

## Review of “How to Grow a Robot”

In this book the author lays out the path thinks we should take to achieve human level artificial intelligence and artificial general intelligence (AGI). The author mentions that AGI has not been able materialize despite decades of research and hypothesizes that the most likely route to achieving it is through artificial human intelligence (AHI). Researchers in a field called Cybernetics, which was created to explain behavior in humans and machines, have attempted to solve this problem by modelling human behavior and trying to imitate it.

Throughout the chapter, the author discusses many new directions considered by AI researchers such as the importance of situated, embodied, and embedded agents or more generalized as enaction, and through these ideas formed the field of developmental robotics, which recognized that the growth of an agent is very important for AGI. Like we saw in the previous paper on grounded language learning, in which agents were trained on a curriculum, where they were trained on different simple tasks successively before a complex task and produced much better results than training solely on the complex task.

The main problem in this book is the idea that in order to have a human-level AI, it must be trained through real-world embodiment. While it is important for an AI which interacts with people in the real world via a humanized robotic shell, it is hard to imagine how such a robot would be more useful than one that exists virtually or one that exists for certain tasks. If a robot exists in the physical world, it still has to abide by the laws of the physical world, a human-looking robot would not magically be able to push over buildings or do things that humans could not do. If we continue to move toward a future of automation, which is to some people

the end goal of AGI, we have to consider whether it is better for a humanized robot to control a mechanical bulldozer instead of a fully autonomous bulldozer. Similarly, ordering a meal in a restaurant which has a human-like robot preparing your meal would not nearly as efficient as a robot system which was made to prepare meals very quickly.

Requiring a robot to be put into a humans shoes does not bring us closer to the utopian dream of robots working hand-in-hand to produce everything for us. Human-level AI applied to specific systems would be much more beneficial and infinitely more practical. Going back to previous examples, consider the scenario where you need to level a forest to install a neighborhood. Presently, this would be done by hiring someone who is experienced in bulldozing and all of the quirks that come with it, so for example if there were unlevelled ground in front of the driver, the driver would know how to approach it. Now consider the options for how this would be done after the creation of AGI.

From my perspective there are two approaches that can be taken to accomplish this task. We can either build a bulldozer which has been pretrained virtually before being applied to the real world where it would then familiarize itself with its controls and the real-world experiences required for a bulldozer operator, such as bulldozing many fields and updating itself when it makes a mistake; or we can build a human-like robot to use existing bulldozers to level the forest. While the latter option would be more universal, it would be more practical to create a machine which can use AGI to understand exact instructions and have perfect knowledge on how form the land than to put a human robot to the same task.