

Evolutionary Innovations and Where to Find Them: Routes to Open-Ended Evolution in Natural and Artificial Systems

Tim Taylor

Monash University

Faculty of Information Technology

tim@tim-taylor.com

Abstract This article presents a high-level conceptual framework to help orient the discussion and implementation of open-endedness in evolutionary systems. Drawing upon earlier work by Banzhaf et al. (2016), three different kinds of open-endedness are identified: *exploratory*, *expansive*, and *transformational*. These are characterized in terms of their relationship to the search space of phenotypic behaviors. A formalism is introduced to describe three key processes required for an evolutionary process: the generation of a phenotype from a genetic description, the evaluation of that phenotype, and the reproduction with variation of individuals according to their evaluation. The formalism makes explicit various influences in each of these processes that can easily be overlooked. The distinction is made between intrinsic and extrinsic implementations of these processes. A discussion then investigates how various interactions between these processes, and their modes of implementation, can lead to open-endedness. However, an important contribution of the article is the demonstration that these considerations relate to exploratory open-endedness only. Conditions for the implementation of the more interesting kinds of open-endedness—expansive and transformational—are also discussed, emphasizing factors such as multiple domains of behavior, transdomain bridges, and non-additive compositional systems. In contrast to a traditional Darwinian analysis, these factors relate not to the generic evolutionary properties of individuals and populations, but rather to the nature of the building blocks out of which individual organisms are constructed, and the laws and properties of the environment in which they exist. The article ends with suggestions of how the framework can be used to categorize and compare the open-ended evolutionary potential of different systems, how it might guide the design of systems with greater capacity for open-ended evolution, and how it might be further improved.

Keywords

Open-ended evolution, types of open-endedness, transdomain bridges, non-additive compositional systems, evolutionary novelty, evolutionary innovation

I Introduction

In this article I identify different routes by which open-endedness (OE) can be introduced into the design and implementation of an evolutionary system.

I begin by presenting a definition of three different kinds of open-endedness. My treatment of the topic expands upon the approach recently proposed by Banzhaf et al. [1]. In their article, the distinction is made between *scientific models*, which are “descriptive models of part of the existing world,” and *engineering models* (including software design models), which are “prescriptive or

normative models of a system to be built in the world” [1, p. 135]. One of the main aims of their article was to develop a descriptive scientific (meta-)model to illustrate their definitions of open-endedness. They express the hope that “such a definition of OE in terms of models and meta-models will help the design of normative engineering models for implementing ALife” [1, p. 136].

The aim of the current contribution is to make progress towards exactly that goal—the development of an engineering model to guide the design and implementation of artificial evolutionary systems that possess the capacity for various kinds of open-endedness.¹

Having clarified what I mean by open-endedness, I then introduce a formalism for describing the key processes that must be present in any evolutionary system. The formalism makes explicit some important dependences and interrelationships that are otherwise easy to overlook.

Equipped with the necessary preliminaries, I then utilize the formalism to identify the various routes by which open-endedness can be accommodated in the design of an evolutionary system. It is found that the formalism only helps directly in the investigation of one type of open-endedness. I therefore continue the discussion with an analysis of potential factors involved in the other kinds of open-endedness as well.

Throughout the article I demonstrate how the presented framework² helps orient the study of open-endedness within the context of existing literature. I close the discussion by outlining how the framework could be used as a tool for analyzing and improving the open-endedness of existing artificial evolutionary systems, and offering some suggestions for further developments of the approach.

I consider the main contributions of the article to be the analysis of how the many different topics from the theoretical biology literature fit into the overall picture of open-ended evolution (as summarized in Figure 4), the finding that these only relate to one type of OE (Section 4.2), the discussion of ways of achieving the other types of OE (Section 4.3), the suggestions for extending the analysis with a more sophisticated treatment of behavior (Section 5), and also the simple schematic representation of Banzhaf et al.’s classes of open-endedness, upon which the discussion is based (Figure 1).

2 State Spaces, Novelties, and Open-Endedness

The idea of a *possibility space* or *state space* to represent the range of all possible forms of an individual in an evolutionary system is a widely employed concept (e.g., [1, 5, 12]). Indeed, they are simplifications of the concept of *adaptive landscapes* first proposed by Sewall Wright (for genotypes) in 1932 and by G. G. Simpson (for phenotypes) in 1944 [14]. State spaces are simpler than adaptive landscapes in that they lack a representation of the *adaptive value (fitness)* of each point in the space. I use the simpler concept of state space in the following discussion because it is sufficient for the purpose of the discussion; I consider how fitness comes into the picture later in the article.

While it is easy to use state spaces and adaptive landscapes to describe particular, well-constrained systems comprising a small number of clearly defined variables, it is nontrivial to apply them to elaborate and potentially open-ended systems. In these cases it can be problematic to enumerate and quantify all relevant variables to be used as dimensions of the space.³ However, even if it can be difficult to quantitatively describe a specific complex evolutionary system, state spaces can still be useful *intuition pumps* [13]—this is my intention in using them here.

To present the following ideas in more concrete terms, I have chosen to illustrate state spaces defined according to the ideas of *models* and *meta-models* set out in Banzhaf et al.’s recent treatment of

¹ This article concentrates specifically on open-ended evolution. In places reference is made to the more general concept of *open-endedness* (which admits that open-ended dynamics may be observed in other types of systems beyond evolutionary ones [35]), but the discussion presented here assumes an evolutionary context.

² Note that throughout the article I use the term “framework” to describe the whole approach to understanding open-endedness outlined here, which includes the definition of three kinds of open-ended evolution (OEE) (Section 2), the formalism describing the three basic evolutionary processes (Sections 3 and 4.2), and the discussion of mechanisms for expansive and transformational OEE (Section 4.3). I use the term “formalism” to refer specifically to the subset of the framework relating to the basic evolutionary processes (Sections 3 and 4.2).

³ But note that methods for inferring latent variable models can be employed to generate more meaningful *latent spaces*.

open-endedness [1]. Central to their approach is the idea that the behavior of a system can be described by a scientific (descriptive) model. The model is expressed in terms of a set of concepts, and those concepts can themselves be described by a meta-model. The meta-model describes a set of concepts that can be used to build a variety of specific models that use the same concepts in different ways. Readers unfamiliar with Banzhaf et al.'s contribution may benefit from reading it in order to fully understand what is summarized in this section.

Banzhaf et al. identify three different kinds of novelty that may occur in a system, defined according to whether the novelty necessitates changes in the system's model or meta-model. Their approach closely resembles Boden's ideas of three different kinds of creativity that have been developed over several decades [4, 5]. The distinction between novelties that fall within the system's current model and those that necessitate a change in the model (or meta-model) can also be seen in Waddington's pioneering work on open-ended evolution from 1969:

the ... requirement, that the available genotypes must be capable of producing phenotypes which can exploit ... new environments, requires some special provision of a means of creating genetic variation.... It is important to emphasize that the new genetic variation must not only be novel, but must include variations which make possible the exploration of environments which the population previously did not utilize.... It is not sufficient to produce new mutations which merely insert new parameters into existing programmes; they must actually be able to rewrite the programmes. [43, pp. 116–118]

I therefore adopt the general idea of three different kinds of novelty here (see below), without necessarily committing to Banzhaf et al.'s specific approach.

