Memory for Category Information Is Idealized Through Contrast With Competing Options

Tyler Davis and Bradley C. Love
University of Texas at Austin

Abstract

We suggest that human category formation relies on contrastive learning mechanisms that seek to reduce prediction error. In keeping with this view, manipulating category contrast leads to systematic distortions in people's memory for category information. Simply by changing the basis of comparison (i.e., the available response options), we can systematically distort people's perceptions of novel, energy-source, and political categories. Our proposal explains perceived variations in category members' typicality, including cases in which average items are judged as highly typical and cases in which extreme or ideal members are judged as highly typical of the category. Although straightforward, our account spans findings from studies in goal-derived, cross-cultural, and object-based categorization and suggests ways in which society's perception of key issues is distorted by political discourse.

Categories help us make sense of novel situations and objects by enabling us to generalize from past experiences. Categories, such as dogs and cats or merlot and shiraz, guide expectations about new examples. For example, dogs are generally larger than cats, so one will have different expectations about an animal's size depending on whether it is classified as a dog or a cat. Accounts of human conceptual structure aim to explain how various forces influence category formation. These forces can occur within categories (e.g., how similar dogs are to one another) or between categories (e.g., how similar dogs are to cats).

Members of categories vary in typicality, and these differences can provide clues to the underlying structure of categories. Features of category members that are perceived as most typical include a willingness in people to project properties from that member to other category members (Osherson, Smith, Wilkie, López, & Shafir, 1990), ease of processing (Posner & Keele, 1968; Rosch, Simpson, & Miller, 1976), and perceived attractiveness (Winkielman, Halberstadt, Fazendeiro, & Catty, 2006).

A number of theories have been developed to account for differences in typicality. One class of theories focuses on how category members are distributed and how the resulting similarity relations among items shape perceptions of typicality (e.g., family resemblance—Rosch & Mervis, 1975; geometric similarity—Nosofsky, 1986). A second class of theories focuses

on cultural influences (Bang, Medin, & Atran, 2007; Medin & Atran, 2004). Theories that are structured around distributional or statistical measures of similarity are successful at explaining cases in which the most typical members are those that are closest to the central tendency (i.e., average) of their category. However, these theories fail to explain results from studies on cultural and goal-derived categories in which members that are ideal are judged to be most typical (Atran, 1999; Barsalou, 1985; Lynch, Coley, & Medin, 2000). For example, in the category "diet foods," items that minimize calories, such as celery, are rated as most typical, even though these extreme items are not average for their category.

We suggest that a common learning mechanism may underlie both cases, when highly typical items are statistically average and when they are ideal. Both cases are consistent with the proposal that human category acquisition involves the use of contrastive learning mechanisms that seek to reduce prediction error. This view holds that perceived category averages (i.e., centers of typicality) are not determined solely by the statistical properties of the category but can shift depending on situational demands. We predict that when categories contrast

Corresponding Author:
Tyler Davis, Department of Psychology, 1 University Station A8000, University of Texas, Austin, TX 78712
E-mail: thdavis@mail.utexas.edu

along a psychological dimension, the perceived center of the category will shift to reduce the chances of making a categorization error. Such shifts can lead to idealization. For example, under our proposal, people's perceived average of diet foods moves away from non–diet foods, to avoid the possible selection of non–diet foods when dieting.

To illustrate, consider Categories A and B, depicted in Figure 1. Items in these categories vary on two dimensions, horizontal and vertical, but the categories contrast with one another only on the horizontal dimension. Our account predicts that each category's perceived average will simultaneously be pulled toward its members (consistent with family resemblance) and repelled by members of contrasting categories (consistent with idealization). When there is no contrast between categories, such as on the vertical (noncontrast) dimension, the sum of the forces pulling the representations in a particular direction balances, and the perceived average remains at the true average. However, when the categories are differentiated along a dimension, such as with the horizontal (contrast) dimension, the sum of the forces pulling the perceived averages away from contrasting categories becomes greater, leading to idealization. For example, in Figure 1, both extreme Category A members and all of Category B's members exert forces on the perceived Category A average to move further from Category B than the true central tendency. The additional between-category forces upset the balance within Category A, resulting in a perceived average that is idealized with respect to the contrasting category. The computation of these ideas is presented in Sakamoto, Jones, and Love (2008) and the Model Appendix in the Supplemental Material available on-line. This analysis of representational change is in the spirit of other work that explains processing in terms of force dynamics (e.g., Spencer, Perone, & Johnson, 2009).

