Editorial Introduction to the ALIFE 2018 Special Issue

Artificial Life Next Generation Perspectives: Echoes from the 2018 Conference in Tokyo

Olaf Witkowski*
Cross Compass Ltd.
Cross Labs
Tokyo Institute of Technology
Earth-Life Science Institute
Institute for Advanced Study
University of Tokyo
olaf@crosslabs.org

Takashi Ikegami
University of Tokyo

Nathaniel Virgo
Tokyo Institute of Technology
Earth-Life Science Institute
Max Planck Institute for
Mathematics in the Sciences

Mizuki Oka
University of Tsukuba

Hiroyuki Iizuka
Hokkaido University

Artificial life is a research field devoted to the theoretical study of features of living systems, such as evolution and the brain. The field has developed philosophical concepts such as autopoiesis and emergence, alongside a large range of computational and experimental setups, from evolutionary simulations to robotics and chemical experiments.

The complexity and diversity of the artificial life field is crucial to its community. Many researchers consider the community a real source of creativity and free-minded exchange of ideas on important questions. For ideas that don't fit neatly into a single "mainstream" field of science, there is value in examining and discussing them in a context free from departmental or disciplinary constraints, with the purpose of reaching a better knowledge of the fundamental mechanisms that govern living systems.

Workers in the field of artificial life convene yearly at a conference, which most recently took place in 2018 in Tokyo. The ALIFE 2018 conference highlighted eight categories of topics:

- Perception, cognition, behavior
- Bio-inspired, cognitive, and evolutionary robotics; swarms
- Ecological and social systems
- Artificial chemistry, origins of life, computational biology
- Complex dynamical systems and networks
- Evolution of language, computational linguistics
- Philosophy of mind, philosophy of science

* Corresponding author.

- Artificial-life–based art
- Synthetic biology and wet artificial life
- Education and society issues

In this issue, we present a selection of recent studies on these topics, which were presented at ALIFE 2018 and which reflect the very rich heterogeneity of the field, and many of which are contributions from young scientists. These articles have been selected from a very diverse pool of submissions. The conference featured 112 papers, selected out of 212 submissions, of which 64 were accepted as oral presentations and 48 as poster presentations.

We focused on top-rated, promising, and impactful studies by young contributors to artificial life research. The articles cover a set of concepts key to artificial life, and ground them in concrete models and outputs to better describe those concepts and make them actionable for the whole field. A simple classification of the current research paradigm in this volume can be described as follows: (i) autopoiesis, (ii) evolution and evolvability, (iii) emergence, and (iv) embodiment.

The first three articles examined autopoiesis and embodiment. Randall Beer [1] investigated the initial-state dependence of the Game of Life to show how the initial distribution generates gliders but also annihilates and regenerates them as time passes. This is a prototypical life evolution example. This work is in the line of embodied autopoiesis, which is reconsidered by studying gliders and their peripheral configurations.

Tamaru, Yui, and Hashida [9] demonstrated the self-organized motion of a pine cone, emphasizing its unexpected locomotion, which realizes a sensory-motor coupling. The motion of gliders and pine cones is different, but has that common property, resulting from environmental conditions (humidity for pine cones, and the bit configuration for gliders).

Masumori et al. [5] discussed the concept of autopoiesis in terms of artificial and biological neural focused on networks. Their work proposes a new learning principle of neural networks called stimulus avoidance. If a network can avoid being stimulated from the outside, it does so by organizing its behavior to avoid the source of stimuli in the environment. If not, the network tends to inactivate its sensory part to prevent the information from the sensor neurons from spreading. So, effectively, the sensory inputs are removed.

In contrast with Randy Beer's work [1], neural networks can self-organizingly change their boundary conditions to acquire robustness. A neural network protects itself from environmental stimuli, whereas gliders are totally fragile against external perturbation.

The next articles examined aspects of evolutionary dynamics.

Dolson et al. [2] proposed a set of metrics for analyzing evolutionary histories in digital evolution experiments, with the hope of boosting the computational evolution community's interest in using such measures to improve our understanding of evolutionary dynamics. The article offers some simple use cases, and points to code available online to compute the measures.

Khajehabdollahi and Witkowski [3], by using restricted Boltzmann machines (RBMs) (installed agents), demonstrated that a critical state of an RBM is not always adaptive. The critical states should be selected through evolutionary dynamics for a particular task.

Liard et al. [4] examined whether complexity is driven primarily by selection or a non-selective force, which is a central question in artificial life. Using the Aevol system for digital organisms, they designed a simple world and a simple genome, and applied directional selection pressure for a protein that can function to produce a triangle. They explored how an information metric for complexity at genomic and phenotypic levels appears to increase at low mutation rates, and as well, explored robustness and evolvability.

Moore et al. [6] gave a tangible example of how artificial life models are a powerful tool to help society, in particular in the biomedical realm, to test hypotheses in great detail. Although methods such as genome-wide association study (GWAS) have allowed identifying thousands of genetic risk factors for hundreds of diseases, the genetic risk factors detected have very small effect sizes and tell little about the overall heritability of diseases. With the help of a heuristic simulation, the authors were able to ask questions about the complexity of genetic architectures, in particular whether common diseases are partly driven by non-additive interactions. The simulations were highly consistent with the magnitude and distribution of univariate effects seen in real data, as is corroborated by a comparison with large-scale studies of sporadic breast cancer.

Veenstra et al. [10] studied how senescence, or intrinsic mortality, influences the evolvability of a population. They used the hierarchical if-and-only-if function as a deceptive fitness landscape, together with a spatial artificial life model. They showed an explicit relationship between mutation and mortality rate in a few different application scenarios. Their results support the premise that programmed death of individuals can have a beneficial effect on the evolvability of the entire population.

