Mind, a Machine?

Peter Cariani • Harvard Medical School, USA • cariani/at/mac.com

> Upshot • Written by recognized experts in their fields, the book is a set of essays that deals with the influences of early cybernetics, computational theory, artificial intelligence, and connectionist networks on the historical development of computational-representational theories of cognition. In this review, I question the relevance of computability arguments and Jonasian phenomenology, which has been extensively invoked in recent discussions of autopoiesis and Ashby's homeostats. Although the book deals only indirectly with constructivist approaches to cognition, it is useful reading for those interested in machine-based models of mind.

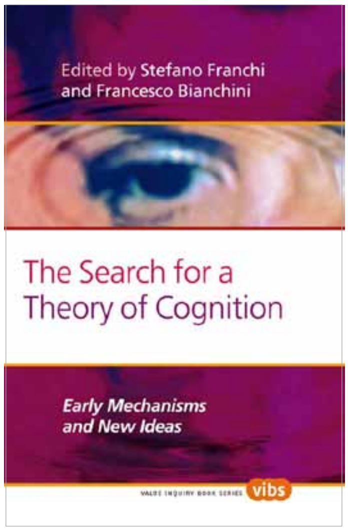

Review of *The Search for a Theory of Cognition: Early Mechanisms and New Ideas* edited by Stefano Franchi and Francesco Bianchini.
Rodopi, Amsterdam, 2011.
ISBN 978-90-420-3427-3. 375 pages.

How we think is the ultimate subject of this book, as seen through the history of cognitive science, artificial intelligence (AI), cybernetics, Saussurian linguistics, mathematics, and logic. The main focus is on the intellectual history of cognitive science, and on the role of cybernetics and symbolic AI on its evolution.

Like Stefano Franchi's previous edited volume (Franchi & Güzeldere 2005), this book will be of interest to historians of science and technology, philosophers, and cognitive scientists, i.e., anyone interested in contemporary reappraisals of cognitive science. The essays within the book are thoughtful, well-researched, and well-written by recognized experts in their respective areas. Each essay contains a good deal of fresh historical observations and conceptual content. The writing styles are generally straightforward, with a relatively little postmodern posturing.

The book's contents concentrate on the influences of cybernetic adaptive systems and digital programmable computers on theories of cognition. The cybernetics of World War II, with its impressive analog servomechanisms and automated anticipatory fire control systems provided the first promises of electronic brains. In the postwar era, Ashby's homeostat and Grey Walter's autonomous robotic turtle captured the public imagination. For the next three decades, computational cognitive science, whose model of mind was fashioned in the image of the digital electronic computer, dominated technologically-inspired approaches to cognitive theory. At the end of its reign, computational cognitive science was succeeded by the connectionist revival of neural networks in the mid-1980s with the renewed promise of brain-like electronic devices.

Rather than empirical predictive models that come from studies of brain and behavior, these engineering approaches to brain and mind provided strategic thinking about how brains might work, as well as constructive demonstrations of proofs-in-principle. Mind engineering attempted to show by fabrication and demonstration how systems with particular functional organizations might be capable of mental functions such as perception, coordinated action, decision-making, anticipation, goal-seeking, calculation, and reasoning. Cybernetics, artificial intelligence, and neural networks successively created impressive information-processing artifacts that could, by resemblances to animal behavior and deliberative processes, plausibly claim to be models of brain and mind.

Psychogeography of cognitive theory

The editors' introduction, "On the historical dynamics of cognitive science: a view from the periphery," sets a postmodern, mildly oppositional tone, with the book's twelve chapters grouped into four sections: "The cybernetic suburb," "AI's peripheries," "Margins of computation," and "At the threshold of computability." The core of this opposition is "computationalist-representationalist" cognitive science, whose disciplines included symbolic AI, linguistics, analytical philosophy, computer science, and theoretical cognitive psychology, and whose paradigmatic hegemony lasted roughly from the mid-1950s to the mid-1980s. The margins, defined by psychogeographic distance from MIT computationalist orthodoxy, include cybernetics, alternative conceptions of computation, non-Chomskian linguistics, and the cognitive sciences in Europe and the Eastern Block.

