Creating a Truly Conscious AI – A Novel Approach

Catarina Cunha

The relatively new field of artificial intelligence (AI), which is defined as intelligence performed by machines, is crucial for progress in many disciplines in today's society, including medical diagnostics, electronic trading, robotic process automation in finance, healthcare, education, transportation and many more. However, until now, AIs were only capable of performing very specific tasks such as low-level visual recognition, speech recognition, coordinated motor control and pattern detection. What we still need to achieve is a form of everyday human-level performance that is based on common sense, where AI are able to carry out adaptable planning and task execution and possess meaning-based natural language capabilities and generation. These are considered to be “conscious” or “creative” activities that are naturally part of our daily lives and which we execute without great mental effort. Developing conscious AI will allow us to gain knowledge and further our understanding about how consciousness works. In order to develop conscious and creative AI, machines must be self-aware; however, we hypothesize that current AI developments are skipping the most important step which will lead to AGIs: introspection (self-analysis and awareness).

Current AI development strategies

We have already managed to develop computers that can learn without being explicitly programmed to do so. This is known as machine learning, an application of AI which uses algorithms and models to perform a task. The next goal in the field of AI research is to develop artificial general intelligence (AGI) which would enable machine intelligence to accomplish any intellectual task that humans are able to perform¹⁻⁶. Another field designed to achieve this ambition is evolutionary robotics (ER), a methodology which uses evolutionary computation (algorithms designed for global optimization that are inspired by biological evolution) to develop controllers and/or hardware for truly autonomous robots. Artificial neural networks (which are discussed later in the article) and other forms of reinforcement learning have been used to some success in the context of ER⁷⁻⁹.

New developments such as cloud computing (leading to increased computing power), more sophisticated algorithms, the falling costs of data storage and acquisition and the increased general availability of data provide the ideal soil for the growth of AGIs. Renowned futurist Ray Kurzweil, known for his highly accurate predictions made in the late 90's for the year 2009, forecasts that by 2029, computers will have human-level intelligence¹⁰⁻¹³. Numerous articles have discussed this question: if an AGI is able to learn independently and is truly creative, could it eventually become aware of itself? I propose that, on the contrary, we need to first create a self-aware AI before we can develop an AGI.

From our standpoint, current AI and robot developments have one problem which prevents them from becoming AGIs: their initial information system network is built to mainly process input information. However, researchers are now focusing on perfecting robotic body representation in a 3D space (an important factor in producing a robot which can interact with its environment in a meaningful way), which presents a step in the direction of what we hypothesize is crucial for achieving creative and conscious AGIs.

In order to develop a truly creative and conscious AGI, we must analyze further the concept of consciousness, which most hypothesize is exclusive to humans. French philosopher René Descartes famously proposed, "*Cogito, ergo sum*", which can be translated into, "*I think, therefore I am*". This is further encapsulated by the statement, "*We cannot doubt of our existence while we doubt*"¹⁴ – to doubt is to think, and to think is to exist.

There are many definitions and concepts of consciousness, but we will take it to refer to an organism existing in an awake state and experiencing inner awareness. In other words, it is aware of its own representation and performs intentional self-monitoring and evaluation, otherwise known as introspection¹⁵. The definition of "self" is to be consciously aware of one's own being. In order to learn about ourselves and our consciousness, we need to go back in time and look into how we became multicellular organisms.

How did we evolve into conscious organisms?

The first multicellular organisms were the Ediacaran biota, which were tubular and frond-shaped and mostly immobile; these organisms existed without predators and flourished

for ca. 300 million years until the beginning of the Cambrian Period. Mark McMenamin, whose contributions are significant to the study of the Cambrian explosion and the Ediacaran biota, named it "The Garden of Ediacara"¹⁶. Researchers were unable to find fossils from the Ediacaran period with sensory input systems such as light receptors. As far as we know, the organisms which existed at that time did not interact with one another and there is no evidence of engagement with their environment; this means that the first information transmission (or "nervous") system was for controlling internal coordination and not sensory-motor control¹⁷. Therefore, the sensory-motor system came after the internal coordination system, and the nervous system was originally built for self-coordination and not interaction with the environment. The main challenge in producing such a system was the need for an organism to possess inner awareness in order to be able to perform intentional self-monitoring and evaluation to organize and compartmentalize different cell functions and thus become a completely functioning organism. Only the rise of predators introduced the necessity of automated introspection integration to create a fast and efficient response mechanism for survival; this led to the reinforcement of increasingly complex and distinct strategies which developed into the diversity and evolution of the cultures we live in today. Therefore, conscious self-awareness was what allowed the Ediacaran biota to organize themselves from unicellular organisms into functioning complex multicellular organisms, and this should also be the first objective in creating AGIs and sentient robots. To do so, we should first give AIs the task of introspection integration and self-organization through neural network simulations – a field which surprisingly predates the development of computers⁶.

How artificial neural nets could become conscious

In 1969, Marvin Minsky and Seymour Papert published the book *Perceptrons*, inside which they expressed their frustrations towards the limitations of neural networks¹⁸. However, much has changed since then, and the neural network field interests many due to new and recent advances. In 1943, neurophysiologist Warren McCulloch and logician Walter Pitts were the first researchers to develop an artificial neuron¹⁹. However, the technology available to them at the time did not provide enough resources for major progress.

