



HAL
open science

Open-Ended Evolution: Perspectives from the OEE Workshop in York

Tim Taylor, Mark Bedau, Alastair Channon, David Ackley, Wolfgang Banzhaf, Guillaume Beslon, Emily Dolson, Tom Froese, Simon Hickinbotham, Takashi Ikegami, et al.

► **To cite this version:**

Tim Taylor, Mark Bedau, Alastair Channon, David Ackley, Wolfgang Banzhaf, et al.. Open-Ended Evolution: Perspectives from the OEE Workshop in York. *Artificial Life*, 2016, 22 (3), pp.408-423. 10.1162/ARTL_a_00210 . hal-01371116

HAL Id: hal-01371116

<https://hal.science/hal-01371116v1>

Submitted on 3 Oct 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Open-Ended Evolution: Perspectives from the OEE Workshop in York

Abstract We describe the content and outcomes of the First Workshop on Open-Ended Evolution: Recent Progress and Future Milestones (OEE1), held during the ECAL 2015 conference at the University of York, UK, in July 2015. We briefly summarize the content of the workshop's talks, and identify the main themes that emerged from the open discussions. Two important conclusions from the discussions are: (1) the idea of *pluralism* about OEE—it seems clear that there is more than one interesting and important kind of OEE; and (2) the importance of distinguishing observable *behavioral hallmarks* of systems undergoing OEE from hypothesized *underlying mechanisms* that explain why a system exhibits those hallmarks. We summarize the different hallmarks and mechanisms discussed during the workshop, and list the specific systems that were highlighted with respect to particular hallmarks and mechanisms. We conclude by identifying some of the most important open research questions about OEE that are apparent in light of the discussions. The York workshop provides a foundation for a follow-up OEE2 workshop taking place at the ALIFE XV conference in Cancún, Mexico, in July 2016. Additional materials from the York workshop, including talk abstracts, presentation slides, and videos of each talk, are available at <http://alife.org/ws/oeel>.

Tim Taylor*†
Mark Bedau‡
Alastair Channon§
David Ackley††
Wolfgang Banzhaf‡‡
Guillaume Beslon§§
Emily Dolson***
Tom Froese†††
Simon Hickinbotham‡
Takashi Ikegami‡‡‡
Barry McMullin§§§
Norman Packard***
Steen Rasmussen††††
Nathaniel Virgo‡‡‡‡
Eran Agmon§§§§
Edward Clark†
Simon McGregor***
Charles Ofria***
Glen Ropella†††††
Lee Spector‡‡‡‡‡
Kenneth O. Stanley§§§§§
Adam Stanton§
Christopher Timperley†
Anya Vostinar***
Michael Wiser***

Keywords

Open-ended evolution, ongoing evolution, perpetual novelty, adaptive evolution, dynamical hierarchies, major transitions

* This paper was primarily written by the first three co-authors [T.T., M.B., A.C.] based upon material presented and discussed by participants of the First Workshop on Open-Ended Evolution (OEE1) during the European Conference on Artificial Life 2015 at the University of York, UK. Additional material and comments were provided by the other presenters at the workshop [D.A., W.B., G.B., E.D., T.F., S.H., T.I., B.M., N.P., S.R., and N.V.] and by the other coauthors, who contributed to discussions at the workshop or afterwards online.

* Contact author.

† University of York, UK. E-mail: tim@tim-taylor.com (T.T.)

‡ Reed College, USA. E-mail: mab@reed.edu (M.B.)

§ Keele University, UK. E-mail: a.d.channon@keele.ac.uk (A.C.)

†† University of New Mexico, USA.

‡‡ Memorial University of Newfoundland, Canada.

§§ Université de Lyon, France.

*** Michigan State University, USA.

††† National Autonomous University of Mexico.

‡‡‡ University of Tokyo, Japan.

§§§ Dublin City University, Ireland.

**** ProtoLife, Inc., USA.

†††† University of Southern Denmark.

‡‡‡‡ Earth-Life Science Institute, Japan.

§§§§ Indiana University, USA.

***** University of Sussex, UK.

††††† Tempus Dictum, Inc., USA.

‡‡‡‡‡ Hampshire College, USA.

§§§§§ University of Central Florida, USA.

I Introduction

From the first experiments with digital evolution in the 1950s to the increasingly sophisticated simulations of the present day, the concept of *open-ended evolution* (OEE) has been a central concern for artificial life (ALife) researchers [51]. Loosely defined, an open-ended evolutionary system is one that is capable of producing a continual stream of novel organisms¹ rather than settling on some quasi-stable state beyond which nothing fundamentally new occurs. Some definitions of OEE further require that the maximum complexity of organisms in the system increases over time, or that ecosystem complexity increases. Understanding open-ended evolution remains a holy grail in ALife—and yet there remains little agreement within the community on precise definitions and measures.

There has been progress on a variety of fronts concerning OEE in the past decade. To take stock of and document recent work, and to identify key milestones for the immediate future, a workshop on *Open-Ended Evolution: Recent Progress and Future Milestones (OEE1)* was held at the European Conference on Artificial Life (ECAL 2015) at the University of York, UK, in July 2015.² The workshop aimed to create a common framework for discussing and evaluating research on open-ended evolution, and to catalyze further progress. In particular, a follow-up workshop (OEE2) will take place at the ALIFE XV conference in Cancún, Mexico in July 2016.³ The Cancún workshop will be followed by a special issue on open-ended evolution in the *Artificial Life* journal, including a comprehensive review paper on work on OEE.

The York workshop had two sessions. The first session consisted of 14 short presentations that addressed (one or more of) five tasks:⁴

1. Define key concepts concerning open-ended evolution.
2. Produce actual models that do (or do not) generate interesting kinds of open-ended evolution.
3. Find and use operational empirical—and, ideally, quantitative—measures of key kinds of open-ended evolution.
4. Demonstrate examples of (kinds of) open-ended evolution in models or natural systems.
5. Identify critical future research milestones in open-ended evolution.

The second session was an open discussion among the speakers and other attendees about open-ended evolution, focusing on the key hallmarks of various kinds of OEE and the hypothetical mechanisms that could produce those hallmarks. Section 2 below briefly summarizes the short presentations, and Section 3 highlights some central themes that emerged from the open discussion.

2 Summary of Short Presentations

In “Karl Popper, artificial life, and the curious tale of the hopeful behavioral monster,” **Barry McMullin** highlighted Karl Popper’s philosophy of evolutionary epistemology and its relevance to ALife and OEE. Focusing particularly on Popper’s work “Evolution and the Tree of Knowledge” (a chapter of his book *Objective Knowledge: An Evolutionary Approach* [37]) based upon a lecture delivered in 1961, McMullin outlined Popper’s thought experiment on how mutations in an agent’s “central propensity structure” (a hierarchical control system defining its set of skills and behaviors)

1 We use to term *organism* here to include both biological organisms and individuals in artificial evolutionary systems in software, hardware, or wetware.

2 <http://alife.org/ws/oeel>

3 <http://alife.org/ws/oee2>

4 Videos and presentation slides from this session can be accessed from the workshop website.

could guide the future evolution of the agent, leading to apparently “goal-directed” evolution. This perspective suggests that the existence of hierarchical control organization and the continuing feasibility of inheritable change at the highest control levels (including emergence of higher, newly dominating levels) may be critical to the substantive openness of evolution of complex function.

