“The Action of the Brain”
Machine Models and Adaptive Functions in Turing and Ashby

Hajo Greif

[1] Munich Center for Technology in Society, Technical University of Munich, Germany
[2] Department of Philosophy, University of Klagenfurt, Austria
hajo.greif@tum.de

Abstract Given the personal acquaintance between Alan M. Turing and W. Ross Ashby and the partial proximity of their research fields, a comparative view of Turing’s and Ashby’s works on modelling “the action of the brain” (in a 1946 letter from Turing to Ashby) will help to shed light on the seemingly strict symbolic / embodied dichotomy: while it is a straightforward matter to demonstrate Turing’s and Ashby’s respective commitments to formal, computational and material, analogue methods of modelling, there is no unambiguous mapping of these approaches onto symbol-based AI and embodiment-centered views respectively. Instead, it will be argued that both approaches, starting from a formal core, were at least partly concerned with biological and embodied phenomena, albeit in revealingly distinct ways.

Keywords: Artificial Intelligence, Cybernetics, Models, Embodiment, Darwinian evolution, Morphogenesis, Functionalism

1 Introduction

Not very much has been written to date on the relation between Alan M. Turing and W. Ross Ashby, both of whom were members of the “Ratio Club” (1949–1958).[1] Not much of the communication between the two seems to have been preserved or discovered either, the major exception being a letter from Turing to Ashby that includes the following statement:

> In working on the ACE [an early digital computer] I am more interested in the possibility of producing models of the action of the brain than in the practical applications of computing. [. . .]
> It would be quite possible for the machine to try out variations of behaviour and accept or reject them in the manner you describe and I have been hoping to make the machine do this. This is possible because, without altering the design of the machine itself, it can, in theory at any rate, be used as a model of any other machine, by making it remember a suitable set of instructions. (Turing, 1946, 1 f)

[1] The best historical accounts of the Ratio Club and Turing’s and Ashby’s roles therein are Husbands and Holland (2008), Holland and Husbands (2011).

A comparative view of Turing's and Ashby's work on modelling "the action of the brain" (Turing, 1946) will help to elucidate the modelling properties of machines with respect to human thinking and behaviour. It will be safe to say that Turing was committed to formal "symbolic simulations" and Ashby to material "working models", with corresponding modes of reference to their target systems. This part of the analysis will be largely in line with Peter Asaro's account of "Computers as Models of the Mind" (2011). However, in terms of Turing's and Ashby's fundamental views of the nature of what is modelled, the picture gets more complex: despite the respective foci of their models on the functions of machines and biological systems, both approaches were in some important respects concerned with biological and embodied phenomena. Both relied on theories of these phenomena, but they relied on competing theories in distinct ways. I will go through Turing and Ashby twice in order to make these points clear, first outlining their takes on modelling (Section 2), then their biological credentials (Section 3) and finally their implications (Section 4).

2 Formal and Material Models

There are various key motives shared between Turing's and Ashby's work that would figure in either AI and cybernetics. Both Turing and Ashby believed that "the action of the human brain" can be subject to a method of modelling that casts it in a strict mathematical description and breaks it down into elementary routines in such a way that the model could be implemented in some kind of machine, in principle at least.

The shared motive of devising machine-implementable models rests on the premise that the behaviour of some system can be described or imitated by a system of altogether different physical make-up. The notion that an identical set of logical operations can be realised in physically variant systems has its paradigm in Turing's universal computing machines, which were initially *theoretical* machines (1936). This proposition has come to be known as "machine state" or "Turing Machine" functionalism (this terminology being introduced by Putnam, 1975).

However, first, while Ashby certainly embraced multiple realisability for his machine models, the functionalism he employed and the analogies it implied ultimately were different in kind from Turing's machine state functionalism. Instead, Ashby's functionalism was a biological, essentially Darwinian one, as shall be demonstrated in Section 3. Second, there is a number of ways in which Turing and Ashby differed with respect to the manner in which machines shall serve as models, and what the paradigm of machines to do that modelling is. These differences will be briefly outlined in the present section.

