Mission Statement

John Laird and Robert Wray
Center for Integrated Cognition

The mission of the **Center for Integrated Cognition** (CIC) is to be a world leader in organizing, supporting, and pursuing basic research on machines that think and learn the way people think and learn and that easily and naturally interact and work with people. To do this, we investigate and create end-to-end AI systems that sense and act; reason; plan; represent, remember, use, and create knowledge; interact with language; and learn robustly and flexibly about their experiences and interactions. The CIC approach leverages decades of research on *cognitive architecture* – the computational *components* underlying intelligence and how they are *integrated*.

A hallmark of human intelligence is that it is broad, flexible, and deep. Humans have a complex internal mental life that includes hierarchical and relational models of the environment, deliberate problem solving, anticipation, planning, analogy, and many types of reasoning (spatial, temporal, visual). We don't just sense patterns and react to them - we *think!* We can reason about the thoughts of others (theory of mind), use knowledge of our past and facts about the world, and communicate via gestures and natural language. We are continually learning, and we learn extremely fast from our own experience, but also from others – all within a single integrated "mind."

CIC research concentrates on AI research informed and inspired by the underlying principles and computational mechanisms that humans use to interact with and learn from the world around us, with a special emphasis on working with and learning from people. Our focus is building minds – creating systems that reason; plan; represent, remember, use, and create knowledge; interact with language; learn robustly and flexibly about their experiences and interactions, ...

Why Think, Interact, and Learn Like People?

The human mind is an existence proof of integrated intelligence. That alone is reason to use it for inspiration in AI research, but there are many more.

1. The human mind is still one of the great mysteries of science. Only through the development of comprehensive, computational models of the processes and structures underlying cognition and its connection to perception and action will we have a complete understanding of it. CIC research integrates research across the fields of cognitive science, contributing to new scientific understanding and a Common Model of Cognition (Laird, Lebiere, & Rosenbloom, 2017).

2. Humans are unique in our abilities to work together and learn from each other. We have unique capabilities for establishing and maintaining common ground and joint intentions that are missing in animals, including primates (and machines). Our hypothesis is that the best path to creating truly collaborative machines is to replicate the capabilities of humans, and in ways comparable to how humans achieve those capabilities. Machines that think like people will be easier for us to understand, internally model, predict, and ultimately trust. The human capabilities that support collaboration are not limited to a specific mental module or isolated skills (although some specific skills such as gaze tracking appear to be critical). Instead, we combine many disparate cognitive capabilities such as communicating through language; modeling other humans' goals, beliefs, and capabilities; reasoning about space, time, and objects; and using and learning

multiple forms of knowledge. This collection of capabilities supports not only our interaction skills, but also our ethical and moral frameworks (Tomasello, 2019).

3. The human mind is still unparalleled in its ability to work across wide ranges of different tasks, learn new tasks very quickly, and make progress even in novel situations. Current approaches to Machine Learning (ML)/Deep Learning (DL) are limited to specific tasks, requiring huge training corpora, and must be prepared by hand to work on those specific problems. If the problem changes just a little, they must be retrained, and they are limited in transferring what they have learned to new tasks. Our goal is to pursue systems that have the flexibility, robustness, and adaptability of humans. There is no question that there are many tasks for which current ML/DL approaches are appropriate, and we will use such specialized systems as *tools* within our AI systems, but we are targeting cognitive capabilities that go far beyond those approaches.

Why Work with People?

The future of AI will involve working with, helping, teaching and learning from people. Not only are collaboration and coordination going to be critical capabilities to achieve, attempting to achieve them will drive research in many aspects of AI. For example, collaboration is one of the most challenging problems for any intelligent system, requiring integration of all of its perceptual and cognitive capabilities. There are also strong arguments that collaboration itself is a core aspect of cognition that is necessary for developing any truly intelligent machine. Individual cognitive capabilities provide a starting point, but human inherent social nature drives us to be learning (and teaching) machines, which in turn is responsible for the development of language, technology, and culture – achievements that are impossible for individual intelligent systems. The research of the CIC not only attempts to understand the fundamental computational structures processes that enable and support social behavior, but also how to create AI systems that productively utilize them to interact with, learn from, and teach people.

