The Biological Foundations of Enactivism: A Report on a Workshop Held at Artificial Life XV

Eran Agmon*,**
Columbia University

Matthew Egbert†
Harvard University
University of Auckland

Nathaniel Virgo‡
Tokyo Institute of Technology

Abstract This is a report on the Biological Foundations of Enactivism Workshop, which was held as part of Artificial Life XV. The workshop aimed to revisit enactivism's contributions to biology and to revitalize the discussion of autonomy with the goal of grounding it in quantitative definitions based in observable phenomena. This report summarizes some of the important issues addressed in the workshop's talks and discussions, which include how to identify emergent individuals out of an environmental background, what the roles of autonomy and normativity are in biological theory, how new autonomous agents can spontaneously emerge at the origins of life, and what science can say about subjective experience.

Keywords
Enaction, modeling, autonomous agents, origins of life

1 Introduction

The Biological Foundations of Enactivism Workshop was a one-day event that was held as part of Artificial Life XV on July 4, 2016 in Cancún, Mexico. It continued a tradition of enaction-themed workshops that have run alongside or as part of Artificial Life and related conferences, including: the *Workshop on Simulation Models of Autonomous Systems*, held in 2007 at Artificial Life X [4]; *Agency in Natural and Artificial Systems*, held as part of the International Conference on the Simulation of Adaptive Behavior, 2008 [15]; and *WAAT: The Workshop on Artificial Autonomy*, held at the European Conference on Artificial Life 2011.

While enactivism is a growing force in cognitive science, this workshop aimed to revisit enactivism's contributions to biology. More specifically, the workshop's goals were to (1) revitalize the discussion of autonomy, with the goal of grounding it in quantitative definitions based in observable physical phenomena, and (2) revisit the topic of the origins of life from an enactive perspective, by focusing on the question of how a new autonomous agent can emerge where none existed before; or perhaps equivalently, how properties such as autopoiesis or normativity can first emerge from the abiotic physical world.

* Contact author.
** Department of Biological Sciences, Columbia University, New York, NY 10027. E-mail: agmon.eran@gmail.com
† Department of Earth and Planetary Sciences, Harvard University, Cambridge, MA 02138; Department of Computer Science, University of Auckland, Auckland, New Zealand. E-mail: mde@matthewegbert.com
‡ Earth-Life Science Institute (ELSI), Tokyo Institute of Technology, 2-12-1-IE-1 Ookayama, Meguro-ku, Tokyo 152-8550, Japan. E-mail: nathanielvirgo@gmail.com

 Artificial Life 24: 49–55 (2018) doi:10.1162/ARTL_a_00253

2 Background

Enactivism was originally founded on the notion of autopoiesis (later generalized to autonomy), which identifies living agents with their self-producing mode of organization [17, 16]. That is, life is conceived of in terms of a network of ongoing processes (such as chemical reactions) that constantly re-produce the conditions required for those same processes to occur.

Thus, traditionally, the definition of autonomy has rested upon abstractions such as "processes," "organizational closure," and "distinguishable unities." A key question of the workshop was how to ground these abstractions in terms of what can be measured in real physical systems. Participants aimed to focus discussion on how a self-producing mode of organization can arise from the physics and chemistry of our world. Other questions included: What are the relevant processes in living systems, and how can they be measured? How is organizational closure actually realized in our physical world? Is autonomy absolute, or does it come in degrees?

3 Presentations

A call for abstracts resulted in a number of submissions, which were reviewed by the workshop's organizers for suitability, and selections were invited to be presented as 20–30-minute talks. These talks were framed by invited contributions including presentations from three of the main-track conference's keynote speakers—Randall Beer, Ezequiel Di Paolo, and Mark Bickhard—an indication perhaps, of enactivism's strong and growing relevance within the artificial life community. Here, we provide a brief summary of a subset of the presentations. The workshop covered a broader range of topics than those included below, and our selection reflects a focus on a particular set of topics for this report.

