

# The Cognitive Systems Paradigm

**Pat Langley**

Computer Science and Engineering  
Arizona State University  
Tempe, Arizona, USA

Thanks to Paul Bello, Ron Brachman, Nicholas Cassimattis, Ken Forbus, John Laird and others for discussions that helped refine the ideas in this talk.

# The Original AI Vision

The early days of artificial intelligence research were guided by common vision:

- *Understanding and reproducing, in computational systems, the full range of intelligent behavior observed in humans.*

This paradigm was adopted widely from the field's founding in the 1950s through the 1980s.

However, the past 20 years have seen a very different AI emerge that has largely abandoned these initial goals.

Why have most researchers stepped back from the discipline's original aspirations?

# Why Has AI Gone Astray?

We can track this sea change in the AI community to a number of important factors:

- Increased computer speed and storage has aided simple-minded CPU-intensive and memory-based approaches;
- Emphasis on quantitative performance metrics has encouraged incremental progress on standardized problems;
- Influence of mathematics has led to “theorem envy” and to an optimality obsession, encouraging a focus on simple tasks;
- Commercial success on narrowly defined problems has fostered research on similarly limited tasks.

Taken together, these trends have transformed AI into a field that has adopted greatly restricted goals.

# Why Has AI Gone Astray?

Maslow (1966) postulates some other reasons why a scientific field can become narrow and conservative:

... these “good”, “nice” scientific words – prediction, control, rigor, certainty, exactness, preciseness, neatness, ..., quantification, proof, ... – are all capable of being pathologized when pushed to the extreme. [They] may be pressed into the service of safety needs [to] become ... anxiety-avoiding ... mechanisms ... for detoxifying a ... frightening world as well as ways of ... understanding a fascinating ... world.

But Maslow notes that science need not proceed in this way:

... healthy scientists [can] enjoy not only the beauties of precision but also the pleasures of sloppiness, casualness and ambiguity... They are not afraid of hunches, intuitions, or improbable ideas... All of this is exemplified in the greater versatility of the great scientist, of the creative, courageous, and bold scientists.

# Cognitive Systems

The field's original challenges of still remain and provide many opportunities for research.

However, because “AI” has become associated with such limited aspirations, we need a new label.

We will use *cognitive systems*, a term coined by Brachman and Lemnios (2002), to refer to the discipline that:

- *designs, constructs, and studies computational artifacts that exhibit the full range of human intelligence.*

We can further distinguish this paradigm from what has become mainstream AI by describing its key characteristics.

# Feature 1: Focus on High-Level Cognition

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

People share basic capabilities for categorization and empirical learning with dogs and cats, 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 the key charge of cognitive systems research.

## Feature 2: Structured Representations

Another distinctive aspect of cognitive systems research concerns its reliance on *structured representations*.

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

Computers and humans are *general symbol manipulators* that:

- Encode information as list structures or similar formalisms
- Create, modify, and interpret this relational content
- Incorporate numbers only as annotations on these structures

Our paradigm assumes that *physical symbol systems* (Newell & Simon, 1976) of this sort are key to human-level cognition.

## Feature 3: Systems Perspective

Research in our 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 interact and fit together
- Cognitive architectures that offer unified theories of mind
- Integrated intelligent agents that combine capabilities

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

Otherwise, we will remain limited to the *idiot savants* developed by the mainstream AI community.

## Feature 4: Influence of Human Cognition

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

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

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

We still have much to gain by following this strategy, even when an artifact's operation differs in its details.

Moreover, human capabilities provide *challenges* for cognitive systems researchers to pursue.

## Feature 5: Exploratory Research

Cognitive systems research also differs from mainstream AI in its approach to *evaluation*.

Although quantitative experiments remain welcome, our paradigm also encourages:

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

These evaluation styles encourage *exploratory research*, which is appropriate given how little we understand about the mind.

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

# Fostering the Movement

We can encourage cognitive systems research through the usual types of academic activities:

- Organizing an annual refereed conference
- Publishing a refereed, archival journal
- Hosting invited symposia on related topics
- Teaching courses and tutorials in the area
- Holding summer schools to train new researchers
- Providing readings and Web resources to the community

When combined, these activities will raise cognitive systems to a visible and vital intellectual movement.

Working together, we can create a new *Zeitgeist* that recaptures the spirit of AI's founders.

# Plans for a Conference and Journal

We plan to follow up this initial meeting two specific activities:

- a new *Conference on Advances in Cognitive Systems* that:
  - meets annually, probably sometime in the fall
  - may be colocated with other AI-related events
  - includes a careful but enlightened review process
- an electronic journal, *Advances in Cognitive Systems*, that:
  - publishes papers accepted for conference presentation
  - also welcomes submissions outside conference season
  - will have initial volume with essays that define the field and selected papers from this meeting

Together, the conference and journal will give us a venue to exchange ideas and a place to publish research results.

# Creating a Research Agenda

Our new community must also identify challenges that can drive research; some natural candidates include:

- Mechanisms for flexible, scalable inference
- Flexible problem solving and formulation
- Deep processing of language and dialogue
- Models of emotion and moral cognition
- Reasoning about others' mental states
- Metacognitive reasoning systems

We must also work with program managers to secure funding for innovative research on these topics.

# The Road Ahead

Although cognitive systems adopts the original aims of AI, its modern incarnation is relatively new.

To ensure its success as a scientific discipline, we must:

- Clarify and defend its distinctive characteristics
- Create a community of broad-minded researchers
- Identify research challenges and make progress on them
- Establish venues for communication and publication
- Recruit, train, and place promising new researchers
- Never abandon the audacious goals we have set ourselves

The current meeting is only the first step on the road toward a broader and deeper understanding of high-level cognition.

End of Presentation