Turing: Turing based his models on his mathematical theory of computation. His original quest was for solving the "Entscheidungsproblem" (decision problem) in Gödel (1931): Turing's idea was that the operations required for evaluating whether a given logical proposition can be proven true or false within the calculus to which that proposition belongs could be implemented in an appropriate kind of theoretical machinery, christened the "Logical Computing Machine", by means of interchangeable sets of formal instructions, that is, programs (Turing, 1936). This method could be applied to any field of inquiry that allows for a translation of complex logical propositions into a set of elementary, machine-executable operations. The physical characteristics of any real-world machine to implement these functions are underspecified by this requirement, as long as the requisite formal characteristics are in place. As a matter of historical fact, however, the machines to accomplish this task turned out to be digital, electronic, programmable computers. While Turing's own work made major conceptual and practical contributions to the development of these machines, and while computers are the paradigmatic logical computing machines, this does not imply that digital computers are the only conceivable machines of this sort.

The theoretical import of Turing's models lies fully within the realm of mathematics, while their empirical import lies in demonstrating the scope and force of his theory of computability in (thought-) experimental fashion in a variety of fields. His self-ascribed primary empirical interest was in the action of the brain, but his most substantial contribution to any field outside computer science was his mathematical theory of morphogenesis, that is, the patterns of organic growth (Turing, 1952),[2] to which I will pay detailed attention in Section 3. With respect to cognitive phenomena, Turing placed his inquiry on two separate levels: in "Computing Machinery and Intelligence", he professed in "drawing a fairly sharp line between the physical and the intellectual capacities of a man" (1950, p. 434) in the design of his "imitation game". In order to grant fair play to machines in this game of simulating human conversational behaviour, he suggested to disconnect conversational abilities from any underlying organic traits. However, when Turing moved on to a consideration of possible mechanisms responsible for these conversational abilities, he introduced his proto-connectionist "B-type unorganised machines" that exemplify structures and processes in the brain on an abstract level (Turing 1948; see also Copeland and Proudfoot 1996).

Either way, Turing considered the phenomena in question chiefly in their *form*, and thereby to the extent they are accessible to the computational method. More precisely, his quest was for descriptions of the behaviour of a target system that can either directly serve as, or be transformed into, input variables for a set of equations which can then be solved by applying computational routines, so that the output either directly describes or predicts a further behaviour of that target system, or can be transformed into such a description or prediction. Any

[2] In the introduction to a posthumous collection of Turing's writings on morphogenesis, Peter T. Saunders claims that Turing (1952) "is still very frequently cited (more than the rest of Turing's works taken together [...])" (Saunders, 1992, p. xvi). If *Google Scholar* and citation counts are resources to go by, the parenthetical part of this statement is an exaggeration, but Turing (1952) still ranks approximately 10 and 20 percent higher in number of citations respectively than the other two of his most-referenced works, Turing (1950) and Turing (1936): https://scholar.google.de/citations?user=VWCHlwkAAAAJ&hl=en&oi=ao (accessed March 28th, 2018). The *Thomson Reuters* and *Scopus* databases have an incomplete record of the original editions, hence cannot be used for comparison.

system whose behaviour can be formalised so as to be amenable to this kind of procedure could be subject to the computational method.

