

Teaching and Learning Through Pedagogical Environment Design

Emily G. Liquin (emily.liquin@nyu.edu)

Nicole Luzuriaga (nl2210@nyu.edu)

Todd M. Gureckis (todd.gureckis@nyu.edu)

Department of Psychology, New York University
New York, NY 10003 USA

Abstract

People often rely on knowledgeable teachers to help them learn. Sometimes, this teaching is direct: teachers provide instructions, examples, demonstrations, or feedback. But other times, teaching is more subtle: teachers construct the physical environment in which a learner explores. In the present research, we investigate this more subtle form of teaching in an artificial grid-based learning environment. How do people construct the physical environment to teach, and how does the (pedagogical) design of the physical environment affect people's learning? Study 1 shows that people pursue multiple approaches to pedagogical environment design. Study 2 shows that learners make systematic, often accurate inferences from pedagogically designed environments, even in the absence of exploration. Together, these studies add to our understanding of the myriad ways in which experts communicate their knowledge to novices—a capacity that is part of what makes human intelligence unique.

Keywords: pedagogy; learning; teaching; environment design; guided play

Introduction

Humans can learn a great deal through their own exploration and observation, but we often rely on knowledgeable teachers to help us. This is because learning through exploration and observation alone, without pedagogical guidance, is often an insurmountable challenge. For example, one could place a large number of objects into a large number of fluids without ever successfully discovering the hidden forces (i.e., gravity, buoyancy) that make objects float.

Prior work has investigated how teachers teach and how learners learn using instructions (Popp & Gureckis, 2020), examples (Shafto et al., 2014), demonstrations (Ho et al., 2021), and feedback (Ho et al., 2017), as well as how learners evaluate teachers' pedagogical actions (Bass et al., 2022). In all these cases, teaching is direct: the teacher provides information that will help the learner infer the desired concept.

However, there is another subtle way teachers can shape their students' learning: by constructing the learning environment in which the student explores. For example, a museum exhibit might provide a learner with particular objects and fluids—grouped in ways that would be especially informative for learning about buoyancy, but giving the learner freedom to take any actions they choose. This sort of “pedagogical environment design” is also used by video game designers: players encounter objects in specific arrangements (e.g., a floating platform carefully placed above a scary-looking opponent).

Often, this occurs without explicit instruction but is designed to encourage particular actions (jump on the platform) and inferences (the opponent is dangerous; the platform is safe).

The goal of the present research is to understand this subtle form of teaching: how do people construct learning environments with the goal of teaching, and how is learning affected by the design of the environment? Answering these questions will shed light on what role knowledgeable teachers can play in facilitating learners' successful exploration and learning.

Prior research

Prior research has suggested that pedagogical environment design may be an effective way of teaching. For example, guided play—which involves adult construction of the learning environment, followed by child-directed play with adult guidance—is beneficial for children's learning (Fisher et al., 2013; Weisberg et al., 2013, 2016). In addition, expert-designed “playful learning landscapes” can provide structured opportunities for children to play and learn in real-world settings (Bustamante et al., 2019).

However, it is less clear *how* the environment's design influences learning. One possibility is that an environment's design makes certain actions more salient than others (see Weisberg et al., 2014), and these salient actions provide important learning opportunities. For example, in a video game, placing an opponent directly in the player's path gives the player no choice but to interact with it (and thus learn how to defeat it). Consistent with this possibility, research on nudges has shown that the design of the physical environment can affect people's choices in non-learning contexts (Mertens et al., 2022; Thaler & Sunstein, 2009). In addition, a physical object's design might highlight particular affordances (Gibson, 1977, 1979), and thus prompt certain ways of interacting with the object (Norman, 1988). Similarly, teachers might use the environment's design to guide learners towards the actions that are best for learning. If this is the case, the environment plays an *indirect* role in learning—people will learn from the actions that they take, and the environment's role is to guide people towards certain actions over others.

