Cognitive Psychology 105 (2018) 9–38

Contents lists available at ScienceDirect

Cognitive Psychology

journal homepage: www.elsevier.com/locate/cogpsych

Intuitive experimentation in the physical world[☆,☆☆]

Neil R. Bramley[a,*], Tobias Gerstenberg[b], Joshua B. Tenenbaum[b], Todd M. Gureckis[a]

[a] *Department of Psychology, NYU, New York, NY, United States*
[b] *Center for Brains, Minds and Machines, MIT, Cambridge, MA, United States*

A R T I C L E I N F O

Keywords:
Active learning
Mental simulation
Experimental design
Physical understanding

A B S T R A C T

Many aspects of our physical environment are hidden. For example, it is hard to estimate how heavy an object is from visual observation alone. In this paper we examine how people actively "experiment" within the physical world to discover such latent properties. In the first part of the paper, we develop a novel framework for the quantitative analysis of the information produced by physical interactions. We then describe two experiments that present participants with moving objects in "microworlds" that operate according to continuous spatiotemporal dynamics similar to everyday physics (i.e., forces of gravity, friction, etc.). Participants were asked to interact with objects in the microworlds in order to identify their masses, or the forces of attraction/repulsion that governed their movement. Using our modeling framework, we find that learners who freely interacted with the physical system selectively produced evidence that revealed the physical property consistent with their inquiry goal. As a result, their inferences were more accurate than for passive observers and, in some contexts, for yoked participants who watched video replays of an active learner's interactions. We characterize active learners' actions into a range of micro-experiment strategies and discuss how these might be learned or generalized from past experience. The technical contribution of this work is the development of a novel analytic framework and methodology for the study of interactively learning about the physical world. Its empirical contribution is the demonstration of sophisticated goal directed human active learning in a naturalistic context.

What makes physics physics is that experiment is intimately connected to theory. It's one whole.

— LENE HAU

Much of what we believe about the world, we infer from passive observation and inductive reasoning. For example, if we see that the ground is wet, we might infer that it has been raining. However, we also continuously shape our experience by actively interacting with the world. To determine if a container holds water or sand, we might shake it and observe the resulting forces and sounds. From a causal learning perspective, our actions can be seen as *interventions* (Pearl, 2000) that help reveal how the world works. As such, our everyday actions can share some of the characteristics of scientific experiments: comparing the outcomes of different manipulations while controlling for confounding factors. For example, we might lift two suitcases to judge which is heavier, perhaps switching sides to control for our hand-dominance; drop a rock, or shout, down a well to judge how deep it is; or bounce a squash ball

☆ The authors thank to Jack Valenti and Victor Wang for help with video coding and David Lagnado, Anselm Rothe and Tomer Ullman for useful comments, and Hongyi Zhang for initial code. Preliminary analysis of Experiment 1 appeared in a non-archival conference paper Bramley, Gerstenberg, and Tenenbaum (2016).

☆☆ NB is supported by a Moore Sloan Data Science Environment postdoc position at NYU as well as a James S. McDonnell Scholar Award to TMG. TG and JT are supported by the Center for Brains, Minds & Machines (CBMM), funded by NSF STC award CCF-1231216 and by an ONR grant N00014-13-1-0333. TMG is supported by BCS-1255538 from the National Science Foundation and the John S. McDonnell Foundation Scholar Award.

* Corresponding author.
E-mail addresses: neil.bramley@nyu.edu (N.R. Bramley), tger@mit.edu (T. Gerstenberg), jbt@mit.edu (J.B. Tenenbaum), todd.gureckis@nyu.edu (T.M. Gureckis).

https://doi.org/10.1016/j.cogpsych.2018.05.001
Accepted 11 May 2018
0010-0285/ © 2018 Elsevier Inc. All rights reserved.

to estimate if it is warm enough for play. A key aspect of such behaviors is that they combine an intuitive understanding of how the physical world works, with actions that exaggerate, isolate, or bring into sharper relief a particular physical property of interest. Sometimes we perform these everyday experiments ourselves, while at other times we learn from observing others performing similar actions. Often, success in these endeavors is contingent on acting (or watching someone act) in appropriately "experimental" ways.

In this paper, we investigate how people learn about latent physical properties when interacting with virtual "microworlds". The microworlds are simulated environments on a computer screen. In these worlds, objects' movements are determined by a physics engine that approximates real-world physical laws (Ullman, Spelke, Battaglia, & Tenenbaum, 2017). In our experiments, we allow participants to freely grab and move the objects in the microworld. Our classification and quantification of their action strategies gives new insight into how people decide to act in the physical world to reveal information. Although it seems intuitively obvious that people engage in systematic behaviors during learning, it remains unclear how they decide which strategies to invoke, how complex they are, and how informative they are compared to other things they could have done. We presume people are effective because this matches our intuition, but it is an important scientific question to formally quantify and describe these abilities.