Moreover, Turing professed in formally "describing the mode of behaviour" (Turing, 1946, p. 2) of a learning system on an *individual* level in the first place, treating it as a self-contained entity connected to an environment through a number of in- and output channels. The environment of that entity remains underspecified, and is mostly conceived of as input from the experimenter. When Turing considered the action of the brain, he purposefully limited his focus on that organ proper (for this observation, see Hodges, 2008, p. 85). Such focus on self-contained, individual entities was arguably guided by a methodological presupposition: as the original topic of Turing's inquiry were elementary recursive operations within a calculus, any empirical test for the force and scope of the computational method would, if not necessarily then naturally, commence with relations of this kind in the target system.[3] The notion of arithmetical routines repeatedly using their own output as input for their next round of application (hence "calling on themselves") may count as the paradigm of Turing's computational method. The method's focus is on what happens to an initial, expressly restricted, input over the course of repeated computational steps, unlike, for example, the equations describing the time evolution of a dynamical system, which take an open-ended sequence of states of the environment as their input values. This basic methodological presupposition, rather than the higher-order question of embodiment, might be the first indicator of the schism that would later develop between cognitivist AI and cybernetics.

Ashby: Ashby's models and their target systems differ from Turing's, first, in his quest being for the *origins* of adaptive behaviour of organisms and other systems with respect to their environments (Ashby, 1947, 1960). He built his homeostat as a system that was supposed to actually learn, and to share a set of core features of functional organisation with any other, natural or artificial, learning system that has to cope with changing environmental variables. The overarching systematic goal of his research was to explain "whence come the patterning properties of the nervous system" (Ashby, 1928-1972, p. 6117, entry of June 13th, 1959). In shaping those patterning properties, interactions with the environment, including the behaviour of other organisms, played a crucial role, and were thus fully and expressly incorporated in Ashby's machine model. Hence, second, the best available evidence of the validity of the model lies in its ability to function in real, variable environments, and is best exemplified by a physical machine, the homeostat. There was both a didactic and a systematic purpose in having a physical implementation of the model. Third, Ashby's methodological choice was to describe the time evolution of a system, as defined by the observer. That description consisted in tracking a succession of states as "lines of behaviour" in a phase space or "field" of the values of the selected variables, given certain

[3] For anti-individualistic views of Turing's approach, see the reading of Turing (1948) proposed by Herold (2003) and the claim that Turing machines are situated systems by virtue of their tapes being part of their local environments (Fabry, 2018).

initial states and a sequence of environmental inputs. This approach has fairly little in common with the recursive operations applied to delimited and largely self-contained systems that was favoured by Turing.

These marks of distinction from Turing's approach find their common roots in the Darwinian paradigm to which Ashby committed himself, and hence in a theory of the evolution of biological systems. Ashby took this framework to generalise to all sorts of systems capable of adaptive behaviours or, in Ashby's terminology, attaining equilibrial states. His assumption was that adaptivity and goal-directed organisation emerge from processes that are not goal-directed themselves but include random variation and deterministic selection, and he tried to single out the basic natural mechanisms by which it is accomplished. Ashby's machines modelled organism-environment relations as relations of negative feedback, in which changes in environmental variables provoke counter-effects in the machine, and vice versa. If the change in environmental variables pushes some of the variables within the machine beyond a threshold of stability, the machine, by means of "step-mechanisms", randomly produces new states (as the counterpart of variation in Darwinian evolution) and matches them against what the environment provides (as the analogue of natural selection) until it re-enters a domain of stable values (resulting in an analogue of fitness).

Ashby's approach to modelling was formal and mathematical inasmuch as the theory of feedback mechanisms, equilibria, and stability can be articulated in rigorous mathematical fashion. However, Ashby employed a formal, mathematical apparatus primarily in instrumental fashion, making it subserve the broader purpose of a general science of organisms and other systems. Hence, the theoretical import of Ashby's modelling did not lie within the realm of mathematics. Nor were his models *computational* under any interpretation that would approximate Turing's notion. Moreover, Ashby's approach to modelling was also genuinely material, not merely in terms of the model's physical implementation. There are isomorphisms supposed to hold between the functional status of a machine and the functional status of a target system in such a way that transformations in the target system are matched by transformations in the machine model that are analogous in terms of the functional states involved. If there are an irritation and a negative feedback in target system upon a certain input, there should be an irritation and a negative feedback in the machine model as well, with measurable correspondences. Arguably, the materiality and non-computationality of Ashby's models in conjunction made them less universal and less adaptable than Turing's computer models, at least once they were computer-implementable in fact (which, as Asaro 2011 observes, makes the choice of modelling approaches a matter of available resources in part).

