

Inferring Individual Differences Between and Within Exemplar and Decision-Bound Models of Categorization

Irina Danileiko (idanilei@uci.edu)

Michael D. Lee (mdlee@uci.edu)

Department of Cognitive Sciences, University of California, Irvine
Irvine, CA 92617 USA

Abstract

Different models of categorization are often treated as competing accounts, but specific models are often used to understand individual differences, by estimating individual-level parameters. We develop an approach to understanding categorization that allows for individual differences both between and within models, using two prominent categorization models that make different theoretical assumptions: the Generalized Context Model (GCM) and General Recognition Theory (GRT). We develop a latent-mixture model for inferring whether an individual uses the GCM or GRT, while simultaneously allowing for the use of special-case simpler strategies. The GCM simple strategies involve attending to a single stimulus dimension, while the GRT simple strategies involve using unidimensional decision bounds. Our model also allows for simple contaminant strategies. We apply the model to four previously published categorization experiments, finding large and interpretable individual differences in the use of both models and specific strategies, depending on the nature of the stimuli and category structures.

Keywords: category learning; exemplar models; decision bound models; General Recognition Theory; Generalized Context Model; Bayesian inference, latent-mixture model

Introduction

The ability to categorize is widely regarded as a cornerstone of cognition (Murphy, 2002). It allows people to use innate or learned concepts to give meaning to stimuli, and provides a basis for the key cognitive capabilities of generalization and prediction. Existing theories of categorization vary in their assumptions about how people represent categories, and make decisions when categorizing stimuli. Accordingly, there are many different formal models of categorization, including prototype, exemplar, decision-bound, and rule-based models, as well as various hybrid models (Kruschke, 2008). Hybrid models involve representational assumptions that combine two or more approaches such as prototype with exemplar, or exemplar with rule-based representations.

A prominent exemplar model is the Generalized Context Model (GCM: Nosofsky, 1986), which assumes people store exemplars of each category in memory, attend to the relevant dimensions of the stimuli, and categorize a stimulus using similarity-based generalization from these exemplars. A prominent decision-bound model is General Recognition Theory (GRT: Ashby & Townsend, 1986), which assumes that people use a decision bound to partition the stimuli into discrete categories. Categorization seems likely to depend heavily on psychological components and processes such as memory capacity, attentional control, decision-making biases, and so on, all of which may vary across people.

Accordingly, it seems reasonable to expect meaningful individual differences in categorization, and this expectation is supported by model-based and empirical evidence (e.g., Bartlema, Lee, Wetzels, & Vanpaemel, 2014; Soto, Vucovich, Musgrave, & Ashby, 2015). Previous work has studied different types of strategies within models like the GCM and GRT. For example, previous GRT modeling has emphasized the possibility that different people might use different decision bounds, including special cases like unidimensional horizontal or vertical bounds (Ashby & Gott, 1988; Maddox & Ashby, 1993).

In this paper, we present an approach to inferring the general models and specific strategies people use in categorization tasks. We do this by allowing for individual differences between the GCM and GRT models, and for individual differences in specific strategies, like unidimensional bounds, possible within each model. We develop a latent-mixture modeling approach that infers the model and strategy each person is using. Building on previous work in which both the GCM and GRT have been implemented as Bayesian graphical models (Lee & Wagenmakers, 2013; Danileiko, Lee, & Kalish, 2015), we implement our latent-mixture approach also as a graphical model, allowing for fully Bayesian inference. We apply our model to four existing categorization experiments—all involving stimuli that can be represented in terms of two underlying psychological dimensions—but with various types of stimuli and category structures. We find evidence for large individual differences both between and within models. We finish by discussing the implication of our model and results for future research in understanding how people represent categories.

Latent-Mixture Model

Our latent-mixture approach assumes that each subject uses one categorization model or specific strategy within that model, and that the overall data set is therefore a mixture of these specific components. We also allow for the possibility of contaminant subjects, who are guessing or repetitively assigning stimuli to the same category. Instead of filtering these people out, we model the contaminant behavior as another mixture component (Zeigenfuse & Lee, 2010).

