The Next Physical Wave

Subin Kim

July 9, 2024

The reason the Industrial Revolution was able to change the world was that it enabled the operation of machines that replaced human physical labor. Chat-GPT is gaining attention because it replaces human intellectual labor and increases work efficiency. I believe that within the next 10 years, physical labor will again be replaced by robots. However, these robots will not be like the heavily human-dependent machines of the Industrial Revolution era. Instead, they will be capable of thinking for themselves, generalizing, and solving various tasks autonomously. To create such robots, I think it is necessary to conduct research on creating the "brain" of a robot.

Most existing AI-based robotics research has primarily focused on"imitation learning." In imitation learning, when teaching a robot a desired task, a human creates example data of successfully completing the task, essentially saying, "try to follow this." While imitation learning can teach specific tasks with high accuracy, it has clear limitations: a large amount of example data is required. In many cases—for instance, with Boston Dynamics' new humanoid"Atlas," which was unveiled two months ago and can perform gymnastics in ways humans cannot—it is impossible to create such data. Therefore, while existing industrial robots which performed simple repetitive tasks could be commercialized using imitation learning, it is difficult for this method to serve as the brain of a robot that can generalize and solve multiple tasks.

A learning method that can replace this is reinforcement learning. Reinforcement learning is intuitively aligned with how humans learn. When an agent performs well, it receives a reward and learns actions to maximize this reward. The advantage of this method is that if the reward function is appropriately set, tasks that were previously impossible can be learned through the reward function.

Based on this idea, I conducted research on reinforcement learning as an undergraduate intern and master's student in Professor Sung-Jin Ahn's lab at the Department of Computer Science at KAIST. I researched based on "world models." A "world model" refers to a virtual model that an agent can use to simulate based on the real environment. For example, if a person is familiar with the geography of a place, they can anticipate that "if I turn the corner, there will be a refrigerator" without actually turning the corner. Furthermore, we trained the agent to divide the environment into several regions and learn a macroscopic model for moving between regions and a microscopic model for detailed movement within regions. This is similar to how people use major routes like airplanes or trains first and then use taxis or other means to find detailed destinations. This research was accepted as a joint first author at ICML 2024, a world-renowned conference in the AI field.

I believe that to create generalizable robots, two things are necessary. This is also the common factor of models in other fields like Chat-GPT for LLM or Dall-E for vision: **Data (=Knowledge/Experience, Input) and Neural Networks (= Brain, Algorithm)**. Chat-GPT was trained using an enormous amount of data, and generative AIs like Dall-E have shown tremendous performance improvements recently with techniques like diffusion methods. In the robotics field, there is also a project called Open X-Embodiment, where 21 institutions worldwide are gathering data on robots performing 160,266 tasks.

Based on this trend, I want to research two things: 1) Although a large amount of data can be collected, it is often inconsistent and noisy. I want to study ways to standardize this data into a single format for training using LLMs. 2) Additionally, agents trained with this data have limitations in performing tasks beyond the data. I want to research how to enable trained agents to combine what they have learned in new ways to perform new tasks not present in the existing data.