

To Teach Better, Learn First

Oana Stanciu (Stanciu.Oana@phd.ceu.edu)¹

Máté Lengyel (m.lengyel@eng.cam.ac.uk)^{1,2}

József Fiser (FiserJ@ceu.edu)¹

¹Department of Cognitive Science, Central European University
7 Oktober 6 utca, Budapest 1051, Hungary

²Computational and Biological Learning Lab, Department of Engineering, University of Cambridge
Trumpington street, Cambridge CB2 1PZ, UK

Abstract

There has been little cross-fertilization between research on active learning and teaching, despite extensive conceptual similarities. The current study aims to bridge the gap by showing that engaging in active learning can influence subsequent teaching performance. In a one-dimensional boundary teaching task, participants who first took the role of an active learner went on to become better teachers than participants who did not. In order to disentangle the effect of active selection of samples from their information content, the performance of active learners was compared to that of yoked passive learners. While prior passive learning also significantly boosted teaching performance, it did so to a lesser extent. However, in paired comparisons, teachers with active learning experience did not differ significantly from their yoked-passive learning counterparts. Based on the current results we cannot argue for a teaching benefit specific to active learning as opposed to a more general improvement caused by experiencing the task from the learner's perspective. However, we suggest that this is a promising line of inquiry using more complex learning and teaching tasks.

Keywords: teaching; active learning; evidence selection

Introduction

Perhaps the most enduring debate in the education literature, as well as around kindergartens and classrooms, concerns the virtues of exploratory play in contrast to the canonical, largely passive mode of teacher-led instruction (Bruner, 1961; Mayer, 2004). The discussion has been naturally phrased in terms of the relative benefits and disadvantages that the learner incurs when learning from self-guided discovery compared to direct instruction. However, the complementary, and equally important, link between efficient self-guided learning and good teaching has remained largely unexplored.

The common thread running between teaching and active learning is easy to identify when comparing their formal descriptions. Recent rational-agent models have conceptualized teaching as a recursive process in which the teacher and the learner reason about each other. Specifically, the teacher selects training samples for the learner such that, given the learner's prior knowledge and inference making mechanisms, these samples would lead the learner to the desired conclusion efficiently, i.e. by requiring the smallest number of samples (Shafto, Goodman, & Griffiths, 2014). Conversely, the learner interprets the observed samples assuming they were generated by this pedagogical process (as opposed to randomly). Similarly, an ideal active learner will also sample the

environment strategically. However, they will do so by directing their information gathering (e.g. by moving their eyes to explore a visual scene or choosing interventions on the environment) in order to maximize their expected information gain (Yang, Wolpert, & Lengyel, 2018). There are two ways in which active learning can be advantageous. First, observations collected in a strategic way will be more informative for any learner (not just the one sampling information); for instance, by avoiding irrelevant or redundant evidence. Second, and more importantly, there is an added advantage specific to the active learner stemming from the fact that they sample information in light of their prior knowledge and the hypotheses that they wish to test. This effect was demonstrated in experiments in which the data selected by an active learner was also presented to a yoked "passive" learner, and, despite the observations being matched, active learners performed better than their yoked passive counterparts (Markant & Gureckis, 2014).

Thus, both being a good teacher and a good active learner rest on the same general ability to evaluate the potential value of a new piece of evidence relative to a current state of knowledge and a task. Nonetheless, there are important differences. First, teaching brings the added complexity of selecting data for the use of another agent, who might differ widely from the teacher in their state of knowledge and inference making. In line with this, Bass, Shafto, and Gopnik (2017) have linked Theory of Mind (ToM) development to children's pedagogical sampling ability. Second, the active learner does not have access to the target hypothesis, and thus can only select data that minimize uncertainty. However, Yang, Vong, Yu, and Shafto (2019) recently proposed a reconceptualization of active learning as self-teaching by envisioning a learner who simulates an uninformed teacher whose task is limited to providing queries. In this framework, the self-teacher does not optimize for expected information gain, although this will often be the collateral result. Thus, despite differences, it is still feasible to think about teaching and active learning as two highly related cognitive processes.

Given the computational similarity of teaching and active learning, is it possible that they are also integrated through linked processes in human behavior? In other words, would it be possible to hone teaching skills through active learning?




Figure 1: Example image array from the teaching task. In this trial, food items were sorted from left to right in ascending order of their vitamin B content. The black vertical bar represents the daily recommended dose of vitamin B, which is the boundary the participant had to teach. In this case, the participant clicked on the two images closest to the boundary, which were automatically labelled.