In his introductory presentation, "How can the enactive approach inform biology?", **Eran Agmon** advocated for more scientific engagement with the enactive concept of organizational closure (briefly described in the background section above), suggesting that biologists could look for relevant processes in biological individuals and evaluate the extent to which they are organized in closures. Enactive biologists can begin by looking for the relevant processes of minimal biological individuals, measuring them, and showing whether they manifest as closures – how closures are identified still needs to be determined. Given that this is achievable, they can ask what types of organizationally closed systems can be considered agents, and how these can be differentiated from non-agents. With individuality and agency grounded in biology, scientists can then look for other forms of autonomy, such as sensorimotor habits, and investigate the relation of biological and non-biological autonomy, and whether non-biological forms can arise independently of biology.

Ezequiel Di Paolo identified conflicting requirements for idealized forms of self-production and self-distinction, arguing that the former involves maximal environmental interaction (where the system uses all possible interactions to support self-production), and that the latter suggests minimal environmental interaction. Di Paolo suggested that agents are thus engaged in a constant struggle,—the system at some times favoring openness, at some times closure—and that levels of interiority and exteriority are formed by agents to support this operation [10].

Randy Beer presented recent research on the origins of gliders in the Game of Life (GoL) cellular automaton. This work investigates the relationship between the proportion of on cells in a randomly seeded GoL, and the resulting number of gliders. Glider origins is part of Beer's ongoing work where the sometimes difficult concepts that are central to enaction are made more explicit and accessible through their modeling and interpretation in the alternative "physics" of Conway's GoL (see, e.g., [6]). Other work within this broader project include the characterization of gliders' autopoiesis [8], cognitive domain [7], coupled interactions, and the ontogenies of more chemically realistic protocells [2, 3].

Mark Bickhard contrasted interactivism and enaction, focusing on the process metaphysics of emergent organizations. Bickhard proposed that in contrast to self-organization, self-maintenance requires function, in which the environment is utilized to bring about conditions of maintenance. Bickhard then described representational truth values, in which agents can recognize indications of possible actions within a context, and thus make truth-based decisions. This he argues, underlies the emergence of normative representations, allowing agents to develop webs of normative contingencies that they navigate to maintain their organization [9].

Inman Harvey presented a talk entitled "I Don't See What the Problem Is," where he emphasized the relational nature of autonomy, and observed that the emergence of autonomous systems, such as hurricanes, is commonplace and not a great mystery. Harvey called for operational tests to distinguish between autonomy and agency, suggesting that in their absence the distinctions being made are often meaningless. In a prepared response, **Matthew Egbert** argued that in addition to autonomy there are likely other *relational* descriptions of life (such as agency) that will be useful in understanding how life works and that these are worth looking for. In response to the call for operational tests, Egbert suggested that operational tests for agency might look similar to the methods used by experimentalists to distinguish between metabolism-based and metabolism-independent forms of bacterial chemotaxis, such as those performed by Adler [1] in the 1960s.

Mario Villalobos, in a talk entitled "Enactivism, Autopoietic Theory, and the Trivial Nature of Living Beings," rejects the recent "normative turn" in enaction [5], arguing that living systems are trivial thermodynamic machines and that their normativity, teleology, and intentionality do not exist in any "real" sense, but are merely ascribed properties or illusions. In this vein, Villalobos called for a return to the non-teleological interpretation of autopoiesis by Maturana [18].

Matthew Egbert presented a recent publication [12] that shows how agents, seen from an enactive perspective as systems that regulate their environment in response to their own viability dynamics, can adapt to changes in their own organization, resulting in more evolvable systems, providing an example of how enactive theory can contribute to evolutionary biology.

Dobromir Dotov, in collaboration with **Tom Froese**, proposed that non-well-founded set theory might be a useful tool in distinguishing between the variety of related but distinct notions of closure that have emerged in the enactive and artificial life communities (e.g., autocatakinetics, operational closure, and Rosen's closure to efficient cause). Such a method of formal comparison between these different models can systematize the influences of enactivism and identify what properties they share and where they diverge.

