« 8 » Proponents of the enactive approach agree that the homeostat's double feedback architecture made an important contribution, but at the same time they struggle to overcome its limitations (Ikegami & Suzuki 2008). Franchi briefly refers to evolutionary robotics models inspired by this ultrastability mechanism, but he does not mention that further progress has been difficult. Ezequiel Di Paolo (2003) showed that implementing Ashby's mechanism is a significant step toward more organism-like robots, but the *contingent* link between (internal) homeostatic mechanisms and (external) behavior is a continuing source of problems.

« 9 » Switching the connections of an ultrastable system may in some cases lead to an adaptive re-organization of original behavior, as Franchi's model of two self-connected homeostat units shows, but it is far from guaranteed in more complex systems. It is possible that the switch is not sufficient to "break" the essential variables, thereby failing to provoke an adaptive response at all, or, if adaptation does occur, the new equilibrium may not reestablish the original behavior. From Ashby's perspective, this failure to recover the original behavior is not a problem, since he simply equates all equilibriums with survival – without asking how they might differ in terms of desirability (an idea that is later echoed by Maturana's claim that autopoiesis does not admit gradations).

« 10 » However, this neutrality is not found in nature. For example, inverted goggle experiments can be safely conducted without any irreversible consequences to the participants. How then is the appropriate reconfiguration achieved? There have been a couple of attempts at addressing the practical problems stemming from contingent dependencies with improved evolutionary robotics models, in particular by more closely tying the desired behavior to the satisfaction of the homeostatic conditions (Iizuka & Di Paolo 2008; Iizuka et al. 2013). However, these attempts have met with only partial success.

« 11 » We need to move beyond the traditional fixation on equilibrium dynamics in order to make real progress on these issues. According to Franchi, an Ashbyan organism is continuously required to behave so as to go to equilibrium, since "its always possible failure to do so will necessarily result in the homeostat's death" (§10). However, at the level of its physical body, an organism is always far-from-equilibrium with respect to its environment. Falling into an equilibrium is the same as dying, because the organism would lose its ability to do the work of self-producing its own material identity, i.e., the very process that ensures that the double feedback loop between behavior and internal homeostasis is intrinsically connected within a whole. Only non-living matter can be in physical equilibrium with its environment. Franchi conflates the living with the non-living in another way when he notes that the homeostat "will continue acting the way it normally does until an outside force compels it to change course of action" (§26). This is precisely the way in which Newton's first law of motion describes the behavior of objects. However, living beings – subjects – can change their behavior even in the absence of a change in external conditions.

Overcoming the autonomy/heteronomy dichotomy

« 12 » Franchi enticingly hints at the possibility of a "reevaluation of the traditional distinction between heteronomous and autonomous behavior" (Abstract). Is perhaps the *relational* view of life defended by the enactive approach (e.g., Di Paolo 2009) such a reevaluation? As I see it, there is a natural development of ideas from first-order cybernetics via second-order cybernetics toward enactive cognitive science. The early approach was too heteronomous, while the second-order approach was too autonomous, but the relational stance of the enactive approach formulates a dialectic between these two extremes (Froese 2011; Froese & Stewart 2013). Relatedly, it aims for a middle ground between the extreme of computationalism, which explicitly represents goals, and the extreme of Ashby's (and Maturana's) insistence on absolute non-contingency. Crucial for these two theoretical shifts is to go beyond traditional theories that portray life's ultimate purpose, no matter whether by active or passive means, as equilibrium, stability, or survival. As Hans Jonas recognized, only when we face up to the essential precariousness of living existence can we hope to understand the meaning of life:

> "the survival standard itself is inadequate for the evaluation of life. If mere assurance of permanence were the point that mattered, life should not have started out in the first place. It is essentially precarious and corruptible being, an adventure in mortality, and in no possible form as assured of enduring as an inorganic body can be. Not duration as such, but 'duration of what?' is the question." (Jonas 2001: 106)

Tom Froese graduated with a D.Phil. in cognitive science from the Univ. of Sussex. He is currently a postdoctoral researcher at the National Autonomous University of Mexico. He investigates the relationship between life, mind and sociality using a variety of methods, including evolutionary robotics, systems theory and phenomenology.