Intuitively, taking the perspective of the learner prior to teaching should be a useful experience. It could allow the teacher to better understand, even if implicitly, how a learner would make inferences to solve the task at hand based on the data provided, which in turn would help refine the data selection process.

Taking this reasoning one step further, having the experience of being an active learner prior to teaching should generate robust insights about how to select good examples for teaching in similar tasks. Additionally, if both tasks rely on a core ability to sample environmental data efficiently, the transfer could occur automatically during learning, without the knowledge or expectation that the acquired information will need to be used for teaching in the future. Furthermore, active learning should improve teaching performance beyond passive learning (even when the same information content is acquired) if the active selection of data was the crucial driver of the learning effect, rather than the benefit of familiarity with the teaching task or taking the perspective of a learner.

Experiment

In order to test the hypothesis that active learning improves teaching performance, we designed a simple task in which participants were required to both learn a one-dimensional categorization boundary, and teach it, in counterbalanced order. Thus, there were two independent groups of participants in our design, those who learned actively first and then taught, and those who first taught and then performed active learning. In addition, to probe whether the effect learning on teaching performance was specific to active learning, a yoked control group performed the same teaching task after learning passively from watching the active learners labeled queries.

Method

Participants Eighty-eight participants (54 female, $M_{\text{age}} = 24$ years, range = 18 - 42 years old) were recruited from the local population through the university online research participation system and the student union. Ethical approval was obtained from the United Ethical Review Committee for Research in Psychology (EPKEB) in Hungary.

Tasks All tasks (active learning, passive learning, and teaching) consisted of three trials. In each trial, participants were shown eight images in a horizontal array such as the one displayed in Figure 1. Participants were told that the images were sorted left-to-right according to a given "key" feature. For instance, animals were sorted according to their speed relative to body size or the average amount of time they sleep, or foods were sorted by their carbon footprint or their vitamin content. Images belonged to one of two categories (which were clearly marked at the extremes of the image array) according to whether their key feature was below or above a "boundary" (threshold value) which lied between two adjacent images (i.e. at one of seven possible locations). Unknown to the participants, the true boundary locations which dictated the category membership of the images were uniformly sampled in each trial from all the possible locations.¹

The categories used for the learning and teaching tasks were randomly selected for each participant. Images and category cover stories were only presented once throughout the entire experiment.

In **active learning trials**, participants first saw the image array alongside the description of the categories and the boundary, following which they were told that their task was to find the boundary by querying two images. An image could be queried by clicking on it, which immediately revealed its category membership through the color of the frame drawn around it. After the second query, participants were asked to pinpoint where they thought the boundary was located, again by clicking on one of the seven possible boundaries. Participants received feedback on whether they were correct, un-

¹Participants were provided with a description of a seemingly objective classification boundary (e.g. that slow and fast animals were separated by the speed of the average human scaled by size). These descriptions were intentionally chosen such that the participants were unlikely to have any strong priors about the location of the boundary. Knowing the participants' prior was essential because it determined the optimal query choice in active learning. The six categories and boundary descriptions used in the experiment were chosen based on a pilot in which participants were asked to select the boundary location by relying merely on their prior knowledge. The distribution of chosen boundaries (across participants) for the items included in the current experiment was not significantly different from the uniform distribution.

lucky (they selected a boundary compatible with the labelled images that was not the true boundary) or incorrect (selected an incompatible boundary).

The **passive learning trials** had the same structure, except that the labels of two images were sequentially revealed to the participants before they had to make their decision about the location of the boundary. Crucially, for each passive learning participant, the images labelled corresponded to the queries of a previous active learning participant.

In **teaching trials**, participants were shown the boundary separating the two categories and were asked to teach it to another participant who they were told would take part in the experiment at a later time. It was made explicit that the other participant would be presented with the same set of sorted images. The participant only needed to click on an image to mark it as an example, and it was automatically labelled. Mirroring the learning tasks, participants were only allowed to provide two examples, which is the number of examples sufficient to fully specify the correct boundary. Intuitively, selecting two adjacent images with different labels is sufficient to identify the boundary in this task.

Materials All the images were selected from the MultiPic databank of standardized color drawings of concrete concepts (Duabeitia et al., 2018).