To provide a real-world corollary, imagine that Category A is a diet food and Category B is a non–diet food. The horizontal dimension on which the categories contrast is caloric content, and the vertical dimension represents a property on which the categories match (e.g., the manufacturer). The forces from within the diet food category balance, and in isolation this would lead to the perceived and true category averages matching. However, non–diet food category members repel the diet food category along the calorie dimension. Combining both the within- and between-category forces results in a perceived average for diet foods that is lower in calories than the true central tendency.
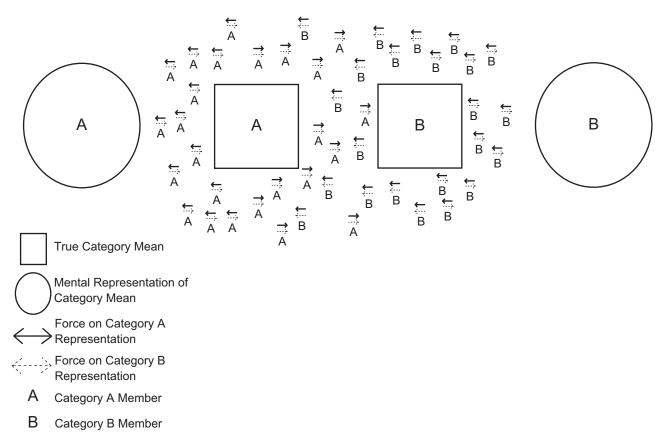

Fig. 1. An illustration of how idealization arises from error-driven learning models. The small *A*s and *B*s represent stimuli presented during learning, which vary in their typicality to the category mean. Here, they vary on the horizontal and vertical dimensions. Our account predicts that each category's perceived average will simultaneously be pulled toward its members (consistent with family resemblance) and repelled by members of contrasting categories (consistent with idealization).

In keeping with this view, we find that manipulating category contrast leads to systematic distortions in people's memory for category information. Simply by changing the basis of comparison, we can systematically distort people's perceptions of novel, energy source, and political categories. Our learning account explains variations in category members' typicality, including cases in which average and ideal category members are judged as highly typical. This account spans findings from studies in goal-derived, cross-cultural, and object-based categorization and suggests ways in which society's perception of key issues can be distorted by political discourse. According to our view, high-level influences on conceptual structure, such as culture, shape category representations by modulating relevant contrasts (Love & Gureckis, 2005).

Overview of Experiments

In the current studies, subjects learned to assign items to four categories by trial and error. Each category can be thought of as occupying a quadrant in a two-dimensional space (see Fig. 2). Each category matches one category and mismatches another along a vertical and horizontal dimension. On category-learning trials, subjects were presented with a stimulus, chose a possible category from the available response options, and received feedback indicating the actual category membership of the stimulus. In Experiment 1, the stimuli were political categories (i.e., hypothetical supporters of presidential primary candidates) or energy sources. Experiment 2 used novel perceptual stimuli.

