Progress and Challenges in Interactive Cognitive Systems

Pat Langley

Institute for the Study of Learning and Expertise
Palo Alto, California

Thanks to Paul Bello, Ron Brachman, Ken Forbus, John Laird, Allen Newell, Paul Rosenbloom, and Herbert Simon for discussions that helped refine ideas in this talk.

The Cognitive Systems Paradigm

The Vision of Artificial Intelligence

The field of artificial intelligence was launched in 1956 at the Dartmouth meeting; its audacious aims were to:

- Understand the mind in computational terms;
- Reproduce all mental abilities in computational artifacts.

This view continued through the mid-1980s, but recent years have seen adoption of very different goals.

Most AI researchers are now content to work on narrowly defined tasks that involve little intelligence.

In fact, many have forgotten the essential difference between AI and pattern recognition.

The Cognitive Revolution

During the 1950s/1960s, breakthroughs in both AI and cognitive psychology resulted from:

- Rejecting behaviorists' obsession with learning on simple tasks and information theory's focus on statistics;
- Studying problem solving, language understanding, and other tasks that involve *thinking* (i.e., *cognition*);
- Emphasizing the role of *mental structures* in supporting such complex behaviors.

Yet many modern AI researchers have abandoned the insights of the cognitive revolution.

Why have so many retreated from the field's initial aspirations?

Reasons for the Shift

This change in AI’s focus has occurred for a number of reasons, including:

- Commercial successes of ‘niche’ AI
  - Encouraging focus on narrow problems
- Faster processors and larger memories
  - Favoring blind search and statistical schemes
- Obsession with quantitative metrics
  - Encouraging mindless ‘bake offs’
- Formalist trends imported from computer science
  - Favoring simple tasks with optimality guarantees

Together, these have drawn many researchers’ attention away from AI’s original vision.

The Cognitive Systems Movement

Yet many of the original challenges still remain and offer many opportunities for research.

Because "AI" has altered its meaning, we will define *cognitive systems* (Brachman & Lemnois, 2002) as the field that:

- *Designs, constructs, and studies computational artifacts that exhibit complex, human-like behavior over the full range of activities we regard as intelligent.*

We can distinguish this paradigm from most current AI work by six major characteristics.

See *Advances in Cognitive Systems* (http://www.cogsys.org/).

Feature 1: Focus on High-Level Cognition

One distinctive feature of the cognitive systems movement is its emphasis on *high-level cognition*.

Dogs and cats can recognize objects, execute routine skills, and learn empirically, but only humans can:

- Understand and generate language
- Solve novel and complex problems
- Design and use complex artifacts
- Reason about others' mental states
- Think about their own thinking

Computational replication of these abilities is a core charge of cognitive systems research.

Feature 2: Symbol Structures

Another key aspect of cognitive systems research is its reliance on *symbol structures*, including stored *knowledge*.

The insight behind the 1950s AI revolution was that computers are not mere number crunchers.

Computers and humans are *general symbol manipulators* that:

- Encode content as list structures or similar formalisms
- Create, modify, and interpret this structured content
- Use numbers mainly as annotations on these structures

The paradigm assumes that representing, and reasoning over, rich symbolic structures is key to human-level cognition.

Feature 3: Systems Perspective

Research in the paradigm is also distinguished by approaching intelligence from a *systems perspective*.

While most AI efforts idolize component algorithms, work on cognitive systems is concerned with:

- How different intellectual abilities fit together and interact
- Integrated intelligent agents that combine these capabilities
- Cognitive architectures that offer unified theories of mind

Such systems-level research provides an avenue to artifacts that exhibit the breadth and scope of human intelligence.

Otherwise, we will be limited to the *idiot savants* so popular in academia and industry.

Feature 4: Influence of Human Cognition

Research on cognitive systems also draws ideas and inspiration from findings about *human cognition*.

Many of AI's early insights came from studying human problem solving, reasoning, and language use, including:

- How people represent knowledge, goals, and beliefs
- How humans draw inferences and achieve goals
- How people acquire new structures from experience

We still have much to gain from this strategy, even when our artifacts differ in their operational details.

Human capabilities also offer *challenges* for cognitive systems researchers to pursue.

Feature 5: Heuristics and Satisficing

Another important assumption of cognitive systems work is that intelligence relies on *heuristic methods* that:

- Are not guaranteed to find the best or even *any* solution but
- Greatly reduce search and make problem solving tractable
- Apply to a broader range of tasks than methods with guarantees

They mimic high-level human cognition in that they *satisfice* by finding acceptable rather than optimal solutions.

