



February 7, 2018

Intelligent Cognitive Assistants (ICA) Workshop Summary and Research Needs

Collaborative Machines to Enhance Human Capabilities

Abstract:

The research needs and opportunities to create physico-cognitive systems which work collaboratively with people are presented.

Workshop Dates and Venue:

November 14-15, 2017, IBM Almaden Research Center, San Jose, CA

Workshop Websites:

<https://www.src.org/program/ica>

<https://www.nsf.gov/nano>



Oakley, John / SRC

John.oakley@src.org

Table of Contents

Executive Summary.....	2
Workshop Details.....	3
Organizing Committee	3
Background.....	3
ICA-2 Workshop Outline	6
Research Areas Presentations	7
Session 1: Cognitive Psychology	7
Session 2: Approaches to Artificial Intelligence	11
Session 3: System Architectures for ICA.....	14
Session 4: Data and Modeling	17
Session 5: Natural Interfacing	21
Output from Break-out Sessions.....	24
Next Steps	27
Research Needs.....	27
Concluding Remarks.....	30
Prepared By.....	31
Appendix A: Report Feedback and Additional Thoughts	31
Appendix B: Workshop Agenda.....	34
Appendix C: References.....	36
Appendix D: Recommended Reading/Viewing.....	39

Acronym List

AI	Artificial Intelligence
ASSIST	Advanced Self-powered Systems of Integrated Sensors and sysTems
CAI	Cognitive Artificial Intelligence
CHAOS	Cognitive Holistic Assistant Operating System
EKG	Electrocardiogram
GPS	Global Positioning System
GPU	Graphics Processing Unit
IoT	Internet of Things
ICA	Intelligent Cognitive Assistants
JITAI	Just-In-Time Adaptive Intervention
ML	Machine Learning
PCAST	President’s Council of Advisors on Science and Technology
POV	Point Of View
SOC	System On Chip
TDC	Thermo-Dynamic Computing
UX	User eXperience

Executive Summary

The next evolution of the human-computer collaboration will soon become reality. The Intelligent Cognitive Assistant (ICA) initiative will accelerate the creation of systems that demonstrate the harnessing of machine intelligence to work collaboratively, enhancing human cognitive and physical capabilities. ICA will lead to the creation of new platforms to assist people in working, learning, interacting with new cyber-physical systems, redefine the human-computer interface, transport, healthcare, and other activities. Humans face a daunting list of challenges every day, and the future ICA applications will lend a helpful digital hand. These new applications will be individually adaptable and developed to augment the human performance and productivity, establishing new fields such as Cognitive Artificial Intelligence (CAI) at the human-technology frontier.

This document summarizes the discussions and findings from the ICA-2 Workshop sponsored by the National Science Foundation and Semiconductor Research Corporation on November 14-15, 2017. The workshop brought together a multi-disciplinary group of experts to shape the vision of the future of CAI and determine a comprehensive set of research needs for the ICA program. The workshop consisted of representatives from 18 industry companies, 15 universities, and 4 government agencies.

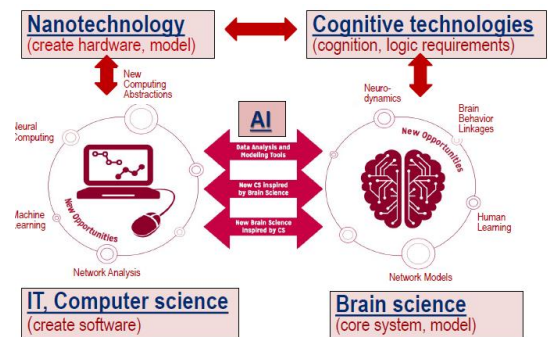


Figure 1. Convergence of Foundational Technologies for Brain-like Computing and ICA (Mihail C. Roco)

The workshop was divided into multiple research topic areas with a diverse set of experts in each area. Each session had a distinct focus with a very engaged audience discussion resulting in the emergence of six common themes as listed below. Many of these themes were echoed in the research needs identified in the break-out sessions:

1. **Contextual knowledge:** Aggregate existing stimuli with past history to create context
2. **Holistic AI:** Combine methods and technologies to enable complex reasoning tasks
3. **Social Science:** Humans are complex, emotional, dynamic, connected and interdependent
4. **Natural Human-ICA Interaction:** Needs to be seamless, comfortable, and trustworthy
5. **Edge Processing:** Significant improvements in user local, low power architectures
6. **Exception cases:** Should be the operational norm not the outlier

The program goals, future applications, and research space for the Intelligent Cognitive Assistant (ICA) initiative are vast, from production to social life. Despite that immensity, a pool of experts was able to identify these common themes as the near-term foci of research that will drive this vision. We are on the precipice of creating Cognitive Artificial Intelligence which will enrich the lives of humans over their entire lifespan from childhood to senectitude. Enhancing human physico-cognitive capabilities, while respecting social, ethical and legal concerns, is a main goal. Future ICA applications will be the new foundational technology driving the next wave of economic growth and improving the quality of life. **A new public-private partnership (government, industry, and academia) must be created to drive fundamental research on these multi-disciplinary themes.**

“Collaborative Machines to Enhance Human Capabilities”

Workshop Details

Organizing Committee

Ro Cammarota / Qualcomm	Sylvia Downing / Intel	Leslie Faiers / SRC
Brent Hailpern / IBM	Ken Hansen / SRC	Bruce Horn / Intel
Greg Leeming / Intel	Ron Kaplan / Amazon	Arvind Kumar / IBM
Lama Nachman / Intel	John Oakley / SRC – Chair	Jose Pineda De Gyvez / NXP
Mihail C. Roco / NSF	Vijay Saraswat / IBM	Reid Simmons / NSF
Howard Wactlar / NSF	David Wallach / SRC	

Background

This Intelligent Cognitive Assistants workshop (ICA-2) is a follow-on to an earlier workshop (ICA-1) held on May 12-13, 2016. The goal of the first meeting was to determine the guiding values and end user applications to be covered; one of the key take-aways was the need for intelligent agents which enhance the human-computer interface. This follow-up ICA workshop was focused on reinvigorating the ICA program with industry involvement and refining a list of ICA research needs. This list of research needs will be the basis for a joint NSF/SRC partnership for university research.



Figure 2. Defining Cognitive AI for ICA (Mihail C. Roco)

An earlier report, “Converging Technologies for Improving Human Performance” (Ref 1, NSF-DOC Report 2002, Springer 2003), recommended “personal assistant, broker” or in the first phase “digital assistant” as one the visionary projects that will be realized within twenty years.

In 2015, PCAST recommended support for brain-like computing: “Create a new type of computer that can proactively interpret and learn from data, solve unfamiliar problems using what it has learned, and operate with the energy efficiency of the human brain.” (OSTP / NNI Grand Challenge, <http://www.nano.gov/futurecomputing>, 2015, Ref 2). As a follow up, NSF proposed using brain-like computing as the logic unit for ICA.

The Computing Community Consortium (2016, Ref 3) sponsored the report “Accelerating Science: a grand challenge for AI”, where the suggested areas of research would support cognitive artificial intelligence development.

The ICA initiative is the result of convergence of fast advancing fields in science and engineering that would allow efficient implementation of cognitive artificial intelligence, non-von Neumann computing architectures, progress in cognitive psychology, problem-solving algorithms, and human-machine-environment effective interfaces (Ref 4 and 5).

Earlier research on ethical, legal and other societal aspects will be a condition for acceptance and progress of ICA including human-ICA and human-ICA-human interactions. Artificial intelligence and

“Collaborative Machines to Enhance Human Capabilities”

brain-computer interfaces must respect and preserve people's privacy, identity, agency and equality according to the report from a 2017 NSF workshop (Ref 6). Future AI systems need management transparency (explainable systems and actions).

The ICA vision is at the confluence of two national programs (National Nanotechnology Initiative, www.nano.gov; and National Information Technology R&D, www.nitrd.gov), as well as two OSTP initiatives (BRAIN and National Robotics Initiative), three NSF Big Ideas (Human-Technology Frontier, Growing Convergence Research; and Harnessing Data Revolution), and a NAE Grand Challenge “Reverse-engineer the Brain”.

In order to change paradigms and be able to address both general and exception cases, the report aims at encouraging exploration of new concepts, such as those illustrated in Ref 7, 8 and 9.

In 2016, NSF and SRC together with 11 companies sponsored the October 2016 workshop focused on the requirements of “Intelligent Cognitive Assistants” (ICA-1). The goal is to create systems that are highly useful to humans, specifically on the application of harnessing machine intelligence to augment human cognition and mimic human problem-solving capabilities. Research will focus on exploring scenarios for developing the novel architectures, concepts, and algorithms that will be required for ‘assistants’ to perceive, interpret, reason, learn, evolve, and interact in an energy-efficient manner, and in this way to provide actionable information, physical support, and informed advice to their human users.

Three application scenarios were described – life-long education, group work, and elder care – that incorporate sensitivity to these research parameters in complex social environments, and which require interdisciplinary research to fully address. These life cycle scenarios can also be used directly as a ‘roadmap’ to guide new research towards addressing the most challenging needs as humans individually evolve as well as interact collectively throughout their lifetimes. In addition, because of the breadth of the pending impact of the application of machine intelligence and more specifically, ‘Intelligent Cognitive Assistants’, on such a wide swath of the global public, government has a moral and ethical responsibility to ensure that such research is aimed at the common good.

The same guiding values from the ICA-1 workshop were used in the ICA-2 workshop; these are:

1. Enhance and serve, not replace, human capabilities.
2. Adapt with flexibility to dynamic, real-world environments
3. Cultivate “trust” among humans and machines
4. Facilitate "natural" interactions
5. Incorporate multi-disciplinary, multi-stakeholder perspectives

The ICA-2 organizing committee identified five primary research areas and five ICA applications. These areas were used to solicit industry, government, and academic involvement in the ICA workshop.

Table 1. Suggested Research Topic Areas

Cognitive Psychology	<ul style="list-style-type: none"> ▪ Perception vs Cognition ▪ Case Studies of Existing Digital Assistants ▪ Increasing Emotional Intelligence
Approaches to Artificial Intelligence	<ul style="list-style-type: none"> ▪ Beyond Five Tribes of Machine Learning ▪ Better Understanding of Common Sense Reasoning (Essence of Things) ▪ Data Selection for Reinforcement Learning
System Architectures and Devices	<ul style="list-style-type: none"> ▪ Moving Beyond von Neumann Computing Architecture ▪ Local Preprocessing for Speed, Security, and Personal Context ▪ Adaptable, Scalable, and Flexible ▪ Devices, sensors, actuators
Data and Modeling	<ul style="list-style-type: none"> ▪ Data Sources where Biases are Modeled and Encoded ▪ APIs and Toolkits ▪ Handling Sparse/Incomplete Data Sets
Natural Interfacing	<ul style="list-style-type: none"> ▪ Natural Language Processing ▪ Gestures, Perception, Vision, Sound, etc. ▪ Augmented Reality to Increase Trust ▪ Hardware for natural interfacing

Table 2. Proposed Applications from Organizing Committee

Personal Adaptive Tutor	<ul style="list-style-type: none"> ▪ Adapts to each individual students/employees ▪ Awareness to things such as mood, surroundings, situation, etc ▪ Leverage traits such as response time to assess knowledge absorption
Elder Care	<ul style="list-style-type: none"> ▪ Personal/Healthcare assistant to support specific needs ▪ Knowledge of medicine habits and connections to doctors
Cognitive Assistants in Smarter Communities	<ul style="list-style-type: none"> ▪ Address social and ethical aspects of cities and communities ▪ Team Assistant to support collaborative work environments
Context-aware Assistant (E.g. Nail gun training)	<ul style="list-style-type: none"> ▪ Seeing it would make basic information available in AR/HUD - controls, safety, etc. ▪ Pointing at others would display warnings ▪ Automatic counting of nails, time, etc.
Behavioral Models for Autonomous Vehicles	<ul style="list-style-type: none"> ▪ Model behaviors of drivers and pedestrians in the context of "Intention & Negotiation"
Other applications for consideration	<ul style="list-style-type: none"> ▪ Global Instability ▪ Environmental sustainability ▪ Climate Change ▪ Augment factory workforce to reduce repetitive tasks

ICA-2 Workshop Outline

To facilitate the creation of the ICA research program, a SRC - NSF sponsored workshop was held at the IBM Almaden Research Center in San Jose, CA on November 14-15, 2017. The agenda and presentation materials are available at <https://www.src.org/calendar/e006378/> (log-in required). The goal of the workshop was to gather a multi-disciplinary group of experts from industry, academia, and government together to discuss relevant, high priority research needs. The outcome of the workshop will be a set of research needs which are the basis for a Joint SRC - NSF research partnership beginning in 2018. Over 60% of the participants were from industry, representing 18 companies in software, hardware, social media and various application areas. About 30% of participants were experts from academia, and 10% from government organizations and professional societies.

