Harvey, I. (2013)
Standing on the broad shoulders of Ashby.
Open Peer Commentary on: "Homeostasis for the 21st century?
Simulating Ashby simulating the Brain" by Franchi, S.
Constructivist Foundations 9(19):102-104.

175 OPC on "Homeostats for the 21st Century? Simulating Ashby Simulating the Brain" by Stefano Franchi

Standing on the Broad Shoulders of Ashby

Inman Harvey • University of Sussex, Brighton UK • inmanh/at/gmail.com

Upshot: It is a mistake to characterise as passive Ashby's view of life (from the 1950s), abstractly modelled in part by the homeostat; one should distinguish the stasis of homeostasis from the activity of the (model) organism. Likewise mistaken is the accusation of contingency; one should distinguish the purposeless mechanism from the purposeful (model) organism. There is no basic conflict between Ashby's view and later developments in a similar tradition; technical advances are not the same as foundational gaps.

1. I had some trouble recognising the nature of Stefano Franchi's criticisms of Ross Ashby, since they seemed to largely ignore those aspects of his work that make me rate him a giant of the twentieth century. I am a critical fan of Ashby, since I think he got quite a few technical issues wrong; but in my view, as a man of his time he asked the right sort of questions, phrased in novel form by framing cognition in terms of dynamical systems, and provided inspiration for many who built on and extended his ideas. At first sight Franchi's reading of Ashby's writings focusses on some heteronomous-autonomous distinction (§12), whereas Ashby, as far as I was aware, made little or no obvious comment on this. Ashby's focus on adaptation seems orthogonal to that distinction.

2. It can be a revelation to see how some complex, many-layered piece of writing can be interpreted differently through different eyes. Though my response here is directed at the target article, some extra background to Franchi's perspective is given in Franchi (2011); with these two articles in mind I welcomed the excuse for re-reading Ashby's *Design for a Brain*.

3. I was immediately struck, as always, by the directness and honesty of his writing; and the somewhat dated style reflects the date it was written, 1952 for the first published version. This placed it 2 years after Turing's (1950) *Mind* paper on "Computing Machinery and Intelligence", and 4 years before the 1956 Dartmouth conference usually thought of as defining (the GOFAI version of) AI based on a computational theory of mind. This dates and places Ashby's book at a major fork in the cognitive road, pointing clearly down the left branch (relatively rather neglected for some decades thereafter) defining life and cognition in terms of biology and adaptivity, using concepts from cybernetics and dynamical systems theory; whereas the right branch (soon to become the multi-lane highway of classical AI or GOFAI) defines life and cognition in terms of logic and rationality, using concepts from computational theory.

4. In his review Franchi wants to contrast Ashby with "the autonomy-autopoiesis framework" that we can tag with the name Francisco Varela. Now all would agree that Ashby and Varela are somewhere along the same left-hand branch of this major fork, and most would agree that even where not explicitly constructivist this branch is at least constructivism-friendly. I would claim that the main differences between Ashby and Varela relate basically to their earlier and later positions on this branch; the latter extending and going beyond the ideas of the former but with no fundamental contradictions between them. However Franchi claims some unbridgeable and foundational gap.

5. Franchi starts (§3) by positing a distinction between simulating objects and simulating concepts. I cannot make sense of "simulating concepts", unless perhaps this can be rephrased as "simulating objects at a rather abstract level." In any modeling exercise one chooses a subset of all the properties of an object, one simplifies according to the needs and context. So when modelling traffic flow I might choose – depending on the motivation for the simulation – to simulate the specific properties (and idiosyncracies) of my car, or of any car of its class, or of some more abstract vehicle-in-general. These are shades of grey, a spectrum from less to more abstract rather than some black-and-white step-change from object to concept; but we can agree with Franchi that Ashby's

models are at the more abstract end of the scale. At the end of Chapter 2 of *Design for a Brain* Ashby says:

" …this book will attend closely to certain idealised cases… Maybe it will be found eventually that not a single mechanism in the brain corresponds *exactly* to the types described here; nevertheless the work will not be wasted if a thorough knowledge of these idealised forms enables us to understand the workings of many mechanisms that resemble them only as approximations." (Ashby 1960: page 29, section 2/17)

6. Franchi's main issue (§6) with Ashby is with what he calls Ashby's first thesis: "(1) A complex system running to equilibrium generates a complex behavior." As this stands, this makes no claims about life or cognition or autonomy or learning or adaptation, it seems (if "generates" is altered to "may generate") an innocuous uncontroversial statement about a class of mathematical models. But Franchi interprets this as a "very general and broad thesis about life" and sees the Homeostat as an exemplar for this that has non-action, or "*going back to sleep*" as its core principle.

7. An initial problem with Franchi's criticism is that Ashby (1960) does not present the homeostat, or homeostasis, as representing life or cognition in general. He makes it very clear that he is concerned very specifically with *adaptation*; what type of mechanism could explain how a curious and naive kitten, on first putting its paw into the fire and getting burnt, learns not to do this next time? But if we allow this extension to life-in-general, the central issue that Franchi raises is that in his eyes the homeostat (and by extension the Ashbian organism) is (§10) "…essentially a *passive* machine whose activity is a byproduct of its search for non-action. It is also a *contingent* machine insofar as its "search" for equilibrium will involve essentially random processes."