Much of the flexibility in human intelligence comes from its use of heuristic methods.

Feature 6: Exploratory Research

Cognitive systems research also differs from mainstream AI is that it encourages *exploratory research*, such as:

- Demonstrations of entirely new functionality
- Novel approaches to well-established problems
- Analyses of challenging cognitive tasks
- Architectures and frameworks for integrated intelligence

Studies must still make clear claims and support them, but many forms of evidence are possible.

Competitions and ‘bake offs’ encourage incremental progress and premature convergence on narrow problems.

Three Hypotheses for Cognitive Systems

Newell and Simon (1976) proposed two hypotheses that underlie most work on cognitive systems:

- The ability to encode, manipulate, and interpret *symbol structures* offers the means for general intelligent action.
- *Problem solving* involves *heuristic search* through a space of states (symbol structures) generated by mental operators.

We offer a third claim – the *social cognition hypothesis* – that deals with interactive behavior:

- Intelligence requires the ability to represent, reason over, and use *models of other agents' mental states*.

Humans are inherently social animals, and many key cognitive faculties involve thinking about others.

Examples of Interactive Cognitive Systems

SHRDLU (1970)

SHRDLU (1970) was an early AI system that answered questions in English and executed complex commands.

Although limited in scope, it reproduced many aspects of human dialogue.

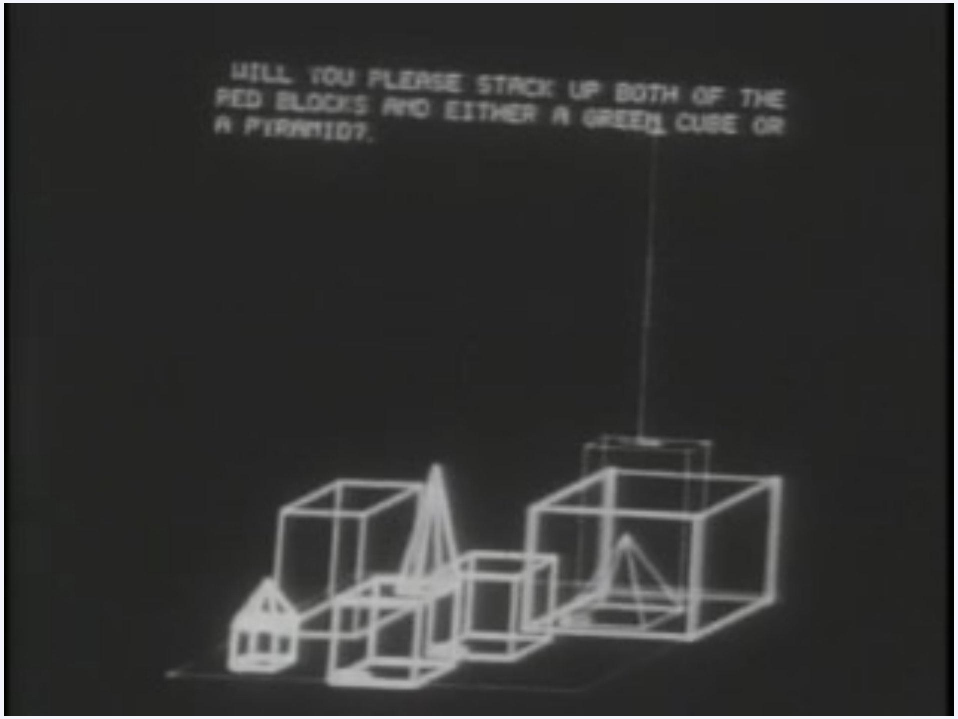

Person: Pick up a big red block.
Computer: *OK.*
Person: Grasp the pyramid.
Computer: *I don't understand which pyramid you mean.*
Person (changing his mind): Find a block which is taller than the one you are holding and put it into the box.
Computer: *OK.*
Person: What does the box contain?
Computer: *The blue pyramid and the blue block.*
Person: What is the pyramid supported by?
Computer: *The box.*
Person: How many blocks are not in the box?
Computer: *Four of them.*
Person: Is at least one of them narrower than the one which I told you to pick up?
Computer: *Yes, the red cube.*

Carnegie Learning's Algebra Tutor (1999)

This tutor encodes knowledge about algebra as production rules, infers models of students' knowledge, and provides personalized instruction.

The system has been adopted by hundreds of US middle schools.

Studies have shown that it improves student learning in this domain by 75 percent.