Contrast was manipulated across conditions in both experiments by varying the response options available on category-learning trials. All other aspects of the task were the same across subjects. In all conditions in Experiment 1 and in Experiment 2's unidimensional condition, the stimuli were contrasted on a single dimension. For example, subjects were offered, on each trial, a choice of two categories that mismatched along a single (horizontal or vertical) dimension. For a subject, the key dimension that discriminated between category-response options, referred to as the *contrast* dimension, was held constant
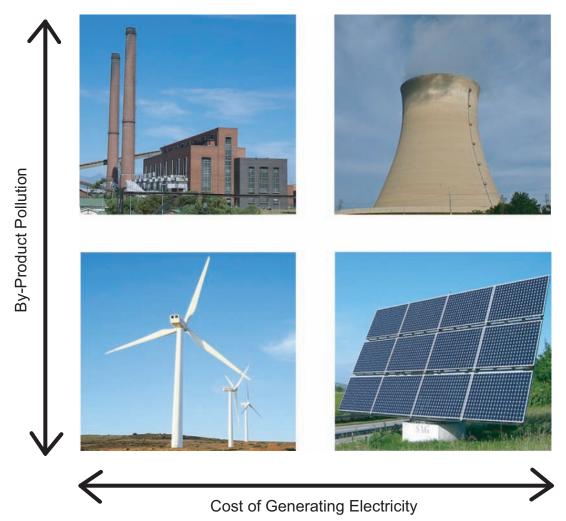

Fig. 2. An illustration of the category structure for the energy-source condition in Experiment 1. The four photographs depict (clockwise from the top left) a coal plant, a nuclear plant, a solar panel, and a windmill. The four categories vary on the dimensions of by-product pollution and cost of generating electricity.

across trials. For example, for the energy-source categories with cost as the contrast dimension (see Figs. 2 and 3), the available choices on each learning trial would be coal versus nuclear or wind versus solar. In this example, the idealization prediction is that coal and wind would appear cheaper than they actually are and nuclear and solar would appear more expensive. Experiment 2 also included conditions in which both stimulus dimensions were contrasted.

For Experiment 1 and 2's unidimensional condition, our model predicts idealization only along the contrast dimension, but not on the noncontrast dimension. In conditions in Experiment 2 that contrast categories on both stimulus dimensions, we therefore predict idealization along both dimensions. These predictions are the same for real-world categories for which subjects have prior expectations and completely novel categories (Experiment 2).

Whether idealization occurred along a dimension was assessed using two measures during category learning—accuracy and response times (RTs)—and two measures postlearning—reconstructions (from memory) of category averages and typicality ratings (Experiment 2 only). Idealization was consistent with higher accuracy, faster RTs, and higher typicality ratings for items that were extreme along a contrast dimension and stimulus reconstructions were idealized in comparison with the true category average.

Method
Subjects

In Experiment 1, 97 undergraduates were randomly assigned either to the energy-source condition ($n = 50$) or to the political-categories condition ($n = 47$). In Experiment 2, 188 undergraduates were randomly assigned to the unidimensional ($n = 95$), free ($n = 45$), or mixed ($n = 48$) conditions.

Materials

In both Experiments 1 and 2, subjects completed a category-learning task in which they learned four categories by trial and error. In Experiment 1, the categories were either energy sources (as depicted in Fig. 2), which varied in terms of their percentile cost of generating electricity and amount of by-product pollution, or political categories, which consisted of hypothetical supporters of John McCain, Barack Obama, Hillary Clinton, or Mike Huckabee and varied in terms of their percentile age and income (for an example of a category member, see Fig. 3). In Experiment 2, these categories were made up of rectangles that varied along two dimensions, height and the position of a vertical line segment along the lower edge, and the categories were simply labeled A, B, C, and D.

Procedure

Subjects in both experiments completed a category-learning task followed by a reconstruction task. In Experiment 2, they also completed a typicality-rating task. Before the category-learning task, subjects were informed that they would learn about four categories of items and that perfect accuracy might not be attainable. Subjects were trained on the category-learning task in blocks of 40 trials until they achieved 80% accuracy within a block or completed five total blocks. Each category served as the target category on 10 trials per block. During a trial, a stimulus from the target category was displayed, along
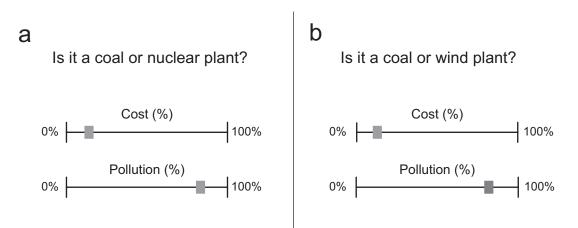

