

Maximum Individual Complexity is Indefinitely Scalable in Geb

Alastair Channon

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

Abstract Geb was the first artificial life system to be classified as exhibiting open-ended evolutionary dynamics according to Bedau and Packard's evolutionary activity measures and is the only one to have been classified as such according to the enhanced version of that classification scheme. Its evolution is driven by biotic selection, that is (approximately), by natural selection rather than artificial selection. Whether or not Geb can generate an indefinite increase in maximum individual complexity is evaluated here by scaling two parameters: world length (which bounds population size) and the maximum number of neurons per individual. Maximum individual complexity is found to be asymptotically bounded when scaling either parameter alone. However, maximum individual complexity is found to be indefinitely scalable, to the extent evaluated so far (with run times in years and billions of reproductions per run), when scaling both world length and the maximum number of neurons per individual together. Further, maximum individual complexity is shown to scale logarithmically with (the lower of) maximum population size and maximum number of neurons per individual. This raises interesting questions and lines of thought about the feasibility of achieving complex results within open-ended evolutionary systems and how to improve on this order of complexity growth.

Keywords

Open-ended evolution, biotic selection, ongoing growth of complexity, diversity, indefinite scalability

1 Introduction

Perhaps the most important outcome of the First Workshop on Open-Ended Evolution (OEE) was the distillation of previously disparate thoughts about OEE into something close to a consensus about what the term OEE should mean: “Loosely defined, an open-ended evolutionary system is one that is capable of producing a continual stream of novel organisms rather than settling on some quasi-stable state beyond which nothing fundamentally new occurs. Some definitions of OEE further require that the maximum complexity of organisms in the system increases over time, or that ecosystem complexity increases” [27, p. 409]. Crucial in arriving at and refining this loose definition was recognition of “the importance of distinguishing observable *behavioral hallmarks* of systems undergoing OEE from hypothesized *underlying mechanisms* that explain why a system exhibits those hallmarks” [27, p. 407]. While there are a range of opinions about what underlying mechanisms may be necessary or sufficient for a system to exhibit OEE, there is something close to a consensus about the observable behavioral hallmarks of OEE: “*ongoing adaptive novelty*” [p. 415] (for example, new adaptations, new kinds of entities, major transitions, or the evolution of evolvability) and (some definitions of OEE further require) “*ongoing growth of complexity*” [p. 416] of the most complex entities in the evolving population, or of interactions among entities [27].

2 Ongoing Adaptive Novelty and the Accumulation of Adaptive Success

At the core of open-ended evolution is the ongoing evolution of *adaptive* novelty: “new components flowing into the system and proving their adaptive value through their persistent activity” [5, p. 230] (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 ever more novel components. Ongoing progress, an unbounded accumulation in adaptive success, is also core to OEE.

2.1 Unbounded Accumulation of Adaptive Success

Logically, unbounded accumulation of adaptive success can occur in an evolutionary system through an unbounded increase in either (minimum, mean, median, or maximum) adaptive success per component or the *diversity* of adaptive components, or both. In a system of evolving entities (organisms, creatures, agents, etc.) and considering lower-level components (for example genes), an increase in the diversity of adaptive components (i.e., the number of different adaptive components) can occur through an increase in either the number of different adaptive components per entity (a simple measure of entity *complexity*) or the number of different entities (i.e., the diversity of entities), or both.

