ARTIFICIAL INTELLIGENT LIFE: A NEW PERSPECTIVE ON ARTIFICIAL GENERAL INTELLIGENCE

Anonymous authors
Paper under double-blind review

ABSTRACT

We propose artificial intelligent life (AILife) as a new perspective to approach artificial general intelligence, similar to a living organism in reality. Unlike machine learning approaches that focus on reward functions and mathematical optimizations, AILife seeks to develop an artificial organism that learns by the mechanism of biological neurons. AILife is composed of a biological-like neuron system to learn from interactions with the world, a sensory system to feel the world, and actuators to perform activities. We show a toy example to explain AILife.

1 INTRODUCTION

Artificial General Intelligence (AGI) Goertzel & Wang (2007) refers to intelligent systems that reason, learn, and adapt to new scenarios in a way comparable to human intelligence. In the last decade, machine learning paradigms, especially deep learning, have been widely regarded as the most promising pathways to AGI, because they have been successful in various tasks. AlphaGo Silver et al. (2016) beats the best human Go players, and ResNet Szegedy et al. (2017) surpasses human vision for image recognition. But they still face many critical problems to reach AGI, such as lacking generalization for multi-tasks Ruder (2017) and rare/unseen scenarios Ribeiro et al. (2016).

To date, human intelligence still is the only general intelligence. It is developed during interactions with the real world. Infants perceive space and time by instinctively exploring the world through watching, touching, and crawling Casasanto et al. (2010), and gradually develop more complicated capabilities. Many attempts at machine learning have been proposed to mimic such a learning process, e.g, reinforcement learning Matsuo et al. (2022) and autonomous intelligent agent LeCun (2022). However, they cannot perceive space and time as humans do, and are fundamentally different from human intelligence. Specifically, to solve a task, machine learning methods design mathematical models for optimization, while living organisms (including human) learn through complex communications among biological neurons, which are triggered by interacting with the world Shine et al. (2019). Thanks to the advances in neuroscience, the basic mechanism of neuron communication within living organisms has been revealed Hawk et al. (2018), and that enables the use of the same principles to approach AGI.

Inspired by that, we propose a new perspective to approach AGI, artificial intelligent life (AILife), aiming at developing intelligence similar to living organisms. To achieve that, AILife involves two key elements: 1) the interaction, AILife lives in a virtual/physical world and interacts with the world following pre-defined activities (instincts), and 2) the learning mechanism, AILife learns by neuron communications triggered by interactions, with biological-like artificial neuron systems.

2 ARTIFICIAL INTELLIGENT LIFE

AILife aims at developing intelligence by mimicking a living organism in reality. That is, an AILife instance interacts with the virtual/real world based on pre-defined activities, and such interactions trigger neuron communications within biological-like neuron systems of AILife to develop intelligence. To do that, an AILife instance includes an artificial neuron system to learn from interactions by the mechanism of neuron communication, a sensory system to perceive the world, and actuators to perform activities. The sensory signals and actuator actions form the input-to-output circle for the neuron system to learn, similar to the conditioned stimulus in living organisms.

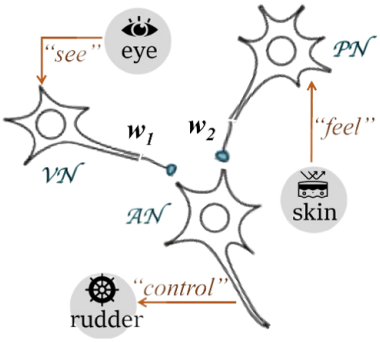

Figure 1: The neuron system of $\mathcal{P}$.

We provide a toy example (*smart pinball*) to explain AILife (code is available on Github[1]). We define a virtual one-dimensional world, which has two barriers. A pinball lives between barriers, and its pre-defined activity (instinct) is to constantly move with an invariant speed. The physical rule of this world is that barriers cannot be crossed, and thus the pinball has to change direction when hits them. We expect the pinball to be *smart* to realize this rule by interacting with the barriers, so that it changes direction before hitting them.

To do that, we design *smart pinball* ($\mathcal{P}$) based on the principle of AILife. Specifically, $\mathcal{P}$ contains: 1) a neuron system implemented by spiking neural networks (SNN) Nunes et al. (2022), which is capable of mimicking the biological learning process Taherkhani et al. (2020); 2) a sensory system that includes an "*eye*" to see a short distance in the front and a tactile "*skin*" to feel the hit; and 3) an actuator ("*rudder*") to change direction. By convention, we assume the neuron system with fixed neurons and adjustable connections Wang et al. (2020). A neuron in SNN has inputs to receive signals from other neurons and outputs to fire signals. The signals are transmitted as discrete spikes. When a neuron receives a spike from another neuron, its *membrane potential* increases, based on the weight of the connection. The neuron fires when the *membrane potential* reaches a threshold. After firing, the *membrane potential* returns to the initial value. The learning of SNN is based on Spike-Timing Dependent Plasticity (STDP) Shouval et al. (2010), which mimics the strengthen/weaken mechanism of connections among biological neurons.