However, the design of the environment might also have a *direct* influence on learning, by pointing to certain inferences over others. For example, Walker et al. (2020) found that when a novel toy's design contained slots for two objects (versus no slots), children and adults more readily inferred a rela-

tional causal rule governing the toy's behavior (e.g., that pairs of the same object must be placed on the toy to produce an effect). How could an environment directly point to particular inferences? For other forms of teaching, a learner's inference from a teacher's example or demonstration is based on the assumption that the teacher intentionally selected an action that would communicate the correct inference to the learner (Butler & Markman, 2012; Bonawitz et al., 2011; Shafto et al., 2014). Thus, we might expect a learner's inferences from a physical environment to be made on the basis of a similar communicative assumption: that the teacher specifically designed the environment to communicate the correct inference. This assumption enables the environment to serve as evidence in support of particular inferences (e.g., Why would the toy designer put two slots on the toy, if not to communicate that it needs two objects to activate?).

The present research

In the present research, we investigate (1) how teachers design pedagogical environments, and (2) how learners make inferences from pedagogically designed environments. In both cases, we look for behavioral signatures of the two distinct ideas outlined above: that pedagogically designed environments support certain actions over others, thus indirectly affecting learning (which we call an "**action-based approach**"), and that pedagogically designed environments serve as evidence supporting certain inferences over others due to communicative assumptions, thus directly affecting learning (which we call a "**communicative approach**"). We predict that environment design has both direct and indirect effects on learning. Thus, we do not view these as competing accounts, but rather as complementary approaches.

To investigate these approaches, we adapt an artificial grid-based learning task from Ho et al. (2021), in which the participant navigates through a colored grid to reach a goal (see Fig. 1). Some colors in the grid are dangerous (cost points) and other colors are safe (have no effect on points). This task is especially amenable to studying teaching and learning through pedagogical environment design because the learning goal (i.e., which colors are safe and which are dangerous) is easy to describe to participants, and the space of possible environments (i.e., how colors are arranged in the grid) is finite but nonetheless quite large.

In Study 1, participants are tasked with generating a grid-based environment, with the goal of teaching another (future) participant whether each color is safe or dangerous. We investigate the extent to which participants' designed environments are structured in a way that encourages particular informative actions (action-based approach) and communicates particular inferences (communicative approach).

In Study 2, participants are shown several of the designed environments from Study 1, and they are tasked with inferring the reward structure of each environment. We test whether participants make accurate inferences from pedagogically designed environments, even in the absence of direct exploration. In addition, we test whether these inferences are

driven by communicative assumptions (in line with the communicative approach). Data and analyses for both studies can be found here: <https://osf.io/7s2ey/>

Study 1: Teaching

How do people design environments with the goal of teaching? In the context of our learning task (see Fig. 1), a purely action-based approach would encourage the learner to interact with each color, thus ensuring the learner will learn about each color through their own actions. In contrast, a purely communicative approach would attempt to communicate through the environment's design whether each color is dangerous or safe. For example, an obvious path to the goal that passes through only one color might communicate that the color along that path is safe (why would the path be so obvious if the color were dangerous?). Study 1 tests whether teachers' designed environments align with these approaches.

Methods

Participants We recruited 91 participants from Prolific, who were paid \$2.50 for participating in the 10-minute study. An additional 9 participants were excluded because there were issues saving their data, they failed to pass two attention checks (see Procedure below), or they reported that they were colorblind. Participants were randomly assigned to design an environment so that it was pedagogical ($N = 33$), easy ($N = 29$), or difficult ($N = 29$). Due to limited space, we only report the procedure and results for the 33 participants in the pedagogical condition here.

Participants ranged in age from 19 to 71 ($M = 37$) and included 13 women, 18 men, and 2 non-binary/genderqueer individuals. Participants were 76% White, 9% Black, 3% Asian, 3% Middle Eastern/North African, 3% Pacific Islander/Native Hawaiian, and 9% Hispanic/Latinx of any race. All participants were located within the United States.

Procedure Participants were introduced to a novel game (see Fig. 1), in which they used the arrow keys to control a blue circle on a 6x5 grid. Their goal was to traverse the grid to reach a yellow square, worth 10 points. To reach the goal, participants needed to cross a field of 20 colored squares. Each time the blue circle hit a "dangerous" color, the participant lost two points. Other colors were "safe" and had no effect on points. Participants were familiarized with this game through verbal instructions, example videos, a comprehension test, and four practice games with unique colored grids.

