The impact of biased hypothesis generation on self-directed learning

Doug Markant (markant@mpib-berlin.mpg.de)
Center for Adaptive Rationality
Max Planck Institute for Human Development
Lentzeallee 94, Berlin, Germany

Abstract

Self-directed learning confers a number of advantages relative to passive observation, including the ability to test hypotheses rather than learn from data generated by the environment. However, it remains unclear to what extent self-directed learning is constrained by basic cognitive processes and how those limits are related to the structure of the to-be-learned material. The present study examined how hypothesis generation affects the success of self-directed learning of categorical rules. Two experiments manipulated the hypothesis generation process and assessed its impact on the ability to learn 1D and 2D rules. Performance was strongly influenced by whether the stimulus representation facilitated the generation of hypotheses consistent with the target rule. Broadly speaking, the findings suggest that the opportunity to actively gather information is not enough to guarantee successful learning, and that the efficacy of self-directed learning closely depends on how hypothesis generation is shaped by the structure of the learning environment.

Keywords: self-directed learning, category learning, active learning, information search, hypothesis generation

Self-directed learning (SDL) is typically characterized by an interaction between external and internal search processes. Active information collection and exploration of the external environment are hallmarks of SDL (e.g., a student deciding how to study or playing with a new toy to learn how it works). Information resulting from this external search then influences an internal belief updating process, often conceptualized as the sequential generation and evaluation of new hypotheses, which in turn drives subsequent information gathering. This ongoing interaction is central to theories of conceptual discovery, including scientific inquiry (Klahr & Dunbar, 1988), explanatory reasoning (Johnson & Krems, 2001), and sensemaking (Weick, 1995).

Although previous research has described a number of benefits of SDL relative to passive learning conditions, it is less clear how its efficacy is constrained by basic cognitive and perceptual processes (Gureckis & Markant, 2012). Previous work suggests that, in particular, SDL may be strongly limited by failures or biases in the hypothesis generation process. For instance, in concept learning tasks people make effective information search decisions when given a set of hypotheses to discriminate, but often fail to generate alternative hypotheses on their own (Tweney et al., 1980). Biased hypothesis generation may be especially impactful in real-world domains involving large or ill-defined hypothesis spaces. In education, some researchers have argued against pedagogical practices which emphasize self-directed discovery because students often fail to generate the target concept in the absence of guidance that narrows the hypothesis space (Mayer, 2004). Thus, understanding the hypothesis generation process is critical to being able to predict whether SDL will be effective depending on the nature of the learning problem.

Self-directed category learning

A recent study by Markant and Gureckis (2014) suggests that, indeed, the benefits of SDL depend closely on the structure of the target concept. The authors compared passive *reception* and self-directed *selection* in a perceptual category learning task (see stimuli in Figure 1, left), in which the target classification rule was either one-dimensional (1D, a criterion on a single feature dimension) or two-dimensional (2D, requiring the integration of both feature values). Reception-based learners observed stimuli generated from bivariate normal distributions corresponding to each category, whereas selection-based learners chose items to learn about by specifying the feature values. Selection led to higher accuracy than reception among those participants learning 1D rules. In contrast, in the 2D case selection was no better than reception, and participants in both conditions were well-described as responding according to simpler 1D hypotheses.

The goal of the present study was to examine whether this divergence between selection-based learning of 1D and 2D rules was caused by a biased hypothesis generation process. The experimental manipulations were motivated by an account of hypothesis generation as involving the sampling of salient cues and relations in the environment (Trabasso, Bower, & Gelman, 1968; Wallach, 1962). Accordingly, the success of SDL should depend on whether the feature representation facilitates the generation of hypotheses that are of the same form as the target rule. Poor learning of 2D rules in Markant and Gureckis (2014) may have been driven by stimuli with highly distinct dimensions that encouraged the generation of 1D hypotheses, rather than 2D hypotheses that involve integrating two dimensions. In contrast to existing theories of category learning that assume inherent differences in the difficulty of acquiring 1D and 2D rules (e.g., Ashby, Paul, & Maddox, 2011), under this proposal the ability to learn either type of rule through SDL depends on whether the feature representation encourages the generation of hypotheses of the same form.

