Interactive Cognitive Systems and Social Intelligence

Pat Langley PATRICK.W.LANGLEY@GMAIL.COM
Institute for the Study of Learning and Expertise, 2164 Staunton Court, Palo Alto, CA 94306 USA

Abstract

Research on cognitive systems adopts the aims and assumptions of classical AI research, emphasizing the construction of intelligent agents that exhibit complex behavior. In this paper, I review the cognitive systems paradigm and two widely adopted hypotheses – physical symbol systems and heuristic search – that underpin it. I also introduce a third claim – the social cognition hypothesis – that intelligence requires the ability to represent and reason about others' mental states. I review a number of cognitive systems, both historical and recent, that focus on interaction and that exhibit this capacity to varying degrees. Examples come from work on dialogue systems, synthetic experts, believable agents, tutoring systems, interactive robots, and instructable game players. In closing, I identify some issues in social cognition that deserve greater attention and pose some challenges that can drive future research on interactive cognitive systems.

1. Introduction

Artificial intelligence was launched as a field at the 1956 Dartmouth meeting. The aims were audacious: to understand the human mind in computational terms and to reproduce all known mental abilities in computational artifacts. Many of the latter involved single-agent activities, such as proving theorems, generating plans to achieve physical goals, understanding written stories, and designing novel structures. But other tasks involving interaction – through natural language and other means – were also a central part of the vision, influenced by Turing's (1950) early proposal that conversational ability might be used in tests for intelligence.

In this paper, I review the cognitive systems movement, which adopts the same goals as early AI researchers, and its efforts to develop interactive intelligent systems. I start by reviewing assumptions that underlie the paradigm and how they differ from what has become the AI mainstream. After recounting two hypotheses about intelligence that it adopts, I propose a new hypothesis about the key role of social cognition. Next I review both early interactive systems developed in this tradition and more recent work on interactive agents, including examples that involve synthetic experts, believable agents, tutoring systems, instructable game players, and robotic dialogues. Finally, I identify open research problems for the field and propose challenge tasks that would address them.

2. The Cognitive Systems Paradigm

Early research in AI adopted a number of assumptions that distinguish it from most recent work. The cognitive systems movement, named after a major DARPA initiative launched by Brachman and Lemnios (2002), pursues the field's original goal of constructing computational artifacts that address the full range of human intelligence. The paradigm also adopts similar postulates that guide its research efforts. Langley (2012) discusses six of its key assumptions:

- *High-level cognition*. Research focuses on abilities that are distinctively human, such as understanding language, complex reasoning, and solving novel problems, as opposed to those shared with dogs and cats, like perception, categorization, and motor control.
- *Structured representations*. The AI revolution revealed that computers are not mere number crunchers; they are general symbol processors that can encode, interpret, and manipulate rich mental structures, including substantial amounts of domain knowledge.
- *Systems perspective*. The paradigm makes progress not by focusing on component algorithms, but rather by developing integrated systems that clarify how intelligence arises from interactions among cognitive abilities.
- *Human cognition*. Researchers draw many of their ideas, inspiration, and challenges from theories and empirical findings in cognitive psychology, especially the study of human language, reasoning, and problem solving.
- *Heuristics and satisficing*. The movement assumes that intelligent behavior relies heavily on heuristic methods that do not guarantee the best or even any solution, but that typically reduce search and make complex cognition tractable.
- *Exploratory research*. Cognitive systems adopts a flexible approach to evaluation that encourages demonstrations of new functionality, respects new approaches to well-established problems, values analyses of challenging tasks, and favors architectures for integrated intelligence.

Taken together, these assumptions lead to a research style that differs substantially from what has become the AI mainstream and that has far more in common with the field's earliest days. This does not mean there is no value in work that violates some of these tenets, but it does mean that research in the cognitive systems paradigm has both a venerable history and importance of its own.

