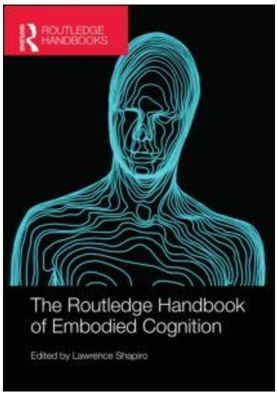

The Routledge Handbook of Embodied Cognition

Lawrence Shapiro

The Embodied Dynamics of Problem Solving

Publication details
https://www.routledgehandbooks.com/doi/10.4324/9781315775845.ch16
James A. Dixon, Damian G. Kelty-Stephen, Jason Anastas
Published online on: 01 May 2014

16

THE EMBODIED DYNAMICS OF PROBLEM SOLVING

New structure from multiscale interactions

James A. Dixon, Damian G. Kelty-Stephen, and Jason Anastas

Problem solving has been a rich area for psychology. Nearly all major cognitive phenomena occur in the context of problem solving, including memory search (Hélie and Sun, 2010), analogical mapping (Gick and Holyoak, 1980), reasoning (Thibodeau and Boroditsky, 2011), association (Grabner *et al.*, 2009), abductive inference (Langley, Laird, and Rogers, 2009), and priming (Slepian, Weisbuch, Rutchick, Newman, and Ambady, 2010), as well as many social phenomena, such as group facilitation (Liker and Bókony, 2009) and inhibition (Diehl and Stroebe, 1987) effects. In a sense, problem solving provides us with a microcosm of the key issue for psychology, "How does an organism successfully engage in goal-directed action?" Indeed, if we had a fully worked out theory of problem solving, it would provide a foundation on which to build theories of cognition. Currently, however, psychology and cognitive science place little emphasis on problem solving as a major area in the field, although posing problems to organisms, typically humans, is the stock-in-trade of experimental psychologists. Most studies in psychology pose a problem to their participants, in the form of task instructions and constraints, with the hope that participants will solve the problem using the processes the researcher wishes to study. Missing from current practice is a theory (or even a reasonable hypothesis) about how an organism could configure itself into some novel spatio-temporal structure in a way that accomplishes a novel task. Of course, we recognize that experimental tasks might be considered graded in terms of their novelty at a particular moment in time. For example, recalling a list of words might be analogous to recalling items on a shopping list and thus weakly novel. Naming the color of ink in which words are written, while trying not to read the words (the classic "Stroop" task), might be considered somewhat more strongly novel. At the far end of this continuum, might be tasks that are explicitly intended to require an innovative solution, such as the mutilated checkerboard problem (Kaplan and Simon, 1990) or traveling salesman problem (Applegate, Bixby, Chvatal, and Cook, 2011).

Early giants in the field, such as Köhler (1947) and Duncker (1945), appreciated that an organism's ability to generate innovative solutions in the face of novel problems was a deep theoretical issue. Indeed, as we will argue below, taking this aspect of problem solving seriously provides a powerful constraint on how we should conceive of cognition, in general. Consider, for example, how an organism could know, in advance of actually doing it, that configuring itself in a particular way would allow it to obtain a desired goal. Organisms from all five kingdoms are capable of this, but for now let us take the familiar example of a primate staring at a banana that is suspended from the ceiling, well out of reach. After trying out a few behaviors, such as jumping and climbing on objects, the primate stacks some nearby boxes, creating a platform from which he or she can reach the banana (Köhler, 1956). Let us stipulate that, because this primate has been reared in controlled conditions, we can be quite certain of the novelty of the problem and the solution. The primate in this example has generated a new structure, the complex behavior of stacking the boxes and climbing upon them, in the service of a goal. The question is: how could a theory of cognition explain such an event?

First, we consider the ability of computationalism-representationalism to address novel cognitive structures. We then discuss embodied cognition as an emerging position that may provide an account of novelty in cognition. Finally, we illustrate how such an approach might begin to be fleshed out using examples for our own recent work.

Computation over representations

Clearly, the dominant theoretical framework for cognition assumes that cognition is the processing of "information," and that information processing must be computations over representations (Fodor, 1981; Newell, 1994). The assumptions of this framework are so deeply embedded in how cognitive scientists and psychologists currently think about cognition that it often appears that they are unaware that computationalism-representationalism (henceforth, CR) is actually a theoretical position, not a self-evident fact. At the risk of treading too-well-worn paths, recall that CR takes the strong position that representations are internal states of arbitrary form that stand in some relation to external events in the world, and that these states are operated on by a set of rules. The output of this process yields another symbol (or symbols) the meaning of which is relevant to the organism, e.g. "It's a frog," "Scratch your left ear," "She's not listening," etc. It is important to point out here that if one takes CR as a theoretical approach to explaining cognition, then the whole theory has to be fleshed out in representations and computations. For example, one cannot insert an intelligent agent at the back end of the process that interprets the symbol, provides the necessary background knowledge, decides what actions to take, or otherwise provides the intelligent, goal-directed behavior we set out to explain (Dennett, 1981).