Experiment 1

The first experiment examined the effect of perceptual feature representations on self-directed category learning. People were predicted to be more likely to generate 1D hypotheses when stimulus features are highly separable or it is difficult to combine information about their relative magnitudes, a bias that should facilitate the learning of 1D target rules.

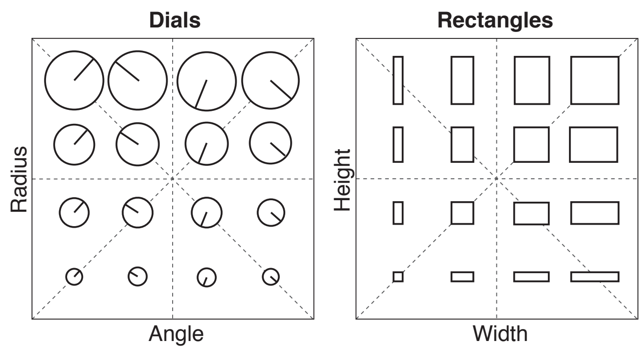

Figure 1: Example stimuli in the DIAL (left) and RECT (right) conditions in Experiment 1. Dotted lines indicate the set of target classification rules. Axis-aligned lines correspond to 1D rules, diagonal lines correspond to 2D rules.

Conversely, people may be more likely to generate 2D hypotheses when the two feature values are easily compared or combined, leading to faster learning of 2D target rules. In both cases, a mismatch between the feature representation and the form of the target rule (e.g., highly separable dimensions when learning a 2D rule) should lead to a longer internal search process and slower learning.

The task involved classifying shapes defined by two continuous feature dimensions into two categories, *A* and *B* (Figure 1). Each participant was assigned to learn either a 1D or 2D rule involving one of two types of stimuli, *dials* or *rectangles*. Dials were composed of a circle that varied in radius and a line segment that varied in orientation, whereas rectangles were defined in terms of their height and width. Participants learned to categorize stimuli through self-directed selection, such that on each trial they chose a combination of features and observed the true category label for that stimulus.

Previous studies involving dial stimuli have shown that the two feature dimensions (radius and angle) are highly separable (Shepard, 1964), whereas rectangle stimuli are perceived in an integral manner, with the dimensions of shape and size more salient than width and height (Macmillan & Ornstein, 1998). Based on this perceptual distinction, participants in the RECT conditions were expected to be more likely to generate 2D hypotheses that involved integrating both features, and as a result, to learn 2D rules more effectively than when stimuli are represented as dials.

Participants and materials

One-hundred twelve participants were recruited from the MPI subject pool (66 female, 44 male, 3 no gender given; $M_{age} = 25.1$, $SD = 5.2$) and were paid 7€ for participation in addition to a bonus based on their performance. Participants were randomly assigned to one of four conditions: 1D-DIAL, 1D-RECT, 2D-DIAL, or 2D-RECT. In addition, they were randomly assigned one of four possible variants of the target rule corresponding to rotations of the rule in the stimulus space.

Test items corresponded to a grid of 256 stimuli that uniformly tiled the feature space. This grid was partitioned into 8 blocks, each of which included 8 stimuli from every quadrant of the stimulus space, for a total of 32 items in each test block. The order of test items within each block and the order of the eight test blocks were randomized for each participant.

Procedure

The experiment was presented in a web browser using the *psiTurk* software (Gureckis et al., 2015). There were 8 blocks in total, each comprised of 16 learning trials followed by 32 test trials. Participants received .02€ for every correct classification during test trials, for a total possible bonus of 5.12€.

