

Explaining Categorization Response Times with Varying Abstraction

Loes Stukken (loes.stukken@psy.kuleuven.be)

Department of Psychology. Tiensestraat 102

3000 Leuven, Belgium

Gert Storms (gert.storms@psy.kuleuven.be)

Department of Psychology. Tiensestraat 102

3000 Leuven, Belgium

Wolf Vanpaemel (wolf.vanpaemel@psy.kuleuven.be)

Department of Psychology. Tiensestraat 102

3000 Leuven, Belgium

Abstract

We use the Exemplar-Based Random-Walk model (EBRW) to extend the Varying Abstraction Model (VAM). Unlike the VAM which is designed to account for categorization proportions, this Varying Abstraction-Based Random-Walk (VABRW) model is able to predict categorization response times. The extension is especially useful in situations where response accuracies are not very informative for distinguishing between models. Application of the VABRW to data previously gathered by Nosofsky and Palmeri (1997) provides additional evidence for the view that people use partial abstraction in category representations.

Keywords: response times; categorization; abstraction; cognitive models

Introduction

Because of its crucial role in almost every aspect of cognitive processing, categorization has been studied extensively in the past 35 years. Many different models have been proposed to account for the categorization results and tested using empirical data collected in category learning tasks. The vast majority of these experiments and models have focused on categorization decisions and on typicality ratings, neglecting the corresponding processing times. This is surprising, as response times are arguably the most prominent dependent variable in the whole discipline of cognitive psychology. Response times have been argued to provide a window into understanding the nature of cognitive representations and decision processes (Nosofsky & Palmeri, 1997).

The few existing attempts of modeling categorization response times all start from the exemplar view on category representation (Smith & Medin, 1981). In this view, a category is assumed to be represented by memory traces of all the exemplars that were previously encountered as a member of the category. For example, Lamberts (2000) has extended Nosofsky's (1986) Generalized Context Model (GCM) to the Extended Generalized Context Model for response times and Nosofsky and Palmeri's (1997) Exemplar Based Random Walk Model (EBRW) combines elements of the GCM with Logan's (1988) instance based model of automaticity. Both approaches showed the

usefulness of extending models that only account for categorization decisions to models that also predict response times.

The exemplar view on category representation is not the only view that has gained support. The most notable alternative to the exemplar view in accounting for categorization decisions is the prototype view, which holds that a category is represented by a summary of its members. Though the exemplar model seems to outperform the prototype model in general (e.g., Nosofsky, 1992; Vanpaemel & Storms, 2010), in some conditions the prototype representation seems to be supported (Minda & Smith, 2001; Smith & Minda, 1998). Moreover, models have been proposed that leave room for partial abstraction in category representation (Anderson, 1991; Griffiths, Canini, Sanborn, & Navarro 2007; Love, Medin & Gureckis, 2004; Rosseel, 2002; Vanpaemel & Storms, 2008). These models provided a broader window on representational abstraction suggesting that also levels of abstraction that are intermediate between exemplar and prototype models are viable representations that deserve consideration.

Despite more than three decades of research on categorization, the question about the nature of category representation still awaits conclusive evidence. Moreover, the field seems to be converging to the view that human conceptual structure is sufficiently flexible to adopt highly abstract representations as well as exemplar representations, shifting the focus from establishing the single representation that always underlies categorization decisions, to identifying the conditions in which one representation is more likely to be used than another. In this paper, we contribute to this debate by presenting a general framework to account for decision times. More specifically, we generalize the EBRW model to encompass a broad set of representations each assuming different levels of abstraction, as implemented in the Varying Abstraction Model (VAM; Vanpaemel & Storms, 2008). In what follows, we first describe the way the VAM and EBRW model are combined into the Varying Abstraction-Based Random-Walk model (VABRW). Next we apply the VABRW to data collected by Nosofsky and Palmeri (1997).

VABRW

The VABRW is most easily described in two steps: the different representations it assumes, and the processes and mechanisms that, when added to these representations, give rise to the categorization behavior of interest.

VAM Representations

Following the GCM, the VAM starts with the assumption that the members of a category can be represented as points in a multidimensional psychological space. Category representations are constructed by dividing these points in clusters and averaging the coordinates of the points within these clusters. The resulting coordinates define the set of subprototypes that make up the category representation.

For example, if cluster Q_j contains n_j stimuli, the coordinate value of μ_{jk} for subprototype j on dimension k can then be calculated as follows:

$$\mu_{jk} = \frac{1}{n_j} \sum_{x_i \in Q_j} x_{ik},$$

where x_{ik} is the coordinate value of stimulus i on dimension k .

In general, there is no restriction on how the members of a category should be divided in clusters. Thus, the VAM encompasses all the representations that can be constructed from a category, including the exemplar representation as the least abstract and the prototype representation as the most abstract representation. If, for example, a category consists of four members, 15 different category representations can be created (see Figure 1).