Figure 1 presents a schematic graphical model that summarizes our approach. Each subject's categorization data y is generated by either the GCM, GRT, or a contaminant process. Within each of these general models, there are specific possibilities. For the GCM, either the original model with general




Figure 1: Schematic graphical model representation of the latent-mixture approach. A GCM or ,GRT, or contaminant categorization process generates the observed behavior of each individual. Within each model, special-case strategies involving the nature of selective attention, the decision bound, or the contaminant response probabilities are considered.

selective attention w is used, or attention is focused on only one of the stimulus dimensions, $w = 0$ or $w = 1$. For the GRT, either a general diagonal bound is used α^D, β^D , or unidimensional horizontal β^H or vertical β^V bounds are used. For the contaminant processes, either a category is repeatedly chosen θ^R or a guess is made on each trial θ^G .

GCM and Strategies

The GCM is an exemplar model that assumes people store all stimuli in memory and categorize a new stimulus by comparing it the stored stimuli. It is based on a similarity comparison between the presented stimulus and every other stimulus, using the concept of psychological distance (Shepard, 1957). If the stimuli are points in a two-dimensional coordinate space, psychological distance is defined as $d_{ij} = [\sum_{k=1}^N w_k |x_{ik} - x_{jk}|^r]^{\frac{1}{r}}$, where x_{ik} is the value of coordinate point x_i on dimension k , N is the number of dimensions, and r is either equal to 1 or 2 for separable-dimension or integral-dimension stimuli, respectively. The selective attention parameter w controls the level of attention given to one stimulus dimension. The distance is used to calculate the similarity $\eta_{ij} = (e^{-cd_{ij}})^\gamma$ so as the distance between points gets larger, their perceived similarity decreases exponentially. The generalization parameter c controls the steepness of the generalization gradient. The response determinism parameter γ_k controls probabilistic or deterministic responding (Ashby & Maddox, 1993). The final probability based on these processes can be affected by the bias toward each category. In our implementation, we assume there is no bias, and so set $b = \frac{1}{2}$. This means that the probability of responding J to

stimulus i is equal to $\theta^{GCM} = \sum_{j \in C_j} \eta_{ij} / \sum_{K=1}^m (\sum_{k \in C_K} \eta_{ik})$.

In the full GCM, selective attention can range between 0 and 1. The special cases of $w = 0$ and $w = 1$ correspond to attending to just one of the two stimulus dimensions, and constitute theoretically interesting strategies. For example, if one stimulus dimension is shape and the other is color, one person might attend only to the shape dimension and place circles in one category and squares in the other. However, another person might attend only to the color dimension and place red shapes in one category and blue shapes in the other. A third person might attend to both dimensions and categorize red circles separately from blue squares. These possibilities correspond to the three GCM components included in our model, represented by the $w = 0$, $w = 1$, and w nodes respectively.

GRT and Strategies

The GRT model assumes that instead of storing each stimulus in memory, people partition the stimulus space into response regions divided by boundaries. Response probabilities are determined by these decision bounds, based on which region a noisy perception of the presented stimulus, $x_{pi} = x_i + \epsilon_p$, belongs. Our model considers only linear decision bounds, although quadratic bounds have also been considered in the literature (Ashby & Maddox, 1992). A linear bound is defined as a discriminant function of the two dimensions satisfying the implicit line equation $h(x_1, x_2) = b_1 x_1 + b_2 x_2 + c$.

GRT assumes that there is variability in people's memory of the location of the bound. To account for this, the function is adjusted to include criterial noise ϵ_c . The function is compared to a bias parameter δ which captures bias toward a category. If $h(x_{pi}) + \epsilon_c$ is smaller than δ , the response is category A. If it is larger than δ , the response is category B. If it is equal to δ , the response will be a guess between A and B. We again assumed no bias so that $\delta = 0$. Thus, the probability of a category A decision for stimulus i is $\theta^{GRT} = P(h(x_{pi}) + \epsilon_c < 0)$.

