Recovery from Misleading Featural and Relational Information

Maureen Gray

SURF 2015 Summary Report

Introduction

The human experience would be immensely more complicated, perhaps even impossible to navigate, without the help of categorization. Categorization greatly reduces the complexity of the perceptual world by fitting each individual object we see into categories of like objects. This allows the seamless inference of properties so that we do not have to wonder whether that small, fluffy, long-eared animal we see poses a threat to us — we know immediately that the animal can be classified as a rabbit, and infer from our experiences with other rabbits that it is harmless. My research explores some of the differences between two particular kinds of categories: featural and relational. A featural category is one whose membership is determined by common features of category members. *Rabbit* is a good example of a featural category. We use the features (long ears, fur, etc.) to categorize things as rabbits. A relational category, in contrast, is a category whose membership is determined by a common relational structure rather than shared properties (Gentner and Kurtz, 2005). For example, the members of the relational category *barrier* include walls and fences, along with more abstract things like race, or socioeconomic status. The category is relational because it is defined by the relationship between two things — i.e., that something is being blocked from something else.

The motivation for my experiment is to elucidate some of the ways in which people acquire both kinds of categories. This is very important to understand, because relational concepts play an essential role in human cognition. Reasoning, problem-solving, and the productivity of language are all possible as a result of the uniquely human capacity for relational thought (Gentner, 1983; Hummel and Holyoak 1997; Penn et al., 2008). In spite of their importance, surprisingly little is known about how relational categories are structured, how they are acquired, and how they relate to featural categories (Kittur et al., 2006).

The question I am seeking to answer with my research is how the type of category (relational vs. featural) affects a person's willingness to alter that category in the face of contradictory evidence. Specifically, I am running a category learning experiment that investigates the extent to which relevant featural information can supplant once-relevant relational information, and vice versa. In the experiment, subjects learn to categorize simple stimuli into one of two categories: either Category A or Category B. In Phase 1 of the experiment, there are two cues that are always predictive of category membership: one salient, and ultimately incorrect, and another, more subtle and ultimately correct. After the subject has made enough categorizations correctly and reached criterion (seven out of eight categorizations correct for two consecutive blocks), the subject unknowingly enters Phase 2 of the experiment. At this point, the exemplars change so that the salient cue is no longer predictive of category membership, and only the subtle cue remains. The measure of interest is how well the learner recovers from the apparent changes in the rules of the game. Specifically, we measured how

many categorizations the subject made correctly out of a block containing eight trials. The primary manipulations are what kinds of information appear to be predictive at first (relational or featural) and what kind of information is genuinely relevant in later trials (again, relational or featural). These manipulations are also compared with a baseline condition in which there are no "false" cues at all, and the subjects are presented with only the genuinely relevant information from the start.

Accomplishments this summer

This summer I analyzed and interpreted behavioral data collected during the Spring semester and planned further experiments based on these analyses to be executed in the Fall of 2015. Data analysis included constructing graphs that allowed visual comparison of all six conditions and performing Bayesian Estimation to evaluate the credibility of the differences we observed in the data. Also, I have begun preparing to present my results at the Psychonomics conference in November of 2015. At this conference, I will present my research at a poster session.

Results and Discussion

The behavioral data we analyzed was subjects' categorization accuracy across a block of 8 trials, and the blocks of interest were those in Phase 2 of the experiment: when the easy, salient cue is no longer reliable and subjects must use the subtle cue to categorize correctly. Specifically, we looked at how long it took subjects to reach criterion in Phase 2 of the experiment and whether or not subjects' ability to do this was affected by whether the salient cue was relational or featural. Again, criterion was defined as an accuracy of 87.5% (7/8 trials) correct for two consecutive blocks.

Analysis of the data revealed some interesting trends. Subjects that learned salient relational cues in Phase 1 of the experiment recovered to criterion in Phase 2, regardless of whether the genuine cue was relational or featural. By contrast, subjects that learned salient featural cues in Phase 1 of the experiment did not recover to criterion in Phase 2, regardless of whether the genuine cue was relational or featural. Bayesian Estimation that compared these two groups of subjects showed that the difference between the mean accuracy was credible (highest density interval does not include 0 with 20,000 burn-in steps).

So, it appears that a featural cue in Phase 1 of the experiment negatively affects learning in Phase 2, but a relational cue in Phase 1 does not have this same detrimental effect on learning in Phase 2. Before interpreting this finding further, it is important to note that features and relations are processed in entirely different ways. Features can be perceived and processed automatically, without attention. Relations, on the other hand, require attention to be perceived.

Vendetti, Wu, and Holyoak (2014) showed that it is possible to induce a relational mindset that influences processing of relational stimuli later on. We have proposed that featural learning in Phase 1 puts the subject in a featural mindset: that is, one where stimuli are processed holistically. This mindset is a contrast to Vendetti et al.'s relational mindset, where separate pieces of the stimulus *must* be processed in order to calculate the relation between them. In the holistic featural mindset, on the other hand, it is unnecessary to process each piece of the stimulus separately — it can be evaluated and classified based on just a single feature. Processing stimuli in any way could lead subjects to form a sort of processing bias that alters how they process future stimuli. In the case of this experiment, the results suggest that the processing mindset induced by dealing with featural information is less flexible and more difficult to overcome.

Future Directions

These results opened the door to many more experimental questions, but before tackling them, it is of utmost importance to ensure that this effect is indeed real. Therefore, I intend to replicate this experiment in an analogous paradigm but with different, more realistic stimuli. Thanks to the work I was able to do this summer, this direct replication is ready to run. Both the original experiment and its direct replication work with features that take different metric values (like absolute size, or darkness). To see if these results obtain with different kinds of features, my advisor and I have planned a third experiment in an analogous paradigm using categorical features (e.g., shape).

Finally, I will be presenting these results at the Psychonomics Conference in November of this year. Once all of these follow up experiments have been executed and analyzed, my advisor and I intend to write up the results and submit them to a journal for publication.