As discussed in previous OEE Workshops, one of the most general and widely accepted hallmarks of open-ended evolution is the presence of *ongoing adaptive novelty* [38]. The three different kinds of novelty therefore give rise to three different kinds of open-endedness. The three classes of novelty and their corresponding classes of open-endedness are:⁴

1. *Exploratory novelty*: A novelty that can be described using the current model (e.g., by recombining existing components, or changing the values of existing parameters). Potential examples from biology include the production of a new combination of alleles on a genome, and a change in the number of vertebrae in a new vertebrate species. A potential example from an ALife evolved virtual creature system would be the appearance of a creature with limbs of a length that is different to what has been observed before.
Exploratory open-endedness: The ongoing production of adaptive exploratory novelties.
2. *Expansive novelty*: A novelty that necessitates a change in the model but still using concepts present in the current meta-model. Potential examples from biology include synthesis of a new chemical species that has not been used in previous metabolic reactions, and the introduction of a new species of an existing genus that can occupy a new ecological niche. A potential example in a virtual creature's ALife system might be the evolution of legged locomotion where only snake-like locomotion had existed previously; this might represent a new instance of an existing meta-model concept of "terrestrial forward locomotion."
Expansive open-endedness: The ongoing production of adaptive expansive novelties.
3. *Transformational novelty*: A novelty that introduces a new concept, necessitating a change in the meta-model. Potential examples from biology include a major transition in

⁴ Banzhaf et al. used the terms *variation*, *innovation*, and *emergence*, respectively, in place of the terms used here [1]. I have chosen to introduce new terminology because the existing terms (especially innovation and emergence) are already widely used in many different contexts and with many different meanings. Furthermore, the new terms nicely fit the concepts of open-endedness described below and illustrated in Figure 1. My terms fit closely with Boden's concepts of *exploratory*, *combinational*, and *transformational creativity* [5]. As an example of the potential for confusion when using Banzhaf et al.'s terms, de Vladar et al. have recently used the term *innovation* [12] to describe novelties that most closely match Banzhaf et al.'s *emergent novelties*.

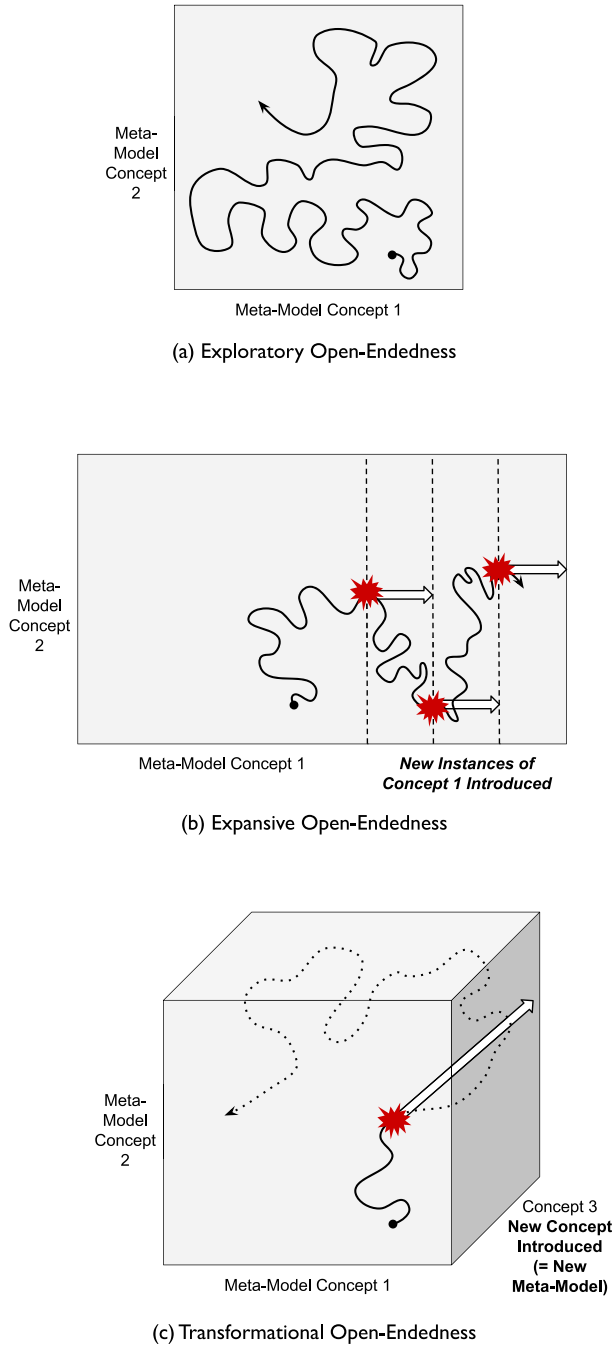


Figure 1. Types of open-endedness in a state space described by a model and its associated meta-model. See text for details.

individuality [25], the appearance of winged flight, and the appearance of visual sensory systems. A potential example in a virtual creature ALife system might be the evolution of flight where only terrestrial locomotion had existed previously.

Transformational open-endedness: The ongoing production of adaptive transformational novelties.

Note that I indicate “potential” examples in the list above, because according to Banzhaf et al.’s approach each type of novelty is defined relative to a given model and meta-model [1].

Banzhaf et al. define their three classes of novelty in terms of the system’s *current* model and meta-model. This means, for example, that once one major transition has been witnessed, the concept of major transition is then added to the meta-model, so any subsequent major transitions are not regarded as transformational. In contrast, I suggest that novelty be defined relative to the initial model and meta-model applied to an evolutionary system at its inception. In that case, after a transformational novelty appears for the first time, any further instances of the same kind of novelty will also be labeled transformational (and likewise for expansive novelties).⁵ A defining feature of expansive and transformational novelties, and hence the reason to label subsequent examples in the same class, is their ability to open up new adjacencies in an expanded state space [12, 24] (see further discussion in Section 4.3), and this occurs each time such a transition arises, not just the first time.

Furthermore, Banzhaf et al. chose not to classify the ongoing production of *exploratory* novelties as a type of open-endedness. In contrast, I have chosen to do so because, even though it takes place within a state space of fixed and finite size, that size might well be immense. Indeed, the number of possible combinations of entities and interactions described by a model might easily be so astronomical that an evolutionary process could not possibly visit all adaptive points in the space within the lifetime of the universe. This raises the distinction between *effective* OE and *theoretical* OE [1, p. 144],⁶ my interest in this article is in effective OE.

If we use a state space diagram to represent all possible entities and interactions describable by a system’s model and its associated meta-model, we can represent the three different kinds of open-endedness as shown in Figure 1.⁷

2.1 Genetic and Phenotypic State Spaces

In Figure 1, open-endedness is represented as an ongoing traversal of the space of possible organisms. In evolutionary systems, an organism’s phenotype and behavior are derived from a genetic description contained in its genome. The process of generating the phenotype from the genotype is defined by the organism’s genotype-phenotype (G-P) map. As discussed below, this map may be more or less complex, and more or less explicit in the system’s design.

We can split the representation of phenotypic state space (P-space) and genetic state space (G-space) into two separate diagrams. When considering open-endedness, we are ultimately interested in whether the system has the capacity for the ongoing production of adaptive phenotypes in P-space. However, the ability of an evolutionary system to explore P-space is fundamentally affected by the nature of the G-P map as the evolutionary processes of reproduction and variation of the genome explore different points in the genetic state space (G-space).

