

Strategies and representations in physical inference

Kevin A Smith (k2smith@mit.edu), Joshua B Tenenbaum (jbt@mit.edu)

Department of Brain and Cognitive Sciences, MIT

**Erin Anderson (erinanderson2014@u.northwestern.edu), Susan Hesospos (hespos@northwestern.edu),
Lance Rips (rips@northwestern.edu)**

Department of Psychology, Northwestern University

Chaz Firestone (chaz@jhu.edu)

Department of Psychological and Brain Sciences, Johns Hopkins University

Jessica B Hamrick (jhamrick@google.com)

DeepMind

Keywords: Physical Inference; Representation; Perception; Logical Rules; Simulation; Meta-reasoning

Overview

Human cognition is not thought to be a monolithic entity, but instead is often considered to include a range of strategies and representations that can be applied selectively in different domains. For instance, we are thought to have both intuitive and deliberative processes for decision making (Kahneman, 2011) and both approximate and precise representations of numeracy (Feigenson, Dehaene, & Spelke, 2004). But how do these different strategies and representations develop, and how do we select which ones to use?

In this symposium, we will discuss this question in the domain of physical inference. We use this domain because it has been well studied across many branches of cognitive science – including development, perception, neuroscience, and artificial intelligence – and researchers from across these fields have proposed a wide range of strategies and representations that support this ability. Some suggest that our physical inferences are based on simulating possible outcomes using mental models of the world (Battaglia, Hamrick, & Tenenbaum, 2013; Shepard & Metzler, 1971; Smith & Vul, 2013). Others suggest that we apply logical reasoning to discover what is entailed from the relationships between objects (Davis, Marcus, & Chen, 2013; Forbus, 1983). And yet others suggest that we can extract physical information from a scene using bottom-up perceptual processes (Biederman, Mezzanotte, & Rabinowitz, 1982; Firestone & Keil, 2016; Firestone & Scholl, 2017).

This symposium brings together experts in physical inference with backgrounds in developmental psychology, psychophysics, and machine learning for the goal of understanding the different strategies and representations the mind uses for physical scene understanding. These speakers will discuss questions such as how different representations develop in infancy (*Anderson*), how we can distinguish between cognitive mechanisms for physical inference (*Firestone*), how we combine different strategies in our physical judgments (*Smith*), and how we allocate cognitive resources to those strategies (*Hamrick*). In a concluding panel, the speakers will discuss how to build a theory of physical inference that combines these different strategies and representations.

Infants' reasoning in dual physical domains: Physical inference without the instruction manual

Erin Anderson, Susan Hesospos, & Lance Rips

To understand the basis of our physical inferences, we can study infants' reactions to possible and to utterly impossible events. These reactions lay the foundation for how humans represent everyday entities and substances. Decades of research on infant cognition have shown a puzzling imbalance: infants expect principled behavior from solid objects but seem to have no expectations (or highly undependable ones) for non-solid substances like sand (Chiang & Wynn, 2000; Huntley-Fenner, Carey, & Solimando, 2002).

In this talk, we will discuss work that confirms that infants do represent both solid objects and nonsolid substances like liquids or sand. First, we present experiments that show that 5-month-old infants expect objects and substances to behave in different ways (Hesospos, Ferry, & Rips, 2009; Hesospos, Ferry, Anderson, Hollenbeck, & Rips, 2016). We further demonstrate that infants have expectations specific to substances when the outcomes are not contrasted with object behavior (Anderson, Hesospos, & Rips, 2018). Together, these studies provide evidence that the guiding principles for substances and solids develop in parallel over the first year, and they raise questions about when and how infants co-ordinate reasoning across these domains.

Doing physics by eye and by hand

Chaz Firestone

When we appreciate that a stack of dishes will collapse or a tower of blocks will topple, we interpret observable events in terms of unobservable physical forces. How? Classic and contemporary work on this question typically treats such inferences as species of *higher-level cognition*, akin to solving physics riddles by reasoning and calculation. Here I explore a very different possibility: that such inferences are rooted in automatic visual processing. I discuss evidence that, just like visual processing of color or motion, physical inferences are (a) spontaneous, (b) fast, (c) attention-grabbing, and (d) phenomenologically distinctive (Firestone & Scholl, 2016, 2017). I extend this work to show that such inferences are

also (e) *reflexive*, and surprisingly intransigent to explicit beliefs or instructions. Recent work suggests that physical inferences flexibly incorporate new information about a scene, such as an instruction to treat gray blocks in a tower as 10x heavier than green blocks (Battaglia et al., 2013). However, I show that when observers evaluate such scenes by making continuous mouse movements (Freeman, Dale, & Farmer, 2011), their response trajectories reveal otherwise: even observers who successfully update their considered judgments about a scene are nevertheless ‘pulled’ toward initial judgments that resist this new information. I suggest that physical understanding may not be a single process, but rather one with dissociable stages: a fast, reflexive, “perceptual” stage, and a slower, flexible “cognitive” stage.