One of the many limitations of CR is that the symbols and the rules are fixed. That is, there's no natural way for the system to generate new symbols or new rules. No matter how many times a rule is invoked or a symbol used as input or output, novel rules and representations do not occur. This point is easy to appreciate, if one imagines the steps taken in the operation of a Turing machine: advance tape, read symbol, write symbol, advance tape. These operations are all there is to this form of computation. The machine can only read-write symbols it already knows and only execute rules it already has stored. Indeed, it is reasonable to think of computation as a highly efficient method of expressing what is already known (and encoded into the computing system).

Two major escape hatches might save CR from this seeming incompatibility with the facts of cognition. First, one might propose that the system actually contains all the rules it will ever use, but that they unfold over the life of the organism through some unspecified developmental process. Novelty here is just a mirage. The system already knows everything there is to know, but rolls out that knowledge when the time is right. This would clearly be a miracle of very high order. But even if one accepts that biology and evolution might be capable of pulling off something so extraordinary, it seems obvious that this explanation only works for objects and situations with an evolutionary history. Humans, for example, both create and adapt to cultural artifacts (e.g. golf clubs, cars, keyboards), objects that have no history on an evolutionary timescale. Thus, our abilities to configure ourselves into a biomechanically efficient striker of a tiny white ball or a non-self-propelled controller of a complicated high-speed machine are clearly not going to be explainable this way.

The second way one might attempt to save CR from the conundrum of novelty is to suggest that the computational system essentially operates on itself (Goldberg and Holland, 1988). This account proposes that there exists a class of symbols, call them second-order symbols, which stand in relation to the system's own internal rules or parts of those rules. For example, a rule such as "if *C* is true, then *P*" might be represented by two second-order symbols, *X* which would stand for phrase "if *C* is true," and *Y* which would stand for "then *P*." These second-order symbols would then be manipulated by their own set of rules. The key role of these rules would be to recombine the second-order symbols into new combinations, thus allowing the system to invent new rules. Thus, the system might make a combination that specifies, "if *C* is true, then *Q*." While such a modification of CR is plausible, it is hopelessly underpowered. The size of the set of possible combinations is just the product of the number of unique phrases of each type. The second-order rules generate novel first-order rules, but from a limited and completely prescribed set.

Note that this approach cannot be expanded to generate new symbols (i.e. representations) that stand in relation to the objects, events, etc., in the environment. The reason is easy to appreciate. CR systems are only in touch with the environment through their current set of symbols. All elements in the environment that are not encoded by those symbols are undetectable for that system. They literally have no means of entry. So, in the absence of a special front-end, an explicitly non-CR front-end that knows about the environment in some other way, CR systems cannot generate new symbols. This point begins to sail close to the now standard (and quite devastating) critique of CR that emphasizes its failures in the realm of semantics (Bickhard and Terveen, 1996). While we appreciate the force of those arguments, here we have focused briefly on issues that bear more especially on the generation of new structure.

Embodied cognition

Unlike CR, embodied cognition does not yet have a rigorous definition, and thus evaluating its potential requires some speculation about what exactly it entails. At a minimum, embodied cognition proposes that the states and actions of the body play a causal role in processes that are traditionally considered cognitive.

Embodied cognition has the potential to offer a serious alternative to CR. However, realizing this potential requires a radical rethinking of the nature of cognition. If embodied cognition is to move the field beyond its current state, it must embrace the deep implications of cognition being a physical (i.e. embodied) system. Attempting to slide embodiment beneath CR, as a grounding for representations, will simply inherit all the problems of CR while resolving none of them.

What then are the deep implications of cognition being a physical system? We propose the following fundamental properties of biological systems as an initial starting place for understanding cognition.

Dissipative structures

All organisms, including humans, are ensembles of dissipative structures (Kondepudi, 2012). Dissipative structures are self-organizing systems that form and persist in response to the flow of energy and matter. That is, they emerge to degrade energy and their survival continually requires energy degradation (Kondepudi, 2008). That biological entities are dissipative structures is non-controversial, although the theory has been primarily developed in physico-chemical systems. Importantly, dissipative structures have a number of properties that help get a theory of cognition off the ground. First, they occur spontaneously in living and non-living systems, so their origin is cashed out from first principles, that is, thermodynamic laws. One need not license their existence through appeals to special processes or hidden agents. Second, dissipative structures are exquisitely sensitive to changes in their environments. Any change in circumstance that changes the flow through the structure can have consequences for the whole structure, including generating new structures. Thus, explaining how a system can rapidly and seamlessly adjust to changes in the environment becomes a question of tracking how exogenous energy fluctuations impact the flow through the system. Inherently intractable issues such as "encoding," "sensory integration," and the like are rendered nugatory in this framework. Third, dissipative structures exhibit a rudimentary form of end-directedness; they form such that their spatio-temporal properties degrade energy and, perhaps, maximize the rate of entropy production. Thus, the question of how inert matter becomes active is answered directly by the theory of dissipative structures.

