Cognitive Mechanics: Natural Intelligence Beyond Biology and Computation

Jon Doyle
Department of Computer Science
North Carolina State University
Box 8206, Raleigh, NC 27965-8206, USA
`Jon_Doyle@ncsu.edu`

Abstract

Framing traditional disputes between computational and biological approaches to understanding cognition in terms of a divide between natural and engineering inspirations artificially restricts the notion of nature to biology. One need not restrict natural inspiration to biology, however. One can find common ground between the two approaches in mechanics, which underlies and provides important extensions to both conceptions.

Introduction

The computational conception of cognition has come in for criticism on several grounds, including undesirable brittleness in capacities; the large efforts needed to manually formalize and represent human knowledge; the nonuniformity evident in needing different learning methods for each different representation and task; and a lack of direct connection with neurophysiological realizations (though some count this as a strength rather than as a weakness). These weaknesses of the computational conception are often contrasted with the strengths of biological conceptions, which in the popular artificial neural network frameworks have boasted of robust behaviors insensitive to small perturbations in situation or mental state; relatively universal and uniform learning methods, at least in comparison to computational conceptions; and structures that fit well with both animal physiology and evolutionary theories.

Critics of biological conceptions of cognition in turn have pointed out that unstructured neural conceptions lack ready interpretation or explanation in linguistic terms; that uniform learning methods sometimes require long training times in comparison with specific symbolic methods; and that universality of learning methods is trivial to obtain. These weaknesses of biological conceptions are in turn contrasted with notable strengths of computational conceptions, namely expressive and independently intelligible representations of statements in natural and formal human languages; ready provision of explanations of mental states and their changes; straightforward cognitive computations that appear close to some introspective accounts of human experience; and direct provision for innate abilities and structure.

I believe that some of the respective strengths of the biological and computational approaches come from the same natural source, namely mechanics, and that one can seek fruitful combinations of the two and add further strengths by looking to concepts of mechanics that have not played much of a role in either approach to date.

Cognitive Mechanics

The brief position statement offered here draws on *Extending Mechanics to Minds* (Doyle 2006), which applies modern axiomatic rational mechanics to psychology and argues that mechanical concepts, such as force, mass, and work, provide a useful vocabulary supplementing ones from computation and neurophysiology for describing and characterizing cognition and other mental processes.

Only a few of the axioms of modern rational mechanics (Noll 1958; 1973; Truesdell 1991) embed assumptions about the continuum nature of space and time. One can separate these continuum assumptions from the rest of the content to obtain axioms that characterize almost all of the familiar properties of mechanical concepts in a way that is independent of continuum assumptions. The resulting mechanics reproduces standard mechanics as a special case, in which the familiar continuum assumptions need apply only to physical space and time.

The broadened mechanics also applies to nontraditional mechanical systems that characterize time, space, mass, and force with discrete structures that are mathematically quite similar to the vector space structures used to characterize traditional mechanics. Some common formalizations of discrete cognitive architectures studied in artificial intelligence can be recast in these structures and observed to satisfy the mechanical axioms. Minds that satisfy the axioms of mechanics are mechanical systems, so use of mechanical concepts to describe such minds involves no use of metaphor.

Hybrid combinations of discrete and continuous mechanical systems, such as persons with bodies and minds that each form mechanical systems, can also form mechanical systems in the broadened mechanics.

Just as physical materials appear in different kinds with different characteristics, so also do mental materials making up different kinds of minds. In particular, cognitive architectures with different characteristics can correspond to different mechanical materials. Indeed, physical materials appear

---

in such variety that a particular mental material can be more similar to a particular physical material than that physical material is to some other physical material.

Transcending Rigid Kinematics

Development of Turing's (1936) mathematical model of an intelligent human mathematician engaged in calculation produced the modern theory of symbolic and numeric computation, in which discrete symbols and rules govern discrete movement through finite or infinite machine states. In developing its application to cognition, one can replace symbols with logical or quasi-logical statements, behavioral rules by inference rules, and machine configurations by memory lists and structures, and one can combine such discrete structures with finite numerical values to obtain Bayesian networks and update rules.

