Investigating the Mechanisms of Embodied Intelligence with Evolvable Modular Soft Robots

Eric Medvet[1,*], Giorgia Nadizar[1], Michel El Saliby[1], Francesco Rusin[1] and Berfin Sakallioglu[1]

[1] *Università degli Studi di Trieste, Trieste, Italy*

Abstract

Recent scientific and technological progresses in the development of artificial intelligence (AI) resulted in the availability of tools able to assist humans in increasingly sophisticated tasks, as, e.g., summarizing a textual information captured by photo of a document upon a spoken request. The capabilities of these tools are making the boundary between task-specific and general AI progressively fuzzier: apparently, we are hence approaching the so called artificial general intelligence (AGI). However, despite their complexity, these tasks very often deal only with abstract information, *i.e.*, there is no direct interaction with the physical world. We argue that an alternative path to AGI can be found by considering those tasks where an agent is immersed in, and has to interact with, an environment, *i.e.*, those where intelligence is *embodied.* In this brief survey, we review some recent research works where we used evolved modular soft robots as a mean for investigating the conditions for the arising of embodied intelligence. Modularity and softness constitute an opportunity for intelligence to be distributed and supported by the body in dealing with environment changes. Evolution is the way the agents can adapt to the environment. We discuss several experiments that we designed to answer specific research questions.

Keywords

Evolutionary robotics, Body-brain evolution, Morphological computation

1. Introduction

Artificial intelligence (AI) has become a popular term used by the general population to indicate a broad set of tools able to perform complex tasks which were considered inaccessible to machines until very recent times. The present and future impact of AI on society will be certainly relevant and pervasive, not only on the economy [1]. However, a sharp definition of AI is missing. Namely, defining what is really the intelligence in AI is an unsolved problem.

Several efforts are being made for defining AI. Historically, the criterion for a task to require intelligence was often whether humans outperformed "mechanisms" in solving it. This pattern regularly failed, with some well known milestones. In the 70s, the chess senior master and academic Eliot Hearst said that "the only way a current computer program could ever win a single game against a master player would be for the master, perhaps in a drunken stupor while playing 50 games simultaneously, to commit some once-in-a-year blunder". In 1997, the chess world champion Gary Kasparov was indeed defeated by Deep Blue, a "computer" built and programmed by IBM for the purpose of playing chess. Not much later, the game Go started to be considered the next frontier for intelligence, as it requires more than brute computational power, which was how Deep Blue and similar programs excelled. It was thought that the qualities that marked out the master Go player were the hallmarks of human intelligence: adaptation, intuition, and the ability to plan for the future. Eventually, in 2016, Google's AlphaGo AI defeated Go world champion Lee See-dol.

Not only games have been considered as "battle fields" for assessing intelligence. The ability to produce creative content started being considered the next challenge, with the implicit assumption that such a process would have been easily mastered by humans, but poorly performed by machines. Indeed, this idea is rooted in the history of AI: the very same Turing test is built on the intuition that

Ital-IA 2025: Convegno Nazionale CINI sull'Intelligenza Artificiale, June 23–24, 2025, Trieste, Italy
emedvet@units.it (E. Medvet); giorgia.nadizar@dia.units.it (G. Nadizar); michel.elsaliby@phd.units.it (M. El Saliby); francesco.rusin2@phd.units.it (F. Rusin); berfin.sakallioglu@phd.units.it (B. Sakallioglu)
0000-0001-5652-2113 (E. Medvet); 0000-0002-3535-9748 (G. Nadizar); 0009-0003-7402-3483 (M. El Saliby); 0009-0001-2308-8514 (F. Rusin); 0009-0009-7835-6481 (B. Sakallioglu)

the somehow creative task of fooling a human mimicking a human interaction can be used to tell apart non intelligent and intelligent agents. However, it is today rather clear that AI *excels* also in the creative content generation [2].