Fig. 3. Illustrations of how category contrast was manipulated on the learning trials in the energy-source condition in Experiment 1. On each trial, subjects were offered a choice of two categories, and their task was to indicate which category corresponded to indicated values on two dimensions. The choices mismatched along a single dimension, the *contrast* dimension, and matched along the second dimension. For a given subject, the contrast dimension was held constant. Subjects responded by hitting the key associated with one of the response options and were then shown the correct (target) category label. In (a), percentile cost of generating electricity is the contrast dimension, and amount of by-product pollution is the noncontrast dimension; that is, coal plants and nuclear plants differ in cost but not pollution. In (b), pollution is the contrast dimension, and cost is the noncontrast dimension; that is, coal plants and wind plants differ in pollution but not cost. In both examples, the correct choice is "coal plant."

with response options that included the target category. Subjects responded by choosing one of the response options and were then shown the correct (target) category label.

On each trial of the category-learning task, subjects were shown a stimulus from one of the categories and were asked to categorize it based on the available response options. That is, in Experiment 1, on each trial, subjects were shown a bar graph depicting an energy source in terms of its cost and pollution (expressed as percentiles) or depicting a candidate in terms of the age and income of his or her supporters. In Experiment 2, on each trial, subjects saw a rectangle of varying dimensions.

The values for stimulus dimensions were sampled on each trial from one of four partially overlapping and uncorrelated Gaussian category distributions. Extreme samples that were more than 2.0 standard deviations from the mean on either dimension were discarded, and the stimulus was resampled. Each category was 2.5 standard deviations from its neighboring category on each dimension. The standard deviation of each dimension was 12 percentiles in Experiment 1 or 2.4 mm in Experiment 2.

The key difference across conditions was the response options made available to subjects during the category-learning task. In the energy-source condition in Experiment 1, the response options available on each trial were determined based on whether subjects were in a condition in which the pollution dimension was the contrast dimension or the cost dimension was the contrast dimension. For example, for conditions in which the pollution dimension was the contrast dimension, subjects would have to choose, on each trial, from categories whose distributions differed only on the pollution dimension. On a trial in which the stimulus came from the coal plant category, subjects in this condition would be asked to choose whether the stimulus was a coal or wind plant. The same logic was used in the political-categories condition and for constructing the unidimensional condition in Experiment 2.

Experiment 2 also included two conditions in which both dimensions were contrast dimensions. Although the stimuli used in these conditions were novel perceptual figures, we describe the conditions using the energy-source stimuli for consistency with the figures and the above description. In the mixed condition, trials alternated within a subject between those in which the response options contrasted categories that differed on the pollution dimension and trials in which the response options contrasted categories that differed on the cost dimension. For example, for coal plants, half of the trials would ask subjects to choose between coal and wind (pollution was the contrast dimension), and half would ask subjects to choose between coal and nuclear (cost was the contrast dimension). In the free condition, subjects could choose from all categories (i.e., solar, wind, coal, and nuclear) on every trial.

After the category-learning task, subjects completed a reconstruction task in which they were instructed to adjust each dimension of a stimulus to match their memory for a given category average. For example, subjects would be asked to adjust the value of a coal plant on the cost and pollution dimensions until it appeared average for the category of coal plants. The reconstruction phase consisted of 12 trials that were organized as three blocks in which each category was queried once. Finally, subjects in Experiment 2 also completed a typicality-rating task in which subjects rated stimuli for how typical they were of a given category. The typicality phase consisted of two blocks of 40 trials (10 stimuli from each category).