A simple example of G-space, along with its relationship to P-space, is shown in Figure 2. Note that the dimensionality of the G-space is not necessarily the same as that of the P-space: the relationship is determined by the G-P map, which can be of arbitrary form and can also depend upon the system’s global laws of dynamics and the local context in which the phenotype is generated.⁸

As shown in Figure 2, small moves in G-space might result in large moves in P-space, depending upon the nature of the G-P map. Furthermore, different mappings will lead to different paths in P-space for a given set of moves in G-space. Hence, the nature of the mapping, and whether or

5 Cf. Boden’s distinction between I-creativity and H-creativity [5].

6 On the distinction between these, Banzhaf et al. say “Should we be looking for systems able to continually produce open-ended events, or ‘simply’ for systems able to produce a sufficient number of open-ended events? We thus distinguish systems that are theoretically open-ended from those that are effectively open-ended. The former may be demonstrable in a mathematical universe, but questionable in a finite universe; the latter are questionable in a mathematical universe ... but may be demonstrable in a physical universe” [1, p. 144].

7 Any real system of interest will obviously have far more than the two conceptual axes shown in the figure, and it is not clear how different instances of a concept can be mapped onto a scalar scale in the general case. Hence, these diagrams are not meant to be taken too literally, but are nevertheless useful to communicate an intuitive idea of the different kinds of OE.

8 In terms of the formalism to be introduced later (Section 3.1), the G-P map is the function M_L in Equation 1, the global laws of dynamics are represented by the subscript L of that function, and the local context is represented by the function’s parameters c_a and c_b .

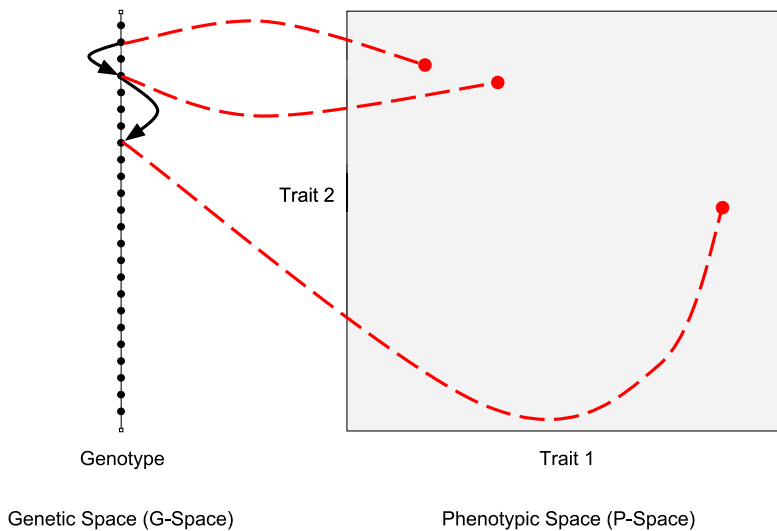


Figure 2. Genetic space and its relation to phenotypic space. The mapping from G-space to P-space is defined by the G-P map. The mapping might be such that small moves in G-space can sometimes result in large moves in P-space. Note that the dimensionality of G-space might be different to that of P-space.

not the mapping itself has the ability to evolve over time, will fundamentally affect the system's ability to explore different areas of P-space.

2.1.1 Relationship between G-Space and Open-Endedness in P-Space

In some cases, a system might exhibit effective transformational open-endedness in P-space even with a fixed G-space; in Section 4.3.2 we will see how this might happen.⁹ While effective transformational open-endedness is possible in a fixed G-space, one might think that a more obvious way to achieve it is to allow the number of genes on the genome to grow—leading to an expanding G-space. If the size of the genome can potentially expand without limit, we have what is referred to in the evolutionary biology literature as an *indefinite hereditary replicator* [25]. All else being equal, a larger genome can (but does not necessarily) specify a more complicated phenotype. While this can indeed be the case, the ability of a larger genome to specify expansive or transformational novelties depends upon the capacity of the additional genes to specify new traits. This can be achieved (as in the fixed G-space case) through the methods to be discussed in Section 4.3.2.

3 Evolutionary Processes

Having introduced some concepts and definitions relating to open-ended evolution in the previous section, I now discuss some high-level general features of evolutionary systems, and introduce a formalism to describe them. This will then provide a framework that can be used to explore different ways in which open-endedness may be introduced into an evolutionary system, which we will do in Section 4. I do not claim that the formalism presented in this section is particularly novel or of wide applicability beyond the current discussion; its main purpose is to emphasize various parameters and routes of interaction between processes that are not usually explicitly denoted. Within the

⁹ As we will see in Section 4.3.2, this can come about where there is a *non-additive compositional complexity* in the building blocks of the phenotype, or the presence of a *transdomain bridge*.

context of this article it provides a useful tool for considering various interactions of interest in an evolutionary system, and how they relate to open-endedness.

Considering evolutionary systems in general—including, for example, biological evolution, genetic algorithms, evolutionary robotics systems, and systems of self-reproducing computer code—we can discern three fundamental processes that any such system must instantiate in some form or other:¹⁰

1. The *generation* of the phenotypic behavior of an individual from its informational (genetic) description.
2. The *evaluation* of phenotypes to determine which ones get to reproduce. In its most general form the evaluation also determines the schedule of reproduction (rate and number of offspring) and lifetime of the individual.
3. The *reproduction with variation* of successful individuals.

The explicitness and complexity of implementation of each of these processes varies significantly from one type of system to another. In some cases a process might be implemented *extrinsically* as a special-purpose hard-coded mechanism acting upon the system, whereas in other cases the process might be provided *intrinsically* by a mechanism implemented within the system itself. Intrinsic mechanisms may rely exclusively upon the general laws of dynamics of the system (e.g., a simple self-replicating molecule), or they may be under sophisticated evolved control provided by the organisms themselves (e.g., the generation and reproduction mechanisms of modern biological organisms).

In some cases it may be easy to overlook the presence of a particular process; for example, in systems such as Tierra [33] and Avida [28], one might think there is no process of generation from genotype to phenotype, but a closer look shows that the phenotype comes about through the action of the system's (virtual) CPU that executes the instructions present in a program's genotype [36]. One way or another, these three processes are implemented by *all* evolutionary systems.

A schematic overview of how the three processes act upon a population of individuals is shown in Figure 3. Each process is explained in more detail below, and a formalism is introduced to make explicit various aspects of each process and interrelationships between the processes.¹¹

3.1 Generation

The process of generation can be represented in a very general form as follows:

$$p = M_L(g, c_a, c_b) \quad (1)$$

where g is the genotype, p is the resulting phenotype, M is the function that generates p from g , that is, the genotype-phenotype (G-P) map, L indicates fixed global laws (“laws of physics”) acting upon the system (which may contribute to determining the outcome of the generation process, e.g., self-organizational processes arising from laws of physics and chemistry in the biosphere, or the CPU interpretation of instructions in Tierra), c_a indicates the local abiotic context (environmental conditions) in which the generation process occurs, and c_b indicates the local biotic context (influence of other organisms on the process).