Bedau, Snyder, and Packard’s classification of long-term evolutionary dynamics [5] provided the field’s first “objective, quantitative test of success” [5, p. 236]: a means to distinguish those systems that exhibit unbounded evolutionary dynamics according to this classification scheme from those that do not, rather than a definitive definition of open-ended evolution. The classification scheme is based on elegantly simple evolutionary activity statistics that can be computed for any evolving system with an available record of its components’ existence times, making the scheme widely applicable across artificial and natural systems. It uses cumulative *evolutionary activity*, based on *adaptive persistence* (specifically, the length of time that a component has existed, discounting any periods of absence), as “a measure of the continual adaptive success of the components in the system” [5, p. 230], that is, as a measure of the accumulation of a component’s adaptive success; and it uses the sum of component activities (for those components present that are in use) as a measure of the system’s accumulation of adaptive success, termed *total cumulative evolutionary activity*. Ongoing adaptive novelty is determined through *new activity*: the sum of *newly adaptively significant* components’ activities, divided by the component diversity (the number of components present that are in use). A component is considered adaptively significant if its activity is above a threshold that screens out most non-adaptive activity, as determined through the use of a *shadow* system that mirrors the real system in every detail except that where selection (artificial or natural) operates in the real system, neutral (random) selection is employed in the shadow. In line with the logic above, total cumulative evolutionary activity can be considered as the product of component diversity and mean evolutionary activity per component. The classification scheme requires ongoing adaptive novelty and unbounded total cumulative evolutionary activity (unbounded component diversity or unbounded mean evolutionary activity per component) for a classification of unbounded evolutionary dynamics.

Earth’s biosphere was classified, through fossil record data on taxonomic families, as exhibiting open-ended evolutionary dynamics according to Bedau, Snyder, and Packard’s classification scheme [5]. Bedau et al. reasoned that it was not necessary to include a shadow mechanism in this analysis; they considered “this normalization to be accomplished de facto by the fossil record itself,” arguing that “the mere fact that a *family* appears in the fossil record is good evidence that its persistence reflects its adaptive significance” (emphasis added), as “[s]ignificantly maladaptive taxonomic families would likely go extinct before leaving a trace in the fossil record” [5, p. 229].

2.2 Component Activity Normalization and an Enhanced Classification Scheme for Unbounded Evolutionary Dynamics

Channon [8, 9] presented improvements to Bedau et al.'s classification scheme. Resetting the system's shadow (including components and activity history) to be identical to the real run immediately after each snapshot (when an entry is made in the component existence record) allows us to compare inter-snapshot changes in activity in the real run with changes we would expect from neutral (random) selection, the result being an improved generic shadowing mechanism. The shadow can then be used to normalize (exclude non-adaptive) evolutionary activity at the component level ("component activity normalization" [p. 253]), giving a measure of each component's *adaptive* evolutionary activity and so also component-normalized (adaptive) measures of both ongoing adaptive novelty and total (and mean or, better, median) cumulative evolutionary activity.

Stout and Spector [26] attempted to "break" the original and enhanced classification schemes by achieving a classification of unbounded dynamics in "intuitively unlikelike" [p. 137] 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" [p. 141]. 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" [5, p. 229] is generally accepted. But for artificial systems, Stout and Spector's findings support the argument for employing the enhanced classification scheme, at least for cases (choices of component class) in which components can be maladaptive.

Geb is a two-dimensional environment populated with agents that move around and interact with each other. Its evolution is driven by biotic selection, with no (or negligible) abiotic selection [11, 12]. It was designed following the principles of Harvey's Species Adaptation Genetic Algorithms (SAGA) framework for incremental evolution [15–18], but with coevolutionary feedback arising (rather than being specified: cf. 'get cube' [24]) via biotic selection rather than abiotic fitness functions—that is (approximately), via natural selection rather than artificial selection. In *Geb*, selection results from interactions that are activated by the agents' genetically specified neurocontrollers. So selection varies as the population and environment of individuals evolve. This allows for the possibility of feedback in selection, and of that driving the ongoing evolution of adaptive novelty. Novel adaptations reported include behaviors such as following, fighting, fleeing, mimicking, and novel artifacts such as matching input and output channels in agents' neurocontrollers. *Geb* was the first artificial life system to be classified as exhibiting open-ended evolutionary dynamics according to Bedau and Packard's evolutionary activity measures [5, 7] and is the only one to have been classified as such according to the enhanced classification scheme [8, 9, 27]. It was shown to exhibit ongoing adaptive novelty (positive component-normalized new activity per component) and unbounded component-normalized median component evolutionary activity.

3 Ongoing Growth of Complexity

One of the most interesting questions that OEE systems can address is whether or not OEE can be the cause of an unbounded increase in maximum individual (or group or system) complexity. This, of course, requires a definition of complexity, and also that our definition of OEE not already include ongoing growth of complexity [27], which would prevent us from being able to address (or even ask) this important question.