Neural networks attempt to deduce the essential features of neurons and their interconnections in order to simulate these features in a computer. Since our knowledge of neurons is incomplete and our computing power is limited, our current models are gross idealizations of the neuronal networks in our brains²⁰. However, we hypothesize that this should not present a problem if we apply our approach of introspection design for further developments. In theory, we can give our neural networks (or neural nets) the task to learn and recognize their own components completely and self-organize themselves in order to become more efficient based on the resources provided. Ideally, the neural nets will come up with their own organization and strategies which means they will be self-aware and also creative. After the neural nets achieve this, we can introduce some input information for them to learn and subsequently allow them to distinguish themselves from the environment; by introducing them to one another, we may be able to create a form of culture. These neural nets will be able to create new and improved versions of themselves, a major ambition of large technology companies such as Google. However, if we really want to achieve this goal, we should start thinking about the risks that could come with conscious and creative AGIs. A good place to start would be Isaac Asimov's "Three Laws of Robotics", a set of rules which first appeared in his 1942 short story *Runaround*, and was later included in his 1950 collection *I, Robot*:

First Law – A robot may not injure a human being or, through inaction, allow a human being to come to harm.

Second Law – A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.

Third Law – A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.²¹

For now, AIs are under our full control and only perform tasks that we have specifically programmed them to do. In the future, however, AGIs may achieve cognitive capabilities indistinguishable from those of a human. In such cases, we need to develop a consensus on the ethical rules and laws which AI must follow, because technological singularity may change the definition of consciousness and even what it means to exist. Conscious and

creative AGIs will bring many benefits to our society, and we must also recognize and prepare for any potential changes – good or bad – they may bring.

1. Dias, R. D., Gupta, A., & Yule, S. J. (2018). Using Machine Learning to Assess Physician Competence: A Systematic Review. *Academic Medicine*. doi:10.1097/ACM.0000000000002414
2. Pellegrini, E., Ballerini, L., Hernandez, M., Chappell, F. M., Gonzalez-Castro, V., Anblagan, D., . . . Wardlaw, J. M. (2018). Machine learning of neuroimaging for assisted diagnosis of cognitive impairment and dementia: A systematic review. *Alzheimer's and Dementia (Amsterdam, Netherlands)*, 10, 519-535. doi:10.1016/j.dadm.2018.07.004
3. Liu, D., Cheng, D., Houle, T. T., Chen, L., Zhang, W., & Deng, H. (2018). Machine learning methods for automatic pain assessment using facial expression information: Protocol for a systematic review and meta-analysis. *Medicine (Baltimore)*, 97(49), e13421. doi:10.1097/MD.00000000000013421
4. Liakos, K. G., Busato, P., Moshou, D., Pearson, S., & Bochtis, D. (2018). Machine Learning in Agriculture: A Review. *Sensors (Basel, Switzerland)*, 18(8). doi:10.3390/s18082674
5. Cust, E. E., Sweeting, A. J., Ball, K., & Robertson, S. (2019). Machine and deep learning for sport-specific movement recognition: a systematic review of model development and performance. *Journal of Sports Sciences*, 37(5), 568-600. doi:10.1080/02640414.2018.1521769
6. Kohonen, T. (1988). An introduction to neural computing. *Neural Networks*, 1(1), 3-16. doi:[https://doi.org/10.1016/0893-6080\(88\)90020-2](https://doi.org/10.1016/0893-6080(88)90020-2)
7. Gigliotta, O., Bartolomeo, P., & Miglino, O. (2015). Neuromodelling based on evolutionary robotics: on the importance of motor control for spatial attention. *Cognitive Processing*, 16 Suppl 1, 237-240. doi:10.1007/s10339-015-0714-9
8. Harvey, I., Di Paolo, E., Wood, R., Quinn, M., & Tuci, E. (2005). Evolutionary robotics: a new scientific tool for studying cognition. *Artificial Life*, 11(1-2), 79-98. doi:10.1162/1064546053278991
9. Scheper, K. Y. W., & de Croon, G. (2017). Abstraction, Sensory-Motor Coordination, and the Reality Gap in Evolutionary Robotics. *Artificial Life*, 23(2), 124-141. doi:10.1162/ARTL_a_00227
10. Kurzweil, R. (2006). Reprogramming biology. *Scientific American*, 295(1), 38.
11. Kurzweil, R. (2011). Interview with Ray Kurzweil. Interview by Vicki Glaser. *Rejuvenation Research*, 14(5), 567-572. doi:10.1089/rej.2011.1278
12. Kurzweil, R. (2005). Human 2.0. *New Scientist*, 187(2518), 32-37.

13. Kurzweil, R., & Grossman, T. (2009). Fantastic voyage: live long enough to live forever. The science behind radical life extension questions and answers. *Studies in Health Technology and Informatics*, 149, 187-194.
14. Descartes, R., & Veitch, J. (1853). *Discourse on the method of rightly conducting the reason, and seeking truth in the sciences* (2nd ed.). Edinburgh: Sutherland and Knox.
15. Jack, A. I., & Shallice, T. (2001). Introspective physicalism as an approach to the science of consciousness. *Cognition*, 79(1-2), 161-196.
16. McMenamin, M. A. S. (1986). The Garden of Ediacara. *Palaios* 1(2), 178-182.
17. Dennett, D. (2019). Review of Other Minds: the octopus, the sea and the deep origins of consciousness: Peter Godfrey-Smith, Farrar, Straus and Giroux, NY, 2016). *Biology & Philosophy*, 34(1).
18. Minsky, M., Papert, S. A., & Bottou, L. (2017). *Perceptrons: An Introduction to Computational Geometry*. Cambridge, Massachusetts: MIT Press.
19. Abraham, T. H. (2002). (Physio)logical circuits: the intellectual origins of the McCulloch-Pitts neural networks. *Journal of the History of the Behavioral Sciences*, 38(1), 3-25.
20. Davalo, E., & Naim, P. (1991). *Neural networks*. London: The MacMillan Press.
21. Asimov, I. (1963). *I, Robot*. New York: Doubleday.