¹⁰ The three core processes of a Darwinian evolutionary process are often stated as *variation*, *differential reproduction*, and *inheritance*. This ignores the process of *generation* (of phenotypic behavior from genetic description) stated in the list given here, which is important in the current context. On the other hand, we collapse the processes of *variation* and *inheritance* into a single process *reproduction with variation* as a simplifying step in this discussion. This is a valid simplification as long as we are only dealing with systems where the main source of variation among individuals arises during the reproduction of a parent(s) to generate an offspring.

¹¹ The formalism is by no means complete, but it does at least emphasize influences on, and interrelationships between, the three processes. Some weaknesses of the formalism are discussed at the end of the article, along with suggestions for future improvements.




Figure 3. Schematic overview of key processes that must be implemented by any evolutionary system. Note that the timing and duration of each process do not necessarily need to be the same for each organism in the system, and the total number of individuals does not necessarily need to be constant from one generation to the next.

3.2 Evaluation

The evaluation of a phenotype to determine its evolutionary significance, that is, if and when it reproduces, how many offspring it produces, whom it mates with, and how long it lives, can be represented very generally as follows:

$$(l, \bar{s}_r, \bar{p}_m) = E_L(p, c_a, c_b) \tag{2}$$

where E is the evaluation function, L indicates fixed global laws acting upon the system (which may contribute to determining the outcome of the evaluation, e.g., laws of aerodynamics determining the ability of a bird to fly), p is the phenotype, c_a and c_b are the local abiotic and biotic contexts (as above), l is the resultant lifetime of the phenotype as determined by the evaluation process, \bar{s}_r is a vector representing the phenotype's resultant *reproduction schedule* (i.e., the number and timing of applications of the reproduction process on the individual), and \bar{p}_m is a vector representing the phenotype's resultant *mate set*, that is, the mate(s) that will participate in the individual's reproduction process (in the most general case, this set may be empty or of any nonempty size).

3.3 Reproduction with Variation

Finally, the reproduction process can be represented in general form as follows:

$$g' = R_L^{\bar{s}_r}(p, \bar{p}_m) \tag{3}$$

where R is the reproduction function, L indicates fixed global laws acting upon the system (which may contribute to determining the outcome of the reproduction process, e.g., by specifying global

mutation rates), p is the phenotype, \bar{p}_m is the mate set as determined by the evaluation process, \bar{s}_r is the reproduction schedule as determined by the evaluation process, and g' is the resultant new genotype.¹²

The reproduction function may incorporate any of a variety of different procedures depending upon the evolutionary system under consideration, including mutations of various kinds, recombination, gross chromosomal rearrangements (GCRs), error correction mechanisms, and so on.

Note that R is stated as a function of p rather than g . It is assumed here that p has access to the original g that created it, so that R could produce the new g' by simply copying g . But using p in the function allows for a more general representation that can also describe the transmission of acquired characteristics from p to g' (Lamarckian evolution) if relevant.

4 Routes to Achieving Open-Endedness

Having covered the three different kinds of open-ended evolution and a general formalism with which to describe the key processes of an evolutionary system, I now use the formalism to identify various routes by which open-endedness can be introduced in the design of an evolutionary system. As the formalism encapsulates a fairly standard Darwinian view of evolution (with extra emphasis on inputs and parameters of processes), the idea is to see how far this view will get us in explaining how open-endedness may arise in an evolutionary system.

In a simple system of noninteracting individuals that reproduce with variation according to static evaluation and reproduction functions, the individuals will evolve towards a local optimum in the adaptive landscape. At this point, stasis (or at least quasistability) will be reached. We are therefore looking for routes by which the organisms (p in the equations) can influence the three evolutionary processes described in the formalism, to enable them to escape from this situation and maintain an ongoing exploration of P-space. These could involve an organism causing a change in the implementation of one of the functions defined in Equations 1–3, or causing a change to one of the parameters of those equations. All such routes are mapped out in Figure 4; the top half of the figure (the line labeled “Organism 1” and the arrows above it) shows how an organism p can affect the evolutionary processes associated with its own lineage, and the bottom half (the line labeled “Organism 2” and the arrows above it) shows how other organisms can affect the parameters of the evolutionary processes of organism 1.¹³ Each of these routes is discussed in more detail in Section 4.2.

As revealed in the following discussion, an analysis of open-endedness based upon the formalism only really addresses issues concerning *exploratory* open-endedness. This illustrates why traditional approaches to modeling evolutionary systems based upon the processes of generation, evaluation, and reproduction with variation do not provide much insight into the more interesting kinds of open-endedness, namely, expansive and transformational open-endedness. In the following discussion I also suggest routes by which these other two types of open-endedness can be achieved, although these are more tentative suggestions offered without the support of the formalism.

Before discussing the different routes specifically, I begin with some general comments on the distinction between intrinsic and extrinsic implementations of the evolutionary processes.

4.1 Intrinsic and Extrinsic Implementations

A cross-cutting issue in the quest for open-endedness described in the following discussion is the extent to which each of the specific processes is defined *intrinsically* within the system by being implemented through the components and dynamics of the system itself. In contrast, all existing artificial evolutionary systems define some or most of these processes *extrinsically* to the evolving

¹² In Equation 3, \bar{s}_r is presented as a superscript to indicate that it determines *when* R is applied.

¹³ The figure does not show an arrow from organism 1's p to the c_b parameters of M_L and E_L because c_b is, by definition, the context caused by *other* organisms.

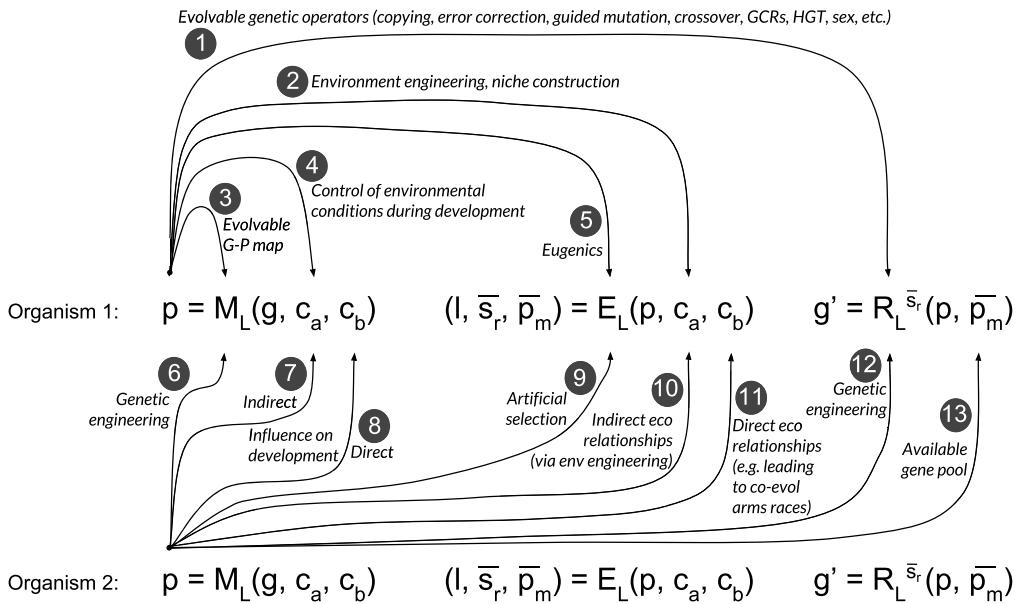


Figure 4. Potential routes to exploratory open-endedness in an evolutionary system.

system as a special-purpose hard-coded mechanism. Banzhaf et al. refer to extrinsically implemented mechanisms as *shortcuts* [1, p. 146].