The first half of the book deals with cybernetics, AI, and Cognitive Science, with essays on Ashby's homeostat adaptive device (Stefano Franchi), Wiener's wartime attempts to formulate a cybernetic theory of the mind (Peter Galison), homeostat- and computer-inspired models and simulations of mental processes (Peter Asaro), and postwar histories of cybernetics and artificial intelligence in Italy (Claudio Pogliano) and the Soviet Union (Slava Gerovitch). The second half of the book deals with theories of computation, with essays on Chinese room arguments and Turing tests (Francesco Bianchini), formal automata theory and regular expressions (Christopher Kelty), and computability- and quantum-inspired extensions to classical computation (Maurizio Matteuzzi, Soloman Feferman, Rossella Lupacchini, Giorgio Sandri).

Left out of the dualistic core-periphery construction and most of the essays, are other, natural science components of the mind sciences that operated in parallel with engineering approaches. These include cognitive psychology, developmental psychology, psycholinguistics, cognitive neuroscience, and cognitive neurology. These parts of the mind sciences have been somewhat more resistant to full-blown machine metaphors because of their focus on specific, observable aspects of cognition in human subjects. The two thrusts, experimental engineering and natural science, need each other. The natural sciences need engineering visions of how brains might work, and mind engineering needs to be firmly grounded in the details of what biological brains can do and how they are put together. Biological brains and natural minds are the touchstones not only because these are by far the most capable informational engines that we know, but also because these are the only systems that we know with certainty can support conscious awareness. Looking to the future, a new theory of cognition should integrate the two, by searching for general engineering principles that are inspired by the observations and models of the behavioral and brain sciences.

Cognition as computation

Although many of the essays do, as the book's title suggests, effectively mix analysis of history and anticipations of the future, readers should not expect wholesale systematic critiques or programmatic prescriptive alternatives (e.g., Bickhard & Terveen 1995 and Dreyfus & Dreyfus 1988). Some general, foundational criticisms do come up in the book, but mainly implicitly, amidst historical description. The purpose of remembering the past is to anticipate (and make) the future. A general survey of what went wrong with the computationalist project is therefore entirely in order in the search for a new theory of cognition. What follows below is my own list of what went astray (see also Cariani 2011, 1989).

A common pitfall of theoretical programs involves the successive isolation of theory from empirical observation, beginning with the redefinition of what counts as data, and culminating in what counts as an explanation. A healthy research program stays in touch with and grounded by its empirical data, and in the face of shortcomings of theory, revises theory, not data. Willingness to reject beliefs in the face of contrary evidence is what distinguishes critical from dogmatic thought.

A major inherent, flawed dynamic of computationalist approaches to cognition was the tendency to redefine all problems in terms of discrete, symbolic computations, to truncate problems, Procrustean-style, to fit the logic-calculating capabilities of deterministic automata. In the case of Chomskian linguistics, language was reduced from an all-purpose representational system for communicating semantic meanings and pragmatic intents between human actors to purely syntactic, idealized, production grammars generated by formal automata. As problems become increasingly virtualized, computational theories explain more and more about the capabilities and limitation of pure computation, and less and less about how our own, brain-produced minds work.

The rise of the digital computer as the dominant metaphor for mind and the dominant research program for artificial models of mind marked a fundamental change in philosophy from the Aristotelian, materially-embedded, analog, interactionist, goal-directed, adaptive, self-organizing systems of cybernetics to the Platonic, digital, symbolic, logical, virtual-world representational systems of symbolic AI. The ascendant computationalist paradigm redefined problems of mind away from perception, action, motivation, affect, learning, and concept formation in real-world contexts to intelligence, reasoning, and knowledge representation in virtual worlds. The behavioristic Turing-test imitation games that were proposed as evaluations of success, discussed by Asaro and Bianchini, proved to be theoretical dead ends, symptoms of a field that has given up trying to understand fundamental requisites for cognition and moved on to work on technical problems of increasingly narrower scope.

The pervasive redefinition of fundamental problems was coupled with a profound neglect of empirical evidence and insights from biology, neuroscience, and psychology related to how biological brains are thought to work. Rather than attempting to simulate biological brains using computational mechanisms that would in some, even distant, sense parallel neural processes or distributed neural networks, say along the lines of David Marr's computational neuroscience or Douglas Hofstadter's sub-symbolic computations, computational cognitive science opted for black box, rule-based, symbol-processing models that had little direct connection to either brains or concrete behaviors.