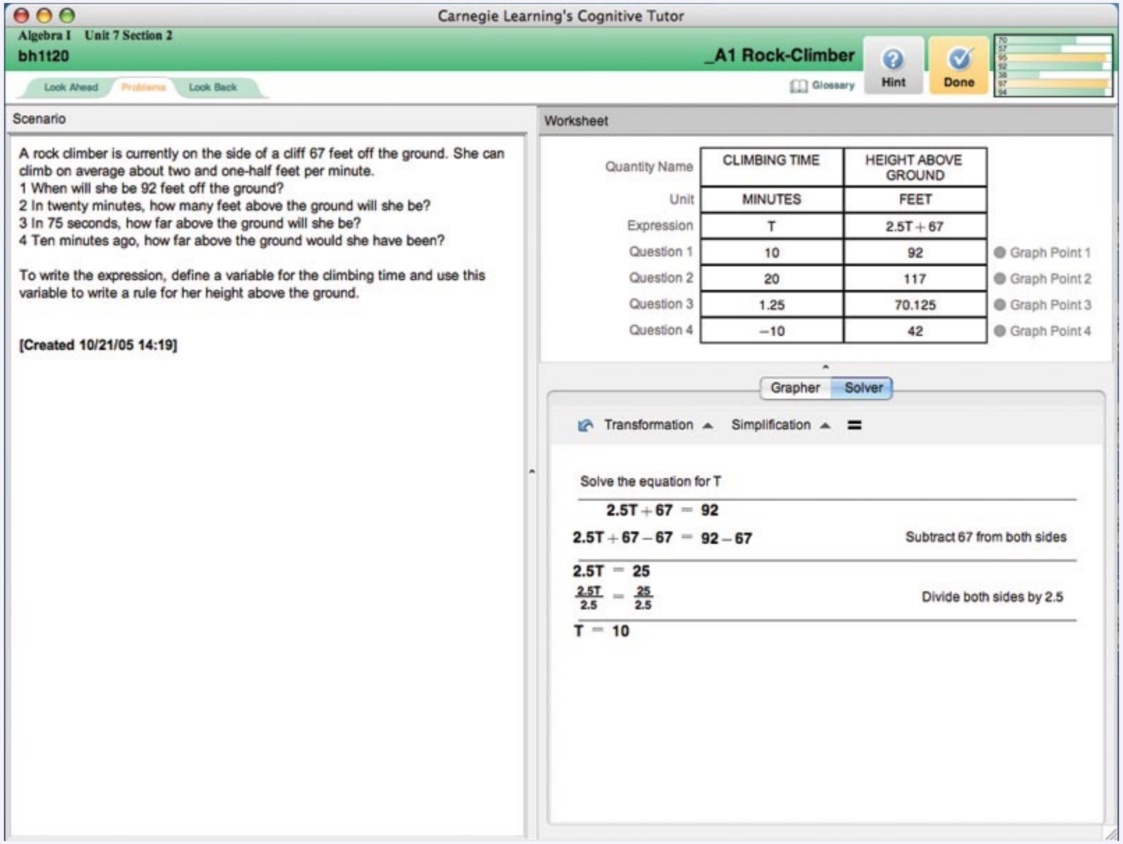

TacAir-Soar (1997)

The TacAir-Soar system reproduces pilot behavior in tactical air combat.

It combines abilities for spatio-temporal reasoning, plan generation / recognition, language, and coordination.

The system flew 722 missions during the STOW-97 simulated training exercise.