Before proceeding, I should mention the excitement among many AI researchers about recent progress on statistical methods for learning on tasks that involve pattern recognition and reactive control (e.g., Schmidhuber, 2015). Although this work constitutes clear progress over earlier approaches to such tasks, they have marginal relevance to problem solving, multi-step reasoning, language understanding, or other forms of high-level cognition. They focus on categorization during perception and stimulus-response schemes for action, which are important to any physical agent but which are not abilities we associate with the term *intelligence*. In fact, these ideas have been linked historically with behaviorism and information theory, two movements that the cognitive revolution of the 1950s rejected (Miller, 2003). They also ignore the central role of high-level cognition in human learning (Langley, 2016), which acquires knowledge far more rapidly than by mere induction. The cognitive systems paradigm builds on these insights, in contrast to recent statistical methods.

3. Three Hypotheses About Intelligence

Newell and Simon (1976), in recounting their early contributions to AI and computer science, argued that every scientific field is built on qualitative hypotheses which lay the foundation for later work. They gave examples of such 'laws of qualitative structure' from the history of science, including the cell doctrine in biology, the atomic theory of matter, the germ theory of disease, and plate tectonics in geology. These hypotheses provided the context for more detailed and precise accounts of phenomena that appeared later in each field.

They also made two analogous claims about the nature of intelligent behavior. The first, which they called the *physical symbol system* hypothesis, stated that:

- *A physical symbol system has the necessary and sufficient means for general intelligent action.*

They defined a *symbol* as some persistent physical pattern that remains stable unless modified, and a *symbol structure* as an organized set of symbols. List structures stored on a digital computer are the classic example, but sentences on paper and equations on a blackboard are equally valid instances. A *physical symbol system* has processes for creating, modifying, and interpreting such symbol structures. Symbolic processing of this sort is the fundamental idea that has enabled most successes in artificial intelligence over its sixty-year history.

Newell and Simon's second claim elaborated on the first, focusing on structures and processes that underlie the ability to solve novel problems. This *heuristic search* hypothesis stated that:

- *Problem solving involves search through a space of candidate solutions guided by heuristics.*

The problem solver represents candidate situations, actions, and solutions as symbol structures, and it generates, modifies, and tests these structures, making it a form of physical symbol system. The search metaphor has been adopted widely by AI researchers, with many successful breakthroughs depending on it, but the notion of *heuristics* has fared less well in the mainstream community, where many have come to view them with disdain because they lack formal guarantees.

Both hypotheses are reflected in core assumptions of the cognitive systems paradigm. The first relates closely to the notion of structured representations and knowledge, an important class of symbol structures, whereas the second maps directly onto its reliance on heuristics and satisficing to make problem solving tractable. However, they also suggest a third idea,[1] not reflected extensively in early AI research, that I will call the *social cognition* hypothesis:

- *Intelligence depends on the ability to represent models of other agents' mental states, generate and reason over such models, and use them for informed interaction.*

This claim builds directly on the other two, in that it posits social cognition relies on symbol processing and heuristic search, but it seems worth calling out because so many facets of human intelligence involve interaction with others. To my knowledge, this hypothesis has never been stated explicitly in the AI literature, although it has obvious links to Turing's early proposal (however problematic) for testing computational intelligence through extended conversation. This third postulate suggests interesting challenges for research on cognitive systems and, as we will see, implicit concern with this idea has also led to clear intellectual progress.

4. Classic Work on Interactive Cognitive Systems

In this section, I review three reasonably early computational artifacts that fit my characterization of interactive cognitive systems and that exhibit facets of social intelligence. In each case, I describe the system's abilities and underlying assumptions, what made it interesting and important intellectually, and the issues that it still left unaddressed.