The presentation by **Wolfgang Banzhaf** on “Open-Endedness and Novelty in Evolution” started with the observation that the notion of novelty in a system must be defined with respect to a particular model. Banzhaf identified three different types of novelty: (1) novelty within a model (*variation*), (2) novelty that changes the model (*innovation*), and (3) novelty that changes the meta-model⁵ (*emergence*). He then addressed the question of whether OEE requires *unbounded novelty* or *unbounded complexity*. Observing that the universe is limited and hence cannot afford an unbounded increase in *levels of complexity*, and also that all combinatorial possibilities at any one level are bounded, he argued that novelty can still be practically unbounded if the number of levels of complexity in the system is allowed to grow (as novelties grow exponentially with complexity). Hence, Banzhaf’s position is that OEE does not require unbounded complexity, but that unbounded novelty is sufficient. He concluded with some comments on the competing roles of exponential growth and competition due to resource constraints in natural selection, and the analogous situation in hierarchical systems whereby climbing the levels of complexity introduces exponentially more possibilities, but exploration of these possibilities is restricted by resource constraints on the number of individuals that can populate higher levels. Banzhaf’s talk was based upon a forthcoming paper [3].

After highlighting some of the many different concepts associated with the term OEE in the literature, in “Requirements for Open-Ended Evolution in Natural and Artificial Systems” **Tim Taylor** proposed a high-level classification of these issues in the form of five basic requirements for a system to exhibit OEE: (1) robustly reproductive individuals, (2) individuals capable of producing more complex offspring, (3) mutational pathways to other viable individuals, (4) a medium allowing the possible existence of a practically unlimited diversity of individuals and interactions, and (5) drive for continued evolution. For each requirement, Taylor explained why it was important, what theoretical issues it encompassed, and what practical issues were involved in implementing a system to meet the requirement. The talk was based upon a paper [50] presented at the EvoEvo Workshop⁶ at the same ECAL 2015 conference.

Guillaume Beslon started his talk “Is Biological Evolution Open-Ended?” by observing that the vast majority of literature on OEE comes from the ALife community and not from evolutionary biology. Historically, the mathematical models of evolutionary biologists have focused on stable states. Moreover, selection is thought to commonly act as a stabilizing force on genetic diversity. However, although the concept of OEE is largely lacking in the biological literature, the concept of novelty pervades it in many forms. Of all kinds of biological novelty, Beslon identified *coevolution* and *major transitions* as the two being most closely related to the concept of OEE. He proposed that the most important idea of open-endedness was the emergence of novelty leading to new levels of individuality (i.e., major transitions). However, he conjectured that biology cannot be open-ended with regard to major transitions, arguing that as higher levels of organization are inevitably populated by smaller population sizes, this leads to decreasing probability of fixation of beneficial mutations. A saving grace for computational systems, according to Beslon, was that this limitation could be overcome by tricks such as suitable fitness landscapes (although whether they were open-ended would still be an open question). A version of Beslon’s argument can also be found in the forthcoming paper mentioned by Banzhaf [3].

In “Normalised Evolutionary Activity Statistics and the Need for Phenotypic Evidence,” **Alastair Channon** noted that there was widespread agreement that OEE involves the continued evolution of new adaptive traits. As this can be achieved trivially,⁷ he argued that OEE must also involve a

⁵ A meta-model is a description of the *kinds of things* that might be present in a model.

⁶ <http://evoevo.liris.cnrs.fr/evoevo-workshop/>

⁷ For example, evolving variable-length bit strings in a selective environment that favors longer strings would result in the continued evolution of new adaptive traits.

sustained increase in some measure of accumulated adaptive success. However, he questioned the inclusion of increasing complexity as a hallmark of OEE, as this would preclude the possibility of addressing important questions such as whether or not OEE can be the cause of increasing maximal complexity (whether individual, group, or system complexity) or what conditions might be necessary or sufficient for this. Channon then described his Geb system [15], based upon Harvey's SAGA principles [19] with the addition of coevolutionary feedback arising via biotic selection rather than being specified by abiotic fitness functions, followed by a description of Bedau et al.'s work on evolutionary activity statistics [4, 14]. Two useful features of these measures, he suggested, were that they are widely applicable, and that the key metric (cumulative *evolutionary activity*, based on *adaptive persistence*) is "a measure of the continual adaptive success of the components in the system" [7], that is a measure of accumulated adaptive success. When applied to Geb, these measures classify the system as producing unbounded evolutionary dynamics (OEE). Despite that, it becomes increasingly difficult (over evolutionary time) to visually observe the behaviors that evolve. Channon identified three critical future milestones for the field: (1) more systems classified as OEE according to the evolutionary activity statistics, in order to refine definitions and tests for the hallmarks of OEE; (2) evidence of complex artifacts or behaviors arising from evolutionary changes (rather than from a very small number of mutations from a hard-coded ancestor); and (3) evidence of long evolutionary sequences of evolved artefacts or behaviors (a result that has not been conclusively observed in work to date).

In "Indefinite Scalability for Open-Ended Evolution," **David Ackley** agreed with Beslon that major transitions are the most important aspect of OEE, and that if we accept that, then a finite system cannot be open-ended: Successive major transitions produce larger and slower individuals until ultimately producing a population size of one that lives forever. He argued that conventional component-based evolutionary activity measures of OEE are problematic because they require us to identify the components of interest beforehand—if we treat components as priors rather than observables, we will be unable to detect major transitions. To avoid this problem and treat evolutionary components as observables, models should be defined at the level of physics and chemistry, not at the level of biological components. But this raises the question, what kind of physics and chemistry is appropriate? Ackley's answer is that satisfactory models should, in principle, be *indefinitely scalable*. This rules out the whole class of deterministic, synchronous models (such as Game-of-Life-type systems), and suggests that OEE models should embrace nondeterminism. This approach could create a unifying research strand between different ALife projects that implemented different kinds of indefinitely scalable systems. Ackley concluded his talk by proposing a research challenge to develop a statistical OEE measure based upon identifying potential evolutionary components at a given scale by near-perfect spatial autocorrelation of elements, study of the phase space defined by the "life lines" of such components over time, and application of the same technique at different scales in the system. Ackley presented further details of his concept of indefinitely scalable architectures in a paper at the main ECAL 2015 conference [1].

In "Emergence of Emergence," **Norman Packard** discussed current work with Nicholas Guttenberg and others on the evolution of *coding* in the transition from prebiotic systems to biotic evolutionary systems. The central question being tackled is: What dynamical processes lead generically to sequestration of information into units that have long-term stability, control fast time-scale dynamics, and can serve as evolvable elements? The goal is to understand this transition well enough to be able to *engineer* systems that will naturally implement information sequestration, evolvability, and robustness. Packard observed that evolutionary dynamics is very different from attractor dynamics: It behaves somewhat like an attractor on the short term, but over a longer term, instabilities lead to the generation of innovation. Their work augments the language of dynamical systems theory with concepts capable of describing such phenomena. In particular, the concept of dynamical *canals* is used in place of attractors. A mechanism that seems to produce this kind of system generically is one involving an alternation between unstable (or neutrally stable) dynamics and contracting, fixed-point dynamics. Alternation forces the system to produce *information bottlenecks*,

which seem to imply the emergence of informationally stable components that become proto-code. In addition to Packard's own work on this origin-of-life transition, colleagues are working on applying these ideas to other transitions, including multicellularity, ecological niche formation, and the evolution of cognitive mechanisms.