On the background of this discussion, the differences between Turing's and Ashby's approaches can be located on two general levels: *First*, the primary though not the exclusive focus was on formal versus material modelling respectively, with diverging roles assigned to the mathematical methods involved. *Second*, the functionalism implied by Ashby's argument was of a different sort than Turing's: it primarily considered adaptive, biological or biologically based, functions of brains and other systems to be modelled rather than the logical properties of machine states. Ashby apparently focused on, while Turing intentionally skipped over, the question of where and how the goal-directed organisation of organism or machine in their concrete environments originate.

3 Adaptive Functions vs Laws of Form

I will now argue that Turing's and Ashby's views on biology provide a key to understanding their differences in approach on either of the aforementioned levels. Both the preferred type of modelling and the kind of functionalism chosen are deeply informed by their views of the relevance of Darwinian evolution.

The arguably most articulate and most influential computational model devised by Turing concerned the biological processes of morphogenesis (see Turing, 1952), which built on Sir D'Arcy Thompson's, at its time, influential work *On Growth and Form* (1942). Turing sought to apply and test Thompson's account of the generation of organic patterns, from an animal's growth to the grown animal's anatomy, from the dappledness or stripedness of furs to the arrangement of florets of a sunflower and the phyllotaxis, that is, the ordering of leaves on a plant's twigs. Turing's question, like Thompson's, was how such intricate and differentiated patterns develop from genetically homogenous cells. Was there a general mechanism of pattern formation that could be formally described? The formalism of linear and non-linear differential equations used by Turing was expressly impartial to the actual biochemical realisation of pattern formation. It would only provide some clues as to what concrete reactants, termed "morphogens" by Turing, one should look out for. Answering questions of de facto realisation would be a central task for the many biologists, chemists and others who followed Turing's lead. Still, his account of morphogenesis was as close to a direct modelling relation to a natural target system as it would get in Turing's entire work – closer in embodied detail than his proto-connectionist endeavours, and certainly much closer than his imitation game.

In his morphogenetic inquiries, Turing did *not* inquire into any adaptive function, in Darwinian terms, of the patterns so produced. These patterns may or may not serve an adaptive function. Computationally modelling their formative processes does not contribute to explaining that form's function. Whether the florets of a sunflower are patterned on a Fibonacci series, as they, in fact, are, or whether they are laid out in grid-like fashion, as they *cannot* possibly be according to the mathematical laws of form, is unlikely to make a difference in terms of selective advantage. In turn, however, natural selection may not offer a path to a grid-like pattern in the first place, while allowing for, and perhaps enabling, but arguably not determining the Fibonacci pattern.

This seeming indifference towards adaptive functions might be referred to Asaro's observation that the computational mode of modelling is distinctly indirect, and cannot establish the modelling relation by itself (2011). It would have to incorporate a reasonably elaborated theory of the target system that assigns a suitable calculus and suitable inputs in order to attain any degree of sophistication. A reasonably elaborated theory of this kind was amenable to Turing's computational method in D'Arcy Thompson's laws of form but apparently not in Darwinian variation and natural selection.

