

Language & Common Sense

Integrating across psychology, linguistics, and computer science

Joshua K. Hartshorne (jharts@mit.edu) & Joshua B. Tenenbaum (jbt@mit.edu)
Department of Brain & Cognitive Sciences, 77 Massachusetts Avenue
Cambridge, MA 02139 USA

Keywords: language; common sense; world knowledge; pragmatics; intuitive theories; cognitive architectures; natural language processing

Introduction

The language understanding that underlies state-of-the-art Internet search, machine translation, and dictation software is undeniably impressive. Equally undeniable is that these systems do not really understand language. What is missing? One candidate is common sense. Language is a mechanism for moving ideas from one mind to another – ideas that are meant to be understood in the context of pre-existing, interlocking beliefs. Thus, fully exploiting its power may require sophisticated, explicit representations of world knowledge – that is, common sense.

That understanding language requires deploying knowledge about the world is not a new observation (cf. Winograd, 1972). However, new opportunities for significant progress have been suddenly opened up by recent, rapid advances in the science of common sense, along with related advances in machine vision, natural language processing, and computational tools for developing more precise cognitive theories (Liang & Potts, 2015; Sonka, Hlavac, & Boyle, 2014; Tenenbaum, Kemp, Griffiths, & Goodman, 2011).

This workshop brings together researchers from across the cognitive sciences – including developmental and cognitive psychology, linguistics, natural language processing, artificial intelligence, and robotics – to disseminate recent findings, discuss approaches to future progress, and set an agenda for the field. Given the interdisciplinary nature of the participants and of the research challenges faced, the Annual Meeting of the Cognitive Science Society is an ideal venue for these conversations.

Goals and Scope

Language, as a vehicle for conveying thoughts and beliefs, is highly compressed. In many cases, resolving the resulting ambiguity seems to require knowledge of the world. Our knowledge of summertime activities suggests that the *bank* in *Sally dove off the bank into the river* is probably not a financial institution. Our knowledge of current market prices implies that *these pencils cost \$100* is probably an exaggeration for emotional effect (Kao, Wu, Bergen, & Goodman, 2014). That *the elephant on the wall* probably describes a hanging picture whereas *the elephant on the truck* describes a translocating pachyderm may be inferred from our knowledge of physics and zoology.

These observations leave open *how* world knowledge is incorporated in language understanding. Do we make use of

online, domain-general reasoning, or do we have language-specific strategies and representations? To the degree that online reasoning is involved, how can it be characterized (e.g., as rote heuristics, as inference over intuitive theories, as distributional probabilities, etc.)?

Related issues arise in relating language to perception. We talk about what we see, but again the information is highly compressed: *The elephant was on the truck* is consistent with a staggering range of visual input. How do the semantic representations of language relate to the mid- or high-level representations used in perception? What role does knowledge of the world play (elephants are more likely to be on certain parts of certain trucks)?

The different cognitive sciences have approached different aspects of these questions in different ways. This workshop brings together ten researchers from across the cognitive sciences to disseminate and discuss successes (and challenges), and also to help formulate an agenda for the field: What phenomena and challenge problems should be explored? How can progress in formalizing pragmatics and semantics be used to formalize theories of common sense (and *vice versa*)? Which questions raised in linguistics may have answers in psychology or computer science (and *vice versa*)?

Winograd Schema

In addition to presentations covering a range of topics at the intersection of language and common sense, the workshop will contain a special session on Winograd Schema, so named for Winograd's (1972) classic demonstration of commonsense reasoning's influence on pronoun interpretation:

- (1) The city council denied the protesters a permit because they **advocated** violence.
- (2) The city council denied the protesters a permit because they **feared** violence.

Most readers agree that the pronoun *they* refers to the protesters in (1) but the city council in (2). This seems to derive from our understanding of city councils, protesters, and the permitting process: City councils rarely advocate violence, and even if they did, that would be poor reason for them to deny protesters permits.