Next, we told participants that they would be designing a new round of the game for another person, who wouldn't know which colors were safe and which colors were dangerous. The participant was told which colors were which: colors were randomly assigned to be either safe or dangerous, with at least one color of each type. 16 participants had one safe color, and 17 participants had two safe colors. Participants were instructed: "You should design your round so it helps the other participant learn as much as possible about which colors are safe and which colors are dangerous." Then,

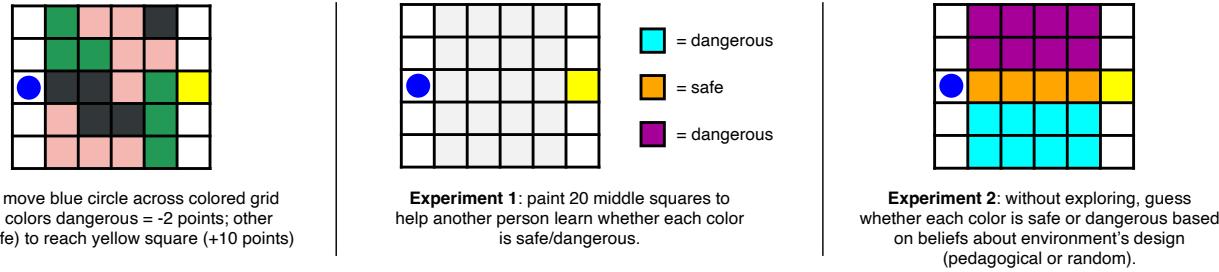


Figure 1: Overview of novel game and tasks for Studies 1 and 2.

participants painted the three colors onto an empty grid, with two requirements: each color had to be used at least once, and all 20 middle tiles had to be colored (see Fig. 1).

After designing their round, participants rated how helpful for learning their round was (1 = Not at all, 7 = Very much). Participants were also prompted to describe in text what they thought made their round helpful for learning.

Finally, participants provided demographic information and reported on their experience completing the study. During these tasks, they completed two attention checks: selecting the middle option on a seven-point scale and identifying which instructions they had read from three options.

Results

We investigate whether participants designed environments in a manner consistent with the two proposed approaches.

To do so, we first coded participants' free-response justifications for why their round was helpful for learning. 42% of participants mentioned trying to force a learner to cross all three of the colors (e.g., "I made sure that all colors must be encountered at least once."), consistent with an action-based approach. In contrast, only 9% of participants mentioned trying to highlight an obvious safe path (e.g., "There is a clear path to the yellow target. It is a helpful guide."), consistent with a communicative approach. Due to the smaller proportion of participants who provided communicative justifications, the following analyses of participant-designed grids contrast participants who provided action-based justifications ("action-based grids", $N = 14$) against participants who provided any other justification ("non-action-based grids", $N = 19$). Fig. 2 includes examples of the grids designed by eight participants: four action-based and four non-action-based.

Do people take an action-based approach? An action-based approach would prompt interaction with all three colors, as this would enable a learner to infer the category (safe/dangerous) of each color through their own exploration. Thus, for each designed environment, we investigated what proportion of paths through the 6x5 grid crossed all three colors at least once. The total number of possible paths through a 6x5 grid is quite large—if an agent never returns to an already-visited cell, there are 37,609 unique paths through the grid. However, many of these paths are inefficient or redundant, e.g., crossing the same cells in a different order. Rather than analyzing all paths, we only considered the 271 paths that crossed 10 or fewer unique cells (the 75th percentile of unique cells that participants crossed in their own exploration while learning the game). For each designed grid, we simulated the results of taking these 271 possible paths, focusing on which colors were encountered on each path.

On average, 77% of paths through a given designed grid crossed all three colors ($SD = 23\%$). However, this differed between participants who provided action-based justifications and those who did not. In a linear regression model, justification type was significantly related to the proportion of paths crossing all three colors, $b = -0.24$, 95% CI $[-0.38, -0.10]$, $t(31) = -3.42$, $p = .002$, with an average of 91% ($SD=14\%$) of paths crossing all three colors for action-based grids but only 67% ($SD = 23\%$) of paths crossing all three colors for non-action-based grids. This suggests that those who provided action-based justifications largely executed the action-based strategy successfully: the majority of paths through the designed grids crossed all three colors, thus providing an opportunity to learn about each color through exploration.