Why Now and Why an independent research center?

There have been significant advances in the past decade in the exploration of mind-like computational systems. Researchers now routinely interact with their systems by talking with them, gesturing, sketching, etc. Systems learn not just to perform a task better, which has been a long-time hallmark of integrated cognitive systems, but to remember and to learn from past experiences, to build models of the world from experience and instruction and reading, and even to learn to modulate their own cognitive performance. However, this progress has been attenuated in two important ways. First, the tidal wave of function-based learning has drowned out the successes of integrated cognitive systems approaches. Many members of the AI community are not aware of recent advances in integrated cognitive systems; the public clearly is not. The lack of awareness has led to a chilling if not actual winter for funding and recruitment of the next generation of cognitive AI scientists. More disturbingly, is has led the majority of researchers in AI to abandon problems that require the generality, flexibility, and adaptability possible with cognitive AI.

A second factor is that the field of integrated cognitive AI, like other integration sciences, is fragmented across disciplines. This fragmentation reduces the visibility of the research (cognitive AI research is published in AI, cognitive psychology, linguistic, and application/domain-specific conferences), curtails the sharing of infrastructure (software, data, methods) essential to successful research programs, and increases the difficulty of cohesive community perspective, concepts, and action (different jargons, different departmental priorities, etc.).

The CIC's mission and approach is distinct from academic and industry research and is intended to mitigate these issues and fill key gaps in AI research. The CIC is focused on research of wholistic or end-to-end AI systems and the various interactions and tradeoffs that one encounters in creating such systems. Academic AI research broadly tends to specialize on algorithms and components, rather than this systems-level perspective. Industry use of AI also contrasts with the mission because the typical goal is to develop a solution for a specific set of application requirements. This focus is true for both commercial research and in military investment at the applied research and advanced development stages. CIC research is focusing on understanding and evaluating AI in the context of more complete systems but taking a systematic approach to the design space rather than driving toward point solutions for particular applications.

The CIC seeks to mitigate these factors and to extend and amplify the impact of cognitive AI. Today, there are isolated islands of research at universities, where individual faculty are working on different cognitive architectures, without common benchmarks or shared methodologies. The CIC will not only take a lead in cognitive AI research, but it will also develop and nurture a community of integrated cognitive AI researchers via four interrelated thrusts:

1. **CIC Research:** The CIC will have a robust and active internal research program, with the goal, together with its external research partners, of being the leader in Cognitive AI research.
2. **Research Community Integration and Synthesis**: The CIC will create programs and initiatives designed to develop a community of researchers that share software, problems, environments, and evaluation. While specific programs will be chosen and detailed outside of the mission statement, examples of activities that fall under this thrust are likely to include:
    - Hosting researcher sabbaticals at the CIC for periods of sustained interaction;
    - Research and development of standards and consensus artifacts that facilitate common understanding (e.g., a "gene ontology" for integrated cognitive AI);
    - Representing the goals and progress of the cognitive AI community to funders and advocating for robust investment in cognitive AI; and
    - Creation of an academic workshop focused on integrated cognitive AI.
3. **Educational Outreach**: The CIC will seek to provide educational resources to support recruitment into the field, mentoring, conceptual maturity of the research goals and progress, and improved understanding of integrated cognitive AI in other disciplines, as well as the public at large. Examples of activities that fall under this thrust could include:
    - Creation of tutorials and workshops ("summer schools") designed to introduce the field to new researchers;
    - Targeted outreach to policy and governmental initiatives to offer perspective and encourage the inclusion of cognitive AI in requirements analysis and investment and program planning;
    - Production capabilities and assistance to enable researchers to communicate research results to the larger public, via both journalistic methods (press releases) and direct publication (YouTube, twitter videos, etc.); and
    - Preparation of instructional units (instructional videos, games and exercises with supporting software, etc.) to enable the introduction of integrated AI concepts in undergraduate and secondary school curricula.

