Natural science: Active learning in dynamic physical microworlds

Neil R. Bramley[1] (neil.bramley@ucl.ac.uk), **Tobias Gerstenberg**[2] (tger@mit.edu)
Joshua B. Tenenbaum[2] (jtb@mit.edu)
[1]Experimental Psychology, UCL, London, WC1H 0DS, UK
[2]Brain and Cognitive Sciences, MIT, Cambridge, MA 02139, USA

Abstract

In this paper, we bring together research on active learning and intuitive physics to explore how people learn about "microworlds" with continuous spatiotemporal dynamics. Participants interacted with objects in simple two-dimensional worlds governed by a physics simulator, with the goal of identifying latent physical properties such as mass, and forces of attraction or repulsion. We find an advantage for active learners over passive and yoked controls. Active participants spontaneously performed several kinds of "natural experiments" which reveal the objects' properties with varying success. While yoked participants' judgments were affected by the quality of the active participant they observed, they did not share the learning advantage, performing no better than passive controls overall. We discuss possible explanations for the divergence between active and yoked learners, and outline further steps to categorize and explore active learning in the wild.

Keywords: active learning; intuitive physics; causality; probabilistic inference; mental simulation.

The great majority of research on human and machine learning has focused on *passive* situations, where evidence is fixed or preselected. Participants are typically invited to make judgments based on carefully pre-chosen evidence; and machine learning algorithms compete for predictive accuracy on pre-existing datasets. In contrast, Nature's successful learners are necessarily embedded in the world they must learn about and exploit. Thus, it is the norm for human learners to exert some degree of *active* control over the evidence they see. To understand human learning then, one must also understand the myriad decisions about where to attend, and what action to take, that control and manage the flow of incoming evidence. An effective active learner will be able to bootstrap their learning and improve the utility of the information they receive by tailoring it to resolving their subjective uncertainty. On this view, we can think of the little actions in everyday life as small experiments, ranging from the automatic (e.g. cocking one's head to better locate the origin of a sound), to the deliberate (lifting a suitcase to judge its weight; shaking a present to try and guess its contents; holding a pool cue to one eye, or spinning it, to gauge its straightness). A common element in these examples is that they create situations that bring into sharper relief the physical properties of interest.

In this paper we explore this naturalistic type of learning by looking at how people learn about physical laws and properties, such as magnetism and object mass. The structure of the paper is as follows. We first survey the literatures on active learning and intuitive physics, then describe experiments that contrast passive learners with active and yoked learners. Finally, we look closely at the types of actions that active participants performed to reveal the microworlds' hidden physical properties.

Active learning

Human active learning has largely been studied in simple situations where the space of possible actions is limited and the hypothesis space is well-defined. Examples include category rule learning (Gureckis & Markant, 2009) and games like "Guess Who" (Nelson, Divjak, Gudmundsdottir, Martignon, & Meder, 2014) and "Battleships" (Markant & Gureckis, 2010). A related line of research has explored active causal learning, where participants can intervene on causal systems (Bramley, Lagnado, & Speekenbrink, 2015; Coenen, Rehder, & Gureckis, 2015; Lagnado & Sloman, 2004). Since many causal structures cannot be distinguished by co-variational data alone (Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003), the concept of intervention captures a key aspect of real world active learning that goes beyond simply asking the right questions. The learner's actions can effectively create idealized situations that would rarely happen under normal circumstances, and thus uncover the true causal relationships. However, the "causal systems" explored in these studies are invariably causal Bayesian networks (Pearl, 2000) where time and space are abstracted away, and actions are limited to idealized interventions.

In general, these studies found that people select actions that are more informative than random selection, but that also tend to be more stereotyped and repetitive than those prescribed by models of optimal active selection. Bramley et al. (2015) propose that learners tailor their actions to their own limited learning capacities, testing only a subset of the possible hypotheses at any given time. If bounded learners fail to consider the true hypothesis, they can fail to generate the necessary evidence to support it, and thus perform worse than passive learners (Markant & Gureckis, 2010). This is a common problem for active learning algorithms that do not consider the whole hypothesis space (MacKay, 1992). People are typically found to be bad at balancing the costs of active learning against its benefits, typically oversampling, e.g. taking too many actions when they carry cost (Markant & Gureckis, 2012). In real-time active learning, this might lead participants to allocate too much of their limited attentional resources to controlling rather than learning (Sweller, 1994).

