

Routes to Open-Endedness in Evolutionary Systems

Tim Taylor^{1,2}

¹Independent Researcher, Edinburgh, U.K.

²Associate Examiner, University of London Worldwide, London, U.K.

tim@tim-taylor.com

Abstract

This paper 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., three different kinds of open-endedness are identified: exploratory, expansive, and transformational. These are characterised in terms of their relationship to the search space of phenotypic behaviours. 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, it is demonstrated 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 behaviour, transdomain bridges, and non-additive compositional systems. In contrast to a traditional population genetics 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 paper ends with suggestions of how the framework can be used to categorise and compare the open-ended evolutionary potential of different systems, and how it might guide the design of systems with greater capacity for open-ended evolution.

Introduction

In this paper 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. (2016). In their work, Banzhaf et al. make the distinction 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” (Banzhaf et al., 2016, p. 135). One of the main aims of their paper 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” (Banzhaf et al., 2016, 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.

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 dependencies and interrelationships that are otherwise easy to overlook.

Equipped with the necessary preliminaries, I then utilise 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 this approach only helps in the investigation of one type of open-endedness. At the end of the paper I therefore discuss potential factors involved in the other kinds of open-endedness as well.

State Spaces, Novelty 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. State spaces are simpler than adaptive landscapes because 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 as it is sufficient for the purpose of the discussion; I consider how fitness comes into the picture later in the paper.

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 non-

trivial 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.¹ However, even if it can be difficult to quantitatively describe a specific complex evolutionary system, state spaces can still be useful *intuition pumps* (Dennett, 2013)—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. (2016)'s recent treatment of open-endedness. Central to their approach is the idea that the behaviour 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.

Banzhaf et al. (2016) 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 (Boden, 2004, 2015).

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* (Taylor et al., 2016b). 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:²

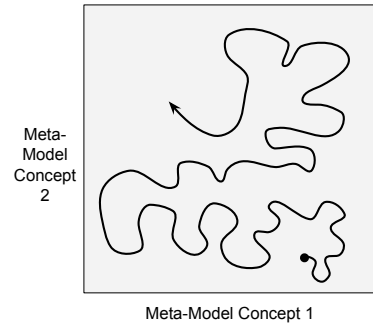
1. *Exploratory Novelty*: A novelty that can be described using the current model. Potential examples from biology include the production of a new combination of alleles on a genome, and a change in the number of vertebra in a new vertebrate species.

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 cur-

¹Although methods for inferring latent variable models can be employed to generate more meaningful *latent spaces*.


²Banzhaf et al. (2016) used the terms *variation*, *innovation* and *emergence*, respectively, in place of the terms used here. 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 (2015)'s concepts of *exploratory*, *combinational* and *transformational creativity*. As an example of the potential for confusion when using Banzhaf et al. (2016)'s terms, de Vladar et al. (2017) have recently used the term *innovation* to describe novelties that most closely match Banzhaf et al. (2016)'s *emergent* novelties.



(a) EXPLORATORY OPEN-ENDEDNESS



(b) EXPANSIVE OPEN-ENDEDNESS



(c) TRANSFORMATIONAL OPEN-ENDEDNESS

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

rent 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.

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 individuality, the appearance of winged flight, and the appearance of visual sensory systems.

Transformational Open-Endedness: The ongoing production of transformational exploratory novelties.

Note that Banzhaf et al. (2016) 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 (Banzhaf et al., 2016, p. 144–145); my interest in this paper 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 1.³

Note that I indicate “potential” examples in the list above, because according to Banzhaf et al. (2016)’s approach each type of novelty is defined *relative to a given model and meta-model*.

Furthermore, Banzhaf et al. (2016) 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 is 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 labelled transformational (and likewise for expansive novelties).⁴ The defining

³Any 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.

⁴cf. Boden (2015)’s distinction between I-creativity and H-creativity.

feature of these 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* (de Vlarar et al., 2017; Longo et al., 2012).

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 behaviour 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.

