Report on Cognitive Neuroscience Roadmapping Workshop

Held on May 11, 9:30 to 1:30

Purpose of Workshop

The purpose of this workshop was to bring together a cross-section of intellectual leaders across the Institute of Cognitive Science to construct a research roadmap outlining strategic areas for future interdisciplinary cognitive neuroscience research that ICS members and partners are uniquely poised to carry out. This event is intended to lead to the formation of a few teams that want to tackle one or more of these areas together. A secondary goal is to find out what types of support (financial, administrative, moral, etc.) ICS members and partners need in order to embark on this type of ambitious interdisciplinary research. Towards this end, participants were asked to discuss the following questions:

1. What are the "high leverage" challenges or research questions that you would like to investigate over the next 5 years? How are these questions or topics related to each other, or not?
2. For each high leverage area, what do we need to "figure out" to better understand this challenge or scientific phenomena? What are our working hypotheses? What sorts of expertise, data or facilities are needed to make progress? Are there other people or partners we need to recruit?

The definition of a “high leverage” challenge or research question was based on Bryk et al (2016)[1]. High leverage challenges or research questions are those that will significantly impact lives, dramatically influence needed resources to address problems, or contribute to our understanding of variability in critical outcomes.

Workshop Participants

Alaa Ahmed	Marta Ceko	Al Kim
Yoni Ashar	Tim Curran	Philip Kragel
Marie Banich	Phillip Gilley	Tammy Sumner
Cinnamon Bidwell	Pavel Goldstein	Tor Wager
McKell Carston	Matt Jones	Karli Watson

Workshop Processes

All workshop materials are available in our google drive folder. Participants were provided with a study packet to review prior to the workshop. This packet contained overview materials (slide decks, short documents) summarizing research taking place in ICS and INC and outlining new opportunities in the area of embedded and embodied cognition. The purpose of these documents was to provide a common starting point for participants without having to devote significant portions of the workshop to introductions and presentations.

After brief introductions, our time was spent in a series of idea generation activities (small groups) and synthesis discussions (whole group). In the first half, we split into three groups where each group was tasked with discussing question 1 and developing one or more fishbone diagrams depicting how the identified topics, questions, challenges related. After 20 minutes, we rotated and repeated this exercise with new group members. We then spent 50 minutes as a whole group discussing and comparing the resulting diagrams. In the second half, due to time and energy constraints, we modified the agenda to only have one longer round of group discussions, followed by a longer period for lunch and synthesis. These discussions focused on question

[1] Learning to Improve, by Bryk, Gomez, Grunow, and LeMahieu, Harvard Education Press, 2016.

2 and groups were provided with a recording template. Our final whole group discussion considered what resources or types of support we need to conduct this research; we also completed individual surveys.

Outcomes - Question #1

We identified seven high leverage challenges or research strands:

1. *Mindset*: understanding how prior knowledge and expectations influence cognition, performance, and behavior change
2. *Meaning making*: understanding individual differences
3. *Learning Optimization*: how can we make learning more efficient and effective, how can we predict when something has been learned, how do emotions and affect influence learning, can we predict robust learning (learning that generalizes)
4. *Social Fairness*: what factors contribute towards helping people to be more resilient to social injustice, how does social fairness and related factors contribute towards social pain or chronic pain, how do these factors influence learning
5. *Human optimization*: how can we build on next generation wearables, functional skin, and other human augmented computational environments to optimize learning and behavior change, what types of industry partnerships are needed or enabled, how to link our understanding of neural circuity with behavior, what new approaches to machine learning are needed to work with smaller data sets
6. *Improving well-being by reducing chronic pain*: what are the risk factors, what is the influence of environment factors such as social unfairness, or mindsets such as openness of mind, can we predict outcomes or reduce drug dependencies
7. *Changing behavior to improve health*: how can new devices such as wearables be leveraged to provide more ecologically valid measures or to provide more active prompting and feedback to support behavior change, can we track individual differences to provide more personalized support, how can we partner with manufacturers to provide quality, can we use data from these devices to better understand extended cognition

White board depictions of each corresponding fishbone diagram are below. You can click on the picture to see larger versions!