Learning trials On each learning trial a stimulus was randomly generated (with feature values drawn from a uniform distribution over their respective ranges) and displayed onscreen. The participant could then simultaneously vary the features by moving the mouse, with vertical mouse movement controlling one feature dimension and horizontal movement controlling the second feature dimension. The mapping between mouse directions and feature dimensions was randomized for each participant. After adjusting the feature values as desired, the participant pressed the spacebar to query the selected stimulus, after which the category label (*A* or *B*) was displayed until the participant pressed a button to complete the trial.

Test trials In each test trial a single test item was displayed at the center of the display. The participant categorized the item by pressing *A* or *B* on the keyboard at their own pace. No feedback was provided during test trials. At the conclusion of each test block, participants were told the proportion of items they classified correctly during that block.

Results

Classification accuracy Accuracy across blocks in shown in Figure 2A. A 2×2 ANOVA on overall classification accuracy (averaged over blocks and rule variants) was performed with stimulus (DIAL vs. RECT), rule (1D vs. 2D), and stimulus $\times$ rule interaction as factors. There was no main effect of stimulus type ($F(1,108) = 2.4$, $p = .12$), but there were significant effects of rule type ($F(1,108) = 22.6$, $p < .001$) and

Table 1: Classification accuracy.

		Pairwise comparisons (Tukey HSD)		
Experiment 1	*M* (*SD*)	1D-RECT	2D-DIAL	2D-RECT
1D-DIAL	.94 (.04)	* .13 [.07, .19]	* .23 [.17, .29]	* .06 [.01, .12]
1D-RECT	.81 (.08)		* .09 [.03, .15]	* .07 [.02, .13]
2D-DIAL	.71 (.09)			* .17 [.11, .23]
2D-RECT	.88 (.09)			
Experiment 2	*M* (*SD*)	1D-REL	2D-ABS	2D-REL
1D-ABS	.95 (.06)	* .08 [.02, .15)]	* .18 [.12, .25]	* .08 [.02, .16]
1D-REL	.86 (.11)		* .09 [.03, .17]	.00 [–.07, .06]
2D-ABS	.76 (.10)			* .09 [.02, .16]
2D-REL	.86 (.12)			

