

Relations between Associative and Structured Knowledge in Category-based Induction

Eoin Travers (etravers01@qub.ac.uk)

Jonathan Rolison (j.rolison@qub.ac.uk)

School of Psychology, Queen's University Belfast
Belfast, BT7 1NN UK

Aimée K. Bright (a.bright@qmul.ac.uk)

Psychology Division, School of Biological and Chemical Sciences, Queen Mary University of London
London E1 4NS, UK

Aidan Feeney (a.feeney@qub.ac.uk)

School of Psychology, Queen's University Belfast
Belfast, BT7 1NN UK

Abstract

Theories of category-based human inductive reasoning typically rely on either associative or structured knowledge about relationships between categories. Here, we test a prediction, derived from a hybrid theory that utilizes both kinds of knowledge representation, that participants will experience conflict on a reasoning task in which associative and structured knowledge support different responses.

Participants completed a triad task that tested their ability to generalize a genetic property from a target species to a taxonomically related response. The strength of association between the target and an alternative non-taxonomic (i.e., foil) response was manipulated across trials. Analysis of participants' mouse cursor trajectories revealed that they were initially drawn toward strongly associated foil responses, even when they ultimately chose the correct (taxonomic) option.

Keywords: Category-based induction; Knowledge; Response dynamics; Cognitive conflict;

Introduction

Inductive reasoning is among people's most important cognitive skills, allowing them to draw on prior knowledge to make predictions under uncertainty. Induction is both simple and complex. Induction is simple in the sense that we can easily and automatically associate causes with their effects, generalize properties to a category from a single instance, or from one instance to another, and select actions in complex situations by recognizing commonalities with past experiences. Conversely, inductive reasoning can be complex: how we generalize a property from a given exemplar depends on the nature of the property in question, the circumstances under which the property is observed, and the nature of the relationships between categories.

Theories of category-based induction may be distinguished in a similar manner. Simple, or *associative*, models (Kruschke, 1992; Rescorla & Wagner, 1972; Rogers & McClelland, 2004; Sloman, 1993; Sloutsky & Fisher, 2008) rely on similarity, contiguity, or co-occurrence between instances, and are often modelled using connectionist neural networks. Conversely, structured

relational models take advantage of more sophisticated knowledge about the world, including directional causal relationships, and domain specific rules (Griffiths & Tenenbaum, 2009; Heit, 1998; Kemp & Tenenbaum, 2009; Murphy & Medin, 1985; Osherson, Smith, Wilkie, Lopez, & Shafir, 1990).

Evaluating the strengths and weaknesses of structured and relational models, Bright and Feeney (2015) proposed that induction relies on two forms of knowledge: simple, associative representations that are retrieved easily and automatically, and more complex, structured relational knowledge, including causal relationships, and domain specific intuitive theories, that require cognitive effort. A prediction of this hybrid model, investigated in the present article, is that when associative and structured knowledge come into conflict it often becomes necessary for the reasoner to inhibit an incorrect inference, generated automatically from associative knowledge, in order to reason on the basis of more complicated relational information. To test this prediction, we employed a well-established mouse-tracking paradigm (Freeman, Dale & Farmer, 2011), allowing us to monitor participants' moment-by-moment movements toward responses that are cued by associative and structured knowledge in a forced-choice inductive reasoning task.

Knowledge types in category-based induction

Associative knowledge features most prominently in connectionist, or neural network-based, models of category-based induction (Rogers & McClelland, 2004; Sloman, 1993; Sloutsky & Fisher, 2008). In Sloman's (1993) influential feature-based model, the known features of a category can be represented as an activation vector applied to the input nodes of a feed-forward neural network. Each node is activated when its corresponding feature is possessed by the target category. The network can be presented with the premises of an inductive argument by training it to activate its output node when presented with the features of categories which do have a novel property.

The outcomes of the trained network (i.e., its inferences) can be elicited by probing the input nodes with the features of a novel category. The activation of the output node corresponds to the degree to which the network believes the novel category will have the property in question.

