A procedure for testing for Tokyo Type 1 Open-Ended Evolution (draft edition)

Alastair Channon
School of Computing and Mathematics
Keele University
a.d.channon@keele.ac.uk

Introduction

While there are a range of opinions about what underlying mechanisms may be necessary or sufficient for a system to exhibit Open-Ended Evolution (OEE), there is something close to a consensus about the observable behavioral hallmarks of OEE. The report from the first workshop on OEE summarized these in the *York categories of OEE* (Taylor et al., 2016; Packard et al., 2019):

York 1: Ongoing generation of adaptive novelty:
- **(a)** Ongoing generation of new adaptations
- **(b)** Ongoing generation of new kinds of entities
- **(c)** Emergence of major transitions
- **(d)** Evolution of evolvability

York 2: Ongoing growth of complexity:
- **(a)** Ongoing growth of entity complexity
- **(b)** Ongoing growth of interaction complexity

Following the second and third workshops on OEE, the editorial introduction to the Open-Ended Evolution II Special Issue (Packard et al., 2019) presented the *Tokyo categories of OEE*, in which York types 1a, 1b, 2a and 2b were merged into Tokyo type 1; York types 1d and 1c became Tokyo types 2 and 3; and a new category (type) was added, with the caveat that "perhaps semantic evolution is a special case of Tokyo type 1 OEE":

Tokyo 1: ongoing generation of Interesting new kinds of entities and interactions
Tokyo 2: ongoing generation of Evolution of evolvability
Tokyo 3: ongoing generation of Major transitions
Tokyo 4: ongoing generation of Semantic evolution

OEE can be studied in Nature (Bedau et al., 1997); in systems with ongoing human intervention (such as the global economy, Internet traffic and systems involving user evaluation or interaction); and in autonomous artificial systems, that is systems with no ongoing human (or other external) intervention. In the context of autonomous artificial systems, Tokyo type 1 OEE is considered (by this author at least) as a necessary foundation for Tokyo types 2, 3 and 4 OEE. It is difficult to conceive of an autonomous artificial system achieving Tokyo type 2, 3 or 4 OEE without first achieving Tokyo type 1 OEE. So the construction *and testing* of autonomous artificial systems to achieve Tokyo type 1 OEE is the most immediate challenge for research into the open-ended evolution of autonomous artificial systems.

This paper brings together five methods of analysis to form a *procedure for testing for Tokyo type 1 OEE*:

1. Evolutionary activity statistics (Bedau et al., 1997);
2. Component-normalised evolutionary activity statistics (Channon, 2003, 2006);
3. Long-term evolutionary dynamics classification (Bedau et al., 1998);
4. Analysis of Indefinite Scalability in Diversity and Complexity (Ackley and Small, 2014; Channon, 2019);
5. Analysis of the Order of Indefinite Scalability.

The procedure involves proceeding through these five steps sequentially. It is presented here as simply as possible, isolated from the complexities of any particular evolutionary system, and with a clear rationale for each step. The title of this extended abstract includes "draft edition" as the intention is to encourage discussion of this at The Fourth Workshop on Open-Ended Evolution (OEE4) and to then refine the procedure and this paper for publication as a "first edition". This continued use of the term "edition" is intended to convey that the procedure will be developed furthr in the years to come, with the construction and evaluation of evolutionary systems aiding in that development.

Step 1: Compute basic evolutionary activity statistics

At the core of OEE is the ongoing evolution of *adaptive* novelty (York type 1 OEE): "new components flowing into the system and proving their adaptive value through their persistent activity" (Bedau et al., 1998) (components could be, for example, genes, organisms or species). However, an evolutionary process could continue to generate adaptive novelty but lose what had previously been evolved at the same or a faster rate, cycling or idling with a limited extent of adaptive success. Ongoing adaptive novelty alone would provide for a poor definition of OEE, for a trivial system could generate evermore novel components. Ongoing progress, an unbounded *accumulation of adaptive success*, is also core to OEE. Evolutionary activity statistics provide a measure of exactly that: "a measure of the continual adaptive success of the components in the system" (Bedau et al., 1998), based on *adaptive persistence*.

$$\Delta_i(t) = \begin{cases} 1 & \text{if component } i \text{ exists at } t \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

$$a_i(t) = \begin{cases} \sum_{\tau=0}^{t} \Delta_i(\tau) & \text{if component } i \text{ exists at } t \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