The central claim of **Nathaniel Virgo** in "Open-Ended Fitness Landscapes" was that open-endedness is a property of *fitness landscapes*, and not of the process of evolution itself. He characterized OEE in terms of increasing phenotypic complexity, and argued that ecological factors (e.g., changing environments, coevolution, and niche construction) might *not* be necessary for OEE, or at least that this is a hypothesis worth taking seriously. To evaluate this claim, he suggested we should focus on understanding how to create more "lifelike" fitness landscapes of high-dimensionality, containing many qualitatively different "solutions," where fitter solutions also generally tend to be more complex, and where those solutions can be reached through a sequence of small changes. By comparison with the biological-physical case, Virgo argued that this kind of fitness landscape required the existence of many *degrees of freedom* (DoFs), which he characterized as the "capacity [of a system] to be changed in some nontrivial way." Complex systems with many DoFs, he suggested, enable the existence of many qualitatively different solutions and the capacity to move between those solutions. Virgo then hypothesized that many nontrivial landscapes have small regions of evolvability, and that evolutionary systems might evolve towards such regions through a process of the *evolution of evolvability*. His tentative conclusion was that the requirements for OEE in computational systems might just involve larger search spaces, more nontrivial fitness functions, larger populations, weaker selection pressure, and more computer time.

In "Empirical Measurements of Door-Opening Evolution of Technology," **Mark Bedau** described recent work with colleagues on studying open-ended evolution within the context of cultural rather than biological evolution. Specifically, the work investigates the evolution of human technology. Technological evolution differs from biological evolution in many ways, including the presence of hyper-parental reproduction, intentionally directed progress, and indirect (human-mediated) reproduction. Furthermore, one can identify populations of technology *adopters*, technology *designers*, technological *innovations*, and technological *products* as four distinct, but interrelated components. Technological evolution is therefore different from biological evolution in nontrivial ways, but, Bedau argued, we should not restrict ourselves to studying OEE just in biological systems. He identified the concept of *reach*—intuitively, the idea of an invention that has descendants that are very different to itself—as an important aspect of technological OEE. In recent preliminary work, the evolution of technological innovations was studied by using text-mining techniques on historical patent records to extract relevant traits in each record. Dimensionality reduction and clustering techniques were used to study the reach of different traits in particular genealogies of technologies. These new results build on earlier work on operationalizing the study of technological evolution [43, 11, 13], and they open the door to the empirical study of many questions about open-ended evolution in nature.

In "The OEE Measure—Will It Blend?" **Simon Hickinbotham** questioned whether existing evolutionary activity methods reduced the complexity of a system too much to produce a simple measure, and whether they really highlight the relevant features relating to a system's open-endedness. He argued that (improved) evolutionary activity measures were useful for making sense of the huge amounts of data produced by computational evolutionary systems, and, more importantly, they allow us to rigorously compare different systems and thereby demonstrate when improvements in ALife systems have been achieved. Hickinbotham then introduced his new *quantitative non-neutral* (QNN) evolutionary activity measure. He highlighted some of its attractive features as being that (1) it produces a single numerical value, (2) it is based solely on population data (like other evolutionary measures), and (3) it can be applied to systems with intrinsic or extrinsic fitness. The application of the QNN measure to the Tierra and Stringmol systems was described, with discussion of how it was used to guide improvements in the design of each system: Further details can be found in papers presented during the main conference [23, 22] and in a subsequently published article [20]. Hickinbotham concluded by suggesting that we need more measures to address different

aspects of OEE, and that once we had developed an adequate suite of such measures, there was the potential to create a *meta-evolver* for OEE.⁸

Steen Rasmussen in “Minimal Life and Open-Ended Evolution” conjectured that high-dimensional systems with rich object complexity and/or diversity enable the emergence of higher-order functionalities, and that these are *necessary* for OEE. However, simply adding complexity and diversity is not a *sufficient* condition. He stresses the existence of two different ways to increase complexity in a physical system: through the aggregation of things from the environment, and through the evolution of new encoded entities. His view is based upon his work over many years with *protocells*—minimal self-reproducing molecular machines comprising a metabolism, genes, and a container in a given environment. Protocells utilize self-organization and self-assembly processes to maintain their organization, and are driven by a metabolism feeding on free energy and resources from the environment. A full chemical protocell system has not yet been synthesized in the lab, but simulation results show that published protocell designs apparently lack the ability to evolve in an open-ended manner. For example, in the case of Rasmussen’s own protocell design, simulations show that the system’s evolution is limited to the optimization of its metabolic rate. Both experimental and simulation results show that a richer environment is necessary to expand the system’s evolutionary potential. Real chemical systems demonstrate the emergence of higher-order functionality at multiple hierarchical levels, and Rasmussen described simulation results in which similar higher-order functionality had emerged [38]. This was achieved by adding to the complexity of the lowest-level elements in the system. A similar approach might therefore be viable in the protocell systems. However, Rasmussen pointed out that *just* adding complexity to the system in an unprincipled way was likely to lead to “black tar” rather than any interesting higher-order behavior—the addition of complexity must be done with care. This leads to an as yet unanswered question: Are there principles to guide us in adding complexity at the right places in the system, or are we essentially left to experiment by trial and error?

Emily Dolson started her presentation “Understanding Complexity Barriers in Evolving Systems” with an informal definition of open-endedness as the ability of a system to “keep doing interesting things.” Dolson discussed how we might more accurately define both “keep doing” and “interesting things.” She suggested that there is fairly general agreement that “keep doing” means unbounded rather than asymptotic behavior of a measure. With regard to what measures to use, that is, what constitutes “interesting things,” she argued that it might be productive to flip the question around, and ask what kinds of *barriers* might *prevent* a system from exhibiting open-endedness. Dolson proceeded to describe four barriers that she and her colleagues had come up with: (1) *change potential*—how much we expect the population composition to change during an interval; (2) *novelty potential*—how many entirely new strategies we expect to arise during an interval; (3) *complexity potential*—how much we expect the greatest individual complexity to increase during an interval; and (4) *ecosystem potential*—how much we expect “meaningful” diversity to increase during an interval. This breaks down the concept of OEE into separate aspects, each of which suggests more clearly focused lines of research for advancing our understanding of open-endedness. Dolson went on to discuss the relationships between these factors: To have novelty potential, a system requires change potential, and to have complexity potential or ecosystem potential, a system requires novelty potential. She acknowledged that other barriers might also exist, and, in particular, she and her colleagues are currently considering including a barrier of *major transition potential* in their picture.⁹

In “A New Design Principle for Open-Ended Evolution,” **Takashi Ikegami** discussed evolution in the context of web-based systems. Specifically, he reported work with colleagues on studying the dynamics of a social network site¹⁰ where users can upload photos and other users can attach

⁸ That is, an evolutionary system in which the individuals were themselves evolutionary systems, with selection based upon the individual systems’ capacities for OEE.

⁹ Dolson and her collaborators elaborated on these issues, including major transition potential, in a web article published after the workshop, <https://thewinnower.com/papers/2309-what-s-holding-artificial-life-back-from-open-ended-evolution>.

¹⁰ <http://roomclip.jp/>

tags to the photos to describe their content. By using a variety of mathematical techniques to analyze the use and evolution of tags in the system over a period of three years, Ikegami argued that the increasing vocabulary of tags observed over time generated a self-maintaining system in which certain types of tags stimulate users to create new combinations, and which prompted users to upload new photos to be annotated by those tags. Furthermore, a phase-transition-like event was observed in the system's activity, involving a sudden increase both in typical social network size and in tagging activity of each user. Based upon these observations, Ikegami and colleagues hypothesize that OEE in such systems can be driven by the users' collective activities. The work discussed in this talk is described in more detail in a late-breaking paper presented at the main ECAL 2015 conference [33].