If learners' actions are tailored to their idiosyncratic learning trajectories, the evidence they generate may be less useful for other learners, who are considering different hypotheses while observing the active learners' choices (Markant & Gureckis, 2014). This view is broadly (Lagnado & Sloman, 2004; Sobel & Kushnir, 2006), but not always (McCormack, Bramley, Frosch, Patrick, & Lagnado, 2016), supported by experiments that include yoked conditions, where one participant observes the tests performed by another. Intuitively, the divergence between information that is in principle available, and what participants can actually learn will be much larger in more complex and naturalistic situations, where only a fraction of the total evidence can plausibly be attended to.

Intuitive physics

Early research into intuitive physics focused on documenting how people's understanding of some aspects of physics, such as ballistic and curvilinear motion, is sometimes systematically biased (e.g. McCloskey, Caramazza, & Green, 1980). More recent research has demonstrated how some of these biases may be explained if we assume that 1) our physical understanding is only approximately Newtonian, and 2) we are often fundamentally uncertain about some important aspects of the physical scene (e.g., the masses of the objects involved in a collision, Sanborn, Mansinghka, & Griffiths, 2013).

Battaglia, Hamrick, and Tenenbaum (2013) have argued that people's understanding of physics is best understood in analogy to a physics engine used to model physically realistic scenes. Accordingly, people have a physics simulator in their mind that they can use to approximately predict what will happen in the future (Smith & Vul, 2013), reason about what happened in the past (Smith & Vul, 2014), or simulate what would have happened if some aspect of the situation had been different (Gerstenberg, Goodman, Lagnado, & Tenenbaum, 2015). The results of these experiments are consistent with the view that people have a rich intuitive theory of physics that supports approximately accurate mental simulations of key aspects of physical scenes. However, these experiments do not address the question of how we get there – how do people acquire their intuitive physical theories?

Intuitive theories can be expressed as probabilistic programs (Gerstenberg & Tenenbaum, to appear). Program induction is a thorny problem, but one where human-like performance has been demonstrated by sophisticated Bayesian machinery embodying principles of causality and compositionality (Lake, Salakhutdinov, & Tenenbaum, 2015). Ullman, Stuhlmüller, Goodman, and Tenenbaum (2014) explored intuitive physics learning by looking at how people learn about different latent physical properties of 2D "microworlds" similar to the ones shown in Figure 1. The worlds were bounded by solid walls and contained a number of colored pucks with differing weights, surfaces with differing levels of friction, as well as local (magnet-like) forces between pucks and a global (gravity-like) force pulling all the pucks in one direction. In each clip, the pucks bounced around, attracting and repelling each other, being slowed down by the friction, and being pulled by the global force. Participants identified the correct global force around 70% of the time and were much better at detecting local attraction (82%) than repulsion (53%). Ullman et al. argued that repulsion is more difficult to identify because pucks that repelled one another were rarely close enough to exhibit strong repulsion, while attracting pucks rapidly approached one another and stuck together, thus offering stronger evidence of the latent force.

Ullman et al. modeled participants' judgments by assuming a mixture of an Ideal Observer Model (IOM) and a Simulation Based Approximation Model (SBAM). The IOM compares the observed objects' trajectories to simulations of expected trajectories under the different possible worlds. The SBAM compared statistics about each clip such as the pucks' average positions, velocities and pairwise distances, to the summary statistics of repeated simulations under the different possible worlds. For instance, objects in worlds with a global force towards south generally tend to be closer to the southern wall of the world. A mixture model that combined both IOM and SBAM explained participants' judgments well.