Analysis

For purposes of the analyses, each stimulus's dimension values were mean-centered, so that a stimulus located at the category average had zero-valued coordinates. Positive values indicated ideal stimuli, whereas negative values indicated stimuli located away from the mean and toward the contrasting category.

Accuracy, RTs (standardized within subjects), and typicality data were examined as a function of stimulus position within a category.[1] Items were grouped within subjects into five different bins organized from nearest to the contrasting categories to the most ideal (in *z* scores): less than –1, –1 to –0.5, –0.5 to 0.5 (region of true average), 0.5 to 1, and greater than 1. For unidimensional conditions, within-subject means were calculated for each bin on the contrast dimension by averaging over the noncontrast dimension, and likewise for each bin on the noncontrast dimension by averaging over the contrast dimension. For the free and mixed conditions in Experiment 2, means were calculated by taking the average at each bin over both dimensions. These groupings were entered into linear trend analyses to test whether RT, typicality, and accuracy increase as items become further from opposing categories (i.e., ideal) along contrast dimensions.

Results

Overall performance levels in both experiments were consistent with our expectations (see Table 1). Conditions that had equivalent complexity (e.g., equal response options) were similar in terms of how quickly subjects reached the 80% learning criterion and overall accuracy in the first and final learning blocks. Experiment 2's more complex free and mixed conditions were more difficult (e.g., lower accuracy and slower learning), but a majority of subjects exceeded chance performance, and thus we were able to continue to evaluate our research hypothesis.[2]

In both experiments, all measures converged on the same conclusion: When a dimension is contrasted during categorization, items that are ideal along this dimension are associated with greater accuracy, faster RTs, and higher typicality ratings. Items reconstructed from memory are also more idealized along contrasted dimensions. These patterns held in both experiments, regardless of whether stimuli were novel or from real-world categories.

Table 1. Performance on the Category-Learning Task in Experiments 1 and 2

Experiment and condition	*n*	Proportion of subjects exceeding chance	Number of blocks to criterion	Mean accuracy: first block	Mean accuracy: final block
Experiment 1					
Energy source	50	.90	1.91	.74	.85
Political categories	47	.94	2.42	.66	.82
Experiment 2					
Unidimensional	95	.81	2.41	.66	.85
Free	45	.80	3.75	.35	.58
Mixed	48	.77	3.50	.60	.79

Note: Number of blocks to criterion is the average number of blocks it took for subjects to reach the criterion of .80 correct (subjects who did not reach the criterion were excluded). Mean accuracy was calculated excluding subjects who failed to achieve better-than-chance performance within five learning blocks. In the unidimensional condition, stimuli were contrasted on a single dimension. In the mixed condition, the dimension on which stimuli contrasted alternated between trials. In the free condition, subjects could choose from all four categories on every trial.

Reconstruction results for Experiments 1 and 2 are shown in Figure 4. In Experiment 1, items were significantly idealized on the contrast dimension in comparison with the true category average, ts $\geq$ 4.01, p_{rep}s $\geq$ .99.[3] The opposite was true on the noncontrast dimension, where distortion tended to be toward contrasting categories, ts $\geq$ 2.17, p_{rep}s $\geq$ .90. The same pattern held and was significant in Experiment 2's unidimensional condition. Experiment 2's free and mixed conditions, which contrasted both dimensions, showed significant idealization on both dimensions, ts $\geq$ 4.61. p_{rep}s $\geq$ .99.