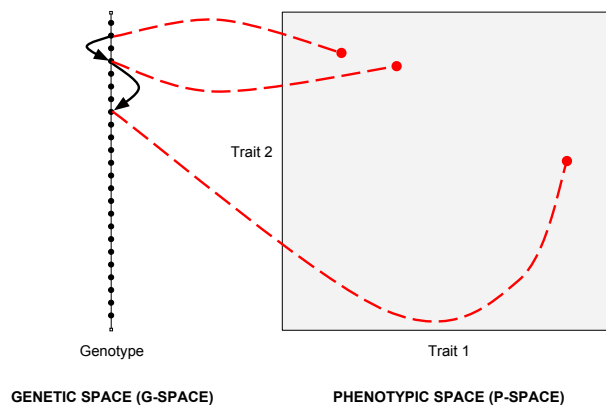


Figure 2: Genetic Space and Relation to Phenotypic Space. The mapping from G-space to P-space is defined by the G-P map (M_L in Equation 1). The mapping might be such that small moves in G-space can sometimes result in large moves in P-space. Note that the dimensionality of G-space might be different to that of P-space.

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 (the M_L function in Equation 1 introduced later on), which can be of arbitrary form and can also depend upon the system’s global laws of dynamics (L in Equation 1) and the local context in which the phenotype is generated (c_a and c_b in Equation 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. 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. These topics will be discussed in more detail later.

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* (Maynard Smith and Szathmary, 1995). All else being equal, a larger genome can (but does not necessarily) specify a more complicated phenotype. While this can indeed be the case, it depends on the capacity of the additional genes to specify new traits. This can be achieved (as in the fixed G-space case) through non-additive compositional complexity in the building blocks of the phenotype, or through transdomain bridges (to be discussed later).

Evolutionary Processes

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:

1. The *generation* of the phenotypic behaviour of an individual from its 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 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. In some cases it may be easy to overlook the presence of a particular process; for example, in systems such as Tierra (Ray, 1991) and Avida (Ofria and Wilke, 2004), 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 (Taylor, 2001). 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 in shows 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.




Figure 3: Schematic overview of key processes that must be implemented by any evolutionary system. Note that the timing and duration of each process does 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.

Generation

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

$$p = M_L(g, c_a, c_b) \tag{1}$$

where g is the genotype, p is the resulting phenotype, M is the function that generates p from g , i.e. the genotype-phenotype (G-P) map, L indicates fixed global laws acting upon the system (which may contribute to determining the outcome of the generation process, e.g. self-organisational 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).

Evaluation

The evaluation of a phenotype to determine its evolutionary significance, i.e. 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:

$$(l, \overline{s_r}, \overline{p_m}) = E_L(p, c_a, c_b) \quad (2)$$

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 context (as above), l is the resultant lifetime of the phenotype as determined by the evaluation process, $\overline{s_r}$ 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 $\overline{p_m}$ is a vector representing the phenotype's resultant *mate set*, i.e. 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 non-empty size).

Reproduction with Variation

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

$$g' = R_L^{\overline{s_r}}(p, \overline{p_m}) \quad (3)$$

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, $\overline{p_m}$ is the mate set as determined by the evaluation process, $\overline{s_r}$ is the reproduction schedule as determined by the evaluation process, and g' is the resultant new genotype.

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.

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 show how

the formalism can be used to identify various routes by which open-endedness can be introduced in the design of an evolutionary system. These routes are illustrated in Figure 4.

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 modelling 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, i.e. 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.

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 system as a hard-coded mechanism. Banzhaf et al. (2016) (p. 146) refer to extrinsically implemented mechanisms as *shortcuts*.

The importance in using an intrinsic evaluation process in computational models of biological evolution has been recognised for a long time, e.g. (Packard, 1988), and indeed was a feature of some of the earliest implementations of computational evolutionary systems (Barricelli, 1957; Conrad and Pattee, 1970). 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.