The ICA-2 workshop began with opening remarks about the vision of the ICA initiative and opportunity for research. The ICA-2 workshop was structured in five sessions on several topics determined by the organizing committee. Each topic area session was chaired by an industry leader, started with a 30 minute keynote, followed by multiple 5 minute talks from each panelist, a 30 minute round table with audience interaction, and concluded with an informal discussion period. Each topic area session had a good mix of industry, academia, and government representation from many different fields of study.

After the five topic sessions, there were five parallel, research area focused break-out sessions where the attendees spent 90 minutes brain storming on the research needs for their research area, followed by 60 minutes to prioritize the research needs, and concluded with each topic chair presenting the research needs from each group to all attendees.

Research Areas Presentations

Session 1: Cognitive Psychology

Topic Chair Lama Nachman / Intel

<i>Keynote Speaker</i>	Bruce Horn / Intel	Intelligent Cognitive Assistants with Complementary Learning
<i>Panelists</i>	Shiwali Mohan / PARC	Interactive Cognitive Assistants: Design, Deployment, & Analysis
	Jonathan Gratch / USC	Challenges in Artificial Emotional Intelligence
	Misha Pavel / Northeastern	The Ultimate ICA: Approximating Human Intelligence
	David Reitter / Penn State	Bounded Cognitive AI and Dialogue

Intelligent Cognitive Assistants with Complementary Learning – Bruce Horn / Intel

“In not too many years, human brains and computing machines will be tightly coupled and the resulting partnership will think as no human brain has ever thought!”, JCR Licklider, 1960

Context is the circumstances that form the setting for an event, statement, or idea, and in terms of which it can be fully understood. Future ICA systems will need to combine Artificial Intelligence approaches with inspiration from psychology, cognitive science and neuroscience to achieve conceptual learning, fast instance learning, and lifelong learning. The multi-faceted learning capability is a key step in Cognitive AI development. In the future, the digital assistants will customize themselves to their human partner by having a personal knowledge graph, social graph, activity graph, and world knowledge graph. A personal intelligent agent will act in your best interest, anticipate your needs/wants, runs on-device/not in-cloud, and acts on your behalf. Some key metrics are trust, access, reciprocity, explainability, and transparency.

One example created by Intel and Oakley, Inc., is the Oakley Radar Pace. This very ambitious project merged many different technologies to build a brand natural interaction platform. This platform has a spoken conversational interface and is a running and cycling coach where context is key.



Figure 3. Oakley PACE: Example Natural Interface Platform (Bruce Horn)

“Collaborative Machines to Enhance Human Capabilities”

Interactive Cognitive Assistants: Design, Deployment, & Analysis – Shiwali Mohan / PARC



To advance Artificial Intelligence, there will need to be advances in Cognitive Science. Cognitive architectures are blueprints for intelligent behavior (perceive-decide-act). Future AIs will be expected to dynamically learn new tasks. The cognitive models will be instrumental in collaborating with humans; coaching human activities will help users achieve their skill learning and health goals. Some key metrics are usability, acceptability, correctness, efficiency, confidence, and flexibility of instruction.

Figure 4. Example of a Personal Digital Coach (Shiwali Mohan)

Challenges in Artificial Emotional Intelligence - Jonathan Gratch / USC

Recognizing emotions can reveal many things about a user, but still needs additional context to have a full understanding. Humans have been trained to self-regulate their expressions, and most expressions have nothing to do with emotion. Emotions are heavily dependent on context and are sometimes used to evoke emotion in the observers. Studying the dynamics of human interaction with social machines is a prime facet on future Cognitive AI research. One key question is “should a machine express emotion?”

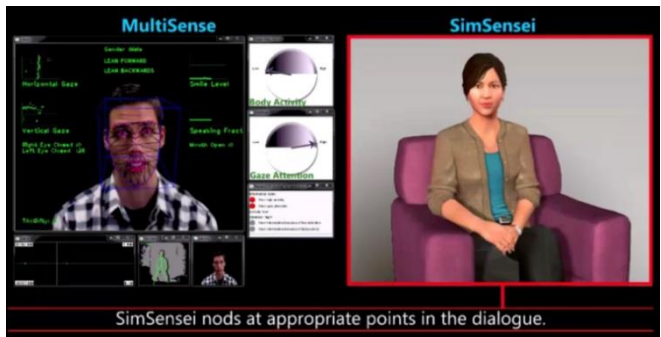


Figure 5. Ongoing Research in Emotion Intelligence (Jonathan Gratch)

The Ultimate ICA: Approximating Human Intelligence - Misha Pavel / Northeastern

“Assisting humans is more challenging than landing F/A-18 on a carrier ship in a storm.”

Understanding types of intelligence is another key aspect. Intelligence comes in many forms such as Crystallized Intelligence (trained skills) and Fluid Intelligence (reasoning). Developing the framework for fluid AI which handles anomalies without corruption is a key focus of future research. With ever increasingly ‘smart’ devices, there are many unobtrusive, continuous sensing systems which can assess the user. One interesting find from studies is many users which repeatedly train with digital assistants become accustomed to the training and perform better at a given task **before and after** the training.

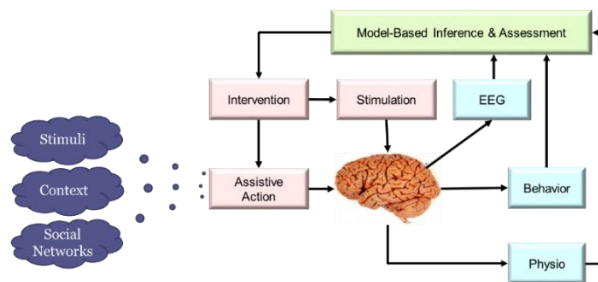


Figure 6. Vision: A framework for Intelligent Cognitive Assistant (Misha Pavel)

“Collaborative Machines to Enhance Human Capabilities”

Bounded Cognitive AI and Dialogue - David Reitter / Penn State

Soon, intelligent assistants will be able to outperform human capacities for communication and memory in many ways. For instance, they will directly access systems that computers have and humans do not have, such as large-scale, reliable, reproductive memory, and fast, parallel procedural processing. In other words, humans and computers have different limitations. To allow both to interface efficiently, it will become important for A.I. to structure interactions to suit human needs and to limit information display to provide only as much information as is necessary and processable. Because intelligent assistants will ideally use non-scripted, open-ended interactions as compared to the static graphic user interfaces we now know, they need good models of their human users. Specifically, they need to be aware of information bottlenecks, or generally, bounded cognition. Cognitive science can, today, explain how we contextualize our thinking and our communications. Bounds to cognition are important: we do not think of everything at all times, but rather pay attention to important aspects, and keep just a few concepts in working memory. We expect cognitive assistants to do the same in their interactions with humans.

Limiting scope, attention, and information density is key to managing bounded cognition. In an example from natural-language dialogue between humans, it can be seen that human dialogue partners similarly manage the information density of their utterances, shifting strategically from information contributor to recipient and back, and leveling information density to avoid peaks in cognitive demand. Intelligent assistants should aim to do the same, and designers of such assistants are advised to make strategic use of human limitations. In that way, we can harness the power of advanced AI, which outperforms humans in specific ways, while interfacing with humans to augment their broader cognition.




Figure 7. Humans strategically manage information density throughout each topic episode, taking on different roles before converging (empirical data shown from two large-scale conversation datasets) (David Reitter)

Session 1 Round Table Discussions

Are humans willing to give future Intelligent Agents total authority? The emergence of GPS has revolutionized navigation but has also limited human's ability to self-orientate. Will the next generation of CAI do the same? Giving a future ICA application suggested control will need to be done with extreme care to engender trust and improve human-computer collaboration.

How do you measure the performance of a given ICA application? It is highly dependent on the goal of application and to a lesser extent the user preference. Trust is an emotional construct and a key metric, but it is hard to measure and very easy to lose. If an application says it was "sorry", would it actually change its future actions to prevent a future occurrence?

What are the challenges in moving to add "Cognition" to AI? The third wave of AI has focused on perception and has exceeded humans in several cases. The frontal lobe of the brain handles the slower, more deliberate thought processes. Today's personal assistants appear to be thinking but are actually executing canned responses. Understanding the way human's think is crucial to creating the CAI. On the forefront of research is the concept of 'systematic forgetting' to mimic how human memory functions.

Moving out of the lab into the real world, today's systems break down very quickly. Determining emotional context becomes increasingly difficult as more stimuli are given to the human user. Research has shown that expression does not directly map to emotion (sometimes smiling is a defense mechanism instead of a display of happiness); personal context is key. Emotional responses are dynamic, not static; using standard predictive coding methods works on the assumption of static response to a given stimuli.

When future ICA applications read more of the user's biometrics, what are the ethical implications of machines having this additional data? One suggestion is keeping this information local as a requirement for future applications, but as humans we share some things with one daughter but not with other, so this is very context specific for each user. If the user knew the amount of biometrics which are being captured, would the user change their responses? Humans typically cannot predict their own emotional response to sharing information until they have already shared that information.

Session 2: Approaches to Artificial Intelligence

Topic Chair Brent Hailpern / IBM

<i>Keynote Speaker</i>	Charles Ortiz / Nuance	Holistic Conversational Assistants
<i>Panelists</i>	Todd Hylton / UCSD	On Thermodynamics and The Future of Computing
	Julie Adams / Oregon State	Intelligent Adaptive Assistants
	David Traum / USC	Off to See the Wizard
	Craig Knoblock / USC	Task Learning to Build Intelligent Cognitive Agents

Holistic Conversational Assistants - Charles Ortiz / Nuance

Developing “holistic AI systems” will be a key challenge of future Intelligent Assistants. These are systems that require a hybrid collection of AI technologies (deep learning, symbolic AI, statistics, etc), all of which must operate in a holistic fashion to solve problems of interest to a user. Examples of such systems include future personal conversational assistants as well as future robotic systems that require integrated reasoning and language capabilities. Development will primarily focus on examples from the former case and describe work in support of task-centered dialogues that involve the coordination of many complex modules, along with models of multi-agents with collaborative support.

