

Jacob Devasier

4/2/2021

Review of The Society of Mind

The Society of Mind provides a historical perspective on the brain and theories about how it functions. Tasks that seem simple for humans to do, such as stacking blocks or drinking tea, are actually quite complex when it comes time to program. If you were to take a drink of tea, you might need an agent telling your hand how grasp the cup, an agent ensuring that the tea does not spill as you move it to your lips, an agent to swallow the tea in your mouth, and so on. This idea is the main focus in these sections of the book. Humans are able to learn a variety of different tasks that which are an abstraction of many other tasks leading to an oversimplification of the task at hand.

The author introduces the idea of a hierarchical control system inside our brains that abstract more complex tasks, i.e., drinking tea would be an abstraction of the processes that are mentioned above. The primary problem that arises in this model is that the hierarchical structure introduces blocks, i.e., the seeing agent in the move-see block example. To solve this problem, loops and cross-connected rings can be implemented as mentioned by the authors.

Though I think there is something the authors may have looked past before adjusting the model when the problem arose. When humans do an abstract task like drinking tea, we pull from many shared resources to do so, for example our eyes. In the model described in the book, recursive dependencies caused problems, i.e., the seeing-moving problem. Humans do not work synchronously, so when we need to perform the task, we use our seeing ability to provide the information to multiple different agents

performing different independent tasks. For example, when walking on a sidewalk on a busy road, it is important to watch out for hazards at the same time. Obviously humans do this simultaneously, we don't take one step, scan our surroundings, take another step, scan again, and so on. These are tasks that run in the background of our mind.

This asynchronous model is still supportive of the idea of abstracted actions. It is true that there are many actions that we take that are combinations of lower-level actions which may also be abstracted from even lower-level actions. The hierarchical structure still exists, but without the dependencies for certain resources and without the synchronous methodologies implied in the book.

A related issue that arises is the idea of conflicts moving upwards through the hierarchy. This idea feels like it is missing some crucial aspects and runs into the following example: If I cannot decide whether I want to build a stack of blocks one block higher or knock it over, my brain would likely not come to the conclusion of doing neither instead. While it is apparent that conflicts will arise, the way in which they are settled seems different from moving up in the hierarchy and potentially allowing another agent to take over.

One point that the author mentioned that is very strong and quite obvious once considered, is the idea that when two agents strongly want something, negotiations can happen between them to settle the dispute. If we imagine a probabilistic approach where the next task is chosen by an external agent based on some score, we can imagine that there will be times when some of the scores may be very similar, i.e. the array $\langle 0.4, 0.4, 0.1, 0.05, 0.05 \rangle$ where the agent is trying to determine whether to do

task 1 or task 2. There may be some agreement that can be made between the respective agents that cause the array to look more like $\langle 0.7, 0.2, 0.1, 0, 0.5 \rangle$.