Simple associative architectures are capable of capturing many aspects of human performance (Rolison, Evans, Dennis, & Walsh, 2012). For instance, like human subjects, an associative network can be trained to rate similar categories (which share many features) as more likely to share a novel property, and properties that are present in a diverse range of categories as more likely to be found in a novel category. Adding additional layers to the neural networks allows them to account for further characteristics of human inference, such as sensitivity to different property types, such as “is a” and “has a” properties (Rogers & McClelland, 2004), as well as basic context-dependent inferences (Sloutsky & Fisher, 2008).

Murphy and Medin (1985) argue, however, that similarity-based approaches fail to capture the full flexibility of people's intuitive theories about the relationships between categories in specific domains. In particular, participants have been shown to be sensitive to *property effects* when reasoning inductively, such that the strength of an argument is dependent on the kind of property projected. (Heit & Rubinstein, 1994; Shafto, Coley, & Baldwin, 2007; Shafto, Kemp, Baraff, Coley, & Tenenbaum, 2005). For instance, transmittable properties such as infectious diseases are thought to be shared by animals that are related ecologically, such as predators and prey in a food chain, whereas biological properties such as genes are shared only by animals that are close together in their evolutionary taxonomic tree. Such intuitions, in one domain, are captured by the similarity-coverage model (Osherson et al., 1990), which uses a taxonomic tree to capture intuitions about how properties are shared by related species in the natural world. More recently, structured Bayesian models have been introduced (Griffiths & Tenenbaum, 2009; Kemp & Tenenbaum, 2009; Shafto, Kemp, Bonawitz, Coley, & Tenenbaum, 2008; Tenenbaum, Griffiths, & Kemp, 2006), which are capable of describing flexible human performance in a range of domains. These models require that, for each domain, a specific structure is generated to capture relationships between categories, such as food chains or taxonomic trees, along with a generic probabilistic process by which properties can be transmitted. In the biological domain, such structures include unidirectional causal links connecting prey to predators, or distance from a common ancestor in a biological taxonomy.

Bright and Feeney (2015) argue that neither associative nor structured models are sufficient to account for all of the phenomena observed in category-based induction. They propose a hybrid theory in which both associative and structured knowledge can be used in reasoning. They provide evidence that the two kinds of knowledge can be dissociated. Namely, measures of the strength of the association between two categories predicts participant

ratings of the strength of inductive arguments made under cognitive load, and under time pressure. Conversely, a measure of structured knowledge predicts ratings of argument strength otherwise

Conflict in category-based induction

If both associative and structured knowledge play a role in induction, a natural question is how associative and structured knowledge interact when they support conflicting beliefs. For instance, upon learning that a biological property is true of salmon, does one decide that this is also true of grizzly bears (strongly associated, but no structured means of transmission for biological properties) or of goldfish (weakly associated, but related taxonomically)? Clearly, the decision depends on what kind of knowledge is recruited, with purely associative knowledge in this case leading to a non-normative inference. One possibility is that one or other representation is activated, depending on available time and cognitive capacity, in what Evans (2007) labels a “preemptive conflict resolution” model. Alternatively, both representations may compete, either with associative knowledge being recruited by default, which must be inhibited in order for structured representations to come online (“default interventionist models”), or with both representations activated in parallel (“parallel-competitive models”), leading to a conflict. Bright and Feeney (in prep.) offer evidence that associative and structured knowledge do conflict during category-based induction. In a triad task (Gelman & Markman, 1986), in which participants were asked which of two target species was most likely to share a biological property given that it was found in a third base species, participants were more likely to fail to select a structurally (i.e. taxonomically) related target when the alternative response was strongly associated to the base. Crucially, participants were less able to inhibit the association-driven response under cognitive load, or when lacking in semantic inhibitory control or working memory capacity.