1. Other postulates about intelligence are certainly possible, but we will not attempt to discuss them in this article.

4.1 The SHRDLU System

One of the most visible and influential early AI efforts was Winograd's (1972) SHRDLU, which interacted with human users in natural language text. Exchanges focused on a simulated blocks world that humans could view on a graphics display and to which the system had direct access. Users guided the conversation by typing English sentences. These could include commands like "Find a block which is taller than the one you are holding and put it into the box" and questions like "Is there anything which is bigger than every pyramid but is not as wide as the thing that supports it?" These inputs required not only the ability to parse quite complex structures and extract their meanings, but also to draw inferences about relationships and execute multi-step activities. The innovative system also handled simple anaphora, disambiguated word senses, and had basic memory for its previous interactions.

SHRDLU was an important advance because it integrated sentence-level understanding, reasoning about domain content, execution of multi-step activities, and natural interaction with human users. Nothing of its sort had existed before, and it offered a proof of concept that such an integrated intelligent system was possible. This accomplishment relied on some important simplifying assumptions. SHRDLU operated in a narrow and well-defined domain, which limited the knowledge needed to understand sentences and carry out instructions, and it was constrained to grammatical English sentences and vocabulary relevant to the blocks world. The system had complete access to the entire state of simulated environment, and its actions always had the expected effects. Nevertheless, it was an impressive achievement that fostered further work on intelligent agents.[2]

4.2 The Practical Algebra Tutor

Interactive cognitive systems have obvious applications in education, where individual tutoring appears to play an important role in learners' success. One system, the Practical Algebra Tutor (Koedinger et al., 1997), presented students with word problems and traced their steps as they attempted to answer questions through algebraic manipulation. The system encoded knowledge about the domain as production rules, not only the target structures but also plausible misconceptions, which it used to track students' behavior, infer what content they had mastered, and decide where they needed assistance to provide personalized instruction. Experiments with students in three Pittsburgh high schools showed that its use led to substantial improvement on standardized tests, and a commercial descendent of the system has since been adopted by hundreds of US campuses.

This tutoring system offered another substantial advance for the field. Like the previous example, it integrated a number of capabilities into a single cognitive system that interacted with humans in a complex and challenging setting. The main limitations concerned its focus on procedural knowledge, which applies to some educational topics but not others, and its highly stylized approach to interacting with students. Constructing the knowledge base that guided instruction was also a challenge, but less than in some applications because it was constrained and partly codified in textbooks. More important was that its models of students' mental states specified only what knowledge they had mastered or lacked, and that it ignored important social factors like emotions and motivation.

2. There have been few attempts to build on Winograd's work, an important exception being PLAYMATE (Christensen et al., 2010), a robotic agent that carried out similar interactions in a physical setting and conversed in spoken English.

A final limitation was the system's particular interaction style, which has been partly addressed by more recent work on tutorial dialogue systems (Graesser et al., 2001); these converse with students in spoken language, giving personalized instruction based on their answers to questions.

4.3 The TacAir-Soar System

Another effort – even more audacious and impressive – resulted in TacAir-Soar (Jones et al., 1999), an AI system that reproduced the behavior of human pilots on tactical air combat missions. This operated, along with human participants, in a realistic simulated environment used for military war exercises. The system incorporated a large knowledge base that encoded established doctrine for air combat, implemented in the Soar architecture and organized as a hierarchical task network. TacAir-Soar integrated abilities for reasoning about space and time, generating its own plans and recognizing those of others, executing plans reactively, and coordinating with teammates using constrained natural language text. The system flew 722 missions during the STOW-97 simulated training exercise, with human pilots reporting that the synthetic agents were indistinguishable from people.

TacAir-Soar was another important advance for artificial intelligence. The project showed how to organize large amounts of knowledge, about many different topics, and use it efficiently enough to simulate the behavior of highly trained humans on a very challenging task. Moreover, the system interacted in real time not with a single adversary but with multiple opponents and with multiple allies, communicating with the latter as needed to coordinate complex joint activities. TacAir-Soar was efficient, effective, and robust, and human pilots judged it an excellent teammate in a domain far richer than chess. One drawback that emerged from the project concerned the time and expense needed to construct the knowledge base that provided this expertise. Also, despite its expertise as a fighter pilot, the system could not operate outside this complex but circumscribed domain.