Recently, Chollet proposed the Abstraction and Reasoning Corpus, based on a new definition of intelligence as skill-acquisition efficiency, which takes into account scope, generalization difficulty, priors, and experience [3]. While this proposal is strongly motivated and theoretically solid, it still considers intelligence as mostly related to the manipulation of information, as all the aforementioned challenges (chess, Go, creativity). We instead argue that at least *one version* of intelligence is the one required to successfully interact with physical environments. More precisely, we believe that an agent that adapts to its dynamic environment, possibly actively modifying it, in order to persist is exhibiting some form of intelligence. A commonly accepted name for this form is *embodied intelligence* [4].

In this brief summary we review a few recent research studies we conducted with the aim of investigating the characteristics, enabling factors, and limitations of embodied intelligence. For most of these studies, we considered the scenario of (simulated) modular soft robots which are subjected to evolution. Modular soft robots are particularly well suited for studying embodied intelligence. Their body, composed of many simple and identical modules, allows for investigating the "location" of intelligence, which may be distributed across modules (which hence resembles a form of *collective intelligence* [5]) or centralized. Their softness makes the body contribute to forming the behavior of the agent, possibly even more than the brain, a form of *morphological computation* [6]. Finally, when subjected to evolution, robots can adapt to the environment by promoting more successful bodies, brains, or behaviors. While evolution is often viewed as a mere form of optimization (with very broad applicability), in robotics it can also serve as a powerful experimental tool for research [7].

2. Background: evolutionary computation and voxel-based soft robots

2.1. Evolutionary computation (EC)

Evolutionary computation (EC) is about designing, studying, and using evolutionary algorithms (EAs) and is considered a population-based, bio-inspired form of optimization.

Given an optimization problem where the goal is to find an $s^\star \in S$, the *solution space*, which maximizes (or minimizes) a *fitness function* $f : S \to \mathbb{R}$, a typical EA is an iterative process that works as follows. First, it builds an initial population $S_\text{pop} \subset S$ by sampling a probability distribution $P_\text{init} \in \mathcal{P}_S$ over S (with $\mathcal{P}_S$ being the set of probability distributions over S). Then, until some termination criterion is met, the EA repeats the following steps: (1) it builds an offspring $S_\text{offspring}$ starting from S_pop by repeatedly selecting one or two solutions and obtaining a new solution through a stochastic genetic operator $o_\text{mut} : S \to \mathcal{P}_S$ (mutation) or $o_\text{xover} : S \times S \to \mathcal{P}_S$; (2) it merges S_pop and $S_\text{offspring}$, obtaining a larger population; (3) it trims S_pop back to the initial size by repeatedly selecting and removing one solution. The selection criteria for the first phase (reproduction) and the third phase (survival) are typically stochastic and based on the fitness $f(s)$ of an individual s. Common options are tournament selection and worst selection.

When S is not a "trivial" space, *i.e.*, one for which defining meaningful initialization P_init and genetic operators $o_\text{mut}, o_\text{xover}$ is hard (e.g., the space of modular soft robots), it is common to search in a space G, called the *genotype space*, rather than directly in S, to map it to S through a mapping function $\phi : G \to S$. A common option for G is $\mathbb{R}^p$, for which many reasonable $P_\text{init}, o_\text{mut}, o_\text{xover}$ exist. ϕ, often along with its $P_\text{init}, o_\text{mut}, o_\text{xover}$, is usually called the *representation* of solutions.

2.2. Voxel-based soft robots (VSRs)

A voxel-based soft robot (VSR) is an assembly of soft cubes linked together, each with the ability of contracting or expanding its volume, hence resulting in a hopefully interesting behavior of the entire robot. While VSRs can be actually fabricated [8], a vast body of research is and has been done with a simulated version of them, often in a two-dimensional environment [9] and in discrete time, which makes the computation lighter: for this reason, they are also known as virtual soft robots [10].