Winograd Schema like (1-2) are of growing interest in artificial intelligence, where it has been suggested that they provide a sophisticated, alternative Turing Test, given that they involve both language and common sense (Levesque, Davis, & Morgenstern, 2012). They also provide an enticing opportunity for cross-disciplinary dialog: There is a deep, robust literature on such phenomena in both psychology and linguistics, including (recently) computational models. In keeping with the overall theme of the workshop, this session

will consist of presentations by a computer scientist (Charles Ortiz), a linguist (Andrew Kehler), and a psychologist (Joshua Hartshorne).

Workshop Organization

The workshop will be organized around a set of thirty-minute presentations (including Q&A) and panel discussions. The presentations will range from theoretical overviews to detailed discussion of specific phenomena. The panel discussions and coffee breaks will help spur discussion about promising avenues for future research and help build a common vocabulary and agenda.

Workshop Organizers

Joshua K. Hartshorne is a Ruth L. Kirschstein NRSA post-doctoral fellow in the Department of Brain and Cognitive Sciences at MIT and an incoming assistant professor at in the Department of Psychology at Boston College. His research focuses on how cognitive representations constrain and inform language, both in online processing (Hartshorne, O'Donnell, & Tenenbaum, in press) and during development (Hartshorne, Pogue, & Snedeker, in press). **Joshua B. Tenenbaum** is Professor of Cognitive Science and Computation at MIT. His recent work focuses on computational models of commonsense reasoning and intuitive theories (Tenenbaum et al., 2011; Battaglia, Hamrick, & Tenenbaum, 2013).

Target Audience

The target audience for this workshop overlaps significantly with the target audience of CogSci. The workshop incorporates two themes (language and commonsense reasoning) that are central across many of the cognitive science disciplines (artificial intelligence, linguistics, psychology, etc.). Moreover, the workshop approaches these themes from a multidisciplinary perspective, as seen in the disciplinary diversity of the participants. Because the presentations will be geared towards an interdisciplinary audience, they should be approachable by a broad cognitive science audience.

Confirmed Speakers

[David Barner](#)

Department of Psychology
University of California-San Diego

Nancy Chang

Google

[Noah D. Goodman](#)

Department of Psychology
Stanford University

[Joshua K. Hartshorne](#)

Department of Brain & Cognitive Sciences
Massachusetts Institute of Technology

[Andrew Kehler](#)

Department of Linguistics
University of California-San Diego

[Percy Liang](#)

Department of Computer Science
Stanford University

[Charles L. Ortiz, Jr.](#)

Natural Language and Artificial Intelligence Laboratory
Nuance Communications

[Jeffrey M. Siskind](#)

School of Electrical and Computer Engineering
Purdue University

[Stefanie Tellex](#)

Computer Science Department
Brown University

[Joshua B. Tenenbaum](#)

Department of Brain & Cognitive Sciences
Massachusetts Institute of Technology

References

- Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *PNAS*, 110, 18327-332.
- Hartshorne, J. K., O'Donnell, T. J., & Tenenbaum, J. B. (*in press*). The causes and consequences explicit in verbs. *Language, Cognition, and Neuroscience*.
- Hartshorne, J. K., Pogue, A., & Snedeker, J. (*in press*). Love is hard to understand: The relationship between transitivity and caused events in the acquisition of emotion verbs. *Journal of Child Language*.
- Kao, J. T., Wu, J.Y., Bergen, L., & Goodman, N. D. (2015). Nonliteral understanding of number words. *PNAS*, 111, 12002-07.
- Liang, P., & Potts, C. (2015). Bringing machine learning and compositional semantics together. *Annual Review of Linguistics*, 1, 355-376.
- Levesque, H. J., Davis, E., & Morgenstern, L. (2012). The Winograd Schema challenge. *Proceedings of Commonsense Reasoning 2012*.
- Sonka, M., Hlavac, V., & Boyle, R. (2014). *Image processing, analysis, and machine vision*, 4th Edition, Stamford, CT: Cengage Learning.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331, 1279-1285.
- Winograd, T. (1972). Understanding natural language. *Cognitive Psychology*, 1-191.