5. Recent Progress on Interactive Cognitive Systems

Research in this arena did not halt at the end of the last century. Efforts to develop computational artifacts that demonstrate key features of social intelligence have continued. In this section I examine some more recent examples, again describing their abilities, the reasons they are interesting, and the simplifying assumptions on which they have relied.

5.1 The Façade System

Another important example of an interactive cognitive system is Mateas and Stern's (2005) Façade, which actually comprised multiple intelligent agents that operated in a simulated apartment. Users could view this environment, including avatars for two of the agents, through a graphical interface, and they could interact with them using natural language text. The agents could understand these sentences, generate responses or initiate their own utterances, control their bodies' gaze, expression, and gestures, and exhibit emotional responses and distinct personalities. A third agent served as a high-level manager that modulated the avatars' behavior to achieve certain dramatic goals. Many users have felt they were interacting with a genuine but troubled couple.

Façade was another intellectual breakthrough for the field, offering the first believable agents that interacted with users in a natural manner over an extended period. These agents integrated a number of abilities that are central to human interaction and demonstrated their integration in a dramatic setting. They infer simple models of the user's beliefs and goals, they attempt to manipulate the user to achieve their own ends, and they display a variety of emotions in the same situations as humans. One drawback of Façade, as with TacAir-Soar, was the time and effort needed to create its agents' knowledge bases. Moreover, although they exhibited impressive abilities in social cognition, these were nevertheless limited to the setting in which they operated.

5.2 Human-Robot Interaction

Even more recently, Talamadupula et al. (2014) have reported a mobile robotic system that interacts with a human teammate in the context of disaster relief scenarios, which involve finding and treating injured people in dangerous spaces for which available maps may be unreliable. Interaction takes place through spoken natural language, with the robot accepting commands and answering questions it receives from the teammate and informing him of important events as they occur. Based on these communications, the robotic agent interprets the human's goals and intended plans, generates its own top-level goals to drive behavior, interleaves generation of plans with their execution in pursuit of these goals, and even learns the meanings of new terms to extend its vocabulary and improve its ability to interact with others more effectively.

This system is not as advanced as others I have considered, since the project has not had access to the same level of engineering resources, but it offers a proof of concept that interactive cognitive systems have an important role to play in robotics. Like the previous examples, this effort has shown that one can integrate language processing, plan understanding, plan generation, reactive execution, and social cognition to support coordinated behavior in pursuit of complex joint activities. Results with the system have been limited to reasonably narrow domains of application and to only a few participants, but it shows that the cognitive systems approach holds considerable promise for supporting human-robot coordination, an important topic that has been receiving increased attention.

5.3 Instructable Game Players

In other recent work, Hinrichs and Forbus (2014) have described an agent that learns to play games from instruction and demonstration. The human teacher provides natural language text about participants, layout, entities, and moves, which the system interprets using generic knowledge about games to resolve ambiguities. The instructor also demonstrates these ideas with diagrams on a sketch pad, which the program maps onto known concepts and activities. Once it has acquired this content, the system can use it to play the game by making legal, but not always good, moves at appropriate times. The developers demonstrated their artifact's abilities on Tic Tac Toe, which it learned from only 11 instructions, and Hexapawn, which it acquired from 16 utterances.

Like the previous effort, this system is still in its early stages, but it demonstrates a novel capacity – learning to play games – by integrating natural language processing, sketch understanding, and multi-step reasoning, as well as the ability to use the resulting knowledge in practice. The computational artifact uses multi-modal interaction and background knowledge to constrain interpretations of the tutor's input, letting it acquire complex procedures that take reasonable actions in a variety of different game-playing contexts. Kirk and Laird (2014) have reported a similar system that learns – and transfers across settings – entities and rules of games and puzzles without sketch understanding but with the ability to ask clarification questions. These efforts offer a radically different – and far more human-like – approach to learning in games than assumed by techniques for reinforcement learning (e.g., Tesauro, 1995) that have become popular in some circles.