Do people take a communicative approach? A communicative approach would highlight safe paths and de-emphasize dangerous paths, thus communicating that the salient paths are safe. First, we test whether such safe paths exist—for each grid, are there any paths that do not include dangerous colors? Again, this differed between participants who provided action-based justifications and those who did not. In a logistic regression model, justification type was significantly related to whether or not a grid provided any safe paths, $OR = 22.50$, 95% CI $[4.15, 191.82]$, $z = 3.28$, $p = .001$. Only 14% of action-based grids had at least one safe path, while 79% of non-action-based grids had a safe path.

Next, we tested whether safe paths were somehow emphasized relative to dangerous paths. This could occur in (at least) two ways: (1) a safe path could be the shortest path to the goal, and/or (2) a safe path could be visually distinct from dangerous paths (e.g., by including only a single color on the safe path, but multiple colors on dangerous paths).

In line with possibility (1), the shortest path contained

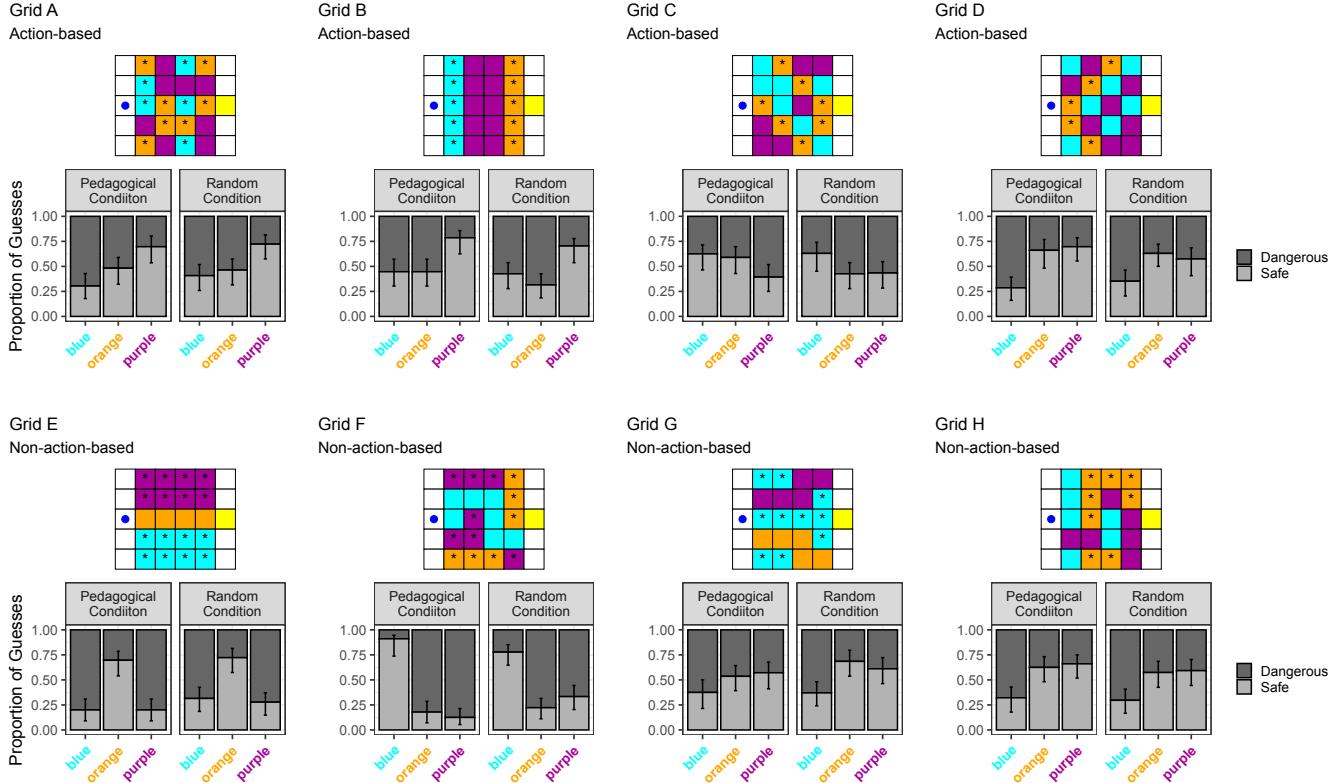


Figure 2: Grids generated by participants in Study 1 who provided action-based justifications (Grids A-D) and any other justifications (Grids E-H). Asterisks (*) indicate dangerous cells. Below each grid, the proportion of Study 2 participants who guessed safe/dangerous for each color within each grid, with bootstrap 95% confidence intervals.