Nathaniel Virgo argued that there are some nonliving natural phenomena that share important features with living organisms. He focused on the example of a hurricane as an emergent self-producing and self-delimiting network of processes. He argued that the differences between organisms and these abiotic examples are all differences in degree rather than in kind, and that they all stem from life's ability to accumulate features over evolutionary time. This led to a suggestion that understanding *open-endedness* is currently a key challenge for enactive accounts of the origins of life [14].

4 Discussion

The talks were stimulating, but perhaps the highlight of the workshop was the discussion that took place after the talks and continued informally throughout the week. One gathering, in which an assembly of many of the workshop participants emerged in the hallway and maintained itself for a couple of hours, was later referred to as the Second Workshop. The topics discussed in these informal meetings included those described above, but also ventured into other areas, which we will try to recapitulate here.

4.1 Identifying Individuals

A recurring topic of conversation, and indeed a long-standing problem for autopoiesis, is concerned with the ontological status of biological individuals—is individuality a measurable and objective quality of living matter, or is it a subjective construct, created by scientists out of practical consideration? This is not a merely philosophical dilemma, but a scientific one with far-reaching implications. Participants who were convinced individuality is a natural and objective quality suggested that we need to determine appropriate algorithms for extracting the system-environment boundaries that support individuality. Those who identified individuality as a subjective construct leaned towards a methodology for formalizing more inherently observer-dependent notions, in which the partitioning of the system into an agent and its background is always a choice that could be made differently by different observers. This is analogous to (and perhaps an example of) the concept of coarse-graining in physics.

One difficulty facing the quantification of individuality is the constant change seen in living agents. An individual's identity is not material—a living agent's material turnover guarantees that the matter present within an agent at one point in time is different from the matter present at a different point in time. Identity is not location—agents are motile and will be found in different locations at different points in time. Identity is not a closed set of states, as living agents likely never return to the same exact state twice. Morphology also varies throughout development, and so the shape and distribution of biological processes is not the factor determining individuality. Yet perhaps there is still persistence at a more abstract level of organization, and individuality is a certain type of coordinated pattern that is preserved amongst the ever-changing physical components. We believe that looking for empirical methods for identifying and measuring these organizations should be a primary goal for enactivists.

4.2 The Ontological Status of Autonomy and Norms

Among enactivists, the autonomy of a system is not a description of its material components, but instead a description of how its components interact or relate. A number of other central concepts in enaction are similarly *relational properties*. The relational nature of these concepts has caused some to question their ontological status. Is autonomy "real"? Do norms "exist" in the world, or are they just ascribed properties, mere artefacts of how we have chosen to describe the system? These topics were the center of numerous discussions.

For some, skepticism stems from the observation that at different levels of description, the property of interest (e.g., autonomy) disappears, suggesting that it is merely an artefact of how we are choosing to describe the system. For others, this observation is not a concern, as they believe that either (i) all observations involve artefacts of how we choose to describe them and/or (ii) other established scientific descriptions suffer from a similar defect of disappearing when described at a different scale. As a case in point of the latter, we raised the example of temperature, a statistical description of the kinetic energy of atoms that loses meaning when the system is seen at the level of an individual atom, never mind its constituent quarks. Despite not existing at the micro level of atoms, temperature is a central property of thermodynamics, with laws governing its relations with pressure, volume, and chemical reaction affinities.

The view that concepts such as autonomy and norms are macroscopic in this way was appealing to most of the participants. However, one of enactivism's goals is to ground first-person experience [13], and if this can only be done with respect to an observer-dependent coarse-graining, then it leads to a potentially problematic circularity. Within this view it seems hard to talk about an agent's perspective from the agent's own point of view, rather than from the point of view of an external observer. Resolving this issue seems like an important challenge for future work.