Received: 2 October 2013
Accepted: 18 October 2013

For Biological Systems, Maintaining Essential Variables Within Viability Limits Is Not Passive

Matthew Egbert
University of Hertfordshire, UK
mde/at/matthewegbert.com

> Upshot • The target article proposes that Ashby's investigations of the homeostat and ultrastability lead to a view of living systems as heteronomous, passive "sleeping" machines and thus are in fundamental conflict with concepts of autonomy developed by Jonas, Varela and others. I disagree, arguing that (1) the maintenance of essential variables within viability limits is not a passive process for living systems and (2) the purpose of Ashby's investigations of the homeostat was to investigate adaptivity, a subject that is related to, but clearly distinct from, autonomy. As such, I find Ashby's work on adaptivity to be neither in opposition to nor in direct support of modern concepts of biological autonomy and suggest that a productive way forward involves the investigation of the intersection between these two fundamental properties of living systems.

« 1 » Stefano Franchi distinguishes between *models of objects*, which are used to evaluate the accuracy of our understanding of real-world systems, and *models of concepts,* which are used to refine, explore or explain theoretical constructs. He observes, rightly, that the two approaches are not incompatible and can in fact be complementary, provided that they are not conflated.[2] I agree with Franchi that the homeostat is most productive when it is considered as a model of a concept rather than a model of the brain. However, I may disagree with Franchi over what concept the homeostat is intended to be a model *of.*

« 2 » I see the homeostat as a system for investigating adaptivity. "How does the brain produce adaptive behaviour?" is the opening sentence of Ross Ashby's *Design for a Brain* (1960), and, in my view, the goal of the book. Franchi appears to have a different interpretation, seeing the homeostat as targeting a broad view of life: "Ashby's device is a proxy for a view of life as generalized homeostasis, whose principles it embodies" (§11). I will argue below that if the homeostat is seen as targeting "just" adaptivity, and not a broader concept of life in general, a more consistent view emerges of Ashby's research, and of the subsequent research that integrates Ashby's work with concepts of biological autonomy.

« 3 » Franchi's argument that the homeostat is passive and heteronomous begins with William Grey Walter's comment that the homeostat is a sleeping machine "because its ideal function is to go back to sleep as quickly as possible" (§10). Franchi observes that this does not take into account the complexity of the behaviour that the homeostat can generate but, other than that, he seems to agree:

> "the main source of resistance to Ashby's thesis about life lies […] in its effort to derive action from non-action [and this] means something philosophically far more disturbing: the homeostat will act as much as is needed in order to *go back to sleep*. It will go to extraordinary lengths, it will take whatever action it can – just in order to go back to rest (i.e., to equilibrium)." (§10, emphasis in original)

« 4 » Here, Franchi is suggesting that because the homeostat "derives action from non-action," it is passive; but there is some confusing language used here that I believe leads to an incorrect view of the homeostat as passive. In particular, there is a conflation with the idea of (1) a variable being at a steady state (a concept that is more general than that of equilibrium and that allows for stability, i.e., an absence of change, at conditions that are far away from equilibrium) and (2) the system being inactive, passive, or "asleep." The homeostat is capable of the former, but is described in the quote above as doing the latter. It is not the same thing to say:

A: "The homeostat will act as much as is needed in order to go back to sleep." (Franchi §10)

B: The homeostat will act as much as is needed in order for an essential variable to remain (roughly) the same. (My interpretation of Ashby)

« 5 » In *Design for a Brain*, Ashby clearly presents a system that performs as described in (B), but this system is being described by Franchi according to (A). The error in (A) is apparent when we recognise that the maintenance of an essential variable within viability limits can require substantial dynamic activity, or as Ashby puts it: "the constancy of some variables may involve the vigorous activity of others" (Ashby 1960: 67). We can consider, as a metaphor, the act of balancing a vertical pole on your finger tips – the pole remains roughly vertical, but it would be a mistake to describe the system balancing the pole as passive or asleep – constant effort is needed to counteract the system's inherent tendency *not* to remain at that state.