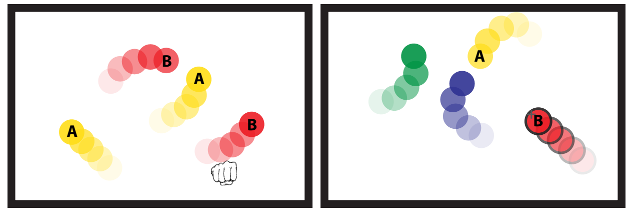

(a) "Punch" condition in pilot study. (b) "Grab" condition in experiment.

Figure 1: Schematic displays of two "microworlds". Note: In the pilot there were two pucks of each "type" i.e. two yellow "A"s and two red "B"s while in the experiment there were two target pucks and two distractor pucks and all were randomly colored.

In the current work we build on these results, exploring how people interact with physical microworlds and how this impacts on their learning of the different physical properties.

Pilot study: From Passive to Active

For our pilot study we adapted the setup from Ullman et al. (2014). However, rather than showing participants prechosen replays, we generated the simulations on the fly to allow for active conditions in which participants could exert control over the scene. We chose two setups that differed in the extent to which participants had fine-grained control over the scene. In the *"active punch"* condition, participants controlled a fist that allowed them to roughly knock other objects around, mimicking the clumsy actions of a baby yet to develop fine motor skills. In the *"active grab"* condition, we allowed learners to use the mouse to grab the pucks with the mouse and drag them around, staging more precisely orchestrated interventions.

We were interested in whether *active* participants would be able to use these forms of control to better identify the forces than the *passive* participants; or conversely if the costs of controlling while learning would lead to worse performance. We expected the active learning advantages to be greater in the fine-grained *grab* condition, and the costs to be higher for the *punch* condition where effective control was harder.

Methods

Participants Sixty participants were recruited through Amazon Mechanical Turk (34 male, age 33.5 ± 9.7). They were paid at a rate of $6 per hour.

Materials The experiment was programmed in Javascript using a port of the Box2D physics game engine. The microworlds were displayed in a 600×400 pixel frame, with 1 m in the world corresponding to 100 pixels on the screen.

Each world was bounded by solid walls with high elasticity (90% of energy retained per collision) – and contained four pucks (2 yellow, 2 red, each with radius .25 meters, mass 1 kg and elasticity 75%). Each world either had a global force of $1\,\mathrm{m/s^2}$ in one of the four compass directions, or no global force. Each world also had up to three distinct local forces, one between the yellow pucks, one between the red pucks, and one between pucks of differing colors. Each of these could either be attractive ($3\,\mathrm{m/s^2}$), repulsive ($-3\,\mathrm{m/s^2}$), or no force.[1] The pucks' initial positions were random but non-overlapping, with initial velocities in the x and y direction drawn from $\mathrm{Unif}(-10,10)\,\mathrm{m/s}$. Whenever all pucks' velocities fell below $.15\,\mathrm{m/s}$, the simulation froze and the window went black for 500 ms before the positions and velocities of the pucks were redrawn. Each world was simulated for 30 seconds at 60 frames per second.

Conditions Participants were randomly assigned to one of three learning conditions, *passive* ($N = 21$), *active punch* ($N = 20$), *active grab* ($N = 19$, see Figure 1):

1. **Passive** Participants observed the microworlds unfold without being able to interact.
2. **"Active punch"** In addition to the four pucks, this condition featured a "fist" (see Figure 1a). The fist was the same size as the pucks but was heavier (10 kg) and less elastic (50%). The fist was initially located in the middle of the screen but strongly attracted to the location of the participant's mouse.[2]
3. **"Active grab"** In this condition, participants could grab pucks and drag them around with the mouse. Grabbed pucks retained their properties (i.e. mass and local forces and location and momentum) but became strongly attracted to the position of the mouse. When released they would continue on their current trajectory but no longer be attracted to the mouse.

Worlds Each participant either passively observed or actively interacted with 18 microworlds. The set of worlds comprised all combinations of the six possible within-color local force patterns [None-None, Attract-Attract, Repel-Repel, None-Attract, None-Repel, Attract-Repel] and the three possible between-color local forces [None, Attract, Repel]. Half of the microworlds also had a global force in one of the four compass directions. Object colors and direction of the global force were counterbalanced.