Holistic AI systems will require end-to-end integration of many computation elements like perception, language, cognition, and action. Today’s assistant model is typically one-shot and uses a spaghetti model to add new capabilities. Human dialog must be studied as a linguistic ontology where context is constructed as part of the dialog progression; human conversation is not a website search string. Future holistic AI systems must be tested with real world data; one suggestion is for NSF to sponsor a challenge problem to stretch the boundaries of future Cognitive AI systems.

Goal: **Conversational interaction in context!!**

Humans change their minds, juggle multiple tasks, have personal preferences



Figure 8. Human Conversation is Very Complex (Charles Ortiz)

On Thermodynamics and the Future of Computing - Todd Hylton / UCSD

A primary problem in computing today is that computers cannot organize themselves: trillions of degrees of freedom doing the same stuff over and over... Today's mechanistic approach to AI is ill-suited to complex, real-world problems. One concept for the future of computing is Thermodynamic Evolution. The idea that thermodynamic potential in the environment drives the creation and maintenance of organizations to relieve that potential by increasing the thermodynamic entropy in the greater environment. A Thermo-Dynamic Computer (TDC) is a system that evolves its organization to transport free energy in response to electrical and information potential in its environment. A TDC learns to gather diffuse sources of energy in the environment, concentrate them, and deliver them back to the environment where they become even more diffuse (entropy is increased). Evolution of a Thermodynamic Computing system can be biased through programming, training and rewarding.

Intelligent Adaptive Assistants - Julie Adams / Oregon State

The goal of the Intelligent Agent is to develop a seamless synergy between humans and (semi-) autonomous machines in order to achieve complex tasks as effectively and efficiently, if not better, than human only teams. Some key technologies to be developed in CAI are real-time sensor fusion/assessment, understanding team and task states, modeling and predicting of current and future states, and the ability to autonomously adapt.

Off to See the Wizard - David Traum / USC

Many of today's digital assistants are one-shot task engines which are hand authored and somewhat brittle. Current deep learning systems have superior performance to known and similar tasks and often have spectacular or weird failures that may even look like success. To advance the cognitive AI, new directions need to be investigated to achieve Intelligent, Empathic, and Ethical assistants:

1. If I only had a brain: include models of user goals (not just utterances) with plans to achieve them, use of conversational implications, and common-sense interfacing
2. If I only had a heart: include models of emotion and empathy (build rapport with user, sensitive to user preferences) and help when user would welcome advice (do not break flow)
3. If I only had a nerve: include system goals separate from user goals, distinct models for different users, and say "no" when appropriate (weigh short-term vs long-term user desires)

Task Learning to Build Intelligent Cognitive Agents - Craig Knoblock / USC

Today, most people use personal assistants (SIRI, Alexa, Cortana, Google, etc) for fixed, known tasks like to set an alarm. What is missing is the ability to dynamically learn new tasks. The typical user desires the Intelligent Agents to execute tasks for them but it should be able to handle new tasks, adapt to changing situations, and customize knowledge to individual users. To achieve this goal, the ontology of knowledge and tasks needs to be understood with the ability to adapt over time. There needs to be a combination of machine learning, inputs source discovery, and directed task learning.

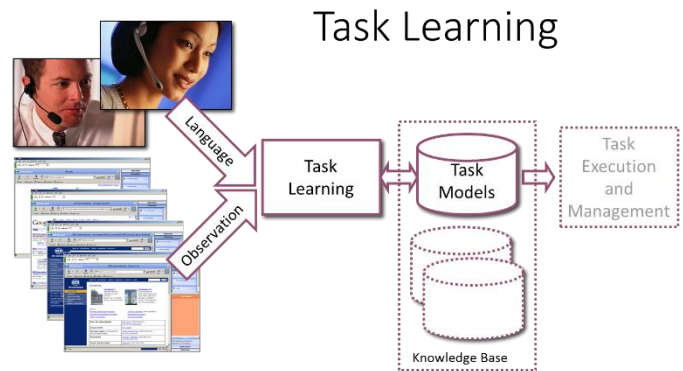


Figure 9. Task Learning is Key for Future Cognitive AIs (Craig Knoblock)

Session 2 Round Table Discussion

How will ICA redefine the definition of “user experience”? Future applications will move the UI beyond colors/fonts/buttons and move natural language processing to becoming directed search-like strings. The expanding interfacing modals make it more challenging and requires rigorous study and human trials by giving them options to use/not use modals in various combinations (mind reading will be a new challenging modality to handle). The user should not need to learn a new language to interface with future applications.

A few examples were given of previous semi-autonomous systems and issues with how humans interfaced with them. As the ICA application become more ubiquitous, the ways humans interact with the intelligent systems and the world will dramatically change and will continue to evolve other time.

Much of the world knowledge is fact based, not task based. This means that the existing ontology of knowledge does not scale well to CAI. One idea is to crowd source the development of a pantheon of task based knowledge. Both types of knowledge we be required for future ICA applications (“meta-reasoning”). Most dialog occurs because two sides do not know what the other side wants or needs. The exception is the norm for most conversations. Aggregating knowledge with intent is a key focus of CAI.

Future ICA applications will be tasked with team collaboration, personal education, and aging-in-place. Different situations create different challenges, maybe decisions need to be temporal salient. Multiple ICA applications will need to collaborate with the users as well as each other. The time value of knowledge and action is a challenge for ICA research.

What would you do if there was a billion dollars for research? One answer was an executive assistant which shapes your daily life, but future Intelligent Agents will simplify almost every human daily life. Having contests are great ways to accelerate the state-of-art.

“Collaborative Machines to Enhance Human Capabilities”

Session 3: System Architectures for ICA

Topic Chair Arvind Kumar / IBM

Keynote Speaker	Jeff Krichmar / UC Irvine	Efficient, Predictive Coding and Thermodynamic Computing
Panelists	Geoffrey Burr / IBM	Are non-von Neumann Architectures Essential?
	Ada Poon / Stanford	What More Can We Learn from Biological Systems That Would Aide Us in Designing ICAs?
	Roozbeh Jafari / Texas A&M	System Architectures For ICA: Role of Wearables
	Emre Neftci / UC Irvine	Neuromorphic Learning Machines

Efficient, Predictive Coding and Thermodynamic Computing - Jeff Krichmar / UC Irvine

Cognitive computing and cognitive architectures have derived inspiration from many sources over several decades. There are two main directions that have dominated developments in this field for decades: understanding and mimicking the brain (e.g. neuromorphic computing or computational neuroscience), and optimization with large-scale parameter refinement (e.g. machine learning). A lack of holistic methods and a strong reductionist, data-driven approach, hampers the brain-directed approach. In contrast, the fields of artificial neural networks, machine learning and artificial intelligence focus on well-defined benchmarks or problem statements, are largely agnostic about the technique used, and thus tend to be over specialized. Both approaches have merit, but that they lack a critical component of understanding necessary for diverse, capable cognitive computing systems that rival natural systems.

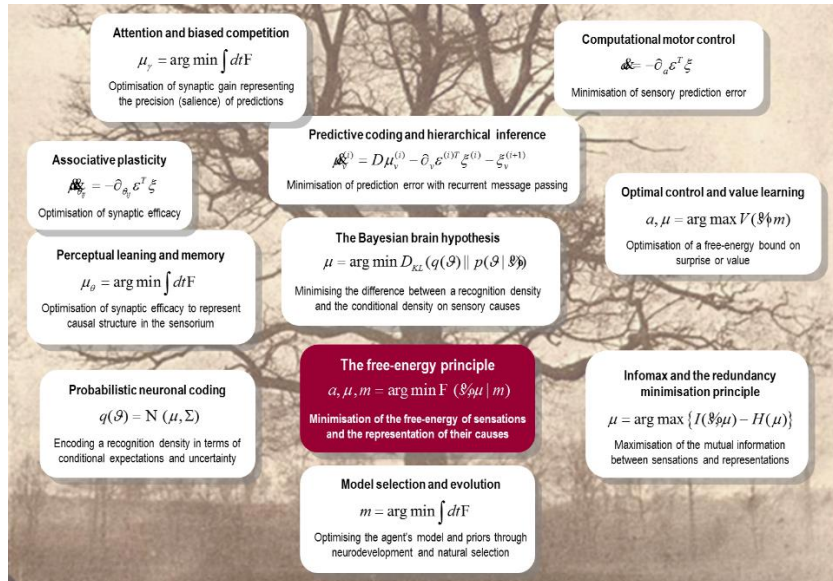


Figure 10. The Free Energy Principle: Unifying Brain Theory (Karl Friston, 2010)

A potential architecture for developing ICA is grounded in the hypothesis that thermodynamics should be the principal concept in future computing systems. It has been suggested that for any self-organizing system that is at equilibrium with its environment, it must minimize its free energy. The idea of minimizing free energy has close ties to many existing brain theories, such as Bayesian brain, efficient coding hypothesis, predictive




Figure 11. Possible Architecture for ICA (Jeff Krichmar)

“Collaborative Machines to Enhance Human Capabilities”

coding, cell assemblies and Infomax, as well as an evolutionary inspired theory called Neural Darwinism or neuronal group selection. In these theories, plasticity is modulated by value, and value systems control which actions lead to evolutionary fitness, that is, predicting outcomes that lead to positive value and avoid negative value. From a thermodynamic perspective, value is associated with those actions that minimize the increase of entropy (e.g., feeding, predicting outcomes, gathering information). We propose a thermodynamic computing architecture that is a closed loop system where the control (algorithm) is closely coupled with the body (robots) and the world (environment).

Are non-von Neumann Architectures Essential? - Geoffrey Burr / IBM

Looking beyond today’s Deep Learning paradigm is a key aspect of research. Artificial Intelligence hardware is based on neural network theories which rely on large training sets for accurate classification. The mathematics of these neural networks (like matrix multiplication and inversion) are handled in Graphics Processing Units (GPU), but are limited by data movement. Future systems may perform computation in the storage areas using weighted scaling approaches. Because the von Neumann architecture has been used for many generations, these future systems will likely first be combined with von Neumann systems until there is a complete understanding of the needs.

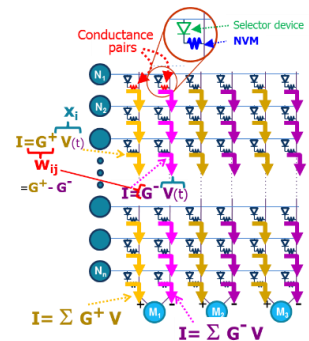


Figure 12. In-memory Computation Example (Geoffrey Burr)

What More Can We Learn from Biological Systems That Would Aide Us in Designing ICAs? - Ada Poon / Stanford

Studying biological systems can give insights to new system architectures. For example, each brain neuron is a self-contained processing element. A human brain has 86 billion neurons with 100 trillion synaptic connections and connect in ways that humans are just beginning to understand. The human brain is not connected in a limited 2D way as are modern computer systems. Future computing platforms will need to move to a 3D interconnect paradigm with asynchronous communication. One concept is event driven, wireless computing nodes which interconnect beyond normal 2D fabrics.

System Architectures for ICA: Role of Wearables - Roozbeh Jafari / Texas A&M

Future Intelligent Agents will need to leverage wearables to facilitate natural interactions between human and machine and decouple the traditional user interface paradigm. Existing wearables are seamlessly collecting data but the fusion and context needs to be better understood. Some of the gaps in today’s systems are information delivery overload (ignoring user context or lacks timeliness) and lack of intuitive UI beyond display/keyboards. Future Intelligent Assistants must be able to summarize information quickly and provide the right information at the right time to improve situational outcomes.

