



University of the
West of England

Winfield, A. F. (2014) Estimating the energy cost of (artificial) evolution. In: Sayama, H., Rieffel, J., Risi, S., Doursat, Ren and Lipson, H., eds. *Artificial Life 14: 14th International Conference on the Synthesis and Simulation of Living Systems*, New York, 31 July - 2 August 2014., pp. 872-875 Available from: <http://eprints.uwe.ac.uk/23446>

We recommend you cite the published version.

The publisher's URL is:

<http://dx.doi.org/10.7551/978-0-262-32621-6-ch143>

Refereed: Yes

(no note)

Disclaimer

UWE has obtained warranties from all depositors as to their title in the material deposited and as to their right to deposit such material.

UWE makes no representation or warranties of commercial utility, title, or fitness for a particular purpose or any other warranty, express or implied in respect of any material deposited.

UWE makes no representation that the use of the materials will not infringe any patent, copyright, trademark or other property or proprietary rights.

UWE accepts no liability for any infringement of intellectual property rights in any material deposited but will remove such material from public view pending investigation in the event of an allegation of any such infringement.

PLEASE SCROLL DOWN FOR TEXT.

Estimating the Energy Cost of (Artificial) Evolution

Alan F.T. Winfield

Bristol Robotics lab, UWE Bristol, UK
alan.winfield@uwe.ac.uk

Abstract

This short discussion paper sets out to explore the question: what is the energy cost of evolving complex artificial life? The paper takes an unconventional approach by first estimating the energy cost of natural evolution and, in particular, the species *Homo Sapiens Sapiens*. The paper argues that such an estimate has value because it forces us to think about the energy costs of co-evolution, and hence the energy costs of evolving complexity. Furthermore, an analysis of the real energy costs of evolving virtual creatures in a virtual environment, leads the paper to suggest an artificial life equivalent of Kleiber's law – relating neural and synaptic complexity (instead of mass) to computational energy cost (instead of real energy consumption). An underlying motivation for this paper is to counter the view that artificial evolution will facilitate the technological singularity, by arguing that the energy costs are likely to be prohibitively high. The paper concludes by arguing that the huge energy cost is not the only problem. In addition we will require a new approach to artificial evolution in which we construct complex scaffolds of co-evolving artificial creatures and ecosystems.

Introduction

Despite more than 20 years of good progress in evolutionary robotics the most complex robots evolved to date are – if we are honest with ourselves – not very complex.

Yet, within the wider discourse on predicted advances in super intelligent robotic and AI systems leading (perhaps) to a technological singularity (Eden et al., 2012), there is frequently an assumption that artificial evolution will do much of the heavy lifting in their development. For instance Chalmers (2010), in his philosophical analysis of the technological singularity, writes:

If we produce a AI by artificial evolution, it is likely that soon after we will be able to improve the evolutionary algorithm and extend the evolutionary process, leading to AI+.

I believe the assumption that artificial evolution will facilitate the technological singularity (i.e. the development of an Artificial General Intelligence followed by an intelligence explosion) to be mistaken, for several reasons. This

paper focusses on one: the energy cost of evolving complexity. My contention is that this cost is likely to be colossal and well beyond the resources that may be realistically available in the near or medium term future. Of course I am not suggesting that the energy cost of artificially evolving human-equivalent AI (AGI) will be directly comparable to the energy cost of naturally evolving humans from scratch. I do, however, contend that an estimate of the latter is a useful argument in countering the optimism that increasingly appears to characterise opinion in AI researchers (Goertzel et al., 2010).

This paper proceeds as follows. First is an attempt to estimate the upper and lower bounds of the energy cost of human evolution. Then a consideration of the energy costs of artificial evolution, proposing an Artificial Life version of Kleiber's law and speculating on how the energy cost might scale with neural and synaptic complexity. The paper concludes with a short discussion around the evolution of complexity, noting that artificial evolution will need new approaches which combine artificial selection and niche construction.

The energy cost of evolving humans

An upper bound Let us estimate an upper bound by considering the energy delivered by the Sun, neglecting geothermal energy sources since these probably account for a very small proportion of the energy used by natural evolution.