Procedure Participants were instructed about the setup of the microworlds, what judgments they had to make, and – if they were in an active condition – how they could interact with the pucks. Participants first saw two practice trials, and then 18 test trials in randomized order.

On each trial, participants answered 4 questions: One question about the direction of the global force (response options: "North", "East","South","West","None" and "Don't know"), and one question each about the local forces between pucks of the same color (red and yellow), and pucks of different color (response options: "Attract", "None","Repel" and "Don't know"). Participants took on average 22.4 ± 9.3 minutes to complete the experiment.[3]

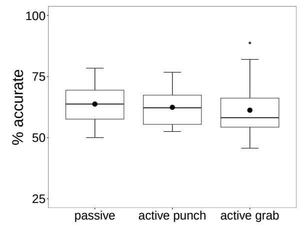

(a) Boxplots of accuracy by condition.

Ground truth \ Judgment	Attract	None	Repel	Dunno
Attract	**86%**	4%	7%	3%
None	**40%**	28%	25%	7%
Repel	11%	16%	**66%**	7%

(b) Confusion matrix for local force judgments.

Figure 2: Pilot study performance plots.

Results

Participants in the *passive*, *active punch*, and *active grab* condition answered on average 64%, 62% and 61% of the questions correctly. Chance performance was approximately 30%.[4] Thus, judgments were well above chance in all three conditions. However, there was no main effect of condition on performance $F(2,57) = .38, p = .67$. As Figure 2a shows, both the highest and the lowest performing participants were in the *active grab* condition, suggesting that an active learning advantage for this scenario was at least possible although not generally achieved.

On the global force question people were worse at identifying when there was no force, with accuracy of only 42% when the right answer was "none" compared to an average of 85% when the right answer was one of the compass directions. The accuracy difference for identifying "none" vs. one of the other forces interacts with condition $F(4,294) = 2.6, p = .03$, with only 40% in the *active punch* and 30% in *active grab* condition identifying when there was no global force compared to 57% in the *passive* condition. For the local force questions, accuracy differed considerably depending on the ground truth (Figure 2b). Participants in the *active grab* condition were better than *passive* and *active punch* participants at identifying repel forces with an overall accuracy of 77% compared to 68% and 70%.

Due to the simulation restarting whenever all the pucks fell below a certain velocity ($2.0 \pm .8$ times per trial on average), participants in the *passive* condition actually experienced significantly more puck motion than the active participants. We can see this in terms of the total distance traveled by the four

[1] Local forces scaled with the inverse squared distance between the objects in line with Newton's universal law of gravitation. Thus the current local force L exerted on object o_1 by object o_2 (and the reverse) was given by $\frac{3}{d^2}$.

[2] We opted for strong attraction rather than simply copying the position of the mouse because this allowed the controlled object to interact reciprocally with the other objects in collisions rather than behaving as if it was infinitely heavier than the other objects.

[3] Complete specification of the settings of the Box2D simulator and demos of both experiments are available at `ucl.ac.uk/lagnado-lab/el/apl`

[4] Any "don't know" responses were treated as judgments spread evenly across the remaining 3 or 5 options. Random responding would be correct with probability $\approx \frac{1}{4} \times \frac{1}{5} + \frac{3}{4} \times \frac{1}{3} = .3$

colored pucks over the trials of 168 ± 47, 98 ± 55, 85 ± 73 meters for *passive*, *active punch* and *active grab* conditions.

For the participants in the *active grab* condition, more time spent manipulating the pucks was positively related to performance $F(1,17) = 7.2, p = .015$.

Interim Discussion

While participants' overall accuracy was not affected by learning condition, performance in the *active grab* condition was more variable, depending on how much learners controlled the pucks. *Active* learners' lower accuracy on the global force question indicates that controlling individual pucks may have led them to neglect the global properties of the scene (i.e. that the pucks congregated on one side of the world). *Active* learners were more accurate at detecting repulsion, perhaps because they were able to force repellent pucks closer together and so experience stronger evidence for repulsion.