This is closely related to the open question: Can or under what conditions does biotic (natural) selection lead to a sustained increase in maximum entity complexity? In particular this is relevant to the debate as to whether or not it is natural selection that led to the increase in maximum organism complexity observed in nature. Within both biology and artificial life, doubts have been raised as to natural selection's role as the drive for increasing complexity, and arguments have been made on both sides [4, 10, 20, 21, 28, 29], with the suggestion put forward that non-adaptive evolutionary

forces (such as mutation, recombination, and genetic drift) or mathematical/statistical constraints may be the primary drives, through either a passive increase in the variance of complexity in the presence of a lower bound, or a constraint-driven drive toward complexity.

One unsatisfactory general measure of complexity is the number of components in an entity. A more satisfactory one is the number of different components, sometimes referred to as the diversity of components. The number of different components is still not a very satisfactory measure of complexity, just as it is not a very satisfactory measure of diversity, but it does lead us toward the general idea that complexity at one level of analysis (for example, individual, species, or system) can be considered as the diversity of components at the level(s) below (for example, genes, genes or individuals, genes or individuals or species). The same desirable tweaks to discount redundancy (such as counting only adaptive components or measuring information content), and to include behaviors and interactions as well as artifacts, apply to both. Also, considering diversity at only some component levels (for example, just at a very low level: diversity of atoms in a biological organism, logic gates in a computer's processors, neuron types in a neurocontroller, etc.) risks missing other diversity in a system and correspondingly underestimating the system's complexity; again, both diversity and complexity suffer from the same problem.

Within biology, it has long been understood that eukaryotic genome size (length) does not correlate with organism complexity [22], but assumed that the number of distinct genes (i.e., the diversity of genes) that an organism makes use of is a valid measure of its complexity [6]. This assumption was called into doubt [14] following the first complete sequencing and analysis of plant and human genomes. However, subsequently Schad, Tompa, and Hegyi [23] demonstrated that organism complexity correlates significantly with gene number (and more closely with proteome information content) in the absence of plant genomes. More recently still, Chen, Bush, et al. [13] reached the same finding. They also found that specific protein-protein interaction and alternative splicing indices were better predictors; these have no analogue in Geb.

In line with the logic above (Section 2.1), in Bedau, Snyder, and Packard's classification of long-term evolutionary dynamics [5], the class of systems with unbounded evolutionary dynamics can be divided into three subclasses: (a) those with unbounded diversity of adaptive components but bounded adaptive success (cumulative evolutionary activity, based on adaptive persistence) 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.

Yet, while adaptive success per component can be truly unbounded (if measured based on adaptive persistence and over unbounded time), the diversity of adaptive components (both the number of different components per entity and the diversity of entities) is necessarily bounded: in artificial systems by unavoidable physical limits such as computer memory, and in nature (whether considering the biosphere or the universe) 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 (of entities or of adaptive components per entity) is needed.

3.1 Indefinite Scalability

Ackley's concept of *indefinite scalability*, "defined as supporting open-ended computational growth without requiring substantial re-engineering" [1, p. 606], now enables us to address this. The key criterion 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.

However, it is not possible (in finite system time) to establish that a metric (for example, a measure of adaptive success per component) is truly unbounded. And it is not possible (over a finite number of increases in system parameter(s)) to establish that a metric (for example, a measure of

diversity) is infinitely scalable. Further, an increase in parameter(s) may require a longer system (run) time before a greater scale (higher-value metric) is achieved.

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, that is, that no end to it has 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.

3.2 Indefinitely Scalable Complexity

If the diversity of adaptive components is accepted as a simple measure of system complexity, and the number of different adaptive components per entity as a simple measure of entity complexity, then it follows trivially from the logic above (Section 2.1) first, that indefinitely scalable system complexity can occur through indefinite scalability in either entity complexity or entity diversity, or both; and second, that indefinite scalability in the accumulation of adaptive success can occur through indefinite scalability in either adaptive success per component or system complexity, or both.