In computationalist theory, connection to the external environment via sensing (perception) and effecting (action), i.e., the essence of being "situated" and "embodied," was replaced with abstract truth-values in possible-world scenarios. The making of meaning, which involves construction of internal, anticipatory-predictive models from the structure of sensorimotor experiences, was reduced to logical inferences on truth-values. Cybernetic notions of agency, i.e., autonomous action directed at pursuing embedded purposes, were similarly eliminated from consideration. The goal-seeking, purposive, steering mechanisms and the adaptive, learning processes that they drive were replaced by attempts to codify knowledge through clever programming. Communication in service of cooperative coordination of behavior was replaced with formal grammars.

In the end, the computationalist program attempted (and failed) to reduce semantic meanings and pragmatic purposes to purely syntactic, computational operations. In my opinion, we are still digging ourselves out from under the train wreck of these systemic platonic distortions that have plagued philosophy and the mind sciences from the 1930s onwards.

The world did get something for money spent on AI research, although it was not the theory of mind that was promised would usher in a new era of intelligent machines. As Kelty concludes in his perceptive chapter on regular expressions, although artificial intelligence failed as cognitive theory, it did develop new and powerful tools for computer science that have enabled the miracles of the large-scale informational networks that we enjoy today.

Computational extensions

Still somewhat under the spell of platonism are considerations of Gödelian-Turing computability issues and alternative theories of computation in the book's last segment. These are coherently discussed by the authors, who are experts in these fields. For this reason alone, these essays will be of interest to many readers. But I must confess that I cannot see how these alternatives bear any relevance whatsoever to cognitive theories.

I have recently argued in the pages of this journal (Cariani 2012b) that computability issues have no significance whatsoever for finite automata such as digital electronic computers. These issues only come up for potentially-infinite systems, such as arithmetic on the natural numbers. There is no Gödelian undecidability for strictly finite formal systems and no halting problem for Turing machines with limited tapes. The full set of natural numbers cannot even be physically realized in a finite universe.

Likewise, if the brain is a finite automaton of some sort, computability is similarly irrelevant, and the Lucas-Penrose arguments are rendered meaningless. If the brain has discrete functional states, being a physical system, it has but a finite number of them. The functional states of the brain, to the very limited extent that we understand them, appear to be based on mass-statistics of spiking activity in neural populations that have little resemblance to the static, discrete states of Turing machines or electronic digital computers. Memory in the brain is not indefinitely extendable, as it is in imaginary Turing machines.

All of the proposals to augment the realm of Turing computation (discussed in Matteuzzi's and Lupacchini's respective chapters) involve adding state-transition indeterminacies or making the machine's behavior contingent on the outputs of measuring devices. Turing called these additions "oracles" (Turing 1939), but such oracle machines are no longer carrying out reliable, formal operations that one can call "computations" in anything faintly resembling the original, 1930s sense of the term. The resulting augmented machines are no longer deterministic and no longer functionally equivalent to formal systems.

It is also trivially true that robotic systems equipped with sensors and effectors in addition to computational organs do carry out functions that are qualitatively different from those of a Turing machine (Cariani 1989, 2011). In contrast to disembodied and isolated Turing machines, biological brains and artificial robots can sense the world and act on it to influence its course of events.

Symbolic AI vs. cybernetics

More could and should be said about the history of the split between symbolic AI and cybernetics in a book such as this (only briefly mentioned in the editors' introduction). Unfortunately, the transition has often been treated as an evolutionary transformation of one "mechanization of mind" technoscience program rather a struggle between two competing, antithetical paradigms (Dupuy 2000 is frequently guilty of this).

Symbolic AI was born in 1956 at a conference at Dartmouth College, where the universalistic, platonic, core tenets of evangelical, high-church computationalism were codified. What was to become the party line and movement raison d'être is the (specious) assertion that everything can be fully specified and simulated on a digital computer and therefore purely computational approaches can subsume all others, rendering them superfluous. It was here that many distinctions were buried: analog vs. digital (the subject of Asaro's contribution), hardware vs. software, sensing vs. effecting vs. computing, and learning vs. direct specification.

By the end of the 1960s, amidst funding contractions associated with the Vietnam War, advocates of symbolic AI gained hegemony over government funding mechanisms and proceeded to defund all existing competing paradigms: cybernetics, trainable machines and neural networks, bionics, autonomous robotics, systems theory, and self-organizing systems. As the symbolic AI research program ran dry in the early 1980s with brittle failures of expert systems, openings in funding space permitted research to be revived in massively parallel computers, connectionist neural networks, and autonomous robotics. Cybernetics, systems theory, self-organizing systems, bionics, and evolutionary robotics (Harvey et al. 2005) are in various stages of full reincarnation, many of them under the rubric of artificial life.

