UNIVERSITÉ LIBRE DE BRUXELLES

ULB

DÉPARTEMENT DE LA COMMUNICATION ET DES RELATIONS EXTÉRIEURES

[PRESS RELEASE]

Neuro-evolutionary Robotics: a gap between simulation and reality

Publication in Nature Communication: Neuro-evolutionary robotics is an approach to realize collective behaviors for swarms of robots. A comparative study of the most popular neuro-evolutionary methods shows that the control software produced by most of the analyzed methods gives good results in simulation. Unfortunately, in experiments with real robots, all the differences were practically erased and all the control software produced by the different methods gave unsatisfactory results. This clearly indicates that the real problem that neuro-evolutionary swarm robotics must solve is the so-called reality gap problem.

Service Communication – Communication Recherche com.recherche@ulb.be

T : +32 (0)2 650 48 50 - GSM : + 32 (0)473 97 22 56

Neuro-evolutionary robotics is an attractive approach to realize collective behaviors for swarms of robots. Despite the large number of studies that have been devoted to it and although many methods and ideas have been proposed, empirical evaluations and comparative analyses are rare.

A publication in the journal Nature Communication, produced by Mauro Birattari and his team at the research center IRIDIA, École Polytechnique de Bruxelles, Université libre de Bruxelles, compares some of the most popular and advanced neuro-evolutionary methods for offline design of robot swarms.

"Concretely, these methods could enable the development of humanoid robot behavior, but neuro-evolutionary robotics is not yet routinely adopted in real-world applications," explains Mauro Birattari.

All of these methods use evolutionary algorithms to generate a neural network that controls the robots, i.e., a neural network that takes sensor readings as input and outputs actuator commands. These methods use computer simulations to generate a neural network appropriate for the specific mission that the robots must accomplish. Once the neural network is generated (in simulation), it is installed on the physical robots and deployed in the target environment.

When comparing the different methods, the researchers observed a kind of "overfitting": the design process becomes too specialized to the simulation environment, and the neural network produced fails to "generalize" to the real world. The problem is the so-called reality gap, i.e., the difference between reality and the simulator used in the design process. Although the simulator is fairly accurate, differences are inevitable.

"For example, if robots need to move back and forth between two areas, one solution that the evolutionary process might find in simulation is to produce a neural network that makes the robot move along a circular path that touches both areas. This solution is very elegant and works very efficiently in simulation. Yet, when applied to the real robots, this solution would fail miserably: for example, if the real diameter of (one of) the robot's wheels differs slightly from the nominal value, the radius of the trajectory will be different; the trajectory will no longer pass through the two given zones as desired and as predicted by the simulation," illustrates Mauro Birattari.

The Chocolate Solution

Although counter-intuitive, the solution seems to be to reduce the "power" of the design method: adopt a method that can produce a limited range of behaviors. This clearly means that one will have to accept to get worse results in simulation. This method will not perform as well in simulation as a "powerful" one because it will not be able to exploit all the characteristics of the simulator; yet, the result will be more general, less specialized to the simulator, and therefore more likely to generalize well to reality. The simpler, the better!

Chocolate seems a good illustration of this idea. Chocolate is a method that researchers at the IRIDIA Center proposed a few years ago and that does not belong to neuro-evolutionary robotics but that, in a similar way to neuro-evolution, automatically generates control software for robots, under the same conditions. Chocolate operates on pre-existing software modules that are low-level behaviors (e.g., I go in the direction of the light, I stop, I move away from perceived peers...) and conditions for moving from one low-level behavior to another (e.g., I am surrounded by peers, the color of the floor I am on is black...).

Instead of playing with a very powerful neural network capable of producing a wide range of behaviors, Chocolate plays with predefined building blocks that are (comparatively) much "coarser". The working hypothesis is that, by doing so, the risks of "overfitting" will be reduced.

This work is a key contribution of the DEMIURGE project, funded by the European Research Council through an ERC Consolidator Grant awarded to Mauro Birattari of IRIDIA, the artificial intelligence laboratory of the Université libre de Bruxelles. The DEMIURGE project focuses on the automatic design of robot swarms.

Reference:
Empirical assessment and comparison of neuro-evolutionary methods for the automatic off-line design of robot swarms
Nature Communications
Ken Hasselmann, Antoine Ligot, Julian Ruddick, and Mauro Birattari

https://www.nature.com/articles/s41467-021-24642-3

DOI : 10.1038/s41467-021-24642-3

Press contact:
Mauro Birattari
IRIDIA – Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle
Université libre de Bruxelles
mauro.birattari@ulb.be