Although the above results provide some support for a proposal that associative and structured knowledge can compete during inductive reasoning, these conclusions are drawn from analysis of participants' responses – the end product of the reasoning process – and thus constitute only an indirect measure of the underlying processes. The mouse-tracking paradigm (Spivey, Grosjean, and Knoblich, 2005; Freeman et al., 2011), on the other hand, provides a powerful tool for measuring these processes as they unfold during cognition. Monitoring the location of the mouse cursor whilst participants are choosing between choice options located on opposite sides of the computer monitor, this method allows us to track the time-course of reasoning that leads to an inference. Mouse-tracking has been used to reveal parallel competition effects on a range of simple cognitive and perceptual tasks (i.e. Freeman, Ambady, Rule, & Johnson, 2008; O’Hora, Dale, Piironen, & Connolly, 2013; Spivey et al., 2005), in which participants are shown to be attracted simultaneously to competing response

options. In more complicated tasks, participants also have been found to exhibit more discrete “changes of mind” tendencies by switching between choice options mid-trial (Dale & Duran, 2011; Freeman, 2013). In the present study, we use this technique to test for conflict between associative and structured knowledge in the triad task. Participants were asked to choose between projecting a biological property from a base species to a correct target species belonging to the same taxonomic group, or to an unrelated foil species. Critically, the strength of the association between the base and the foil species is varied within subjects. If, as suggested by Bright and Feeney (2015), responses cued by associative knowledge must be inhibited in order to reason on the basis of structured relations, we should expect to observe an initial attraction toward the foil that is proportional to the strength of the associative connection between the foil and the target. On the other hand, if participants recruit one or other form of knowledge, we should not expect to find an initial attraction to the foil, regardless of the strength of association between the target and foil.

Method

Stimuli

Participants were presented with a version of the inductive triad task (Gelman & Markman, 1986). On each trial, participants were informed that a particular gene is possessed by a given base species and were asked to decide which of two candidate target species was most likely to possess the same gene (see Figure 1). The correct response was the species belonging to the same taxonomic group as the base (mammals, birds, insects, reptiles, or plants). The foil response belonged to a different taxonomic group than the base and was weakly, moderately, or strongly associated with the base. The strength of association was determined by prior testing. Across 27 experimental trials, nine base species were each presented three times, paired with the same correct response species but a different foil species on each occasion. An additional 27 filler trials were included, in which the property to be generalized was susceptibility to a given disease.

Design and Procedure

Forty four undergraduate students at Queen's University Belfast participated for course credit. Stimuli were custom programmed using the OpenSesame software package (Mathôt, Schreij, & Theeuwes, 2011) and were presented on a computer monitor. The 27 experimental trials were presented in three blocks of nine trials each, interspersed with nine filler trials. Trials were randomly assigned to each block with the constraints that each base species appeared once in each block and each block contained three weakly, three moderately, and three strongly associated foil trials. Trial order within blocks was randomized with the constraint that the same base could not appear twice within three trials.

On each trial, participants were first primed with the kind of property they were to reason about: “gene” for experimental trials or “disease” for fillers. This prime appeared in the center of the monitor for one second. For each experimental trial, participants were then informed that the given gene (i.e. “Gene r3P”) is found in the bodies of one of the two target species, which appeared as labeled images in the top left and right corners of the screen (Figure 1). The two species were randomly assigned to the left and right positions on each trial and appeared for 1.6 seconds each. The targets then remained visible and participants were asked which species they believed was most likely to possess the gene, given that it was possessed by another species. Participants were then instructed to click a “START” button located in the bottom center of the monitor, after which a fixation cross appeared for 1.5 seconds, which was then replaced by a labeled image of the base species (Figure 1). At this point, the mouse cursor was reset to the center of the start button and participants were given five seconds to respond by selecting one of the two target species labels with their mouse cursor. Participants were given five seconds to respond following presentation of the base category. Additionally, in line with previous mouse-tracking research, on trials in which participants did not move the cursor away from the start button within 1.5 seconds of the onset of the base, they were shown a message reminding them that they were under time pressure. This was done to encourage participants to make their decision while the mouse cursor was in motion.