At present, no one artificial device is universally held up as a guide for brain and mind. If pressed for an answer, most cognitive scientists would say that the brain is best described in terms of some sort of neural network. As Maniglier perceptively points out in his essay on linguistic theory, the combinatorics of mental attributes, concepts, and linguistic constructs are the main arguments in favor of Fodorian, symbol-based, computational-representational systems. Connectionism has never been able to provide a satisfying explanation for how its networks could realize these operations without long and extensive training. In my opinion, connectionism is the beginning and not the end-stage of neural network research. Looking toward the future, we need new kinds of more flexible, temporally-coded neural networks to tame these combinatorics (Cariani 2012a).

Constructivism and the mind

Alas, there is relatively little in the book that deals directly with constructivism, let alone radical constructivism. As Jean Piaget had said, "The mind organizes the world by organizing itself." Consideration of how the mind constructs itself in the context of its external interactions should be an essential part of any search for a new cognitive theory.

Part of the reason has to do with the book's main subject, the notion of mind as computer. The notion of a limited, situated observer adaptively modifying internal structures in order to cope with external challenges is fundamentally alien to the image of a computer carrying out logical operations on abstract truth-values in virtual worlds. Another part of the reason is that the cybernetics under discussion as a model for mind and brain is mostly the early, first-order cybernetics of Norbert Wiener and Ross Ashby, as exemplified by feedback fire-control systems and the homeostat, whose concern was primarily feedback control and not a theory of the internal organization that makes autonomous observer-actors possible. To be fair, the cyberneticians of the first wave were very well aware of the epistemic implications of self-constructing systems, and arguably, the rudiments of second-order cybernetics are present in the homeostat's self-modifying feedback architecture.

Second-order cybernetics, which blossomed in the aftermath of the split with AI, deals explicitly with the observer-actor and his/her self-modifications in the course of interactions with an external environment. When self-modifying observer-actors interact with others like themselves in shared spaces, the coevolution of concepts amongst the observer-actors can bootstrap communication, coordination, and cooperation. Constructivist approaches include Piagetian developmental psychology, Pask's Conversation Theory, and evolutionary robotics, as well as the origin and evolution of shared grammars and meanings in language (e.g., Steels 2006; Tomasello 2008). Second-order cybernetics does come up in Franchi's essay in the context of the homeostat, autopoiesis, and evolutionary robotics, but its positive, constructivist contributions to a new science of cognition get lost amidst phenomenological complaints (see below).

Maniglier's wide-ranging foray through the generative semiology of language touches on some of these issues, mainly in the context of language acquisition and change. Although rough-going in sections, it is the one essay in the book that engages core theoretical issues in cognitive psychology and psycholinguistics. The discussion oscillates between two poles: the debates between classical symbolic computation and subsymbolic connectionism on one hand, and Saussurian semiology on the other. Maniglier takes issue with Saussure's focus on how language works at any given moment at the expense of attention to the processes by which meanings form and change in response to social interactions. Generally and metaphorically speaking, he seeks to reintroduce "history," "culture," and "society" into semiology, linguistics, and artificial intelligence. As mentioned above, some of these directions are being actively pursued in the contemporary, artificial-life, evolutionary-robotics demonstrations of the origins of cooperation and language communication in public, shared-attention spaces.

Phenomenology and the homeostat

Lastly, there is Franchi's account of the life, death, and resurrection of Ashby's homeostat. Having been engaged in auditory neuroscience for the last two decades, and blissfully unaware of the philosophical revival of the device by Di Paulo, Harvey, Varela, and Weber, this thoughtful and provocative article came as a pleasant surprise for me, despite my profound disagreements with its conclusions.

Franchi accurately describes the workings of the homeostat, delves extensively into the concepts of homeostasis of Bernard, Haldane, and Cannon that inspired it, provides new historical insights into its historical demise, and effectively chronicles its revival.