A related branch of discussion focused on an oft-cited enactive definition of adaptation [11], where (coarsely summarized) adaptation is essentially defined by its slower rate of change. A system's actions or behavior is its fast-time-scale dynamics, and its adaptation is the slower-time-scale transformation of those behavioral tendencies. One criticism of this view that was raised was that the range of time scales associated with adaptation can overlap the time scales associated with the system's non-adaptive action, and that it will thus not always be possible to separate them. For some (but not all!), these kinds of concerns suggest that some of the existing definitions are meaningless and/or in need of improvement, for example, via the specification of operational tests for identifying the presence of the property of interest.

Similar arguments were had pertaining to the ontological status of norms. Mario Villalobos argued that norms are at best epiphenomena, and more likely merely illusions that we should not be fooled into thinking are real, while Ezequiel Di Paolo argued that despite norms' having no localized physicality, they constrain and otherwise influence system behavior in clear and measurable ways, thereby qualifying as "real."

4.3 Origins

One of the workshop's key questions was how these properties—identity, autonomy, and norms—can arise where none existed before. This must have happened at least once on the early Earth, and according to several versions of enactivism it has also happened repeatedly since then, with the emergence of multicellular organisms, autonomy at the behavioral level, the formation of new social structures, and so on. For this reason the origins problem is important not just for enactivism's contributions to the origins of life, but it also directly for understanding cognition.

Traditionally, concepts such as autonomy and autopoiesis have been seen as all-or-nothing binary distinctions. This makes it difficult to think about emergence, since by definition an autonomous system would have to appear fully formed; there can be no "almost autonomous" system that could precede it. This problem led to a discussion of whether enactivism's key concepts can be "softened" to admit degrees of autonomy. Perhaps an organization can be composed of a more flexible, less rigidly defined set of processes, resulting in a weaker form of autonomy that could spontaneously emerge more easily.

A related key concept was the role of the observer in determining whether a system is autonomous, as discussed in Section 4.3 above. On the one hand, if these concepts are observer-dependent, then the problem of emergence is greatly reduced, since the question becomes "at what point in time are we willing to call the system autonomous?" rather than "at what point in time does an (objectively) autonomous system arise?" However, in moving to such a relative point of view one sacrifices some of the central claims of autopoietic theory, and not every participant was willing to go this far. Many still believe an objective definition is possible, in which biological autonomy can be derived from chemical principles without the need for an observer.

4.4 Internal versus External Perspective

Enactivism has a history of interest in the nature of phenomenology, or the experience of being a subject. Participants discussed whether a new type of science would be required for explaining subjective experience. The important question raised is "what can the external, empirical approach say about the internal perspective of a living subject?"

Some participants proposed that we can still look at what the environment *means* for biological individuals without asking about their internal perspective. Scientists can simulate a world with simple physical laws, look for conditions in which individuals emerge, and study how they respond to their environment—how environmental differences influence their behavior and viability. If the individuals behave in a way that strengthens them, they can be said to be adapted to the environment and recognize the significance of given situations. With increasing computational ability and improved analytic methods, this approach may one day extend to the real physical world and real organisms.

The authors of this report identify four distinct perspectives on the topic that arose in the discussion: (1) We cannot take an internal perspective of the subject—to see the world from its subjective perspective—and must proceed with the objective "god's-eye" perspective of the world. (2) We can approach subjectivity with the current scientific method, by systematically investigating living agents and comparing their experiences. (3) We can approach subjectivity, but it will require a new, interpersonal approach to science that takes phenomenology as its foundation. (4) We can approach subjectivity, but it will require a new formal framework that properly takes account of the macroscopic-microscopic distinction.

5 Conclusion

The workshop helped focus concepts developed in enactive cognition on problems facing biology and origins of life. The key points identified were the need for theoretical definitions of individuality, the relation between micro-level chemical properties of living systems and the macro-level phenomena of biological agents that appear to act within the environment according to norms, and how these interactions give rise to internal perspectives and biological significance.