only safe colors in five of the designed grids, all non-action-based. More generally, however, there was no evidence for a difference in the number of safe cells along the shortest path between action-based grids and non-action-based grids, $b = 0.28$, 95% CI $[-0.54, 1.10]$, $t(31) = 0.69$, $p = .49$. In line with possibility (2), there was at least one single-color safe path in 63% of non-action-based grids, compared to 0% of action-based grids (Fisher's exact test, $p < .001$). However, nearly as many non-action-based grids (47%) included a single-color *dangerous* path. A multilevel logistic regression model (with by-participant random intercepts) provided no evidence that single-color safe paths were more likely to be present than single-color dangerous paths within non-action-based grids, $OR = 2.32$, 95% CI $[0.53, 1.93 * 10^7]$, $z = 1.05$, $p = .29$.

Discussion

Together, the results from Study 1 suggest that people take both action-based and communicative approaches to pedagogical environment design. Nearly half of participants took an action-based strategy, designing grids in which most possible paths crossed all three colors. These participants articulated this strategy when justifying the design of their grids. Other participants appeared to take a communicative strat-

egy. However, relatively few participants articulated this as their strategy when justifying the design of their grids, and the evidence that participants highlighted safe paths and de-emphasized dangerous paths was mixed.

Notably, our prompt to participants was (intentionally) vague: we instructed participants to design a grid that would help another person learn, but we did not tell them *how* the person was to learn. The variation in participants' strategies suggests that participants may have made different assumptions about the hypothetical learner: some thought the learner would only learn about each color through interaction with the environment (action-based approach), while other participants thought the learner might take the environment itself as evidence in support of particular inferences (communicative approach). Next, we test whether this latter assumption—that people use the environment as evidence—is justified.

Study 2: Learning

In Study 2, we investigate how people learn from pedagogically designed environments. In particular, we investigate what people infer about the colors in the environment *before* interacting with them, and we test whether these inferences are driven by the assumption that the environment was designed pedagogically by a helpful teacher.

Methods

Participants We recruited 110 participants from Prolific, who were paid \$2.00 for participating in the 8-minute study. An additional 10 participants were excluded according to the same criteria as Study 1. Participants were randomly assigned to one of two instructions conditions: pedagogical (N = 56) or random (N = 54).

Participants ranged in age from 20 to 75 ($M = 39$) and included 55 women, 52 men, and 2 non-binary/genderqueer individuals (one additional participant did not specify gender). Participants were 68% White, 11% Black, 8% Asian, 7% multiracial, and 10% Hispanic/Latinx of any race (one additional participant did not specify race or ethnicity). All participants were located within the United States.

Procedure Participants were introduced to the game using the same verbal instructions, example videos, and comprehension test as Study 1. Then, participants were shown eight grids, corresponding to eight separate rounds of the game. For each grid, participants guessed whether each color was safe or dangerous and indicated their confidence with each guess on a 100-point slider scale, ranging from "not at all confident" (0) to "extremely confident" (100). Notably, participants guessed while looking at the grid, but without actually playing the game. Thus, guesses were made in the absence of any direct exploration (see Fig. 1).

Participants in the pedagogical instructions condition were told that the grids were specially designed by a previous participant, in order to help teach somebody whether each color is safe or dangerous. Participants in the random instructions condition were instead told that the grids were generated by a computer program, which chose a color for each tile using a random algorithm. These prompts were adapted from prior work on pedagogical reasoning (Shafto et al., 2014).

The eight grids were those shown in Fig. 2, sampled from the grids generated in Study 1 (with colors randomly reassigned so that no two consecutive rounds used the same colors). Half of the selected grids were designed by participants who provided action-based justifications, and the remaining half were designed by participants who provided non-action-based justifications.

After providing guesses for the eight grids, participants played the game for two of the grids. These data are not analyzed here. Finally, participants provided demographic information and reported on their experience completing the study. During these tasks, they completed two attention checks: selecting the middle option on a seven-point scale and identifying which instructions they had read from two options (corresponding to the pedagogical/random instructions).