Photosynthesis has been estimated to capture about $3000 EJ$ per year in biomass¹(Myamoto, 1997). It is generally accepted that plants extensively colonised the land during the Devonian period. If we take the mid-point of that period, $\sim 390 MYa$, then the total energy captured by biomass since then amounts to $\sim 1.17 \times 10^{12} EJ$. Field et al. (1998) showed that the sea accounts for about the same primary (photosynthetic) production as the land despite very different physical distribution and producers (marine phytoplankton and land plants), so let us assume $1500 EJ$ per year were captured by the sea's biomass during the long period ($\sim 3B$

¹Total human primary energy use in 2010 was estimated as $539 EJ$.

years) of evolution prior to colonisation of the land. This amounts to $\sim 4.5 \times 10^{12} EJ$. Thus we can estimate the total Solar energy capture by the Earth's biomass since cyanobacteria started photosynthesising as $\sim 5.7 \times 10^{12} EJ$.

This represents an estimate of the total amount of energy *available* to natural evolution, to date. Furthermore, this was the energy available to evolve *all* living things that have ever existed, including humans. Of course not all of that energy was used to power evolution – some of that energy is captured and stored in hydrocarbon deposits – hence this estimate is an upper bound.

A lower bound A different approach will yield a lower bound. By considering humans and working backwards via a series of Last Common Ancestors (LCAs) we can trace an evolutionary path to single celled organisms. Dawkins coins the term *concestors*, and identifies 39 in *The Ancestor's Tale* (Dawkins, 2004). To simplify the energy estimate we choose a subset of concestors, punctuating the stages of evolution from hominids to primates, then mammals, land vertebrates, sea vertebrates, multi-celled animals and finally single-celled animals. For each of these seven stages we need three values: the average daily energy consumption of each individual, the average age of reproduction, and the average number of individuals in our ancestor's population at any given instant. Estimating these values for the more recent stages of evolution is relatively straightforward – especially for daily energy consumption and age of reproduction. For the energy consumption of an organism with a given mass we can refer to Kleiber's law (Dawkins, 2004, p. 422). Population size is much more difficult and here we can do little more than guess: too small a value and there is insufficient genetic diversity, whereas too large doesn't make sense if our ancestor's size and mobility meant it could only access a limited breeding group.

A key assumption of this lower bound estimate is that the species (orders) that branch away from our concestors are not directly implicated in the evolution of humans, because of symbiosis, mutualism or food-chain dependency. In other words we discount the energy cost of the continuing evolution of reptiles (including birds) after concestor 16 – the last common ancestor of mammals and reptiles. Uncontroversially we discount the energy cost of lungfish, from concestor 18, and of the ambulacrarians (including starfish, sea urchins and sea cucumbers) from concestor 25. More controversially we discount the energy cost of the subsequent evolution of insects (which branch from concestor 26) after 570MYa, and plants (which branch from concestor 36) after 900MYa. Humans (and many other ancestral species) have a food-chain dependency on plants and insects, and flowering plants depend on the mutualism of plant and (insect) pollinator. It is for these reasons that the energy estimate here is a lower bound.

Given the estimated (in some cases guessed) values in Ta-

ble 1 and summing the energy costs per epoch we arrive at a lower-bound estimate for the energy cost of evolving humans, of $\sim 8000 EJ$. We shall return to the difficult question of where – between the lower and upper bounds estimated here – the true energy cost of evolving humans might lie.

The energy cost of artificial evolution

The evolution of robots (or AI) with human-equivalent intelligence almost certainly will not require that we recapitulate natural evolution from scratch, either *in materio* or *in silico*. But, from an energy perspective, that doesn't let us off the hook. Evolutionary robotics has, to date, mostly evolved controllers – often based upon simple artificial neural networks – for relatively simple pre-designed robots. The most complex robot controllers evolved to date have perhaps 100 artificial neurons, somewhat less but of the same order as *C. elegans* (nematode roundworm), with 302 neurons and ~ 5000 synapses. Such a controller is typically evolved in a simulated environment and then downloaded into the real robot. The process of artificial evolution may require a population of 100 individuals (genomes), each of which needs to be instantiated as a simulated robot and fitness tested in its simulated environment, for perhaps 1000 generations. If that environment models physics, as well as the robot's sensors and actuators with sufficient fidelity, then a workstation grade PC may complete the task in 10 hours, at a total energy cost of $\sim 9000 KJ$. Setting aside the fact that the robot's artificial neurons are generally very much simpler than *C. elegans* biological neurons, we have an energy cost estimate of evolving an artificial controller for a robot, of roughly comparable neural complexity.