At its base, however, Turing's computational conception, like the earlier calculation-oriented efforts of Babbage and Lovelace, views thinking in terms of movements and configurations of clockwork-like machines. In traditional mechanical clocks, of course, gears shift from one configuration to another in discrete steps. Modern stored-program computers behave much the same, except that the configurations are those of discrete switching circuits. In clocks, as in computers, the motive power that produces these changes of configuration is irrelevant to the operation except when the power fails and the machine stops.

From a mechanical point of view, the computational approach is limited to properties of rigid kinematical systems. The discrete configurations correspond to symbols, and explicitly so in Newell and Simon's (1976) physical symbol system hypothesis about the nature of intelligence. This rigid kinematical approach contributes to the observed brittleness of computational approaches. In the mechanical world in which we live, most things deform a bit when subject to ordinary forces, and few things exhibit perfect rigidity. Turing insisted explicitly on something close to indeformability in his model of computation, basing the restriction to finite tape alphabets on the need to avoid one symbol shading off into another, just as a deformed "j" can resemble an "i".

In contrast, neural models of cognition following from the work of McCulloch and Pitts (1943) and von Neumann (1958) regard the imprecision of symbol boundaries as contributing to robustness. Even though one can sometimes read or design symbols into the structures and states of neural networks, as in (Touretzky and Hinton 1985), one usually expects the meaning of concepts to shift somewhat with changes in one's environment or activities. This sort of deformation is offered as a strength of neural net approaches.

Indeed, models of some types of learning explicitly involve mathematical structures found in the mechanics of deformable materials. Construction of support-vector machines, for example, involves computing a Gram matrix from a set of data points in the course of finding a separating embedding of the original data into a larger space. In mechanics, the intrinsic configuration of a body consists of the set of distances between each pair of body points. Deformations of bodies consist of changes in intrinsic configurations. The Gram matrix of a body of data is one representation of the intrinsic configuration of the data, so the support-vector technique amounts to finding a deformation of the body of data that separates the classes of interest.

Characterizing Realistic Dynamics

Many neural models are tied to a specific level of detail and do not in themselves support abstract and higher-level characterizations of mind. Modern mechanics, in contrast, gives central roles to constitutive assumptions that characterize properties, such as elasticity, incompressibility, polarization, and rigidity, that differentiate the variety of special types of materials and that are critical to understanding behaviors independent of the precise identity of the materials exhibiting them.

One example of an important mechanical property is mental inertia or resistance to change. Neurophysiological answers to questions about how minds change primarily involve schemes of neural parameter adjustments, and to a much smaller extent, ideas about growth or death of neurons and neural structures. Although such physiological changes might mediate mental changes, they do not answer questions about mental change any more than sets of character insertions and deletions answer questions about how Lincoln revised drafts of his Gettysburg address.

Even foregoing questions about motivations of change, neural biology does not provide good answers to questions about the effort, difficulty, or rate of change. These play major roles in understanding realistic notions of rationality. The idealized rational agent assumed in economics and epistemology exhibits perfect epistemic omniscience and consistency, and is capable of assimilating new information and making perfect decisions instantaneously. Realistic rationality, in contrast, suffers limitations: reasoning and deliberation require effort and concentration, actions must be taken despite persistent inconsistency and ignorance, and learning takes time and slows as habits accumulate.

Mechanics provides concepts for understanding these limitations in natural terms (Doyle 2009). Some limitations arise from the mental mass associated with some aspects of memory. All mass generates inertial forces that limit the motion resultant from a given force. Other limitations arise from bounds on the amount of work that can be performed in effecting change. Mechanics measures work in terms of the force applied across the distance traveled. This notion provides a natural measure of effort in mental mechanical systems as well, which in simple kinds of minds corresponds to the number of changes to memory made in the course of reasoning. Some limits on the rate of work stem from bounds on the magnitude of forces acting in a system. Natural materials generate forces bounded by the size and mass of bodies, and such limitations are evident in realistic mental materials as well, in which reasoning rules or other mechanisms change bounded numbers of conclusions and in which sensor dimensionality and Shannon channel capacity limit the magnitude of forces.