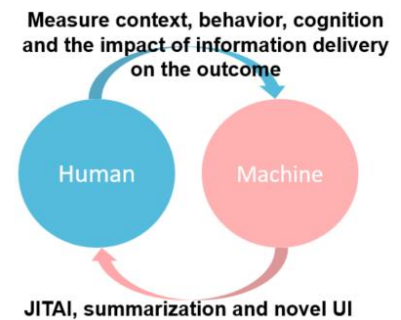


Figure 13. Future ICA Applications will be Seamless (Roozbeh Jafari)

“Collaborative Machines to Enhance Human Capabilities”

Neuromorphic Learning Machines - Emre Neftci / UC Irvine

One of the key requirements for future Intelligent Agents is the ability to process locally with the user at the edge, not in the cloud; this leads to more personalized learning models. One technique for edge learning is neuromorphic computing which enables low-power, massively parallel learning machines. To achieve these goals, only “spikes” of information should be routed between neurons (which have local internal weighing). One research goal is to find the optimal balance between the function, neural, and synaptic dynamics of a given system.

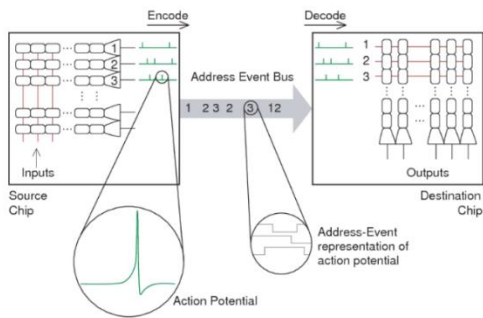


Figure 14. Spiking Neuromorphic Engineering as an Enabler for Low Power Applications (Emre Neftci)

Session 3 Round Table Discussion

Dynamic reconfiguration of the device interconnect is a complex problem. How do we add a level of neuro-plasticity to future devices? Modern systems are good at handling the 2D interconnect problem, but get exponentially worse as the dimensionality increases to 3D or 4D and beyond. Most people think scalability is making a system smaller, but the more interesting question is redefining scalability when if the system gets bigger, it gets better. To measure this requires having good metrics about which problems are being solved- one key metric to minimize is the energy used.

Should we limit our research to evolutionary processes? (For example, a bridge is not functional during each step of its creation, but most think that each step of a CAI development should be functional) Evolution has many missteps, and many variations, and only the optimal solutions survive. Adapting to new situations and stimuli is a requirement for CAI.

The nature of the models used for the predictive systems is another complex problem. Predictive systems will try to determine all the possible outcomes, but predicting the outlier is key; the exception is the metric by which future ICA applications will be judged.

Session 4: Data and Modeling

Topic Chair Danil Prokhorov / Toyota

<i>Keynote Speaker</i>	Davide Santo / NXP	Autonomous Driving Planning: A Convergence of Artificial Intelligence and Cognitive Science
<i>Panelists</i>	Matthias Scheutz / Tufts	Human-Level Performance
	Kagan Tumer / Oregon State	Objectives as Models: Focus On “What” Not “How”
	Danil Prokhorov / Toyota	How to Cope with Deluge of Data?
	Anthony Kuh / NSF	Real-Time Learning for Dynamic Systems

Autonomous Driving Planning: A Convergence of Artificial Intelligence and Cognitive Science - Davide Santo / NXP

The field of intelligent automobiles has seen great progress in distinguishing objects in the environment of an automobile driver, like recognizing and classifying vehicles, landmarks, and traffic signs, among others. A particular challenge is the concept of a Planner, eventually in a SoC, that is capable of planning the path, selecting the maneuvers and deciding the trajectories of an Automated Vehicle starting from autonomous driving level 3, where drivers are still necessary in the car, going eventually to full autonomy level 5. A research application proposes modeling the behavior on the road of drivers and pedestrians in the context of “intention & negotiation.” For instance, understanding and modeling the behavior of the pedestrian’s intention to cross the street ahead of a car in motion, or the driver’s negotiation abilities into merge-in lanes in geo-fenced scenarios like highways. The model should sum up safety levels of risk and cause-and-effect to prevent human fatalities. The resulting model will be encapsulated in a cognitive system to advance autonomous driving.



Figure 15. Perception to “AI Planning” – ‘Cognition’ is the Answer (Davide Santo)

One key concept is that the edge cases are the norm, not the rarity. Future Intelligent Agents need to understand the most contextually important information and must handle new situations effectively and efficiently. The modelling of information is instrumental to understanding how to solve the CAI problem.

How to Cope with Deluge of Data? - Danil Prokhorov / Toyota

In future autonomous and connected vehicle systems, large amounts of data may be generated in just one day but only a small fraction of the data is usable through traditional analytics (e.g., manual

“Collaborative Machines to Enhance Human Capabilities”

analysis). How do future Intelligent Agents handle the deluge of data? There is a need to develop cognitive algorithms to extract various behaviors and situations which are of interest for further improvement of vehicle systems, especially if they are time critical to process. For example, “Strategic Highway Research Program 2” has been collecting driving data since 2010, and had 5M trips, 50M miles, 695 known crashes as of July 2014, and an unknown number of rare events which may be characterized as “near misses”. Future Intelligent Agents must be able to help automating data processing so as to minimize human intervention while maximizing capabilities of data analytics. Some key research topics include developing cognitive planners, developing effective and efficient data processing (automatic extraction of relevant information, e.g., rare, imprecisely described events), building good predictive models, and trusted, life-long learning for cognitive agents.

Human-Level Performance - Matthias Scheutz / Tufts

Driving in the wild among other human drivers is a notable example of how limited and inadequate the bottom-up statistical approaches to driving are (in part, because they focus solely on driving and not at all the other possible situations that could impact driving in human environments). Humans employ



higher-level reasoning and skill to deal with unexpected events. Humans also make judgements about what is the ethical and morally right thing to do, even under time pressure. Future Intelligent Agents will need mental models of human nature, predictive models based on multi-modal interfaces, fast one-shot learning on the fly, fast creative problem solving, normative reasoning, and the ability to extract the relevant data from a given situation. One-shot learning combined with long-term learning is the key to enable future ICA applications.

Figure 16. One-shot Task Learning will be required for ICA (Matthias Scheutz)

Objectives as Models: Focus On “What” Not “How” - Kagan Tumer / Oregon State

Future Intelligent Agents will need to learn which objectives matter at which time based on preferences. In a world of Intelligent Agents, agents will need to interact with other agents as well as their human users. Understanding the objectives of a given task are key to explainability. When another agent does something you don’t approve of, it is typically a lack of knowledge of the objective function the other agent is doing.

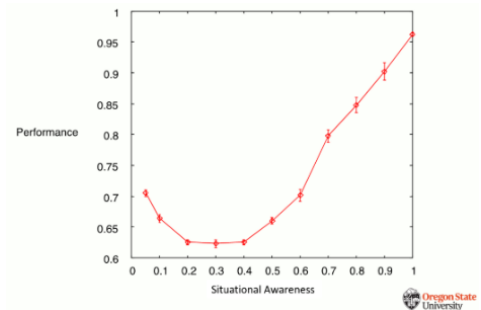


Figure 17. Is a Little Knowledge a Dangerous Thing? (Kagan Tumer)

Many of today’s systems focus on how to perform a given task instead of the objective of a given task: what the agent is trying to do and what to optimize when. One question is how much knowledge is needed- is a little knowledge a dangerous thing? Some studies have shown that more limited knowledge may lead to sub-optimal but acceptable outcomes.

“Collaborative Machines to Enhance Human Capabilities”

Real-Time Learning for Dynamic Systems - Anthony Kuh / NSF

The National Science Foundation has focused on handling the growing “Big Data” problem and is funding numerous projects in machine learning, neural networks, and sparse data processing. Real-time, efficient data reduction and context extraction is a key feature. NSF is focusing on solving the problems of society, and the ICA program will address societal concerns and introduce a new paradigm for contextual computing.



Figure 18. NSF 10 Big Ideas for Future Investment (Anthony Kuh)

Session 4 Round Table Discussion

When there is a failure, what kind of explanation is good enough? First, any explanation is better than what is available in today's machine learning systems. Having a framework which can be explained is the starting point, then refinement will occur to create better explainable models.

What kinds of data should we be concerned with? If we see a truck jack-knife on the other side of the road in the distance, humans would know to be cautious and maybe get off the road but would an Intelligent Agent do the same? The long tail, edge cases are the hard things to solve. One possibility is to somehow improve context the way human drivers do (after all, driving intelligence is a subset of general human intelligence); another is to abstract away specifics to the root objective knowledges and solve *this* abstraction instead.

To maximize safety, an Intelligent Agent could avoid the corner case as much as possible; this should be the most explainable solution as well. In the extreme, a system would do nothing to minimize risk, but this is also an undesirable outcome; a balance of risk and reward is needed to be modelled.

How do you protect all the data and knowledge that future Intelligent Agents collect? Cyber security will be key, humans do not want future ransom-ware to take over. GPS spoofing is already here, how should systems handle cases where the stimuli may be corrupted? Keeping data and processing local is one mechanism to protect data and engender trust.

Humans are very good at applying cross-domain knowledge, the future Intelligent Agents will need to do the same. Understanding the types of knowledge and types of learning need to be refined. A new, erroneous data point should not invalidate previous learning; merging of the learning from error cases is a promising area of research. Using objective functions and commonsense reasoning are mechanisms to handle the erroneous, random stimuli from real world environments. Dynamic switching between these objective functions will be necessary based on the real-time context.

ImageNet enabled the explosion of modern computer vision, what will be the cognitive knowledge equivalent? Image segmentation and classification are easier than understanding the overall scene; the first task will be to develop a rigorous definition of cognitive knowledge.

What is the role of law in Intelligent Agents? Should there be different laws for Intelligent Agents versus humans? Autonomous vehicles may need to imitate behavior of safe human drivers which sometimes yield, rather than stop, at stop signs whenever there's no cross-road traffic.

Session 5: Natural Interfacing

Topic Chair Reid Simmons / NSF

<i>Keynote Speaker</i>	Marilyn Walker / UCSC	Natural Language Generation for Dialog Systems
<i>Panelists</i>	Sarah Shomstein / GWU	Attention and ‘Things-that-make-us-better’
	Ron Kaplan / Amazon	The Natural Interface: Conversation
	Omer Oralkan / NC State	A Hardware Engineer’s POV to ICA: Multimodal Interfaces Enabled by Sensors and Actuators

Natural Language Generation for Dialog Systems - Marilyn Walker / UCSC

When humans interact with one another in dialogue, they express their own personality, but they also adapt their conversational behavior to their dialogue partner. Although there is a great deal of activity, both commercially and in the research community on development of dialogue agents, there are still many challenges to be addressed when building agents that can produce behaviors that vary according to the dialogue context. Some recent research has focused on statistical methods to train an expressive language generator for dialogue, which includes work on both linguistic and nonverbal generation, and includes methods for learning how to say things from corpora. There are many open challenges with building agents who can behave more naturally and engage the user over the course of an extended dialogue interaction.