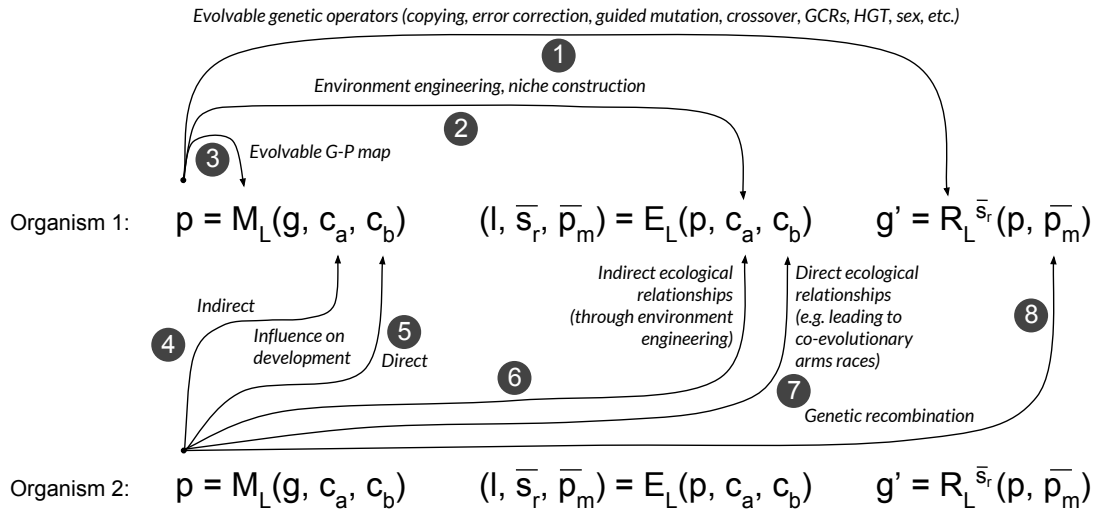


Figure 4: Potential routes to exploratory open-endedness in an evolutionary system.

Exploratory Open-Endedness

In a simple system of non-interacting 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.

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, etc.

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 potentially allows it to evolve such 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.

While these two routes (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 realised via Routes 2, 6 or 7 in Figure 4. These routes all change the context in which the individual is evaluated: Routes 2 and 6 change the abiotic context, and Route 7 changes the biotic context. In Route 2 the change is brought about by the focal individual itself (e.g. by environmental engineering, or niche construction over longer timescales), whereas in Route 6 it is brought about by the other individuals that influence the evaluation function (e.g. environmental engineering by other species). Route 7 represents direct ecological relationships, leading to processes such as co-evolutionary arms races.

Additional routes through which other individuals in the system might promote ongoing exploration of P-space include Routes 4, 5 and 8. Routes 4 and 5 are processes whereby the production of a phenotype from a genotype is affected by the local context: Route 4 represents the local abiotic environment and Route 5 the local biotic environment. Neighbouring individuals can be involved in Route 4 as well as Route 5, through the processes of environment engineering and niche construction. Route 8 represents the determination of which individuals will mate and contribute to the genotype of a new individual.

The role of all of the processes involving interactions with other individuals (Routes 4–8) in promoting ongoing exploration of P-space 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 indi-

viduals can move (actively or passively) around the environment.

It is also conceivable that the evaluation function E_L is itself defined intrinsically by the individual organisms, although if this were completely determined intrinsically it would surrender any notion of objective fitness and likely render the system no more interesting than a continual production of random individuals.⁵

It should also be recognised 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.

Expansive and Transformational Open-Endedness

Expansive and transformational open-endedness both involve the discovery of *door-opening*⁶ states in P-space that open up an expanded space of new adjacencies.

There are various issues involved in how these might come about in an evolutionary system. The following discussion addresses 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?

(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. (Cariani, 1993; Bird and Layzell, 2002).

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

⁵However, having particular components of the function determined intrinsically might conceivably be useful. The aspect that would likely be most productive to determine intrinsically is the procedure that determines $\overline{p_m}$, the mate set used for reproduction, i.e. changing how an individual selects whom to mate with.

⁶To borrow a term from Bedau (Taylor et al., 2016b).

might be achieved would be to open up the system by allowing it to access additional resources on the internet.⁷

(2) Accessing new states Biology 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 (e.g. mechanical, chemical, electrical, responsiveness to electromagnetism, pressure, etc.). Indeed, the distinction between an expansive and transformational novelty can be viewed as the difference between a *door-opening novelty in the same domain* versus 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 (Gould and Vrba, 1982). 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, e.g. multifunctional enzymes (Kacser and Beeby, 1984).⁸ The latter case can be labelled a *intradomain bridge*.

