Imagination, Human and Artificial

Donald Perlis

University of Maryland, College Park
perlis@cs.umd.edu

Abstract

Humans imagine things. We live our lives in great measure by imagining circumstances a bit different from what we find, and then (again using imagination) we explore what it might take to bring those circumstances about, or what it might be like to live in such circumstances. We do this regarding issues large and small, all day long, every day. It is how we operate, and it gives us a huge leg up in detecting and repairing our own confusion as we negotiate this complex dynamic world. It is also quite different from how our artificial systems operate.
.

The Importance of Imagination

I hypothesize that a considerable portion of the gap between human performance and that of our artifacts in general-purpose behavior lies in the difference in the use of imagination. Humans use it routinely, and our artifacts almost not at all. While expert systems often match or surpass human performance, no current architectures come remotely close to us in across-the-board flexibility. There is a substantial effort now in endowing systems with metacognitive skills so that they can note and address their own shortcomings. What I argue here – via an example – is that perceptual imagination can add a great deal to that effort.

An Example

Traditional AI planning ignores perception, imagery, and groupings, and so takes much longer than necessary. By focusing on both perceived and perceptually imagined states, a robot can come closer to human-level performance.

Consider the blocks-world problem illustrated below. A stack of blocks is to be rearranged so that the bottom block ends up on top; nothing else is to change. Of course, to achieve this, all the blocks must be moved. But we are to suppose that the problem specification does not state that the bottom block is to go on top; rather only the initial and goal states are given. These state descriptions could be in the form of a set of five sentences, such as On(B,C), etc.
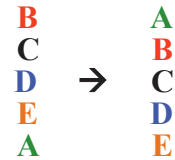
Blocks-World Meets the Real World.

A traditional planner will be swamped by an vast number of possibilities to consider, with five objects each of which can be moved in various ways. To be sure, there are heuristics that can be very helpful. But these must be supplied by the programmer.

On the other hand, an agent that can visually isolate the grouping BCDE as a unit that needs to end up in that same formation, and block A as a separate entity to reposition above the BCDE grouping – that is, an agent that can "see" block A as the only one that needs to end up in a different position – can plan far more efficiently than one that reasons only about pairwise relations between individual blocks.

Not only that: a physical agent must also pay attention to the actual moving of blocks, and for instance will benefit from reasoning that as B, C, D, and E are taken off the stack, it is helpful to place them in convenient locations for ease of rebuilding that grouping afterwards.

These issues are not addressed in traditional AI planning. While this is just one example, I hope that it makes a compelling case for the power of perceptual imagination. Now we only (!) need to make visual imagination and reasoning work together.