Tom Froese began his talk, "Groundlessness Avoids Openness Reduction in Hierarchies of Emergence," with the observation that problems of OEE are of interest to origin-of-life researchers (even if not being addressed by biologists more broadly, as claimed by Beslon). In particular, he highlighted recent work by Peter Strazewski on evolution in chemical systems: Strazewski argues that OEE is more likely if we move away from well-defined systems to *messy* systems with many possible variants (in chemical composition, property, reactivity, shape, size, etc.) [46]. Froese attempted to formalize this intuition by characterizing OEE in terms of a system's emergence of new degrees of freedom (DoFs). He argued that if emergence is defined as collective dynamics resulting from nonlinear coupling between two or more components, then the number of DoFs of the emergent phenomena cannot, in principle, be greater than the sum of the numbers of DoFs of its components. Froese argued that this suggests that as we climb to higher hierarchical levels of complexity in a system, we inevitably witness a decrease in the number of DoFs of the system at those levels. This might not be a problem in systems that have sufficient complexity (many DoFs) at the bottom level. However, an alternative approach to avoiding this limitation would be to assume there is no bottom level—that the system is *groundless* (a line of thought inspired by Michel Bitbol [8]). Froese suggested that an explanation of OEE in the real world might therefore require us to conceptualize reality as a groundless system.

3 Themes from the Open Discussions

Substantial time in the workshop was left for open discussion, and a couple of important conclusions emerged. One was pluralism about OEE. It seems clear that there is more than one interesting and important kind of OEE. This means that those discussing OEE should whenever possible be explicit and precise about the kind of OEE of interest. Every kind of OEE should be identified and defined as precisely as possible, taking care not to lose those kinds that have intuitive appeal but cannot be precisely defined. Each successful definition should be operational and quantitative. But no definition is the one and only right definition of OEE if there is more than one kind of OEE. Some people might be especially interested in one kind of OEE, and others in another kind. Later in this section we attempt to identify the broad categories of kinds of OEE discussed at the workshop.

"Open-ended evolution" refers to a distinctive kind of behavior exhibited by some evolving systems, and different kinds of OEE correspond to somewhat different kinds of behavior. The workshop discussion highlighted the importance of distinguishing observable *behavioral hallmarks* of systems undergoing OEE from hypothesized *underlying mechanisms* that explain why a system exhibits those hallmarks. These "mechanisms" might be merely causally necessary conditions for OEE, or necessary boundary conditions. A sufficiently large population, or a sufficiently long duration, or a sufficiently large evolutionary search space have all been proposed as necessary conditions for the appearance of OEE. Perhaps no single mechanism is causally sufficient to produce OEE, but presumably each kind of OEE is produced by something like a set of individually necessary and jointly sufficient mechanisms. Identifying these key mechanisms for (each kind of) OEE is the question driving much of the research on OEE.

Both hallmarks of OEE and mechanisms for OEE are important, but they are important for different reasons. The hallmarks identify the important distinctive observable signs of (different kinds of) OEE. A given kind of OEE might have more than one behavioral hallmark, and different kinds will have somewhat different hallmarks, so the list of hallmarks of OEE can be expected to be somewhat heterogeneous.

The York workshop constructed an initial list of behavioral hallmarks of OEE. The focus was to generate a comprehensive list, so possible dependences among the hallmarks were secondary. The set of behavioral hallmarks of systems undergoing OEE emphasized in the workshop include these:

1. *Ongoing adaptive novelty* is one important kind of OEE. Novel adaptations come in many kinds, including new properties of entities, new interactions among entities, and new global patterns of behavior. In each case the focus is usually on the evolution of novel *adaptations*. Here are examples of some specific kinds of adaptive novelties:
 - (a) *Ongoing generation of new adaptations* is a very simple kind of OEE, and detecting it was the motivation for the original evolutionary activity statistics of Bedau and Packard [5, 4]. New adaptations could arise through a combination of different evolutionary and ecological mechanisms, such as competitive exclusion, random drift among neutral variants, and kin selection. Adaptation comes in different forms; for example, sometimes it is possible and important to distinguish new instances of a familiar kind of adaptation from qualitatively new kinds of adaptations, and the ongoing generation of qualitatively new kinds of adaptations is more interesting and more challenging to understand. The ongoing generation of new adaptations might seem to involve populations of agents with an unlimited number of different basic adaptive traits. However, practical considerations often impose a finite ceiling on the number of different basic adaptations distinguished in computer models or natural systems. Nevertheless, if evolution can produce finite combinations (sets) of adaptive traits, then the number of potential new adaptive combinations increases dramatically.
 - (b) *Ongoing generation of new kinds of entities* is one way to bring about the ongoing generation of new kinds of adaptations. The emergence of dynamical hierarchies described by Rasmussen et al. [38] is one mechanism for generating new kinds of entities with new kinds of properties. However, since the underlying mechanism in dynamical hierarchies can be merely chemical and physical self-assembly and self-organization, the properties of the entities at different levels in the hierarchy might not be adaptations. But if a dynamical hierarchy incorporates new material and information from the environment, and if the whole hierarchy can reproduce similar daughter hierarchies, then adaptive evolution could arise and start to shape a population of new kinds of entities.
 - (c) A major transition in evolution involves the emergence of a dynamical hierarchy and *does* involve adaptive evolution, and *ongoing major transitions in evolution* constitute another kind of OEE. The major transitions in evolution discussed by Maynard Smith and Szathmary [31] (and recently revisited by Szathmary [47]) are an especially interesting form of dynamical hierarchies, and they are special because each new level in the hierarchy consists of a new population of reproducing and evolving entities. A major transition in evolution is preceded by the evolution of one or several distinct kinds of reproducing entity. Eventually certain groups of those entities come to interact very tightly, and they become members of a new population of higher-level reproducing *wholes*. Entities in the old lower-level population become *parts* of the new wholes, but they cannot reproduce independently. Now the process repeats once more. Certain groups in the population of new wholes come to interact very tightly, and they become new even-higher-level wholes that reproduce and form an even-higher-level

population, and so on. Maynard Smith and Szathmary [31] conclude that the major transitions in evolution they survey are quite contingent; they could easily not have happened, and there may be no more major transitions.¹¹ So, the existence of some major transitions in evolution is not necessarily any kind of OEE. But major transitions can spur many further adaptations and help make evolution open-ended. And ongoing major transitions would be an especially impressive kind of OEE.¹²

- (d) Since major transitions in evolution typically create new kinds of entities with new kinds of adaptations, the transitions are one way in which the ability to evolve can itself evolve. But there are many other ways in which the ability to evolve can itself evolve. Thus, the *ongoing evolution of evolvability* is another kind of OEE. One especially critical step in the evolution of evolvability is the very first step: the *emergence* of the ability to evolve at all (the subject of Packard’s talk at the workshop).
2. *Ongoing growth of complexity* of an evolving system is another kind of OEE, and there are at least two different kinds of complexity to distinguish.
- (a) One focus is the complexity of the entities in an evolving population, and one kind of OEE is the ongoing growth in *complexity of entities* in the evolving population. The property of interest here is the complexity of the *most* complex entities, rather than entities with mean or modal complexity [18]. Further kinds of OEE involve ongoing growth of other global properties of the evolving population, such as diversity or disparity [17]. Note that growth of entity complexity is a side effect of major transitions in evolution, when the old evolving entities become parts of the new ones. But other mechanisms could also produce entities that are more and more complex.
 - (b) Another way in which an evolving system can become more complex is for the *interactions* among the entities in it to become more complex. Ongoing growth of complexity of interactions among entities is another kind of OEE. Even if the internal properties of the entities in a system remain the same, the interactions among entities can become more and more complex, as when food webs among species become more complex.