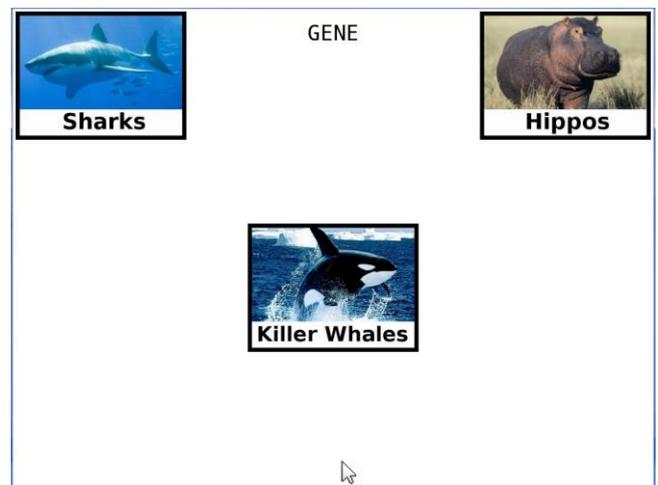


Figure 1: Screenshot of the experiment screen following onset of the base species (“Killer Whales”).

Analysis

Mouse trajectories were normalized to a standard coordinate system, with all trials beginning at point [0, 0], and ending at point [1, 1.5] in the top right corner. Trajectories in which the chosen response was on the left were reflected through the y-axis. For each trial, we calculated the time from target onset to a response (*response time*), the time from target onset to the beginning of the mouse movement

(*initiation time*), the deviation of mouse trajectory away from a straight line to the response, measured in the standard co-ordinate system (*maximum deviation*), and the frequency of changes of mouse trajectory direction on the x-axis (*x-flips*).

Our analysis was restricted to trials on which the taxonomically-related species was chosen. Thus, our data set was unbalanced. Therefore, we conducted random effects linear regression modeling on our data. This analysis accounted for the clustering in our data by allowing for random intercepts at the subject and base species level (Baayen, Davidson & Bates, 2008). Main effects were assessed on the basis of the -2 log-likelihood model fit improvement, tested using the chi-square statistic. Main effects were followed-up with Tukey pairwise comparisons between each group, with p values calculated using the normal approximation.

Log transformations were used for analyses of response time, initiation time, and maximum deviation due to violations of normality. A Poisson regression model was used for the count of x-flips. A logistic random effects model was used for the analysis of choices in each condition.

Results

Results did not differ appreciably between the stimuli blocks, and so data were collapsed across blocks for analysis. Participants selected the correct taxonomically-related response on 81% of trials when the foil was weakly associated with the base, 61% of trials when moderately associated, and 57% when strongly associated. This suggests that participants were influenced by the associative strength of the foil option, such that stronger associations competed with structural knowledge. A logistic mixed effects model indicated a main effect of foil strength ($\Delta AIC = -86.4$, $-2LL \chi^2(2) = 90.4$, $p < .001$). Pairwise comparisons revealed significant differences between the weak and moderate ($t = 7.310$, $p < .001$), and weak and strong foils ($t = 8.489$, $p < .001$), but not between moderate and strong foils ($t = 1.428$, $p > .3$).

Table 1: Descriptive statistics by condition and statistical tests assessing the main effects of condition.

Foil association	RT	IT	MD	X-flips	CoM
Weak	1517	572	0.37	0.34	21.4%
Moderate	1493	583	0.35	0.32	19.3%
Strong	1513	567	0.45	0.41	27.8%
p	.883	.558	.042	.065	.011

Note: RT = Response Time (msec); IT = Initiation Time (msec); MD = Maximum Deviation; CoM = Percentage of trials classed as changes of mind.