The homeostat was a network of four selectable, interconnected circuits that demonstrated the property of ultrastability – almost whatever the form of an external disturbance, the homeostat could eventually find an internal configuration that would allow it to settle into a stable, homeostatic state (see Latil 1956 and Ashby 1960 for contemporary descriptions). I believe that Ashby designed the device (uniselector switch and feedback control circuits with arbitrary, ill-defined, and/or unknown parameters) as a concrete demonstration of his proposed principle, the Law of Requisite Variety (Cariani 2009). Ultrastability was the proximal goal of Ashby's device, but a homeostat could be coupled to a set of evaluative signals in order to constantly strive towards some optimal or satisficing state.

Roughly a decade ago, the homeostat was resurrected from obscurity by Inman Harvey and Ezequiel Di Paulo of the evolutionary robotics group at Sussex University and also by Francesco Varela and Andreas Weber. In discussing its recent reappearance, Franchi criticizes cybernetic conceptions of purposive action that involve goal-directed, feedback-controlled, steering mechanisms, as are found in steam governors, thermostats, and torpedoes. Franchi takes his arguments directly from Hans Jonas, a very influential Heideggerian phenomenologist-ethicist who equates consciousness with life.

In his book *Phenomenon of Life*, Jonas posits that in order to "have purpose," one must first have conscious feelings. Jonas says

> "The answer can be compressed into one statement: living things are creatures of need. Only living things have needs and act on needs… The cybernetical model reduces animal nature to the two terms of sentience and motility, while in fact it is constituted by the triad of perception, motility, and emotion." (Jonas 1966: 126)

Because servomechanisms are not conventionally thought to be either living or conscious entities with feelings, Jonas takes this to mean that they cannot have embedded purposes.

These arguments are very odd, to say the least, and the evidence from neuroscience, neurology, and biology does not appear to support them. First, there are states of general anesthesia in which conscious awareness is abolished while the anesthetized subject remains very much alive. Second, there are behavioral states called "automatisms," of which sleepwalking is an example, in which subjects act purposefully, but without conscious awareness of their acts. Third, why is the phenomenal aspect crucial for emotion, but not also for perception? Would we argue that a machine vision system cannot "see" because seeing requires the conscious experience of seeing?

Jonas' position is clearer in his diatribe against epiphenomenalism (Jonas 1966: 127), which he characterizes as the view that "certain animal actions may be, and probably are, accompanied by states of awareness, but these are scientifically irrelevant under the axiom that external actions can be explained on external, i.e., physical terms alone." Although Jonas explicitly disavows Cartesian dualism of separate but interacting substances, he seems to be endorsing a mind-over-matter view in which physical laws are insufficient, without phenomenal causes, to explain physical actions.

The scientific study of consciousness is devoted to understanding the relation between physical processes, such as neural activity, and conscious awareness. The working assumptions in this field, which are derived from many observations and experimental studies are that:

1 | all phenomenal states and distinctions are produced by particular, appropriately organized patterns of neural activity, and
2 | every phenomenal, supraliminal, experienced difference reflects some difference in underlying neural activity patterns, but not all differences in neural activity produce supraliminal changes in experience.

It has also been found that characteristic patterns of neural activity, such as "readiness potentials," reliably precede conscious percepts, subjectively spontaneous decisions, and actions by hundreds of milliseconds (Libet 2004). What we consciously experience appears to depend on earlier neural activity. Although such potentials do not rule out mind-over-matter causation, they do not appear to support any notion of an autonomous, conscious will that acts independently of neural activity to direct it.

It is perhaps understandable that Jonas and his Heideggerian insistence on the primacy of conscious being might appeal to the mystically-inclined Varela (Weber 2001; Varela, Thompson & Rosch 1991). However, reasons completely escape me why whole enclaves of researchers, including the evolutionary roboticist Ezequiel Di Paulo, should be falling over themselves to incorporate Jonas' apparently incoherent philosophy of life into their theories. Franchi criticizes Di Paulo's valiant, but perhaps impossible, attempt to reconcile Jonasian metaphysics with Varelan autopoiesis (Di Paolo 2005). He judges the interpretations of Di Paolo and Jonas to be incompatible, and accordingly rejects Di Paolo for Jonas.

One has the impression that Jonas is not the right philosopher for the job. Maybe Susanne Langer's (1967) theory would be a better choice in this direction of a naturalized phenomenology of life. Even better would be to scrap Jonas entirely, get out of the exegesis business, and attempt to articulate a coherent phenomenology of cybernetics that provides accounts of purposiveness both in terms of cybernetic organization and internal feeling.