A VSR has a body and a brain. The *body* can be described by a 2-D matrix describing the placement of the modules (voxels). Depending on the scenario, the modules can be identical or can be made of different materials. Typically, the material influences the maximum rate of compression or expansion (which can be zero for rigid voxels). Voxels can host sensors, capable of perceiving the external environments (e.g., proximity sensors) or the voxel itself (e.g., the current relative area). Sensor readings are usually scaled in $[-1, 1]$ so that each voxel instantaneous perception is a vector in $[-1, 1]^p$, with p being the number of sensors.

The *brain* is in charge of determining at each time step the contraction/expansion rate of each voxel: the actual change in volume depends on this control value and also on external forces applied to the voxels (e.g., the contact with other bodies in the environment). In general, the brain of a VSR is hence a dynamical system with $\mathbb{R}^{p_{\text{in}}}$ as observation space and $\mathbb{R}^{p_{\text{out}}}$ as actions space: p_{in} depends on the sensors deployed on the voxels; p_{out} is usually the number of voxels in the VSR. Several options have been explored for realizing the brain, ranging from multilayer perceptrons (MLPs), which are indeed stateless, and recurrent neural networks (RNNs), to symbolic graphs.

A common research scenario is to consider a task, *i.e.*, an environment where the VSR is immersed and a measure of quality of its behavior (e.g., locomotion on an rough terrain), a subset of components of the VSR to be optimized (e.g., brain and sensor placement), and use an EA to solve the resulting optimization problem, *i.e.*, maximizing the behavior quality by “changing” the components.

3. What makes a body good for a brain?

At the core of the idea of embodied intelligence there is the fact that the brain does not interact with the environment directly, but rather through the body. Body and brain need hence to fit each other in order to be effective. From the point of view of the optimization, simultaneous body-brain search is known to be a hard problem [11]. One way to address it is to try to characterize what makes a body good “in general” for a brain.

In [12], we hypothesized that *criticality* could be a good measure for how well a body can be “used” by a brain in different tasks. Intuitively, criticality is a property belonging to dynamical systems close to a phase transition between the ordered and the chaotic regime. VSR bodies are indeed dynamical systems, because of their softness, and can react to external inputs even without a brain.

We first proposed an experimental method to estimate the criticality of a body: in short, it consisted in applying a local stimulus to one voxel (a force impulse) and to measure how long the entire body reacts (known as avalanche). We then repeated the process on different voxels, measured the distribution of the avalanche duration, and defined as criticality the similarity of this distribution to the power law distribution.

With this measure of criticality, we run a first EA to optimize *task-agnostic* VSR bodies, for which the fitness function was the criticality, to be maximized. Then, we considered three tasks (locomotion on a flat terrain, jumping, and escaping from a cave) and run a second EA for each pair consisting of a task and a body evolved in the first phase: in this second optimization, we only optimized the brain of the VSR and the fitness was task-specific.

We compared the results (*i.e.*, effectiveness in each task) of the VSRs with a criticality-driven body (Figure 1a) against those with an hand-designed (Figure 1b) or random (Figure 1c) body. We found that, considering all the three tasks and ranking the bodies based on their overall performance, the first six bodies were criticality-driven.

We believe that this is a particularly interesting result: there is a property related to the ability of a body to give a rich response to stimuli which is a good predictor of how good will the body be for the *average* task with an evolved brain. In other words, criticality is an enabler of embodied intelligence.

In later studies, we found experimentally that the coupling of a body and a brain is particularly fragile: by removing or adding a single voxel, one can make a brain completely ineffective [13]. Similarly,

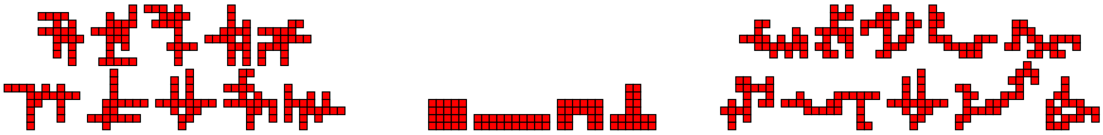

(a) Criticality-driven bodies. (b) Hand-designed bodies. (c) Random bodies.