The emphasis on “ongoing” novelty itself deserves a brief mention. “Ongoing” is better than another common expression used in this context—“perpetual” novelty—because OEE is not actually perpetual although it is ongoing. The discussion in York focused partly on David Ackley’s idea of *indefinite scalability*, after this concept was emphasized in his talk. Ackley [2] defines indefinite scalability “as supporting open-ended computational growth without requiring substantial re-engineering.” The key criteria for indefinite scalability is that, should an upper bound be reached (e.g., in the number of novel entities encountered over the course of evolution or in the diversity or complexity of entities), increasing the values of physical limits (e.g., available matter, population size, or memory) should enable an unbounded sequence of greater upper bounds to be achieved (after sufficient increases in the limits). However, it is not possible in finite system time to establish that a metric is truly unbounded.¹³ And it is not possible—over a finite number of increases in system parameter(s)—to establish that a metric is truly indefinitely scalable. Further, an increase in parameter(s) may require a longer system time before a greater scale (higher value metric) is

¹¹ For alternative perspectives on major transitions, see also [12].

¹² Although, as discussed in several of the talks, each successive major transition produces organisms requiring more resources and existing in smaller populations. Hence, there is an inevitable limit in any finite system on the extent to which the occurrence of successive major transitions can be ongoing.

¹³ Some workshop participants (e.g., McGregor) felt that a sensible null hypothesis might be to assume that ALife systems were unbounded by default, that is, they might just need more time and larger environments to display OEE (an additional complicating factor here is the role of *contingency* in determining the outcome of any one specific run [52]). The research program then becomes a matter of identifying reasons why this might *not* be the case—a view that resonates with Dolson’s talk. On the other hand, others (including various other speakers discussed in Section 2) felt that these systems were missing important enabling conditions for OEE.

achieved. Claims about systems can, though, be expressed and evaluated in terms such as a metric increasing without bound up to a certain system time (or number of generations, etc.); or a metric increasing as system parameter(s) are increased up to certain value(s), where it was necessary to increase these to establish increases in the metric's maximum observed value over successive runs. Furthermore, we can define boundedness of a metric within a system in a rigorous way by fitting mathematical functions to the data and using statistics to ascertain which function is the best fit (e.g., see [57]). If the best-fit function is unbounded, that is a good indication that the system is exhibiting unbounded behavior.

Clarifying the hallmarks of OEE is a crucial step in clearly identifying and distinguishing the different kinds of OEE. After the hallmarks are clear, another crucial step is identifying and testing possible mechanisms that would produce and explain each of the kinds of OEE. Different mechanisms might be proposed to produce or explain a given kind of OEE; the mechanisms could provide competing explanations, or they could provide cooperating mechanisms. Also, a single mechanism might be involved in the explanation of more than one kind of OEE. So, the list of hypothetical mechanisms for a given kind of OEE could be rather heterogeneous. In addition, OEE pluralism means that different kinds of OEE could have different underlying mechanisms. Some mechanisms might be necessary for one kind but not another kind of OEE; other mechanisms might be necessary for every kind.

The discussion in York was weighted towards hallmarks of OEE, but some mechanisms for OEE were also mentioned and discussed. For example, one might think that the evolution of the genetic code is the mechanism behind OEE. Certain mechanisms are very obvious, but often insufficient by themselves. For example, since OEE involves adaptive evolution, natural selection helps explain it, and we already know a lot about evolution by natural selection. The participants in the discussion were divided about whether we already know enough to explain each kind of OEE, with some conjecturing that a fundamentally new mechanism is required for some kinds of OEE, such as major transitions in evolution.

Note that major transitions, the evolution of the genetic code, and the evolution of evolvability in general, are both *kinds of* OEE and *mechanisms for* kinds of OEE. This shows how one and the same thing can appear on the lists of both hallmarks of (one kind of) OEE and mechanisms for (another kind of) OEE.

An important research goal is to document examples of each hallmark and requirement of OEE, both in computer models and in natural systems. Positive examples that demonstrate a kind of OEE in a model or natural system are especially critical, but also important are negative examples of model or natural systems that do *not* demonstrate some kind of OEE.

Open-ended evolution is an ongoing process, so a *single* instance of the behavioral hallmarks of OEE falls short of being genuinely open-ended. A single new adaptation is not OEE, neither is the growth in complexity of one organism, nor one instance of the evolution of evolvability, nor one major transition in evolution. Nevertheless, it can be a significant scientific achievement to document even single instances of some especially challenging hallmarks, such as major transitions in evolution.

The following list summarizes the specific systems described by the speakers and participants in York (and some closely related systems) and the claims made of them regarding OEE:¹⁴

- *Earth's biosphere* has been classified, through fossil data sets at the level of taxonomic families, as exhibiting open-ended evolutionary dynamics according to Bedau and Packard's evolutionary activity measures [6, 7]. Bedau et al. reasoned that it was not necessary to include a shadow mechanism in this analysis, as "the mere fact that a family appears in the fossil record is good evidence that its persistence reflects its adaptive significance" [7].

¹⁴ While this report focuses on research discussed during the York workshop, a task for future work is to compile a more comprehensive list of achievements in OEE to date. Such a list is planned to form part of a comprehensive review paper on OEE to be produced after the OEE2 workshop.

- The *Long-Term Evolution Experiment (LTEE)* (Lenski et al.). The LTEE [28] is the most extensive laboratory study of ongoing biological evolution. Publications from this highly tractable system exhibit dynamics in evolving populations of *E. coli* that appear to be open-ended. Specifically, the LTEE has shown continuous increases in fitness that are best described by an unbounded power law function [57, 29]. Individual populations have shown continuous generation of novelty, such as: new portions of the fitness landscape continually being explored [53], numerous selective sweeps [30], new diversity arising after sweeps [9], and epistatic interactions among mutations where later benefits depend upon earlier mutations [56]. Finally, multiple populations in the LTEE exhibit frequency dependence [42, 24, 41, 30], including a special case [10, 9, 54] shown to be driven by ecological specialization and crossfeeding [55]. Most prominently, the LTEE gained substantial attention when a drastically new phenotype appeared, giving rise to what amounted to a new species [10, 9].
- *Tierra* (Ray). Tierra [39] is perhaps the most well-known example of an early ALife system in which digital organisms—self-replicating computer programs—were free to evolve in an open-ended manner without the guidance of an explicit fitness function. However, each particular run of the system would eventually reach a state of stasis where only selectively neutral variations were seen to emerge [40, 49]. Bedau and colleagues analyzed the dynamics of a Tierra-like system named Evita (but not Tierra itself), and found it to have qualitatively different evolutionary dynamics from those displayed in biological evolution as evidenced by the fossil record [6].¹⁵
- *Avida* (Ofria et al.). Avida [32] is currently the most widely used digital evolution system, and is used to study a wide range of evolutionary and ecological dynamics in populations of self-replicating computer programs. Avida has enabled the evolution of qualitatively novel behaviors such as complex features completely absent in the ancestor organism (ongoing generation of new adaptations) [27], novel collaboration strategies among organisms (ongoing growth of complexity of interactions) [16], and novel ecological interactions through coevolution promoting even greater levels of complexity [58]. Dolson and collaborators are actively testing their complexity barriers in this system, as well as analyzing evolutionary activity statistics. Initial results of the boundedness of fitness growth in simple to complex environments in Avida indicate that fitness continues to increase without an asymptote in the default environment. Many ongoing projects use Avida to evolve cooperation, ecosystems, sexual reproduction, parasitism/mutualism, pleiotropy, intelligence, evolvability, and complexity.
- *Geb* (Channon). Geb was the first ALife system to be classified as exhibiting opened evolutionary dynamics according to Bedau and Packard's evolutionary activity measures [7] and is the only one to have been classified as such according to an enhanced version of these measures developed by Channon [14, 15].¹⁶ Novel adaptations reported in Geb include behaviors such as following, fighting, fleeing, mimicking, and novel artifacts such as matching I&O channels in agents' neurocontrollers. Preliminary (unpublished) results presented at the Artificial Life XI conference in 2008 further indicated that component diversity (a simple measure of system complexity) may be indefinitely scalable (although that term was not yet in use); a more complete study of this is now planned.
- *Picbler's computational ecosystem* [35, 36, 34] is the only other ALife system to date to have been classified as exhibiting open-ended evolutionary dynamics according to Bedau and Packard's

¹⁵ Specifically, they found that the Evita model was bounded in their *component diversity* measure, whereas the fossil data was unbounded. In a separate analysis, Taylor also found that his Cosmos system—an elaboration of the basic Tierra design—also showed evolutionary dynamics similar to those of Evita [48, pp. 122–127].