Special cases of the general GRT model that have previously been emphasized involve unidimensional boundaries corresponding to vertical or horizontal lines. A vertical strategy is defined by an intercept value β^V , a horizontal strategy is defined by an intercept value, β^H , and a general diagonal boundary is defined by a slope and intercept α^D and β^D . These possibilities correspond to the three GRT components included in our model. Although the vertical and horizontal strategies can be viewed as special cases of the diagonal strategy, one way to think about this in the latent-mixture approach is as a single model with a theoretically-rich prior. Including the vertical and horizontal boundaries as special cases corresponds to considering only a diagonal boundary with a prior that places significant density on boundaries with infinite and zero slope.

Contaminant Strategies

The three GCM strategies and the three GRT strategies make up six mixture components in our model. The remaining components capture contaminant subjects. In these cases, the probability values θ^{cont} do not follow from a theoretical

model, but are set directly. For guessing, the response probability is $\theta^G = 0.5$, so that each category response is equally likely on every trial. For a repetitive contaminant behavior, the probability of a category A response is either $\theta_A^G = 0.99$ or $\theta_B^G = 0.01$, depending on which category choice is repeated. Adding these three contaminant possibilities leads to a total of nine components of our latent mixture model, with Figure 1 combining the two repeated contaminant possibilities.

Modeling Results

We implemented the graphical model in JAGS (Plummer, 2003), and used fully Bayesian methods based on MCMC sampling to make inferences.¹ Advantage of this methodological approach include accounting coherently for uncertainty about inferences, both in terms of model use and model-specific parameters, and automatically controlling for the different complexity of the models and strategies considered (Lee & Wagenmakers, 2013).

We applied the model to four previously published categorization experiments. These experiments all involved a series of trials in which subjects viewed a stimulus and placed it into one of two categories, with corrective feedback after each trial. The stimuli used varied across experiments and include rectangles varying in size and interior line segment position (Kruschke, 1993), Shepard circles varying in size and radial lines (Bartlema et al., 2014), Gabor patches varying in frequency and orientation (Zeithamova & Maddox, 2006), and faces (Navarro, Lee, & Nikkerud, 2005). For the first three, there is a natural two-dimensional stimulus representation. For the faces, we assumed a two-dimensional representation based on multidimensional scaling modeling (Okada & Lee, 2016). Details of the experiments, including the number of subjects, blocks, nature of the experiment, and the various conditions, are presented in Table 2. The code for the graphical model, categorization data, detailed analysis of every subject in every experiment and condition, and other supplementary material is available on the Open Science Framework project page for this paper at <https://osf.io/ckwsn/>.

Overall Results

Table 1 summarizes our results by listing how many people are inferred most likely be using each of the possible models and strategies. The individual model-use inferences come are seen in the indicator variable in the JAGS code that assumes a uniform prior over all nine potential models, meaning that in the prior, each person is equally likely to use any of the nine. The “most likely” model for each person is taken from the posterior distribution of the indicator variable. There are four conditions in the Kruschke (1993) experiment: the first two are filtration category structures, in which the stimuli can be categorized correctly by using information from only one dimension, and the second two are condensation category struc-

tures, in which the stimuli can only be categorized correctly by using information from both dimensions. The majority of the 160 participants are inferred to use the GCM exemplar approach, but the specific selective attention strategy varies by condition. The Bartlema et al. (2014) experiment has two conditions, named after the category structures, both of which require information from both stimulus dimensions for correct categorization. The majority of the 65 participants use a decision bound approach. The Zeithamova and Maddox (2006) experiment has four conditions. The unidimensional condition is similar to the filtration condition in the Kruschke experiment and the information-integration condition is similar to the condensation condition. The “+ load” label in Table 1 indicates that that condition also involved a simultaneous working memory load task. The majority of the 170 participants use the decision-bound approach, with the vertical strategy being most common. This experiment involves the most contaminant subjects, who are inferred primarily to be guessing. The Navarro et al. (2005) experiment has four conditions. These involved categorizing faces based on gender, hair color, perceived level of trust, and a random condition with no logical structure. The majority of the 40 subjects use an exemplar approach, with selective attention that considers both available dimensions, but there is large individual variation over both models and strategies across the conditions.

