Why virtual creatures matter

Tired of training neural networks? Try optimizing virtual creatures instead.

The first virtual creatures were evolved by Karl Sims in the early 1990s[1]: morphology and control of autonomous agents were co-optimized so that they ran, jumped, swam, performed phototaxis and fought head-to-head for resources in a virtual world that followed (more or less) the laws of classical mechanics. The creatures were relatively simple, composed of just a handful of jointed, rigid components, but their behaviour was surprisingly rich and lifelike[2].

Karl Sims's experiments started with an objective, such as locomotion, and a population of randomly assembled creatures. Although it was unlikely that any randomly assembled creature would fully satisfy the objective, by replacing the worst-performing designs with slightly and randomly modified copies of the better ones, the population made incremental progress, generation by generation. It was the survival of the fittest; or, in the case of locomotion: the fastest.

In 2000, Lipson and Pollack[3] transferred similarly evolved creatures from simulation to reality with a technology just then emerging: 3D printing. Robot designs were rapidly and safely prototyped as virtual creatures, discarding the truly awful or dangerous designs before testing them in reality. The result were 'robotic lifeforms' that were designed, optimized and built, end-to-end, with almost no human intervention.

Since then, despite many attempts to improve on this seminal work, and despite vast increases in computational power, there has been relatively little progress scaling the complexity, competence and usefulness of virtual creatures and their 3D-printed equivalents.

This stagnancy was ended in 2013 with an important advance by Cheney et al.[4] — winners of the inaugural Virtual Creatures Competition — which introduced the evolution of soft robots composed of thousands of elastic voxels (pictured). This work became highly influential, and for good reason.

Soft robots provide a unique opportunity for virtual creatures research: soft structures can act as shock absorbers, energy stores and even power amplifiers (catapults). However, they are extremely non-intuitive to design and control due to low mechanical impedance (force applied to one part of the robot can propagate in unanticipated ways throughout its body), and because they have

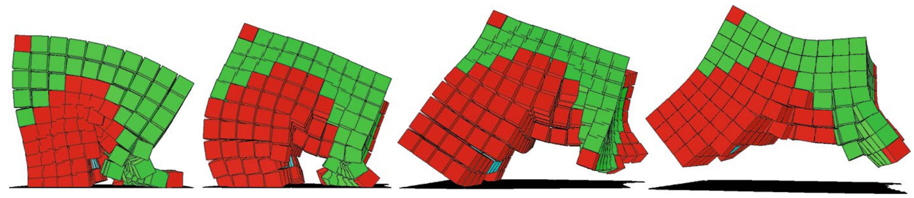

An evolved, galloping creature. Adapted from ref. [4], ACM.

much greater morphological possibility: they can change shape and material properties, as well as their control policies.

By instantiating soft machines in a virtual world, it is possible to rapidly and cheaply evaluate many candidate designs in parallel. This makes exploration of the design space much more efficient than if each proposed design needed to be built in reality, and permits us to use optimization algorithms to automatically discover and refine a robot's shape, material properties and controllers contemporaneously, thus leading to entirely novel solutions outside the purview of human intuition[5].

However, the goal of designing virtual creatures is not necessarily to design the best soft robot. The goal can also be to engage our imagination and creativity — to expand our conception of life and intelligence, and how these things might materialize.

Indeed, Sims, who was primarily interested in computer graphics, viewed his creatures as artistic expressions, but others (the author of this essay included) saw his work as an exciting new paradigm in the science of AI.

Since its inception, the Virtual Creatures Competition has been affiliated with the Genetic and Evolutionary Computation Conference (GECCO), a machine learning conference inspired by evolution and ecology rather than learning and neuroscience. This is because morphological optimization proper — changing the number and placement of mechanical degrees of freedom, not merely tuning the parameters of a predefined structure — has thus far only been demonstrated through evolutionary design.

The winners of this year's competition, Deepak Pathak and colleagues[6], have shattered this assumption by demonstrating how the morphology and controller of a machine can be co-optimized using a gradient-based learning algorithm, without evolution. Instead of evolving a monolithic morphology, Pathak et al. optimize a swarm of elemental agents — autonomous 'limbs' — that, in addition to actuating, can choose at every timestep to either attach to their nearest neighbour (forming an aggregate, symbiotic machine with a shared reward function) or detach, reconfigure and test a new design. Thus morphology is controlled by the same policy that coordinates behaviour.

What comes next for virtual creatures research may be to distil and combine the most useful properties of evolutionary and gradient-based algorithms, and discard the rest. However, this will require some finesse. More powerful optimization, without recourse to testing in the real world, might result in overfit virtual creatures forever trapped inside the computer.

But even completely virtual life can hold real value. And this is precisely where virtual creatures depart from the current paradigm of most machine learning research: virtual creatures are not constrained by their applicability. They are interesting objects of scientific investigation in their own right. More than that, they have the potential to be as beautiful and complex as life itself. ❐

Sam Kriegman

University of Vermont, Burlington, VT, USA.

e-mail: sam.kriegman@uvm.edu

Published online: 9 October 2019

https://doi.org/10.1038/s42256-019-0102-8

References

1. Sims, K. *Artif. Life* **1**, 353–372 (1994).
2. Sims, K. *Evolved Virtual Creatures* (1994); https://archive.org/details/sims_evolved_virtual_creatures_1994
3. Lipson, H. & Pollack, J. B. *Nature* **406**, 974–978 (2000).
4. Cheney, N., MacCurdy, R., Clune, J. & Lipson, H. in *Proc. Genetic and Evolutionary Computation Conference* 167–174 (ACM, 2013).
5. Thompson, A. in *Proc. International Conference on Evolvable Systems* (eds Higuchi, T., Iwata, M. & Liu, W.) 390–405 (Springer, 1996).
6. Pathak D. et al. Preprint at https://arxiv.org/abs/1902.05546 (2019).

Acknowledgements

Thanks to S. Beaulieu and J. Bongard for helpful suggestions.

Competing interests

The author declares no competing interests.