

Between a rock and a hard place: Cognitive Science principles meet AI-hard problems.

Artificial Intelligence and Cognitive Science have always been overlapping disciplines. Early in their history, that overlap was considerable. Herbert A. Simon wrote that “AI can have two purposes. One is to use the power of computers to augment human thinking. . . . The other is to use a computer’s artificial intelligence to understand how humans think.” Conversely, at the foundation of the Cognitive Science Society, Artificial Intelligence was identified as one of the core constituent disciplines. However, over time the two disciplines have increasingly diverged under seemingly incompatible constraints. The difficulty of many problems tackled by AI led it to adopt brute-force or domain-specific solutions that arguably were not cognitively plausible. Conversely, the need for precision and reproducibility increasingly led cognitive science to focus on experimental paradigms that AI did not recognize as hard problems. Recently, in his AAAI presidential address, Tom Mitchell called for a rapprochement between the two disciplines on the basis of convergent evolution.

In this symposium, we will explore what each discipline can expect to contribute to the other in the relatively near future. More specifically, we will ask whether cognitive science approaches can be applicable to AI-hard problems. Conversely, what impact can the challenges of AI-hard problems (and the techniques that have been applied to solve them) have on cognitive science? To improve focus and enliven debate, we will structure the symposium around a number of questions with clear answers outlining sharply divergent positions. These questions and the resulting dual positions include:

1. Is cognitive science relevant to AI problems?
 - A. Yes. The human mind is one of the most flexible, general and powerful thinking devices possible. The right way to solve AI problems is the cognitive way.
 - B. No. Many of the workings of the human mind are specific to its substrate and problems. Cognitive methods have failed before and will fail again.
2. Are “good enough” solutions valuable?
 - A. Yes. Too much focus has been devoted to the search for optimal solutions. Robustness over a broad range of problems and computational tractability are essential to long-range success in cognitive computing.
 - B. No. Optimal performance is essential, even if it results in high computational demands. Moore’s law and distributed computing will provide the scale needed to reach brain-like computing power within a generation.
3. Are multi-level heterogeneous approaches beneficial?
 - A. Yes. Most problems are not amenable to a single approach but instead demonstrate aspects best addressed at different levels, such as numerical techniques, expertise-based pattern matching or meta-level reasoning.
 - B. No. Problems are typically best approached using a single, targeted method (e.g., search in chess) with maximal efficiency. These methods can be hand-selected by human designers.

4. Is adaptiveness an essential component of intelligence?
 - A. Yes. A key component of human intelligence is its ability to adapt to constantly changing circumstances. Hardwired solutions are brittle and in constant danger of becoming obsolete.
 - B. No. Adaptiveness is often a result of insufficient information or ill-defined problems. When a problem is well defined and understood, the best solution can be found and developed once and for all.
5. Are the most efficient solutions problem-specific?
 - A. Yes. Problem-specific solutions will always be the most efficient because they do not have to include the overhead of generality.
 - B. No. Real-world problems are so complex that they include many different aspects which require significant generality to be solved in a robust manner. Instantiation of general knowledge and capabilities and re-use of capabilities across problems are two possible mechanisms to provide the needed cross-problem generality.
6. Is developing specialized, modular components a reasonable way to study general intelligence?
 - A. Yes. The significant progress in both AI and cognitive science have demonstrated that studying individual phenomena largely in isolation is necessary for scientific progress in the field. Recent experimental findings in neuroscience also support modularity as a core principle of neural organization.
 - B. No. A core feature of intelligence is the ability to bring disparate information together. Modular approaches to cognition and intelligence will not scale because assumptions within modules are incompatible across the modules.
7. Can Artificial Intelligence contribute to our understanding of human cognition?
 - A. Yes. Artificial Intelligence addresses many problems central to human cognition in an effective, functional manner. AI solutions reflect fundamental constraints that are applicable to all approaches to those problems, including human cognition.
 - B. No. Artificial Intelligence provides an ad hoc approach to problems that sacrifices generality and plausibility for the sake of performance. Cognitive solutions to problems are fundamentally different from AI solutions.

Participants will be asked to take positions on one or more of these questions by using examples from their own work or synthesizing other works into coherent patterns. Contributions can be in the form of either full-size 6-page papers or short 2-page white papers. The workshop will be structured in a number of sessions each centered around one of these questions. Each session will include two to four presentations arguing both sides of the issue, followed by a brief closing argument and a substantial discussion session between panel presenters and the workshop participants.

The list of questions should not be considered a definitive list. Indeed, we would welcome the identification of additional issues that bring up new connections between the Cognitive Science and Artificial Intelligence communities. However, Contributors who focus their submissions primarily on recent progress in their specific areas of research, unless it is directly and obviously relevant to the general goals of the workshop, will be discouraged.