Figure 1: VSR bodies (some of) generated in three different ways in [12].

changing the properties of the voxel material, as, e.g., softness, friction, affects the brain ability to drive the body, but the impact of the change depends largely on the body shape [14]. Both findings confirmed the relevance of criticality as an estimate of a body capability to host embodied intelligence.

4. What works as a substrate for the embodied intelligence?

As the softness and the modularity of VSR body can help the brain in dealing with the environment, through morphological computation, it is legit to wonder what degree of complexity the brain has to exhibit to make the body successful in different tasks. Indeed, early studies on VSRs employed open-loop controllers as brains: a simple sinusoidal signal (with diverse phases) in each voxel was enough to make the robot exhibit successful periodic behaviors useful for locomotion [15]. We later showed that with closed-loop controllers, in the form of artificial neural networks (ANNs) actually exploiting sensory information, VSRs could exhibit more diverse and effective behaviors. This result confirmed the capability of ANNs as controllers and paved the way for new investigation on what features made ANNs a good substrate for embodied intelligence, in particular in terms of information capacity and plasticity.

In [16] we compared four kinds of ANNs when used as VSR brains: MLPs, RNNs, spiking neural networks (SNNs), the latter in two variants—with and without homeostasis. Interestingly, the four types differ in the number of parameters using for describing the ANN and in the size of the state. Very intuitively, the former can be intended as a proxy for the complexity of the network, while the latter as a proxy for the ability to retain information. For SNNs, homeostasis is a form of plasticity that regulates automatically the threshold for firing a spike in order to mitigate too strong or too weak signals. We considered three hand-designed bodies, each equipped with three sensor configurations, and optimized the parameters of the ANN with an EA using the effectiveness in locomotion on a flat terrain as fitness function. We found that RNNs gave in general the most effective gaits, but the VSRs equipped with SNNs were the ones with the greatest generalization ability. For assessing the latter, we re-evaluated each VSR with an evolved brain on a terrain different than the one it was evolved on.

In later studies, we focused specifically on ANN plasticity, which we instantiated in the form of Hebbian learning [17]. Hebbian learning is a form of unsupervised learning in which the synaptic weights do not remain constant during the life of the agent, but update based on some coefficients and on pre- and post-synaptic signal strength.

In [18] we tackled the research question on whether the brain plasticity given by Hebbian learning can facilitate (distributed) body-brain evolution. We considered a VSR variant in which the brain is an identical ANN in each voxel and a solution representation allowing for the optimization of both the body and the parameters of the brain (the Hebbian coefficients, replicated in each voxel). We used an EA to solve the optimization problem for the locomotion task and compared the results obtained with a simple MLP and with a plastic MLP with Hebbian learning, with different degrees of plasticity. The latter requires many more parameters for a given number of sensors in each voxel (*i.e.*, for the same observation space), and hence corresponds to a larger search space. Nevertheless, plastic MLPs proved to be more effective (Figure 2 left). More interestingly, analyzing the results in detail, we noticed that the VSRs equipped with plastic ANNs did learn: their were faster in later stages of their life than those not learning, but initially slower (Figure 2 right). Finally, and more importantly, the plasticity was differently exploited by MLPs employed in different voxels, *i.e.*, they *specialized* with respect to the role

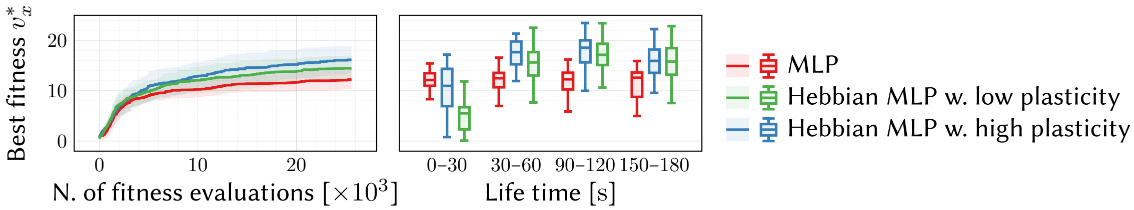

Figure 2: Velocity v_x^* (distribution of) of the best evolved VSR with an MLP as brain, with and without Hebbian plasticity. On the left, during the evolution; on the right, during the life of the single VSR.

that part of the body was playing in the robot.