Kruschke (1993) Results The results from the Kruschke (1993) experiment are shown at an individual level, for selected subjects, in Figure 2. The circles show the eight stimuli. The dark-colored circles show a response of category A while the light-colored circles show a response of category B. The size of the circle shows the number out of the total number of trials that each stimulus was placed in either category. The smallest circle means that stimulus was placed into that category exactly half of the time while the largest circle means that stimulus was placed into that category all the time. The bar graphs on the top of each panel show the uncertainty in the inference about which model and strategy the subject used. Each bar shows the posterior probability for a model or strategy. A tall bar showing one strategy means that we can be more certain of that person’s inferred strategy than when there are shorter bars showing many strategies. The text at the bottom right corner of each panel indicates the inferred most likely strategy. For the general GCM w strategy, the 95% credible intervals and posterior mean for w are listed. For the GRT possibilities, the bound corresponding to the posterior mean is shown as a thick line, and the upper and lower bounds f to the 95% credible intervals are shown as thin lines.

These subject in Figure 2 are chosen to include at least one subject from each condition. The top-left came subject from the Filtration 1 condition, the top-middle subject came from the Filtration 2 condition, the top-right and bottom-left subjects came from the Condensation 1 condition, and the bottom-middle and bottom-right subjects came from the Condensation 2 condition. The first two subjects from the filtration conditions are inferred to be most likely using an exem-

¹Our results are based on 3 independent chains with 100 samples each, collected after discarding the first 500 burn-in ones from each chain, and testing for convergence using the standard \hat{R} statistic.

		Exemplar			Bound			Contam.	
		<i>w</i>	0	1	V	H	D	G	R
Kruschke	Filtration 1	10	-	30	-	-	-	-	-
	Filtration 2	6	30	1	3	-	-	-	-
	Condensation 1	15	4	5	-	4	8	4	-
	Condensation 2	19	6	10	2	-	1	2	-
	Total	50	40	46	5	4	9	6	-
Bartlema et al.	Diagonal	-	3	-	5	15	7	1	-
	Criss-Cross	-	3	4	8	8	4	7	-
	Total	-	6	4	13	23	11	8	-
Zeithamova & Maddox	Unidimensional	-	2	-	31	1	1	5	1
	Unidimensional + load	3	5	7	22	5	2	5	1
	Information-Integration	-	-	-	20	-	11	2	1
	Information-Integration + load	1	9	3	19	0	10	3	-
	Total	4	16	10	92	6	24	15	3
Navarro et al.	Gender	9	-	1	-	-	-	-	-
	Hair	3	3	1	-	3	-	-	-
	Trust	4	1	1	2	-	1	1	-
	Random	4	-	2	-	1	-	3	-
	Total	20	4	5	2	4	1	4	-

Table 1: Number of participants inferred to use an exemplar, decision bound, or contaminant strategy in each data set. (*w*: uniform *w* strategy; 0: *w*=0 strategy; 1: *w*=1 strategy; V: vertical; H: horizontal; D: diagonal; G: guess; R: repeat (either)

plar strategy with $w = 1$ and $w = 0$, with some possibility of the general GCM w strategy. The subjects from the condensation conditions are inferred to be more likely to use either a diagonal boundary or a general GCM w strategy with a mean value close to $w = 0.5$ in one case, and $w = 0.83$ in the other. The last subject is inferred to be a guessing contaminant, with a larger degree of uncertainty.