The above reasoning follows only for the case of complexity equating to diversity of adaptive components (in the system or per entity). It would be perfectly reasonable to use alternative complexity metrics, and to then ask such questions as under what conditions OEE systems exhibit indefinitely scalability in those measures of entity complexity. Likewise it is reasonable to ask under what conditions OEE systems evolve (increasingly) interesting, surprising (not predictable), or impressive artifacts and behaviors [19]. Complexity metrics include simple counts of the number of bases (genome size), genes, cell types, neurons, synapses, species, and behaviors; and measures of information content, again genetic, cellular, neural (for example [29]), ecological, and behavioral.

4 Evaluating Indefinite Scalability in Geb

This work investigates whether or not the observed maximum complexity of any individual is indefinitely scalable in Geb [7–9], where an individual's complexity is measured as the diversity of components in it. Note that if the diversity of components in an individual is indefinitely scalable, then so is the diversity of components in the system, so the question of which subclass (a, b, or c) Geb is in is also being addressed.

As in previous work analyzing Geb's long-term evolutionary dynamics, a component is, in loose terms, an active gene: a gene involved in the agent's neural development; see [9] for details. So, here, an individual's complexity is measured as the number of different genes involved in its neural development.

Two parameters cause diversity to be bounded in Geb: (1) a limit on the maximum number of neurons an agent can have, and (2) the 2D world's length L , as there can be at most L^2 individuals in the population at any one time. These are the two parameters that are scaled. Results are reported below for world lengths 10, 20 (the original system's value for this parameter), 40, and 80; and with the maximum numbers of neurons per individual set at 20 (the original system's value for that parameter), 40, 80, and 160. These ranges avoid edge effects that arise from smaller values of these parameters. 20 runs were carried out for each combination (value pair) of these parameters, and the average (over 20 runs) maximum individual complexity recorded and graphed using a running average of length 100 to reveal underlying trends.




Figure 1. Results when scaling the maximum number of neurons per individual (at different world lengths).

5 Results and Analysis

The basic plots of maximum individual complexity against time indicate that maximum individual complexity may be asymptotically bounded when scaling just the maximum number of neurons per individual (Figure 1). Likewise, maximum individual complexity appears to be asymptotically bounded when scaling just world length (Figure 2).

Figure 3 gives the first indication that maximum individual complexity may be indefinitely scalable—that is, scalable to the extent evaluated so far (with run times in years and billions of reproductions per run)—when scaling both the maximum number of neurons per individual and the world length together.

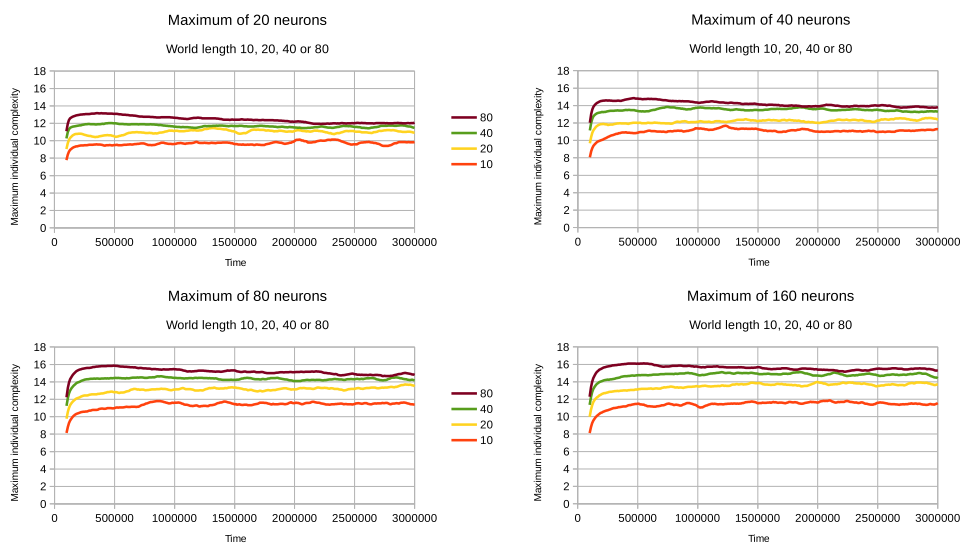


Figure 2. Results when scaling world length (at different maximum numbers of neurons per individual).