5.4 Other Progress on Social Cognition

This list does not exhaust research on the topic of interactive cognitive systems. I have examined only a sample of promising results that can serve as role models for future work. Some other impressive examples include:

- The TRAINS system, an interactive assistant that integrated planning, language processing, and inference to help users create transportation plans through mixed-initiative spoken dialogue (Allen, Miller, Ringger, & Sikorski, 1996);
- COLLAGEN (Rich, Sidner, & Lesh, 2001), which assisted users in operating complex devices through a graphical interface, asking them questions and giving advice as needed until they completed the joint activity successfully;
- The Virtual Humans project (Swartout et al., 2006), which has created a number of synthetic characters that interacted with users in realistic settings by combining language, nonverbal signaling, planning, emotions, and social reasoning;
- The Artificial Receptionist (Bohus & Horvitz, 2009), a synthetic character that integrated spoken dialogue, vision, and inference to welcome visitors to an office building and help them achieve their meeting-related goals;
- PLAYMATE and related robotic systems (Christiansen et al., 2010) interacted with humans in spoken language and operated in physical settings to collaboratively carry out complex tasks.
- The Maryland Virtual Patient (McShane et al., 2012) combined language understanding, reasoning, decision making, episodic memory, and language generation to engage in textual dialogues with a doctor about symptoms, medical treatments, and side effects it experienced.

Like earlier examples, each of these engineering efforts produced an impressive interactive cognitive system that showed specific technical advances and lent credance to the social cognition hypothesis.

Despite their successes, these artifacts share some characteristics that deserve greater attention from the research community. One is that they construct limited models of other agents' beliefs and goals, typically only what is needed for their particular task, which reduces their extensibility. Another issue is that they carry out little reasoning about users' emotions or personalities, which would let them adapt to these individual differences. A related constraint is that their interactions with users are reasonably short, meaning they cannot gain enough experience with them to infer deeper models of this type. I do not intend these as critiques, since each system achieved its aims, but rather as guides for future research.

6. Challenges for Interactive Cognitive Systems

I maintain that previous work on interactive cognition has led to genuine insights about the social character of intelligence, but also that much remains to be done. On what problems should future researchers focus their energies? One feature that the highlighted efforts had in common is that they started with a challenging and audacious task, then designed and implemented a system that achieved it, using available components where possible and developing new ones when necessary. In this section, I pose four challenge problems that could drive additional progress in the area. I do not attempt to define them formally, but only to outline the generic task in each case. I also outline component-level research that would support these integrative efforts.

6.1 Deep Conversational Assistants

Spoken-language dialogue is the natural mode for aiding users on many tasks, such as driving, cooking, and shopping. Systems like Siri and Cortana provide some abilities along these lines, but they are very limited in scope and they do not engage in very long or very deep conversations. A few prototypes have inferred models of their users' long-term preferences, such as Thompson, Göker, and Langley's (2004) destination advisor, which engaged in personalized dialogues based on user profiles it inferred from earlier interactions. However, even these systems have emphasized retrieval from memory rather than social inference and reasoning.

We need additional research on deep conversational assistants that carry out extended dialogues about goal-directed activities. These interactive cognitive systems should take into account the surrounding task context, they should infer joint beliefs and goals – what Clark (1996) has called *common ground* – and they should store and utilize content about previous interactions with the user. Gabaldon, Langley, and Meadows (2014) report one such system that assists people in carrying out complex procedures about which it has more knowledge, inferring the user's beliefs and goals over the course of their dialogue. However, it utilized logical statements for inputs and outputs, and future assistants of this sort should combine deep processing of natural language with reasoning about others' mental states, drawing on social cognition to help achieve their users' aims.