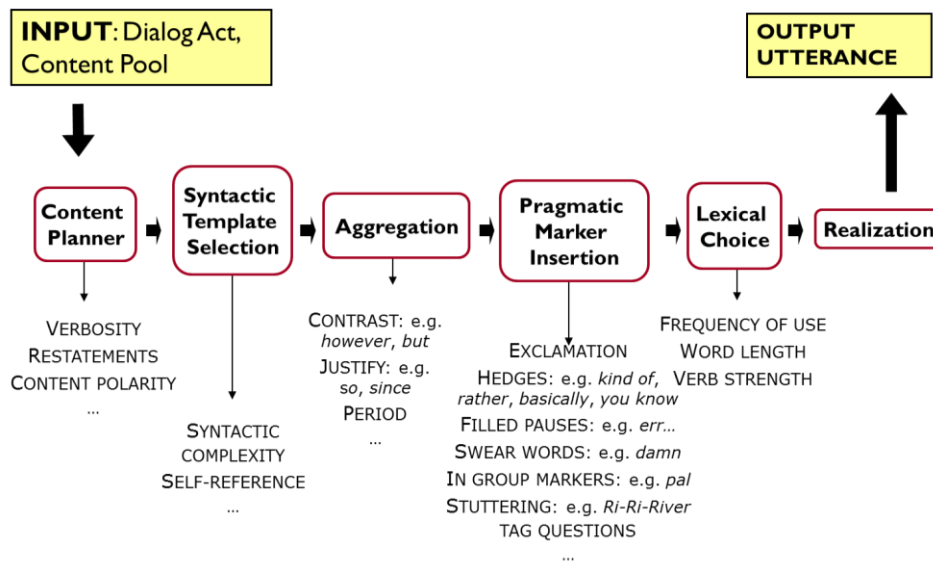


Figure 19. PERSONAGE Architecture: Dialog using 67 Parameters (Marilyn Walker)

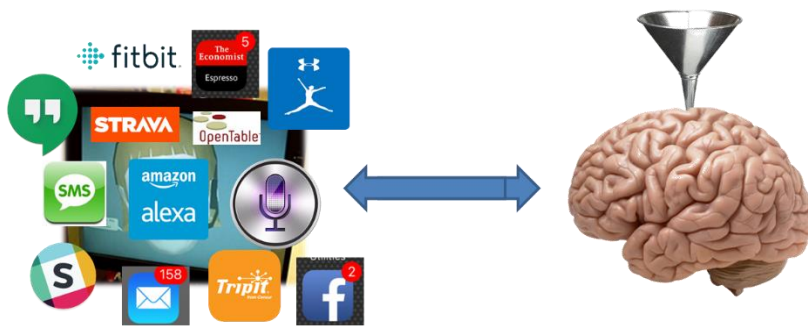
Historical natural language dialogue systems utilize hand-crafted Dialogue Manager and Language Generation which suffer from scalability issues. While machine learning approaches such as Statistical or Neural methods promise to address scalability issues, there are still many open challenges with getting the right training data, and ensuring that the system has a model of what it is trying to achieve. Moreover, human dialog adapts to the conversational partner and is colored by the speaker’s personality, individual style, mood, addressee, and situation; first impressions occur in the first three seconds of interaction. Future dialog systems will be more effective as they adapt to match the user’s

“Collaborative Machines to Enhance Human Capabilities”

personality and emotional state. Most voice web searching today is factoid based and is not dialogue; systems that can complete simple tasks like "setting a timer" are still completely hand-crafted and are highly inflexible. Future natural language interfacing will need a deep semantic representation of meaning and deep understanding of the structure of dialogue and methods for adaptation to the individual user and their personality. On top of natural language, future Intelligent Agents will need to include of other modalities like gesture recognition.

Attention and 'Things-that-make-us-better' - Sarah Shomstein / GWU

There is ongoing research to understand the cognitive of still image, but what's missing is the context. Much of a still image is overlooked by a human observer because the human focuses on the things which are more contextually critical. Humans tend to focus on the differences (the outliers). One



design aspect for future Intelligent Agents is to keep sensitivity to capacity limitations of the user; as the number of Intelligent Agents increases there will be increased competition for the limited user cognitive attention.

Figure 20. Trying to Understanding Limited Human Cognitive Attention (Sarah Shomstein)

The Natural Interface: Conversation - Ron Kaplan / Amazon

Human-Computer conversation should be effective, efficient, and pleasant. Conversation should be incremental, educational, explanatory, repairable, and expeditious. Future Intelligent Agents will need data driven (deep learning by observation) and symbolic (learning by instruction) mechanisms. Contextual knowledge is the key: saying "The chicken is ready to eat" would mean very different things in a barnyard verses a kitchen. Simple failures in understanding can make user distrust future Intelligent Agents.

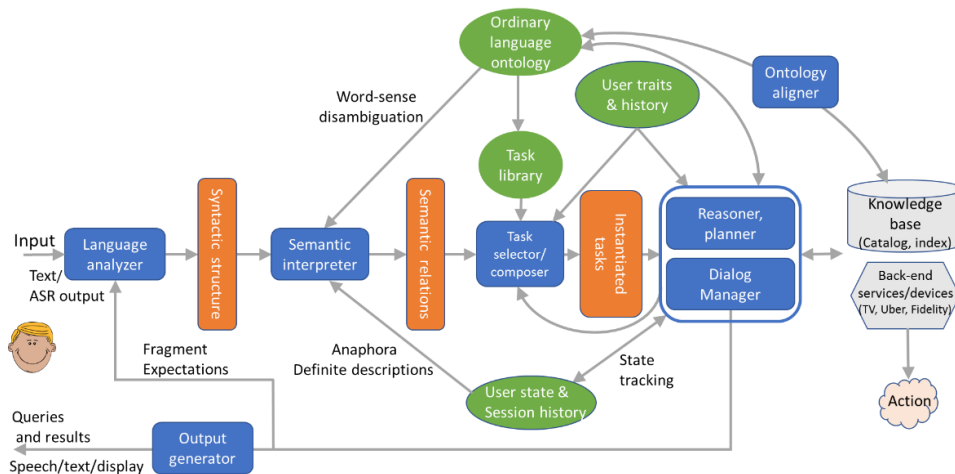


Figure 21. Solving the Natural Interfacing Requires a Combination of Many Resources/Technologies (Ron Kaplan)

"Collaborative Machines to Enhance Human Capabilities"

A Hardware Engineer's POV to ICA: Multimodal Interfaces Enabled by Sensors and Actuators - Omer Oralkan / NC State

The future of natural interfacing is to have many modalities of control beyond displays and language processing like gesture-base interfacing, EKGs on wearables, direct directional neural interfacing, and augmented reality. These interfacing modalities should be small, flexible, seamlessly integrated, low power, low cost, reliable, and provide useful complementary data. There is ongoing research on self-power systems of sensors (ASSIST), ultrasound-based noninvasive neural interfaces, and ultrasonic transducers embedded in glass which will be useful for future Intelligent Agents. Using biometrics will augment the input stimuli from the other interfaces to improve the “natural interface”.

Session 5 Round Table Discussion

History shows that humans adapt to the interfaces used, so why bother making “natural interfaces”? Having different, unique interfaces that are used infrequently can be cumbersome. The goal should be to have a common interface that anyone and everyone uses. It needs to be adaptive and learning capable, a new human assistant learns to work with a boss so should future Intelligent Agents. The better the interface, the less of the cognitive funnel of the user it will require. Context of the user (like biometrics), context of the tasks/goals, objectives of the collaborators, and environmental context are necessary to have a complete view of a situation and make the interface the most “natural”.

Does a good natural interface, make the Intelligent Agent appear to be smarter than it is? There will be a need to manage user expectations. Most of today's examples are built on web searches and appear smarter already. One comment was to improve the back-end processing to be better than web search. People's expectations are managed through dialog failures, people adapt to these failures.

What is a good dialog? There should be formal methods and a ‘world model’ of the users (which will always be incomplete). Proving that a CAI is sufficient requires research and having a human back door is assumed for the near future. A future Intelligent Agent should have mechanisms to detect dead-lock situations and have fall backs as appropriate.

Output from Break-out Sessions

After the research topic area presentations, the attendees were split into research area specific break-out sessions. Each team was asked to create a prioritized list of research needs and then present these needs to the audience as a whole. Each team used different methods, from post-it notes to white boards to electronic collaboration. Below are the research needs identified by the five break-out sessions.

Cognitive Psychology

Build end to end solutions / applications in cognitive assistance, bring current capabilities together, multidisciplinary approach to solving real problems and understanding limitations, metrics, etc. Need access to data, running experiments in the real world, across the world.

<i>Deeper understanding of cognitive/emotional state</i>	<ul style="list-style-type: none"> • Deeper understanding of cognitive state, capacity limitations (attention, memory, etc) and other context. It is important to also consider interaction between different agents and a single human.
<i>Understanding shared mental models</i>	<ul style="list-style-type: none"> • Understanding shared mental models from both sides is important, and even extending this to groups of people. This will help improve trust and efficiency of the communication
<i>Understand the cognitive process of the human</i>	<ul style="list-style-type: none"> • Understand the cognitive process of the human to help inform the design of cognitive assistants (brain inspired assistants)
<i>Metrics to measure the performance of cognitive assistants</i>	<ul style="list-style-type: none"> • The need to come up with metrics to measure the performance of cognitive assistants that bring in the perspective of outcomes beyond simply thinking about the user experience or it being human-like, or UX evaluation needs to incorporate outcomes not just good UX, is “human like” relevant?
<i>Affect recognition in the wild, moving beyond recognition to understanding</i>	<ul style="list-style-type: none"> • Affect recognition in the wild, moving beyond recognition to understanding. Explore the diverse ways to utilize affect (part of the context, Help make the assistant more compelling/ convincing, Feedback channel , help isolate ambiguity and understand impact)
<i>Integration of language and cognitive function</i>	<ul style="list-style-type: none"> • Integration of language and cognitive function, how does a cognitive system mesh with our ability to learn natural language? Language, memory, reasoning and perception. • How do humans detect saliency, how do they talk about it, etc. What can we learn from how humans do it that can influence how we build these systems?

Approaches to Artificial Intelligence

Build end to end solutions / applications in cognitive assistance, bring current capabilities together, multidisciplinary approach to solving real problems and understanding limitations, metrics, etc. Need access to data, running experiments in the real world, across the world.

<i>Train personal or team assistants to learn and execute new tasks</i>	<ul style="list-style-type: none"> • World model of action • Acquiring new tasks (training, rules, crowd-sourcing)
--	--

“Collaborative Machines to Enhance Human Capabilities”

Build innovative, reliable, and scalable task based systems

- Learning primitive tasks
- Self-knowledge of assistant capabilities (reflective system)
- Continuous learning
- Incorporate perception, language, cognition, learning, and action
- Tools and methodologies to build/evolve/debug/test these systems
- Software framework to compose, express, test, deliver such systems
- Integrate multiple mental faculties
- Cognitive holistic assistant operating system – CHAOS

Support collaboration among assistants and with human teams

- Sharing of tasks
- Ecosystem of task models
- Risk management, privacy concerns, etc
- Shared mental models

Define high-value targeted domains to test results or as basis for a challenge problem

System Architectures for ICA

The ability to define, represent and learn context is necessary for ICA, and is not met in current architectures/systems.