Inspection of the mouse cursor data revealed two kinds of mouse trajectory: movements directly toward the taxonomic option, and initial movements toward the foil option that

changed direction toward the taxonomic option mid-trial. Analysis of the distribution of the maximum deviation statistic (Figure 2) revealed two normally-distributed subpopulations of responses, one centered on a deviation close to 0 (measured in the standard co-ordinate system) that corresponded to movements directly toward the taxonomic option, and a second centered around 1.4. The bimodality of this distribution was confirmed by calculating its bimodality coefficient (Freeman & Dale, 2012), yielding a value of .636, well above the threshold of .555 usually interpreted as indicating bimodality. We therefore fitted a two-sample finite mixture model to these maximum deviation values, in order to classify trajectories as either “*changes of mind*” (maximum deviation $> .827$), or “*direct to taxonomic option*”. The two kinds of mouse trajectories are shown in Figure 3.

Condition means and statistical tests for the measures described above are shown in table 1.

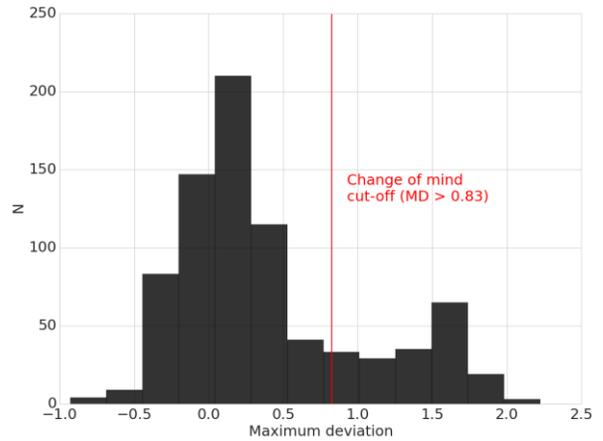


Figure 2: Distribution of the maximum deviation from a straight line for all correct responses.

Significant main effects were observed for maximum deviation ($\Delta AIC = -2.33$, $\chi^2(2) = 6.33$, $p = .042$) and for changes of mind ($\Delta AIC = -5.0$, $\chi^2(2) = 9.012$, $p = .011$), with a marginally significant main effect for x-flips ($\Delta AIC = -1.471$, $\chi^2(2) = 5.471$, $p = .065$). Pairwise comparisons showed significant differences between weakly and strongly associated foils, with greater signs of conflict when strongly associated, for maximum deviation ($t = 2.44$, $p = .038$), and for changes of mind ($t = 2.69$, $p = .023$), and a marginal difference for x-flips ($t = 2.23$, $p = .067$). There was an additional significant difference between moderately and strongly associated foils for changes of mind only ($t = 2.673$, $p = .020$), with more changes of mind for strongly associated foils.

Discussion

Bright and Feeney (2015) showed that both associative and structured knowledge can serve as the basis for inductive reasoning. Bright and Feeney (in prep.) provide evidence

that both kinds of knowledge can conflict during reasoning. Here, we found that participants generalized biological properties from a base species to a target species from the same taxonomic group, rather than to a foil species, on the majority of trials when the target was weakly associated with the foil. The stronger the association between the base and foil species, however, the more likely participants were to generalize the property to the foil instead. This was despite the base and foil species belonging to different taxonomic groups. Analysis of the mouse cursor trajectories revealed that many participants were initially drawn to strongly-associated foil responses, even when they ultimately selected the option that corresponded to structured knowledge.

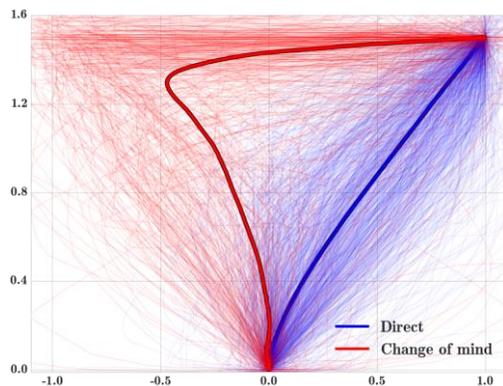


Figure 3: Direct and “change of mind” mouse trajectories. Averaged trajectories are shown in bold.