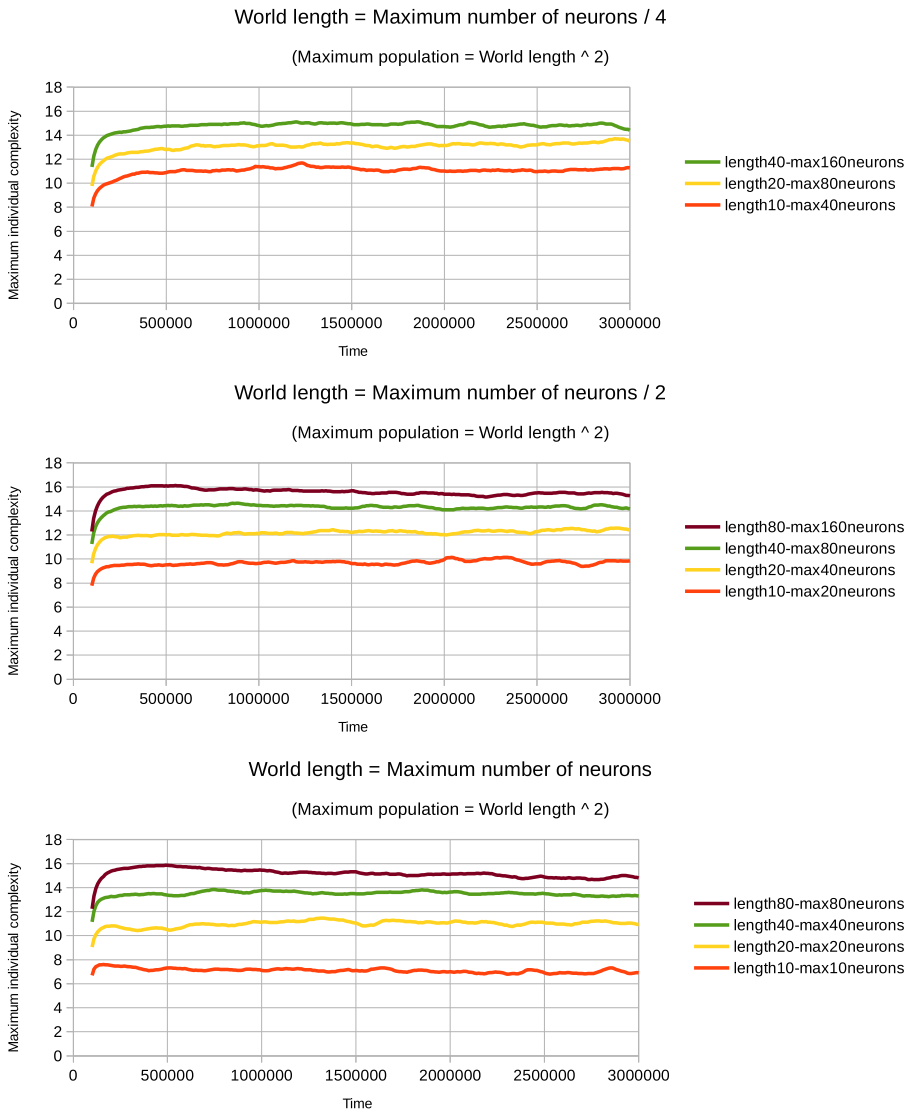


Figure 3. Results when scaling both world length and maximum number of neurons per individual, together.

Figure 4 demonstrates this more conclusively. It shows maximum individual complexity averaged over time steps 2 million to 3 million and over 20 runs, with world lengths scaling in conjunction with the maximum number of neurons per individual as shown. The fitted lines are the result of linear regression on the logarithm of scale (see horizontal axes), with resulting $R^2 = 0.855$, $F_{1,58} = 341$, $P < 10^{-15}$ (top: world length = $20 \times$ scale); $R^2 = 0.794$, $F_{1,78} = 301$, $P < 10^{-15}$ (middle: world length = $20 \times$ scale/2); and $R^2 = 0.741$, $F_{1,58} = 166$, $P < 10^{-15}$ (bottom: world length = $20 \times$ scale/4).

