Minds & Machines (2009) 19:453–463
DOI 10.1007/s11023-009-9169-z

Life and Mind

Margaret A. Boden

Published online: 9 December 2009

Abstract It's sometimes said, and even more often assumed, that life is necessary for mind. If so, and if A-Life promises to throw light on the nature of life as such, then A-Life is in principle highly relevant to the philosophy of mind and cognitive science. However, very few philosophers have attempted to argue for the relation between life and mind. It's usually taken for granted. Even those (mostly in the Continental tradition, including some with a following in A-Life) who have insisted on the linkage have stated it rather than justified it. If an evolutionary account of intentionality is acceptable, then perhaps biological life 'makes room' for mind. But that claim is problematic, since it's not clear that the type of self-organization involved in life-as-such must necessarily include evolution. Even if it does, it's a further step to show that life is strictly necessary for mind.

Keywords Life · Self-organization · Evolution · Intentionality · A-Life

Introduction

When Alan Turing (1950) claimed, in the august pages of Mind, that there's no good reason to deny that computers might be able to think, philosophers were quick to disagree.

Most of them focussed on the concept of thought. Wolfe Mays (1952), for instance, penned the argument that John Searle (1980) would later express as "all syntax and no semantics", adding for good measure that thought implies consciousness (an addition that led Gilbert Ryle to refuse to publish Mays' piece

M. A. Boden (✉)
Cognitive Science University of Sussex, East Sussex, UK
e-mail: M.A.Boden@sussex.ac.uk

in Mind: W. Mays p.c.). But a few appealed also to the concept of life. That is, they relied on the intuition that life is necessary for mind.

Michael Scriven, for example, supplemented the claim that thought implies consciousness, by declaring that "Life is itself a necessary condition of consciousness", and that "Robots… are composed only of mechanical and electrical parts, and cannot be alive" (Scriven 1953, p. 233). However, he didn't say why something made of mechanical and electrical parts can't be alive, whereas something made of biochemical parts can (but see the discussion of metabolism in Section "What is Life?").

Nearly twenty years later, by which time the notion that computers might be capable of thought was no longer surprising (though many people found it no less shocking), Peter Geach said something very similar. AI systems, he insisted, can't have beliefs and intentions because they're "certainly not alive" (Geach 1980, p. 81).

However, neither Geach nor Scriven tried to explain why life is necessary for thought, or for consciousness. And only Scriven offered a reason why robots "cannot" be alive (namely, that they're made of mechanical and electrical parts). As for Geach, his confident "certainly not alive" wasn't glossed in any way: he deemed it obvious not only that mind depends on life but that computers aren't living things. Their reticence wasn't unusual. The same two claims are made fairly often, but with scant argumentation attached.

An apparent exception, if one gets no further than reading the titles in a bibliography, is Hilary Putnam's paper on "Robots: Machines or Artificially Created Life?" (1964). Indeed, the question of the possibility of living robots was here 'twinned' with that of the possibility of robot minds, since Putnam was responding to Paul Ziff's (1959) essay on "The Feelings of Robots".

However, Putnam's paper focussed mainly on consciousness, not life. At one point, he endorsed Ziff's claim that it's an "undoubted fact" that if a robot isn't alive then it can't be conscious. But he ascribed this truth to "the semantical rules of our language", not to any quasi-explanatory relationship between life and mind. He also said (this time, disagreeing with Ziff) that something which is clearly a mechanism might be alive. Again, however, this was linguistic philosophy in action. Elsewhere, Putnam (1962) had heretically recommended changes in word-meaning due to new scientific data (about dreaming, for instance). But in the paper on robots and life, he relied on what current usage allowed one to say without contradiction. The nearest he got to discussing a substantive claim about life was to scorn the suggestion that the primary difference between a robot and a living organism is the "softness" or "hardness" of the body parts (1964, p. 691).

It's hardly surprising that Putnam's desultory discussion didn't lead his fellow philosophers to take an interest in the issue of life-and-mind. Most of them simply took it for granted that these two concepts, or phenomena, are somehow linked—which would imply that if computers aren't alive then they can't be psychological systems. (That's why, in their unquestioned assertions about the relations between life, mind, and computers, Scriven and Geach—and Ziff, too—evidently expected immediate assent.)