Accuracy (see Fig. 5), RTs, and typicality measures (see Fig. S1 in the Supplemental Material available on-line) also revealed evidence of idealization. In Experiment 1 and in Experiment 2's unidimensional condition, subjects were more accurate and showed faster RTs as items became more ideal along the contrast dimension, Fs $\geq$ 6.34, p_{rep}s $\geq$ .94. Subject's accuracy and RTs were less affected by the noncontrast dimension, Fs $\leq$ 4.88, p_{rep}s $\leq$ .91, except for Experiment 2's energy-source condition, in which subjects became less accurate as items became ideal on the noncontrast dimension, $F(1, 44) = 33.14$, $p_{rep} = .99$. Typicality measures (recorded in Experiment 2 only) also showed evidence of idealization, so that typicality increased as items became more ideal on the contrast dimension, $F(1, 76) = 200.48$, $p_{rep} = .99$. This trend was significant, albeit weaker, on the noncontrast dimension, $F(1, 76) = 6.02$, $p_{rep} = .94$. Idealization was
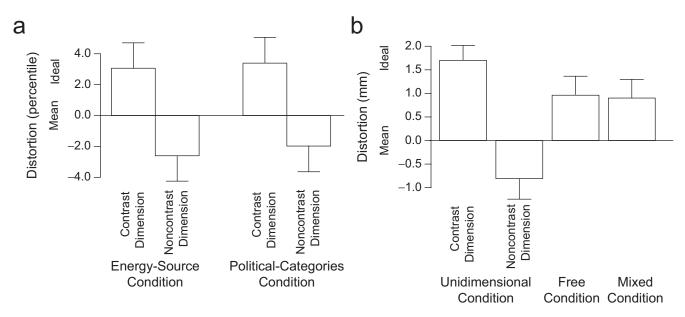

Fig. 4. Results of the reconstruction tasks in (a) Experiment 1 and (b) Experiment 2 as a function of dimension and condition. Reconstructions were scored as the difference between the reconstructed stimulus and the true category mean, with positive numbers indicating idealization and negative numbers indicating a distortion toward contrasting categories. The key dimension that discriminated between two category-response options was the contrast dimension. On the noncontrast dimension, there was no contrast between the two categories. In the mixed condition, the dimension on which stimuli contrasted alternated between trials. In the free condition, subjects could choose from all four categories on every trial. Error bars represent 95% confidence intervals.
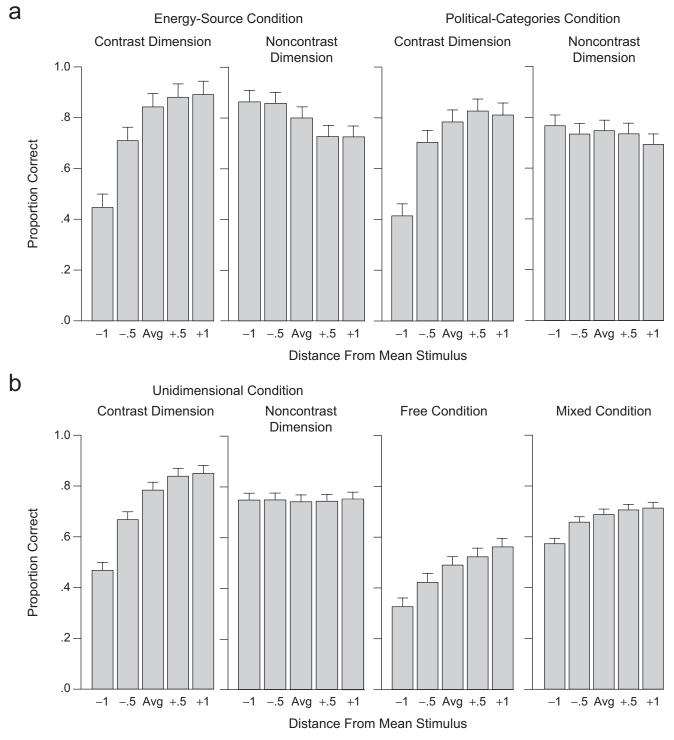
a

b

Fig. 5. Accuracy results for (a) Experiment 1 and (b) Experiment 2. Each bar represents the within-subjects mean proportion correct for stimuli at a particular distance (in *z* scores) from the mean stimulus along a stimulus dimension. Positive values indicate ideal stimuli, whereas negative values indicate stimuli located away from the mean and toward the contrasting category. Results are shown separately for the energy-source condition (contrast and noncontrast dimensions), the political-categories condition (contrast and noncontrast dimensions), the unidimensional condition (contrast and nonconstrast dimensions), the free condition, and the mixed condition. In the free condition, subjects could choose from all four categories on every trial. In the mixed condition, the dimension on which stimuli contrasted alternated between trials. Error bars represent standard errors of the mean.

observed on both (contrast) dimensions in Experiment 2's mixed and free conditions, so that accuracy, RT, and typicality measures increased as items became more extreme, Fs $\geq$ 10.49, p_{rep}s $\geq$ .98.