¹⁶ The enhanced versions of the measures were found by Stout and Spector [45] to be of particular importance to the test's robustness against attempts to “break” the test by achieving unbounded dynamics in “intuitively unlikely” systems.

evolutionary activity measures. To the best of our knowledge, it has not been subjected to the enhanced test.

- *Stringmol* (Hickinbotham et al.). Stringmol is an artificial chemistry system that has been shown to exhibit the ongoing appearance of new chemical species [21]. In some cases the system has been shown to evolve multi-species hypercycles that persist for prolonged periods. Thus, Stringmol demonstrates the ongoing generation of new adaptations. These adaptations affect a species' binding affinity to other species, as well as its reaction rules. Quantitative novelties are certainly arising in the system (e.g., in binding affinities), although it is yet to be established whether any qualitative novelties are arising at the level of the individual chemical species. The appearance of hypercycles also demonstrates growth of complexity of interactions, and a qualitatively new organization of the system.
- *Novelty search* (Lehman and Stanley). Mentioned in Dolson's talk, Lehman and Stanley's *novelty search* technique has attracted considerable interest in recent years [25, 26]. The approach has been shown to generate ongoing generation of new adaptations. However, this is achieved by employing a selection mechanism that specifically looks for novel phenotypes. Hence, by design the approach will produce the *hallmark* of ongoing generation of new adaptations if (and only if) the system has implemented the necessary *mechanisms* for the ongoing generation of such adaptations. Novelty search, by itself, does not take a stand on what kinds of mechanisms are required. Furthermore, it requires a measure of phenotypic distance between individuals to tell whether two individuals exhibit sufficiently different behaviors. Like a fitness function in traditional EAs, this definition of phenotypic uniqueness needs to be carefully chosen. Defining a more general measure, applicable to OEE, appears to be a major research challenge, but potentially a rewarding one. Further work is required to understand the similarities and differences between novelty search and OEE; one line of research along these lines has recently been initiated by Soros and Stanley [44].
- *Dynamical hierarchies* (Rasmussen et al.). Rasmussen and colleagues reported results in a model of a physicochemical system that exhibited dynamical hierarchies. They demonstrated the emergence of two higher orders of entities and interactions on top of the basic first-order elements built into the system. This work was based upon a model of self-assembly rather than evolution; to be of direct relevance to OEE it would need to be augmented with mechanisms for self-replication, variation, and selection of the emergent dynamical hierarchies. Enabling populations of newly emerging dynamical hierarchies to undergo adaptive evolution would unify the processes of self-assembly and self-organization with the process of adaptive evolution, and this could explain one kind of novelty in OEE: the evolution of new kinds of wholes with new kinds of properties. In this context one mechanism driving the ongoing generation of novelty is the availability in the environment of new materials that can aggregate and generate novel properties.
- *Emergence of coding* (Packard and Guttenberg). In his talk, Packard described preliminary work on a model in which alternation between unstable and fixed-point dynamics produced conditions suitable for the emergence of informationally-stable components. While preliminary, the results have relevance for the pre-biotic transition to information-driven systems. Packard reported that colleagues are applying these ideas to models of other major evolutionary transitions too. This work has not yet been published.
- *Patented technology* (Bedau et al.). Bedau suggested that the actual evolution of technology (detected in the patent record) is a real-world system that exhibits a form of OEE that he termed ongoing *door-opening* evolution [11], which occurs when one technological innovation enables a whole new kind of technology to arise and diversify. Bedau conjectured that door-opening innovations are an important mechanism behind the ongoing generation of new kinds of adaptations, and he proposed some first steps to observing and measuring door-opening innovations in the patent record.

- *Social media tags* (Ikegami). In his talk, Ikegami argued that the social media tag system he described represented an example of OEE. With respect to the evolution of new combinations of tags, this would be ongoing generation of new adaptations (at a quantitative level). The role of human users as an integral part of the system, who both supply the new tags and upload new images to be tagged, is a complicating factor in this case.

Of the example systems discussed by the speakers in York, many focused on the ongoing generation of quantitatively new adaptations, where “quantitatively new” means that the adaptations are novel, but identifiable within a determined class of possibilities, and as a result of their identifiability, they may be statistically quantified. In contrast, qualitatively new adaptations lie outside any predetermined class of possibilities. It is clear that qualitatively new adaptations are a part of natural evolutionary processes, but less clear whether and how they might occur in example systems considered so far—indeed, clearer criteria are required for what counts as qualitative, rather than quantitative, novelty in these systems. Sharpening this distinction should lead toward progress in understanding how open-ended evolution manifests properties such as growth of complexity of interactions, ongoing generation of new entities, ongoing generation of new functionalities, and major evolutionary transitions.

4 Conclusion

The workshop in York closed with a better appreciation of what remains to be learned about open-ended evolution, and a clearer picture of the most important open research questions about OEE. These include:

1. What is the best way to categorize, define, and operationally observe each kind of OEE, and how are the different kinds related?
2. What are the most important candidate mechanisms (or necessary conditions) for producing each kind of OEE? Which mechanisms are most plausible?
3. Which kinds of OEE can be demonstrated in specific systems, including analytical models, computer models, laboratory experimental systems, or natural biological communities? What has already been shown in each type of system?

The state of the art on these and other fundamental questions about OEE are the focus of the *Second Workshop on Open-Ended Evolution (OEE2)*,¹⁷ to be held at the ALIFE XV conference in Cancún, Mexico in July 2016.

Acknowledgments

We gratefully acknowledge use of Overleaf, the free, online collaborative LaTeX authoring tool (<https://www.overleaf.com/>), in preparing the draft of this paper.

References

1. Ackley, D. H., & Ackley, E. S. (2015). Artificial life programming in the robust-first attractor. In P. Andrews, L. Caves, R. Doursat, S. H. F. Polack, S. Stepney, T. Taylor, & J. Timmis (Eds.), *Proceedings of the European Conference on Artificial Life 2015* (pp. 554–561). Cambridge, MA: MIT Press.
2. Ackley, D. H., & Small, T. R. (2014). Indefinitely scalable computing = artificial life engineering. In H. Sayama, J. Rieffel, S. Risi, R. Doursat, & H. Lipson (Eds.), *Artificial life 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems* (pp. 606–613). Cambridge, MA: MIT Press.
3. Banzhaf, W., Baumgaertner, B., Beslon, G., Doursat, R., Foster, J. A., McMullin, B., de Melo, V. V., Miconi, T., Spector, L., Stepney, S., & White, R. (2016). Defining and simulating open-ended novelty: Requirements, guidelines, and challenges. *Theory in Biosciences* (in press).