Procedure Participants were pseudo-randomly assigned to one of three groups: active learning followed by teaching ($N = 29$), passive learning followed by teaching ($N = 29$), and teaching followed by active learning ($N = 30$). The experiment was presented on a 27inch screen in a quiet room and lasted for an average of 20 minutes (unspedded). Following the experiment, participants completed an open-ended questionnaire about the strategies that they used to solve the tasks.

Quantifying performance Teaching performance was measured by the information gain, IG_{teach} , which is the amount of entropy by which the teacher reduced the imagined learner’s prior entropy $\mathbb{H}(b)$ by labelling two images:

$$IG_{\text{teach}} = \mathbb{H}(b) - \mathbb{H}(b|s_1, s_2, l_1, l_2)$$

where s , l , and b respectively denote image stimuli, category labels, and potential boundary locations. $\mathbb{H}$ is the Shannon entropy over the possible hypotheses, the prior entropy is $\mathbb{H}(b) = -\sum_{b \in \mathcal{B}} P(b) \log_2 \frac{1}{P(b)}$, where $P(b)$, the learner’s prior over the boundary locations, is assumed to be uniform. The optimal teaching strategy is to label the examples immediately preceding and following the boundary as this will eliminate any uncertainty about the location of the boundary, thus reducing all of the original entropy. On the other hand, selecting an example set that will leave the learner uncertain about the true hypothesis because many potential boundaries compatible with the example set will translate into a lower information gain.

Using the observed information gain to evaluate active learning performance would introduce arbitrariness since it cannot distinguish a learner’s well-planned query from a

lucky one. An ideal learner should choose a query in light of their uncertainty about the labels that will be observed. First, learners should compute the expected information gain of the queries by weighing the posterior entropy by the probability of observing the given labels for the query made and then choose the query that maximizes the expected gain. Therefore, EIG_{learn} , the sum of the expected information gain of the first and second queries, was used instead of observed information gain. The expected information gain of the first query is:

$$EIG_{\text{learn}}(s_1) = \mathbb{H}(b) - \sum_{l_1 \in \mathcal{L}} \mathbb{H}(b|s_1, l_1) \cdot \sum_{b \in \mathcal{B}} P(l_1|s_1, b) P(b)$$

After observing the first label, the prior over the boundary locations is updated, and the expected information gain is computed again relative to the entropy remaining after the first labelled sample:

$$\begin{aligned} EIG_{\text{learn}}(s_2|s_1) &= \\ &= \mathbb{H}(b|s_1, l_1) - \sum_{l_2 \in \mathcal{L}} \mathbb{H}(b|s_2, l_2, s_1, l_1) \cdot \sum_{b \in \mathcal{B}} P(l_2|s_2, s_1, l_1, b) P(b) \end{aligned}$$

Unless otherwise specified, statistical analyses of participants’ responses were performed based on the average measures of IG_{teach} and EIG_{learn} in the three trials of each task.

Decisions about the boundary location In learning trials, after observing two labelled stimuli, participants marked the location of the categorization boundary. Their choice could be assessed based on whether or not the selected boundary was compatible with the labelled images they had seen. However, simply using the proportion of compatible answers (across the three trials) to assess their performance ignores the fact that trials differed in the number of remaining compatible boundaries. To control for this and characterize performance appropriately, we fitted a model that captured the intuition that participants behaved optimally and selected (randomly) from among the remaining compatible boundary locations in some r fraction of trials, while in the rest of the trials they “lapsed” and selected a boundary randomly among all locations:

$$\begin{aligned} P(\text{choice} = b_i|s_1, s_2, l_1, l_2) &= \\ &= r \cdot \mathbb{1}\{b_i \in \mathcal{B}_{\text{compatible}}^{(i)}\} \cdot \frac{1}{|\mathcal{B}_{\text{compatible}}^{(i)}|} + (1-r) \cdot \frac{1}{|\mathcal{B}|} \end{aligned}$$

Thus, $r = 1$ indicates optimal behavior, while $r = 0$ indicates chance performance. We estimated r for each participant by maximum likelihood (under the assumption that trials were *i.i.d.*).

Data analysis Predictions were tested using planned independent t-tests to compare the teaching information gain in the teaching first and learning first conditions. Paired comparisons were used for the two groups who experienced being learners first, the active learners and passive learners.