« 6 » Nowhere is the necessity of activity for the maintenance of essential variables more evident than in living systems. All organisms are far-from-equilibrium dissipative structures that require ongoing acquisition of energetic and material resources to counteract their degradation (Gánti 2003; Schrödinger 1944). If we think about archetypal essential variables, such as body temperature or the levels of sugar and oxygen in the blood, it is clear that these all require ongoing organismic *activity* if they are to be maintained within viability limits. To assume the maintenance of these within viability limits as being passive is a mistake.

« 7 » This leads to why I disagree with the proposition in §§13–14 that Ashby's work opposes the relatively recent frameworks developed for studying autonomy. Ashby developed a framework for defining, describing and studying adaptive behaviour. Hans Jonas, and later Francisco Varela and Humberto Maturana, developed conceptual frameworks for understanding how a system can be autonomous, how its own needs can emerge and how these can be satisfied by the system itself. There is overlap between these frameworks: for instance, essential variables play central roles in both frameworks; however, as observed by Ezequiel Di Paolo (2005), the concept of autopoiesis does not automatically entail adaptivity. Similarly, the concept of adaptivity developed by Ashby does not entail autonomy. Put another way, it is possible to conceive of an autopoietic or autonomous system that is adaptive, or one that is not adaptive; and conversely, it is possible to conceive of an adaptive system that is autopoietic or autonomous, or one that is not. The two properties are orthogonal despite their overlap.[3] In fact, the homeostat is an example of a system that is adaptive, but not autonomous. The "essential variables" of the homeostat (the state of its magnets) are not actually essential to its existence or operation. They are only labelled as essential

2 | As a side note, I do not entirely agree with Franchi's comment that "sometimes, seemingly technical computationally-intensive work in the cognitive sciences is (also) philosophy in disguise […] presented in a technical garb more suited to a positivistic-minded age that shies away from abstract conceptual frameworks unless they are presented under scientific-looking covers" (§4). I see *computational philosophy* in a less duplicitous light, where computational models do not hide philosophy, but are rather provide valuable new methods for philosophers to explain, criticise and investigate the implications of their own frameworks and those of others.

3 | Here, I am arguing that the concepts of adaptivity and autonomy should be kept as distinct. However, it may be that in practice, artificial and/or natural organisms must have some degree of adaptivity (or minimal dynamical robustness to a variable environment) if they are to persist, in which case, all real-world instances of autonomous systems would also be adaptive systems.

because Ashby simply assumes the existence of essential variables, and then investigates how they can be regulated, without delving into the details of the intrinsic dynamics of the essential variables (for example the tendency of blood sugar levels, in the absence of behaviour, to approach non-viable states). When a magnet of the homeostat leaves the predefined viability limits, the system continues to operate. For this reason, the homeostat itself is actually heteronomous, but because the homeostat is a model of an adaptive system and not a model of an autonomous or living system, the heteronomy of the homeostat says nothing about the heteronomy or autonomy of living systems.

« 8 » For Ashby's investigations, it sufficed to consider the essential variables in only very abstract terms – a variable that must be maintained within limits; nothing more. Because the essential variables were not included in detail, the homeostat is not particularly effective at demonstrating the dynamic nature of the maintenance of essential variables. It was therefore possible for the homeostat to be mistaken for a sleeping, passive machine that does everything it can "to do nothing." If essential variables were modelled in more detail and the intrinsic dynamics of the essential variables of dissipative structures such as life were included, it would have been more obvious that Ashbian adaptive regulation, whenever employed by a biological system, must be anything but passive.