While active inference and physical exploration can be studied in naturalistic settings (Karasik, Tamis-LeMonda, & Adolph, 2011; Kretch & Adolph, 2017; Piaget, 1936; Stahl & Feigenson, 2015), it is difficult to accurately measure and control all aspects of natural environments, and to measure participants' actions at a fine grained level. The current studies are designed to leverage some of the unique advantages of observing learning behavior in a virtual environment while also exploring a setting that inherits some of the complexity and dynamics that make the real world a challenging learning domain. Our virtual environments allow us to (1) precisely record and reconstruct every aspect of a learners' interactions and everything else that occurred during a trial (Rieber, 1996); and (2) develop formal models to quantify and objectively evaluate the information content of people's actions. Using this approach we are able to analyze in detail the types of actions people decide to use and how much information they generate for a given goal, helping to better understand the nature of our intuitive physical interactions.

The paper is structured as follows. We first lay out a normative framework for inference about latent physical properties of dynamically interacting objects, and show how we can use this framework to assess the informativeness of actions. We then describe two experiments that compare the inferential accuracy of active learners (who exert control over the objects in the microworld) with passive learners (who simply watch a movie of the microworld without interacting) and yoked learners (who watch videos of a previous active learner's sessions). In addition, we categorize active participants' experimental strategies with the help of our model-based information measures.

To foreshadow, across both experiments we find that active learners use sophisticated control to create situations that are highly informative about target properties (the properties they are incentivised to identify) while minimizing confounding information about other non-target properties of the worlds. We conclude by discussing the scope of our findings more broadly. Specifically, we discuss how physical active learning strategies might be discovered, reinforced across instances, and generalized across contexts; and discuss the important connections between spatiotemporally extended active learning and adaptive control.

1. Active learning in discrete versus continuous and dynamic environments

In studies of "active learning", people shape their learning experience through their own actions (Gureckis & Markant, 2012). To date, active learning has primarily been studied in situations in which the learner's goal is to differentiate between a relatively small number of discrete hypotheses, such as the Wason card selection task (Oaksford & Chater, 1994; Sperber, Cara, & Girotto, 1995; Wason, 1968), category rule learning (Gureckis & Markant, 2009) and games like "Guess Who" (Nelson, Divjak, Gudmundsdottir, Martignon, & Meder, 2014), "Mastermind" (Berghman, Goossens, & Leus, 2009; Best, 1990; Hofer & Nelson, 2016; Goodrich, 2009) or "Battleships" (Gureckis & Markant, 2009; Markant & Gureckis, 2014b, 2012). In these scenarios, participants pick from a fixed set of possible actions or questions in service of a learning goal, usually with each action-outcome pair contributing independently to a set of evidence they can use to make judgments. A subset of this research, on "active causal learning", studies how people infer the underlying causal structure of simple dynamic systems (Bramley, Dayan, Griffiths, & Lagnado, 2017; Bramley, Lagnado, & Speekenbrink, 2015; Coenen, Bramley, Ruggeri, Gureckis, & Todd, 2017; Coenen, Rehder, & Gureckis, 2015; Lagnado & Sloman, 2002, 2004, 2006; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003). In a typical setting, participants perform interventions that manipulate variables within a causal system, and subsequently observe the consequences of their actions on the other variables (Bramley et al., 2017). For example, Coenen et al. (2015) had learners test computer chips to identify which of several possible wiring diagrams correctly described them. They could test them by activating one of the three components and observing whether either or both of the other two components activated as a result.

Evidence for the utility of active learning in these studies is mixed. Many studies find active learners produce more evidence than would be available if they did nothing, or behaved randomly. But other studies have demonstrated cases where people take stereotyped, or heuristic actions (Bramley et al., 2015; Coenen et al., 2015), that can be systematically uninformative (Wason, 1968) or fail to reveal particular kinds of rules (Markant & Gureckis, 2014a) or structures (Bramley et al., 2017). Active learners who consider the wrong hypotheses might produce less relevant evidence than would occur naturally (MacKay, 1992). For example, Markant & Gureckis (2014a) had people learn about both one and two-dimensional category rules. When learning actively about a two dimensional category rule, many participants wrongly expected the rule to relate to a single dimension, and so varied their tests only on that dimension. This resulted in their gathering less diagnostic information and performing worse at test than those who were exposed to a random selection of tests.

Yoking participants to the actions of another provides an additional window on active learning, separating the information that is in principle available, from the process of coming up with tests and updating beliefs. Active learners often outperform their yoked

Download English Version:

https://daneshyari.com/en/article/7272551

Download Persian Version:

https://daneshyari.com/article/7272551

Daneshyari.com