The reasonable methodological choices that entitled Turing to a limited interest in the evolutionary origins of organic patterns have a far-reaching implication in terms of alignment with competing research paradigms: D'Arcy Thompson's laws of form were articulated with an explicit scepticism towards the relevance of adaptation by natural selection in biology, and claimed an autonomy of formative processes in organisms from Darwinian mechanisms (Thompson, 1942). More precisely, D'Arcy Thompson argued that an organism's development is subject to constraints on form that have to be explained, and can be sufficiently explained, by reference to mathematical and physical regularities. Hence, pattern formation cannot be subsumed under a Darwinian account of random variation and natural selection. Natural selection does act on biological forms with respect to their environmental fitness but it cannot generate them, nor is it the only or even the primary constraint on the realisation of possible forms. Instead, D'Arcy Thompson referred to Goethe's archetype theory in this context, picking up on laws of symmetry and the expression of identical forms in otherwise variant organisms, and sharing the observation that some biological forms are more probable to develop than others, whereas still others are genuinely impossible. He puts these Goethean claims and observations on a strictly mathematical and physical footing: organic forms and their transformations are both enabled and constrained by physical laws that can be expressed in mathematical terms. Turing picked up on the form of these mathematical expressions and developed them into a dynamic model of pattern formation.

Even if Turing did not actively endorse D'Arcy Thompon's sceptical view of Darwinism, he at least implicitly went along with it, and chose to work under a paradigm that built on it. Notably, as Boden (2006, pp. 1264-1267) observes, Turing's theory of morphogenesis was enthusiastically received by the embryologists of his day, who were more likely to be attached to theories of archetypes and ontogenetic recapitulation than to notions of Darwinian selection. Moreover, Boden highlights that Turing's mathematical theory of morphogenesis was ultimately proven right by more recent and powerful computer simulation methods, and that it informed those branches of Artificial Intelligence that have come to be known as Artificial Life. However, a similar status can be argued to accrue to principles of Darwinian evolution.

Unlike Turing, Ashby repeatedly referred to the Darwinian notions of adaptation and natural selection (for example, Ashby, 1960, p. 29). In doing so, he exclusively focused on the adaptive aspects of evolution – which is not a matter of course, as we saw, and which was not at all the dominant view in biology at the time of his writing. Still, Ashby believed that random variation of existing traits of a reproducible (or, in Ashby's models, modifiable) system and selection by differential reproduction (or continued functioning versus dysfunction) of that system under the influence of a given set of environmental conditions will not only be the primary but the exclusive path for a system to attain an adaptive or, in Ashby's terminology, equilibrial state. Given that Darwinian evolution is by and large gradual, variation is limited, first, to minor alterations and, second, to acting only on parts of an existing structure, thereby excluding large-scale changes across most or all parts of the system. Third, variation would typically not affect continuously functioning adaptive traits: "organisms are usually able to add new adaptations without destroying the old" (Ashby, 1960, p. 142). With these limiting conditions in place, no other mechanisms than random variation and natural selection were deemed required to explain the goal-oriented structure and behaviour of an organism or other organised system.

Remarkably, however, in his very emphasis on adaptation, Ashby restricted his focus to the origins of adaptive behaviour by learning, not inquiring deeper into "genic" adaptation, so that the organic basis for the production of such behaviour was largely left aside.[4] In fact, Ashby referred to biological evolution and what he called "Darwinian Machinery" only in the more philosophical and speculative of his writings, for example, Ashby (1952, pp. 50-52) and Ashby (1967), and, of course, in many places in his *Journal*. The Darwinian basis appears to be taken for granted in its own, genuinely biological, adaptive organisation by Ashby – while both behavioural and genic adaptation were assumed to be subject to the same general laws of variation and selection.

The first obstacle to an incorporation of the organic level of adaptation into Ashby's model lies in a perennial problem of evolutionary theory: apart from some fast-breeding model populations, observed under laboratory conditions and the partial evidence they provide, one cannot observe natural variation and selection in real-time. Hence, one is not enabled, first, to precisely and unequivocally map organic traits onto environmental conditions so as to define them as adaptations (rather than contingent effects) to precisely these (rather than other) conditions. Second, even if such mapping can be accomplished, there will be no evidence whatsoever on the history and the dynamics of the concrete *process* of adaptation – that is to say, unless one has a very good grasp of population genetics and ecology along with paleontological evidence. These problems were compounded by an interpretation of Darwinian evolution under which "the species is fundamentally aimless (it *finds* its goals as it goes along)" (Ashby, no year, no. 5). In a discussion of the DAMS, a post-homeostat machine model, he actually embraces the view that "variables in the brain should be driven actively by the environment" (Ashby, 1928-1972, p. 3831, entry of May 19th, 1952). Contemporary biology has come to accept different, more differentiated views of the organism-environment relation. Some but not all present-day views appreciate of D'Arcy-Thompson and Turing's morphogenetic laws, while many if not most of them ascribe a more active role to organisms with respect to their environments.