« 9 » For reasons such as these, there needs to be more work modelling the homeostat and its interesting form of adaptation. In the latter section of the target article, Franchi pointed out that Ashby's "environments," when simulated as homeostat units, included an inappropriate, or at least odd, property of self-regulation. This is an important observation that raises questions that can be investigated using models such as that presented in the target article. Similarly, I have argued here that there are assumptions implicit in Ashby's work concerning the nature of essential variables that need to be made explicit and investigated. To understand how adaptive behaviour relates to autonomy and agency, we need to develop a more sophisticated understanding of essential variables, their intrinsic dynamics, the emergence of viability limits and how mechanisms of adaptivity can respond to essential variables to prevent catastrophic system failure. Some work in this area is already underway (Barandiaran & Egbert 2013; Egbert 2013; Egbert, Barandiaran & Di Paolo 2010; Egbert, Di Paolo & Barandiaran 2009), and further developments will not only help us to understand how life differs from nonlife, but also how life could have originated (Ruiz-Mirazo, Pereto & Moreno 2004).

Acknowledgements

I would like to thank Kepa Ruiz-Mirazo for our discussion related to this commentary.

Matthew Egbert: "A prominent theme in my research is the recursive self-maintaining organisation of life, and how this relates to its profound adaptability. This draws me toward the investigation of synthetic protocells and the origin of life, where mechanisms of behaviour are at their simplest, and towards minimally cognitive robotics and simulated agents, where fundamental concepts underlying adaptive behaviour and autonomy can be clearly defined and examined. As a Research Associate at the University of Hertfordshire, I am currently investigating self-maintaining behaviours in autonomous robots and simulated agents." Homepage at http://www.matthewegbert.com

Received: 14 October 2013
Accepted: 18 October 2013

Interpreting Ashby – But which One?

David Vernon
University of Skövde, Sweden
david.vernon/at/his.se

> Upshot • The association of heteronomy with Ashby's work in the target article follows from a direct interpretation of the second edition of Ashby's book Design for a Brain. However, the first edition allows for an alternative – opposite – interpretation that is compatible with autonomy and autopoiesis. Furthermore, a more balanced perspective is suggested to avoid unintentionally giving the casual reader a misleading impression that the homeostat is Ashby's ultimate position on homeostasis and that it is an adequate model of the brain.

« 1 » The target article claims that Ross Ashby's generalized homeostasis thesis entails that living organisms are heteronomous rather than autonomous, being controlled by the environment rather than independently adapting to environmental perturbations. In the following, I will explain why such a conclusion is consistent with the second edition of Ashby's book (1960). However, I will also argue that an interpretation of the first edition (Ashby 1954) can lead one to the opposite conclusion, one that supports autonomy rather than heteronomy and one that is compatible with the principles of autopoiesis. I also wish to highlight the priority of multistability over ultrastability, and the associated limitations of the homeostat and simple ultrastability (which Ashby himself acknowledges, albeit in a more obvious manner in the first edition.).

« 2 » In the following, when referring to work in *Design for a Brain,* first edition (Ashby 1954), I will adhere to Ashby's convention of using, e.g., S. 3/9 to refer to Chapter 3, Section 9. I will add a leading superscript (^{1}S. 3/9 vs. ^{2}S. 3/9) to differentiate between the first and second editions, respectively. I have also retained Ashby's use of italics; words originally in bold face type are <u>underlined</u>.

Heteronomy vs. autonomy

« 3 » In the abstract, Stefano Franchi states that Ashby's thesis "entails that life is fundamentally 'heteronomous.'" While Ashby does not use the term "heteronomy" in either of his books, this conclusion follows naturally from Ashby's development of generalized homeostasis in the second edition (^{2}S. 5/6), according to which an organism and its environment form a single state-dependent system (^{2}S. 3/9, ^{2}S. 3/10), the variables of which include a set of essential variables, the value of which much be kept within certain bounds if the organism is to survive (^{2}S. 3/14).

« 4 » Homeostasis, the process of regulating the essential variables, requires the organism to adapt to its environment to achieve stability, i.e., to keep the essential variables within physiological limits (^{2}S. 5/3). A region in the system's field