Some mental materials exhibit forms of constitutional rigidity, such as cognitive architectures in which some types of conclusions are always drawn automatically regardless of the rules or other knowledge possessed by the reasoner. Elastic and refractory forces are exhibited by other mental materials. In reasoners making use of nonmonotonic default rules and dependency-based revision and restoration operations, for example, the "rest configuration" of defeated default assumptions can be restored on the removal of the assumption defeaters, and the existence of alternate support for conclusions can resist removal of the conclusion in the first place. The common multiplicity of conclusion sets generated by multiple nonmonotonic defaults means that removal of deforming defeaters can let the reasoner rebound to mental states other than those existing prior to imposition of the defeaters. The elasticity characteristic of default reasoning is thus also often accompanied by some amount of plasticity. Rigidity and elasticity can aid in reasoning, but can also impede it. In particular, when reasoning habits generate distracting conclusions that divert the reasoner from its goals, focusing attention on those goals requires generation of forces to counteract the distractions. In such cases, both the creation and removal of the distraction involve work on the part of the reasoner.

Conclusion

One can trace some of the comparative strengths and weaknesses of computational and biological approaches in artificial intelligence to the provision for or lack of important characteristics of mechanical systems. As one might expect from the kinematic appearance of computation, some kinematic characteristics of materials are easier to identify in computational approaches than in superficial biological models, and account for some of the advantages of computational relative to biological approaches. Similarly, some continuum characteristics of materials are more easily identified in neurophysiological models, and account for some of the advantages of biological over computational approaches.

Mechanics provides concepts including mass, force, work, and elasticity that go beyond the standard concepts provided by current computational and biological approaches. Augmenting computational and biological analyses with mechanical concepts promises to illuminate better the full richness of mental materials and phenomena and so improve our understanding of mind. Fulfilling this promise will likely require progress on developing the mathematical analysis of discrete and hybrid mechanics in ways that parallel past progress on the mathematical analysis of continuum physics.

Acknowledgments

The author's preparation of this paper was aided by a SAS Institute Professorship endowment.

References

Doyle, J. 2006. *Extending Mechanics to Minds: The Mechanical Foundations of Psychology and Economics*. London, UK: Cambridge University Press.

Doyle, J. 2009. Mechanics and mental change. In *Evolution of Semantic Systems*, Lecture Notes in Computer Science. Springer-Verlag. To appear.

McCulloch, W. S., and Pitts, W. 1943. A logical calculus of the ideas immanent in nervous activity. *Bull. Math. Biophys.* 5:115–133.

Newell, A., and Simon, H. A. 1976. Computer science as empirical inquiry: Symbols and search. *Communications of the ACM* 19:113–126.

Noll, W. 1958. The foundations of classical mechanics in the light of recent advances in continuum mechanics. In L. Henkin, P. S., and Tarski, A., eds., *The axiomatic method, with special reference to geometry and physics. Proceedings of an International Symposium held at the University of California, Berkeley, December 26, 1957–January 4, 1958*, Studies in Logic and the Foundations of Mathematics, 266–281. Amsterdam: North-Holland. Reprinted in (Noll 1974).

Noll, W. 1973. Lectures on the foundations of continuum mechanics and thermodynamics. *Archive for Rational Mechanics and Analysis* 52:62–92. Reprinted in (Noll 1974).

Noll, W. 1974. *The Foundations of Mechanics and Thermodynamics: Selected Papers*. Berlin: Springer-Verlag.

Touretzky, D. S., and Hinton, G. E. 1985. Symbols among the neurons: Details of a connectionist inference architecture. In *IJCAI*, 238–243.

Truesdell, III, C. A. 1991. *A First Course in Rational Continuum Mechanics*, volume 1: General Concepts. New York: Academic Press, second edition.

Turing, A. M. 1936. On computable numbers with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society Series 2* 42:230–265.

von Neumann, J. 1958. *The computer and the brain*. Silliman Memorial Lectures, Vol. 36. New Haven, Conn: Yale University Press.