This move is not entirely new. Although early cybernetic theories were not specifically formulated as solutions to the mind-body problem, cybernetic accounts of conscious awareness based on feedback loops have been proposed in the past (Powers 1973; Pask 1979). Norbert Wiener, Warren McCulloch, and Jerry Lettvin often referred to Gottfried Wilhelm Leibniz as a spiritual forerunner of cybernetics (e.g., Wiener 1948; McCulloch 1989; Lettvin 1989), and it is not difficult to imagine Leibnizian monads subserved by neural networks and circular causal loops. (As an MIT freshman, I was a student of Lettvin, who often spoke passionately of Leibnizian monads and automatic navigation steering mechanisms for ships that Leibniz had envisioned.) One can also conceive of cybernetic theories of awareness based on autopoiesis-like, self-productive regenerations of circulating sets of mutually-stabilizing neural signals in reentrant neural loops (Cariani 2000).

A cybernetic panpsychist theory might go further to suggest that every purposive, goal-seeking system has inherent physical, organizational, and phenomenal dimensions. Franchi approves of Michel Cabanac's interesting proposals along these lines because he judges them consistent with Jonas. However, it is not clear to me how such a panpsychist theory would handle automatisms (purposive behavior without awareness should not be possible). Before attempting to project awareness onto other kinds of living and nonliving systems, any viable general theory of consciousness first needs to predict correctly those (neural) conditions under which we humans have awareness or lack it.

Depending on one's intents, it may make sense to discuss a purposive system either in terms of its physics (material composition of parts, and actions of mechanisms), its functional organization (embedded goals and steering processes), and/or its putative phenomenal aspects (what it might be like, experientially, to be the system). Thus one way forward is to recognize that there can be multiple, complementary descriptions that explain different aspects of the system at hand, such that no one realm is held to be primary (a position adopted by Langer and rejected by Jonas).

Although I have long admired the Heideggerian-inspired, substantive critiques of symbolic AI that Hubert Dreyfus raised, it was always puzzling to me why this perspective generated so little in the way of concrete alternative programs for research. On the scene in addition to the Jonasian "consciousness = life," one also finds a parallel slogan, "cognition = life" (Stewart 1996), and one wonders if this merging of concepts of autopoietic self-production and informational process also has its origins in Heideggerian thinking. Although I agree with Stewart's trenchant critique of computationalism and his anti-realist adoption of radical constructivist epistemology, I resist conflating life, cognition, meaning, and purposiveness with consciousness so that the ontological primacy of conscious awareness (Being) can be asserted. From where I stand, the world does not seem to fit a mind-over-matter perspective at all. Instead, first-hand awareness, functional organization, and material process appear to me to be complementary aspects of one world, in which material process produces specific types of organization, some of which in turn produce conscious awareness and its contents.

At the end of Franchi's Jonasian polemics, I was left with the distinct feeling that Heideggerian metaphysics itself may be hostile to any explanation of agency that could ever arise from the natural sciences, even one grounded in the neural basis of purposive action in conscious humans and animals. It raises larger questions of the compatibility of Heideggerian ontology with the anti-metaphysical, epistemological stance of radical constructivism. No doubt the question has been debated at many other places and times, and in any case, it is a longer discussion for another day.

Conclusion

The Search for a Theory of Cognition fits in well with other intellectual histories and critiques of cybernetics, AI, and cognitive science such as Boden (2006), McCorduck (1972), Dreyfus (1979), Dreyfus & Dreyfus (1988), Pickering (2010), Latil (1956), Bickhard & Terveen (1995), Heims (1991), Conway & Siegelman (2005), Dupuy (2000) and Stewart (1996). These works all support the intellectual archeology needed for assessing where cognitive science has been and where it needs to go from here. All in all, we think that the editors, Franchi and Bianchini, have achieved what they set out to do with this book.