Results

Do learners make accurate inferences? First, we tested whether the pedagogically designed environments from Study 1 enabled participants to make accurate inferences in the present study. We tested whether accuracy was significantly different from chance performance, using a multilevel

logistic regression model with a fixed intercept, along with random intercepts for participant and for color nested within grid. The intercept was significantly different from zero (i.e., odds ratio significantly different from 1), $OR = 1.59$, 95% CI [1.01, 2.50], Wald $z = 2.25$, $p = 0.02$, indicating greater than chance performance.

However, this varied by grid, $\chi^2(7) = 21.77$, $p < .003$. Using Bonferroni corrected p-values for eight comparisons, accuracy was significantly higher than chance for grids A, B, E, and F (see Fig. 2). Notably, this includes both action-based grids and non-action-based grids (and indeed, there was no evidence that accuracy was predicted by grid type, $\chi^2(1) = 1.38$, $p = .24$). Thus, even "action-based" grids may have included some communicative elements, enabling learners to correctly infer whether each color was safe.

Do learners' inferences rely on communicative assumptions? Next, we tested whether participants' inferences rely on communicative assumptions. If so, we would predict systematic inferences in the pedagogical condition—where assumptions about the teachers' communicative intent are licensed—but not in the random condition—where such assumptions are not licensed.

Each participant guessed the category of three colors for eight grids, a total of 24 judgments. We fit a multilevel logistic regression model to category guesses, with judgment as a fixed effect (i.e., a dummy variable for each grid/color except one, which was the reference group), interacting with condition (pedagogical/random). This model also included by-participant random intercepts. There was no evidence for an interaction between condition and judgment, $\chi^2(23) = 28.02$, $p = .22$, and there was no evidence for a main effect of condition when controlling for judgment, $\chi^2(1) = 0.01$, $p = .92$. Thus, there was no evidence that category guesses varied across conditions (see Fig. 2).

Next, we tested whether confidence ratings differed between conditions. We fit a regression model predicting confidence ratings with condition (pedagogical/random) as a fixed

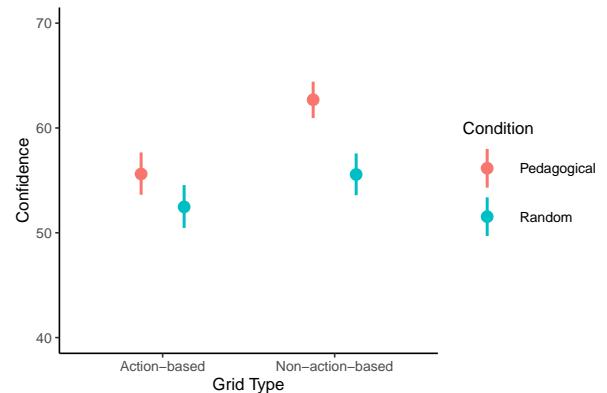


Figure 3: Participants' confidence ratings (averaged across all judgments) in the pedagogical condition and random condition, for grids that were action-based or non-action-based. Error bars indicate bootstrap 95% confidence intervals.

effect and with random intercepts for participant and for color nested within grid. We also controlled for color category guesses (dangerous/safe) with an additional fixed effect. There was no evidence that the model including condition as a fixed effect provided a better fit to the data than a model excluding it, $b = -5.15$, 95% CI $[-12.94, 2.65]$, $\chi^2(1) = 1.69$, $p = .19$. People made reasonably confident predictions about each color's category across both conditions.

Finally, we asked whether the effect of condition on confidence depended on grid type (action-based/non-action-based). Indeed, controlling for category guesses (dangerous/safe), the effect of condition on confidence ratings was moderated by grid type, $\chi^2(1) = 10.46$, $p = .001$ (see Fig. 3). For action-based grids, there was no evidence for an effect of condition on confidence, $b = -3.08$, 95% CI $[-11.53, 5.37]$, $\chi^2(1) = 0.52$, $p = .47$, but for non-action-based grids, the effect of condition was marginally significant, $b = -7.14$, 95% CI $[-14.78, 0.50]$, $\chi^2(1) = 3.37$, $p = .07$. This provides weak evidence for the role of communicative assumptions in learning from designed environments.