A component's *cumulative evolutionary activity* $a_i(t)$ is a measure of the accumulation of its adaptive success. Specifically, it is the length of time that the component has existed, discounting any periods of absence (equation 2). The sum of component activities (for those components present, in use) is a measure of the system's accumulation of adaptive success, termed *total cumulative evolutionary activity* ($A_{\text{cum}}(t)$, equation 3). That is all there is to know to compute these basic evolutionary activity statistics! They can be computed for any evolving system with an available record of its components' existence times, so are widely applicable across artificial and natural systems.

$$A_{\text{cum}}(t) = \sum_i a_i(t) \tag{3}$$

In step 1, these evolutionary activity statistics are computed and a quick check performed to see whether or not total cumulative evolutionary activity is bounded. If it is, there is no potential for a classification of unbounded evolutionary activity in step 3 below and so the procedure ends. Only the biosphere (Bedau et al., 1997, 1998) and a very small number of autonomous artificial systems have demonstrated unbounded total cumulative evolutionary activity (Taylor et al., 2016) (even before normalization using a shadow model). Developing and demonstrating more autonomous systems that exhibit unbounded total cumulative evolutionary activity (and that have good prospects for success at subsequent steps) is a clear priority. Fortunately, the effort required to evaluate a system to this extent is low, given that no random-selection shadow model is required for this first step.

Step 2: Compute component-normalised evolutionary activity statistics

Following success at step 1, step 2 involves first the implementation of a "shadow" model (population and system) that is identical to the real evolutionary system (running in parallel) except that whenever selection operates in the real system, random selection should be employed in the shadow; and, second, resetting the shadow system's components and evolutionary activity history to those of the real (running) system immediately after each snapshot (when an entry is made in the component existence record). Together these provide the data required for the computation of component-normalised evolutionary activity statistics (Channon, 2003, 2006). Specifically, they enable comparison of inter-snapshot changes in activity in the real run with changes we would expect from neutral (random) selection.

$$\Delta_i^{\text{R}}(t) = \begin{cases} 1 & \text{if component } i \text{ exists in the real run at } t \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

$$\Delta_i^{\text{S}}(t) = \begin{cases} 1 & \text{if component } i \text{ exists in the shadow at } t \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

$$\Delta_i^{\text{N}}(t) = \Delta_i^{\text{R}}(t) - \Delta_i^{\text{S}}(t) \tag{6}$$

$$a_i^{\text{N}}(t) = \begin{cases} \sum_{\tau=0}^{t} \Delta_i^{\text{N}}(\tau) & \text{if component } i \text{ exists in the real run at } t \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

The shadow is used to normalize (exclude non-adaptive) evolutionary activity at the component level ("component activity normalization", equations 4-6), giving a measure of each component's *adaptive* evolutionary activity $a_i^{\text{N}}(t)$ (equation 7) and so also component-normalized (adaptive) measures of total, mean and median cumulative evolutionary activity (equations 8-10).

$$A_{\text{cum}}^{\text{N}}(t) = \sum_{i:\,\substack{\text{component i exists}\\ \text{in the real run at t}}} a_i^{\text{N}}(t) \tag{8}$$

$$\bar{A}_{\text{cum}}^{\text{N}}(t) = \frac{A_{\text{cum}}^{\text{N}}(t)}{D^{\text{R}}(t)} \tag{9}$$

$$\widetilde{A}_{\text{cum}}^{\text{N}}(t) = \underset{i:\,\substack{\text{component i exists}\\ \text{in the real run at t}}}{\text{Median}} a_i^{\text{N}}(t) \tag{10}$$

Ongoing adaptive novelty is determined through *new activity* $A_{\text{new}}^{\text{N}}(t)$ (equation 12): the sum of *newly adaptively-significant* components' activities, divided by component diversity (the number of components present, in use in the real run). A component is considered adaptively significant if its activity is above a threshold, taken to be the absolute value of the most negative component-normalised evolutionary activity so as to screen out non-adaptive activity (Channon, 2006).

$$D^{\text{R}}(t) = \#\{i : a_i(t) > 0\} \tag{11}$$

$$A^{\mathrm{N}}_{\mathrm{new}}(t) = \frac{1}{D^{\mathrm{R}}(t)} \sum_{i:\text{component i `new'}} a^{\mathrm{N}}_i(t) \qquad (12)$$

Stout and Spector (Stout and Spector, 2005) attempted to "break" the original (Bedau et al., 1998) and enhanced (Channon, 2003, 2006) classification schemes (step 3 below) by achieving a classification of unbounded dynamics in "intuitively unlifelike" systems. They concluded that component activity normalization is "of particular importance to the scheme's robustness ... canceling out the potential for spurious results arising from the (random) divergence of the real and shadow populations". Bedau et al.'s reasoning that "the mere fact that a family appears in the fossil record is good evidence that its persistence reflects its adaptive significance" (Bedau et al., 1998) (as "[s]ignificantly maladaptive taxonomic families would likely go extinct before leaving a trace in the fossil record" (Bedau et al., 1998)) is generally accepted. But for artificial systems, Stout and Spector's findings support the argument for employing component activity normalization, at least for cases (choices of component class) in which components can be maladaptive.