The same is true today. Analytic philosophers appear to think the life-mind linkage so obvious that, even when they bother to state it explicitly, they don't offer any arguments for it. Those few philosophers who have provided arguments are on, or near to, the Continental side of the intellectual fence. But even Continentals often state the linkage rather than challenging and justifying it (see Section "The Life-Mind Linkage Defended").

My purpose here isn't to overthrow the Geach-Scriven intuitions. For I, too, see computers/robots as quintessentially non-living things. And I, too, suspect—though with less confidence—that mind requires life. Rather, I want to examine the reasons for both of these commonly-held beliefs.

Quite apart from the desirability of philosophical hygiene ("No unexamined assumptions, please!"), this relates to the importance of A-Life, as contrasted with AI, for the philosophy of mind.

The oft-declared 'opposition' between AI and A-Life is in fact largely spurious (Boden 2006, chaps. 4.ix and 15). But there is a clear distinction of research focus. Whereas AI is the computer-based study of psychology (especially human psychology), A-Life concentrates on ethology and biology—including the nature of life as such. It follows that if the Geach-Scriven intuitions are both well-founded, then A-Life cannot generate mind but is, in principle, relevant for understanding it.

The Life-Mind Linkage Defended

Proponents of the life-mind linkage include the existentialist theologian Hans Jonas, and the neurophysiologists Humberto Maturana and Francisco Varela (see Boden 2006, 15.viii.b and 16.x.a and c). Unlike Scriven, Geach, Ziff, and even Putnam, they wrote about this matter at some length.

Jonas wasn't interested in biology for scientific purposes, but approached it in an ethical-theological spirit. In his view, orthodox (molecular and experimental) biology should be replaced by a biological science of a very different kind, because it illustrates the disastrous cultural influence of Cartesianism.

Descartes, he complained, had separated human beings from the rest of Nature, by means of a "spiritual denudation" of the non-human world which had stripped it of any intrinsic value (Jonas 1966, pp. 58–63, 232). Jonas believed that this had helped lead his ex-teacher Martin Heidegger toward Nazism, through attaching more importance to the fact that humans can make free decisions than to considering which values should guide our decisions (Jonas 1990, p. 200). Those values, he said, are themselves grounded in Nature. He meant not only that our values emerge as a result of Nature, but also that Nature is in itself valuable—and therefore worthy of respect (he later became a guru of the environmentalist movement: see Jonas 1984.)

More specifically, values are grounded in life. Embodiment, and in particular metabolism, was seen by Jonas as philosophically crucial (1966, pp. 64–91). So was evolution. Charles Darwin, he said, despite his materialist assumptions, had helped us to understand that all forms of mediation between organism and environment—perception, motor action, emotion, conscious imagination, and self-reflection—emerge as a result of evolution (1966, pp. 38–58). In general, life is essential for the emergence of mind (pp. 99–107).

Indeed, mind is present, or rather prefigured, in all of life. According to Jonas, all self-organized matter is, in a sense, ensouled: life involves "self-concern".

He expressed the life-mind linkage thus: "One way of interpreting [the ascending scale of life] is in terms of scope and distinctness of experience, of rising degrees of world perception…. Another way, concurrent with the grades of perception, is in terms of progressive freedom of action…. [The] 'mirroring' of the world becomes ever more distinct and self-rewarding, beginning with the most obscure sensation somewhere on the lowest rungs of animality, even with the most elementary stimulation of organic irritability as such, in which somehow already otherness, world, and object are germinally 'experienced', that is, made subjective, and responded to" (1966, p. 2).

Whereas freedom, for Descartes, was a God-like immaterial power, for Jonas it is founded in our biology—indeed, in "the blind automatism of the chemistry carried on in the depths of our bodies" (1966, p. 5). But that chemistry, he said, differs from the chemistry of "suns, planets, and atoms" in being embodied as metabolism. The "principle of freedom" common to all living organisms lies in their having a special type of identity and continuity: a stable dynamic form made of an ever-changing material substrate. This both enables and prefigures the human capacity for making decisions: "One expects to encounter [talk of freedom] in the area of mind and will, and not before: but if mind is prefigured in the organic from the beginning, then freedom is. And indeed our contention is that even metabolism, the basic level of all organic existence, exhibits it: that it is itself the first form of freedom" (1966, p. 3).