8. First let us demolish the *passivity* issue. Franchi repeats Grey Walter's jibe that the homeostat is a *machina sopora*, and casts Ashby as claiming that organisms seek quiescence; but this is to misattribute the entity to which the *stasis* in *homeostasis* refers. Within Ashby's framework, what the homeostat tends to maintain, in the face of perturbations, is a continuance of the agent-environment system in a form where the Essential Variables remain within their bounds. But the stasis of that property – a property of the agent-environment system – does not imply any stasis or quiescence in the organism itself. As an example, a bicycle with a homeostat-inspired control mechanism might wobble along, perhaps lurching from side to side, whilst maintaining its Essential Variable (e.g., angle from vertical) within acceptable bounds. The stasis of the EV is not the stasis of the moving bicycle; indeed the continued active control of the bicycle is essential for that stasis – it is no passive machine. An Irish peasant in the 19th century potato famine may have taken the bold decision to emigrate to America; the very active planning and travelling (to be described at the human level) is in no way contradicted by the associated (homeo-)stasis of the peasant's viability as a living organism.

9. Second, let us examine the *contingency* issue. Ashby makes clear that his focus is on adaptation, as with the kitten and fire example; and also makes explicit that he will use no teleological explanation for behaviour – to avoid circular argument. It follows that his mechanism must rely on randomness and contingency in the sense of being non-teleological. But once again we must be careful to distinguish the different levels of explanation: contingency and lack of purpose at the mechanism level is shown to enable purposive behaviour at the organism level. In the homeostat, with little sense of history and no social dimension, this will be the basic low-level purposiveness equivalent to the described adaptive behaviour of a kitten; but this lays the foundation for potentially understanding much more sophisticated purposiveness, for instance extended through personal history and social interaction. It is essential to Ashby's argument that the mechanism be not purposive yet the organism demonstrates purpose; to castigate this as describing a contingent machine is to misunderstand his argument, to miss the subtlety of describing the machine-organism at two different levels.

10. Franchi goes on to replicate some aspects of Ashby's homeostat (§21), but the points raised there do not particularly affect the mistaken central assertion that Ashby was presenting a perspective of life as a passive-contingent phenomenon; he was doing no such thing.

11. Though my own work in evolutionary robotics and cognition (e.g., Harvey et al. 1997) has been immensely influenced by Ashby, I have my own criticisms of him. I think it unfortunate that he glossed over a crucial distinction between two distinct meanings of "Essential Variables": a first (EV-viability) meaning that relates directly to the chances of survival of the organism, and a second (EV-sensory) meaning that relates to some sensory perception used by the homeostatic mechanism as the second form of feedback. Necessarily the second variable cannot be identical to the first,

although some correlation between them is necessary. A related mistake was that Ashby made both these EVs in step-function form (indeed made them one and the same step-function): whereas at least one, if not both, needs to be graded rather than stepped, for instance by equating the EV-viability to the current probability-rate of dying, thus incorporating noise. With the benefit of hindsight, with knowledge gained in catastrophe theory and the notion of phase transitions that has now reached wider public consciousness as "tipping-points", we now realise that step transitions can arise without step-functions (Harvey 2008). A further claim of Ashby's (that Franchi lists as claim 5) is, I believe, unduly and grossly pessimistic; he suggests that a homeostat of 1000 units or more would be almost certainly unstable whereas I believe the converse is true when the interactions are nonlinear and the variables refer to physically bounded quantities (Harvey 2011).

12. Despite these (and many other) technical disagreements, I see Ashby standing as a giant from the mid-twentieth century, with subsequent advances building on his achievements. The directness of his writing on the dynamical systems approach to cognition remains unsurpassed. He was clear in distinguishing between different levels of description. His work necessarily used explanations with circular causation, which can easily be misinterpreted by those who demand that these be recast in terms of linear causation. Varela and others have gone beyond his position, but to suggest this means some foundational gap seems as far-fetched as suggesting a foundational gap between Darwin in the mid-nineteenth century (who knew nothing of genetics) and the twentieth century neo-Darwinian synthesis.

13. Extensions to Ashby's framework include moving from homeostasis to homeorhesis (i.e. stability of process or trajectory, rather than stability of state).

14. One can take an autopoietic perspective on life as a self-maintaining organisation of flows of matter, driven by energy and entropy gradients while maintaining itself in a steady state far from equilibrium despite noise and perturbations. There is a close and natural fit with notions of viability and multiple feedback loops in Gaia Theory and Daisyworld models (Harvey 2011). All these later developments go beyond Ashby's relatively modest homeostat, but are very compatible with the philosophy behind it and fit well with Ashby's analysis of adaptation. Technical advances that imply discarding or revising some of his statements should not be equated with unbridgeable philosophical gaps. Ashby's shoulders are broad and well worth standing on.

References

Ashby W. R. (1960) Design for a brain. Second edition. Wiley, New York.

Franchi S. (2011) Life, death, and resurrection of W. Ross Ashby's homeostat. In: Franchi S. & Bianchini F. (eds.) The search for a theory of cognition: Early mechanisms and new ideas. Rodopi, Amsterdam: 3–52.

Harvey I., Husbands P., Cliff D., Thompson A. & Jakobi N. (1997) Evolutionary robotics: The Sussex approach. Robotics and Autonomous Systems 20: 205-224.

Harvey I. (2008) Homeostasis via chaos: Implementing the uniselector as a dynamical system. In: Bullock S., Noble J., Watson R. A. & Bedau M. A. (eds.) Proceedings of the 11th International Conference on Artificial Life. MIT Press, Cambridge MA: 774.

Harvey I. (2011) Opening stable doors: Complexity and stability in nonlinear systems. In: Lenaerts T., Giacobini M., Bersini H., Bourgine P., Dorigo M. & Doursat R. (eds.) Advances in artificial life, (ECAL 2011). MIT Press, Cambridge MA: 318–325.

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

Received: 3 October 2013, Accepted: 8 October 2013