Bartlema et al. (2014) Results The results from the Bartlema et al. (2014) experiment are shown at an individual level in Figure 3. The top panels come from the diagonal condition and the bottom panels come from the criss-cross condition. The top-left subject is inferred to be using a horizontal boundary, but with some uncertainty about the possible use of a more general diagonal boundary. The bottom-left subject is also inferred to be using a horizontal boundary, but there is a possibility of a diagonal boundary, or a contaminant guessing strategy. The top-right subject is inferred to be using a diagonal boundary, with greater uncertainty. The bottom-right subject is inferred to be using a vertical boundary, also with a high level of certainty.

Zeithamova and Maddox (2006) Results The results from the Zeithamova and Maddox (2006) experiment are shown at an individual level in Figure 4. The top-left subject comes from the unidimensional condition. The top-middle and top-right subjects come from the unidimensional + load condition. The bottom-left subject come from the information-integration condition. The bottom-middle and bottom right subjects come from the information-integration + load condition. In this experiment, very few of the subjects were in-

ferred to be using an exemplar strategy, perhaps as a result of the large number of stimuli required to keep in memory. Even though most subjects were inferred to be using a decision bound, there is still great variation in the specific shape of the boundaries, with varying slopes and intercepts. Two of the subjects selected for Figure 4 are inferred to be using a vertical boundary, even though they come from conditions with different category structures. Similarly, two of the subjects are inferred to be using a diagonal boundary, but one with more uncertainty than the other about the location of the boundary. This experiment also involved a large number of subjects inferred to be contaminants, one of whom is shown in Figure 4. The top-right panel shows one subject who was inferred to be using an exemplar approach with a $w=1$ strategy, although there is large uncertainty about this inference, consistent with the poor categorization performance shown.

Navarro et al. (2005) Results The results from the Navarro et al. (2005) experiment are shown at an individual level in Figure 5. The top-left subject comes from the gender con-

Experiment	# Su	# B	# St	# C	Type of St
Kruschke	160	8	8	4	Rectangles
Bartlema et al.	65	40	8	2	Shepard circles
Zeithamova & Maddox	170	5	80	4	Gabor patches
Navarro et al.	40	8	25	4	Faces

Table 2: Properties of the categorization experiments (Su: subjects; B: blocks; St: stimuli; C: conditions).

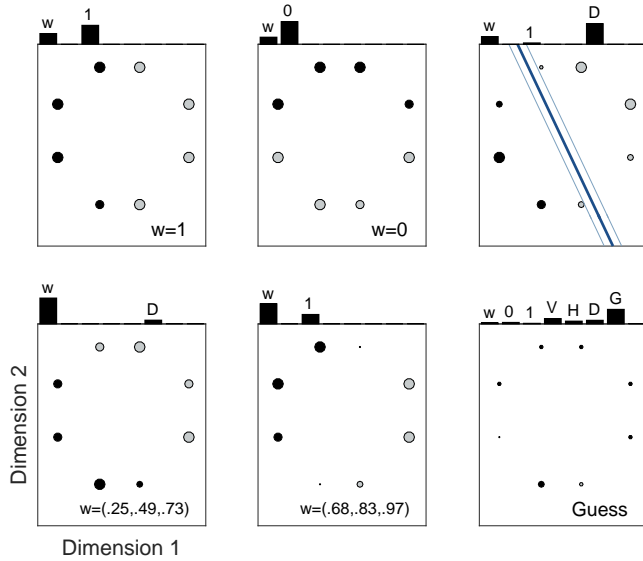


Figure 2: Inferred model or strategy use, and attention values or decision bounds, for selected subjects from the Kruschke (1993) experiment.

dition, the top-middle and top-right subjects come from the hair-color condition, the bottom-left subject comes from the trust condition, and the bottom-middle and bottom-right subjects come from the random condition. Most of the subjects are inferred to be using the GCM, perhaps as a result of the stimuli being faces and not easily separable into psychologically interpretable dimensions. Two of the selected subjects are inferred to be using the general GCM w strategy, with varying mean values depending on the condition. A few subjects are inferred, with less certainty, to be using a decision-bound approach. The random condition has the most contaminants, as for the subject in the bottom-right panel, typically with large uncertainty about model use.