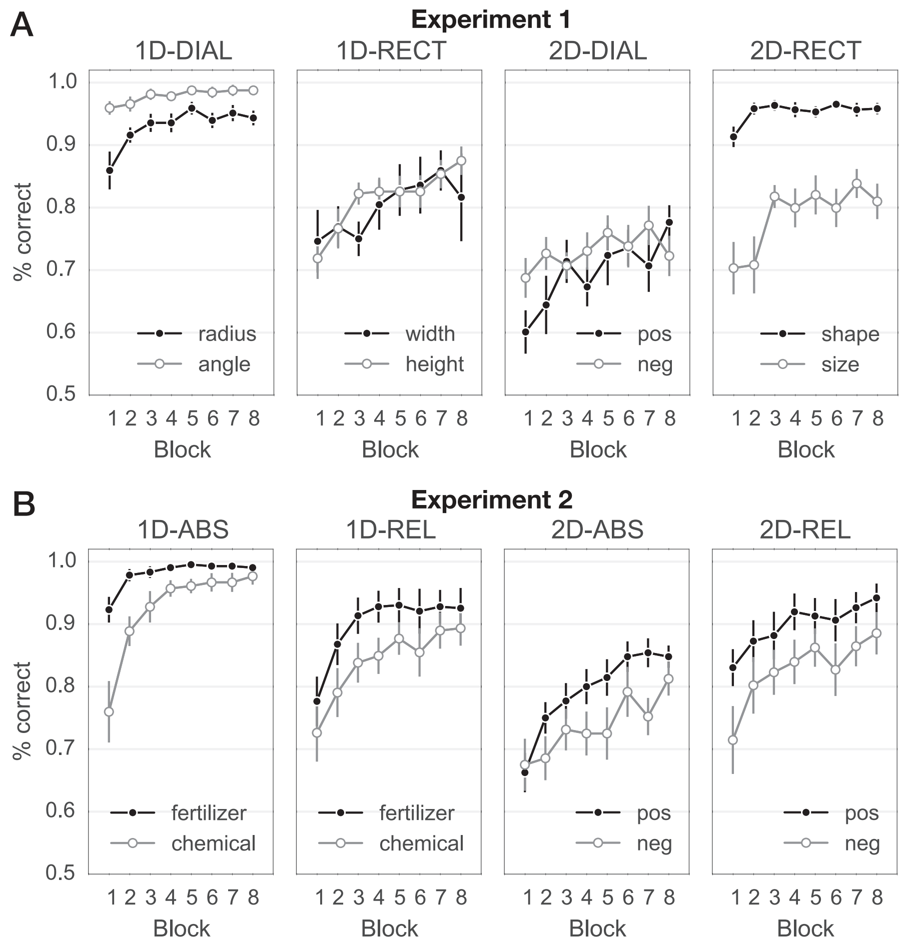

Figure 2: Classification accuracy in Experiment 1 (A) and Experiment 2 (B). Error bars indicate standard errors.

the stimulus × rule interaction ($F(1,108) = 94.4$, $p < .001$). Tukey HSD tests indicated significant pairwise differences in accuracy between all four conditions (see Table 1).

Two-sample t-tests were used to assess effects of rule variants within each condition. For 1D rules, participants were grouped according to the relevant feature dimension (i.e., radius, angle, width, or height). For 2D rules, participants were grouped according to whether the target rule was positively sloped or negatively sloped. In the 2D-RECT condition, positive and negative 2D rules are hereafter referred to as *shape* and *size* rules, respectively. There was no difference between rule variants in the 1D-RECT ($t(11.6) = .34$, $p = .74$) and 2D-DIAL ($t(23.9) = .92$, $p = .36$) conditions. Within the 1D-DIAL condition, participants learning a rule defined on the angle dimension were more accurate than those learning a rule on the radius dimension ($t(19.5) = 4.2$, $p < .001$). Finally, within the 2D-RECT condition, accuracy was higher for participants learning a shape rule as compared to those learning a size rule ($t(13.1) = 6.9$, $p =< .001$).

Modeling classification boundaries The goal of the second analysis was to relate classification performance to the form of hypotheses generated in each condition. Bayesian logistic regression was used to estimate linear decision bounds for each block of test responses, using the *bayesglm* function within the *arm* R package. Four models were estimated for each block: a 2D model with both features as predictors, two 1D models with a single feature as a predictor, and a baseline (intercept-only) model. Decision boundaries were then classified according to the model with the lowest AIC (thirty-two blocks, or 4%, were best-fit by the baseline model, indicating that there were few cases in which a linear decision boundary was not supported by participants' responses). The proportion of 2D boundaries across blocks is shown in Figure 3A.

Logistic regression on the proportion of 2D rules revealed significant effects of stimulus type (RECT: Wald $z = 6.0$, $p < .001$), rule type (2D: $z = 5.8$, $p < .001$), and the stimulus × rule interaction ($z = -4.7$, $p < .001$). As shown in Figure 3A, in the 1D-DIAL and 2D-RECT conditions, nearly all participants respond with a rule of the correct form throughout the task. This included 2D-RECT participants learning a *size rule* who responded with relatively low accuracy, indicating that the poor performance in that condition was not due to a failure to consider 2D hypotheses. In contrast, in both