Nested, multiscale systems

The ensembles of dissipative structures that constitute organisms are organized across a wide range of spatio-temporal scales. These scales are nested such that smaller-scale structures comprise larger scales. In complex biological systems, the number of scales is not known and likely changes over time. Both the spatial structures (i.e. morphology) and temporal structures (i.e. behaviors) that are typically the focus of research in biology and psychology are embedded in the activity of the other scales of the system. Embedded here means that there are causal dependencies amongst structures at different scales. A class of formalisms has been developed to express these multiscale relations in complex systems in physics and related fields (Lovejoy and Schertzer, 2013). It is important to keep in mind that our identification of the behavior (or aspect of morphology) in which we are interested, say the discovery of a new problem-solving strategy, is the application of a particular measurement frame on a continuous and extremely heterogeneous system. The act of measuring the system tempts us not only to reify the behavior, but also to treat it as if it were now a component part in a mechanical system. In psychology and cognitive science, this usually involves casting the behavior as a player on the stage of the general linear model, in which exogenous causes push and pull on non-causal outcome variables. However, the proper metaphysics for understanding biological behavior is dictated by the multiscale, nested nature of the system that produced the behavior, not the convenience and familiarity of the metaphysics of machines.

Commerce in energy and matter

Effects in physical systems, including those within biology, must be understood in terms of the flow of energy and matter. If we are to have an embodied cognition, then it too will require that interactions are the sole result of energy consumption, and thus the quantities we measure should relate directly to energy consumption. This provides a common currency across all scales of the system and environment, as well as explicit connections to thermodynamic laws. It is worth noting here that taking energy (or another thermodynamic variable) as a primary measure does not deny the ultimate goal of psychology and cognitive science to explain relatively macro-scale patterns of behavior. Rather, we propose that these patterns of behavior are the result of thermodynamic (i.e. energy-consuming) processes.

Rule induction in card sorting

Among the many amazing properties of human problem solvers is the seeming ability to detect regularities embedded in sequences of temporal events. From our theoretical perspective, when any organism detects a regularity (i.e. a "rule"), it has reorganized itself such that it is now attuned to a new pattern of variation in the environment. This is a substantial and non-trivial event for the system, and thus we should be able to see evidence for it in the multiscale behavior of the system.

To test this overarching hypothesis, we asked college-age participants to perform a reduced version of the Wisconsin Card Sort Test (WCST) (Grant and Berg, 1948). Participants were asked to sort cards into four categories, each defined by a guide card. After they placed each card, the experimenter provided them with feedback about whether or not that placement was correct. Participants had to induce the correct sorting rule for each of five trials. (In the classic WCST, but not here, the rule also changes within trial). An example of this simple set-up is shown in Figure 16.1.

As a comparison condition, we asked a second group of participants to sort cards according to the same rules, but in this condition the rules were stated explicitly. A new rule was given to the participants (as was the initial rule) at the start of each trial. Thus, this group sorted cards according to the same rules, but did not have to induce the rules. We call this the "explicit" condition to denote the explicit stating of the rules; and the set-up described above, the "induction" condition (see Anastas, Stephen, and Dixon [2011] for details).

For both conditions, we tracked the motion of the participant's dominant hand (i.e. the hand used to sort the cards) at a high sampling rate (60 hertz) and with considerable precision (on the millimeter scale). We used the time series of the motion data to quantify multiscale activity of the system.

A few words are probably in order here about why we think the motion of the hand tells us about activity of the cognitive system with respect to this task. Consider that inferring that a