Discussion

We have presented a latent-mixture model that infers which of the two models—the GCM or the GRT—each person is using, and whether they are using a specific strategy within that model. Our individual differences analysis showed that different people’s categorization behavior can best be explained by different model strategies, depending on the types and number of stimuli involved, and the nature of the category structures. Instead of continuing a debate of a “one model fits all” answer where all behavioral data must be in accordance with one type of model, applications of our modeling approach to individual subject data can potentially reveal multiple models and strategies being used by different people.

Future work could apply our general method, and the specific model we have implemented and demonstrated, to other categorization experiments, exploring how individual differences change with the type of stimuli and category structures involves. It would be interesting to understand individual




Figure 3: Inferred model or strategy use, and attention values or decision bounds, for selected subjects from the Bartlema et al. (2014) experiment.

differences for more complicated real-world stimuli, such as faces, with the goal of understanding how people categorize in everyday life. It is straightforward to extend our model to include other theoretical accounts of categorization behavior, and different specific strategies within them. These could incorporate other categorization models, such as hybrid models that combine prototype with exemplar or rule-based representations. It would also be possible to extend the model to allow for shifts in categorization within an individual, allowing for possibilities like rapid shifts in attention, or the adaptation of an overly simple unidimensional bound to a more general diagonal bound on the basis of feedback. Examination of the strategy shifts that occur can be useful for further predictive modeling of when we can expect participants to switch strategies. Collectively, these extensions allow for broader and deeper investigation of the individual differences in the way people represent and use categories.

Acknowledgments

This work was supported by NSF Award 1431635.

References

- Ashby, F. G., & Gott, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *14*, 33–53.
- Ashby, F. G., & Maddox, W. T. (1992). Complex decision rules in categorization: Contrasting novice and experienced performance. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 50–71.
- Ashby, F. G., & Maddox, W. T. (1993). Relations between prototype, exemplar, and decision bound models of cat-

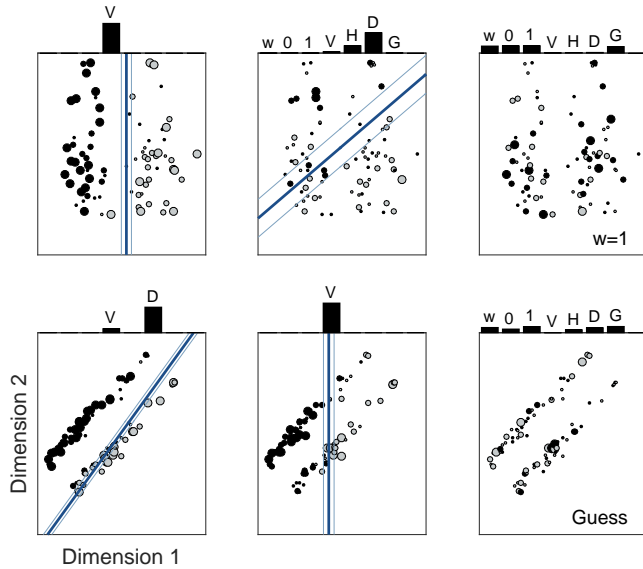


Figure 4: Inferred model or strategy use, and attention values or decision bounds, for selected subjects from the Zeithamova and Maddox (2006) experiment.