There are several possible explanations for the lack of condition differences in overall accuracy. Restarts of the simulation in the passive condition meant that passive participants naturally experienced more balanced clips with more time during which pucks moved and interacted. Active participants had to put in work to create these experiences (evidenced by the higher variance but lower average puck motion) that were seen "for free" in the passive condition. Another possibility is that the setup of the pilot was not well-suited for active exploration. 30 seconds may have been too little time to allow for sequential, controlled testing, especially of four distinct physical properties. Third, end-of-task feedback suggested that errors were often not due to difficulties in detecting the forces but rather because of having to hold the answers in working memory until the end of the clip; or failure to segment the different worlds in memory, mixing up properties experienced in the current versus previously experienced worlds.

Main Experiment

For our main task we used the same setup as in the pilot but made a number of changes to address the issues identified above. Firstly, we improved the match between passive and active conditions by tweaking the settings of the microworlds so that objects rarely came to rest within the the length of a trial. We increased the elasticity of the pucks from 75% to 98%, leading to restarts occurring only in exceptional situations. Additionally, we replaced the *active punch* condition with a *yoked* condition (cf. Lagnado & Sloman, 2004), in which participants were matched with one of the *active grab* participants and observed their mouse movements and controlling actions. To increase the scope for active hypothesis testing, we increased the length of the trials and asked more difficult test questions (see below).

Because active testing is particularly valuable when competing causal explanations cannot be resolved by observational evidence only, we generated confounded evidence by including two distractor pucks along with two target pucks and drew local forces randomly out of attract/none/repel for all pairs of target and distractor objects. This means that it was more important to isolate the target pucks from the distractor pucks to get clear information about the target pairwise force. Instead of including a global force, which was easily identified by passive learners, we varied the relative mass of the two target objects, a property which is more difficult to infer without experiencing curated comparisons and interactions between them. To reduce memory load, we asked two rather than four questions per trial. To make it clear that each world contained new objects we drew random colors for each object and used new labels, cycling through the alphabet, for the target objects. To ensure that participants were motivated to answer the questions as well as they could, we paid a bonus for each correct response. Finally, to get a more fine-grained measure of participants' judgments, we added confidence sliders for each test question and removed the "don't know" option.

Table 1: Experiment design. *Note*: A = attract, N = none, R = repel; masses are in kg.

World	1	2	3	4	5	6	7	8	9
Target force	A	A	A	N	N	N	R	R	R
Target 1 mass	1	2	1	1	2	1	1	2	1
Target 2 mass	1	1	2	1	1	2	1	1	2

We hypothesized that in these worlds *active* participants would outperform *passive* participants, and that *yoked* participants would inherit some, but not all of this advantage.

Methods

Participants Sixty-four participants were recruited from Amazon Mechanical Turk (39 male, age 33.6 ± 10.2). Participants were paid at a rate of \$6 per hour, plus performance-related bonuses (\$0.61 $\pm$ .17) .

Design The first 44 participants were randomly assigned to either the passive ($N = 24$) or the active ($N = 20$) learning condition, and the final 20 participants were yoked 1-to-1 with the active participants. Each participant watched or interacted with 9 microworlds, consisting of all combinations of target force in *attract*, *repel* and *none* and target masses in $[1,1]$kg, $[1,2]$kg and $[2,1]$kg (see Table 1). The five other pairwise forces were drawn uniformly from the three possibilities for each participant on each trial. There were no global forces.

Materials and Procedure We used the same basic set up as in the pilot, but ran the simulations for 45 rather than 30 seconds and increased the elasticity of the pucks from .75 to .98. Rather than two yellow and two red pucks, we drew four random colors for each new world. The two target pucks were labeled with new letters on each trial (e.g. "A" and "B" on trial one, "C" and "D" on trial two, cf. Figure 1b). The distractor pucks were all 1 kg as before but now one of the target pucks could weigh 2 kg. For yoked participants, the cursor of the participant to whom they were yoked (hereafter the *yoker*) was shown with a large "+" symbol whenever it was within the world, and any objects grabbed by the yoker were indicated as in the active condition with a thick black border.