The simplest function of two variables (x and y) that is bounded when each is increased alone but unbounded when both are increased together is $\min(x, y)$: the minimum (lower) of the two values. The function $\text{softmin}(x, y) = -\log(e^{-x} + e^{-y})$ is a smooth approximation to this. It is bounded within $[\min(x, y) - \log(2), \min(x, y)]$.

Figure 5 confirms that maximum individual complexity is asymptotically bounded when scaling just the maximum number of neurons per individual. At each constant world length, the observed




Figure 4. Indefinite scalability of maximum individual complexity (averaged over time steps 2 million to 3 million of 20 runs) when scaling both world length and maximum number of neurons per individual together, through their common dependence on the variable *scale* graphed on the x-axis, as indicated in the legend. Error bars show 95% confidence intervals based on the replicate populations. Shown with linear-fit model and 95% confidence regions for the regression fit.

maximum individual complexities are fitted very closely by the model $a + b \text{softmin}(lw, lm)$, where a and b are the regression parameters, $lw = \log(\text{worldlength})$, and $lm = \log(\text{maxneurons})$. That complexity is bounded when scaling just the maximum number of neurons per individual rules out the possibility that the key finding (that complexity is indefinitely scalable when scaling both the maximum number of neurons per individual and the maximum population size, together) is due to increases in this resource constraint allowing for drift to increasing levels of noise in the neural development process and so to unbounded individual complexity (number of different genes involved in neural development).

Figure 6 confirms that maximum individual complexity is asymptotically bounded when scaling just world length. At each constant maximum number of neurons, the results are fitted very closely by the model $b \text{softmin}(lw + c, lm)$, where b and c are the regression parameters. That complexity is bounded when scaling just maximum population size (world length squared) rules out the possibility that the key finding is due to increasing the sample size over which maximum individual complexity is measured.




Figure 5. Bounded maximum individual complexity (averaged over time steps 2 million to 3 million of 20 runs) when scaling just the maximum number of neurons per individual (at different world lengths). Error bars show 95% confidence intervals based on the replicate populations. Shown with softmin-fit model.




Figure 6. Bounded maximum individual complexity (averaged over time steps 2 million to 3 million of 20 runs) when scaling just world length (at different maximum numbers of neurons per individual). Error bars show 95% confidence intervals based on the replicate populations. Shown with softmin-fit model.

Figure 7 shows that when scaling both the maximum number of neurons per individual and the world length together, the maximum individual complexity is again closely fitted (residual standard error and R^2 equal to those from the linear model above, to at least eleven significant figures) by the model $a + b \text{softmin}(lw, lm)$, which is unbounded. Unsurprisingly, this resembles the linear model very closely: $\text{softmin}(x, y) = x - \log(2)$ when $y = x$, and in the more general case $y = cx$ (for some constant c within $(0, 1]$), $\text{softmin}(x, y)$ is bounded within $[y - \log(2), y)$ and can be empirically shown to be within $y \times (1 - 10^{-5}, 1)$ for $x \geq 20, y \geq 1$ and $c \leq \frac{1}{2}$. Indeed, the linear and softmin regression curves are indistinguishable to the human eye.

6 Conclusions and Discussion

According to this analysis, the only bounds to complexity and diversity are time and computer memory (similarly to nature), and, taking this finding in combination with those from the analysis of evolutionary activity in [9], Geb is in subclass c.




Figure 7. Indefinite scalability of maximum individual complexity (averaged over time steps 2 million to 3 million of 20 runs) when scaling both world length and maximum number of neurons per individual, together. Shown with softmin-fit model. The softmin-fit curves are indistinguishable (to the human eye) from the those of the linear-fit model (Figure 4).