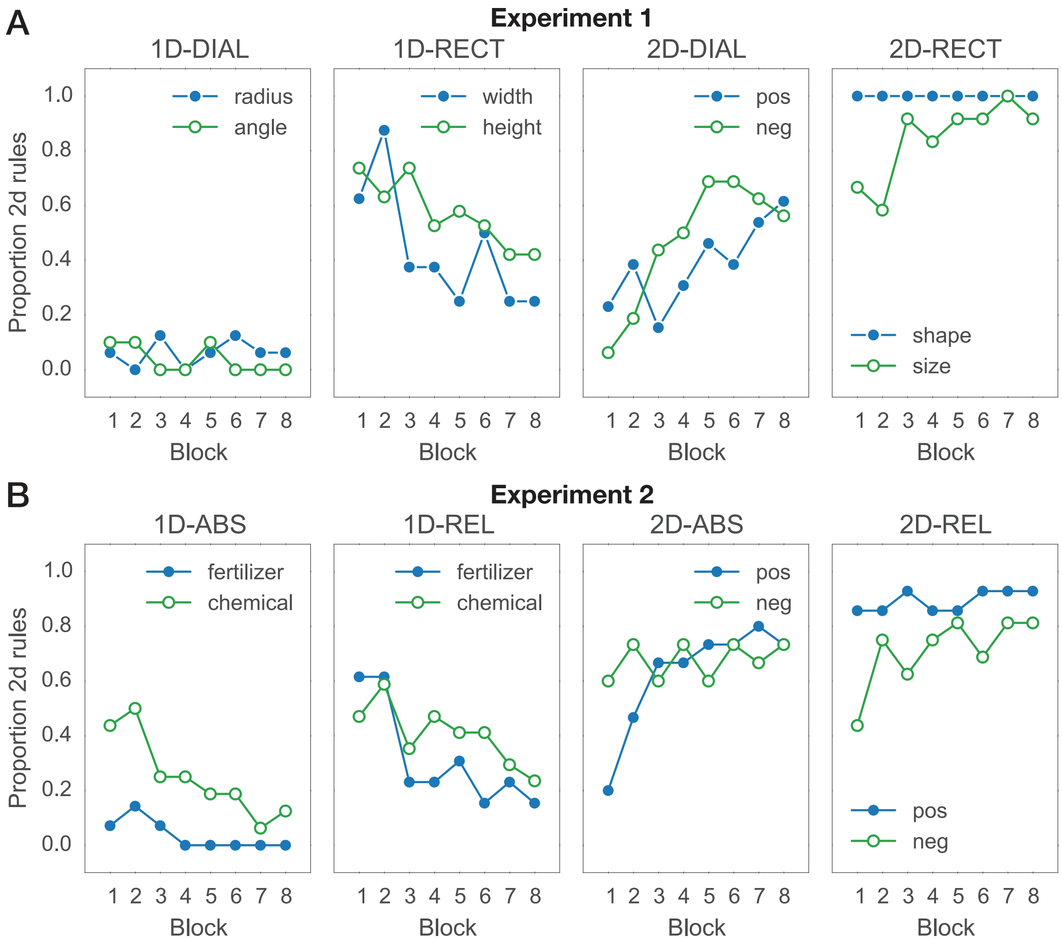

Figure 3: Proportion of best-fit decision boundaries classified as 2D. Within each condition, a separate line is shown for participants learning each type of target rule.

the 1D-RECT and 2D-DIAL conditions, decision boundaries were approximately equally divided between 1D and 2D hypotheses, suggesting that the manipulation of stimulus representation had the predicted effect on the kinds of hypotheses that were generated during learning.

Discussion

The results confirmed the hypothesis that the success of self-directed learning depends on a match between the target rule and salient features of the stimuli. In addition to replicating the difference in performance between 1D-DIAL and 2D-DIAL accuracy observed by Markant and Gureckis (2014), this gap in performance was reversed by changing the stimulus representation in the RECT conditions. The model-based analysis of rule use indicates that poor performance in the 1D-RECT and 2D-DIAL conditions was due, at least in part, to the generation of hypotheses of the wrong form.