4 Constraints, Implications, and Potentials

From the preceding observations, one might infer that there is one remarkable lacuna in both Turing's and Ashby's accounts of the action of the brain: des-

[4] For Ashby's discussion of "Darwinian Machinery", see also Asaro (2008, pp. 166-168).

pite their references to variant traditions on biology, they appear to pay little systematic attention to the question of the origins of the biological mechanisms responsible for that action. This seeming lacuna is not a mere omission, but, first and foremost a constraint imposed by the state of the biological sciences of their time. It may also serve as a diagnostic of Turing's and Ashby's modes of systematic theorising and their implications, which in turn may offer a potential resolution to their seeming opposition.

With respect to the constraints involved, leaving the mechanisms of biological origins out of the picture may, in Turing's reliance on D'Arcy Thompson's view of biological form as well as in Ashby's focus on behavioural adaptation, partly owe to the fact that, at the time of their writing in the late 1940s and early 1950s, the authors were not in a position to rely on what would become known as the "modern synthesis" in evolutionary biology. That synthesis, developed between 1936 and 1947, paradigmatically stated by Julian Huxley (1942) and later summarised by Ernst Mayr (1991, especially Ch. 9),[5] was just about to become the dominant biological paradigm, and took several years more to become part of common knowledge. The modern synthesis amounted to crossing the Darwinian mechanisms of adaptation by random variation and natural selection – which had become somewhat rarefied since Darwin's time – with the statistical laws of mathematical population genetics, so as to produce, for the first time, a comprehensive and strongly empirically grounded paradigm of Darwinian biology. Thereby, a considerable degree of consensus was established in evolutionary biology which, in conjunction with the rise of molecular biology, sidelined D'Arcy Thompson's laws of form along with a variety of epigenetic and vitalist theories.

While Turing did not live to see the full establishment of the modern synthesis, an appreciation of Ernst Mayr's work can be found in Ashby's later writings, see Ashby (1928-1972, p. 6637, entry of December 27$^{\text{th}}$, 1966). Accordingly, an argument for isomorphisms between machines and human cognitive traits that can rely both on mechanisms of genetic replication and variation in population and on the Darwinian concept of functional analogy between phylogenetically distinct traits was not available to Turing, given his preference for a different, competing and at that time still competitive tradition in biology. To Ashby, Darwinian functional analogy was a desirable and straightforward route but, given the state of biology at his time, one he could not consider entirely safe. The perspectives towards a synthesis between their respective approaches might have much improved with an adoption of some of the key insights of the modern synthesis, and of what followed afterwards.

With respect to the systematic implications, if the importance of the organic level of adaptation is generally acknowledged as a precondition of behavioural adaptation in principle, and if the importance of embodiment and environment in adaptive processes is accepted, Ashby's material approach to modelling the action of the brain provides an outline for a biologically informed cognitive science – even though neither his homeostat model nor the state of biological theory available to him were sufficiently equipped to develop it adequately.

[5] A lucid secondary source on the modern synthesis is Depew and Weber (1995, Pt. II).