Representations of context, knowledge, and experience throughout architecture

- Representations may be variably sized and should be changeable
- Adaptability is key to life-long operation

Learning by collecting, synthesizing, and consolidating information

- Enable flexible learning and memory
- Interface between types of learning (...)
- Learning at different time scales (rapid one-shot learning to lifelong learning)
- Different modalities and multi-sensor fusion
- Explainable AI is a must to engender trust

Value/saliency must be built in: bottom-up saliency /top down goal directed behavior

- Will influence connectivity between representations
- Incorporate attentional mechanisms into ICAs

Efficient hardware for representing, storing, and retrieving contextual data

- General purpose vs. special purpose design tradeoffs
- Supporting scalable distributed processing
- Incremental processing of information
- Energy efficiency optimization after achieving functionality

Physical environment and embodiment of agent must be accounted for (e.g., hospital, battlefield, car, school)

- General vs. domain specific agents
- Interfacing must be appropriate to operating environment

Data and Modeling

Advances in control interface and sensors has led to an explosion of data, how will the future Cognitive AI handle the plethora of data? Other considerations: access to large qualities of data is not available to public (e.g. electric utility data, datasets at IT companies)

Investigate at wealth of ML algorithms	<ul style="list-style-type: none"> • Deep Learning, kernel methods (SVM), probabilistic methods, one shot learning, other ML methods (CART, random forests, ensemble learning, boosting), Incorporating feedback
Formalize interfaces for ML	<ul style="list-style-type: none"> • Inputs, outputs, constraints (timing, power, latency), modularity • Benchmark (evaluation metrics), look beyond generalization • Performance, complexity, implementation
Data representation (data comes from diversity of types)	<ul style="list-style-type: none"> • Example: brain codes using spiking neurons • Feature extraction, unsupervised learning, and non-numerical data; can we come up with more uniform way of representing data? (Google: tensor flow)
Pushing to the edge, edge ↔ cloud	<ul style="list-style-type: none"> • Amount of computation at edge (sensors), communication, security
Real-time learning (online learning)	<ul style="list-style-type: none"> • Real environments often nonstationary and decisions must be made in real-time • Can also handle learning with context where we update our estimate/ hypothesis of what we learn • Many applications rely on real-time learning (smart grid, autonomous vehicles) • Use wealth of signal processing /adaptive filter knowledge to extend to ML (deep learning, kernel Methods)

Natural Interfacing

Today's advancements in Natural Language processing is just the first step; the human-computer interface needs to be seamless, trustable, and ubiquitous.

Adaptation & Context Awareness & Personalization	<ul style="list-style-type: none"> • Adaptation: Noticing fatigue and modifying interaction modality. • Context Awareness: <i>Is the chicken ready to eat?</i> • Personalization: <i>I want to watch the game.</i> Preferences are sticky to the person across contexts
Models ↔ Signals	<ul style="list-style-type: none"> • Having a model and making inferences from the inputs and actions to update the models. • Taking those inferences and explaining them and expressing them
Training and Learning	<ul style="list-style-type: none"> • Observational Learning AND Instructable Systems [We need BOTH] • One shot task learning: teach it how to do things • User configurable interfaces
Input Fusion	<ul style="list-style-type: none"> • Take all possible IO channels and determine the context and the user's performance
Output Fusion	<ul style="list-style-type: none"> • The ability to produce an integrated multi-modal output • Channel management and allocation • Attention management

Next Steps

The program goals, future applications, and research space for the ICA program are vast, despite that there are some common themes which emerged from the workshop. Many of these were echoed in the research needs identified in the break-out sessions:

1. **Contextual knowledge:** Contextual knowledge is the key to aggregating the ever-growing volume of stimuli and limiting feedback to the user. Contextual knowledge is also temporal meaning context can change quickly over time. Need to develop a common ‘world model’ of knowledge.
2. **Holistic AI:** Cognitive AI must have a combination of long-term, machine learning and fast, one-shot learning. Future “Holistic AI” systems will combine different reasoning and learning methods with end-to-end intelligence.
3. **Social Science:** Humans are complex, emotional, and dynamic; any future Intelligent Agent will need to be adaptable to a given user. Enhancing human physico-cognitive capabilities, while respecting social, ethical and legal concerns, is a main goal.
4. **Natural Human-ICA Interaction:** Future human-computer interaction needs to be seamless, comfortable, empathic, trustworthy, and transparent. New paradigms are need beyond screen and keyboards. Humans should be able to query and modify future Intelligent Agents.
5. **Edge Processing:** Future Intelligent Agents will need to do much more locally, which requires significant improvements in low power, edge processing. New computing architectures will be needed. User context and knowledge will need to stay local to ensure privacy concerns.
6. **Exception cases:** The exception cases will be the norm by which future Intelligent Agents will be judged (like single occurrence events). This theme is integral to the other themes.

Research Needs

There several on-going research efforts in natural language processing, deep machine learning, and neuromorphic computing. These will augment future ICA applications but will not be the primary focus at the beginning of the ICA program. Instead, the primary focus will be on creating the holistic artificial intelligence, creating a modular ICA system architecture enabling scalability and adaptability, and improving the human-intelligent machine interface. This will be achieved with a wide range of domain expertise and should revolve around the common themes from the workshop. The theme of ‘exception cases’ is integral to the other themes and does not have separate research needs.

Contextual Knowledge

- | | |
|------------|--|
| 1.1 | Develop ontology definition of task/objective based knowledge |
| 1.2 | Aggregate different stimuli based on contextual knowledge (“Is the chicken ready to eat?” is different in a kitchen vs a barnyard) |
| 1.3 | Context is temporal, local, situational, and personal; need metrics for optimization |
| 1.4 | Affect recognition in the real world for contextual understanding |
| 1.5 | Extraction and visualization of objective tasks from stimuli based on context knowledge |
| 1.6 | Use insights from learned data/objectives to solve unfamiliar problems (outlier is the norm) |
| 1.7 | Develop repository database of contextual knowledge to teach ICA applications |
| 1.8 | Need systems to be aware of its own capabilities and limitations (need to understand the contexts which systems can, and cannot, reliably operate) |

Holistic Artificial Intelligence

- | | |
|-------------|--|
| 2.1 | Add task objective knowledge to one-shot learning research |
| 2.2 | Integrate one-shot task learning with machine learning to make holistic AI |
| 2.3 | Develop forgetting model for life-long, continuous, safe learning; data selection for reinforcement learning and temporal relevancy of previous data |
| 2.4 | Develop explainable AI model which connects stimuli to tasks to learned knowledge |
| 2.5 | Self-learning and learning from other Intelligent Agents or crowd sourcing |
| 2.6 | Observational learning based on task objective metrics and user feedback |
| 2.7 | Develop a Cognitive Holistic Assistant Operating System (CHAOS) |
| 2.8 | Develop metrics for measuring ‘correctness’ of AI and actions taken |
| 2.9 | Real-time learning and decision-making machines and systems with dynamic data |
| 2.10 | Combine symbolic, probabilistic and commonsense reasoning (meta-reasoning) |

Social Science

- | | |
|-------------|--|
| 3.1 | Deep understanding of cognitive state and human capacity limitations (attention, memory, empathy) |
| 3.2 | World model of the human user to judge future ICA quality |
| 3.3 | Study human emotional responses and humans interaction with emotional agents (emotional intelligence) |
| 3.4 | Understanding shared mental models of human response, as predictive as possible |
| 3.5 | Metric to measure ICA application performance; how do humans detect ICA saliency? |
| 3.6 | Study human attention/cognition limitations to bound expectations |
| 3.7 | Enhancing human cognitive and physical capabilities for better individual and societal outcomes |
| 3.8 | Embodiments for ICA; should they be robots with hands, some other form, or formless? Will users gain emotional attachment? |
| 3.9 | Group interaction, interdependence and outcomes using ICA applications |
| 3.10 | Ethical, legal and other societal aspects related to use of ICA applications |

Natural Human-ICA Interaction

4.1	True dialog systems for natural interfacing; need for connecting temporal thought snippets
4.2	Integrate natural language and cognitive function; move beyond web search responses
4.3	Multi-sensor fusion with different interface modalities; aggregation of numerous stimuli
4.4	Model different operating environments; agents must adapt to usage paradigm
4.5	Utilize human biometrics and attention model as input stimuli with other modalities
4.6	Develop new paradigms for feedback beyond displays and sounds
4.7	Domain transference - vision case: if you can identify a cat's nose, can you find a dog's nose?
4.8	Sensing and interacting with the surrounding environmental context (physical and social)
4.9	Preventing the data deluge using stimuli fusion, cognitive selection, feature extraction, etc.

Edge Processing and System Design

5.1	Develop heterogeneous systems of neuromorphic, von Neumann, etc computing for portability and scalability
5.2	Tools and methodologies to build/evolve/debug ICA applications
5.3	Develop beyond traditional 2D interconnection between processing elements; like spiking neurons and wireless interconnect
5.4	Develop incrementing pre-processing techniques to minimize input stimuli using context and task objective knowledge
5.5	Minimize response time for known situations, and limit response for unknown situations
5.6	Partition of functions between local computing and cloud based services; confidential data stays local
5.7	Develop techniques which minimize energy usage per decision or action
5.8	Develop non-intrusive, multidimensional, energy efficient sensorial system

Concluding Remarks

There is an urgent need pushed by emerging smart technologies and driven by societal expectations to create a general science and technology foundation for ICA that would identify fundamental design and functional issues, use new methods such as high level multi-domain languages, and be useful for a large variety of application areas. This will allow the creation of new industries and occupations, increased productivity, opportunity for innovation, improved quality of life, and sustained global leadership. With proper attention given to ethical, legal and other societal aspects, the ICA R&D investment promises to create a **general-purpose technology** that may define the national competitiveness in the future decades. This report has identified key research needs in this regard.

Current research already demonstrates that any concentration of control or surveillance creates an unsustainable and problematic situation and lack of trust. This problem is compounded in systems that are designed to be ubiquitous in their application to society at-large. To that end, it is recommended that ICA research address and fundamentally incorporate the extremely challenging **privacy, security, ethical and regulatory challenges** and tradeoffs associated with intimate digital systems. These concerns must be addressed both upfront and throughout any ICA development in order to maintain trust and transparency among humans and future Intelligent Agents.

The ICA devices and platforms that we research now **will operate in the future IoT**, with new action agents, using large databases and various emerging technologies. ICA must be learning, interact with changing surroundings, adapt, and reconfigure to satisfy the task demands in the future knowledge and digital society. It must support team work in production, improve performance and capabilities through real-time machine learning, develop trust within human-machine interactions, and be secure against tampering. Designs may incorporate new system architecture for hardware, actuators, and a wide range of abilities from sensory-perceptual to language use, to learning decision making and predictive strategies.

The long-term goal is to create the **killer-app** to support, enhance and innovate the production, interaction, security and quality of life needs of people living a fast changing and interdependent society. The ICA may lead to disruptive approaches that enhance our ability to work, learn and interact. There is a need also to lay the foundation for improvement of both worker quality of life and employer financial metrics in a manner that does not adversely impact the economic and social fabric of society. Enhancing human physico-cognitive capabilities for better individual and societal outcomes and welfare is a main goal.

A public-private partnership (government, industry, and academia) needs to be created to drive fundamental research on these multi-disciplinary themes. This would ensure various multidisciplinary and multi-stakeholder perspectives, collaborative participation from industry, innovative convergence and emergence of S&T solutions, and overall an accelerated technological and societal progress.