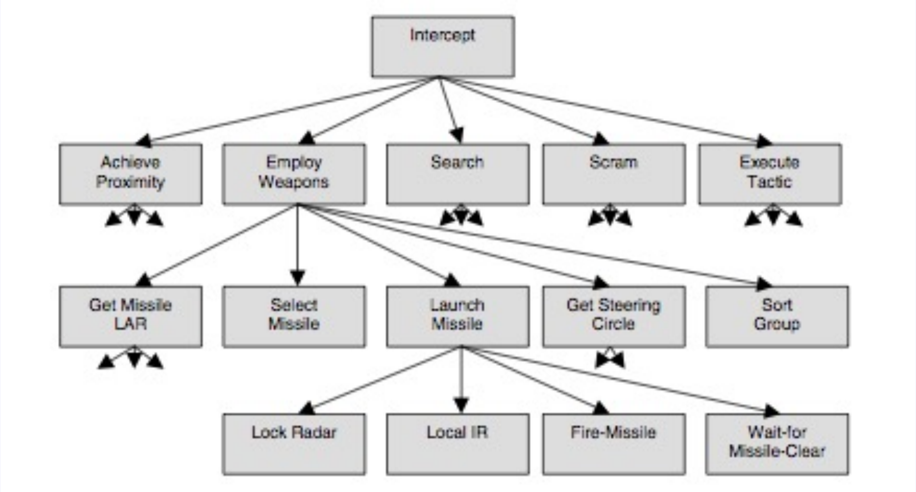

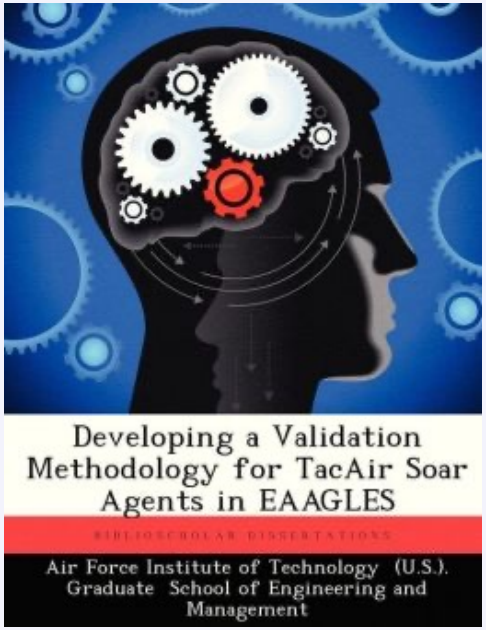

Façade (2003–2007)

Mateas and Stern's Façade is a graphical environment in which characters interact with the user and each other.

The agents understand and generate sentences, control gaze and expression, and they exhibit distinct personalities.

Façade characters use a rich knowledge base to produce inferences, carry out physical activities, and engage socially.

A Conversational Architecture

Gabaldon, Langley, and Meadows (2014) describe DIGA, an architecture for conversational agents that:

- Cleanly separates domain-level from dialogue-level content
- Distinguishes conceptual knowledge from goal-oriented skills

The architecture operates in discrete cycles, during which it:

- Observes new speech acts, including ones it generates itself
- Uses inference to update its model of other agent's beliefs and goals

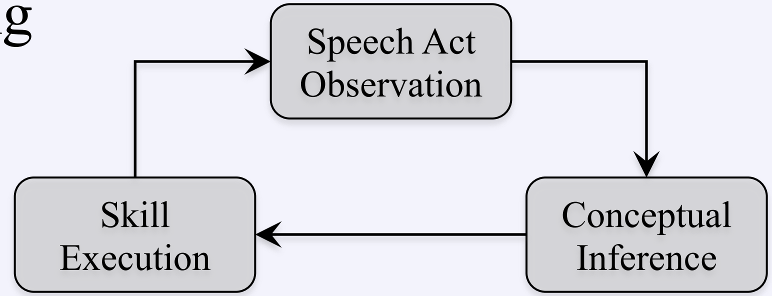

- Executes skills to produce utterances based on this model

At a high level, it operates in a manner similar to production-system architectures like Soar and ACT-R.

Sample Dialogue for DIGA Medic Assistant

M: We have a man injured!

A: Ok. What type of injury?

M: He's bleeding.

A: How bad is the bleeding?

M: Pretty bad. I think it is the artery.

A: Ok. Where is the injury?

M: It's on the left leg.

A: Apply pressure on the leg's pressure point.

M: Roger that.

A: Has the bleeding stopped?

M: No. He's still bleeding.

A: Ok. Apply a tourniquet.

M: Where do I put the tourniquet?

A: Just below the joint above the wound.

M: Ok. The bleeding stopped.

A: Good job.

M: human medic A: advisor

DIGA incrementally updates its model of the *medic's beliefs and goals*, which it uses to generate utterance content.

Interactive Task Learning for Games

Hinrichs and Forbus (2014) describe a system that learns concepts and rules for games from instructions and sketches, after which it can play the games legally but poorly.