Integrating logical rules and simulation

Kevin A Smith & Joshua B Tenenbaum

If we have multiple strategies for performing physical inference tasks, then how do we know which strategy to apply to which task? Do we simply pick a single cognitive strategy, or do we flexibly blend different strategies to solve our problems? I will discuss this choice of strategies in the case of judging whether and how a balance beam will tip – a task that has often been considered to be solved by rule-based decision trees (Siegler, 1976), but shares many features with a task that has been used to demonstrate simulation-based inference (Battaglia et al., 2013). By asking for judgments about how beams stacked with objects of a variety of shapes and materials behave, I demonstrate that people’s judgments cannot be explained by traditional rule-based accounts, but neither can they be explained by simulation alone. Instead, people use a combination of rules and simulation to make inferences about balance beam stability. I further show the mixture of strategies applied to this problem is consistent with a framework in which people rationally trade off between the costs and expected benefits of each strategy to choose how to approach this problem of physical inference.

Meta-reasoning for adaptive physical strategy selection and control

Jessica B Hamrick

One of the most powerful and flexible aspects of cognition is that of mental simulation, which is the mind’s ability to imagine seeing, interacting with, and manipulating objects and scenes, almost as if they were real. While research has shown how the mind is able to use mental simulation to reason about a wide range of physical domains (Battaglia et al., 2013), other research suggests that mental simulation is not always the most efficient strategy and that people do indeed learn alternate strategies and switch between them (Schwartz & Black, 1996). How do people know what strategy will be appropriate for which situation, without actually executing the strategy? And, how do people learn different strategies in the first place? Research on how people choose between strategies more broadly has suggested meta-reasoning as a

framework for strategy selection, which involves choosing the strategy which is expected to best satisfy a speed/accuracy tradeoff (Lieder et al., 2014). Here, I will argue that meta-reasoning can be applied to physical reasoning strategies as well, and will describe a system that learns through experience which strategies to use and how much computation to allocate to each strategy (Hamrick et al., 2017). I will also suggest that people may learn new strategies both through experience as well as through mental simulation (Callaway, Hamrick, & Griffiths, 2017).

References

- Anderson, E. M., Hespos, S. J., & Rips, L. J. (2018). Five-month-old infants have expectations for the accumulation of non-solid substances. *Cognition*, 175, 1-10.
- Battaglia, P., Hamrick, J., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, 110(45), 18327–18332.
- Biederman, I., Mezzanotte, R. J., & Rabinowitz, J. C. (1982). Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive Psychology*, 14(2), 143–177.
- Callaway, F., Hamrick, J., & Griffiths, T. L. (2017). Discovering simple heuristics from mental simulation. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. Davelaar (Eds.), *Proceedings of the 39th Annual Meeting of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.
- Chiang, W.-C., & Wynn, K. (2000). Infants’ tracking of objects and collections. *Cognition*, 77(3), 169–195.
- Davis, E., Marcus, G. F., & Chen, A. (2013). Reasoning from radically incomplete information: The case of containers. *Advances in Cognitive Systems*, 2, 273–288.
- Feigenson, L., Dehaene, S., & Spelke, E. (2004). Core systems of number. *Trends in Cognitive Sciences*, 8(7), 307–314.
- Firestone, C., & Keil, F. C. (2016). Seeing the tipping point: Balance perception and visual shape. *Journal of Experimental Psychology: General*, 145(7), 872.
- Firestone, C., & Scholl, B. J. (2016). Seeing stability: Intuitive physics automatically guides selective attention. *Journal of Vision*, 16(12), 689.
- Firestone, C., & Scholl, B. J. (2017). Seeing physics in the blink of an eye. *Journal of Vision*, 17(10), 203.
- Forbus, K. D. (1983). Qualitative reasoning about space and motion. In D. Gentner & A. Stevens (Eds.), *Mental models* (p. 53-73). New Jersey: LEA Associates, Inc.
- Freeman, J., Dale, R., & Farmer, T. (2011). Hand in motion reveals mind in motion. *Frontiers in Psychology*, 2, 59.
- Hamrick, J. B., Ballard, A. J., Pascanu, R., Vinyals, O., Heess, N., & Battaglia, P. W. (2017). Metacontrol for adaptive imagination-based optimization. In *ICLR 2017*.
- Hespos, S. J., Ferry, A. L., Anderson, E. M., Hollenbeck, E. N., & Rips, L. J. (2016). Five-month-old infants have general knowledge of how nonsolid substances behave and interact. *Psychological Science*, 27(2), 244–256.
- Hespos, S. J., Ferry, A. L., & Rips, L. J. (2009). Five-month-old infants have different expectations for solids and liquids. *Psychological Science*, 20(5), 603–611.
- Huntley-Fenner, G., Carey, S., & Solimando, A. (2002). Objects are individuals but stuff doesn’t count: Perceived rigidity and cohesiveness influence infants’ representations of small groups of discrete entities. *Cognition*, 85(3), 203–221.
- Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Straus and Giroux.
- Lieder, F., Plunkett, D., Hamrick, J. B., Russell, S. J., Hay, N., & Griffiths, T. (2014). Algorithm selection by rational metareasoning as a model of human strategy selection. In *Advances in Neural Information Processing Systems* (pp. 2870–2878).
- Schwartz, D. L., & Black, J. B. (1996). Shuttling between depictive models and abstract rules: Induction and fallback. *Cognitive Science*, 20(4), 457–497.
- Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171(3972), 701-703.
- Siegler, R. S. (1976). Three aspects of cognitive development. *Cognitive Psychology*, 8(4), 481–520.
- Smith, K. A., & Vul, E. (2013). Sources of uncertainty in intuitive physics. *Topics in Cognitive Science*, 5(1), 185–199.