4. **Research Infrastructure**: Since its inception, researchers in the field have recommended the development and maintenance of shared research tools that could speed research and lower the cost of entry for integrated cognitive AI research. The CIC will provide curated and maintained resources for sharing and distributing software, data, and benchmark tasks. Part of the role of professional software developers at the CIC will be to adopt, mature, and maintain various infrastructure. One of the significant benefits of this function of the CIC will be to significantly lower the cost of entry to new researchers in the field. Examples of software resources that the CIC could maintain:
    - Implementations of real and simulated task domains (simulations; interfaces to game engines, mobile-device operating systems, and robotic platforms);
    - Implementations of specific challenge problems and/or benchmark tasks within a specific task domain; and
    - Data warehousing to support comparison of results (over time and against various methods) and accumulation of research results.

The breadth and depth of these activities highlights the important need for an organizational structure beyond an individual researcher. All of these activities can enhance the impact of the research in the field, both helping to speed individual research progress (more people spending more time on research rather than infrastructure) and to compound the impact of individual contributions through greater commonality of purpose and directly building on the advances of others.

Another significant role and benefit of the CIC is continuity. Even the lengthiest awards include only a small number of researchers in the field, and the limited term of the award means researchers must focus on short-term research payoffs. An independent CIC with a clear research mission can take a longer-term perspective. It can invest in important elements of community building to make a sustained, community-wide impact.

Center Organization, Activities, and Personnel

The exact size, organization, and activities of the CIC will be determined by available resources. The CIC will be supported by donations, grants, and contracts for basic research, from individuals, foundations, industry, and government entities.

Activities:

1. Internal Research: The CIC will support a substantial internal research activity on Cognitive AI. This will be carried out by CIC Scientists and CIC Engineers, and coordinated with external CIC projects. We expect there to be 4-6 major projects at a time, depending on available funding.

2. Cognitive AI evaluation development and support: The CIC, together with its research partners will develop challenge problems, test sets, benchmarks, etc. that are relevant to cognitive AI systems, and administer these tests to competing systems. Although there has been explosion in such materials in AI/ML (e.g., OpenAI Gym), the tools and expectations are narrow and do not require the breadth of flexibility and competence across domains that are characteristic of human cognition. We expect this project could be one of the major research thrusts of the CIC.

4. Cognitive AI software support: The CIC will support hosting, maintaining, and distributing cognitive-AI software. This will include software developed at the CIC, but also by external researchers (if they desire). Current systems generally do not have consistent or quality support,

making it challenging for new researchers to build on existing work. The CIC will serve an important role of ensuring these software systems are robust and easy to access, install, and use. Example types of software are cognitive architectures (e.g., Soar, Sigma, ACT-R, Companions), API/middleware for connecting cognitive systems to game engines and robotic platforms, and simulation environments. In general, software developed by the CIC will be open source.

5. Community building and support: The CIC will support hosting, maintaining, and distributing research results for cognitive AI, including research papers, presentations, and wikis.

6. Educational outreach: The CIC will support development of instructional and tutorial materials on the software systems that are supported and distributed. It will also provide onsite tutorials, and potentially summer schools (with student support) on the general research area of cognitive AI and on the specific projects being supported by the CIC.

7. Industry engagement: Long-term, the CIC will explore and likely create an industry consortium group. Members will be provided early access to CIC software and research results. There will also be opportunities for consortium members to have early access to tutorials and other educational materials, and to participate in all CIC activities.

References

Darwiche, A. (2018). Human-Level Intelligence or Animal-Like Abilities? *Communications of the ACM, 61*(Oct), 56-67.

Laird, J. E., Lebiere, C., Rosenbloom, P. S. (2017). A Standard Model of the Mind: Toward a Common Computational Framework across Artificial Intelligence, Cognitive Science, Neuroscience, and Robotics. *AI Magazine*.

Tomasello, M. (2019). Becoming Human: A Theory of Ontogeny. The Belknap Press of the Harvard University Press.