References

Ashby W. R. (1960) Design for a brain. Second Edition. Chapman and Hall, London.
Bickhard M. H. & Terveen L. (1995) Foundational issues in artificial intelligence and cognitive science: Impasse and solution. Elsevier, New York.
Boden M. A. (2006) Mind as machine: A history of cognitive science. Oxford University Press, Oxford.
Cariani P. (1989) On the design of devices with emergent semantic functions. Ph.D Thesis, State University of New York at Binghamton, Binghamton, New York.
Cariani P. (2000) Regenerative process in life and mind. In: Chandler J. L. R. & Van de Vijver G. (eds) Closure: Emergent organizations and their dynamics. Annals of the New York Academy of Sciences, New York: 26–34.
Cariani P. (2009) The homeostat as embodiment of adaptive control. Special issue on W. Ross Ashby, edited by Klir G. & Asaro P. International Journal of General Systems 38: 139–154.
Cariani P. (2011) The semiotics of cybernetic percept-action systems. International Journal of Signs and Semiotic Systems 1(1): 1–17.
Cariani P. (2012a) Creating new primitives in minds and machines. In: McCormack J. & D'Inverno M. (eds) Computers and creativity. Springer, New York: 395–430.
Cariani P. (2012b) Infinity and the observer: Radical constructivism and the foundations of mathematics. Constructivist Foundations 7(2): 116–125. Available at http://www.univie.ac.at/constructivism/journal/7/2/116.cariani.
Conway F. & Siegelman J. (2005) Dark hero of the information age: In search of Norbert Wiener, the father of cybernetics. Basic Books, New York.
Di Paolo E. (2005) Autopoiesis, adaptivity, teleology, agency. Phenomenology and the Cognitive Sciences 4(4): 429–452.
Dreyfus H. L. (1979) What computers can't do. Harper & Row, New York.
Dreyfus H. L. & Dreyfus S. E. (1988) Making a mind versus modeling the brain: Artificial intelligence at a branchpoint. Daedalus 117(1): 15–43.
Dupuy J. P. (2000) The mechanization of the mind: On the origins of cognitive science. Princeton University Press, Princeton.
Franchi S. & Güzeldere G. (2005) Mechanical bodies, computational minds: Artificial intelligence from automata to cyborgs? MIT Press, Cambridge MA.
Harvey I., Di Paolo E., Wood R. & Quinn M. (2005) Evolutionary robotics: A new scientific tool for studying cognition. Artificial Life 11(1/2): 79–98.
Heims S. J. (1991) The cybernetics group. MIT Press, Cambridge MA.
Jonas H. (1966) The phenomenon of life: Toward a philosophical biology. Harper & Row, New York.
Langer S. K. (1967) Mind: An essay on human feeling. Johns Hopkins Press, Baltimore.
Latil P. de (1956) Thinking by machine. Houghton Mifflin, Boston.
Lettvin J. (1989) Introduction to volume I. In: McCulloch R. (ed.) Collected works of Warren McCulloch. Intersystems Press, Salinas CA: 7–20.
Libet B. (2004) Mind time. The temporal factor in consciousness. Perspectives in cognitive neuroscience. Harvard University Press, Cambridge MA.
McCorduck P. (1972) Machines who think. W. H. Freeman, San Francisco.
McCulloch W. S. (1989) Reflections on the many sources of cybernetics. In: McCulloch R. (ed.) Collected works of Warren McCulloch, 1989. Intersystems Press, Salinas CA: 21–49. Originally appeared in 1974 in ASC Forum 26(22): 25–16.
Pask G. (1979) Consciousness. Journal of Cybernetics 9: 211–258.
Pickering A. (2010) The cybernetic brain: Sketches of another future. University of Chicago Press, Chicago.
Powers W. (1973) Behavior: The control of perception. Aldine, New York.
Steels L. (2006) Experiments on the emergence of human language. Trends in Cognitive Sciences 10(8): 347–349.
Stewart J. (1996) Cognition = life: Implications for higher-level cognition. Behavioural Processes 35: 311–326.
Tomasello M. (2008) Origins of human communication. MIT Press, Cambridge MA.
Turing A. (1939) Systems of logic based on ordinals. Proceedings of the London Mathematical Society 45(2): 161–228.
Varela F., Thompson E. & Rosch E. (1991) The embodied mind. MIT Press, Cambridge MA.
Weber A. (2001) The "surplus of meaning." Biosemiotic aspects in Francisco J. Varela's philosophy of cognition. Cybernetics and Human Knowing 9(2): 11–29.
Wiener N. (1948) Cybernetics, or control and communication in the animal and in the machine. Wiley, New York.

Peter Cariani considers himself a philosophically-inclined natural scientist who has trained and worked in both "wet" and "dry" sciences. He has published works on neural coding of auditory and musical qualities (pitch, timbre, consonance, rhythm), temporal coding of sensory information, cybernetics of self-organizing percept-action systems, functional emergence and creativity, semiotics, and the neural basis of consciousness.

Received: 5 July 2012
Accepted: 9 July 2012