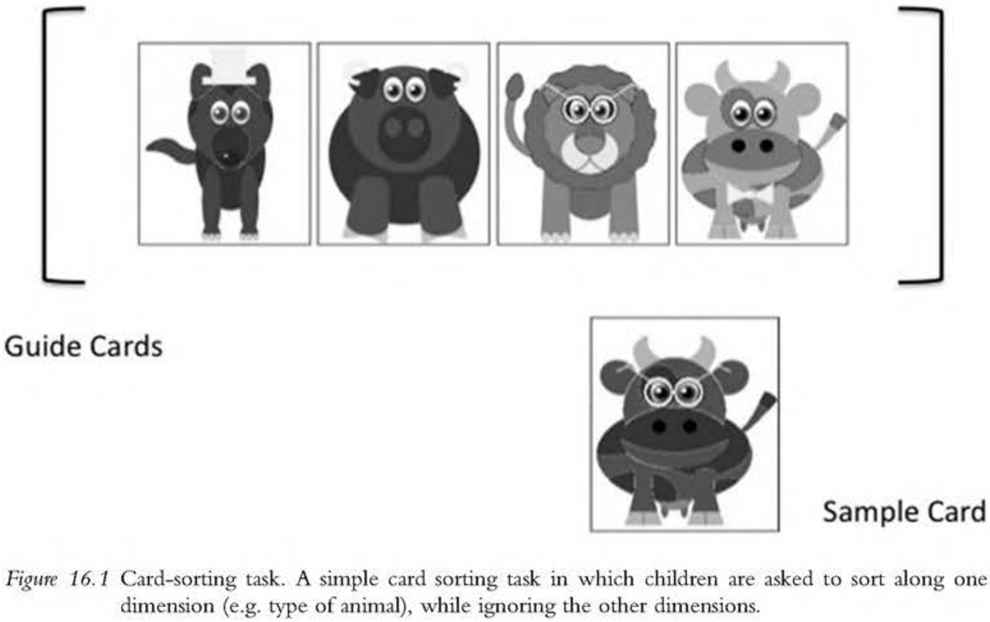

Figure 16.1 Card-sorting task. A simple card sorting task in which children are asked to sort along one dimension (e.g. type of animal), while ignoring the other dimensions.

participant has a "rule" involves applying a particular measurement frame to the pattern of motions he or she generates, relative to the problem context. Relatively macro-scale temporal relations in the participant's motions are taken as evidence of having acquired the rule. While it is tempting to talk about "rules" and other cognitive structures as if they existed as disembodied entities, their measurement in scientific practice requires an instantiation in the physical world. (One might also note that a commitment to materialism makes a parallel argument for the physical instantiation of rules.) If rules are then to be embodied entities that are produced by real, multiscale physical structures, they should show interesting micro-scale behavior, especially during their formation. Because hand motions are intrinsic to the instantiation of the rules in the present task, they should provide a quantitative window onto the relevant system activity.

A rejoinder to this argument, it should be noted, is that hand motions are run by the motor system, and that a different system is in charge of cognitive structures such as rules. This second system tells the motor system what to do, in general terms, and the motor system works out the details. Implicit in this description of two systems (rather than one) is the idea that motor processes and cognitive processes are insulated from each other. The systems communicate over some channel but do their work independently. This strong architectural assumption, an implication of Simon's near-decomposability thesis (Simon, 1962), is rarely tested, despite the fact that it is the essential licensure for all component approaches to cognition.

There are a few ways to empirically evaluate the component assumption. Some involve quite quantitatively advanced methods, usually in the realm of time-series analysis, such as iterative amplitude adjusted Fourier transform (Ihlen and Vereijken, 2010). Others are more pragmatic in that they simply ask whether the component that is proposed to be causally downstream (e.g. motor) is related to upstream components (e.g. cognitive) in unexpected ways. In a sense, most of the surprising effects in embodied cognition are of this latter type. We have examined both types of evidence for the card-sorting data. Although we focus on the latter type here, we note that more quantitatively advanced methods strongly suggest that card-sorting is not the product of a component-dominant process (Stephen, Anastas, and Dixon, 2012).

To explore the relationships between the microstructure of the motion of the hands and the phenomenon of induction, we briefly and non-technically consider a quantitative measure of multiscale activity, because some such measures are necessary to tell the rest of the story. The motion data obtained during each participant's card sorting contain a remarkable amount of information about the activity of the system during the task. One analysis of interest asks how the activity of the system is distributed over temporal scales. Physical systems that are composed of many interacting spatio-temporal scales will show a systematic relationship between temporal scale (i.e. the length of time over which the measurement is taken) and the degree of activity. We used detrended fluctuation analysis (DFA) (Peng *et al.*, 1994) to estimate the relationship between temporal scale and activity. This relationship is typically (and testably) a power law, quantified in DFA as a Hurst exponent (H). Values of H near 1 are consistent with long-term correlations in the time series, and have been taken as an index of the degree to which the system is "poised" to respond to changes. Not too long ago, researchers began to find $H \sim 1$ in a wide variety biological domains, suggesting that many processes in biology and psychology were not only poised to respond to changes, but appear to return repeatedly to that poised state (Bak, 1999). A related body of work shows that as a complex system approaches a structural or functional change, it goes to a poised (or "critical") state, showing properties much like a phase transition (Hinrichsen, 2006).