The importance of using an intrinsic evaluation process in computational models of biological evolution has been recognized for a long time (e.g., [29]) and indeed was a feature of some of the earliest implementations of computational evolutionary systems (e.g., [2, 9]). Here I consider the benefits of implicit implementations not just of the evaluation process, but also of the generation and reproduction processes. The key benefit of processes instantiated intrinsically by being explicitly implemented within the system itself is that it allows the possibility that the implementation—the process—can itself change. This opens the door for the G-P mapping, the evaluation processes, and the reproduction and variation processes to evolve as the system unfolds.

While it is possible to imagine extrinsically coding not just a process but also mechanisms for changing the process, such a process would still only be able to change and evolve in the hard-coded ways provided by the extrinsically defined change mechanism. In contrast, for intrinsically implemented processes, not only might the process evolve, but *the evolvability of the process* might itself evolve.

4.2 Exploratory Open-Endedness

As stated above, in a simple system of noninteracting individuals that reproduce with variation according to static evaluation and reproduction functions, the individuals will evolve towards a local optimum in the adaptive landscape beyond which the variational methods of the reproduction function can no longer take them to a state with higher fitness.

To introduce the possibility of ongoing adaptive novelty we therefore need mechanisms that can “shake up” the system. This can be achieved by allowing for intrinsic means for ongoing modification of the adaptive landscape experienced by an individual (i.e., changes relating to the evaluation function E_L), of the topology of genetic space (i.e., changes relating to the reproduction and variation function R_L), or of the mapping between genotype and phenotype (i.e., changes relating to the generation function M_L). As explained above, the routes for implementing these kinds of mechanisms within our formalism are identified in Figure 4. We now look at each of these in more detail.

If the reproduction and variation processes (R_L) are themselves implemented intrinsically by the individual organisms (route 1 in Figure 4), the individuals might be able to jump out of the local optimum by bringing new areas of G-space and P-space into reach of the variational operators. Route 1 represents topics in the literature concerning the *evolution of evolvability* (evo-evo), such as evolvable genetic operators, including copying processes, error correction, mutator genes for guided mutations, crossover mechanisms, gross chromosomal rearrangements, horizontal gene transfer, and so on (e.g., [10, 32]).

Another route to improving evolvability is provided by route 3: allowing the G-P map (M_L) to evolve by having it intrinsically implemented by the individual organisms. As discussed earlier (see Figure 2), the nature of the G-P map dictates which regions of P-space can be easily explored. Implementing the G-P map intrinsically may allow it to evolve so that mutations are more likely to produce adaptive variations in P-space. Route 3 represents topics in the literature such as evo-devo, facilitated variation, and developmental robustness (e.g., [16, 26]).

An alternative possibility for an organism to influence the result of the G-P mapping process is for it to exert some control over the environmental conditions under which the process occurs (in the case where the mapping is sensitive to such conditions); this is represented by route 4 in Figure 4. However, in biological systems this is more likely to be a route whereby organisms *reduce* variability (e.g., canalization [42]) rather than increase it.

While changes to M_L and R_L via routes 1 and 3 (evo-evo and evo-devo) might be sufficient to prevent the system becoming stuck in a local optimum state, they are not in themselves sufficient to achieve an ongoing exploration of P-space, as the system will still halt when all individuals have reached a statically defined optimum fitness. In order to provide an ongoing drive for exploring P-space, ongoing change in the adaptive landscape is required. This can be achieved by inducing ongoing changes to an individual's evaluation function E_L , which can be realized via routes 2, 5, 9, 10, or 11.

Routes 2, 10, and 11 all change the context in which the individual is evaluated: Routes 2 and 10 change the abiotic context, and route 11 changes the biotic context. In route 2 the change is brought about by the focal individual itself (e.g., by environmental engineering [20], or niche construction over longer time scales [27]), whereas in route 10 it is brought about by the other individuals that influence the evaluation function (e.g., environmental engineering by other species). Route 11 represents direct ecological relationships, leading to processes such as coevolutionary arms races [40].

Beyond changing the context in which evaluation occurs, it might also be possible to intrinsically change the evaluation function (E_L) itself (i.e., to change the fundamental factors determining an individual's longevity and fecundity). This might arise when an intelligent species that practices artificial selection (e.g., farming or eugenics) evolves within the system. These mechanisms are represented by routes 5 and 9.¹⁴

Regarding the other two functions (M_L and R_L), it would be very unusual to have a system where one organism could directly change another organism's implementation of these functions. However, this is at least conceivable, and perhaps the human species is approaching that ability with advances in genetic engineering (routes 6 and 12).

Additional routes through which other individuals in the system might promote ongoing exploration of P-space include routes 7, 8, and 13. Routes 7 and 8 are processes whereby the production of a phenotype from a genotype is affected by the local context: route 7 represents the local abiotic environment, and route 8 the local biotic environment. Neighboring individuals can be involved in route 7 as well as route 8, through the processes of environment engineering and niche construction. Route 13 represents the available gene pool provided by potential mates in the local context, that is, the raw material upon which the genetic operators might act in route 1.

¹⁴ Note that this case of an intrinsically defined evaluation function is different to the case of evolutionary algorithms that explicitly avoid imposing notions of objective fitness, such as *minimal criterion novelty search* [22]. These algorithms still have an extrinsic evaluation function at their core (in the case of novelty search, this includes the method by which the novelty of a new organism is calculated).

The role of all of the processes involving interactions with other individuals via the parameters of the functions (routes 7, 8, 10, 11, and 13) can be boosted if the local context experienced by an individual and its descendants changes over time. An obvious route for achieving this is through the provision of a spatial environment and means by which individuals can move (actively or passively) around the environment.

It should also be recognized that the formalism developed here is not exhaustive, as it concentrates only on processes that affect individuals. It does not explicitly deal with population-level effects that are also relevant in promoting ongoing exploration of G-space and P-space. Topics from the evolutionary population dynamics literature such as finite sampling, drift, adaptive radiations, and neutral networks are additional mechanisms by which the ongoing exploration of the adaptive landscape might be promoted.

To summarize the preceding discussion, the routes to open-endedness depicted in Figure 4 represent mechanisms by which organisms can promote ongoing evolutionary activity by modifying the adaptive landscape, the topology of genetic space, or the nature of the G-P map.

However, while these routes promote ongoing activity within a given P-space, none of them cause the expansion of P-space itself. In other words, the routes to open-endedness suggested by a fairly standard Darwinian analysis, as represented by the formalism, relate to *exploratory* open-endedness only; they do not directly help us in our search for *expansive* or *transformational* novelties.

4.3 Expansive and Transformational Open-Endedness

We will now discuss concepts not explicitly covered by the formalism that may be required to produce the other kinds of open-endedness involving expansive and transformational novelties. These more interesting kinds both involve the discovery of *door-opening*¹⁵ states in P-space that open up an expanded space of new adjacencies, as exemplified by the red flashes in Figure 1(b) and (c).

There are various issues involved in how these might come about in an evolutionary system. The following discussions address two of the most important questions:

1. Where does this extra space of possibilities come from?
2. How can the evolutionary system access the new states via intrinsic mechanisms?