Recent work by Auerbach and Bongard (2014) explores the influence of the environment on the evolution of morphological complexity in virtual machines, in work that is representative of the state-of-the-art in the co-evolution of morphology and control system. This elegant research demonstrates that increasing morphological complexity is actively driven by environmental complexity. This work provides further evidence of the energy cost of artificial evolution; the exploration of artificial organisms of low complexity (compared with biological organisms) evolved for one behaviour only – efficient locomotion across a ridged 'icy' landscape, is reported to have cost 100 CPU-years of computational effort on a supercomputing cluster.

One objection to the approach outlined in this paper is that there is no proper basis for comparison, no equivalence, between biological evolution and artificial evolution. Of course the processes and mechanisms are profoundly different (except for the meta-level equivalence of the Darwinian evolutionary operators: variation, selection and heredity), but there is an ineluctable truth: artificial evolution still has an energy cost. Virtual creatures, evolved in a virtual world, have a real energy cost. And we can estimate that energy cost. For the *C. elegans* equivalent example outlined above

Epoch	Hominid	Primate	Mammal	Land Vertebrate	Sea Vertebrate	Multi-celled organism	Single-celled organism
Concestor no (Dawkins (2004))	1	9	16	18	25	32	–
Indiv. energy cost per day (KJ)	8500	4000	1000	100	100	10	0.0001
Time to reproduction (years)	15	5	0.2	1	1	1	0.003
Life to repro. energy cost (KJ)	46537500	7300000	73000	36500	36500	3650	0.0001
Epoch (M years)	6	20	150	200	200	100	3000
Generations per epoch	400000	4000000	750000000	200000000	200000000	100000000	1.095E+12
Population size per generation	1000	1000	10000	10000	10000	100000	100000000
Energy cost per epoch (EJ)	186.15	292	5475	730	730	365	109

Table 1: Lower bound energy calculation

each simulated robot has a real energy cost of about $9J/hr$, which interestingly is about 2000 times greater than the energy cost of a very small (1mg) organism, $0.004J/hr$. It is clear that ‘larger’ artificial creatures, i.e. with more artificial neurons, must incur a greater computational energy cost.

In general, if the energy cost of simulating and fitness testing a virtual creature is e , then the energy cost of evolving that creature will be $E = gpe$, where g is the number of generations required and p the population size. Energy cost e is clearly a function of the complexity of that virtual creature, but how might e scale with complexity? Kleiber’s law (Dawkins, 2004, p. 422) relates the mass of an organism to its energy consumption and, plotted on logarithmic axes, shows a remarkably consistent linear relationship from micro-organisms to the largest animals. Perhaps a similar relationship might exist between, say, neural complexity and energy cost e for virtual creatures: an artificial life equivalent of Kleiber’s law?

Figure 1 imagines such a plot, of neural complexity against energy cost e . We cannot yet plot such a relationship since we have, to date, only one or two points at the very bottom of the artificial neural complexity scale. But, if we assume that a human-equivalent AI will require roughly comparable neural complexity to *Homo Sapiens*², with 85×10^9 neurons and $10^{14} - 10^{15}$ synapses (and noting that neural complexity must take account of the number of synapses, since neural connections incur a computational energy cost), then e for an artificial creature of this synaptic complexity could be $10^{10} - 10^{12}$ times greater than for something equivalent to *C. elegans*. But this scale factor is still likely to be too low because fitness testing of increasingly complex artificial creatures will take longer and incur greater energy cost. It seems likely that the gradient of our ALife version of Kleiber’s law will be greater than 1.

Discussion

An important consideration is the question of what, exactly, do we mean by the evolution of complexity. Levins and Lewontin (1985) articulate the significant difficulty of

²almost certainly not a safe assumption, but it’s all we have to go on here.

measuring complexity and demonstrating its increase during evolution. Adami et al. (2000) explore the same question by developing an information theoretic approach to biological complexity. We should rightly be wary of any suggestion of a monotonic increase in complexity during evolution. And, returning to the energy cost of evolving humans, much of the structural and morphological complexity of hominids – of vascular and nervous systems, skeletons and sense organs – was established early in our evolutionary history. But then somehow the happy coincidence of dextrous hands with opposable thumbs, acute forward-facing binocular vision and big brains with a neocortex gave rise to a late explosion of phenotypic complexity in the last 6M years, resulting in what Mithen (1996) calls the architecture of the modern mind.