The studies by Masumori et al. [5], Khajehabdollahi and Witkowski [3], and Beer [1] all addressed the important theme of embodiment, that is, how an embodied mind "as it could be" might develop a first-person perspective in a different but even more direct way. Tamaru, Yui, and Hashida [9] also addressed how intelligence can be embedded in a direct physical task, where an embodiment computes how to advance.

The general theme of this conference was "Beyond AI." Since around 2010, AI has rapidly become a major trend, and the use of big data has gradually invaded every area of science. AI does not need to visualize big data, nor does it have to solve any equations. Just from the big data itself, AI can extract patterns and forms to predict future outcomes. A missing piece of this new era driven by AI is embodiment and evolutionary perspectives, which are also the main foci of this volume.

The fundamental principles of life are missing from current AI systems, which bear little of the richness of living systems. From the point of view of the ALife community, it is fairly evident that the post-AI era will need to invent something completely new, which can be done by leveraging research topics that are already active in the ALife community.

If one is serious about designing systems that mimic human intelligence, the priority should be put on a theory of the nature of living (if not human) systems, on their emergence, and on their evolution through time. The understanding of such principles can then contribute to an effort to design intelligence, which could even come as an inevitable side effect of homeostasis and evolution in a complex and dynamic environment.

Although this approach may seem like wishful thinking, possibly sprouting from a whole community's hallucination due to its dedication to tackling life as it could be, perhaps the rich interdisciplinarity of the field may manage to protect it from biases. Actually, looking at the major conferences about diverse types of machine learning, AI scientists themselves have generally started looking in the direction of ALife, in a search for creative ideas.

AI research currently sees limitations, suggesting the arrival of a next AI winter, in which AI will pass on its popularity to other sectors of science. The first signs are the absence of new paradigms in machine learning, and the shift from the invention of new algorithmic principles to their use: We are now merely applying all the key inventions at our disposal, such as backpropagation, neural architectures, and powerful computation, to solve application problems in other fields. AI needs a next source of creativity, which will bring about the next step towards artificial general intelligence (AGI) or artificial superintelligence (ASI).

In a sense, ALife has been pushing in the creative direction all along, in particular through its search for open-ended evolution (OEE), which is also the study of the production of novelty in nature. The extraordinary source of inventiveness and diversity seems to be key to the nature of living systems. In 2019, *Artificial Life* journal published two special issues (see [7] and [8] for overviews) that focused on OEE as an essential part of artificial systems as well as real living systems, from many different angles. Those issues and the current one could not be better aligned, for in order to capture the I in AI, one needs to focus on what is at its roots: the fundamental principles of life, including the juice that makes it maintain its creative productivity.

ALife as a field has managed to develop a unique creative ecology for the study of living systems from first principles, by building them piece by piece from the most diverse substrates. Before this proves useful for enabling AI's next steps, ALife will probably undergo changes itself, following its two major conferences, which have now been merged into one since the Tokyo conference. Not only a central concept in this day and age, ALife also is a symbol of transdisciplinary, diverse, and open-ended scientific endeavor, which exemplifies the creativity required to make significant leaps in understanding our universe, and us as cognitive beings living, evolving, and growing in it.

References

1. Beer, R. D. (2020). An investigation into the origin of autopoiesis. *Artificial Life*, *26*(1), 5–22.
2. Dolson, E., Lalejini, A., Jorgensen, S., & Ofria, C. (2020). Interpreting the tape of life: Ancestry-based analyses provide insights and intuition about evolutionary dynamics. *Artificial Life*, *26*(1), 58–79.
3. Khajehabdollahi, S., & Witkowski, O. (2020). Evolution towards criticality in Ising-neural agents. *Artificial Life*, *26*(1), 112–129.
4. Liard, V., Parsons, D. P., Rouzaud-Cornabas, J., & Beslon, G. (2020). The complexity ratchet: Stronger than selection, stronger than evolvability, weaker than robustness. *Artificial Life*, *26*(1), 38–57.
5. Masumori, A., Sinapayen, L., Maruyama, N., Mita, T., Bakkum, D., Frey, U., Takahashi, H., & Ikegami, T. (2020). Neural autopoiesis: Organizing self-boundary by stimulus avoidance in biological and artificial neural networks. *Artificial Life*, *26*(1), 130–151.
6. Moore, J. H., Olson, R. S., Schmitt, P., Chen, Y., & Manduchi, E. (2020). How computational experiments can improve our understanding of the genetic architecture of common human diseases. *Artificial Life*, *26*(1), 23–37.
7. Packard, N., Bedau, M. A., Channon, A., Ikegami, T., Rasmussen, S., Stanley, K. O., & Taylor, T. (2019). Open-ended evolution and open-endedness: Editorial introduction to the open-ended evolution I special issue. *Artificial Life*, *25*(1), 1–3.
8. Packard, N., Bedau, M. A., Channon, A., Ikegami, T., Rasmussen, S., Stanley, K. O., & Taylor, T. (2019). An overview of open-ended evolution: Editorial introduction to the open-ended evolution II special issue. *Artificial Life*, *25*(2), 93–102.
9. Tamaru, J., Yui, T., & Hashida, T. (2020). Autonomously moving pine-cone robots: Using pine cones as natural hygromorphic actuators and as components of mechanisms. *Artificial Life*, *26*(1), 80–89.
10. Veenstra, F., Gonzalez de Prado Salas, P., Stoy, K., Bongard, J., & Risi, S. (2020). Death and progress: How evolvability is influenced by intrinsic mortality. *Artificial Life*, *26*(1), 90–111.