4.3.1 Question 1: Expanding the State Space

Regarding question 1, in the biological world the answer is that the extra space was always there in the complexity of the laws of physics and chemistry—it is just a matter of biological systems evolving to make use of the existing complexity (by methods pertaining to question 2). Engineered *physical* evolutionary systems can also make use of this existing complexity—indeed, the most impressive instances of transformational novelties arising in artificial systems are found in physical systems (e.g., [3, 6]).

This situation exemplifies the fact that our definitions of novelty and open-endedness, as presented in Section 2, are defined *relative to our model* (and meta-model) of the system. The state space of the actual system has not expanded, but an expansive or transformational novelty reveals a *deficiency in the model* of the system regarding its ability to describe the actual system. These kinds of novelties therefore require an expansion of the model.

In the case of computational evolutionary systems, the same solution of providing a world with rich possibilities for complexity in its laws of dynamics and interactions is also an option. It is notable that most existing ALife work with computational evolution takes place in very impoverished virtual environments. But there is also another possibility with computational systems: to dynamically increase P-space as the system unfolds. One route by which this might be achieved would be to open up the system by allowing it to access additional resources on the Internet (e.g., stock trading agents with the ability to discover and utilize new online data sources to improve their performance).¹⁶

¹⁵ To borrow a term from Bedau [38].

¹⁶ This idea has been discussed by Boden among others [5]. See [37] for many pointers to how this might be implemented.

4.3.2 Question 2: Accessing New States

Consideration of the mechanisms involved in biological evolution suggests at least two general ways in which Question 2 can be addressed:

(a) *Domains, exaptations, and transdomain bridges.* Components in physical systems possess multiple properties in different domains (mechanical, chemical, electrical, responsiveness to electromagnetism, pressure, etc.). Indeed, the distinction between an expansive and a transformational novelty can be viewed as the difference between a door-opening novelty in the same domain and a door-opening novelty in a different domain, respectively. In this view, the distinction between expansive and transformational novelty depends upon an observer's ontology of domains; this is a more specific interpretation of the picture of models and meta-models introduced earlier.

A common mode by which innovations arise is *exaptation*, where a structure originally selected for its properties in one domain coincidentally has adaptive properties in a different domain, which then become a new focus of selection [18]. In this situation, the multi-property component has acted as a *transdomain bridge* to open up a new domain for potential exploitation by the organism—this would represent a transformational novelty. This mechanism can also produce expansive novelties if the components have multiple properties within the same domain, for example, multifunctional enzymes [21].¹⁷ The latter case can be labeled a *intradomain bridge*.

Another example, provided by Dawkins, is the evolutionary appearance of segmented body plans in animals [11]. While the first segmented animal might have been unremarkable in terms of its functionality, and just an exploratory novelty, it gave rise to a radiation leading to a whole new phyla with new possibilities for behavior (i.e., expansive and/or transformational novelties), such as new possibilities for locomotion arising from the free movement afforded to organisms with a segmented spinal cord. In close alignment with our terminology of “door-opening novelties,” Dawkins describes discoveries of this kind as “watershed events ... that open floodgates to future evolution” [11, p. 218].¹⁸

Most computational evolutionary systems lack significantly multi-property components, and therefore miss this route to transformational novelty. The examples that currently benefit the most from this route to novelty are those in which the evolving agents are embedded within a simulated physical environment [39].

(b) *Non-additive compositional systems.* An alternative route for accessing new states hinges on the mechanism by which a phenotype is generated from a genotype (M_L). To take a very general view, we can see this process as the construction of a structure and/or behavior by the specific arrangement of a number of components drawn from a given set of component types. I will call this mode of construction a *compositional system*.¹⁹ Examples from biology range from the construction of a protein from amino acids drawn from a set of 20 different types, to the construction of an termite colony from termites drawn from a set of different castes. Examples from ALife include the construction of a neural network controller from a given number of neurons and connections. In many cases, particularly in biology, there may be hierarchical levels of composition; see [1] for an extensive discussion of levels and hierarchies.

Compositional systems can arise in many different domains, such as chemistry, physics, and information systems. To a first approximation, we could view prokaryotic life as an exploration of compositional *chemistry* and multicellular eukaryotic life as an exploration of compositional *physics*.²⁰ Furthermore, animals with nervous systems, and ALife agents with evolved controllers, engage in the exploration of compositional *information* systems.

¹⁷ The importance of multifunctional components for biological evolvability and robustness has been argued by various authors (e.g., [17, 45]).

¹⁸ The distinction between the exploratory discovery of a door-opening state and the potential it introduces for expansive or transformational novelties in function is similar to G. Wagner's distinction between (his conception of) *novelties* and *innovations* [44]. I further discuss this distinction at the end of the article.

¹⁹ I use the term *compositional* rather than Boden's [5] *combinational* to emphasize that the size of structures may increase, and that the specific arrangement and connections between components might be important.

²⁰ This is obviously a gross simplification, as all domains of life utilize both chemistry and physics.

Note that the ability of a lineage to concurrently explore multiple compositional domains is in itself an enabler of exploratory open-endedness, as it can prevent evolution from getting stuck in a local optimum in any one domain by providing an *extradimensional bypass*.²¹

We can distinguish between *additive* and *non-additive compositional systems*. For *additive* systems, the functionality of the resulting product is an amplification of the existing function of the components (e.g., joining a number of batteries in series to create a new battery with a greater voltage). For *non-additive* systems, the act of composition can introduce new functionality depending upon the specific arrangement and connections between the parts (e.g., composing a computer algorithm out of a specific set of subroutines and individual instructions). In some non-additive compositional systems such as biomolecular chemistry, this can also be a route to accessing new domains (e.g., with the production of a photoreceptor protein such as rhodopsin from its amino acid sequence).²²

While additive compositional systems result only in exploratory novelty, they can play an important role in enabling later expansive or transformational novelties. To take the battery example mentioned above, the creation of a new battery with higher voltage does not introduce new functionality in itself, but the higher voltage might make other processes and reactions possible that were not previously achievable. Another example is the previously discussed case of the evolutionary appearance of segmented body plans in animals. So additive compositional systems can create door-opening states that subsequently lead to expansive or transformational novelties.

In contrast, *non-additive* compositional systems can lead *directly* to expansive or transformational novelties. For example, building proteins from amino acids can produce new molecules possessing *expansive* novelties in the chemical reaction repertoire: “Once a new molecule appears for the first time in the chemosphere new interactions and further adjacencies emerge” [12, p. 4]. As already mentioned above, protein building can also potentially lead to *transformational* novelties, as might be the case with the production of a photosensitive rhodopsin molecule.

I close this section with a few remarks about compositional systems and the evolution of complexity. While the complexity of organisms and interactions does not necessarily increase in an evolutionary system, those that employ compositional systems in the production of phenotypes have a clear capacity for *cumulative compositional complexity* as evolution builds upon what has gone before. This capacity would appear to be particularly pronounced in non-additive compositional systems, where new compositions can offer direct routes to expansive and transformational novelty. Furthermore, compositional systems able to cumulatively produce hierarchical organizations are particularly suitable as a basis for the evolution of complexity [34]. Increases in complexity in these cases will be aided by the usual drivers of complexity discussed in the evolutionary biology literature, such as coevolutionary arms races [40].