Step 3: Long-term evolutionary dynamics classification

After determining long-term trends in component-normalised evolutionary activity statistics, including *new activity* and total, mean and median *cumulative evolutionary activity*, the system's long-term evolutionary dynamics can be classified. The hallmark of unbounded evolutionary dynamics is ongoing positive new evolutionary activity $A^{\mathrm{N}}_{\mathrm{new}}(t)$ in combination with unbounded total ($A^{\mathrm{N}}_{\mathrm{cum}}(t)$) and median ($\widetilde{A}^{\mathrm{N}}_{\mathrm{cum}}(t)$) cumulative evolutionary activity (Bedau et al., 1998; Channon, 2006).

A classification of unbounded evolutionary dynamics, using component-normalised evolutionary activity statistics, provides a test for York type 1a OEE (without ruling out other types). Only one autonomous artificial system has demonstrated this (Taylor et al., 2016). Developing more autonomous systems that do is a further clear priority; any such system would be a significant new contribution to the field.

Further, the one autonomous artificial system we have at this level (Channon, 2006) lacks behavioral transparency, preventing the direct observation of artifacts and behaviors far beyond the early stages of evolution. A very significant advance would be made by the development of an autonomous artificial system that demonstrates unbounded evolutionary dynamics (using component-normalised evolutionary activity statistics) *and* in which long (evolutionary) sequences of evolved artifacts or behaviors and the evolution of more complex artifacts and behaviors can be *seen*, evidenced by phenotypes rather than by metrics. This highlights the need to develop future systems such that behavioral descriptions are as easy to generate as possible, for example by constructing systems such that behaviors will be transparent to human observers.

Step 4: Analysis of Indefinite Scalability in Diversity and Complexity

If an evolutionary system exhibits unbounded evolutionary dynamics in step 3, it would be natural to want to know whether or not the system also exhibits ongoing growth in maximum individual (or group or system) complexity, i.e. York type 2 OEE.

The diversity of components (the number of different components) in an individual (or group or species) is one measure of its complexity. This is particularly appropriate when a component is analogous to a gene. Schad, Tompa and Hegyi (Schad et al., 2011) demonstrated that organism complexity correlates significantly with gene number in the absence of plant genomes, and Chen, Bush et al. (Chen et al., 2014) reached the same finding.

In Bedau, Snyder and Packard's classification of long-term evolutionary dynamics (Bedau et al., 1998), the class of systems with unbounded evolutionary dynamics is divided into three subclasses: (a) those with unbounded diversity of adaptive components but bounded adaptive success (cumulative evolutionary activity) per component; (b) those with bounded diversity but unbounded adaptive success per component; and (c) those with unbounded diversity and unbounded adaptive success per component. Where complexity is measured as diversity of components, ongoing growth in system complexity is equivalent to unbounded diversity, so these subclasses relate directly to ongoing growth in system complexity, i.e. York type 2 OEE, and subclass c is implied by ongoing growth in any of individual, group or system complexity when observed together with unbounded adaptive success per component.

While adaptive success per component can be truly unbounded (over unbounded time), the diversity of adaptive components is necessarily bounded: in artificial systems by unavoidable physical limits such as computer memory, and in nature again by physical limits such as number of atoms. A claim of unbounded diversity in the biosphere is really a claim that diversity is not practically bounded, or that it has not reached the upper bound yet. A more precise notion than "unbounded" diversity is needed. Ackley's concept of *indefinite scalability* (Ackley and Small, 2014) provides this. The key criteria for indefinite scalability is that should an upper bound be reached, increasing the values of physical limitations (such as available matter, population size or memory) should enable an unbounded sequence of greater upper bounds to be achieved (after sufficient increases in the limitations); in the case of diversity this means an unbounded sequence of greater upper bounds on diversity.