Prepared By

Ken Hansen / SRC	John Oakley / SRC	Mihail C. Roco / NSF
------------------	-------------------	----------------------

Appendix A: Report Feedback and Additional Thoughts

From: Tom Skalak / Allen Institute

This is an excellent and very broad treatment of opportunities in digital individual cognitive assistants (ICAs) and digitally-assisted human-societal-industrial endeavors.

The inclusion of “Physical enhancement” along with cognitive assistance broadens it greatly! – E.g. mechanized farming, sport, warfare, cooking, exercise, etc. why not consider some of these more explicitly, given the chosen scope.

What about “simplifying” our increasingly global complexity - not “assisting the human to cope”, but actually simplifying human existence for optimal thriving or human excellence/fulfillment/happiness?

Good to attempt knowledge transfer from brain science to AI - it has been very limited to date. It will be complex, with 1,000s of genes contributing to “simple traits”, and no extant neural circuits at the large scale needed to reverse engineer the human cognitive process - yet. More basic science at Systems level is still needed as part of effort.

Definitely will need emphasis on special human traits- including social cognition, empathy, risk tolerance, and collective action/collaboration at a distance across non-familial, non-genetically-related lines! The latter quality may be the key aspect separating human collaboration from many other species- so build it into future cognitive systems.

Are “exception cases” the new jargon for “creativity?” I find it odd and somewhat off-putting to parse all of human existence into either “past patterns analyzable by machine learning” or “exception cases”.

I suggest opening up a section on “assisting in anticipating and creating the desired or imagined human/planetary future”. The word creativity does not appear in the executive summary at all. This seems like a missed opportunity for such an ambitious program of study and development.

Not sure about the claim that this is “the killer app” that will drive the next wave of economic growth. Right now, AI apps pushed to consumers and business are becoming intrusive rather than assistive in character, forcing an undesirable shift in data entry and attention capital consumption onto individual humans for everything from elementary school education to arts and simple adult products such as banking, insurance, and reading/learning experiences. So a major pivot to such apps is already underway. It is true that improved digital assistants will be developed, but the horse is out of the barn already. What about “hope”, loneliness, altruism, friendship and family, among other human attributes? An “app” per se, as we know apps today, is unlikely to be a truly transformative moment at this stage in digital history. Perhaps a larger vision for CAI could be the centerpiece here.

“Collaborative Machines to Enhance Human Capabilities”

Other economic drivers in the physical domain, such as clean cheap solar energy and adequate food supply and efficient building methods may be larger in economic and social impact.

There is no question that improved human-machine interfaces are critically needed. My current smart phone's spelling assistant exhibits significantly poorer performance (costs me more time for the same result) than my prior version, as one immediate challenge to this industry. Searching ranked lists on a screen is not an optimal human interaction.

I'd suggest adding another "convergence" with the NSF's 10 big ideas: "Understanding the Rules of Life". Certainly we are still far from understanding human brain cognition, memory, or theory of mind (self-awareness and consciousness). Also, learning about non-human nervous system or even somatic tissue (non-brain tissue) information processing could inform elements of this program. For example, cephalopods appear to have evolved highly distributed sensing-processing-actuating Systems, representing a kind of distributed cognition or intelligence that could inform the next generation of ICAs. The tentacles can think and act independently of the head. This came about through evolution. Could be very useful as a design principle for digital assistants. Let's recognize it.

Parsing the universe of inputs to humans as "Tasks, sensing, and data" seems a limited starting point. A task would seem to denote already-specified operations to be executed. This implies all tasks are well-defined or have a prior/past form. Data is all about the past. Can the "task" box be enlarged to include design goals, creativity, play, imagination, or other social activities?

Adapting to real-world environments and natural interactions are two key goals here. The two most important research topic frontiers for ICAs would seem to be "enhancing emotional intelligence" and "natural interfacing". Market applications for data analytics will likely advance most of the other topics and applications, independent of government/public support.

Of the listed "applications", the first 5 seem to be existing areas of intensive industrial effort today. What's new?

Perhaps the most impactful of these could be the "personalized tutor", because today's school systems seem to have lost the ability to provide individualized guidance to learning, although they do serve a human inspirational and mentoring role to a degree that no ICA has yet approached. "Knowledge absorption" as a central goal seems a bit off-Target- the killer digital app should be a "guide to learning", not force human knowledge absorption. The Oakley-Pace technology example presented by Horn is an example of the coaching and mentoring function, to a degree. Could inspiration, role-modeling, and mentoring be added to the sub goals of this application?

The rest of the applications seem fairly routine extensions of today's control systems, expert systems, and digital assistants.

Autonomous vehicles are fraught with major economic and unresolved public-private issues that may exceed the actual technical challenge here, such as usage rights for massive taxpayer-funded infrastructure (built highways with \$T sunk public costs, airspace, etc) by private entities, which may infringe on constitutional or individual rights.

"Collaborative Machines to Enhance Human Capabilities"

Could a couple of more forward-looking applications be prioritized for public-private partnerships?

“Crystallized intelligence” (by Pavel) ought to include the mature, adult capacity to use experiential learning and associative thinking (metaphor) to create new ideas. Fluid intelligence is believed to “peak” at the age of 25-35, while the ability to act widely as a result of accumulated experience grow throughout life, even into old age for many healthy elderly humans. Let’s be sure to recognize that perhaps 80-90% of elderly people can think quite well- despite media attention the expected wave of population-wide cognitive impairment.

The “holistic system” concept (by Ortiz) is very important to long -term success of ICAs. The discussion of what holistic entails should be deepened and clarified. There may be a good short-term application here.

In the thermodynamic system concept (by Hylton), it is worth noting that living systems (and their brains) probably cannot be fully characterized as “concentrating energy sources” - they also use energy to create long-lived organized structures and information storage, access, and asymmetries. As one example, planarian (flatworms) have solved the problem of immortality, essentially being capable of maintaining the same body structure (an energy-consuming process) over millions of years. This far exceeds any computer or silicon-based memory system built to date, and also affects the time course of entropy generation beyond most computing timescales.

Traum’s inclusion of “empathy” as a design goal is very important.

Improving the transfer of neuromorphic computing concepts into actual devices is a good idea. It needs more attention.

Perhaps instead of “coping with the deluge of data” - which is today’s thinking - this program ought to include research on how ICAs would “prevent the data deluge” in the first place and simplify sensing, cognition, and acting? Thus might lead to disruptive and transformative changes in major industries, with benefits for humans and society.

Focusing on the “what” of objectives versus tasks per se (Tumer) is very important for long term success of ICAs. Appropriate layers of knowledge or definition of perceived objectives is a hard problem facing all organizations worldwide.

“Failure tolerance” is a key issue (Roundtable 4). This is not an engineering problem. It is a societal and cultural issue. The larger design question should be elevated well out of just the immediate trend for autonomous vehicles. Today’s power companies are not liable for damages or personal injuries caused by storm damage to power lines or poles. Yet this is simply a cost-safety trade off or design issue. It is possible to design power systems that withstand 100-year storms, but it would be costly. Society tolerates some risks but not others. Markets and engineering designs follow.

Breakout session topics:

- All the “cognitive psychology” ideas are important and need more research.
- In AI, perception and adaptability to real world is a priority.

“Collaborative Machines to Enhance Human Capabilities”

- In Systems Architecture, adaptability and flexible memory are important.
- In Data Modeling, feature extraction and data representation, and edge computation, are key.
- In Natural Interfacing, this whole section is very important- but as written here seems to focus on near-term, granular design details, rather than the larger issues. Among these larger issues are context awareness, empathy (missing here again), and preventing data overload.

“Next Steps”: items 1-5 seem quite on the mark. Item 6 on “exceptions” needs clarification.

Research needs seem appropriate in all sections. The notion of “correctness” of ICA actions is a major idea- it implies again a task orientation for all the research and assessment metrics. This seems non-optimal, if the larger goal is to realize the full long-term potential of ICAs, because a major part of their impact in the world will be to support human creativity and the future design of imagined societal structures, economies, and planetary systems including a sustainable ecosystem. Given this, the summary recommendation ought to articulate aspects of this forward-looking challenge, rather than only the current goals of managing data for achieving greater efficiencies in human-centered or human-machine integrated “tasks”. There are huge market forces driving this latter aspect of productivity enhancement, so public-private partnerships could aim at the longer term challenges.

Appendix B: Workshop Agenda

Day 1 - Tuesday - November 14, 2017			
Start Time	End Time	Title	Speakers
7:30 AM	8:30 AM	Registration & Continental Breakfast	
8:30 AM	9:00 AM	Welcome, Overview	Ken Hansen / SRC
		Workshop Goals, Day 1 Agenda	John Oakley / SRC
		Need for ICA Research	Mihail C. Roco / NSF
Cognitive Psychology - Topic Chair : Lama Nachman / Intel			
9:00 AM	9:30 AM	Keynote I - User Experience in World of Intelligent Agents	Bruce Horn / Intel
9:30 AM	10:15 AM	Open Panel Discussion I	Shiwali Mohan / PARC Jonathan Gratch / USC Misha Pavel / Northeastern David Reitter / Penn State
10:15 AM	10:45 AM	Network/Informal Discussions	
Approaches to Artificial Intelligence - Topic Chair : Brent Hailpern / IBM			
10:45 AM	11:15 AM	Keynote II - Holistic Conversational Assistants	Charles Ortiz / Nuance
11:15 AM	12:00 PM	Open Panel Discussion II	Todd Hylton / UCSD Julie Adams / Oregon State David Traum / USC Craig Knoblock / USC
12:00 PM	12:15 PM	Network/Informal Discussions	
12:15 PM	1:00 PM	Lunch	

“Collaborative Machines to Enhance Human Capabilities”

[Feb 7, 2018] ICA-2: Intelligent Cognitive Assistants Workshop Summary and Research Needs

System Architectures and Devices - Topic Chair : Arvind Kumar / IBM			
1:00 PM	1:30 PM	Keynote III - Efficient, Predictive Coding and Thermodynamic Computing	Jeff Krichmar / UC Irvine
1:30 PM	2:15 PM	Open Panel Discussion III	Geoffrey Burr / IBM Ada Poon / Stanford Roosbeh Jafari / Texas A&M Emre Neftci / UC Irvine
2:15 PM	2:45 PM	Network/Informal Discussions	
Data and Modeling - Topic Chair : Danil Prokhorov / Toyota			
2:45 PM	3:15 PM	Keynote IV - Behavioral Model of Road User	Davide Santo / NXP
3:15 PM	4:00 PM	Open Panel Discussion IV	Matthias Scheutz / Tufts Kagan Tumer / Oregon State Daniel Prokhorov / Toyota Anthony Kuh / NSF
4:00 PM	4:30 PM	Network/Informal Discussions	
4:30 PM	5:00 PM	Day 1 Panel Summaries	Topic chairs
5:00 PM	7:00 PM	Network/Social Hour	

Day 2 - Wednesday - November 15, 2017			
Start Time	End Time	Title	Speakers
7:30 AM	8:25 AM	Continental Breakfast	
8:25 AM	8:30 AM	Day 2 Agenda	John Oakley / SRC
Natural Interfacing - Topic Chair : Reid Simmons / NSF			
8:30 AM	9:00 AM	Keynote V - Natural Language Generation for Dialogue	Marilyn Walker / UCSC
9:00 AM	9:45 AM	Open Panel Discussion V	Sarah Shomstein / GWU Ron Kaplan / Amazon Omer Oralkan / NC State
9:45 AM	10:15 AM	Network/Informal Discussions	
Breakout Sessions			
10:15 AM	11:45 AM	Breakout Sessions - Research Needs	Lead by Topic Chair w/ Everyone's participation
11:45 AM	12:30 PM	Lunch	
12:30 PM	1:30 PM	Breakout Session Summaries / Priorities	Lead by Topic Chairs w/ Everyone's participation
1:30 PM	2:00 PM	Cross Alignment of Research Needs	Topic Chairs
2:00 PM	2:15 PM	Wrap-Up	John Oakley / SRC
2:15 PM	2:15 PM	Adjourn	