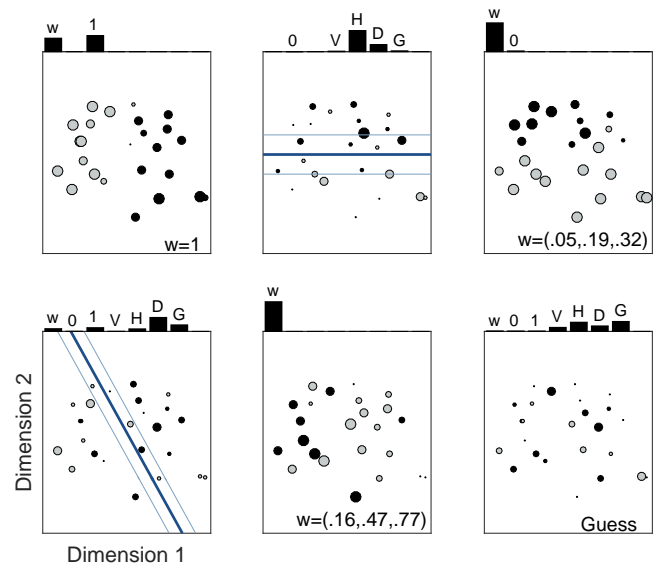


Figure 5: Inferred model or strategy use, and attention values or decision bounds for selected subjects from the Navarro et al. (2005) experiment.

egorization. *Journal of Mathematical Psychology*, 37, 372-400.

- Ashby, F. G., & Townsend, J. T. (1986). Varieties of perceptual independence. *Psychological Review*, 93, 154-79.
- Bartlema, A., Lee, M., Wetzels, R., & Vanpaemel, W. (2014). A Bayesian hierarchical mixture approach to individual differences: Case studies in selective attention and representation in category learning. *Journal of Mathematical Psychology*, 59, 132-150.
- Danileiko, I., Lee, M. D., & Kalish, M. L. (2015). A Bayesian latent mixture approach to modeling individual differences in categorization using General Recognition Theory. In D. C. Noelle et al. (Eds.), *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (p. 501-506). Cognitive Science Society.
- Kruschke, J. K. (1993). Human category learning: Implications for backpropagation models. *Connection Science*, 5, 3-36.
- Kruschke, J. K. (2008). Models of categorization. In R. Sun (Ed.), *The Cambridge Handbook of Computational Psychology* (pp. 267-301). New York: Cambridge University Press.
- Lee, M. D., & Wagenmakers, E.-J. (2013). *Bayesian Cognitive Modeling: A Practical Course*. Cambridge University Press.
- Maddox, W. T., & Ashby, F. G. (1993). Comparing decision bound and exemplar models of categorization. *Perception and Psychophysics*, 53, 49-70.
- Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.
- Navarro, D. J., Lee, M. D., & Nikkerud, H. (2005). Learned

categorical perception for natural faces. In B. G. Bara, L. W. Barsalou, & M. Bucciarelli (Eds.), *Proceedings of the 27th Annual Conference of the Cognitive Science Society* (pp. 1600-1605). Mahwah, NJ: Erlbaum.

- Nosofsky, R. M. (1986). Attention, similarity and the identification-categorization relationship. *Journal of Experimental psychology: General*, 115, 39-57.
- Okada, K., & Lee, M. D. (2016). A Bayesian approach to modeling group and individual differences in multidimensional scaling. *Journal of Mathematical Psychology*, 70, 35-44.
- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In K. Hornik, F. Leisch, & A. Zeileis (Eds.), *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*. Vienna, Austria.
- Shepard, R. N. (1957). Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space. *Psychometrika*, 22, 325-345.
- Soto, F. A., Vucovich, L., Musgrave, R., & Ashby, F. G. (2015). General recognition theory with individual differences: a new method for examining perceptual and decisional interactions with an application to face perception. *Psychonomic Bulletin & Review*, 1-24.
- Zeigenfuss, M. D., & Lee, M. D. (2010). Finding the features that represent stimuli. *Acta Psychologica*, 133, 283-295.
- Zeithamova, D., & Maddox, W. T. (2006). Dual task interference in perceptual category learning. *Memory & Cognition*, 34, 387-398.