¹⁷ <http://alife.org/ws/oe2>

4. Bedau, M. A., & Brown, C. T. (1999). Visualizing evolutionary activity of genotypes. *Artificial Life*, 5(1), 17–35.
5. Bedau, M. A., & Packard, N. H. (1991). Measurement of evolutionary activity, teleology, and life. In C. Langton, C. Taylor, D. Farmer, & S. Rasmussen (Eds.), *Artificial life II*, vol. X of *Santa Fe Institute Studies in the Sciences of Complexity* (pp. 431–461). Boston, MA: Addison-Wesley.
6. Bedau, M. A., Snyder, E., Brown, C. T., & Packard, N. H. (1997). A comparison of evolutionary activity in artificial evolving systems and in the biosphere. In P. Husbands & I. Harvey (Eds.), *Proceedings of the Fourth European Conference on Artificial Life* (pp. 125–134). Cambridge, MA: MIT Press.
7. Bedau, M. A., Snyder, E., & Packard, N. H. (1998). A classification of long-term evolutionary dynamics. In C. Adami, R. K. Belew, H. Kitano, & C. E. Taylor (Eds.), *Artificial life VI: Proceedings of the Sixth International Conference on Artificial Life* (pp. 228–237). Cambridge, MA: MIT Press.
8. Bitbol, M. (2007). Ontology, matter and emergence. *Phenomenology and the Cognitive Sciences*, 6(3), 293–307.
9. Blount, Z. D., Barrick, J. E., Davidson, C. J., & Lenski, R. E. (2012). Genomic analysis of a key innovation in an experimental *Escherichia coli* population. *Nature*, 489(7417), 513–518. <http://dx.doi.org/10.1038/nature11514>.
10. Blount, Z. D., Borland, C. Z., & Lenski, R. E. (2008). Historical contingency and the evolution of a key innovation in an experimental population of *Escherichia coli*. *Proceedings of the National Academy of Sciences of the USA*, 105(23), 7899–7906. <http://www.pnas.org/content/105/23/7899.abstract>.
11. Buchanan, A., Packard, N. H., & Bedau, M. A. (2011). Measuring the evolution of the drivers of technological innovation in the patent record. *Artificial Life*, 17(2), 109–122.
12. Calcott, B., & Sterelny, K. (Eds.) (2011). *The major transitions in evolution revisited*. Cambridge, MA: MIT Press.
13. Chalmers, D., Francis, C. C., Pepper, N., & Bedau, M. A. (2010). High-content words in patent records reflect key innovations in the evolution of technology. In H. Fellermann, M. Dörr, M. M. Hanczy, L. L. Laursen, S. Maurer, D. Merkle, P.-A. Monnard, K. Stoy, & S. Rasmussen (Eds.), *Artificial life XII: Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems* (pp. 838–845). Cambridge, MA: MIT Press.
14. Channon, A. (2003). Improving and still passing the ALife test: Component-normalised activity statistics classify evolution in Geb as unbounded. In R. K. Standish, M. A. Bedau, & H. A. Abbass (Eds.), *Artificial life VIII: Proceedings of the Eighth International Conference on Artificial Life* (pp. 173–181). Cambridge, MA: MIT Press.
15. Channon, A. (2006). Unbounded evolutionary dynamics in a system of agents that actively process and transform their environment. *Genetic Programming and Evolvable Machines*, 7(3), 253–281. <http://dx.doi.org/10.1007/s10710-006-9009-3>.
16. Goldsby, H. J., Dornhaus, A., Kerr, B., & Ofria, C. (2012). Task-switching costs promote the evolution of division of labor and shifts in individuality. *Proceedings of the National Academy of Sciences of the USA*, 109(34), 13686–13691.
17. Gould, S. J. (1989). *Wonderful life: The Burgess Shale and the nature of history*. New York: W. W. Norton.
18. Gould, S. J. (1996). *Full house: The spread of excellence from Plato to Darwin*. New York: Harmony Books. (Published in the UK under the title *Life's Grandeur*.)
19. Harvey, I. (2001). Artificial evolution: A continuing saga. In T. Gomi (Ed.), *Evolutionary robotics. From intelligent robotics to artificial life*, vol. 2217 of *Lecture Notes in Computer Science* (pp. 94–109). Berlin: Springer. http://dx.doi.org/10.1007/3-540-45502-7_5.
20. Hickinbotham, S., Clark, E., Nellis, A., Stepney, S., Clarke, T., & Young, P. (2016). Maximizing the adjacent possible in automata chemistries. *Artificial Life*, 22(1), 49–75.
21. Hickinbotham, S., Clark, E., Stepney, S., Clarke, T., Nellis, A., Pay, M., & Young, P. (2010). Diversity from a monoculture: Effects of mutation-on-copy in a string-based artificial chemistry. In H. Fellermann, M. Dörr, M. M. Hanczy, L. L. Laursen, S. Maurer, D. Merkle, P.-A. Monnard, K. Stoy, & S. Rasmussen (Eds.), *Artificial life XII: Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems* (pp. 24–31). Cambridge, MA: MIT Press.
22. Hickinbotham, S., & Stepney, S. (2015). Conservation of matter increases evolutionary activity. In P. Andrews, L. Caves, R. Doursat, S. H. F. Polack, S. Stepney, T. Taylor, & J. Timmis (Eds.), *Proceedings of the European Conference on Artificial Life 2015* (pp. 98–105). Cambridge, MA: MIT Press.