The order of complexity growth has significant implications for the prospects of achieving complex results from an open-ended evolutionary system *within feasible time scales*. The analysis above shows that, in Geb, maximum individual complexity scales logarithmically with (the lower of) maximum population size and maximum number of neurons per individual. This would be sufficient for achieving the evolution of more complex artifacts and behaviors (arising from evolutionary changes rather than from a very small number of mutations from a hard-coded ancestor) than have been *seen* (evidenced by phenotypes rather than by metrics) to date. It would also be sufficient to achieve nontrivial *long (evolutionary) sequences* of evolved artifacts or behaviors. Again, we have not *seen* these yet, evidenced by phenotypes evolved within an ALife system. In terms of these two goals, Geb is lacking in its behavioral transparency, preventing the direct observation of artifacts and behaviors much beyond those noted in Section 2.2, as discussed in [9]. 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.

The evolution of artifacts and behaviors of much greater complexity, for example comparable to those in nature, within feasible time scales, will almost certainly require a higher order of complexity scaling. How to achieve this is an open question. Perhaps the most promising line of thought here relates to establishing the requirements for evolution to itself generate (perhaps an open sequence of) major transitions. Following its very earliest phases, our universe has evolved from a sparse fog of hydrogen and helium atoms. Its history includes the emergence of complex molecules, replicators, single-cell and multicellular life, brains, sociality, users and manufacturers of simple and compound tools, cultural learning, and technology, to highlight just a few of the major transitions [2, Table 1]. Aunger divides big history into four eras: material, biological, cultural, and technological [3, Table 2]. Some evolutionary innovations increase the evolvability (capacity for adaptive evolution) of their lineages [25]. It is not computationally feasible (even if we knew how) for an OEE simulation to start from a sparse fog of hydrogen and helium and transition to a biological-level era, so it is clearly necessary to skip over or engineer in at least some complex features that arose through major transitions in our universe. Geb can be used as an example to inform decisions about engineering in such features, through its demonstration of a feature set that is sufficient for achieving open-ended evolutionary dynamics and indefinitely scalable complexity.

References

1. Ackley, D. H., & Small, T. R. (2014). Indefinitely scalable computing = artificial life engineering. In H. Sayama, J. Rieffel, S. Risi, R. Doursat, & H. Lipson (Eds.), *Artificial life XIV* (pp. 606–613). Cambridge, MA: MIT Press.
2. Aunger, R. (2007). Major transitions in ‘big’ history. *Technological Forecasting and Social Change*, 74(8), 1137–1163.
3. Aunger, R. (2007). A rigorous periodization of ‘big’ history. *Technological Forecasting and Social Change*, 74(8), 1164–1178.
4. Bedau, M. A. (2009). The evolution of complexity. In A. Barberousse, M. Morange, & T. Pradeu (Eds.), *Mapping the future of biology: Evolving concepts and theories* (pp. 111–130). Dordrecht: Springer Netherlands.
5. Bedau, M. A., Snyder, E., & Packard, N. H. (1998). A classification of long-term evolutionary dynamics. In C. Adami, R. K. Belew, H. Kitano, & C. E. Taylor (Eds.), *Artificial life VI* (pp. 228–237). Cambridge, MA: MIT Press.
6. Bird, A. P. (1995). Gene number, noise reduction and biological complexity. *Trends in Genetics*, 11(3), 94–100.
7. Channon, A. (2001). Passing the ALife test: Activity statistics classify evolution in Geb as unbounded. In J. Kelemen & P. Sosik (Eds.), *Advances in artificial life: Proceedings of the Sixth European Conference on Artificial Life* (pp. 417–426). Berlin, Heidelberg: Springer-Verlag.
8. Channon, A. (2003). Improving and still passing the ALife test: Component-normalised activity statistics classify evolution in Geb as unbounded. In R. K. Standish, M. A. Bedau, & H. A. Abbass (Eds.), *Artificial life VIII* (pp. 173–181). Cambridge, MA: MIT Press.