Post-hoc analyses were conducted to ensure that variables extraneous to the predictions did not have a meaningful impact on performance or modulate the reported effects. The design of the experiment lends itself naturally to mixed model analysis, since it allows fitting trial level data (without aggregation) and can describe variation arising from the experimental design. Starting from a baseline fixed effects only model with the experimental condition as a predictor of teaching performance, we sequentially fitted and compared models using two additional fixed factors, learning performance and trial number (and their interactions with the condition), as well as random intercepts for participant and trial identity (i.e. dimension used for classification of the objects). Fixed effects were tested using log-likelihood ratio tests for nested models with the same random effects structure. Non-nested models were compared using the Bayesian Information Criterion (BIC) and Akaike Information Criterion (AIC). Similarly, random effects (fitted via maximum likelihood) were tested using log-likelihood ratio tests while keeping the fixed effects model identical. Given that the mixed-effects analysis confirmed the results of the planned comparisons on the aggregated trial data, we will focus on these comparisons in the Results section for brevity and clarity.

Results

Descriptives Despite the surface level simplicity of the teaching task, a large proportion of participants ($\approx 60\%$) did not perform it optimally (i.e. did not choose the two images on either side of the boundary as the teaching samples). However, prior active learning made it easier to gain insight into the optimal solution for teaching. More than half of active learners, 17 out of 29 participants, performed at ceiling level by consistently providing example sets compatible with only one categorization boundary. In contrast, only 11 of 29 participants in the yoked passive learning group, and 7 out of 30 of the participants who did not complete a learning task before teaching managed to select the optimal example sets.

Teaching performance across conditions As predicted, participants who were active learners before being teachers outperformed those who started directly with teaching, on average providing .63 bits, 95% CI [.22, 1.05], of additional information to their (fictitious) learners (see Figure 2). The group difference was highly significant in an independent t-test, $t(57) = 3.04$, $p = .01$, Bayes Factor (BF)² = 10.81 in favor of the alternative hypothesis.

Learning passively before teaching conferred a smaller, but still significant, benefit relative to foregoing learning. Passive learning increased teaching information gain by an average of .45 bits, 95% CI [.05, .85], $t(57) = 2.26$, $p = .03$, $BF = 2.16$ in favor of the alternative.

While we found strong evidence in support of the differ-

²Bayes Factors were calculated for a null model that assumes a zero standardized difference between groups, and a Cauchy alternative with a prior scaled to an effect size of .7, following Rouder, Speckman, Sun, Morey, and Iverson (2009).




Figure 2: Teaching and learning performance across the three conditions. Each dot represents the information gain for one participant, averaged across the three trials of each task. Crosses represent the 95% confidence intervals for the group means. Dotted lines represent the expected mean information gain from teaching as a function of expected information gain. The maximum information gain for the task is 2.81 bits. The asterisks mark significance levels in independent t-tests (* $p < .05$, ** $p < .01$).

ences between the groups completing the learning and teaching tasks in different orders, a possible concern was that these differences were not induced by the experimental manipulation *per se*. Specifically, if there are prior differences in learning performance favoring the group that completed the active learning task first, and learning performance is correlated with teaching performance, then the condition effect could be just an artifact. In order to eliminate this possibility, a regression was performed on teaching performance with both the group (active learning before / after teaching³), learning performance, and their interaction as predictors. The group difference remained significant, $\beta = .62$, $p = .01$, when controlling for expected information gain in learning, which was not a significant predictor of teaching ability, $\beta = .08$, $p = .81$, nor did it interact with the group effect, $\beta = .68$, $p = .3$. Figure 2 shows, for each condition, the estimated (non-significant) slopes for information gain from teaching as predicted by ex-

³The same pattern of results was found for the difference between the group learning passively and then teaching, and the one teaching before active learning.




Figure 3: Teaching performance for the active-passive learning dyads. Each dot represents the information gain from teaching for one dyad. In dyads situated under the diagonal identity line, the active learner was the better teacher. A small Gaussian scatter was applied to make overlapping dots visible.

pected information gain for learning. Coupled with the fact that the difference in active learning performance between the two groups was not significant, $t(57) = 1.77$, $p = .08$, $BF = 1.28$ in favor of the null hypothesis, this suggests that the effect of the manipulation was not mediated by prior differences in active learning performance. To investigate this issue further, the two groups were repeatedly resampled with replacement such that the learning performance between groups could be matched and fixed at different levels. Comparing the teaching performance across these resampled groups confirmed the advantage of those who completed the learning tasks prior to the teaching task (the 95% CI of the mean of the resampled groups' differences did not include a null effect).