In the 2D-RECT condition, performance diverged strongly depending on the form of the target rule, with higher accuracy among participants learning the *shape rule* than the *size rule*. However, rapid learning of the shape rule may be unsurprising given that the category boundary coincided with a simple relational comparison of the two features (i.e., whether the shape is taller than it is wide). The same general pattern was reported by Ashby and Gott (1988, Exps. 1 and 2) using similar stimuli (perpendicular line segments) under passive training. They found that participants were close to optimal performance when learning a positive, 2D classification rule that required comparing the lengths of the two line segments. When tasked with learning a 1D rule with the same stimuli, participants still responded according to 2D rules that integrated the two features. The authors concluded that the ease of comparing the two features led people to adopt 2D decision rules despite being able to separately attend to individual features. The present results show that this interference generalizes to other types of stimuli, and, more surprisingly, has a persistent effect under self-directed conditions despite learners' control over the training experience. This effect is particularly striking in the 1D-RECT condition in which, even on the last block of training, more than a third of participants responded according to some form of 2D hypothesis.

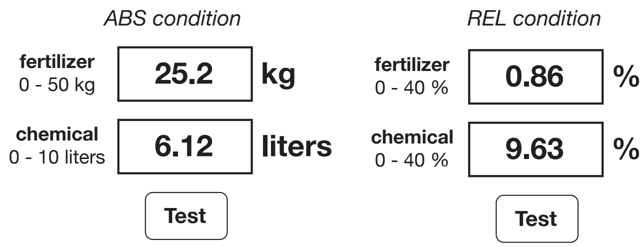

Figure 4: Example training trial displays in the ABS (left) and REL (right) conditions in Experiment 2.

Experiment 2

The goal of the second experiment was to evaluate the generalizability of the previous findings to an abstract domain in which perceptual biases were minimized. The task was designed to be structually equivalent to that of Experiment 1, while involving stimuli defined by abstract numerical features. As in Experiment 1, hypothesis generation was biased by manipulating the ease with which the two feature values could be integrated (see Figure 4). In the *relative* (REL) condition, the two features had the same range and were described in the same units. The common scale was predicted to cause participants to generate hypotheses based on integrating information about the two features. In the *absolute* (ABS) condition, the features had different ranges and were described in terms of different units. Like the DIAL condition in Exp. 1, the ABS condition was predicted to increase the likelihood of generating 1D rules involving a single dimension. Moreover, there should be an interaction between the stimulus representation and the target rule, such that SDL is most effective when the feature description facilitates the generation of the correct form of rule.

Participants and materials

One-hundred twenty people were recruited from the MPI participant pool (54 male, 62 female, 4 no gender given; age: $M = 25.2$, $SD = 3.3$) and were paid in the same manner as in Experiment 1. Participants were randomly assigned to one of four conditions: 1D-ABS, 1D-REL, 2D-ABS, or 2D-REL.

The goal of the task was to learn how the amount of two substances (a *Chemical* and a *Fertilizer*) affected whether a patch of virtual farmland would experience a successful (S) or failed (F) crop. Stimuli were defined by two continuous dimensions corresponding to the quantities of each substance. Within each target rule condition (1D or 2D), participants were assigned one of four possible variants of the target rule corresponding to different rotations of the rule in the stimulus space. In addition, the target rule was offset such that the classification boundary did not bisect the stimulus space.

In the ABS conditions, the two substances were defined in terms of different units (kg or liters) and had different ranges (one feature ranged from 0 to 50 while the second ranged from 0 to 10). In the REL conditions, both dimensions had the same range (0 to 40) and were labeled as percentages of a mixture applied to the soil (see Figure 4).

Test items A set of 32 items were generated by tiling the stimulus space at even intervals. Test sets were generated by randomly perturbing the location of each coordinate by a small amount (5% of the range on each dimension) while ensuring a constant number of items from each category. The order of test items and the order of the eight test blocks were randomized for each participant.

Procedure

The participants' goal was to learn how different feature combinations affected whether a patch of farmland would experience a successful or failed crop. Aside from the cover story and stimuli, the structure of the task was identical to that of Experiment 1. Participants completed 8 blocks, alternating between 16 learning trials followed by 32 test trials.