Our findings provide evidence against theories of inductive reasoning that describe either an associative *or* structured knowledge account. Rather, our present findings suggest that both forms of knowledge are engaged during reasoning, and that both can influence a single decision. At the outset, we raised the question of how precisely the two forms of knowledge interact. Although our data do not provide a definitive answer, the patterns in the mouse trajectory data shown in Figure 3 suggest that the majority of trajectories went directly to the taxonomically related target without any evidence of conflict. To the extent that participants detect conflict between the choice options, such trajectories are consistent with pre-emptive conflict resolution (Evans, 2007). “Change of mind” responses, on the other hand, indicate online resolution of conflict. Further research will be required to determine (a) why conflict is sometimes resolved pre-emptively and sometimes online, and (b) whether, when conflict is resolved online, both types of knowledge are activated in parallel or in sequence.

Bayesian models of inductive reasoning claim that people represent structured relations between categories when reasoning, appropriate to the domain in question. By placing these structured representations in conflict with associative knowledge, we have shown that in order to reason in a way consistent with a Bayesian account, it is sometimes necessary for people to inhibit the associative representations which come to mind more easily. This is

consistent with Bright and Feeney’s (2015) demonstration that reasoning is consistent with the structured Bayesian model when people have adequate time and mental resources, but is driven by simpler associative knowledge otherwise.

Our results may be challenged by theorists who favor purely associative models of induction (Rogers & McClelland, 2004; Sloman, 1993; Sloutsky & Fisher, 2008). Neural network models have been shown to capture some context sensitivity effects by means of input nodes encoding contextual features (Sloutsky & Fisher, 2008). However, the “change of mind” movements which characterized our mouse trajectory data are difficult to explain within this framework. Simple feedforward neural networks of the type used in models of induction are static, in that they are probed once, and produce a single output pattern, providing no mechanism for reversals during a trial. More complex recurrent networks, on the other hand, with input and output changing over time, can capture the evolution of choices. However, extensive mouse-tracking research has demonstrated that conflict in such networks is continuous, with participants partially drawn toward two competing responses, selecting one response but curving toward the alternative (i.e. Freeman & Ambady, 2011; Freeman et al., 2008; Spivey et al., 2005). Discrete reversals have been demonstrated on tasks thought to involve the sequential operation of two processes (Barca & Pezzulo, 2015; Dale & Duran, 2011; Freeman, 2013; Freeman & Dale, 2012; Hindy & Spivey, 2008; Tomlinson, Bailey, & Bott, 2013), suggesting that our results reflect the initial activation of associative knowledge and the subsequent retrieval of structured knowledge.

To conclude, we believe that neither associative nor structured models alone are capable of describing the processes underlying human inductive reasoning. Instead, people draw upon two forms of knowledge representation, one associative, and easily accessed, and one structured, and requiring mental effort to utilize. Making use of structured knowledge appears to require the inhibition of associative information, and as a result, participants were more likely to select a foil response if it was strongly associated with the base. Uniquely, our mouse trajectory results reflect the online inhibition of association-driven responses, necessary to reason according to structured knowledge.

Acknowledgments

Eoin Travers is funded by a PhD studentship from the Department for Employment and Learning in Northern Ireland.

References

- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59, 390–412.