Discussion

We provide a learning account of how conceptual structure emerges that predicts the idealization effects observed in

studies of goal-derived categories and in the cross-cultural and expertise literature. This model predicts that idealization arises from contrastive learning mechanisms that minimize prediction error. When categories are contrasted along a dimension, opposing categories exert repelling forces that cause the categories to move apart along that dimension. This does not occur when dimensions are not contrasted, and thus this account retains the ability to predict cases in which typicality is structured around a category's central tendency. Likewise, when the forces from multiple contrasting categories roughly balance, representational distortion should be minimal. Using this theory, we were able to predict the pattern of subjects' categorization accuracy, RTs, typicality ratings, and stimulus reconstructions for categories in which subjects had prior beliefs and for completely novel categories.

Our work adds to a growing body of work that favors a transfer-appropriate processing account of category learning (Love, Medin, & Gureckis, 2004; Markman & Ross, 2003). Rather than acquiring veridical representations, people tailor category representations to the demands of the learning task. In our studies, these demands were subtly manipulated by changing the available choice options. Work comparing different induction tasks (e.g., unsupervised, inference, and classification learning) has converged on a similar conclusion (Love et al., 2004; Markman & Ross, 2003).

Previous work on caricature effects in natural and artificial categories closely relates to the present results. Goldstone and colleagues (Goldstone, 1996; Goldstone, Steyvers, & Rofosky, 2003) have found that caricatures have an advantage in RT and accuracy measures when categories are interrelated and central tendencies have an advantage when categories are isolated. Our work suggests that this view of category interactions should be elaborated to account for the nature of the stimulus dimensions (i.e., contrastive or noncontrastive). Other recent work suggests that the psychological representations of stimuli that show typicality-related advantages in ill-defined categories, such as natural categories (Ameel & Storms, 2006) and perceptual dot patterns (Palmeri & Nosofsky, 2001), are caricatures in relation to opposing categories, and not central tendencies. Our work suggests a mechanistic explanation for how these effects arise; the representations of these categories shift to minimize the possibility of error in categorization.

The majority of category-learning models do not anticipate our results. For example, Nosofsky's (1986) generalized context model (GCM) is not sensitive to category contrast during learning. The GCM veridically stores each stimulus experienced during training as an exemplar, irrespective of contrast. Although GCM could account for our results through a number of means (e.g., modified decision rules that take into account contrast, shifting exemplars, etc.), none of these solutions readily follow from GCM's formalism.

Instead, the present results favor models that use error-driven learning mechanisms in which representational shifts can occur as a result of the two-category forces we identify (see Fig. 1 for the basic concept and the Model Appendix in the Supplemental Material available on-line). Such mechanisms can be readily incorporated into a variety of models, including exemplar (e.g., GCM) and prototype models.

One intriguing aspect of the present research is the possibility that impression formation can be affected by compelling contrasts. The research we report here indicates that simply manipulating the available choice options can influence people's representations of categories. In politics, these contrasts are often created by rhetoric that omits certain choice options and introduces others (e.g., "You either support our policies and are a patriot or you are with the terrorists"). Language can be used to enshrine contrasts (e.g., the political left and right). Somewhat disturbingly, our work, coupled with previous efforts (Winkielman et al., 2006), suggests that such contrasts lead to caricatures that people should find easier to process and more compelling than the true categories.