Another example, provided by Dawkins (1988), is the evolutionary appearance of segmented body plans in animals. 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 behaviour (i.e. expansive and/or transformational novelties). Dawkins describes discoveries of this kind as “watershed events... that open floodgates to future evolution” (Dawkins, 1988, p. 218).⁹

Most computational evolutionary systems lack significantly multi-property components, and therefore miss out on this route to transformational novelty.

⁷This idea has been discussed by Boden (2015) among others. See (Taylor et al., 2016a) for many pointers to how this might be implemented.

⁸The importance of multifunctional components for biological evolvability and robustness has been argued by various authors, e.g. (Goldenfeld and Woese, 2011; Whitacre, 2010).

⁹The distinction between an exploratory discovery of a door-opening state and the potential it introduces for expansive or transformational novelties in function is similar to Wagner (2015)'s distinction between (his conception of) *novelties* and *innovations*.

(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 behaviour 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*.¹⁰ 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 (Banzhaf et al., 2016) 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 degree of approximation, we could view prokaryotic life as an exploration of compositional chemistry and multicellular eukaryotic life as an exploration of compositional physics.¹¹ Furthermore, animals with nervous systems, and ALife agents with evolved controllers, engage in the exploration of compositional information systems.

Note that the ability of a lineage to concurrently explore multiple compositional domains is in itself an enabler or exploratory open-endedness, as it can prevent evolution from getting stuck in a local optima in any one domain by providing an *extradimensional bypass*.¹²

We can distinguish between *additive compositional systems* 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 serial 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. as is the case with the production of a photoreceptor protein such as rhodopsin from its amino acid sequence).

While additive compositional systems result only in exploratory novelty, they can play an important role in en-

abling 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 leading 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 its chemical reaction repertoire: “Once a new molecule appears for the first time in the chemosphere new interactions and further adjacencies emerge” (de Vladar et al., 2017, p. 4).

I close this section with a few remarks about 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 organisations are particularly suitable as a basis for the evolution of complexity (Simon, 1962). Increases in complexity in these cases will be aided by the usual drivers of complexity discussed in the evolutionary biology literature, such as co-evolutionary arms races and evolutionary ratchets.

Final remarks

The framework presented above can act as a guide for categorising and comparing the OE potential of existing systems. For example, von Neumann’s CA implementation of a self-reproducing system concentrated heavily on the role of the laws of dynamics (L) in its intrinsic implementation of $E_L(p, c_a, c_b)$, but ignored the organism’s local context (c_a and c_b), making the system very brittle to perturbations. Tierra 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 (Channon, 2006) 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 (Lehman and Stanley, 2011) 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

¹⁰I use the term *compositional* rather than Boden (2015)’s *combinational* to emphasise that the size of structures may increase, and that the specific arrangement and connections between components might be important.

¹¹This is obviously a gross simplification, as all domains of life utilise both chemistry and physics.

¹²This concept was introduced by Conrad (1990) and later named by Gavrillets (1999).

systems could be improved.

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 categorisation 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.

Acknowledgements

I am grateful to Peter Turney, Alastair Channon, Adam Stanton, James Borg and Bradly Alicea for recent discussions which have helped in developing the ideas set out in this paper.