9. 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.
10. Channon, A. (2008). A measure for natural selection's contribution to the origins and maintenance of organismal complexity. In S. Bullock, J. Noble, R. Watson, & M. A. Bedau (Eds.), *Artificial life XI* (p. 758). Cambridge, MA: MIT Press.
11. Channon, A., & Damper, R. I. (1998). Perpetuating evolutionary emergence. In R. Pfeifer, B. Blumberg, J.-A. Meyer, & S. W. Wilson (Eds.), *From animals to animats 5: Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior* (pp. 534–539). Cambridge, MA: MIT Press.
12. Channon, A., & Damper, R. I. (2000). Towards the evolutionary emergence of increasingly complex advantageous behaviours. *International Journal of Systems Science*, 31(7), 843–860.
13. Chen, L., Bush, S. J., Tovar-Corona, J. M., Castillo-Morales, A., & 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.
14. Hahn, M. W., & Wray, G. A. (2002). The g-value paradox. *Evolution & Development*, 4(2), 73–75.
15. Harvey, I. (1992). Species adaptation genetic algorithms: A basis for a continuing SAGA. In F. J. Varela & P. Bourguine (Eds.), *Towards a practice of autonomous systems: Proceedings of the First European Conference on Artificial Life* (pp. 346–354). Cambridge, MA: MIT Press.
16. Harvey, I. (1993). Evolutionary robotics and SAGA: The case for hill crawling and tournament selection. In C. G. Langton (Ed.), *Artificial life III* (pp. 299–326). Reading, MA: Addison-Wesley Longman.
17. Harvey, I. (1997). Artificial evolution for real problems. In T. Gomi (Ed.), *Evolutionary robotics: From intelligent robots to artificial life (ER'97)*, proceedings of the 5th International Symposium on Evolutionary Robotics (pp. 127–149). Ontario, Canada: AAI Books.
18. Harvey, I. (1997). Cognition is not computation; evolution is not optimisation. In W. Gerstner, A. Germond, M. Hasler, & J.-D. Nicoud (Eds.), *Proceedings of the 7th International Conference on Artificial Neural Networks (ICANN'97)* (pp. 685–690). Berlin, Heidelberg: Springer-Verlag.
19. Lehman, J., & Stanley, K. O. (2012). Beyond open-endedness: Quantifying impressiveness. *Artificial Life*, 13, 75–82.
20. Lynch, M. (2007). The frailty of adaptive hypotheses for the origins of organismal complexity. *Proceedings of the National Academy of Sciences of the U.S.A.*, 104, 8597–8604.
21. Miconi, T. (2008). Evolution and complexity: The double-edged sword. *Artificial Life*, 14(3), 325–344.
22. Moore, G. P. (1984). The c-value paradox. *BioScience*, 34(7), 425–429.
23. Schad, E., Tompa, P., & Hegyi, H. (2011). The relationship between proteome size, structural disorder and organism complexity. *Genome Biology*, 12(12), R120.
24. Sims, K. (1994). Evolving 3D morphology and behavior by competition. In R. A. Brooks & P. M. Maes (Eds.), *Artificial life IV* (pp. 28–39). Cambridge, MA: MIT Press.
25. Sterelny, K. (2011). Evolvability reconsidered. In B. Calcott & K. Sterelny (Eds.), *The major transitions in evolution revisited* (pp. 83–100). Cambridge, MA: MIT Press.
26. Stout, A., & Spector, L. (2005). Validation of evolutionary activity metrics for long-term evolutionary dynamics. In H.-G. Beyer, U.-M. O'Reilly, D. V. Arnold, W. Banzhaf, C. Blum, E. W. Bonabeau, E. Cantu-Paz, D. Dasgupta, K. Deb, J. A. Foster, E. D. de Jong, H. Lipson, X. Llorca, S. Mancoridis, M. Pelikan, G. R. Raidl, T. Soule, A. M. Tyrrell, J.-P. Watson, & E. Zitzler (Eds.), *Proceedings of the 2005 Genetic and Evolutionary Computation Conference (GECCO 2005)* (pp. 137–142). New York: ACM.
27. 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., & Wisner, M. (2016). Open-ended evolution: Perspectives from the OEE workshop in York. *Artificial Life*, 22(3), 408–423.
28. Yaeger, L. S. (2009). How evolution guides complexity. *HFSP Journal*, 3(5), 328–339.
29. Yaeger, L. S. (2014). Evolution of complexity and neural topologies. In M. Prokopenko (Ed.), *Guided self-organization: Inception* (pp. 415–454). Berlin, Heidelberg: Springer-Verlag.