In short, "mind is prefigured in organic existence as such" (5). Life and mind are ontologically inseparable: "the organic even in its lowest forms prefigures mind, and… mind even on its highest reaches remains part of the organic" (1).

It's evident from these quotations that, besides trying to explain why all the minds we know about are found in living things, Jonas was trying also to say what life is.

The same applies to Maturana and Varela (1980), who defined life as "autopoiesis in the physical space". This concept is close to, though not identical with, metabolism (Boden 2000). Autopoiesis in general, they said, is the continuous self-production of an autonomous entity. The boundaries, components, and internal relations of "an autopoietic machine" (i.e. a living organism) are produced and maintained by a network of self-organizing processes (1980, p. 79). The system thus "pulls itself up by its own bootstraps and becomes distinct from its environment through its own dynamics, in such a way that both things are inseparable" (ibid.).

As biologists, Maturana and Varela had a much better grasp than Jonas did of the scientific issues involved (although their prose was even more convoluted and obscure than his). They were less concerned with ethics, and more with the adaptive functionality of organisms. (This isn't to say that they were functionalists: in theorizing the mind/brain, they explicitly rejected talk of input, output, computation, internal representations, and even feature-detectors–which had been discovered by a team including Maturana himself: Lettvin et al. 1959.) Accordingly, where Jonas had credited all living things with "freedom" and "self-concern", they credited them instead with "cognition"—calling their major book "Autopoiesis and Cognition".

All three terms, in my view, should be taken with a large pinch of salt when associated with the concept of life. There is indeed an important sense in which living things, at base thanks to metabolism, have a degree of autonomy. But to call this "freedom", or even a "prefiguration" of human free choice, demands a far more detailed justification than Jonas provided.

Similarly, living things are, indeed, pre-adapted to the specifics of their habitat, and capable (in varying degrees) of adapting to it appropriately when it changes. To call this "cognition", however, is to push the term too far. When an oak-tree loses its leaves in the autumn, this is an adaptive response, nicely adjusted to the environmental conditions. But knowledge (cognition), properly so called, it is not. What ethologists term innate releasing mechanisms, such as the hawk-like stimulus that prompts the newly-hatched grouse chick to crouch, are somewhat more persuasive. Even so, the longstanding philosophical debate about the coherence of the notion of "innate knowledge" (e.g. Edgley 1970) shows how problematic it is to ascribe knowledge to creatures (newborn babies, as well as grouse) lacking evidence, hypothesis, judgment, or even learning.

However, to say that a claim should be taken with a pinch of salt is not to throw it off the dining-table altogether. For one thing, linguistic usage can be extended in the light of scientific advance—as Putnam (1962) had rightly insisted. So terms such as "freedom" or "knowledge" might justifiably come to be ascribed way beyond their usual boundaries. (And recent empirical research, while showing that the seemingly innocuous term "innate" is in fact highly problematic, has provided ample data for various types of prefigurement in newborn animals: Elman et al. 1996.)

More to the point, for present purposes, an unconvincing argument is better than no argument at all. The three writers I've mentioned here did at least make a start in explaining the link between life and mind, instead of taking it for granted.

What is Life?

As we've seen, part of their project (besides establishing the life-mind link) was to say what life is. There's still no universally agreed definition of it. Indeed, one of the aims of A-Life research, as expressed by Chris Langton (1989), is to arrive at one. Nevertheless, ten characteristics are mentioned repeatedly in attempts at defining life: self-organization, autonomy, emergence, development, adaptation, responsiveness, evolution, reproduction, growth, and metabolism.

One could say that only one characteristic is crucial, for self-organization (the spontaneous appearance of new levels of order) covers all the others as special cases. It's no accident, then, that Jonas and Maturana and Varela referred continually to self-organization in their discussions of life.