If however, as in Turing, adaptive functions are deemed of secondary theoretical importance, while trial-and-error learning and neuronal patterns became topics of his computational models, the possible relevance vs. irrelevance of adaptive functions, for matters of consistency, should be allocated to different levels of his inquiry. The distinction between the formal nature of the computational method and the materiality and embodiment of its target systems will be of methodological importance when it comes to computationally modelling embodied phenomena and their adaptive functions. After all, Turing was right in claiming a degree of universality for his computational method and its applicability in science that could not be attained by other approaches to modelling, including Ashby's. Nothing in his approach rules out the possibility of computationally modelling adaptive processes once a sufficiently elaborated theory is in place. At the same instance, nothing in Turing's arguments requires that the operations of his theoretical machines directly correspond to, let alone are identical with, the action of the brain or the development of organisms. If one prefers to argue that cognition and biological pattern formation *are* computational processes *sensu strictu*, one will have to look somewhere else than Turing.

With these constraints and implications stated, it will be possible venture beyond what Turing and Ashby could de facto accomplish: from opposite angles, they charted routes towards solutions to present-day issues in biology and cognitive science alike. In conjunction, mathematical principles of biological pattern formation and the mechanisms of genetic replication, expression and control discovered since will provide information on the bounds on genetic variation and phenotypical variance. By the same token, they will help to identify mechanisms of transmission and use of both genetic and developmental information in structuring organic patterns and organism-environment relations. Conversely, the effects of the structures thus produced might be subject to natural selection or analogous mechanisms of retention of reproducible properties. Hence, apart from natural selection proper, organisms become able to use information actively and directly, or they become able to modify their environments in such a way as to adapt them to their needs.

It is not too surprising then that, on the one hand, some of the more recent heterodox accounts of evolution, such as Goodwin (1994), expressly supplement Darwinian mechanisms with D'Arcy Thompson and Turing's morphogenetic laws and similarly minded strands of complexity theory. It is perhaps more surprising that even as like-minded an approach as developmental systems theory (Oyama, Griffiths, & Gray, 2001) fails to do so. On the other hand, selection-based self-organisation of Ashby's variety encountered a renaissance in some subfields of Artificial Intelligence (for example, Beer & Williams, 2015; Harvey, Husbands, & Cliff, 1994), or is integrated with Turing's approach in inquiries into the evolution of information processing, from molecules to human beings, under the heading of "meta-morphogenesis" (Sloman, 2013, 2018). Hence, what fist may seem like a tension between two diverging approaches may actually converge on various levels of inquiry, and may always have been meant to do so by Ashby and Turing themselves.

Acknowledgements

This paper has been developed from a research proposal currently under review with the Austrian Science Fund (FWF), project ref. P 31136. Particular thanks go out to the reviewers of this paper, Aaron Sloman and Torben Swoboda.

References

Asaro, P. M. (2008). From mechanisms of adaptation to intelligence amplifiers: the philosophy of W. Ross Ashby. In P. Husbands, O. Holland, & M. Wheeler (Eds.), *The mechanical mind in history* (pp. 149–184). Cambridge/London: MIT Press.

Asaro, P. M. (2011). Computers as models of the mind: on simulations, brains, and the design of computers. In S. Franchi & F. Bianchini (Eds.), *The search for a theory of cognition. Early mechanisms and new ideas* (pp. 89–114). Amsterdam/New York: Rodopi.

Ashby, W. R. (1928-1972). Journal. The W. Ross Ashby Digital Archive. In 25 volumes. Retrieved March 28, 2018, from http://www.rossashby.info/journal/index.html

Ashby, W. R. (1947). The nervous system as physical machine: with special reference to the origin of adaptive behaviour. *Mind*, *56*(221), 44–59.

Ashby, W. R. (1952). Can a mechanical chess-player outplay its designer? *British Journal for the Philosophy of Science*, *3*, 44–57.

Ashby, W. R. (1960). *Design for a brain. The origin of adaptive behaviour* (2nd ed.). New York/London: John Wiley & Sons.

Ashby, W. R. (1967). The place of the brain in the natural world. *Currents in Modern Biology*, *1*, 95–104.