“Collaborative Machines to Enhance Human Capabilities”

Appendix C: References

[ICA-1] Jon Candelaria et al., “Intelligent Cognitive Assistants Workshop Summary and Recommendations”, SRC-NSF Report, https://www.nsf.gov/crssprgm/nano/reports/2016-1003_ICA_Workshop_Final_Report_2016.pdf

[O1_Oakley] John Oakley, “ICA Workshop Goals” (Day 1, 08:35)
<https://www.src.org/calendar/e006378/e006378-icaworkshop-john-oakley.pdf>

[O1_Roco] Mihail C. Roco, “Need for ICA Research” (Day 1, 08:45)
<https://www.src.org/calendar/e006378/e006378-icaworkshop-mike-roco.pdf>

[K1_Horn] Bruce Horn, “Intelligent Cognitive Assistants with Complementary Learning” (Day 1, 09:00)
<https://www.src.org/calendar/e006378/e006378-icaworkshop-bruce-horn.pdf>

[P1_Mohan] Shiwali Mohan, “Interactive Cognitive Assistants: Design, Deployment, & Analysis” (Day 1, 09:30)

[P1_Gratch] Jonathan Gratch, “Challenges in Artificial Emotional Intelligence” (Day 1, 09:35)
<https://www.src.org/calendar/e006378/e006378-icaworkshop-ionathan-gratch.pdf>

[P1_Pavel] Misha Pavel, “The Ultimate ICA: Approximating Human Intelligence” (Day 1, 09:40)
<https://www.src.org/calendar/e006378/e006378-icaworkshop-misha-pavel.pdf>

[P1_Reitter] David Reitter, “Bounded Cognitive AI and Dialogue” (Day 1, 09:45)
<https://www.src.org/calendar/e006378/e006378-icaworkshop-david-reitter.pdf>

[K2_Ortiz] Charles Ortiz, “Holistic Conversational Assistants” (Day 1, 10:45)
<https://www.src.org/calendar/e006378/e006378-icaworkshop-charles-ortiz.pdf>

[P2_Hylton] Todd Hylton, “On Thermodynamics and The Future of Computing” (Day 1, 11:15)
<https://www.src.org/calendar/e006378/e006378-icaworkshop-todd-hylton.pdf>

[P2_Adams] Julie Adams, “Intelligent Adaptive Assistants” (Day 1, 11:20)
<https://www.src.org/calendar/e006378/e006378-icaworkshop-iulie-adams.pdf>

[P2_Traum] David Traum, “Off to See the Wizard” (Day 1, 11:25)

[P2_Knoblock] Craig Knoblock, “Task Learning to Build Intelligent Cognitive Agents” (Day 1, 11:30)
<https://www.src.org/calendar/e006378/e006378-icaworkshop-craig-knoblock.pdf>

[K3_Krichmar] Jeff Krichmar, “Efficient, Predictive Coding and Thermodynamic Computing” (Day 1, 13:00)
<https://www.src.org/calendar/e006378/e006378-icaworkshop-jeff-krichmar.pdf>

[P3_Burr] Geoffrey Burr, “Are non-von Neumann Architectures Essential?” (Day 1, 13:30)
<https://www.src.org/calendar/e006378/e006378-icaworkshop-geoffrey-burr.pdf>

[Feb 7, 2018] ICA-2: Intelligent Cognitive Assistants Workshop Summary and Research Needs

[P3_Poon] Ada Poon, "What More Can We Learn from Biological Systems That Would Aide Us in Designing ICAs?" (Day 1, 13:35) <https://www.src.org/calendar/e006378/e006378-icaworkshop-ada-poon.pdf>

[P3_Jafari] Roozbeh Jafari, "System Architectures for ICA: Role of Wearables" (Day 1, 13:40) <https://www.src.org/calendar/e006378/e006378-icaworkshop-roozbeh-jafari.pdf>

[P3_Neftci] Emre Neftci, "Neuromorphic Learning Machines" (Day 1, 13:45) <https://www.src.org/calendar/e006378/e006378-icaworkshop-emre-neftci.pdf>

[K4_Santo] Davide Santo, "Autonomous Driving Planning: A Convergence of Artificial Intelligence and Cognitive Science" (Day 1, 14:45) <https://www.src.org/calendar/e006378/e006378-icaworkshop-davide-santo.pdf>

[P4_Prokhorov] Daniel Prokhorov, "How to Cope with Deluge of Data?" (Day 1, 15:15) <https://www.src.org/calendar/e006378/e006378-icaworkshop-danil-prokhorov.pdf>

[P4_Scheutz] Matthias Scheutz, "Human-Level Performance" (Day 1, 15:20) <https://www.src.org/calendar/e006378/e006378-icaworkshop-matthias-scheutz.pdf>

[P4_Tumer] Kagan Tumer, "Objectives as Models: Focus On "What" Not "How"" (Day 1, 15:25) <https://www.src.org/calendar/e006378/e006378-icaworkshop-kagan-tumer.pdf>

[P4_Kuh] Anthony Kuh, "Real-Time Learning for Dynamic Systems" (Day 1, 15:30) <https://www.src.org/calendar/e006378/e006378-icaworkshop-anthony-kuh.pdf>

[K5_Walker] Marilyn Walker, "Natural Language Generation for Dialog Systems" (Day 2, 08:30) <https://www.src.org/calendar/e006378/e006378-icaworkshop-marilyn-walker.pdf>

[P5_Shomstein] Sarah Shomstein, "Attention and 'Things-that-make-us-better'" (Day 2, 09:00) <https://www.src.org/calendar/e006378/e006378-icaworkshop-sarah-shomstein.pdf>

[P5_Kaplan] Ron Kaplan, "The Natural Interface: Conversation" (Day 2, 09:05) <https://www.src.org/calendar/e006378/e006378-icaworkshop-ron-kaplan.pdf>

[P5_Oralkan] Omer Oralkan, "A Hardware Engineer's POV to ICA: Multimodal Interfaces Enabled by Sensors and Actuators" (Day 2, 09:10) <https://www.src.org/calendar/e006378/e006378-icaworkshop-omer-oralkan.pdf>

Ref 1. M.C. Roco and W.S. Bainbridge, "Converging Technologies for Improving Human Performance", NSF-DOC Report, Washington, D.C., 2002, Printed by Springer, 2003, http://www.wtec.org/ConvergingTechnologies/Report/NBIC_report.pdf

Ref 2. OSTP, "A Nanotechnology-inspired grand challenge for future computing. Oct 20, 2015, <http://www.nano.gov/futurecomputing>

"Collaborative Machines to Enhance Human Capabilities"

[Feb 7, 2018] ICA-2: Intelligent Cognitive Assistants Workshop Summary and Research Needs

Ref 3. Computing Community Consortium, “Accelerating science: a grand challenge for AI”, CCC Report, 2016. <https://cra.org/ccc/events/symposium-accelerating-science-grand-challenge-ai/>

Ref 4. W.S. Bainbridge and M.C. Roco, “Science and technology convergence, with emphasis for nanotechnology-inspired convergence”, J. Nanopart. Research, Vol. **18**, Paper 211, 2016, 19p (Case study – Intelligent Cognitive Assistants on pages 13-14)

Ref 5. W.S. Bainbridge and M.C. Roco, “Handbook of science and technology convergence”, Springer Reference, Boston and Berlin, 2016, 1156 pages

Ref 6. Rafael Yuste et al., “Four ethical priorities for neurotechnologies and AI”, Nature **551**, 159–163 (09 November 2017), doi:10.1038/551159a
<https://www.nature.com/news/four-ethical-priorities-for-neurotechnologies-and-ai-1.22960>

Ref 7. Karl Friston, “The free-energy principle: a unified brain theory?” (Jan 13, 2010)
<http://www.fil.ion.ucl.ac.uk/~karl/The%20free-energy%20principle%20A%20unified%20brain%20theory.pdf>

Ref 8. X. Wang and Y. He, “Learning from Uncertainty for Big Data: Future Analytical Challenges and Strategies”, IEEE SMC Magazine, Volume 2, Issue 2 (Apr 2016)

Ref 9. Yang Xu and David Reitter, “Information density converges in dialogue: Towards an information-theoretic model”, Cognition (2018) http://www.david-reitter.com/pub/xu_reitter_cognition2017.pdf

Ref 10. Gerd Altmann, “Girl Forward Digital” (Mar 2017) <https://pixabay.com/en/girl-forward-digital-digitization-2181709/>

Appendix D: Recommended Reading/Viewing

Jennifer Chu, “Engineers design artificial synapse for brain-on-a-chip hardware” (Jan 22, 2018)
<http://news.mit.edu/2018/engineers-design-artificial-synapse-brain-on-a-chip-hardware-0122>

HRLaboratory, “Learning Cognitive Affordances through Natural Language Instructions” (Dec 14, 2017)
<https://www.youtube.com/watch?v=hj1aZYnTRNg>

Cliff Kuang, “Can A.I. Be Taught to Explain Itself?” (Nov 21, 2017)
<https://www.nytimes.com/2017/11/21/magazine/can-ai-be-taught-to-explain-itself.html>

Yann LeCun, “Facebook’s AI Boss: ‘In terms of general intelligence, we’re not even close to a rat’” (Oct 27, 2017) <http://www.businessinsider.com/facebook-ai-boss-in-terms-of-general-intelligence-were-not-even-close-to-a-rat-2017-10>

Eileen Brown, “Most US workers want to see more AI and robots in the office” (Oct 24, 2017)
<http://www.zdnet.com/article/four-out-of-five-workers-want-to-see-more-ai-and-robots-in-the-office/>

Massimiliano Versace, “Does Artificial Intelligence Require Specialized Processors?” (Oct 20, 2017)
<https://thenewstack.io/ai-hardware-software-dilemma/>

Luke Dormehl, “A history of artificial intelligence in 10 landmarks” (Sep 23, 2017)
<https://www.digitaltrends.com/cool-tech/history-of-ai-milestones/>

NSF, “The Rules of Life” (Aug 8, 2017)
https://www.nsf.gov/discoveries/disc_summ.jsp?cntn_id=242752

Elon Musk on AI at National Governors Association 2017 Summer Meeting (Jul 15, 2017)
<https://www.c-span.org/video/?c4676816/elon-musk-artificial-intelligence-a-fundamental-risk-existence-civilization>

Severin Lemaignan, “Artificial cognition for social human–robot interaction: An implementation” (Jun 2017) <https://www.sciencedirect.com/science/article/pii/S0004370216300790>

NSF, “10 Big Ideas for Future NSF Investments” (Mar 2017)
https://www.nsf.gov/about/congress/reports/nsf_big_ideas.pdf

Goncalo Pereira, “Integrating social power into the decision-making of cognitive agents” (Dec 2016)
<https://www.sciencedirect.com/science/article/pii/S0004370216300868>

USCICT, “SimSensei & MultiSense: Virtual Human and Multimodal Perception for Healthcare Support” (Feb 7, 2013) <https://www.youtube.com/watch?v=eiczMs6b1Q4>