A practical (and the most literal) interpretation of indefinite scalability is that the sequence of greater upper bounds (on increasing the values of physical limitations) continues to an unknown length, i.e. that no end to it has been been found. It is therefore best to qualify any empirical claims by quantifying the extent to which indefinite scalability has been established. Claims about systems can be expressed and evaluated in terms such as a metric (for example a measure of adaptive success per component) increasing apparently without bound *up to* a certain system time (or number of generations, etc.); or a metric (for example diversity) increasing *up to* certain value(s) of system parameter(s) being reached, where it was necessary to increase these to establish increases in scale (for example of diversity) over successive runs.

So, in step 4 the aim is to demonstrate a sequence of greater upper bounds on diversity (on increasing the values of physical limitations) that increases without any known bound, qualifying the extent (for example number of generations or values of physical limitations) to which this has been established. This step (step 4) is likely to be developed further in future years, as to date only one system (Channon, 2019) has been evaluated at this level.

Step 5: Analysis of the Order of Indefinite Scalability

While that system exhibited indefinite scalability in diversity and complexity, the analysis in (Channon, 2019) also revealed that complexity scaled logarithmically with (the lower of) maximum population size and maximum number of neurons per individual. It was noted that the evolution of artifacts and behaviors of much greater complexity, for example comparable to those in nature, within feasible timescales, will almost certainly require a higher order of complexity scaling. Achieving a higher order of complexity growth (within a system exhibiting indefinite scalability in complexity) can be considered a grand challenge for Tokyo Type 1 Open-Ended Evolution. Promising approaches to this grand challenge include also achieving one of more of the other Tokyo Types of OEE. Indeed, this can be seen as one answer to *why* these other types of OEE are important, providing a unified view of Open-Ended Evolution.

References

Ackley, D. H. and Small, T. R. (2014). Indefinitely scalable computing = artificial life engineering. In Sayama, H., Rieffel, J., Risi, S., Doursat, R., and Lipson, H., editors, *Artificial Life XIV*, pages 606–613, Cambridge, MA. MIT Press.

Bedau, M. A., Snyder, E., Brown, C. T., and Packard, N. H. (1997). A comparison of evolutionary activity in artificial evolving systems and in the biosphere. In *Proceedings of the Fourth European Conference on Artificial Life*, pages 125–134. MIT Press.

Bedau, M. A., Snyder, E., and Packard, N. H. (1998). A classification of long-term evolutionary dynamics. In Adami, C., Belew, R. K., Kitano, H., and Taylor, C. E., editors, *Artificial Life VI*, pages 228–237, Cambridge, MA. MIT Press.

Channon, A. (2003). Improving and still passing the ALife test: Component-normalised activity statistics classify evolution in Geb as unbounded. In Standish, R. K., Bedau, M. A., and Abbass, H. A., editors, *Artificial Life VIII*, pages 173–181, Cambridge, MA. MIT Press.

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.

Channon, A. (2019). Maximum individual complexity is indefinitely scalable in Geb. *Artificial Life*, 25(2):134–144.

Chen, L., Bush, S. J., Tovar-Corona, J. M., Castillo-Morales, A., and Urrutia, A. O. (2014). Correcting for differential transcript coverage reveals a strong relationship between alternative splicing and organism complexity. *Molecular Biology and Evolution*, 31(6):1402–1413.

Packard, N., Bedau, M. A., Channon, A., Ikegami, T., Rasmussen, S., Stanley, K. O., and Taylor, T. (2019). An overview of open-ended evolution: Editorial introduction to the open-ended evolution ii special issue. *Artificial Life*, 25(2):93–103.

Schad, E., Tompa, P., and Hegyi, H. (2011). The relationship between proteome size, structural disorder and organism complexity. *Genome Biology*, 12(12):R120.

Stout, A. and Spector, L. (2005). Validation of evolutionary activity metrics for long-term evolutionary dynamics. In Beyer, H.-G., O'Reilly, U.-M., Arnold, D. V., Banzhaf, W., Blum, C., Bonabeau, E. W., Cantu-Paz, E., Dasgupta, D., Deb, K., Foster, J. A., de Jong, E. D., Lipson, H., Llora, X., Mancoridis, S., Pelikan, M., Raidl, G. R., Soule, T., Tyrrell, A. M., Watson, J.-P., and Zitzler, E., editors, *Proceedings of the 2005 Genetic and Evolutionary Computation Conference (GECCO 2005)*, pages 137–142, New York. ACM.

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., and Wiser, M. (2016). Open-ended evolution: Perspectives from the OEE workshop in York. *Artificial Life*, 22(3):408–423.