Ashby, W. R. (no year). Aphorisms index. The W. Ross Ashby Digital Archive. Retrieved March 28, 2018, from http://www.rossashby.info/aphorisms.html

Beer, R. D. & Williams, P. L. (2015). Information processing and dynamics in minimally cognitive agents. *Cognitive Science*, *39*, 1–38.

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

Copeland, B. J. & Proudfoot, D. (1996). On Alan Turing's anticipation of connectionism. *Synthese*, *108*(3), 361–377.

Depew, D. J. & Weber, B. H. (1995). *Darwinism evolving. Systems dynamics and the genealogy of natural selection*. Cambridge: MIT Press.

Fabry, R. E. (2018). Turing redux: enculturation and computation. *under review*.

Gödel, K. (1931). Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für Mathematik*, *38*, 173–198.

Goodwin, B. (1994). *How the leopard changed its spots. The evolution of complexity*. Princeton Science Library edition, Princeton: Princeton University Press 2001. New York: Scribners.

Harvey, I., Husbands, P., & Cliff, D. (1994). Seeing the light: artificial evolution, real vision. In D. Cliff, P. Husbands et al. (Eds.), *From animals to animats 3: proceedings of the third international conference on simulation of adaptive behavior* (pp. 392–401). Cambridge: MIT Press.

Herold, K. (2003). An information continuum conjecture. *Minds and Machines*, *13*, 553–566.

Hodges, A. (2008). What did Alan Turing mean by "machine"? In P. Husbands, O. Holland, & M. Wheeler (Eds.), *The mechanical mind in history* (pp. 75–90). Cambridge/London: MIT Press.

Holland, O. & Husbands, P. (2011). The origins of British cybernetics: the Ratio Club. *Kybernetes*, *40*, 110–123.

Husbands, P. & Holland, O. (2008). The Ratio Club: a hub of British cybernetics. In P. Husbands, O. Holland, & M. Wheeler (Eds.), *The mechanical mind in history* (pp. 91–148). Cambridge/London: MIT Press.

Huxley, J. (1942). *Evolution. The modern synthesis.* London: Allen & Unwin.

Mayr, E. (1991). *One long argument. Charles Darwin and the genesis of modern evolutionary thought.* Cambridge: Harvard University Press.

Oyama, S., Griffiths, P. E., & Gray, R. D. (Eds.). (2001). *Cycles of contingency. developmental systems and evolution.* Cambridge/London: MIT Press.

Putnam, H. (1975). *Mind, language and reality. Philosophical papers, vol. 2.* Cambridge: Cambridge University Press.

Saunders, P. T. (1992). Introduction. In P. T. Saunders (Ed.), *Collected works of A.M. Turing: morphogenesis* (Vol. 3, pp. XI–XXIV). Amsterdam: North-Holland.

Sloman, A. (2013). Virtual machinery and evolution of mind, part 3: meta-morphogenesis: evolution of information-processing machinery. In S. B. Cooper & J. van Leeuwen (Eds.), *Alan Turing – his work and impact* (pp. 849–856). Amsterdam: Elsevier.

Sloman, A. (2018). Huge but unnoticed gaps between current AI and natural intelligence. In V. C. Müller (Ed.), *Philosophy and theory of Artificial Intelligence III* (in this volume). SAPERE. Berlin: Springer.

Thompson, D. W. (1942). *On growth and form* (2nd ed.). Cambridge: Cambridge University Press.

Turing, A. M. (1936). On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, *s2-42*, 230–265.

Turing, A. M. (1946). Letter to W. Ross Ashby of 19 November 1946 (approx.) The W. Ross Ashby Digital Archive. The W. Ross Ashby Digital Archive. Retrieved August 4, 2017, from http://www.rossashby.info/letters/turing.html

Turing, A. M. (1948). *Intelligent machinery: a report by A.M. Turing.* National Physical Laboratory. London.

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

Turing, A. M. (1952). The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society, B*, *237*, 37–72.