Attempting to estimate the energy cost of evolving humans, and establishing approximate upper and lower bounds on this cost, exposes a deeply interesting question: how much of the Earth’s biota was necessary for the evolution of humans? Biological complexity apparently arises from an evolutionary arms-race in which organisms both adapt to and exploit niches in their ecosystem and – in so doing – co-create that ecosystem. As Levins and Lewontin (1985) point out the organism is both the subject and the object of evolution. Niche construction is the process by which organisms continuously modify their own and others’ niches; from a niche construction perspective “evolution consists of mutual and simultaneous processes of natural selection and niche construction” (Laland et al., 2000). So, the answer to our question ‘how much of the Earth’s biota was necessary for the evolution of humans?’ is something we cannot know, since unpicking the immense tangle of co-evolving species and niches is almost certainly impossible.

But what seems clear is this. To evolve artificial life, or AI, of significantly greater complexity than anything so far achieved will require a new approach to artificial evolution in which we construct complex ‘scaffolds’ of co-evolving artificial organisms and ecosystems. Learning how to sustain artificial scaffolds for long enough to make real progress will be a significant long-term challenge, requiring both ingenuity and energy. From an energetic point of view the Kleiber’s law like relationship between neural and synaptic




Figure 1: The computational energy cost of artificial neural complexity, after Kleiber's Law.

complexity and the energy cost of evolution suggested in the previous section is almost certainly a gross oversimplification. To evolve artificial creatures (or robots) of significant neural complexity will require that we co-evolve multiple 'species' within complete eco-systems in a parallel process of artificial selection and niche construction, in order to promote the evolution of greater levels of complexity and capability. Even if we succeed in understanding how to engineer such scaffolds, the energy costs are likely to be many orders of magnitude greater than Fig. 1 might suggest. Artificial evolution is not the silver bullet that advocates of the technological singularity might suppose it is.

Acknowledgements

The development of these ideas have greatly benefited from discussion with several people including Susan Blackmore, Jolyon Troscianko, Andy Radford and Marina Strinkovsky.

References

Adami, C., Ofria, C., and Collier, T. (2000). Evolution of biological complexity. *Proc. Nat. Acad. Sci. USA*, 97:4463–4468.

Auerbach, J. and Bongard, J. (2014). Environmental influence on the evolution of morphological complexity in machines. *PLoS Computational Biology*, 19 (1): e1003399.

Chalmers, D. (2010). The singularity: A philosophical analysis. *Journal of Consciousness Studies*, 17 (9-10):7–65.

Dawkins, R. (2004). *The Ancestor's Tale*. Weidenfeld and Nicolson, London.

Eden, A., Steinhart, E., Pearce, D., and Moor, J. (2012). Singularity hypotheses: An overview. In Eden, A., Moor, J., Soraker, J., and Steinhart, E., editors, *Singularity Hypotheses: A scientific and philosophical assessment*, pages 1–12. Springer-Verlag, Heidelberg.

Field, C., Behrenfeld, M., Randerson, J., and Falkowski, P. (1998). Primary production of the biosphere: integrating terrestrial and oceanic components. *Science*, 281:1802–1804.

Goertzel, B., Baum, S., and Goertzel, T. (2010). How long till human-level AI? *H+ Magazine*, pages 237–240.

Laland, K., Odling-Smee, J., and Feldman, W. (2000). Niche construction, biological evolution, and cultural change. *Behavioral and Brain Sciences*, 23:131–175.

Levins, R. and Lewontin, R. (1985). *The Dialectical Biologist*. Harvard University Press.

Mithen, S. (1996). *The Prehistory of the Mind*. Thames and Hudson, London.

Myamoto, K. (1997). Chapter 2: Energy conversion by photosynthetic organisms. In *Renewable biological systems for alternative sustainable energy production (FAO Agricultural Services Bulletin - 128)*. Food and Agriculture Organisation of the United Nations.