In the card-sorting task, the motion data from both the induction and explicit conditions had average Hurst exponents that suggested long-term correlations in the data. More importantly, we predicted that participants in the induction condition (who had to discover the rule) would

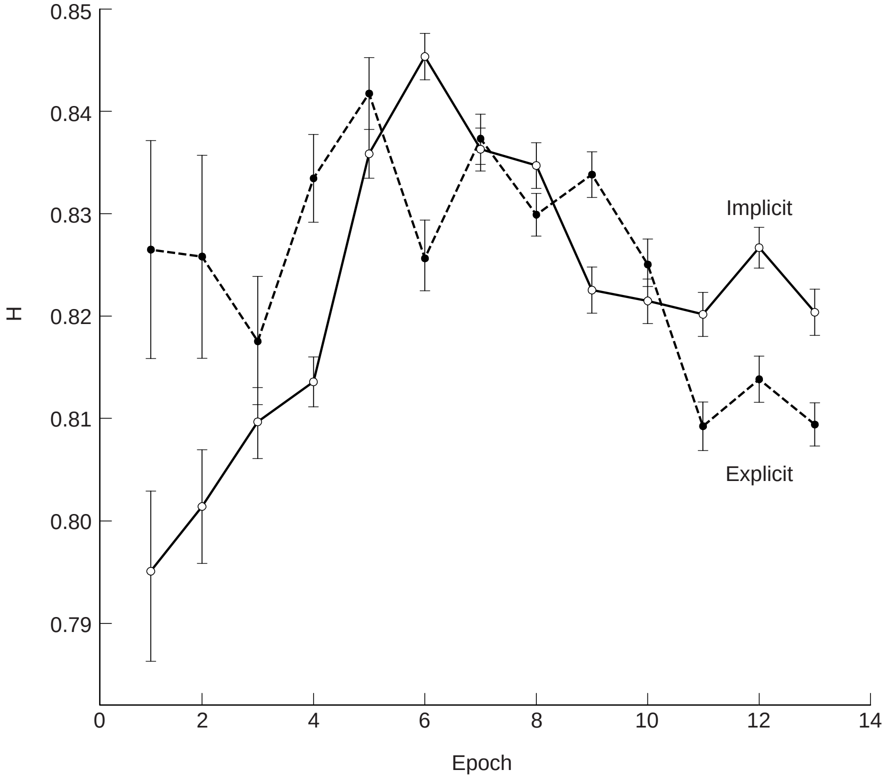

Figure 16.2 Hurst exponents over time during card sorting. The Hurst exponents (*H*) are shown as a function of epoch (time), with separate curves for the "implicit" and "explicit" rule conditions. Errors bars show 1 standard error of the mean.

show an pattern of increase and decrease in their Hurst exponents, indicative of the system becoming poised for a change (increasing *H*), and then settling into a stable regime (decreasing *H*). We also predicted that participants in the explicit condition (who were given the rules) would show a decrease in their Hurst exponents as they settled into a stable organization with the given rule. Figure 16.2 shows the average Hurst exponents for both conditions as a function of time (expressed as epochs; epochs are about 13 seconds in duration). Both these predictions were confirmed: the induction condition showed a strong peak in the Hurst exponent, and the explicit condition showed an overall decrease.

Note that in both these conditions participants were, from a motor perspective, doing much the same thing: sorting cards into four piles. The speed of their sorting motions, the duration of the trials, and a variety of other factors do not significantly differ between conditions, nor do they add predictive power to statistical models of the Hurst-exponent trajectories. The difference between the time series generated in these two conditions is not reducible to some macro-scale difference. Rather, the difference in the over-time trajectories of the Hurst exponents is consistent with the hypothesis that rule induction is an instance of self-organization in a multiscale system. The interactions across scales increase in the induction condition, allowing the system to reach a novel rule, and then decrease as the system stabilizes around that rule. Likewise, in the explicit condition, in which no discovery is necessary, the system just stabilizes as it repeatedly instantiates the rule. "Rule" here, we should emphasize, is a shorthand way of describing a particular, and quite complex, spatio-temporal organization of the system, rather than a simple conditional statement (e.g. "Put the card in the pile that matches on color"). While such a statement is enough to get a participant started sorting, his or her system still must organize to the many constraints of the task, such as spatial positions, variations across individual cards, etc. The rule that the system learns is the amalgamation of all the relevant task constraints.

Gear-system problems

While the card-sorting tasks have the virtue of being easy to manipulate and control, and thus amenable to experimentation, they are quite simple relative to many of the classic paradigms in problem solving. In other work, we have employed a more complex task in which participants spontaneously discover new, higher-order rules without being told that such a rule exists (Dixon and Bangert, 2002; Stephen, Dixon, and Isenhower, 2009). This paradigm has the virtue of being closer to the type of discovery that occurs in many real-world problem situations. In this task, participants are shown a set of interlocking gears. They are asked to predict the turning direction of the final gear in the series, the "target" gear, given the turning direction of the "driving" gear, the gear that provides the force to the system (see Figure 16.3). Participants are asked to solve a