It was broadly agreed that in order to answer these questions within the domain of biology, the central notions of autopoietic and enactive theory will need to be made more precise. Enactive biologists will need to identify relevant chemical processes, so that they can determine whether organizational closure does or does not hold – and thus whether a given system can be objectively identified as an individual. If individuals are identified, they can then be studied relative to their environments to determine how the environment influences their behavior and their organizational closure. At this point, the question of norms will become relevant. Can norms be addressed from a purely objective point of view, or will there always be an element of observer dependence involved in talking about them scientifically—or should the concept be cast aside altogether?

The study of the origins of life would greatly benefit from a new enactive theory of this kind, expressed with greater precision and more suited to empirical application. Knowing in more detail what an autonomous agent is would enable us to talk more precisely about how agency could first arise in the abiotic world, and would allow greater communication between the enactive approach to origins of life and the more empirical approaches taken by biologists and geochemists.

Acknowledgments
This work was partially funded by the ELSI Origins Network (EON), which is supported by a grant from the John Templeton Foundation, though the views expressed here are not necessarily shared by the foundation. EON also kindly sponsored the workshop.

References

1. Adler, J. (1969). Chemoreceptors in bacteria. *Science*, *166*(3913), 1588–1597.
2. Agmon, E., Gates, A. J., & Beer, R. D. (2016). The structure of ontogenies in a model protocell. *Artificial Life*, *22*, 499–517.
3. Agmon, E., Gates, A. J., Churavy, V., & Beer, R. D. (2016). Exploring the space of viable configurations in a model of metabolism–boundary co-construction. *Artificial Life*, *22*, 153–171.
4. Barandiaran, X., & Ruiz-Mirazo, K. (2008). Modelling autonomy: Simulating the essence of life and cognition. *Biosystems*, *91*(2), 295–304.
5. Barrett, N. F. (2015). The normative turn in enactive theory: An examination of its roots and implications. *Topoi*, pp. 1–13.
6. Beer, R. D. (2004). Autopoiesis and cognition in the Game of Life. *Artificial Life*, *10*(3), 309–326.
7. Beer, R. D. (2014). The cognitive domain of a glider in the Game of Life. *Artificial Life*, *20*(2), 183–206.
8. Beer, R. D. (2015). Characterizing autopoiesis in the Game of Life. *Artificial Life*, *21*(1), 1–19.
9. Bickhard, M. H. (2016). Inter- and en-activism: Some thoughts and comparisons. *New Ideas in Psychology*, *41*, 23–32.
10. Di Paolo, E. (in press). The enactive concept of life. In *The Oxford Handbook of 4E Cognition*. Oxford: Oxford University Press.
11. Di Paolo, E. A. (2005). Autopoiesis, adaptivity, teleology, agency. *Phenomenology and the Cognitive Sciences*, *4*(4), 429–452.
12. Egbert, M. D., & Pérez-Mercader, J. (2016). Adapting to adaptations: Behavioural strategies that are robust to mutations and other organisational-transformations. *Scientific Reports*, *6*(Jan.), 18963.
13. Froese, T. (2011). From second-order cybernetics to enactive cognitive science: Varela's turn from epistemology to phenomenology. *Systems Research and Behavioral Science*, *28*(6), 631–645.

14. Froese, T., Virgo, N., & Ikegami, T. (2011). Life as a process of open-ended becoming: Analysis of a minimal model. In T. Lenaerts, M. Giacobini, H. Bersini, P. Bourgine, M. Derigo, & R. Doursat (Eds.), *Proceedings of the European Conference on Artificial Life 2011* (pp. 250–257). Cambridge, MA: MIT Press.

15. Rohde, M., & Ikegami, T. (2009). Editorial: Agency in natural and artificial systems. *Adaptive Behavior*, *17*(5), 363.

16. Thompson, E. (2010). *Mind in life: Biology, phenomenology, and the sciences of mind*. Cambridge, MA: Harvard University Press.

17. Varela, F. J., Thompson, E., & Rosch, E. (2017). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT Press.

18. Villalobos, M. (2013). Enactive cognitive science: Revisionism or revolution? *Adaptive Behavior*, *21*(3), 159–167.