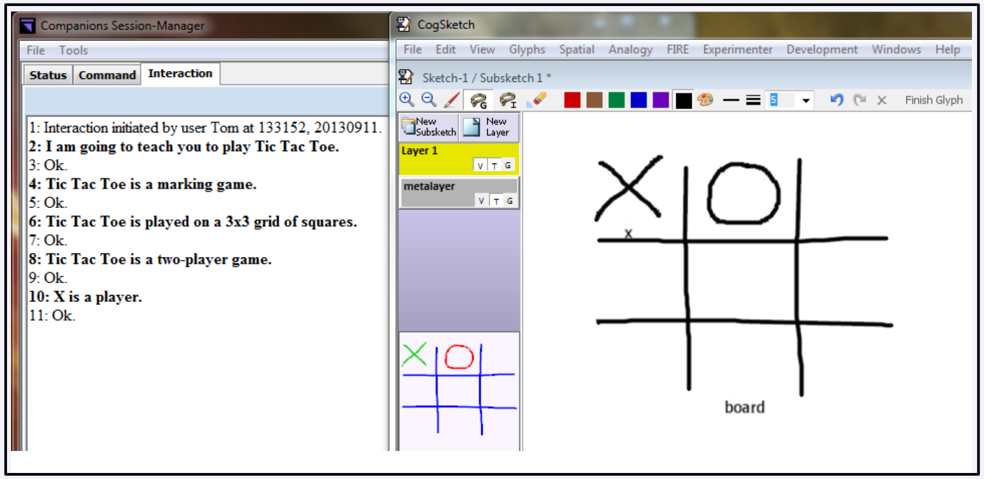

Game	Spatial Concepts	Actions	Goal	Failure
Tic-Tac-Toe	on, under, linear	place	3-in-a-row	
Connect-3	on, under, linear	stack-place	3-in-a-row	
Tower of Hanoi	on, under, smaller	smaller-stack	stacked	
Five Puzzle	on, under, near, diagonal	slide	matching-location	
Frogs and Toads	left, right, on, under	slide-l, slide-r, jump-l, jump-r	side-swap	
Four Queens	on, under, linear	place	all-placed	no-attack
Blocks World	on, under	stack	order-stacked	
Sokoban	on, under, linear, diagonal	push, slide	blocks-in	
Peg Solitaire	on, under, linear	jump-remove	one-left	
Knight’s Tour	on, under, aligned-vert, aligned-horiz	knight-a, knight-b	all-placed	
River Crossing	left, right, aligned	move-l, move-r, carry-l, carry-r	right-bank	fox-goose, goose-beans

Kirk and Laird (2016) report a similar system that learns to play 17 different games from visual demonstrations and instructions, asking questions when necessary. Learning is very rapid, as in humans.

Rapid Learning of Interaction Skills

Frasca et al. (2018) describe DIARC, an architecture that learns *interaction* skills (e.g., passing a knife) from both spoken language instructions and demonstrations.

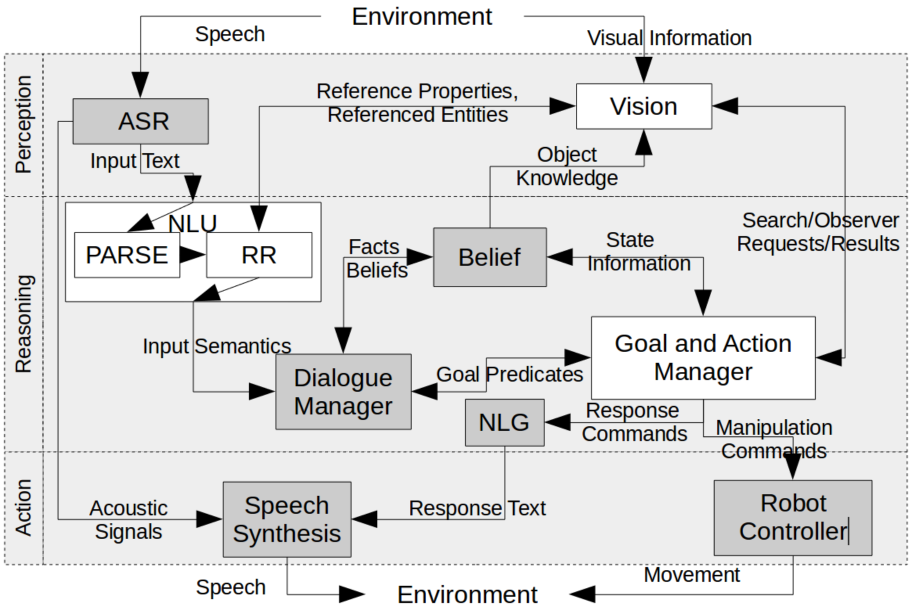

The integrated system acquires complex interaction skills from *single* demonstrations, learning about not only its role, but the roles of other participants.