One could also say—and A-Life researchers typically do so–that the first eight characteristics listed above are abstract (functionalist) notions, defining aspects of what's sometimes called the logical form of life. As such, they are grist to the mill of a computer-based approach. Even growth might be included, if we allow that this term is ambiguous as between physical growth and incresase in size-otherwise-defined (e.g. as length/number of program instructions).

Only metabolism—also heavily stressed by the writers featured in Section "The Life-Mind Linkage Defended"—remains incalcitrant to an abstract, functionalist, interpretation. For it's irredeemably physical. Moreover, it doesn't refer only to energy-use, or even to individual energy-budgeting (both of which can be ascribed literally to computers). Rather, it refers to the use and budgeting of energy in the autonomous construction and maintenance of the living system itself. This, given evolution, will inevitably involve a set of interacting biochemical cycles, of increasing complexity (Boden 1999).

That's why strong A-Life, i.e. virtual life, is impossible. For computers don't metabolize, in the sense just defined. (It doesn't follow that A-Life can't throw significant light on various examples of self-organization, metabolism included. Indeed, it has already done so—see Boden 2006, Chap. 15.)

However, these remarks don't settle our query here. For while self-organization is undeniably a key concept in defining life, it's not at all clear that it's a key feature of mind. This is a conceptual point, not an empirical one; so the growing evidence of self-organization in the brain, though fascinating, is irrelevant (Boden 2006, 14.vi.b and ix.a-d). Only if we were to define mind as mind/brain would this evidence be conceptually germane.

Nor is it clear, pace Jonas, that mind must arise from metabolism. One might say that mind, and intelligence, is necessarily adaptive. But does this mean anything over and above 'efficient', and/or 'well-suited to the specifics of the environment', and/or 'capable of change through learning'? All of those descriptions could conceivably be attached to robots.

In particular (and, again, pace Jonas), it's not obvious that minds must necessarily be generated by evolution. It's not even obvious that life itself must involve evolution.

Maturana and Varela, for instance, argue that evolution (and reproduction too) presupposes autopoiesis, so can't be essential to it (1980, p. 105ff.). Moreover, to include evolution in the list of vital characteristics has some counterintuitive implications (Bedau 1996). Since the concept applies only to populations, an individual oak-tree or lion—usually regarded as paradigm cases of life—can't be regarded as a living thing except by appealing to their ancestry. A non-evolving population, temporarily in evolutionary equilibrium, wouldn't exemplify life either. And creationism becomes incoherent, not simply false.

Nevertheless, all the living things we know of have evolved. Moreover, it's difficult (to put it mildly) to see how highly complex organisms could be generated except via evolution—a point which even Maturana and Varela were happy to grant. In addition, the concept of evolution enables us to inter-relate all known living things, and to explain a host of details about them—which is why it's so often included in the definition of life.

Let's agree, then, that all life (with the possible exception of the most primitive autopoietic unities) has in fact evolved, whether or not we also choose to say that it must have done so. Does this have any bearing on the life-mind linkage?

Evolution and Intentionality

Talk of evolution reminds us of another way of linking life and mind. For intentionality, or meaning, is the key aspect of mind, and some philosophers, such as Ruth Millikan (1984) and David Papineau (1987), have argued that intentionality is grounded in evolution. (So did Jonas, as we've seen; but whereas he asserted it in a rhetorically persuasive fashion, they considered counter-arguments carefully in stating their case.) In a nutshell, what these authors argue is that it is adaptive function, favored and fixed by evolution, which gives meaning to animals' actions—including the linguistic activities of human beings.

I said, in Section "What is Life?", that it's not obvious that minds must be generated by evolution. Not only isn't it obvious (i.e. clear at first sight), but it's still widely contested. Wittgensteinian philosophers, for example, deny the possibility of any naturalistic account of meaning, evolutionary or not (e.g. Morris 1992; McDowell 1994). (Hence Millikan's cheeky book-title, outrageously describing language and thought as "biological" categories.) Moreover, even when the claim is accepted, as the best alternative on offer, it's usually admitted to be problematic.