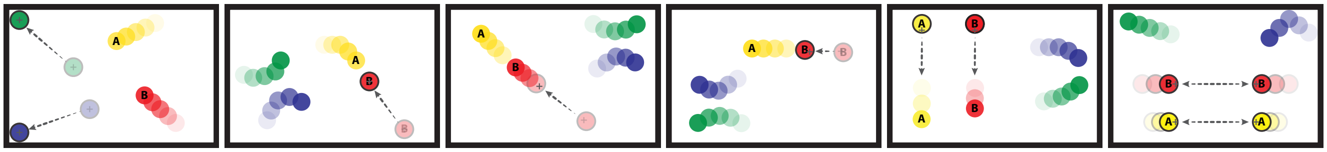

(a) Deconfounding (b) Encroaching (c) Launching (d) Knocking (e) Throwing (f) Shaking

Figure 3: Examples of different interventions participants performed in the *active grab* condition.

Participants first completed instructions relevant to their condition, answered comprehension check questions, and then faced two practice trials followed by the nine test trials. Practice trials were always worlds 1 and 5 (see Table 1). The randomly drawn distractor forces, puck colors and labels differed between the practice and test instances. The two test questions appeared below the world when the time was up. Question order was counterbalanced between participants. At the end of the experiment, participants received feedback about how many of the test questions they got right, and were paid a 5c bonus for each correct answer. The experiment took 19.0 ± 7.3 minutes on average.

Results

Overall accuracy Participants answered 53%, 66% and 54% of questions correctly in the *passive*, *active* and *yoked* conditions respectively (see Figure 4, note chance was $\approx 33\%$ because there were three response options for both questions). Average performance differed significantly by condition $F(2,61) = 3.8, \eta^2 = .12, p = .03$. Post-hoc tests revealed that *active* participants answered significantly more questions correctly than *passive* participants $t(42) = 2.5, p = 0.02$, and (paired) *yoked* participants $t(19) = 2.9, p = 0.02$, with negligible difference between *passive* and *yoked* participants $t(42) = .2, p = 0.83$. Only 4 yoked participants outperformed their active counterparts, with a further 3 answering the same number of questions correctly. Yoked participants' performance was correlated with their active counterparts' $r = .49, p = .03$.

Masses vs. relationships Across conditions, participants were worse at inferring masses than forces $t(63) = -4.8, p < .0001$ and reported lower confidence in mass judgments $66 \pm 25\%$ compared to force judgments $74 \pm 25\%$ $t(63) = -4.2, p < .0001$. Again, participants were less accurate in correctly identifying when there was *no force* between the target pucks (56%) than *repulsion* (70%) or *attraction* (78%), with a main effect of question type $F(2,189) = 7.7, p < .0001$ and significant improvements going from *no force* to *attraction* $t(126) = 3.9, p = .0001$ and *repulsion* $t(126) = 2.4, p = .017$. Force type additionally interacted with condition $F(6,183) = 3.0, p < .0001$. Dummy contrasts with *no force* and *passive* as controls revealed *active* participants were significantly better at identifying *repel* than *passive* participants $t(42) = 3.2, p < .0001$ and a marginal improvement for *yoked* participants as well $t(42) = 1.9, p < .058$. There was no significant relationship between accuracy on the local force question and the number of distractor forces.

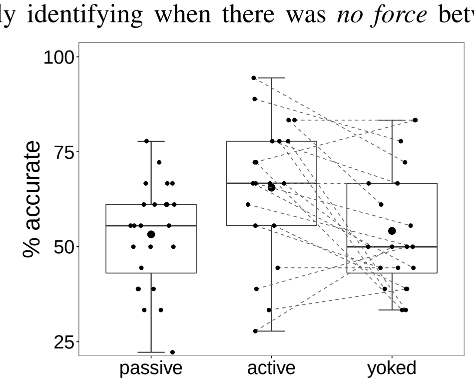

Figure 4: Performance by condition in the main task. *Note*: Large dots indicate condition means. Small dots indicating individual participants in are jittered for visibility. Dotted lines connect active participants with matched yoked participants.