5. Is intelligence centralized or distributed?

The intrinsic modularity of VSRs can be seen as an opportunity to observe forms of collaboration among modules, *i.e.*, voxels. However, differently than in swarm robotics, voxels can be very tightly coupled, as they are attached and are hence "forced" to collaborate. After having introduced in [19] an architecture for the brain that allowed for having a one-fits-all brain with respect to the body, we explored different kinds of modules collaboration.

In [20] we considered the scenario where voxels are assembled together by an external party and need to first detect the shape they are forming and then to select an appropriate brain from a library of brains. We used a neural cellular automaton (NCA) embedded in each voxel for the shape detection phase and a simple MLP for the brain. We first built a few tens of VSR bodies for which we evolved the brain in the form of an identical, replicated in each voxel, MLP—we built the library of brains from this. Then, we trained an NCA to classify the robot shape locally in each cell. Finally, we performed a set of experiments with bodies in the library and light variations of them and showed that voxels (acting as cells of the NCA) were able to choose the proper brain. Each resulting VSR was in general able to achieve its task (locomotion) even when not all the parts correctly detected the overall shape.

Later, in [21], we provided voxels with the ability to actively attach or detach from each other and evolved brains for them (identical for all the voxels involved in the task) under different conditions. Namely, we varied their ability to perceive the environment and the way they were rewarded (in term of fitness function) for how well they collectively perform the task. We considered two tasks (locomotion and piling, were voxels were required to aggregate forming a column) and used an EA for optimizing the weights of an MLP used as brain. We found that fine perception was rarely beneficial, as it brought a larger search space, where finding a good brain was harder for the EA. Moreover, we found that fitness measures promoting selfish behaviors were indeed producing selfish behaviors as, e.g., a single voxel actually "running" with the others staying steady.

6. Concluding remarks and future challenges

With this brief survey, we showed that modular soft robots and evolution can be used together to investigate the origin, characteristics, and limitation of embodied intelligence. While these tools allowed for the discovery of interesting results, there are still open challenges which are worth being tackled.

First and foremost, it is unclear to which degree the phenomena observed in simulation would hold when ported in reality: one way for cope with this uncertainty is to make progress in the physical fabrication of VSRs. Second, while there have been studies on different kinds of adaptation, *i.e.*, change happening at different time scales (e.g., brain plasticity, body development, learning, evolution), an integration of them in a single framework is missing. Third, the vast majority of studies consider VSRs "living" in isolation, while most of the interesting behaviors of intelligent biological agents is due to their co-existence with other agents: research in this direction appears particularly promising.

Acknowledgments

This work has been partially supported by the PNRR research activities of the consortium iNEST (Interconnected North-Est Innovation Ecosystem) funded by the European Union Next-GenerationEU (Piano Nazionale di Ripresa e Resilienza (PNRR) – Missione 4, Componente 2, Investimento 1.5 – D.D. 1058 23/06/2022, ECS_00000043). This manuscript reflects only the authors' views and opinions, neither the European Union nor the European Commission can be considered responsible for them.

Declaration on Generative AI

During the preparation of this work, the authors used Gemini in order to: Paraphrase and reword. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the publication content.