For instance, even sympathetic readers of Millikan's book may be discomfited by her "swamp-man" thought-experiment (Millikan 1984, pp. 93, 337f.; 1996). She admits that her philosophy of intentionality implies that a magically constituted molecule-for-molecule copy of Jo Bloggs' body and brain would have no knowledge, no beliefs, no understanding… despite responding to all our greetings and questions exactly as Bloggs would have done. In short: no evolution, no meaning. She's discomfited by this implication herself. But she refuses to be defeated by it, arguing that the thought experiment is so utterly unrealistic that it's not worth taking seriously. I agree with her. (Similarly, I wouldn't drop thermodynamics simply because it allows for the theoretical possibility, if only for a split second, of a snowball in Hell.) But not all philosophers would.

If the anti-naturalists are right, and even biology (evolutionary theory and/or neurophysiology) can't explain intentionality, then computers certainly can't do so. And if they're wrong, it still doesn't follow that a computer-based discipline can help supply the naturalistic explanation being sought. We've already seen that computers aren't alive. Can programs, or robots, nevertheless be said to evolve? And if so, can they help us to understand the origin of meaning?

Evolution is the gradual change of a population whose individual members reproduce ('asexually' or 'sexually') with inheritance and variation, where some fitness function selects the next breeding-individuals at each generation. The "change" is typically an improvement, with respect to the 'task' implied by the fitness function. Given the abstract nature of this concept (see Section "What is Life?"), the systems concerned need not be organisms: they may be programs, or even robots.

Programs evolve by employing genetic algorithms, or GAs. These enable reproduction—either self-copying or 'copying' from two parents—with variation. They make random changes (mutations, crossovers…) in some of the program-rules at the copying stage, and then apply a fitness function to select the 'best' resulting individuals (plus a few others, to allow potentially advantageous variations to remain in the ‘gene pool’). (Sometimes, it’s the programmer who applies the fitness function; for our purposes, however, the fully automatic cases are more relevant.)

GAs can be used to tackle non-biological tasks, such as optimizing the design of an aircraft engine. But A-Life researchers are typically concerned with biological phenomena. Not all A-Life research involves GAs: much of it studies non-evolutionary types of self-organization. But much of it does, and the examples mentioned below are all taken from ‘evolutionary’ A-Life.

Usually, the researchers pick a single ‘task’ and try to evolve programs that can achieve it, in one or more environments. But, as in biology itself, “evolution” may involve co-evolution of distinct groups, or species—predator and prey, for instance (Cliff and Miller 1995). The first, and still the best-known, example is Tom Ray’s (1992) Tierra model. This program is famous partly because of its results: to Ray’s amazement, it generated parasites, counter-parasites (i.e. hosts resistant to the previous parasites), and super-parasites (which overcame the hosts’ previously evolved resistance); it also proved that gradual genetic change can underlie what looks like saltatory, or “punctuated”, evolution. It’s famous also because of Ray’s claim that virtual “creatures” like those evolved by Tierra are genuinely alive. (He counters the ‘no metabolism’ objection by pointing out that computers use energy; however, we’ve seen that metabolism involves much more than this.)

As for robots, the central controller, or neural-network ‘brain’, can be evolved from a starting-point wherein the connections between the component units (and their nature: excitatory or inhibitory) are random. Similarly, sense-organs—such as whiskers, or the second eye—that aren’t actually needed for the task may lose their connection to the ‘brain’ and become useless, rather like the human appendix (Cliff et al. 1993). Or their anatomy may evolve, so that predators develop narrow-angle, forward-looking, eyes, while the prey develop wider-range, front-and-sideways, vision: think foxes and rabbits, respectively (Cliff and Miller 1995).

What has this to do with intentionality? Two things. On the one hand, if meaning is grounded in evolution then life—which, we’ve agreed, involves evolution (whether necessarily or as a matter of fact)—is at least ‘suitable’ for mind. To that extent, the life-mind linkage is supported. On the other hand, non-living but genuinely evolved robots may offer a foothold for the ascription of meaning that isn’t available to other computer systems.