The second prediction of the study was that active learners would gain a larger benefit from learning before teaching than the yoked passive controls. Active learners fared on average only slightly better in the teaching task than their passive learning counterparts who were shown the same labelled data, with an average difference of .18 bits, 95% CI [-.11, .47]. The dyads' performance is illustrated in Figure 3. The difference was not significant in a paired t-test, $t(28) = 1.29$, $p = .21$, $BF = 2.39$ in favor of the null. It should be noted though that the paired comparison was underpowered (post-hoc power = .24) given the magnitude of the effect size observed.

While there was no significant difference in teaching performance in the planned, marginal comparison between the dyads, Figure 2 suggests that differences may potentially be

present conditional on learning performance. There was no interaction between the three-level condition and learning performance, however, this analysis does not account for the dependence in the active learning and passive learning dyad data. As pairs of active and yoked passive learners had, by design, the same expected learning information gain, we regressed the within dyad difference in teaching performance against learning expected information gain. Learning performance was not a significant regressor of the difference in teaching, $\beta = -.86$, $p = .07$. On the one hand, the predicted within-dyad difference, conditioned on low values of learning performance, was significant (see Figure 4). For instance, the predicted within-dyad teaching difference was .60 bits, $p = .03$, at a one bit expected learning entropy. On the other hand, there was no discernible difference for dyads with high expected information gain. While this is not a strong result, given the low number of dyads and the small effect, it might suggest a potential modulation of the relative benefit of active learning.

Mixed effects analysis The best-fitting model contained the condition, $F(2,85) = 4.30$, $p = .02$, and trial number, $F(2,174) = 6.93$, $p = .01$, as fixed effects, alongside a participant level random intercept ($SD = .70$). The addition of the random intercept was judged meaningful based on the magnitude of the variance at the participant level ($SD = .70$). It also led to a reduction in BIC, from 796.6 for the fixed effects only model to 749.7.

The previous results regarding the condition effect hold, with a significant estimated difference of .63 bits, $se = .22$, $t(85) = 2.94$, $p = .01$, between the active learning first and teaching first conditions. Similarly, no significant difference was found between active and passive learners, estimated difference of .32 bits, $se = .22$, $t(85) = 2.94$, $p = .01$. Additionally, teaching performance improved from the first to the third trial by an estimated .38 bits, $se = .11$, $t(174) = 3.30$, $p = .01$. However, performance improvement from the first to the second trial was not significant, .02 bits, $se = .11$, $t(174) = .17$, $p = .87$.

Decisions about the boundary location In the active learning first condition, the mean of the best-fit individual r values was .79 ($SD = .35$), whereas for those completing the active learning following teaching it was lower, .58 ($SD = .42$). Yoked controls has the smallest average r , .51 ($SD = .38$). Active learners made better inferences about the boundary location than their matched controls as the average within-dyad difference in estimated probability r was .28, $t(28) = 2.99$, $p = .01$, $BF = 7.29$. The order of the active learning task led to marginally significant differences in an independent t-test, $t(57) = 2$, $p = .05$, $BF = 1.37$ in favour of the alternative.

The difference in r within active-passive learning pairs did not correlate significantly with differences in teaching performance, $r(26) = -.28$, $p = .13$.




Figure 4: The difference in teaching information gain within dyads of active and passive learners as a function of the expected information gain for (active) learning. The fitted OLS regression line is shown alongside its 95% confidence bound.

Discussion

It has been proposed that humans have a likely innate sensitivity to pedagogical guidance (Csibra & Gergely, 2009) and a propensity for teaching others. From a normative standpoint, the prevalence of teaching in social groups is to be expected given that learning from others who are knowledgeable, well-intentioned and attuned to the learner is more efficient than self-guided learning. Experimental evidence is also accumulating to suggest that, at least in constrained laboratory settings, the behavior of human teachers matches the predictions of normative models (Shafto et al., 2014). However, while we know that humans are effective and keen teachers, we don't know much about the underlying abilities enabling teaching and how it relates to performance in other tasks, specifically active learning.

In the current study we observed an improvement in teaching performance for participants who engaged in active learning prior to teaching. Three active learning trials, using different stimulus sets than those used for teaching, were sufficient for the majority of participants to gain insight into the optimal solution of the teaching problem on the first attempt. Furthermore, they were able to draw on their experience as learners even though at the time of learning they had not been aware that the teaching task would follow.

The poor performance of participants with no learning experience resonates with previous findings of Khan, Zhu, and Mutlu (2017), who used a boundary teaching task as well, but did not constrain the example set size by their design. It seems that simply asking teachers to generate the minimally sufficient number of examples for optimal teaching was not enough to solicit the optimal solution.