The neuron system of $\mathcal{P}$ (Fig. 1) contains three neurons: 1) a vision neuron ($\mathcal{VN}$) that fires when the "*eye*" sees a barrier; 2) a pain neuron ($\mathcal{PN}$) that fires when the "*skin*" feels hitting a barrier; and 3) an action neuron ($\mathcal{AN}$) that fires with threshold δ to control the "*rudder*" to change direction. The connection of the three neurons are $\mathcal{VN} \rightarrow \mathcal{AN}$ (with weight w_1) and $\mathcal{PN} \rightarrow \mathcal{AN}$ (with weight w_2). We set the initial $w_1 \ll \delta$ assuming there is no causal relationship between vision and action in the beginning, and the initial $w_2 > \delta$ as feeling pain naturally drives $\mathcal{P}$ to change direction.

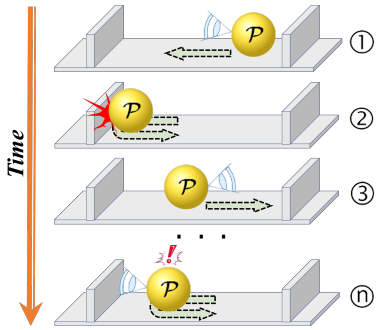

Figure 2: The learning process of $\mathcal{P}$.

The learning process of $\mathcal{P}$ is shown in Fig. 2. Initially, $\mathcal{P}$ does not know barriers cannot be crossed. Thus $\mathcal{P}$ keeps moving in one direction (Fig. 2 ①) till it hits a barrier. The hit causes $\mathcal{PN}$ to fire, which further causes $\mathcal{AN}$ to fire, due to $\mathcal{PN} \rightarrow \mathcal{AN}$ with large w_2; and that makes $\mathcal{P}$ change the moving direction (Fig. 2 ② and ③). The key point is that, at the same time, $\mathcal{VN}$ also fires as $\mathcal{P}$ can see the barrier in front. According to STDP, the connection $\mathcal{VN} \rightarrow \mathcal{AN}$ is strengthened (w_1 increases) because $\mathcal{VN}$ and $\mathcal{AN}$ express a causal relationship. After a few rounds, when $\mathcal{VN} \rightarrow \mathcal{AN}$ grows strong enough ($w_1 > \delta$), the firing of $\mathcal{VN}$ can cause $\mathcal{AN}$ to fire to change direction (Fig. 2 ⓝ). That means, $\mathcal{P}$ has learned to change direction when it sees a barrier.

3 Discussion

In this paper, we propose AILife as a new perspective to approach AGI. As preliminary work, we provide a toy example to explain AILife, expecting to attract interested researchers on exploring many exciting challenges ahead. For example, can AILife develop more complicated intelligence in more complex worlds? can an AILife instance interact with other AILife instances to evolve social activities? can AILife learn from highly abstractive information such as languages to gain high-level cognitive abilities? We envision solving these challenges requires the incorporation of more sophisticated computing technologies and cutting-edge neuron sciences. Meanwhile, we believe the development of AILife also can benefit neuroscience, because it provides a platform that allows for the customized construction of artificial organisms and neuron systems.

[1] https://github.com/brcai/SmartPinball

URM STATEMENT

The authors acknowledge that at least one key author of this work meets the URM criteria of ICLR 2023 Tiny Papers Track.

REFERENCES

Daniel Casasanto, Olga Fotakopoulou, and Lera Boroditsky. Space and time in the child's mind: Evidence for a cross-dimensional asymmetry. *Cognitive science*, 34(3):387–405, 2010.

Ben Goertzel and Pei Wang. A foundational architecture for artificial general intelligence. *Advances in artificial general intelligence: Concepts, architectures and algorithms*, 6:36, 2007.

Josh D Hawk, Ana C Calvo, Ping Liu, Agustin Almoril-Porras, Ahmad Aljobeh, María Luisa Torruella-Suárez, Ivy Ren, Nathan Cook, Joel Greenwood, Linjiao Luo, et al. Integration of plasticity mechanisms within a single sensory neuron of c. elegans actuates a memory. *Neuron*, 97(2):356–367, 2018.

Yann LeCun. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review*, 62, 2022.

Yutaka Matsuo, Yann LeCun, Maneesh Sahani, Doina Precup, David Silver, Masashi Sugiyama, Eiji Uchibe, and Jun Morimoto. Deep learning, reinforcement learning, and world models. *Neural Networks*, 2022.

Joao D Nunes, Marcelo Carvalho, Diogo Carneiro, and Jaime S Cardoso. Spiking neural networks: A survey. *IEEE Access*, 10:60738–60764, 2022.

Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. " why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1135–1144, 2016.

Sebastian Ruder. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098*, 2017.

James M Shine, Michael Breakspear, Peter T Bell, Kaylena A Ehgoetz Martens, Richard Shine, Oluwasanmi Koyejo, Olaf Sporns, and Russell A Poldrack. Human cognition involves the dynamic integration of neural activity and neuromodulatory systems. *Nature neuroscience*, 22(2): 289–296, 2019.

Harel Z Shouval, Samuel S-H Wang, and Gayle M Wittenberg. Spike timing dependent plasticity: a consequence of more fundamental learning rules. *Frontiers in computational neuroscience*, 4: 19, 2010.

David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.

Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017.

Aboozar Taherkhani, Ammar Belatreche, Yuhua Li, Georgina Cosma, Liam P Maguire, and T Martin McGinnity. A review of learning in biologically plausible spiking neural networks. *Neural Networks*, 122:253–272, 2020.

Xiangwen Wang, Xianghong Lin, and Xiaochao Dang. Supervised learning in spiking neural networks: A review of algorithms and evaluations. *Neural Networks*, 125:258–280, 2020.