Confidence judgments differed by condition $F(2,61) = 5.3, \eta^2 = .15, p = .007$, with *active* participants significantly more confident on average than *passive* $t(42) = 2.8 p = .006$ or *yoked* participants $t(38) = 2.9, p = .006$. Confidence was positively correlated with accuracy $F(1,62) = 10.6, \eta^2 = .15, p = .002$ but did not interact with condition.

Natural experiments Active participants experienced slightly fewer between-puck collisions than passive participants, 59 ± 14 compared to $65 \pm 9, t(42) = 2.0, p = 0.056$. However they experienced significantly more collisions between the two target pucks 15.0 ± 8.1, compared to $9.8 \pm 4.4, t(42) = 2.7, p = 0.01$. 13.2 ± 7.8 of collisions in the active condition took place while one of the two target objects was being controlled by the participant.

Time spent controlling objects was positively related to final performance for active and yoked participants $F(1,38) = 4.8, \eta^2 = 11, p = 0.04$. Therefore, a key question is what kinds of experiments active participants used to find answers to the test questions. Space constraints prohibit a full analysis in the present paper, but we want to share some of the strategies that participants discovered (see Figure 3 and `ucl.ac.uk/lagnado-lab/aplc`):

(a) **Deconfounding** Even though participants mainly manipulated the target pucks, they also sometimes manipulated the distractor pucks. Many of these manipulations involved moving the distractor pucks out of the way and leaving them at rest in a far corner.

(b) **Encroaching** Participants grabbed one target puck and brought it toward the other target puck. This simple strategy allowed participants to infer whether and how the two pucks affected one another. In some cases, participants towed one attracting puck with the other, or pushed a repulsive puck around with the other providing a strong and extended demonstration of the force between the pucks.

(c) **Launching** Participants grabbed one of the target pucks and flicked it against the other target puck. This intervention helps to figure out whether one of the targets is heavier than the other.

(d) **Knocking** Similar to *launching*, participants grabbed one of the target pucks and knocked it against the other (without letting it go). This intervention also reveals information about the mass of each object.

(e) **Throwing** Participants grabbed a target puck and then threw it, explicitly avoiding collision with any of the other pucks. By exerting an identical force when throwing each target puck, the results of the intervention help to figure out the mass of each object.

(f) **Shaking** Some participants discovered an effective strategy for comparing the mass of the two target objects. By rapidly shaking each in turn (moving the mouse from side to side) it was possible to see that the heavier object reacted more slugglishly. Its greater momentum takes longer to be counteracted by its attraction to the mouse location.

In line with *encroaching* (Figure 3b), we see evidence that participants in the active condition identified the local forces by bringing the two target pucks close to each other. The lower the average distance between two target objects for an *active* participant, the better they did on the force question $\beta = -.3, F(1,18) = 8.0, \eta^2 = .3, p = .001$ but this had no relationship with accuracy on the mass question $p = .87$. Conversely, in line with the *shaking* strategy (Figure 3f), participants who moved the controlled object around faster did better on the mass question $\beta = 25, F(1,18) = 15, \eta^2 = .45, p < 0.001$, but controlled object speed had no relationship with accuracy on the force question $p = .67$. Yoked participants did not inherit these differences, with no significant relationships between performance on either question and average distance between targets or controlled-object speed.

Discussion

We found a clear benefit for active over passive learning in this experiment. In particular, active participants gathered more evidence about repulsive forces by bringing target objects closer together. The quality of the control exerted by the active participants was an important determinant of the quality of the final evidence available to the yoked participants. However, the substantial drop-off from active to yoked accuracy was consistent with the idea that first-hand knowledge of *what* was being tested (e.g. relationship or mass), *when* and *how*, was likely to be crucial for learning successfully. Since there are too many objects and properties in play to track at once, it helps to align the evidence with the hypotheses currently considered. Another factor might have been that active participants were able to look ahead at the crucial locations in the scenes where diagnostic interactions were expected to occur. Yoked participants lacked the ability to foresee what will happen. Finally, active participants had an additional advantage over yoked participants by receiving direct motor feedback about their interventions. They experienced how quickly they moved the mouse or their finger on the trackpad and thus had an immediate sense for how much force they exerted.

Encroaching and *shaking* permitted simple indirect measures, and accordingly, we found shakers doing better on mass questions and encroachers doing better on relationship questions.