The fact that the active learning benefit, relative to teaching first, was not modulated by the initial active learning performance suggests that active learning can improve teach-

ing across the board, for poor and good active learners alike. However, prior active learning performance may play a role in differentiating teachers in a more complex teaching scenarios. Indeed, the surprising lack of a significant correlation between active learning performance and subsequent teaching performance can be explained by ceiling effects.

The impact of passive learning on teaching, relative to the baseline teaching first group, was smaller than that observed for active learning. However, we did not find a significant effect in the matched comparison between active and yoked passive learners. It is important to note here that the current task can be thought of as an insight problem, which means that there was less scope for observing gradual differences in performance. Further, once insight was achieved in the learning task, the solution was easy to verbalize, allowing the optimal strategy to be explicitly transferred to the teaching task.

On the other hand, for poor performing learning dyads, we observed a difference in the predicted direction. This suggests that in a more complex and ecological task in which the learning is more gradual, and the optimal solution is explicitly unknown to participants, active and yoked passive learners are likely to diverge more in terms of teaching performance. This would provide evidence for a more automatic, implicit link between active learning and teaching. In such a future teaching task it would also be interesting to examine whether the differences between active and passive learners, matched for information content, are moderated by the quality of the queries they both observe. Specifically, it should be tested whether the negative linear trend we observed generalizes to non-insight tasks.

Lastly, it is surprising that those who performed the teaching task prior to the active learning task did not differ in their expected information gain in the learning task, and, if anything, performed poorer than their counterparts who started by active learning. This resonates with previous experimental evidence from the developmental literature that has also highlighted more subtle ways in which being taught can hinder learning, for instance by limiting subsequent exploration (Bonawitz et al., 2011). It is an intriguing idea that, perhaps, not just the experience of being taught, but also teaching itself, can have an effect on exploration. Alternatively, if we assume that the teaching task is more cognitively demanding as it has a meta-cognitive component engaged in reasoning about the learner's knowledge and inference making, results can be explained by the known effect that an easier-to-harder progression of tasks is beneficial for learning, while the opposite order does not provide an appropriate stepping stone for active learning. On the other hand, Yang et al. argue that active learning can be re-formalized to also include a meta-cognitive aspect, reasoning that is applied reflexively to one's own reasoning.

To conclude, active learning proved to be a reliable intervention to improve teaching performance. It is important to investigate if the effect of active learning generalizes to more

complex and more ecologically valid tasks, or even between different learning and teaching tasks. If it does, it will open the way for quantitative inquiries about whether successful teaching benefits from the ability of taking the perspective of an active learner and as such can be improved by prior active learning.

Acknowledgements This research has received funding from the European Research Council (ERC) Consolidator Grant (ERC-2016-COG/726090) awarded to Máté Lengyel.

Yang, S. C.-H., Vong, W. K., Yu, Y., & Shafto, P. (2019). A unifying computational framework for teaching and active learning. *Topics in Cognitive Science*.

Yang, S. C.-H., Wolpert, D. M., & Lengyel, M. (2018). Theoretical perspectives on active sensing. *Current Opinions in Behavioural Science*, 11, 100–108.

References

- Bass, I., Shafto, P., & Gopnik, A. (2017). I know what you need to know: Childrens developing theory of mind and pedagogical evidence selection. In *Proceedings of the 39th annual conference of the cognitive science society* (p. 6).
- Bonawitz, E., 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.
- Bruner, S., J. (1961). The act of discovery. *Harvard Educational Review*, 31, 21–32.
- Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, 13(4), 148–153.
- Duabaitia, J. A., Crepaldi, D., Meyer, A. S., New, B., Pliatsikas, C., Smolka, E., & Brysbaert, M. (2018). MultiPic: A standardized set of 750 drawings with norms for six European languages. *Quarterly Journal of Experimental Psychology*, 71(4), 808–816.
- Khan, F., Zhu, X., & Mutlu, B. (2017). How do humans teach: On curriculum learning and teaching dimension. *Advances in Neural Information Processing Systems 30 (NIPS 2017)*.
- 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), 94–122.
- Mayer, R. E. (2004). Should there be a three-strikes rule against pure discovery learning? *American Psychologist*, 59(1), 14–19.
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16(2), 225–237.
- 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.
- Yang, S. C.-H., & Shafto, P. (2017). Teaching Versus Active Learning: A Computational Analysis of Conditions that Affect Learning. In *Proceedings of the 39th annual conference of the cognitive science society*.