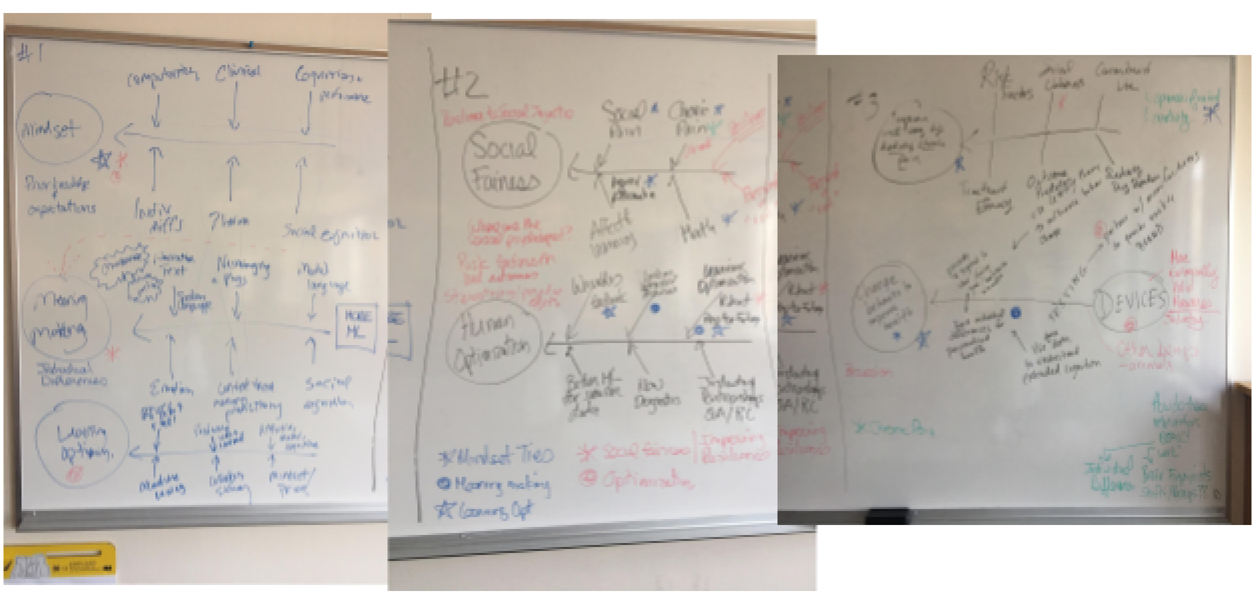

As you can see, there are significant overlaps between between these seven areas. For the second question, we decided to focus on integrating our ideas around two: Mindset and Human Optimization

Outcomes - Question #2

Report from Mindset Group (Lead writers: Tor, Al, and Phillip G.)

The mindset group focused on the role of knowledge and expectations, or "mindset", on sensory, perceptual, and cognitive processing.

1. *Mindset*
2. *Meaning making*
3. *Improving well-being by reducing chronic pain*
4. *Changing behavior to improve health*:

We discussed the idea that cognition is profoundly shaped by prior knowledge and expectations, and we used the term "mindset" to unify the impacts of knowledge and expectations across a wide range of psychological domains that many of us are working in.

We discussed the relationship between mindset and the growing theoretical influence in cognitive science that the brain is a prediction machine. This theoretical framework posits that brains generate predictions about the future state of the world. Predictions allow a brain to keep up with a fast-changing world by pre-activating representations that it will need in the future. Predictions are also checked against the world, yielding prediction error, which serves as a signal that drives learning. The central role of prediction in models of cognition leads to a tight merger of models of memory with models of perception. Like mindset, prediction may also be a rubric that unifies the role of knowledge and expectations across many of our domains of research.

We discussed important clinical implications of mindset in understanding people's approach to pain, drugs, and happiness. People's' perception of pain is strongly modulated by their expectations. Mindset strongly modulate the body's response to drugs--experimental work shows that death following large doses of heroin in heroin-experienced rats occurs at much higher rates if the drug intake occurs in an unfamiliar environment as opposed to a familiar environment (e.g., new cage vs. home cage). This effect suggests that expectations based on context allow the brain to prepare for the drug's impact. Mindset affects how people perceive themselves and others and can lie at the root of depression, when people have inaccurate expectations about how they are perceived.

We discussed the idea of a large-scale study that uses brain activity during language comprehension as a window into the role of mindset in cognition. The study would record brain activity using fMRI and EEG, while participants read and listen to naturalistic texts. We would design texts that deliver meanings of different sorts--emotional, social, and semantic--at different points within the text. By identifying the patterns of neural activity elicited by those different portions of the text, we would characterize the neural circuits associated different domains of cognition. Thus, we would be able to conduct a comprehensive investigation of multiple different dimensions of cognition within a single study, with a single group of participants. The success of this project relies on the flexible ability of language to deliver meanings of many sorts to a comprehender. We refer to this as an important example of “meaning making”. This project would also involve a crucial role for computational linguistics, which can provide quantitative models of meaning at each word of a text.