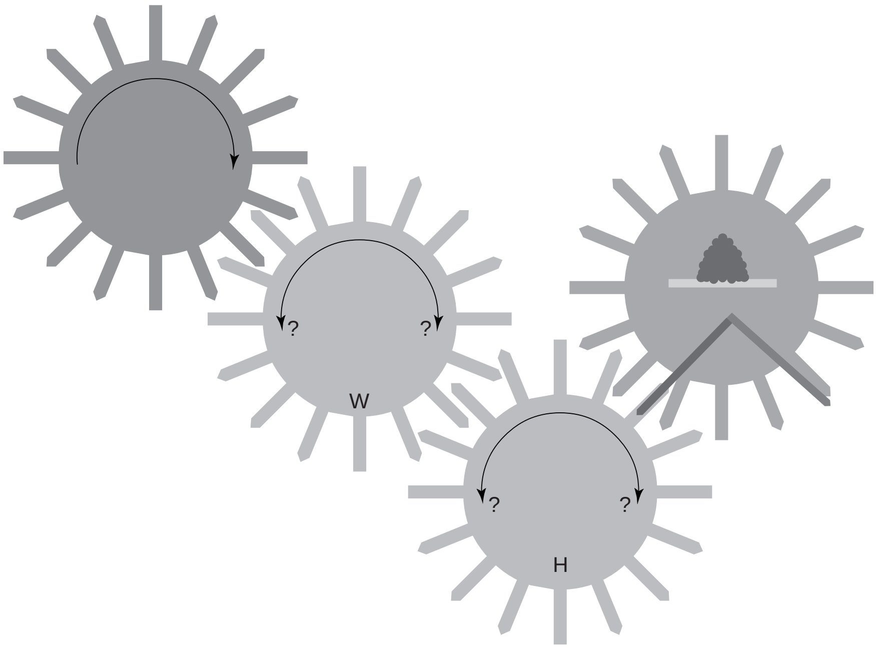

Figure 16.3 An example gear-system problem. Participants were asked to predict the turning direction of the target gear, on the far right, given the turning direction of the driving gear on the far left (clockwise in this example). The intermediate gears were labeled (here, "W" and "H") to facilitate participants talking about them. The task is embedded in a game about a train race. The object of each trial in the game is to position one's train beneath the target gear so as to catch the coal as it slides off. Note that only the target gear is ever seen to move; the other gears are occluded before the gears turn.

variety of problems of this type. The spatial configuration and number of gears in the system varies across problems. The vast majority of college-age participants initially solve this problem by manually simulating the movement of each gear in succession. That is, they trace their finger around the outer edge of each gear, starting with the driving gear, and ending at the target gear (which then gives the solution). We call this strategy "force tracing." The interesting phenomenon here is that, after using the force-tracing strategy some number of times, the majority of participants spontaneously discover (and employ) a higher-order relation to solve the gear problems. Specifically, they discover that adjacent gears alternate directions, and thus solve the problems by sequentially categorizing each gear as "clockwise" or "counterclockwise," a strategy we call "alternation."

Just as in the card-sorting task, we tracked the motion of each participant's hand while they solved the problems. For each participant, we analyzed the motion data on all the trials prior to their discovery of alternation (which were nearly all solved via force tracing). For participants who discovered alternation, we found that the power-law exponent increased to a peak, and then decreased immediately before discovery. Participants who did not discover alternation (i.e. kept using force tracing) showed a much shallower increase in the power-law exponent (Stephen and Dixon, 2009). Figure 16.4 shows the data averaged over participants. The top panel shows the power-law exponents aligned by trial number. As can be seen in the figure, the exponent increases more dramatically for participants who discovered alternation. The lower panel shows the power-law exponent for discoverers only, aligned on the far left on the trial in which they discovered alternation. The trial labeled "–1" is immediately prior to discovery, the trial labeled "–2" is two trials prior to discovery, etc. The figure shows that just prior to discovery there is a decrease in the power-law exponent. Both these effects were confirmed in growth curve analyses.

Conclusions

We have argued that the dominant approach to cognition, CR, cannot handle the phenomenon of new structure, a fundamental aspect of the problem-solving behavior exhibited by humans and other organisms. Embodied cognition has the potential to provide an account of how new structures emerge. The implications of this account could radically reshape how we understand cognition. We propose that such an account must be grounded in thermodynamic laws, while capitalizing on what is already known about how those laws manifest themselves in complex, heterogeneous systems. The work we have reviewed above takes some initial steps in this direction, addressing the phenomenon of new macroscopic cognitive structure as a phase transition. Phase transitions are well understood theoretically, and have been broadly demonstrated empirically in a wide variety of systems (e.g. Cortet, Chiffaudel, Daviaud, and Dubrulle, 2010). Our work shows that the formation of new cognitive structures shows the same signatures as in other embodied, physical systems. That is, an increase in fluctuations as the transition point approaches, indexed by an increase in H (and followed by a decrease in those fluctuations). While it is, of course, possible to create long-term correlated time series using linear methods (and thus consistent with the near-decomposability assumption), it is not clear how such an approach would explain the observed pattern of changes in H. We note that our work shows that these patterns can be experimentally manipulated, as well as observed in spontaneous behavior. Finally, we note that, in our view, embodied cognition is poised at a critical juncture. It can become a side show in the CR cognition circus, an exhibit of somewhat perplexing and bemusing phenomena that run counter to expectation. Or by embracing the deep implications of its commitment to the physical, it can become a truly new approach to cognition.