Declaration of Conflicting Interests

Funding

Supplemental Material

Additional supporting information may be found at http://pss.sagepub.com/content/by/supplemental-data

Notes

1. Only nonoutlier RTs (i.e., those with z scores < 2.5) and correct trials were included. These constraints did not affect any statistical results.
2. Subjects failing to exceed chance performance by the final learning block were removed from subsequent analyses.
3. Test statistics are presented in summary fashion to eliminate redundancy. Degrees of freedom for individual tests can be computed from information provided in Table 1.

References

Ameel, E., & Storms, G. (2006). From prototypes to caricatures: Geometrical models for concept typicality. *Journal of Memory and Language*, *55*, 402–421.

Atran, S. (1999). Itzaj Maya folkbiological taxonomy. In D.L. Medin & S. Atran (Eds.), *Folkbiology* (pp. 119–203). Cambridge, MA: MIT Press.

Bang, M., Medin, D., & Atran, S. (2007). Cultural mosaics and mental models of nature. *Proceedings of the National Academy of Sciences, USA*, *104*, 13868–13874.

Barsalou, L.W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *11*, 629–649.

Goldstone, R.L. (1996). Isolated and interrelated concepts. *Memory & Cognition*, *14*, 608–628.

Goldstone, R.L., Steyvers, M., & Rofosky, B.J. (2003). Conceptual interrelatedness and caricatures. *Memory & Cognition*, *31*, 169–180.

Love, B.C., & Gureckis, T.M. (2005). Modeling learning under the influence of culture. In W. Ahn, R.L. Goldstone, B.C. Love, A.B. Markman, & P. Wolff (Eds.), *Categorization inside and outside of the lab: Festschrift in honor of Douglas L. Medin* (pp. 229–248). Washington, DC: American Psychological Association.

Love, B.C., Medin, D.L., & Gureckis, T.M. (2004). SUSTAIN: A network model of category learning. *Psychological Review*, *111*, 309–332.

Lynch, J.E., Coley, J.D., & Medin, D.L. (2000). Tall is typical: Central tendency, ideal dimensions, and graded structure among tree experts and novices. *Memory & Cognition*, *28*, 41–50.

Markman, A.B., & Ross, B.H. (2003). Category use and category learning. *Psychological Bulletin*, *129*, 592–613.

Medin, D.L., & Atran, S. (2004). The native mind: Biological categorization, reasoning, and decision making in development across cultures. *Psychological Review*, *111*, 960–983.

Nosofsky, R.M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–57.

Osherson, D.N., Smith, E.E., Wilkie, O., López, A., & Shafir, E. (1990). Category-based induction. *Psychological Review*, *97*, 185–200.

Palmeri, T.J., & Nosofsky, R.M. (2001). Central tendencies, extreme points, and prototype enhancement effects in ill-defined perceptual categorization. *The Quarterly Journal of Experimental Psychology*, *54A*, 197–235.

Posner, M.I., & Keele, S.W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*, 353–363.

Rosch, E.H., & Mervis, C.B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, *7*, 573–605.

Rosch, E.H., Simpson, C., & Miller, S.R. (1976). The structural bases of typicality effects. *Journal of Experimental Psychology: Human Perception and Performance*, *2*, 491–502.

Sakamoto, Y., Jones, M., & Love, B.C. (2008). Putting the psychology back into psychological models: Mechanistic vs. rational approaches. *Memory & Cognition*, *36*, 1057–1065.

Spencer, J.P., Perone, S., & Johnson, J.S. (2009). The dynamic field theory and embodied cognitive dynamics. In J.P. Spencer, M.S. Thomas, & J.L. McClelland (Eds.), *Toward a unified theory of development: Connectionism and dynamic systems theory reconsidered* (pp. 86–118). New York: Oxford University Press.

Winkielman, P., Halberstadt, J., Fazendeiro, T., & Catty, S. (2006). Prototypes are attractive because they are easy on the mind. *Psychological Science*, *17*, 799–806.