinated by the widespread existence of symmetry in nature. The Gestaltists were some of the first to suggest that the mind is sensitive to regularities occurring in the environment. However Leyton proposes that symmetry detection is not an end in itself but serves to delineate the causal trajectory of the stimulus, in two specific ways. In unraveling the past of a stimulus object, whether a 2-dimensional shape or a sentence, the symmetries currently present in the object are retained. Thus the "past" of a rectangle (a square) is chosen so as preserve the reflection symmetries that exist in the rectangle. But which of the many pasts consistent with this symmetry preserving condition does the mind identify? The answer to this, too, is given by the current symmetries of the stimulus, or rather the corresponding axes of symmetry. The particular operation that takes the stimulus past into the stimulus present is assumed to have acted along the directions of these symmetry axes. Referring to the example of the rectangle, the direction of the stretching action that transforms the square into the rectangle would be along the vertical/horizontal symmetry axis depending on which two of the four sides are longer. Therefore, a square is derived from the rectangle by (a) identifying the symmetry axes of the rectangle and (b) by performing the inverse of the stretching operation along the appropriate axis of symmetry. Observe that doing so leaves the symmetry present in the rectangle intact. Thus the path from the present to the past involves increasing the symmetry present in the stimulus by removing existing asymmetries at each stage of the process. The recovery of causal history therefore is subject to the presence of asymmetries or inhomogeneities in the stimulus. A perfectly symmetrical object, such as a sphere, carries no information about its "cause". One could conceive of a situation wherein the sphere originates in, say, a humanoid shape. But Leyton's point is that the mind does not see the sphere as such, that a backward trace involving a gradual increase in asymmetry is contrary to the mind's design.

The claim that internal principles of organization determine causal histories which may or may not be veridical suggests that Leyton favors the view that the mind in making sense of the world imposes structure rather than discovers it. While this claim seems reasonable it might require qualification. For instance, it is quite conceivable that the context surrounding the stimulus object exerts some influence on the type of causal history inferred. Further the mind's own history, in other words, the experience of forming causal histories in the past, might come to modulate this process of inferring causal histories. Consider a rock with a relatively smooth surface. The appearance of this rock might give the impression of wear by water or perhaps other hard objects, having originated from a state of greater asymmetry.[2] This does not necessarily contradict SCM's principle of symmetry preservation in going from the object's present to the object's past. One could speculate that Leyton's intention is to have the preservation and minimization principles be applied to the entire system consisting of the actual object, the current context and the mind itself. However, this is a highly non trivial task and perhaps the purpose is to motivate the reader to conduct an exploration along these lines rather than present an actual solution to the problem. In any event, this particular interpretation of Leyton's theory, if accurate, does make the mind and the environment part of a continuum, in a concrete sense.

Having discussed the tracing of the causal history of a stimulus by means of its symmetries, Leyton goes on to speculate that these very same ideas are exemplified in

[2] This issue has been addressed by Michael Leyton in a subsequent book "A Generative Theory of Shape" (New York: Springer-Verlag, 2001).

the art, music and architecture that surrounds us. That is, the principles that determine how the external world is represented also govern the creative process (see the section on Picasso's paintings for a rather unusual treatment of the artist's work). These, Leyton believes, also influence our aesthetic preferences. Certain stimuli are especially appealing, relative to others, because of their particular causal histories. The latter are either too simple or too complex and reflect the corresponding ease or difficulty associated with generating their causal explanations. Much as love is now merely a measure of pheromone-al activity, beauty, SCM seems to suggest, only reflects representational effort on the part of the beholder. Novices and experts in a domain tend to give different representations to the same stimuli. Given that the perceived complexity of a stimulus is a function of its causal trace, experts and novices would also be expected to differ in their perceived complexity and therefore desirability of this stimulus. This, Leyton would suggest, is what lies behind the enormous variability in responses to art exhibits.

Symmetry, Causality, Mind makes for interesting reading. Care has been taken to avoid ambiguity and unnecessary rhetoric, features that are hard to escape when the mind is the elusive subject of discussion. The numerous ideas expounded are intriguing and seem promising in their potential applicability to a wide range of phenomena. Leyton has made a serious effort to flesh out applications to certain domains, while the exploration of many others has been left to the reader as an exercise. However intuitively appealing, theories of the mind must eventually be substantiated by empirical evidence. Experimental support for various hypotheses, particularly in the area of visual cognition, has been provided either directly by Leyton's own research or via re-interpretations of previously existing results. But several others await further explication and verification. At the end of the book, the student of cognition is faced with a decision – to expand on Leyton's work or to attempt a completely different story of the mind. Each choice comes with its pitfalls. The former, however tempting, could lead to a dead end while the latter could result in the reinvention of the wheel.

An original piece of work such as this, especially one that attempts an explanation of everything stands in danger of being completely ignored. SCM, since it first appeared in 1992, has contributed significantly to fields as diverse as aeronautical engineering, computer-aided design (CAD), control theory, meteorology, linguistics, music, neuronal growth models, robotics and even computer vision. However there have been few, if any, attempts to apply the ideas put forth in SCM to the study of cognition. This is especially intriguing as SCM was written with a Cognitive Science (CogSci) audience in mind. Although one might speculate that SCM may have been "ahead of its time", the CogSci tradition of borrowing from disciplines such as Logic, Information Theory, Probability Theory and Graph Theory rules out this possibility. Whatever the cause of the anomaly, SCM has much to offer in the way of a *theory* of mental organization. CogSci as a field is largely characterized by the particular phenomena discovered during the course of empirical investigation. This is rather like characterizing Gravitational Physics as the study of "falling objects" or the study of "tides". While the above may have been a necessary step in the evolution of Physics as a well as CogSci, the ideas in SCM are rich enough to motivate a more "top down" approach to the field, grouping phenomena by the organizational principles they reflect rather than the particular kinds of stimuli used in experimental tasks. For instance the concept of "symmetry" in its most general form may be used to discover the subjective differences in the way experts and novices represent stimuli belonging to a particular domain.

Consider the domain of Chess, and the ability of expert chess players to recall chess patterns far more accurately than novices. One could ask questions such as "What are the symmetries that the expert has extracted over time?", "How does the expert make use of these symmetries to represent chess patterns, i.e., recover their causal history?", "Does this explain the experts' superior recall?", "Is it possible to "recall" a stimulus pattern without being aware of the symmetries of the corresponding domain?" and so on. Indeed SCM does seem to suggest that what the novice perceives as "noise" may, to the expert, appear to reflect a high degree of order. Infants' preference for regular polygons could be explained by the innate availability of simple 2-dimensional symmetries and the resulting ease of representation afforded by these symmetries. The enormous versatility of SCM makes it an invaluable resource for researchers in all areas of Cognitive Science. While the application of the notion of "symmetry" to a domain such as "plan recognition" may not be straightforward, it is by no means an insurmountable task and the benefits of such a course of study are likely to be great. The reader would do well to pick up a copy of Symmetry, Causality, Mind, for its scope is really only limited by the use to which it is put.