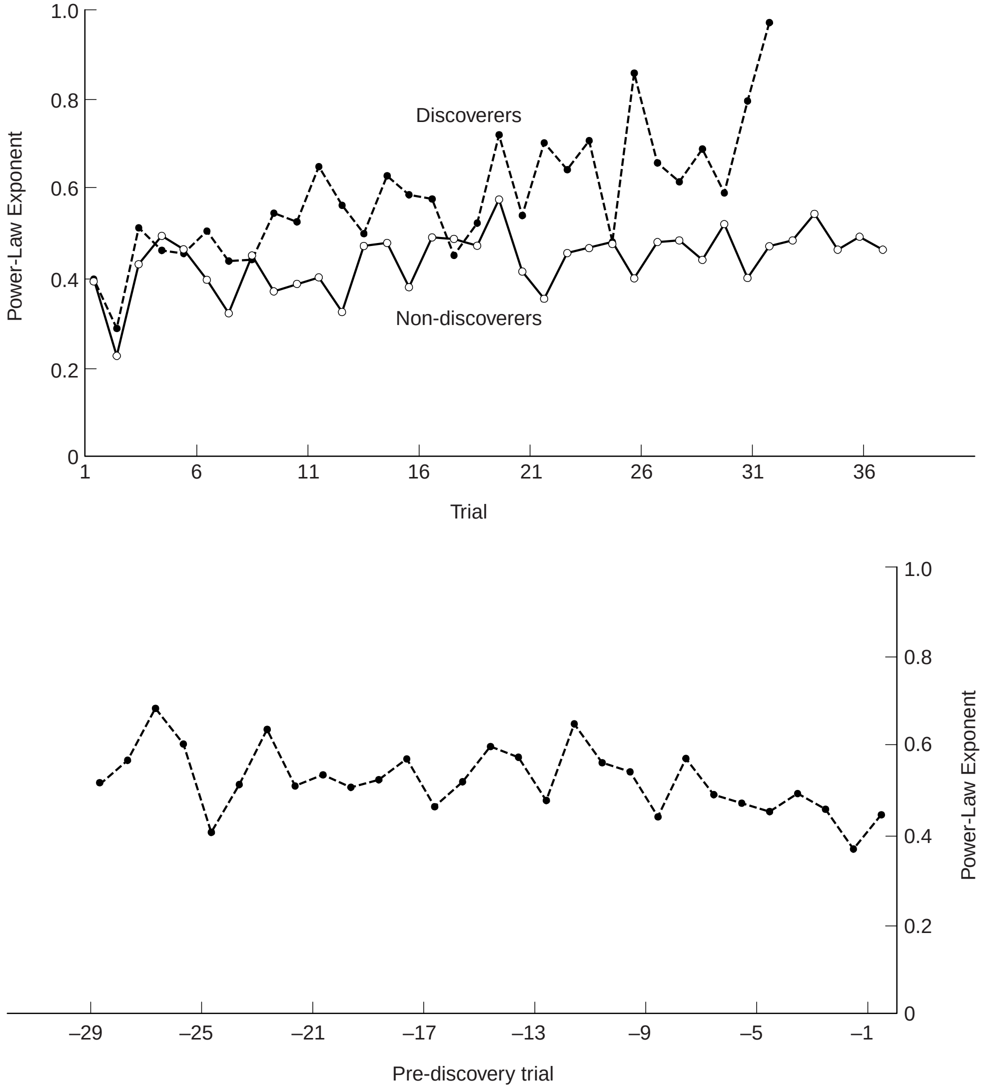

Figure 16.4 Power-law exponents over trials for the gear task. The top panel shows the power-law exponent as a function of trials for participants who discovered alternation, "Discoverers," and participants who did not discover alternation, "Non-discoverers". The lower panel again shows the power-law exponents of "Discoverers," but now the bottom-most point is aligned on the trial on which they discovered alternation. Thus, the lower panel shows the average power-law behavior as participants approach discovery.

References

Anastas, J. R., Stephen, D. G., and Dixon, J. A. (2011). The scaling behavior of hand motions reveals self-organization during an executive function task. *Physica A: Statistical Mechanics and its Applications, 390*(9), 1539–45.

Applegate, D. L., Bixby, R. E., Chvatal, V., and Cook, W. J. (2011). *The traveling salesman problem: A computational study*. Princeton: Princeton University Press.