Learning trials On each learning trial an image of an empty plot of land appeared above two input boxes corresponding to the features. Each input was labeled with the name of the feature, the corresponding units ("kg" and "liters" in the absolute condition; "%" in the relative condition), and the possible range of each feature (Figure 4). Both features were initialized with values drawn from a uniform distribution over the corresponding ranges. Participants could then alter the value of either dimension by entering a new number within the allowed range. They then clicked a button to test the chosen combination of feature values. If the combination led to a successful crop, a new image appeared with fruit on the plot of land and the category label "Success!". If the combination led to a crop failure, a new image without fruit appeared along with the category label "Failure."

Test trials Each test trial began with the presentation of the image of an empty plot of land and a test item. Stimulus values were displayed in the same manner as in learning trials but could not be altered by the participant. Participants clicked on the outcome that they predicted to occur for the displayed feature combination. At the end of each block they were told their proportion of correct predictions.

Results

Classification accuracy Accuracy across blocks in shown in Figure 2B. Two participants (one in the 1D-ABS condition and the other in the 2D-REL condition) were excluded from further analysis because their overall accuracy was more than three standard deviations below the mean of their condition. A 2×2 ANOVA on overall classification accuracy (collapsed across rule variants) was performed with stimulus (ABS vs. REL), rule (1D vs. 2D) and stimulus $\times$ rule interaction as factors. There was no main effect of stimulus type ($F = (1, 114) = .09$, $p = .76$), but there was a significant main effect of rule type ($F(1, 114) = 25.5$, $p < .001$) and a significant stimulus type $\times$ rule type interaction ($F(1, 114) = 23.6$, $p < .001$). Tukey HSD tests indicated significant pairwise differences in accuracy between conditions, with the exception of 1D-REL and 2D-REL conditions (see Table 1).

Two-sample t-tests were used to assess effects of rule variants on overall accuracy within each condition. Overall accuracy was higher for learning rules on the Fertilizer dimension in the 1D-ABS condition ($t(19.9) = 3.2$, $p < .01$) but the differences within the remaining conditions were not significant.

Modeling classification boundaries The same method from Experiment 1 was used to categorize participants' decision boundaries as 1D or 2D. Logistic regression on the proportion of 2D rules revealed significant effects of stimulus type (REL: Wald $z = 5.0$, $p < .001$), rule type (2D: $z = 8.1$, $p < .001$), and a stimulus $\times$ rule interaction ($z = -3.6$, $p < .001$). The proportion of 2D boundaries are shown in Figure 3B, separated by target rule. The manipulation of feature representation had the predicted effect on the generation of 2D rules, albeit to a lesser extent than seen in Experiment 1. In the 1D-ABS condition the proportion of 2D rules was very low, whereas in the 2D-REL condition the proportion of 2D rules consistent with the target is high. Finally, in the 1D-REL and 2D-ABS conditions, there was a higher proportion of decision boundaries that were of a different form than the target rule.[1]

Summary

The present findings show that, in both perceptual and abstract domains, the efficacy of SDL is limited by biases in the hypothesis generation process. In addition to replicating the gap in performance between 1D and 2D rules found by Markant and Gureckis (2014) in both domains, this gap was eliminated through simple manipulations of the stimuli. When the stimulus representation facilitated the generation of hypotheses consistent with the true rule, self-directed learners were more likely to classify items using a hypothesis of the same form. When hypothesis generation was inconsistent with the target rule, self-directed learners were less successful at learning the category structure despite their ability to control the selection of training data. This impairment is perhaps most striking in the 1D-RECT and 1D-REL conditions, in which the learning of simple 1D rules suffered because participants were more likely to respond according to 2D hypotheses.