Discussion

How do people learn from pedagogically designed environments? Study 2 demonstrates that people make confident, often accurate inferences from designed environments—even in the absence of interaction with these environments. However, these inferences did not appear to be based (solely) on communicative assumptions, as there was little evidence that category guesses or confidence varied between the random and pedagogical conditions.

One possible explanation for these unexpected results is that the tendency to interpret environment design as communicative was not successfully overridden by our "random" instructions. Instead, participants might judge whether an environment is randomly or intentionally designed based on intuitive beliefs about randomness (Griffiths et al., 2018). However, it is also possible that people do not base their inferences on communicative assumptions, as we initially predicted. Future research is needed to explore these possibilities. Nonetheless, the results of this study suggest that people readily make inferences from designed physical environments, and they often do so accurately.

General Discussion

In the present research, we investigated (1) how people design the physical environment with the goal of teaching and (2) how people learn from the design of the physical environment. In Study 1, we investigated whether people take action-based and communicative approaches to teaching. The action-based approach predicts that teachers should design environments to prompt learners to take certain actions, while the communicative approach predicts that teachers should design environments that communicate particular inferences. We found evidence for both approaches to teaching.

In Study 2, we investigated learners' inferences from pedagogically designed environments. We found that people make

systematic inferences about the objects in the environment *before* beginning to explore. However, learners' inferences were not solely driven by communicative assumptions. Instead, even learners who believed the environment was generated by a random computer program made systematic inferences that were comparable to the inferences of those who believed the environment was designed by a helpful teacher. These inferences were sometimes (but not always) accurate.

Interestingly, the accuracy of participants' inferences differed systematically across environments. Why might this be the case? First, some environments may not have been designed with communicative intent (i.e., were solely action-based), and thus were not well-suited to communicate accurate inferences. However, it is also possible that some teachers misrepresented learners' beliefs (e.g., see Aboody, Velez-Ginorio, Santos, & Jara-Ettinger, 2023), and thus failed to communicate effectively despite their best attempt to do so. To disentangle these possibilities, our ongoing work focuses on building computational models of teaching and learning through pedagogical environment design.

Many additional questions remain open. For example, when is pedagogical environment design the best method for teaching, and when should more direct teaching methods (e.g., direct instruction, demonstration) be preferred? The literature on guided play emphasizes that child-directed play might increase children's engagement relative to direct instruction (Weisberg et al., 2013, 2016). However, there are cases where active control over one's own learning can be harmful, leading learners to inaccurate inferences (Liquin & Gopnik, 2022; Rich & Gureckis, 2018). It is an open question whether pedagogical environment design can shift learners towards accurate inferences in these cases.

Finally, one notable limitation of the present research is that our artificial grid-based learning environment may differ from real-world learning environments in important ways. For example, our participants generated a single static environment, while real-world teachers can often change the environment over time. This might be particularly impactful for learning: for example, it might be beneficial to interleave learners' encounters with different categories (Kornell & Bjork, 2008) or to present easier learning challenges before more difficult learning challenges (Peng et al., 2018). Future work is needed to investigate how these temporal factors are integrated into pedagogical environment design.

Broadly, this research sheds light on an understudied but important way that humans teach and learn. Parents set up playrooms, teachers construct classrooms, and companies design educational toys, mobile apps, and games. In all these cases, learners interact with a designed physical environment, which may shape their learning in subtle ways. By understanding how people construct environments to teach—and how learners interpret these pedagogically designed environments—we can better understand how humans come to understand the world around them.

Acknowledgements

We thank members of the Computation & Cognition Lab at New York University for their helpful comments on this work, and we thank the anonymous reviewers for their valuable feedback. This work was supported by the National Science Foundation SBE Postdoctoral Research Fellowship under Grant No. 2204021 to E.G.L. and NSF Award No. 2021060 to T.M.G.