We discussed the importance of characterizing inter-individual differences in mindset and meaning making. Although cognitive science research has historically studied human participants as representative of the human brain in general, it is increasingly clear that substantial individual differences in cognition exist within the human population. Understanding the systematic sources of individual differences in cognition is essential to a compelling cognitive science. Mindset may be particularly susceptible to important individual differences. This is obviously true of clinical phenomena, which are fundamentally about some individual brains functioning differently from others. But it may also be true of many cognitive functions, such as predictions about language, which depend on lifetimes of linguistic experiences that vary across people.

At several points in our discussions, we discussed the role of machine learning techniques in developing models of cognition and in the analysis of complex neurophysiological (e.g., EEG) and neuroimaging data (e.g., fMRI). We agreed that while it is clear that multiple aspects of our work would benefit from collaboration with experts ML techniques, the success of such collaborations depends on identifying ML experts who are also sophisticated about models of the mind and brain and motivated to engage deeply with a specific topic in cognitive science.

Report from Human Optimization Group (Lead writers: Tim, Tammy, Pavel)

The Human Optimization Group focused on learning as the primary common theme across high level challenges:

3 – *Learning Optimization*,
5 – *Human Optimization*, and
7 – *Changing behavior to improve health*.

We discussed the hypothesis that we need to look beyond the individual to understand behavior, with a particular emphasis on a consideration of social context. This lead to discussions of being able to study learning in naturalistic social settings, and the importance of using passive, unobtrusive wearable devices to obtain relevant measurements that could include cameras, EEG, and other physiological measurements such as glucose as well as video and audio processing. Can these measurements be used to predict and optimize learning? Other social context issues discussed included the importance of social engagement to encourage compliance and the importance of communicating to the end user in an understandable way that “feels right” to the user.

We discussed the importance of considering multiple forms of learning: education, math, motor skills, clinical outcomes, etc. We also emphasized the understanding of robust forms of learning that generalize beyond the specific training/learning experiences.

We discussed the importance of feedback for learning along with understanding the appropriate type and timing of feedback.

We discussed potential partners:

- Academic:
  - Education
  - Philosophy
  - Computer Science
    - Automatic speech recognition
    - Machine Learning
    - Mike Mozer – Optimization of training/spacing schedules
    - Mike Eisenberg – wearable devices
    - Tam Vu – Devices to optimize sleep quality

- Industry:
  - Sparkfun - www.sparkfun.com/
  - Lena Foundation – NLP - https://www.lena.org
  - Peter Foltz – NLP - analyzing speech patterns to predict Schizophrenia
  - NLP chip company in Boulder

- Other Communities:
  - Biohacker groups
  - Speculative Fiction

Outcomes - Resources and Support Needed

In our group discussion, most of the “needs” coalesced around expertise in machine learning, data mining, and other computational approaches. All seven core areas identified in this workshop require partners with considerable expertise in computational approaches. Several approaches towards building up or seeking out this expertise were identified:

- Hire a research support scientist
- Increase our understanding of existing services (AWS, platforms like TensorFlow, google analytics)

- Use these services to create a “ready to go” environment for targeted research to reduce the entry barrier to getting started
- Use kaggle to outsource algorithm development
- Develop and offer a summer “boot camp” on computational approaches to cognitive science to develop buzz, expertise, and a pipeline of students to participate
- Partner with CS faculty to embed cognitive science examples into machine learning, data mining, and natural language processing courses

Other ideas that came up during our discussion include seeking seed funding for shared projects, working with Terri Fiez to identify industry partners, and developing partnerships with faculty from computer science.

From the survey responses, we learned that participants thought the workshop was helpful for generating ideas, meeting new people, and building and strengthening relationships. Respondents report being motivated to participate by their scientific curiosity, potential new funding streams, the relevance of the topic to their own interests, and the opportunity to contribute towards a larger team-oriented goal.

Barriers to participation were also reported. As expected, time, or lack thereof, is a key barrier. Another potential barrier was lack of leadership to take a team or project forward: many people want to participate but few want to lead! Other respondents note that they do not have the background to lead as their work is a nice complement but not directly related to the core ideas. A few respondents expressed concerns, rather than barriers, such as the need for grants emerging from this process to “pan out” to keep the motivation and momentum going, or at an even more basic level, to ensure that our group processes pan out and lead to substantive outcomes.

In terms of how the Institute can help, participants recommended three things: developing computational support, providing funds for pilot studies, and fostering more opportunities for casual scientific conversations between ICS member.