While the current study provides a valuable first step, there is much more to explore here. In future work we plan to extend the IOM and SBAM models to active data and use them to evaluate the informativeness of different strategies. We also plan to explore the possibility that learners have a generative grammar for constructing natural experiments; and to unpack yoking differences by looking at yoked participants' ability to infer the learning intentions and action plans of active learners.

Acknowledgments Thanks to Hongyi Zhang for initial code, and Tomer Ullman and David Lagnado for helpful comments. TG and JT were 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.

References

Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, *110*(45), 18327–18332.

Bramley, N. R., Lagnado, D. A., & Speekenbrink, M. (2015). Forgetful conservative scholars - how people learn causal structure through interventions. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *41*(3), 708-731.

Coenen, A., Rehder, B., & Gureckis, T. M. (2015). Strategies to intervene on causal systems are adaptively selected. *Cognitive Psychology*, *79*, 102–133.

Gerstenberg, T., Goodman, N. D., Lagnado, D. A., & Tenenbaum, J. B. (2015). How, whether, why: Causal judgments as counterfactual contrasts. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (pp. 782–787).

Gerstenberg, T., & Tenenbaum, J. B. (to appear). Intuitive theories. In M. Waldman (Ed.), *Oxford handbook of causal reasoning.* Oxford University Press.

Gureckis, T. M., & Markant, D. (2009). Active learning strategies in a spatial concept learning game. In *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 3145–3150).

Lagnado, D. A., & Sloman, S. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *30*, 856–876.

Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, *350*(6266), 1332–1338.

MacKay, D. J. (1992). Information-based objective functions for active data selection. *Neural computation*, *4*(4), 590–604.

Markant, D. B., & Gureckis, T. M. (2010). Category learning through active sampling. In *Proceedings of the of the 32nd Annual Meeting of the Cognitive Science Society* (pp. 248–253).

Markant, D. B., & Gureckis, T. M. (2012). Does the utility of information influence sampling behavior? In *Proceedings of the 34th annual conference of the cognitive science society* (pp. 719 – 724).

Markant, D. B., & Gureckis, T. M. (2014). Is it better to select or to receive? learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General*, *143*(1), 94.

McCloskey, M., Caramazza, A., & Green, B. (1980). Curvilinear motion in the absence of external forces: Naïve beliefs about the motion of objects. *Science*, *210*(4474), 1138–1141.

McCormack, T., Bramley, N. R., Frosch, C., Patrick, F., & Lagnado, D. A. (2016). Children's use of interventions to learn causal structure. *Journal of Experimental Child Psychology*, *141*, 1-22.

Nelson, J. D., Divjak, B., Gudmundsdottir, G., Martignon, L. F., & Meder, B. (2014). Children?s sequential information search is sensitive to environmental probabilities. *Cognition*, *130*(1), 74–80.

Pearl, J. (2000). *Causality*. New York: Cambridge University Press (2nd edition).

Sanborn, A. N., Mansinghka, V. K., & Griffiths, T. L. (2013). Reconciling intuitive physics and newtonian mechanics for colliding objects. *Psychological Review*, *120*(2), 411.

Smith, K. A., & Vul, E. (2013). Sources of uncertainty in intuitive physics. *Topics in Cognitive Science*, *5*(1), 185–199.

Smith, K. A., & Vul, E. (2014). Looking forwards and backwards: Similarities and differences in prediction and retrodiction. In *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 1467–1472).

Sobel, D. M., & Kushnir, T. (2006). The importance of decision making in causal learning from interventions. *Memory & Cognition*, *34*(2), 411–419.

Steyvers, M., Tenenbaum, J. B., Wagenmakers, E., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*, 453-489.

Sweller, J. (1994). Cognitive load theory, learning difficulty, and instructional design. *Learning and instruction*, *4*(4), 295–312.

Ullman, T., Stuhlmüller, A., Goodman, N., & Tenenbaum, J. (2014). Learning physics from dynamical scenes. In *Proceedings of the 36th Annual Conference of the Cognitive Science society* (pp. 1640–1645).