References

- Aboody, R., Velez-Ginorio, J., Santos, L. R., & Jara-Ettinger, J. (2023). When naïve pedagogy breaks down: Adults rationally decide how to teach, but misrepresent learners' beliefs. *Cognitive Science*, 47(3), e13257.
- Bass, I., Bonawitz, E., Hawthorne-Madell, D., Vong, W. K., Goodman, N. D., & Gweon, H. (2022). The effects of information utility and teachers' knowledge on evaluations of under-informative pedagogy across development. *Cognition*, 222, 104999.
- Bonawitz, E. B., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition*, 120(3), 322–330.
- Bustamante, A. S., Hassinger-Das, B., Hirsh-Pasek, K., & Golinkoff, R. M. (2019). Learning landscapes: Where the science of learning meets architectural design. *Child Development Perspectives*, 13(1), 34–40.
- Butler, L. P., & Markman, E. M. (2012). Preschoolers Use Intentional and Pedagogical Cues to Guide Inductive Inferences and Exploration. *Child Development*, 83(4), 1416–1428.
- Fisher, K. R., Hirsh-Pasek, K., Newcombe, N., & Golinkoff, R. M. (2013). Taking shape: Supporting preschoolers' acquisition of geometric knowledge through guided play. *Child Development*, 84(6), 1872–1878.
- Gibson, J. J. (1977). The theory of affordances. In R. E. Shaw & J. Bransford (Eds.), *Perceiving, acting, and knowing* (pp. 56–60). Lawrence Erlbaum Associates.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Psychology press.
- Griffiths, T. L., Daniels, D., Austerweil, J. L., & Tenenbaum, J. B. (2018). Subjective randomness as statistical inference. *Cognitive Psychology*, 103, 85–109.
- Ho, M. K., Cushman, F., Littman, M. L., & Austerweil, J. L. (2021). Communication in action: Planning and interpreting communicative demonstrations. *Journal of Experimental Psychology: General*, 150, 2246–2272.
- Ho, M. K., MacGlashan, J., Littman, M. L., & Cushman, F. (2017). Social is special: A normative framework for teaching with and learning from evaluative feedback. *Cognition*, 167, 91–106.
- Kornell, N., & Bjork, R. A. (2008). Learning concepts and categories: Is spacing the “enemy of induction”? *Psychological Science*, 19(6), 585–592.
- Liquin, E. G., & Gopnik, A. (2022). Children are more exploratory and learn more than adults in an approach-avoid task. *Cognition*, 218, 104940.
- Mertens, S., Herberz, M., Hahnel, U. J., & Brosch, T. (2022). The effectiveness of nudging: A meta-analysis of choice architecture interventions across behavioral domains. *Proceedings of the National Academy of Sciences*, 119(1), e2107346118.
- Norman, D. A. (1988). *The design of everyday things*. Basic Books.
- Peng, B., MacGlashan, J., Loftin, R., Littman, M. L., Roberts, D. L., & Taylor, M. E. (2018). Curriculum design for machine learners in sequential decision tasks. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2(4), 268–277.
- Popp, P. O., & Gureckis, T. M. (2020). Ask or Tell: Balancing questions and instructions in intuitive teaching. In S. Denison, M. Mack, Y. Xu, & B. Armstrong (Eds.), *Proceedings of the 42nd Annual Conference of the Cognitive Science Society* (pp. 1229–1235). Cognitive Science Society.
- Rich, A. S., & Gureckis, T. M. (2018). The limits of learning: Exploration, generalization, and the development of learning traps. *Journal of Experimental Psychology: General*, 147(11), 1553–1570.
- Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology*, 71, 55–89.
- Thaler, R. H., & Sunstein, C. R. (2009). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.
- Walker, C. M., Rett, A., & Bonawitz, E. (2020). Design drives discovery in causal learning. *Psychological Science*, 31(2), 129–138.
- Weisberg, D. S., Hirsh-Pasek, K., & Golinkoff, R. M. (2013). Guided play: Where curricular goals meet a playful pedagogy. *Mind, Brain, and Education*, 7(2), 104–112.
- Weisberg, D. S., Hirsh-Pasek, K., Golinkoff, R. M., Kittredge, A. K., & Klahr, D. (2016). Guided play: Principles and practices. *Current Directions in Psychological Science*, 25(3), 177–182.
- Weisberg, D. S., Hirsh-Pasek, K., Golinkoff, R. M., & McCandiss, B. D. (2014). Mise en place: Setting the stage for thought and action. *Trends in Cognitive Sciences*, 18(6), 276–278.