- Barca, L., & Pezzulo, G. (2015). Tracking second thoughts: Continuous and discrete revision processes during visual lexical decision. *PLoS ONE*, *10*, e0116193.
- Bright, A. K., & Feeney, A. (2015). The engine of thought is a hybrid: Roles of associative and structured knowledge in reasoning. *Journal of Experimental Psychology: General*.
- Bright, A. K., & Feeney, A. (in prep.) Conflict resolution in inductive reasoning: The roles of working memory and inhibitory control.
- Dale, R., & Duran, N. D. (2011). The cognitive dynamics of negated sentence verification. *Cognitive Science*, *35*, 983–996.
- Evans, J. S. B. T. (2007). On the resolution of conflict in dual process theories of reasoning. *Thinking & Reasoning*, *13*, 321–339.
- Freeman, J. B. (2014). Abrupt category shifts during real-time person perception. *Psychonomic bulletin & review*, *21*(1), 85–92.
- Freeman, J. B., & Ambady, N. (2011). When two become one: Temporally dynamic integration of the face and voice. *Journal of Experimental Social Psychology*, *47*, 259–263.
- Freeman, J. B., & Dale, R. (2012). Assessing bimodality to detect the presence of a dual cognitive process. *Behavior Research Methods*, *45*, 83–97.
- Freeman, J. B., Ambady, N., Rule, N. O., & Johnson, K. L. (2008). Will a category cue attract you? Motor output reveals dynamic competition across person construal. *Journal of Experimental Psychology: General*, *137*, 673–690.
- Gelman, S. A., & Markman, E. M. (1986). Categories and induction in young children. *Cognition*, *23*, 183–209.
- Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review*, *116*, 661–716.
- Heit, E. (1998). A bayesian analysis of some forms of inductive reasoning. *Rational models of cognition*, 248–274.
- Heit, E., & Rubinstein, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 411–422.
- Hindy, N. C., & Spivey, M. J. (2008). Motor dynamics of task switching. In *Proceedings of the 30th annual conference of the cognitive science society, washington, DC, july 23–26, 2008* (pp. 2474–2479).
- Kemp, C., & Tenenbaum, J. B. (2009). Structured statistical models of inductive reasoning. *Psychological review*, *116*, 20.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological review*, *99*, 22–44.
- Mathôt, S., Schreij, D., & Theeuwes, J. (2011). OpenSesame: An open-source, graphical experiment builder for the social sciences. *Behavior Research Methods*, *44*, 314–324.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*, 289–316.
- Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological review*, *97*, 185.
- O’Hora, D., Dale, R., Piironen, P. T., & Connolly, F. (2013). Local dynamics in decision making: The evolution of preference within and across decisions. *Scientific Reports*, 2210..
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning II: Current research and theory* (Vol. 2, pp. 64–99). Appleton-Century-Crofts.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. MIT Press.
- Rolison, J. J., Evans, J. S. B. T., Dennis, I., & Walsh, C. R. (2012). Dual-processes in learning and judgment: Evidence from the multiple cue probability learning paradigm. *Organizational Behavior and Human Decision Processes*, *118*, 189–202.
- Shafto, P., Coley, J. D., & Baldwin, D. (2007). Effects of time pressure on context-sensitive property induction. *Psychonomic bulletin & review*, *14*, 890–894.
- Shafto, P., Kemp, C., Baraff, E., Coley, J. D., & Tenenbaum, J. B. (2005). Context-sensitive induction. In *Proceedings of the 27th annual conference of the cognitive science society* (pp. 2003–2008). Mahwah, NJ: Erlbaum.
- Shafto, P., Kemp, C., Bonawitz, E. B., Coley, J. D., & Tenenbaum, J. B. (2008). Inductive reasoning about causally transmitted properties. *Cognition*, *109*, 175–192.
- Sloman, S. A. (1993). Feature-based induction. *Cognitive psychology*, *25*, 231–280.
- Sloutsky, V. M., & Fisher, A. V. (2008). Attentional learning and flexible induction: How mundane mechanisms give rise to smart behaviors. *Child Development*, *79*, 639–651.
- Spivey, M. J., Grosjean, M., & Knoblich, G. (2005). Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences*, *102*, 10393–10398.
- Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, *10*, 309–318.
- Tomlinson, J. M., Bailey, T. M., & Bott, L. (2013). Possibly all of that and then some: Scalar implicatures are understood in two steps. *Journal of Memory and Language*, *69*, 18–35.