Searle directed his “all syntax and no semantics” objection to robots, as well as to programs. He said—and he was right—that the ‘meanings’ we ascribe to program-instructions come wholly from us. In principle, one and the same program could be interpreted either in terms of the tax-laws or in terms of dance-steps. There is nothing intrinsic to the program to choose between the two interpretations (or three, or…). With robots, it’s a bit more tricky. After all, it might be a dancing robot, controlled by the tax-or-choreography program—in which case the second interpretation would seem to be clearly more appropriate. Even here, however, Searle could point out that the fact that the robot performed such-and-such a step at a certain time is wholly dependent on its programmer: he/she programmed it to dance the polka, but could have programmed it to tap out the numbers for your tax-return.

With an evolved robot, however, the case is different. Let's take an actual example. A population of robots was being evolved to navigate a particular environment so as to reach a certain spot, irrespective of the starting-point (Harvey et al. 1994; Husbands et al. 1995). One of the things in the environment was a triangle of white cardboard. Sometimes (for evolution is probabilistic, not deterministic), the neural-network controller evolved a 'feature detector' analogous to those discovered in monkeys' brains. This was a mini-network sensitive to a light–dark gradient at a particular orientation. (No other such mini-networks evolved in this case; so a black triangle, or the right side of the white triangle, were in effect invisible.)

Putting this example in intentional terms, activity in the mini-network meant 'light–dark gradient at orientation x', or perhaps 'left side of white triangle', or even 'landmark directing me to veer to the right'. Since the robot lacked language, and hadn't even evolved a rich set of visual discriminations, it's not obvious which of these descriptions one should pick. In other words, if there was content (meaning) here, it was non-conceptual content (which some philosophers deny outright, even in non-human animals: McDowell 1994).

The point, however, is that the mini-network evolved as part of a visuomotor mechanism. The connections of specific visual units in the 'brain' to specific motor units enabled the robot to use the white triangle as a navigation aid in achieving the task which (unknowingly) it had been set. Indeed, the feature-detector's very existence, as well as its function, depended on that evolutionary history—not on any foresight by the roboticists.

Even if one doesn't accept an evolutionary account of meaning in general, it's surely more appropriate to ascribe one of the three meanings suggested above to the mini-network than to say—inspired by Searle—that it could mean just anything. Certainly, it couldn't be used as a tax-calculator: that would require its components and connections—its nature—to be very different. And if one does see evolutionary history as essential to meaning, this conclusion is even firmer.

My claim, here, isn't that an evolutionary semantics must ascribe real (non-conceptual) meanings to this robot, or to complexified versions of it. For although the existence of that specific mini-network depends wholly on evolution, the existence of the (initially, random) neural-network controller itself does not. On the contrary, it depends on deliberate human agency. One might wish to argue that this makes all the difference, that "real" meaning demands evolution all the way down. I shan't pursue that point here (but see Boden 1972, pp. 119, 195). It's enough, for my purposes, to show that if one does ascribe meaning (whether real or metaphorical) to the robot and/or to its feature-detector then there are strict limits on the meanings one can plausibly suggest. Orientation-detector, yes; navigational landmark, probably; one side of a white triangle, perhaps; tax-calculations, certainly not.

Conclusion

I've argued that life implies evolution (though perhaps not of necessity). And I've suggested that evolution-based philosophies of meaning are the most plausible accounts on offer (and that we need some naturalistic semantics or other). If both those claims are true, then life is well-suited for mind, even if it's not actually necessary for it.

Computers, robots included, aren't living things, because they don't metabolize in the required sense. But some of them do, genuinely, evolve. Of those, some are systems to which we naturally ascribe meanings which are not arbitrary with respect to the artefact itself. On the contrary, they are grounded in its specific evolutionary history.

Whether these (non-arbitrary) "meanings" are real or merely metaphorical depends partly on whether our philosophical semantics demands evolution all the way down. For even a robot of the ten-millionth generation, whose behavior (and anatomy) was unforeseen, wouldn't have existed if human beings hadn't embarked on evolutionary robotics in the first place.

Finally, this sub-area of A-Life may help us to model the evolution of non-conceptual content, and (with language added) possibly of conceptual content too. In so doing, and even if that content is merely metaphorical, it may help to clarify our philosophical ideas about mind, and its relation to life.