Bak, P. (1999). *How nature works: The science of self-organized criticality*. New York: Springer.

Bickhard, M. H., and Terveen, L. (1996). *Foundational issues in artificial intelligence and cognitive science: Impasse and solution*. Amsterdam: Elsevier.

Cortet, P.-P., Chiffaudel, A., Daviaud, F., and Dubrulle, B. (2010). Experimental evidence of a phase transition in a closed turbulent flow. *Physical Review Letters, 105*, 214501.

Dennett, D. C. (1981). *Brainstorms: Philosophical essays on mind and psychology*. Repr. ed. Cambridge, MA: MIT Press.

Diehl, M., and Stroebe, W. (1987). Productivity loss in brainstorming groups: Toward the solution of a riddle. *Journal of Personality and Social Psychology, 53*(3), 497–509.

Dixon, J. A., and Bangert, A. S. (2002). The prehistory of discovery: Precursors of representational change in solving gear system problems. *Developmental Psychology, 38*(6), 918–33.

Duncker, K. (1945). On problem-solving (L. S. Lees, Trans.) (Special issue). *Psychological Monographs, 58*(270).

Fodor, J. A. (1981). *The mind-body problem*. New York: W. H. Freeman.

Gick, M. L., and Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology, 12*(3), 306–55.

Goldberg, D. E., and Holland, J. H. (1988). Genetic algorithms and machine learning. *Machine Learning, 3*(2), 95–99.

Grabner, R. H., Ansari, D., Koschutnig, K., Reishofer, G., Ebner, F., and Neuper, C. (2009). To retrieve or to calculate? Left angular gyrus mediates the retrieval of arithmetic facts during problem solving. *Neuropsychologia, 47*(2), 604–8.

Grant, D. A., and Berg, E. (1948). A behavioral analysis of degree of reinforcement and ease of shifting to new responses in a Weigl-type card-sorting problem. *Journal of Experimental Psychology, 38*(4), 404–11.

Hélie, S., and Sun, R. (2010). Incubation, insight, and creative problem solving: A unified theory and a connectionist model. *Psychological Review, 117*(3), 994–1024.

Hinrichsen, H. (2006). Non-equilibrium phase transitions. *Physica A: Statistical Mechanics and Its Applications, 369*(1), 1–28.

Ihlen, E. A. F., and Vereijken, B. (2010). Interaction-dominant dynamics in human cognition: Beyond $1/f$ α fluctuation. *Journal of Experimental Psychology: General, 139*(3), 436.

Kaplan, C. A., and Simon, H. A. (1990). In search of insight. *Cognitive Psychology, 22*(3), 374–419.

Köhler, W. (1947). *Gestalt psychology: An introduction to new concepts in modern psychology*. New York: Liveright Publishing.

——(1956). *The mentality of apes*. New York: Vintage Books.

Kondepudi, D. (2008). *Introduction to modern thermodynamics*. Chichester, West Sussex: Wiley.

——(2012). Self-organization, entropy production, and physical intelligence. *Ecological Psychology, 24*(1), 33–45.

Langley, P., Laird, J. E., and Rogers, S. (2009). Cognitive architectures: Research issues and challenges. *Cognitive Systems Research, 10*(2), 141–60.

Liker, A., and Bókony, V. (2009). Larger groups are more successful in innovative problem solving in house sparrows. *Proceedings of the National Academy of Sciences of the United States of America, 106*(19), 7893–98.

Lovejoy, S., and Schertzer, D. (2013). *The Weather and Climate: Emergent laws and multifractal cascades*. Cambridge: Cambridge University Press.

Newell, A. (1994). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Peng, C. K., Buldyrev, S.V., Havlin, S., Simons, M., Stanley, H. E., and Goldberger, A. L. (1994). Mosaic organization of DNA nucleotides. *Physical Review E, 49*(2), 1685–89.

Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society, 106* (6), 467–82.

Slepian, M. L., Weisbuch, M., Rutchick, A. M., Newman, L. S., and Ambady, N. (2010). Shedding light on insight: Priming bright ideas. *Journal of Experimental Social Psychology, 46*(4), 696–700.

Stephen, D. G., Anastas, J. R., and Dixon, J. A. (2012). Scaling in cognitive performance reflects multiplicative multifractal cascade dynamics. *Frontiers in Physiology, 3*(102).

Stephen, D. G., and Dixon, J. A. (2009). The self-organization of insight: Entropy and power laws in problem solving. *Journal of Problem Solving, 2*(1), 72–101.

Stephen, D. G., Dixon, J. A., and Isenhower, R. W. (2009). Dynamics of representational change: Entropy, action, and cognition. *Journal of Experimental Psychology: Human Perception and Performance, 35*(6), 1811–32.

Thibodeau, P. H., and Boroditsky, L. (2011). Metaphors we think with: The role of metaphor in reasoning. *PLoS One, 6*(2), e16782.