References

[1] M. Farina, P. Zhdanov, A. Karimov, A. Lavazza, AI and society: a virtue ethics approach, AI & SOCIETY 39 (2024) 1127–1140.
[2] E. Zhou, D. Lee, Generative artificial intelligence, human creativity, and art, PNAS nexus 3 (2024) pgae052.
[3] F. Chollet, On the measure of intelligence, arXiv preprint arXiv:1911.01547 (2019).
[4] H. Liu, D. Guo, A. Cangelosi, Embodied intelligence: A synergy of morphology, action, perception and learning, ACM Computing Surveys (2025).
[5] D. Ha, Y. Tang, Collective intelligence for deep learning: A survey of recent developments, Collective Intelligence 1 (2022) 26339137221114874.
[6] V. C. Müller, M. Hoffmann, What is morphological computation? On how the body contributes to cognition and control, Artificial life 23 (2017) 1–24.
[7] S. Doncieux, N. Bredeche, J.-B. Mouret, A. E. Eiben, Evolutionary robotics: what, why, and where to, Frontiers in Robotics and AI 2 (2015) 4.
[8] J. Legrand, S. Terryn, E. Roels, B. Vanderborght, Reconfigurable, multi-material, voxel-based soft robots, IEEE Robotics and Automation Letters 8 (2023) 1255–1262.
[9] E. Medvet, A. Bartoli, A. De Lorenzo, S. Seriani, 2D-VSR-Sim: A simulation tool for the optimization of 2-D voxel-based soft robots, SoftwareX 12 (2020) 100573.
[10] A. Mertan, N. Cheney, Investigating Premature Convergence in Co-optimization of Morphology and Control in Evolved Virtual Soft Robots, in: European Conference on Genetic Programming (Part of EvoStar), Springer, 2024.
[11] N. Cheney, J. Bongard, V. SunSpiral, H. Lipson, On the difficulty of co-optimizing morphology and control in evolved virtual creatures, in: Artificial life conference proceedings, MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info ..., 2016, pp. 226–233.
[12] J. Talamini, E. Medvet, S. Nichele, Criticality-driven evolution of adaptable morphologies of voxel-based soft-robots, Frontiers in Robotics and AI 8 (2021) 673156.
[13] E. Medvet, F. Rusin, Impact of morphology variations on evolved neural controllers for modular robots, in: Italian Workshop on Artificial Life and Evolutionary Computation, Springer, 2022, pp. 266–277.
[14] E. Medvet, G. Nadizar, F. Pigozzi, On the impact of body material properties on neuroevolution for embodied agents: The case of voxel-based soft robots, in: Proceedings of the genetic and evolutionary computation conference companion, 2022, pp. 2122–2130.
[15] N. Cheney, R. MacCurdy, J. Clune, H. Lipson, Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding, in: Proceedings of the 15th annual conference on Genetic and evolutionary computation, 2013, pp. 167–174.
[16] G. Nadizar, E. Medvet, S. Nichele, S. Pontes-Filho, An experimental comparison of evolved neural

network models for controlling simulated modular soft robots, Applied Soft Computing 145 (2023) 110610.
[17] A. Ferigo, G. Iacca, E. Medvet, F. Pigozzi, Evolving Hebbian learning rules in voxel-based soft robots, IEEE Transactions on Cognitive and Developmental Systems 15 (2022) 1536–1546.
[18] A. Ferigo, G. Iacca, E. Medvet, G. Nadizar, Totipotent neural controllers for modular soft robots: Achieving specialization in body–brain co-evolution through Hebbian learning, Neurocomputing 614 (2025) 128811.
[19] E. Medvet, A. Bartoli, A. De Lorenzo, G. Fidel, Evolution of distributed neural controllers for voxel-based soft robots, in: Proceedings of the 2020 Genetic and Evolutionary Computation Conference, 2020, pp. 112–120.
[20] G. Nadizar, E. Medvet, K. Walker, S. Risi, A fully-distributed shape-aware neural controller for modular robots, in: Proceedings of the Genetic and Evolutionary Computation Conference, 2023, pp. 184–192.
[21] F. Rusin, E. Medvet, How Perception, Actuation, and Communication Impact the Emergence of Collective Intelligence in Simulated Modular Robots, Artificial Life 30 (2024) 448–465.