23. Hickinbotham, S., & Stepney, S. (2015). Environmental bias forces parasitism in Tierra. In P. Andrews, L. Caves, R. Doursat, S. H. F. Polack, S. Stepney, T. Taylor, & J. Timmis (Eds.), *Proceedings of the European Conference on Artificial Life 2015* (pp. 294–301). Cambridge, MA: MIT Press.
24. Le Gac, M., Plucain, J., Hindré, T., Lenski, R. E., & Schneider, D. (2012). Ecological and evolutionary dynamics of coexisting lineages during a long-term experiment with *Escherichia coli*. *Proceedings of the National Academy of Sciences of the USA*, 109(24), 9487–9492. <http://www.pnas.org/content/109/24/9487.abstract>.
25. Lehman, J., & Stanley, K. O. (2008). Exploiting open-endedness to solve problems through the search for novelty. In S. Bullock, J. Noble, R. A. Watson, & M. A. Bedau (Eds.), *Artificial life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*. Cambridge, MA: MIT Press.
26. Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2), 189–223. http://www.mitpressjournals.org/doi/pdf/10.1162/EVCO_a_00025.
27. Lenski, R. E., Ofria, C., Pennock, R. T., & Adami, C. (2003). The evolutionary origin of complex features. *Nature*, 423(6936), 139–144. <http://www.nature.com/nature/journal/v423/n6936/abs/nature01568.html>.
28. Lenski, R. E., Rose, M. R., Simpson, S. C., & Tadler, S. C. (1991). Long-term experimental evolution in *Escherichia coli*. I. Adaptation and divergence during 2,000 generations. *The American Naturalist*, 138(6), 1315–1341. <http://www.jstor.org/stable/2462549>.
29. Lenski, R. E., Wiser, M. J., Ribeck, N., Blount, Z. D., Nahum, J. R., Morris, J. J., Zaman, L., Turner, C. B., Wade, B. D., Maddamsetti, R., Burmeister, A. R., Baird, E. J., Bundy, J., Grant, N. A., Card, K. J., Rowles, M., Weatherspoon, K., Papoulis, S. E., Sullivan, R., Clark, C., Mulka, J. S., & Hajela, N. (2015). Sustained fitness gains and variability in fitness trajectories in the Long-Term Evolution Experiment with *Escherichia coli*. *Proceedings of the Royal Society of London B: Biological Sciences*, 282(1821). <http://rspb.royalsocietypublishing.org/content/282/1821/20152292.abstract>.
30. Maddamsetti, R., Lenski, R. E., & Barrick, J. E. (2015). Adaptation, clonal interference, and frequency-dependent interactions in a long-term evolution experiment with *Escherichia coli*. *Genetics*, 200(2), 619–631. <http://www.genetics.org/content/200/2/619.abstract>.
31. Maynard Smith, J., & Szathmáry, E. (1995). *The major transitions in evolution*. Oxford, UK: Oxford University Press.
32. Ofria, C., & Wilke, C. O. (2004). Avida: A software platform for research in computational evolutionary biology. *Artificial Life*, 10(2), 191–229. <http://dx.doi.org/10.1162/106454604773563612>.
33. Oka, M., Hashimoto, Y., & Ikegami, T. (2015). An open-ended evolution in a web system. In *Late breaking papers at the European Conference on Artificial Life*. <https://www.cs.york.ac.uk/nature/ecal2015/paper-159.html>.
34. Pichler, P.-P. (2009). *Natural selection, adaptive evolution and diversity in computational ecosystems*. Ph.D. thesis, University of Hertfordshire.
35. Pichler, P.-P., & Canamero, L. (2007). An evolving ecosystems approach to generating complex agent behaviour. In *Proceedings of the First IEEE Symposium on Artificial Life* (pp. 303–310). New York: IEEE.
36. Pichler, P.-P., & Canamero, L. (2008). Evolving morphological and behavioral diversity without predefined behavior primitives. In S. Bullock, J. Noble, R. A. Watson, & M. A. Bedau (Eds.), *Artificial life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems* (pp. 404–411). Cambridge, MA: MIT Press.
37. Popper, K. (1972). *Objective knowledge: An evolutionary approach*. Oxford, UK: Oxford University Press.
38. Rasmussen, S., Baas, N. A., Mayer, B., Nilsson, M., & Olesen, M. W. (2001). Ansatz for dynamical hierarchies. *Artificial Life*, 7(4), 329–353.
39. Ray, T. S. (1991). An approach to the synthesis of life. In C. G. Langton, C. Taylor, D. Farmer, & S. Rasmussen (Eds.), *Artificial life II* (pp. 371–408). Boston, MA: Addison-Wesley.
40. Ray, T. S. (1992). *Evolution, ecology and optimization of digital organisms*. Technical Report 92-08-942. Santa Fe Institute.
41. Ribeck, N., & Lenski, R. E. (2015). Modeling and quantifying frequency-dependent fitness in microbial populations with cross-feeding interactions. *Evolution*, 69(5), 1313–1320. <http://dx.doi.org/10.1111/evo.12645>.

42. Rozen, D. E., & Lenski, R. E. (2000). Long-term experimental evolution in *escherichia coli*. VIII. Dynamics of a balanced polymorphism. *The American Naturalist*, *155*(1), 24–35. <http://www.jstor.org/stable/10.1086/303299>.
43. Skusa, A., & Bedau, M. A. (2003). Towards a comparison of evolutionary creativity in biological and cultural evolution. In R. K. Standish, M. A. Bedau, & H. A. Abbass (Eds.), *Artificial life VIII: Proceedings of the Eighth International Conference on Artificial Life* (pp. 233–242). Cambridge, MA: MIT Press.
44. Soros, L. B., & Stanley, K. O. (2014). Identifying necessary conditions for open-ended evolution through the artificial life world of chromaria. In H. Sayama, J. Rieffel, S. Risi, R. Doursat, & H. Lipson (Eds.), *Artificial life 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems* (pp. 793–800). Cambridge, MA: MIT Press.
45. Stout, A., & Spector, L. (2005). Validation of evolutionary activity metrics for long-term evolutionary dynamics. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2005)* (pp. 137–142). Berlin: Springer.
46. Strazewski, P. (2015). Omne vivum ex vivo ... omne? How to feed an inanimate evolvable chemical system so as to let it self-evolve into increased complexity and life-like behaviour. *Israel Journal of Chemistry*, *55*(8), 851–864.
47. Szathmáry, E. (2015). Toward major evolutionary transitions theory 2.0. *Proceedings of the National Academy of Sciences of the USA*, *112*(33), 10104–10111. <http://www.pnas.org/content/112/33/10104.abstract>.
48. Taylor, T. (1999). *From artificial evolution to artificial life*. Ph.D. thesis, Division of Informatics, University of Edinburgh. Available online at <http://www.tim-taylor.com/papers/thesis/>.
49. Taylor, T. (2013). Evolution in virtual worlds. In M. Grimshaw (Ed.), *The Oxford handbook of virtuality* (chap. 32). Oxford, UK: Oxford University Press.
50. Taylor, T. (2015). Requirements for open-ended evolution in natural and artificial systems. In *EvoEvo Workshop at the 13th European Conference on Artificial Life (ECAL 2015)*, University of York, UK. <http://arxiv.org/abs/1507.07403>.
51. Taylor, T., Dorin, A., & Korb, K. (2014). Digital genesis: Computers, evolution and artificial life. In *7th Munich–Sydney–Tilburg Philosophy of Science Conference: Evolutionary thinking*. <http://arxiv.org/abs/1512.02100>.
52. Taylor, T., & Hallam, J. (1998). Replaying the tape: An investigation into the role of contingency in evolution. In C. Adami, R. K. Belew, H. Kitano, & C. E. Taylor (Eds.), *Artificial life VI: Proceedings of the Sixth International Conference on Artificial Life* (pp. 256–265). Cambridge, MA: MIT Press.
53. Tenaillon, O., Barrick, J. E., Ribbeck, N., Deatherage, D. E., Blanchard, J. L., Dasgupta, A., Wu, G. C., Wielgoss, S., Cruveiller, S., Médigue, C., Schneider, D., & Lenski, R. E. (2016). Tempo and mode of genome evolution in a 50,000-generation experiment. *bioRxiv*. <http://biorxiv.org/content/early/2016/01/15/036806.abstract>.
54. Turner, C. B., Blount, Z. D., & Lenski, R. E. (2015). Replaying evolution to test the cause of extinction of one ecotype in an experimentally evolved population. *PLoS ONE*, *10*(11), e0142050.
55. Turner, C. B., Blount, Z. D., Mitchell, D. H., & Lenski, R. E. (2015). Evolution and coexistence in response to a key innovation in a long-term evolution experiment with *Escherichia coli*. *bioRxiv*. <http://www.biorxiv.org/content/early/2015/06/17/020958.abstract>.
56. Wielgoss, S., Barrick, J. E., Tenaillon, O., Wisner, M. J., Dittmar, W. J., Cruveiller, S., Chane-Woon-Ming, B., Médigue, C., Lenski, R. E., & Schneider, D. (2013). Mutation rate dynamics in a bacterial population reflect tension between adaptation and genetic load. *Proceedings of the National Academy of Sciences of the USA*, *110*(1), 222–227. <http://www.pnas.org/content/110/1/222.abstract>.
57. Wisner, M. J., Ribbeck, N., & Lenski, R. E. (2013). Long-term dynamics of adaptation in asexual populations. *Science*, *342*(6164), 1364–1367. <http://www.sciencemag.org/content/342/6164/1364.abstract>.
58. Zaman, L., Meyer, J. R., Devangam, S., Bryson, D. M., Lenski, R. E., & Ofria, C. (2014). Coevolution drives the emergence of complex traits and promotes evolvability. *PLoS Biol*, *12*(12), e1002023.