5 Final Remarks

The framework presented above can act as a guide for categorizing and comparing the OE potential of existing systems. For example, von Neumann’s CA implementation of a self-reproducing system [41] concentrated heavily on the role of the laws of dynamics (L) in its intrinsic implementation of the evaluation function $E_L(p, c_a, c_b)$, but ignored the organism’s local context (c_a and c_b), making the system very brittle to perturbations. Tierra [33] implements E_L intrinsically, but M_L is extrinsic and trivial, and the abiotic environment (as represented by the laws of dynamics, L) is very impoverished. Geb [7] features intrinsic E_L applied to non-additive compositional controllers (neural networks), but implements M_L and R_L extrinsically. Most implementations of novelty search [23] implement two or all three key processes (M_L , E_L , and R_L) extrinsically—although in many cases this is applied to non-additive compositional controllers and other compositional systems. A comprehensive examination of existing systems along these lines would provide clear indications of how the OE potential of future systems could be improved.

21 This concept was introduced by Conrad [8] and later named by Gavrillets [15].

22 It is also possible that a system might have a mixture of different types of components, some of which are additive and others non-additive.

As demonstrated in the preceding discussion, the framework can act as a map of the territory of open-endedness. This is useful for showing how the diverse body of relevant existing theory fits into the overall picture, in addition to aiding the categorization and comparison of systems as outlined above. The discussion has revealed that considerations of generation, evaluation, and reproduction with variation indicate routes to *exploratory* open-endedness only; in order to understand the more interesting cases of *expansive* and *transformational* open-endedness, we need to consider not the properties traditionally studied by population genetics, but rather the nature of the building blocks out of which individual organisms are constructed, and the laws and properties of the environment in which they exist.

The presented framework makes explicit various influences and interrelationships between the fundamental processes required for evolution. Nevertheless, as stated above, there are clearly areas where the framework could be further improved. It is currently weak at representing important processes occurring above the level of the individual organism. For example, Figure 4 does not currently capture evolutionary population dynamics concepts such as finite sampling, drift, neutral networks, and so on. The framework could certainly be expanded to be more explicit about such population-level effects.

The current work might also benefit from a more sophisticated treatment of the concept of *behavior*. One approach would be to replace the current use of G-space and P-space with a threefold distinction between *parameter space*, *organization space*, and *action space*. Items in parameter space are informational specifications of particular organizations (configurations of material) in organization space. These organizations are situated in a specific environment that provides boundary conditions (i.e., local environmental context provided by abiotic and potentially other biotic organization) and laws of physics to determine the resulting action of the organization in its environment.²³ Such an approach would make explicit the distinction between novelties in organization space and novelties in action space—the importance of which has recently been highlighted by G. Wagner in his analysis of evolutionary innovations [44].

Under this view, a *behavior* can be defined as the change in state of one organization brought about by the action of another organization in the environment,²⁴ where the *action* of an organization is the result of the application of the global laws of physics to it within the context of boundary conditions provided by the local biotic and abiotic environment.²⁵ Biological function can then be seen as *purposeful* behavior, where *purpose* comes about through evolutionary selection upon evolving organizations.²⁶

To give a concrete example in a mechanical context, imagine that we have one organization comprising a wound spring and a cogwheel, both attached to an axle, situated close to another organization comprising a cogwheel attached to another axle. The positions of the two organizations are such that the teeth of the two cogwheels interlock. Then, as the spring unwinds under the laws of mechanics, it induces a rotational behavior in the first organization (under the boundary condition of a fixed translational position provided by its axle). Furthermore, this also induces a rotational behavior in the second organization (which rotates under its own axle boundary condition).

An advantage of this extended approach is that it makes explicit the processes involved in door-opening events, as exemplified by the red flashes in Figure 1(b) and (c). These involve first the discovery of a special new state in the current organization space, which then produces novel behaviors

23 Similarly, the view of the process of generation discussed here could be elaborated into distinct processes of instantiation of an organism followed by a process of self-maintenance, where both of these are potentially subject to the laws of physics of the system and the local biotic and abiotic context.

24 We also allow the organization being acted upon to be the same as that doing the acting, giving self-directed behavior. Note that this is a general view of action and behavior that is not confined to organisms but could be applied to *any* organization of matter, including subsystems within organisms or even abiotic organizations.

25 This view can be further expanded to allow for elaborate evolved organizations that control their local context to produce very reliable behaviors (we might call these *contrivances*), and for hierarchical organizations comprising multiple contrivances interacting with each other. The distinction between organization, local context, and behavior is just as relevant within a single individual as it is in an individual's interactions with the external world.

26 This view is inspired by Pattee's treatment of laws, initial conditions, measurements, and semantic closure when discussing living systems [30, 31].

in action space. New behaviors representing expansive or transformational novelties can arise if the organization is able to utilize a new feature of the laws of physics that was not previously exploited, or if the organization represents a new boundary condition or contrivance that reliably generates behaviors that would have been very unlikely before the appearance of the new organization.²⁷

To extend the previous mechanical example, a transformational novelty could arise in our cog-wheel system if one or both of the organizations are augmented with a cam and follower rod that hits a metal sheet when at its full reach; the rotation of the cam would convert the rotational movement of the axle to a discrete linear movement in the follower, causing the follower to hit the metal sheet at its full reach. The result would be a discrete regular sound of the rod hitting the metal sheet, which would represent a transformational novelty in the system (assuming our model did not already include the concept of discrete sound generation).

By explicitly representing behavior as an interaction between two (or more) organizations brought about by the action of the laws of physics together with boundary conditions provided by the local biotic and abiotic context, this view can accommodate Kaufmann et al.'s notion of *unprestatability* [24]—for a given organization or structure, any number of behaviors are possible depending upon the application of the laws of physics upon it in a specific local context.

The view set out in this article and in the suggested extensions just discussed offers a new perspective on open-ended evolution—one that fundamentally comprises just two essential processes: the ongoing exploration of a phenotype space (as exemplified in Figure 1(a) and Figure 4), and the discovery of door-opening states in that space that open up an expanded phenotype space (as exemplified by the red flashes in Figures 1(b) and (c)). The former involves many established areas of theoretical biology, as illustrated in Figure 4 and in the accompanying discussion, and the latter relates to the emerging topic in the biological literature of *evolutionary innovation* (e.g., [12, 19]). As the topic of innovation is currently attracting growing attention from biologists, there is rich potential for a profitable two-way exchange of ideas between those interested in biological innovations and those interested in open-endedness in other kinds of evolutionary systems.

Acknowledgments

I am grateful to Peter Turney, Alastair Channon, Adam Stanton, James Borg, Bradly Alicea, and two anonymous reviewers for their comments and discussions, which have helped in developing the ideas set out in this article.

References

1. 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*, 135(3), 131–161.
2. Barricelli, N. A. (1957). Symbiogenetic evolution processes realized by artificial methods. *Methodos*, 9, 35–36.
3. Bird, J., & Layzell, P. (2002). The evolved radio and its implications for modelling the evolution of novel sensors. In *Proceedings of the 2002 Congress on Evolutionary Computation, 2002. CEC'02, Vol. 2* (pp. 1836–1841). Piscataway, NJ: IEEE Press.
4. Boden, M. A. (2004). *The creative mind: Myths and mechanisms*, 2nd ed. (revised, expanded; 1st ed., 1990). London: Routledge.
5. Boden, M. A. (2015). Creativity and ALife. *Artificial Life*, 21(3), 354–365.
6. Cariani, P. (1993). To evolve an ear. Epistemological implications of Gordon Pask's electrochemical devices. *Systems Research and Behavioral Science*, 10(3), 19–33.
7. 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.