This study was motivated by one account of the hypothesis generation process, involving the sampling of salient cues or relations from the environment to form hypotheses about an underlying structure or concept (Trabasso et al., 1968; Wallach, 1962). The present findings suggest that the success of SDL depends on the way that hypothesis generation is shaped by the environment, including how materials guide attention and set the stage for the perception of relevant features or relationships (Goldstone, Landy, & Son, 2010). It is important to note, however, that a number of other hypothesis generation mechanisms have been proposed, including processes based on memory retrieval (Dougherty, Thomas, & Lange, 2010) and local adjustment of existing hypotheses (Bramley, Dayan, & Lagnado, 2015). An important goal of future work is to understand how these generation processes interact in order to guide information search during self-directed learning.

[1] The same analysis was performed on participants' selections during training (rather than test responses) to evaluate whether biased information sampling contributed to these results. However, the proportion of best-fit 2D boundaries indicated that participants' selections strongly supported the correct form of rule in both Exp. 1 (1D-DIAL: .08; 1D-RECT: .02; 2D-DIAL: .93; 2D-RECT: .91) and Exp. 2 (1D-ABS: .06; 1D-REL: .08; 2D-ABS: .86; 2D-REL: .88).

Acknowledgements

The author thanks Katja Munz, Annika Walinski, Jann Wascher, Chantal Wysocki, and Andrii Zakharov for help in translating the experiment and data collection.

References

Ashby, F., & Gott, R. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*(1), 33–53.

Ashby, F., Paul, E., & Maddox, W. (2011). COVIS. In E. M. Pothos & A. J. Willis (Eds.), *Formal approaches in categorization* (pp. 65–87). Cambridge University Press.

Bramley, N. R., Dayan, P., & Lagnado, D. A. (2015). Staying afloat on neurath's boat–heuristics for sequential causal learning. In *Proceedings of the 36th annual conference of the cognitive science society* (pp. 262–267).

Dougherty, M., Thomas, R., & Lange, N. (2010). Toward an integrative theory of hypothesis generation, probability judgment, and hypothesis testing. *Psychology of Learning and Motivation*, *52*, 299–342.

Goldstone, R., Landy, D., & Son, J. (2010). The education of perception. *Topics in Cognitive Science*, *2*(2), 265–284.

Gureckis, T. M., & Markant, D. B. (2012). Self-directed learning: A cognitive and computational perspective. *Perspectives on Psychological Science*, *7*(5), 464-481. doi: 10.1177/1745691612454304

Gureckis, T. M., Martin, J., McDonnell, J. V., Rich, A. S., Markant, D., Coenen, A., … Chan, P. (2015). psiTurk: An open-source framework for conducting replicable behavioral experiments online. *Behavior Research Methods*, 1–14. doi: 10.3758/s13428-015-0642-8

Johnson, T., & Krems, J. (2001). Use of current explanations in multicausal abductive reasoning. *Cognitive Science*, *25*(6), 903–939.

Klahr, D., & Dunbar, K. (1988). Dual space search during scientific reasoning. *Cognitive Science*, *12*, 1–48.

Macmillan, N. A., & Ornstein, A. S. (1998). The mean-integral representation of rectangles. *Perception & Psychophysics*, *60*(2), 250–262.

Markant, D., & 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. doi: 10.1037/a0032108

Mayer, R. (2004). Should there be a three-strikes rule against pure discovery learning? *American Psychologist*, *59*(1), 14–19.

Shepard, R. (1964). Attention and the metric structure of the stimulus space. *Journal of Mathematical Psychology*, *1*(1), 54–87.

Trabasso, T., Bower, G., & Gelman, R. (1968). *Attention in learning: Theory and research*. Wiley.

Tweney, R., Doherty, M., Worner, W., Pliske, D., Mynatt, C., Gross, K., & Arkkelin, D. (1980). Strategies of rule discovery in an inference task. *Quarterly Journal of Experimental Psychology*, *32*(1), 109–123.

Wallach, L. (1962). The complexity of concept-attainment. *The American Journal of Psychology*, *75*(2), 277–283.

Weick, K. E. (1995). *Sensemaking in organizations* (Vol. 3). Sage.