6.2 Richer Nonplayer Game Characters

Another obvious area for interactive cognitive systems is computer games, which Laird and van Lent (2001) have argued is the 'killer application' for artificial intelligence. Synthetic characters have played key roles in this setting for decades, but, despite improved graphical depiction, they typically remain very shallow. Much of the AI research in this area has focused on improved algorithms for finding paths or techniques for executing behavior trees that support conditional reactive control. These efforts, and their parallels in the game industry, have led to more entertaining experiences, but they have not delivered on Laird and van Lent's vision. Façade remains one of the few environments that incorporates deeper forms of synthetic participants.

There remains both a need and an opportunity to develop richer, more human-like nonplayer characters for interactive games. These should infer human players' beliefs and goals, then take them into account in selecting their own actions, so they can better aid their allies and counter their opponents. They should communicate with human players in natural language dialogue, keeping track of earlier discussions and commitments, and they should develop joint plans that involve coordination, monitoring them for unexpected events that require replanning. This will let the agents cooperate with people on extended tasks of common interest and even form long-term relationships based on their interactions, thus providing a much fuller interactive experience.

6.3 Multifunctional Game Players

The General Game Playing competition (Genesereth, Love, & Pell, 2005) has encouraged progress toward general intelligence by providing a venue for demonstrating and evaluating systems that play a variety of games. The competition ensures generality, and sidesteps the knowledge-acquisition bottleneck, by providing entrants with the rules for novel games in a standard formalism. Rather than being fine tuned for one arena, like chess, systems must have the ability to handle any game in the language with reasonable effectiveness. However, the framework focuses on *playing* games, which is only one facet of game-related activity. In contrast, humans can use their domain knowledge in different ways, and we should develop interactive systems with the same versatility.

The field would benefit from an expanded competition which fosters development of systems that are not only general but that are also *multifunctional*. This would require that, given the rules for a particular game, entering systems should do more than simply play it effectively. They would also need to discuss details of their previous games with humans, provide running commentary on games played by others, analyze and converse about particular game situations, and even teach the games to human novices. Evaluating success on each function would be more difficult than counting wins and losses, but it would measure breadth of intellectual ability in addition to generality and domain independence. This would reward research on more versatile cognitive systems that reflect the multidimensional character of intelligence.

6.4 Synthetic Character Actors

Elsewhere (Langley, 2014), I have outlined challenge problems that could drive research on integrated cognitive systems. Some readers have found my proposals for building synthetic entertainers, attorneys, and politicians to be overly fanciful, so I will not repeat them here. However, they suggest another type of challenge problem – devising *synthetic character actors* – that would be more tractable and equally compelling. Computer-generated extras are already used widely in cinema, but these have simple reactive controllers and they lack higher level cognition. In contrast, a character actor must interpret a script and background story, take on the beliefs, goals, emotions, and personality of the role, and use this construct to breathe life into the lines that he or she recites.

Developing synthetic character actors would force the community to delve more deeply into these aspects of cognition in a constrained setting that is nevertheless highly social. The resulting systems would read and internalize brief but radically different scripts, with different background stories, as though auditioning for the parts. This would demonstrate their generality, much as in the General Game Playing competition. Evaluation would be more subjective, but this could be handled by a panel of judges that rates actors on expressiveness, believability, and the like, just as in many human competitions. Moreover, people would enjoy watching and critiquing the auditions, thus exposing cognitive systems technology, and demonstrating its potential, to the broader public.