References

Bedau, M. A. (1996). The nature of life. In M. A. Boden (Ed.), *The philosophy of artificial life* (pp. 332–357). Oxford: Oxford University Press.

Boden, M. A. (1972). *Purposive explanation in psychology*. Cambridge, MA: Harvard University Press.

Boden, M. A. (1999). Is metabolism necessary? *British Journal for the Philosophy of Science, 50*, 231–248.

Boden, M. A. (2000). Autopoiesis and life. *Cognitive Science Quarterly, 1*, 1–29.

Boden, M. A. (2006). *Mind as machine: A history of cognitive science*. Oxford: Clarendon Press.

Cliff, D., Harvey, I., & Husbands, P. (1993). Explorations in evolutionary robotics. *Adaptive Behavior, 2*, 73–110.

Cliff, D., & Miller, G. F. (1995). Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In F. Moran, A. Moreno, J. J. Merelo & P. Chacon (Eds.), *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life (ECAL95)* (pp. 200–218). Berlin: Springer.

Edgley, R. (1970). Innate ideas. In G. Vesey (Ed.), *Knowledge and necessity* (pp. 1–33). London: Macmillan.

Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.

Geach, P. T. (1980). Some remarks on representations. *Behavioral and Brain Sciences, 3*, 80–81.

Harvey, I., Husbands, P. & Cliff, D. (1994). Seeing the light: Artificial evolution, real vision. In D. Cliff, P. Husbands, J.-A. Meyer & S. W. Wilson (Eds.), *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior* (pp. 392–401). Cambridge, MA: MIT Press.

Husbands, P., Harvey, I., & Cliff, D. (1995). Circle in the round: State space attractors for evolved sighted robots. *Journal of Robotics and Autonomous Systems, 15*, 83–106.

Jonas, H. (1966). *The phenomenon of life: Toward a philosophical biology*. New York: Harper Collins. (Page-references are to the 2001 edition. Evanston, Illinois: Northwestern University Press).

Jonas, H. (1984). *The imperative of responsibility: In search of an ethics for the technological age*. Chicago: Chicago University Press.

Jonas, H. (1990). Heidegger's resoluteness and resolve: An interview. In G. Neske & E. Kettering (Eds.), *Martin Heidegger and national socialism: Questions and answers* (pp. 197–203). New York: Paragon House.

Langton, C. G. (1989). Artificial life. In C. G. Langton (Ed.), *Artificial life*. (pp. 1–47). Redwood City, CA: Addison-Wesley. (Revised version in M. A. Boden (Ed.), The philosophy of artificial life (pp. 39–94, 1996). Oxford: Oxford University Press).

Lettvin, J. Y., Maturana, H. R., McCulloch, W. S., & Pitts, W. H. (1959). What the frog's eye tells the frog's brain. *Proceedings of the Institute of Radio Engineers, 47*, 11.

Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living*. Boston: Reidel.

Mays, W. (1952). Can machines think? *Philosophy, 27*, 148–162.

McDowell, J. (1994). *Mind and world*. Cambridge, MA: Harvard University Press.

Millikan, R. G. (1984). *Language, thought, and other biological categories: New foundations for realism*. Cambridge, MA: MIT Press.

Millikan, R. G. (1996). On swampkinds. *Mind and Language, 11*, 105–117.

Morris, M. R. (1992). *The good and the true*. Oxford: Clarendon Press.

Papineau, D. (1987). *Reality and representation*. Oxford: Basil Blackwell.

Putnam, H. (1962). Dreaming and "depth grammar". In R. J. Butler (Ed.), Analytical philosophy (first series) (pp. 211–235). Oxford: Blackwell.

Ray, T. S. (1992). An approach to the synthesis of life. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial life II* (pp. 371–408). Redwood City, CA: Addison-Wesley.

Scriven, M. (1953). The mechanical concept of mind. *Mind, 62*, 230–240.

Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences, 3*, 417–457.

Turing, A. M. (1950). Computing machinery and intelligence. *Mind, 59*, 433–460.

Ziff, P. (1959). The feelings of robots. *Analysis, 19*, 64–68.