²⁷ It is also possible that a new behavior could arise simply by two existing organizations being brought into a new relationship with each other such that they represent new boundary conditions for each other's actions (e.g., two existing organisms carrying out actions that they were already capable of, but doing so in proximity to each other such that their actions affect each other in a new way).

8. Conrad, M. (1990). The geometry of evolution. *BioSystems*, 24, 61–81.
9. Conrad, M., & Pattee, H. (1970). Evolution experiments with an artificial ecosystem. *Journal of Theoretical Biology*, 28, 393–409.
10. Crombach, A., & Hogeweg, P. (2008). Evolution of evolvability in gene regulatory networks. *PLoS Computational Biology*, 4(7), e1000112.
11. Dawkins, R. (1988). The evolution of evolvability. In C. Langton (Ed.), *Artificial Life* (pp. 201–220). Reading, MA: Addison-Wesley.
12. de Vladar, H. P., Santos, M., & Szathmáry, E. (2017). Grand views of evolution. *Trends in Ecology & Evolution*, 32(5), 324–334.
13. Dennett, D. C. (2013). *Intuition pumps and other tools for thinking*. New York: W. W. Norton.
14. Dietrich, M. R., & Skipper, R. A., Jr. (2012). A shifting terrain: A brief history of the adaptive landscape. In E. Svensson & R. Calsbeek (Eds.), *The adaptive landscape in evolutionary biology* (pp. 3–15). Oxford, UK: Oxford University Press.
15. Gavrilets, S. (1999). A dynamical theory of speciation on hole adaptive landscapes. *The American Naturalist*, 154(1), 1–22.
16. Gerhart, J., & Kirschner, M. (2007). The theory of facilitated variation. *Proceedings of the National Academy of Sciences of the U.S.A.*, 104(suppl. 1), 8582–8589.
17. Goldenfeld, N., & Woese, C. (2011). Life is physics: Evolution as a collective phenomenon far from equilibrium. *Annual Review of Condensed Matter Physics*, 2(1), 375–399.
18. Gould, S. J., & Vrba, E. S. (1982). Exaptation—a missing term in the science of form. *Paleobiology*, 8(1), 4–15.
19. Hochberg, M. E., Marquet, P. A., Boyd, R., & Wagner, A. (2017). Innovation: An emerging focus from cells to societies. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 372, 1735.
20. Jones, C. G., Lawton, J. H., & Shachak, M. (1997). Positive and negative effects of organisms as physical ecosystem engineers. *Ecology*, 78(7), 1946–1957.
21. Kacser, H., & Beeby, R. (1984). Evolution of catalytic proteins. *Journal of Molecular Evolution*, 20, 38–51.
22. Lehman, J., & Stanley, K. O. (2010). Revising the evolutionary computation abstraction: Minimal criteria novelty search. In *Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation* (pp. 103–110). New York: ACM.
23. Lehman, J., & Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2), 189–223.
24. Longo, G., Montévil, M., & Kauffman, S. (2012). No entailing laws, but enablement in the evolution of the biosphere. arXiv: 1201.2069.
25. Maynard Smith, J., & Szathmáry, E. (1995). *The major transitions in evolution*. New York: W. H. Freeman.
26. Nuño de la Rosa, L. (2017). Computing the extended synthesis: Mapping the dynamics and conceptual structure of the evolvability research front. *Journal of Experimental Zoology Part B: Molecular and Developmental Evolution*, 328(5), 395–411.
27. Odling-Smee, F. J., Laland, K. N., & Feldman, M. W. (2003). *Niche construction: The neglected process in evolution*. Princeton, NJ: Princeton University Press.
28. Ofria, C., & Wilke, C. O. (2004). Avida: A software platform for research in computational evolutionary biology. *Artificial Life*, 10(2), 191–229.
29. Packard, N. H. (1988). Intrinsic adaptation in a simple model for evolution. In C. Langton (Ed.), *Artificial Life* (pp. 141–155). Reading, MA: Addison-Wesley.
30. Pattee, H. (1995). Artificial life needs a real epistemology. In F. Morán, A. Moreno, J. Merelo, & P. Chacón (Eds.), *Proceedings of ECAL 1995* (pp. 23–38). Berlin: Springer-Verlag.
31. Pattee, H. (1995). Evolving self-reference: Matter, symbols, and semantic closure. *Communication & Cognition: Artificial Intelligence*, 12(1–2), 9–28.
32. Pigliucci, M. (2008). Is evolvability evolvable? *Nature Reviews Genetics*, 9(1), 75.

33. 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.
34. Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6), 467–482.
35. Stanley, K. O., Lehman, J., & Soros, L. (2017). Open-endedness: The last grand challenge you’ve never heard of. O’Reilly Ideas. Available at <https://www.oreilly.com/ideas/open-endedness-the-last-grand-challenge-youve-never-heard-of> (accessed January 2019).
36. Taylor, T. (2001). Creativity in evolution: Individuals, interactions and environments. In P. J. Bentley & D. W. Corne (Eds.), *Creative evolutionary systems* (pp. 79–108). Burlington, MA: Morgan Kaufman.
37. Taylor, T., Auerbach, J. E., Bongard, J., Clune, J., Hickinbotham, S., Ofria, C., Oka, M., Risi, S., Stanley, K. O., & Yosinski, J. (2016). WebAL comes of age: A review of the first 21 years of artificial life on the web. *Artificial Life*, 22(3), 364–407.
38. Taylor, T., Bedau, M., Channon, A., Ackley, D., Banzhaf, W., Beslon, G., Dolson, E., Froese, T., Hickinbotham, S., Ikegami, T., McMullin, B., Packard, N., Rasmussen, S., Virgo, N., Agmon, E., Clark, E., McGregor, S., Ofria, C., Ropella, G., Spector, L., Stanley, K. O., Stanton, A., Timperley, C., Vostinar, A., & Wisser, M. (2016). Open-ended evolution: Perspectives from the OEE workshop in York. *Artificial Life*, 22(3), 408–423.
39. Taylor, T., & Massey, C. (2001). Recent developments in the evolution of morphologies and controllers for physically simulated creatures. *Artificial Life*, 7(1), 77–87.
40. Van Valen, L. (1973). A new evolutionary law. *Evolutionary Theory*, 1, 1–30.
41. von Neumann, J. (1966). *Theory of self-reproducing automata*. Urbana, IL: University of Illinois Press.
42. Waddington, C. H. (1942). Canalization of development and the inheritance of acquired characters. *Nature*, 150(3811), 563.
43. Waddington, C. H. (1969). Paradigm for an evolutionary process. In C. H. Waddington (Ed.), *Towards a theoretical biology*, Vol. 2 (pp. 106–128). Edinburgh: Edinburgh University Press.
44. Wagner, G. P. (2015). Evolutionary innovations and novelties: Let us get down to business! *Zoologischer Anzeiger*, 256, 75–81.
45. Whitacre, J. M. (2010). Degeneracy: A link between evolvability, robustness and complexity in biological systems. *Theoretical Biology and Medical Modelling*, 7(6), 1–17.