6.5 Component-Level Research

As its name suggests, research on cognitive systems emphasizes the creation of integrated computational artifacts from existing components, but it does not rule out work on the elements themselves. Developing interactive intelligent agents of the types just described may also require advances at the component level, including:

- Improved representations for other agents' mental states, including not only their beliefs, goals, intentions, and knowledge, but also descriptions of their emotions, their personalities, and even their moral tenets;
- Better methods for reasoning about such models of other agents based on incomplete observations, using abductive inference mechanisms that introduce plausible default assumptions, rather than deductive techniques;
- Extended approaches to plan understanding that infer, from another agent's behavior, not only its beliefs, goals, and intentions, but also how that agent's decisions themselves take into account its own models of others' mental states;
- Along similar lines, novel methods for generating and executing plans that alter the beliefs, goals, and intentions of other agents, in some cases to help them achieve shared objectives but in others to take advantage of them through omission, misdirection, and even outright deception;
- Incorporating these advances into mechanisms for understanding and generating task-oriented dialogue, especially in situations that involve long-term interaction, which would benefit greatly from detailed models of others' mental states.

We can utilize these in the context of challenge problems like those outlined earlier, but we can also examine them in simpler settings like classic fables (e.g., Meehan, 1977; Pearce et al., 2014), which raise many of the same issues in idealized scenarios but which require subtantially less domain content and thus less knowledge engineering.

7. Concluding Remarks

In this paper, I reviewed the goals and assumptions of the cognitive systems paradigm, along with two widely accepted hypotheses about intelligence that underlie the movement. In addition, I presented a third claim, about the importance of social cognition, that motivates an interest in *interactive* cognitive systems. I reviewed three classic examples of such intelligent agents, along with their contributions and limitations, as well as three more recent instances that show interest in the topic remains active. None of these computational artifacts exhibited every facet of high-level cognition observed in humans, but each nevertheless made subsantial contributions to our understanding of these distinctive mental abilities. In addition, I proposed classes of challenge problems that could foster and guide research on social intelligence. These included deep conversational assistants that help users complete complex tasks, rich nonplayer game characters that become allies through extended interaction, generalized agents that use game knowledge in a variety of different ways, and synthetic character actors that adopt and utilize different personas on demand. Challenges of this sort could drive the research community in exciting new directions.

Critics of the cognitive systems movement – as I have defined it – and its historical antecedents often claim that its assumptions and methods have 'failed'. The examples reviewed in this paper offer clear evidence that the paradigm has succeeded in constructing sophisticated artifacts which reproduce key aspects of human intelligence, including ones that involve social cognition. Moreover, recent progress in the area indicates that the research community remains active and continues to advance our understanding of the mind. Statements that the paradigm has reached a dead end seem highly premature, and attempts to associate the 'cognitive systems' label with statistical techniques are simply inappropriate. Historically, both the pattern recognition and behaviorist movements have been antithetical to the study of high-level cognition (Miller, 2003; Langley, 2016), and there is little evidence their modern incarnations – 'deep' neural networks and reinforcement learning – offer as productive paths toward replicating social intelligence as the cognitive systems framework.

Acknowledgements

The analyses presented here were supported by Grant No. N00014-15-1-2517 from the Office of Naval Research, which is not responsible for its contents. I thank Alfredo Gabaldon, Ben Meadows, Chris Pearce, Ted Selker, and Jim Spohrer for useful discussions about interactive cognitive systems.

References

Allen, J. F., Miller, B. W., Ringger, E. K., & Sikorski, T. (1996). A robust system for natural spoken dialogue. *Proceedings of the Thirty-Fourth Annual Meeting of the Association for Computational Linguistics* (pp. 62–70). Santa Cruz, CA: ACL.

Bohus, D., & Horvitz, E. (2009). Models for multiparty engagement in open-world dialog. *Proceedings of the Tenth Annual SIGDIAL Meeting on Discourse and Dialogue*. London: Association for Computational Linguistics.

Brachman, R., & Lemnios, Z. (2002). DARPA's new cognitive systems vision. *Computing Research News*, *14*, 1.

Clark, H. H. (1996). *Using language*. Cambridge, UK: Cambridge University Press.

Christensen, H. I., Sloman, A., Kruijff, G-J., & Wyatt, J. (Eds.) (2010). *Cognitive systems*. Berlin: Spring-Verlag.