References

- Banzhaf, W., Baumgaertner, B., Beslon, G., Doursat, R., Foster, J. A., McMullin, B., de Melo, V. V., Miconi, T., Spector, L., Stepney, S., and White, R. (2016). Defining and simulating open-ended novelty: requirements, guidelines, and challenges. *Theory in Biosciences*, 135(3):131–161.
- Barricelli, N. A. (1957). Symbiogenetic evolution processes realized by artificial methods. *Methodos*, 9(35-36).
- Bird, J. and Layzell, P. (2002). The evolved radio and its implications for modelling the evolution of novel sensors. In *Evolutionary Computation, 2002. CEC'02. Proceedings of the 2002 Congress on*, volume 2, pages 1836–1841. IEEE.
- Boden, M. A. (2004). *The Creative Mind: Myths and Mechanisms*. Routledge, London, 2nd edition. (revised, expanded. 1st ed., 1990).
- Boden, M. A. (2015). Creativity and ALife. *Artificial Life*, 21(3):354–365. doi:10.1162/ARTL.a.00176.
- Cariani, P. (1993). To evolve an ear. epistemological implications of Gordon Pask's electrochemical devices. *Systems Research and Behavioral Science*, 10(3):19–33.
- 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.
- Conrad, M. (1990). The geometry of evolution. *BioSystems*, 24:61–81.
- Conrad, M. and Pattee, H. (1970). Evolution experiments with an artificial ecosystem. *Journal of Theoretical Biology*, 28:393–409.
- Dawkins, R. (1988). The evolution of evolvability. In Langton, C., editor, *Artificial Life*, volume VI of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 201–220, Reading, MA. Addison-Wesley.
- de Vlarar, H. P., Santos, M., and Szathmary, E. (2017). Grand views of evolution. *Trends in Ecology & Evolution*, 32(5):324–334. doi:10.1016/j.tree.2017.01.008.
- Dennett, D. C. (2013). *Intuition Pumps and Other Tools for Thinking*. WW Norton & Company.
- Gavrilets, S. (1999). A dynamical theory of speciation on holey adaptive landscapes. *The American Naturalist*, 154(1):1–22.
- Goldenfeld, N. and Woese, C. (2011). Life is physics: evolution as a collective phenomenon far from equilibrium. *Annu. Rev. Condens. Matter Phys.*, 2(1):375–399.
- Gould, S. J. and Vrba, E. S. (1982). Exaptation—a missing term in the science of form. *Paleobiology*, 8(1):4–15.
- Kacser, H. and Beeby, R. (1984). Evolution of catalytic proteins. *Journal of Molecular Evolution*, 20:38–51.
- Lehman, J. and Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223.
- Longo, G., Montevil, M., and Kauffman, S. (2012). No entailing laws, but enablement in the evolution of the biosphere. In *GECCO Companion '12: Proceedings of the Fourteenth International Conference on Genetic and Evolutionary Computation*, pages 1379–1392. ACM.
- Maynard Smith, J. and Szathmary, E. (1995). *The Major Transitions in Evolution*. W.H. Freeman, Oxford.
- Ofria, C. and Wilke, C. O. (2004). Avida: A Software Platform for Research in Computational Evolutionary Biology. *Artificial Life*, 10(2):191–229.
- Packard, N. H. (1988). Intrinsic adaptation in a simple model for evolution. In Langton, C., editor, *Artificial Life*, volume VI of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 141–155, Reading, MA. Addison-Wesley.
- Ray, T. S. (1991). An approach to the synthesis of life. In Langton, C. G., Taylor, C., Farmer, D., and Rasmussen, S., editors, *Artificial Life II*, pages 371–408, Boston, MA. Addison-Wesley.
- Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6):467–482.
- Taylor, T. (2001). Creativity in evolution: Individuals, interactions and environments. In Bentley, P. J. and Corne, D. W., editors, *Creative Evolutionary Systems*, chapter 1, pages 79–108. Morgan Kaufman.
- Taylor, T., Auerbach, J. E., Bongard, J., Clune, J., Hickinbotham, S., Ofria, C., Oka, M., Risi, S., Stanley, K. O., and Yosinski, J. (2016a). WebAL comes of age: A review of the first 21 years of artificial life on the web. *Artificial Life*, 22(3):364–407. (doi:10.1162/ARTL.a.00211).
- Taylor, T., Bedau, M., Channon, A., Ackley, D., Banzhaf, W., Beslon, G., Dolson, E., Froese, T., Hickinbotham, S., Ikegami, T., et al. (2016b). Open-ended evolution: Perspectives from the OEE workshop in York. *Artificial Life*, 22(3):408–423. (doi:10.1162/ARTL.a.00210).
- Wagner, G. P. (2015). Evolutionary innovations and novelties: Let us get down to business! *Zoologischer Anzeiger*, 256:75–81.

Whitacre, J. M. (2010). Degeneracy: a link between evolvability, robustness and complexity in biological systems. *Theoretical Biology and Medical Modelling*, 7(6):1–17.