Some Other Examples

Other researchers have also developed cognitive systems with interactive abilities, including:

- COLLAGEN (Rich et al. 2001), which helps users in operating complex devices, asking questions and giving advice as needed
- Tutorial dialogue systems (Graesser et al., 2001) that converse in spoken language, giving personalized instruction.
- The Virtual Humans project (Swartout et al., 2006), which has created many synthetic characters that interact with users.
- The Artificial Receptionist (Bohus & Horvitz, 2009), which welcomes and interacts with visitors in spoken dialogue.
- CWMS, a collaborative problem solver that helps its users analyze situations and generate plans via spoken dialogue (Allen et al., 2018).

These diverse systems show the range of possible applications.

Research Challenges for Interactive Cognitive Systems

Guidelines for Challenge Problems

New problems can foster progress in any area, and productive challenges for interactive cognitive systems should:

- Focus on tasks that require *high-level* cognition
- Benefit from *structured* representations and knowledge
- Require *system-level* integration of capabilities
- Have *human* role models that offer insights
- Be complex enough to need *heuristic* approaches
- Depend centrally on processing *social* structures

They must also move beyond the Turing test by emphasizing *goal-oriented* behavior.

Deep Conversational Assistants

Spoken-language dialogue is the natural mode for providing aid on tasks like driving, cooking, and shopping.

Compared to humans, systems like Siri are primitive, and we need more effective conversational assistants that:

- Carry out extended dialogues about goal-directed activities
- Take into account the surrounding task context
- Infer *common ground* (Clark, 1996) for joint beliefs / goals
- Store and utilize previous interactions with the user

These would carry out deep language processing, reason about others' mental states, and help them achieve their goals.

Rich Nonplayer Game Characters

Synthetic characters are rampant in today's computer games, but they are typically shallow.

We should develop more compelling nonplayer characters that:

- Infer other players' goals and use them toward their own ends
- Interact with human players in constrained natural language
- Cooperate with them on extended tasks of common interest
- Form long-term relationships based on previous interactions

Such agents would generate much richer and more enjoyable experiences for human players.

For this purpose, they must reason about others' mental states.

A Truly General Gamer

Humans use their domain knowledge in different ways, and we need multifunctional systems with the same versatility.

One example might be a system that, given knowledge about a class of games, can:

- Play that class of game in competitions
- Discuss previous games with other players
- Provide commentary on games played by others
- Analyze and discuss particular game situations
- Teach the game to a human novice

This should demonstrate breadth of intellectual ability but avoid the knowledge acquisition bottleneck.

A Synthetic Character Actor

Our society devotes far more attention to its movie stars than to scientists and scholars.

Imagine a synthetic *character actor* with general acting skills and the ability to:

- Read scripts / background stories for very different parts
- Adopt beliefs, goals, emotions and personality for the role
- Audition for the part, breathing life into the lines

Most scenes would involve interaction with other actors, and thus require social cognition.

Requiring the system to take on radically different characters would test its generality.

Some Necessary Components

Cognitive systems involve integration, but we also need research on core abilities for social cognition:

- *Representing* other agents' mental states
- *Reasoning* flexibly about others' beliefs and goals
- Social *plan understanding* from others' observed behavior
- Social *plan generation* to manipulate others' actions
- Understanding and planning in task-oriented *dialogue*
- Cognitive accounts of *emotion*, *morals*, and *personality*

Human-level cognitive systems must exhibit all these capacities, and we need research on each topic.

Summary Remarks

In this talk, I discussed the cognitive systems paradigm, which pursues AI’s original vision, by:

- Stating six distinctive features of research in this area
- Reviewing three hypotheses about intelligent behavior
- Presenting examples of interactive cognitive systems
- Posing challenge tasks for interactive cognitive systems

Research in this emerging field retains the audacity of early AI and promises to keep us occupied for years to come.

Readings on Interactive Cognitive Systems

- Frasca, T., Oosterveld, B., Krause, E., & Scheutz, M. (2018). One-shot interaction learning from natural language instruction and demonstration. *Advances in Cognitive Systems*, *6*, 159–176.
- 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.
- Langley, P. (2012). The cognitive systems paradigm. *Advances in Cognitive Systems*, *1*, 3–13.
- Langley, P. (2012). Intelligent behavior in humans and machines. *Advances in Cognitive Systems*, *2*, 3–12.
- Mateas, M., & Stern, A. (2005). Structuring content in the Façade interactive drama architecture. *Proceedings of Artificial Intelligence and Interactive Digital Entertainment*. Marina del Rey, CA: AAAI Press.
- 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.

Also see *Advances in Cognitive Systems* (http://www.cogsys.org/).