Gabaldon, A., Langley, P., & Meadows, B. (2014). Integrating meta-level and domain-level knowledge for task-oriented dialogue. *Advances in Cognitive Systems*, *3*, 201–219.

Genesereth, M., Love, N., & Pell, B. (2005). General game playing: Overview of the AAAI competition. *AI Magazine*, *26*, 63–72.

Graesser, A. C., VanLehn, K., Rose, C. P., Jordan, P. W., & Harter, D. (2001). Intelligent tutoring systems with conversational dialogue. *AI Magazine*, *22*, 39–51.

Hinrichs, T. R., & Forbus, K. D. (2014). X goes first: Teaching simple games through multimodal interaction. *Advances in Cognitive Systems*, *3*, 31–46.

Jones, R. M., Laird, J. E., Nielsen P. E., Coulter, K., Kenny, P., & Koss, F. (1999). Automated intelligent pilots for combat flight simulation. *AI Magazine*, *20*, 27–42.

Kirk, J. R., & Laird, J. E. (2014). Interactive task learning for simple games. *Advances in Cognitive Systems*, *3*, 13–30.

Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education*, *8*, 30–43.

Laird, J. E., & van Lent, M. (2001). Human-level AI's killer application: Interactive computer games. *AI Magazine*, *22*, 15–25.

Langley, P. (2012). The cognitive systems paradigm. *Advances in Cognitive Systems*, *1*, 3–13.

Langley, P. (2014). Four research challenges for cognitive systems. *Advances in Cognitive Systems*, *3*, 3–11.

Langley, P. (2016). The central role of cognition in learning. *Advances in Cognitive Systems*, *4*, 3–12.

Mateas, M., & Stern, A. (2005). Structuring content in the Façade interactive drama architecture. *Proceedings of Artificial Intelligence and Interactive Digital Entertainment* (pp. 93–98). Marina del Rey, CA: AAAI Press.

McShane, M., Beale, S., Nirenburg, S., Jarrel, B., & Fantry, G. (2012). Inconsistency as a diagnostic tool in a society of intelligent agents. *Artificial Intelligence in Medicine*, *55*, 137–148.

Meehan, J. R. (1977). TALE-SPIN: An interactive program that writes stories. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence* (pp. 91–98). Cambridge, MA: Morgan Kaufmann.

Miller, G. A. (2003). The cognitive revolution: A historical perspective. *TRENDS in Cognitive Sciences*, *7*, 141–144.

Newell, A., & Simon, H. A. (1976). Computer science as empirical enquiry: Symbols and search. *Communications of the ACM*, *19*, 113–126.

Pearce, C., Meadows, B., Langley, P., & Barley, M. (2014). Social planning: Achieving goals by altering others' mental states. *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence* (pp. 402–408). Quebec City, Quebec: AAAI Press.

Rich, C., Sidner, C. L., & Lesh, N. (2001). COLLAGEN: Applying collaborative discourse theory to human-computer interaction. *AI Magazine*, *22*, 15–26.

Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, *61*, 85–117.

Swartout, W. R., Gratch, J., Hill, R. W., Hovy, E., Marsella, S., Rickel, J., & Traum, D. (2006). Toward virtual humans. *AI Magazine*, *27*, 96–108.

Talamadupula, K., Briggs, G., Chakraborti, T., Scheutz, M., & Kambhampati, S. (2014). Coordination in human-robot teams using mental modeling and plan recognition. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 2957–2962). Chicago: IEEE Press.

Tesauro, G. (1995). Temporal difference learning and TD-Gammon. *Communications of the ACM*, *38*, 58–68.

Thompson, C. A., Göker, M. H., & Langley, P. (2004). A personalized system for conversational recommendations. *Journal of Artificial Intelligence Research*, *21*, 393–428.

Turing, A. (1950). Computing machinery and intelligence. *Mind*, *59*, 433–60.

Winograd, T. (1972). Understanding natural language. *Cognitive Psychology*, *3*, 1–191.