The fact that the VAM considers all the representations that can be constructed from a category raises a problem when a category contains a large number of exemplars. In this case, the number of possible category representations within the VAM becomes quite large and this will severely complicate model fitting. A more constrained version of the VAM therefore favors representations that are based on the clustering of similar stimuli (Vanpaemel, 2011) while an even more constrained version considers only one representation at each level of abstraction (Verbeemen, Vanpaemel, Pattyn, Storms & Verguts, 2007). This latter version of VAM uses a K-means clustering procedure (Hastie, Tibshirani, & Friedman, 2001) in order to reduce the number of possible representations per category. The K-means clustering procedure selects the most plausible clustering of category members based on the similarity between the category members, assigning similar members to the same cluster and keeping dissimilar members separate. The K-means clustering procedure would for example select category representations A (exemplar representation), E (out of the six representations with three subprototypes), N (out of the seven representations with two subprototypes) and O (prototype representation) from the representations shown in Figure 1. Within this study we

used this latter version of the VAM to obtain the appropriate category representations.

Predicting Classification Response Time

In order to derive response time predictions from the intermediate category representations of the VAM, we combined the VAM representations with the processing assumptions of the EBRW model. The EBRW model assumes a random walk process in which the categorizer gathers evidence that a particular stimulus belongs to a particular target category. More specifically, the EBRW model assumes that when a to be categorized stimulus is presented, the exemplars of the target categories race to be retrieved, with rates proportional to the similarity of the exemplar to the stimulus. The exemplar that is the first to be retrieved provides evidence in favor of the category to which it belongs. When enough evidence is gathered and a category criterion is reached, the appropriate response is executed.




Figure 1: The 15 possible representations for a category with four members. The subprototypes are represented by the black circles and are connected by lines to the original category members (the white circles). Panel A shows the exemplar representation (no members are merged); Panel O shows the prototype representation (all members are merged together into a single subprototype); Panels B-G show intermediate representations with three subprototypes; and Panels H-N show intermediate representations with two subprototypes.

Although the EBRW model starts from the assumption that a category is represented by an exemplar representation, its mechanisms can be combined with other representational assumptions (see, for example, Nosofsky & Stanton, 2005 for a prototype variant of the EBRW model). To combine the EBRW processes with the representations of the VAM it suffices to assume that upon the presentation of a target stimulus, the subprototypes that constitute the category will enter the race with rates related to the similarity of the stimulus to the subprototype and that the subprototype that is the first to be retrieved provides evidence in favor of the category to which it belongs.

The expected duration of the random walk when stimulus i is presented and thus the predicted response time for stimulus i is given by:

$$E(T|i) = E(N|i) E(T_{step}|i),$$

in which $E(N|i)$ represents the expected number of steps and $E(T_{step}|i)$ the expected time needed to take each individual step. The expected time to take each individual step is computed by:

$$E(T_{step}|i) = \alpha + \frac{1}{(S_{iA} + S_{iB})},$$

in which α is a constant parameter that can be interpreted as the time needed to find the category to which the retrieved exemplar belongs. S_{iA} and S_{iB} are the similarity of stimulus i to category A and category B, respectively, and are calculated by summing the similarities between the stimulus and the subprototypes that constitute category A and category B:

$$S_{iA} = \sum_{\mu_j \in S_A} s(x_i, \mu_j),$$

in which S_A denotes the set of subprototypes representing A. The similarities between the stimulus and the subprototypes can be computed by first determining the distances between the stimulus and the subprototypes that make up the category in the multidimensional psychological space. The distance between stimulus i and subprototype j is calculated by:

$$d(x_i, \mu_j) = \left[\sum_{k=1}^D w_k |x_{ik} - \mu_{jk}|^2 \right]^{\frac{1}{2}},$$

where x_{ik} and μ_{jk} are the coordinate values of stimulus i and subprototype j on dimension k and $0 < w_k < 1$ is an attentional weight associated with dimension k . These distances are then converted to obtain a measure of the similarity from each stimulus to every subprototype

$$s(x_i, \mu_j) = e^{-cd(x_i, \mu_j)},$$

where c is the sensitivity parameter.

The expected number of steps in the random walk is calculated by:

$$E(N|i) = \frac{B}{q_i - p_i} - \frac{A + B}{q_i - p_i} \left[\frac{1 - (q_i/p_i)^B}{1 - (q_i/p_i)^{A+B}} \right],$$

if $p_i \neq q_i$ and

$$E(N|i) = AB,$$

otherwise, where A and B are integers that represent the criteria or the amount of evidence needed to execute an A or B response, and p_i can be computed by:

$$p_i = S_{iA}/(S_{iA} + S_{iB}),$$

and q_i by:

$$q_i = 1 - p_i.$$

The free parameters in the EBRW model are the sensitivity parameter c , the attentional weight w_k , the criteria A and B that represent the amount of evidence needed before a response is executed and the time constant α . Furthermore two scaling parameters (slope k , intercept μ) are added to rescale the model predictions into realistic response times. As VABRW shares its processing assumptions with EBRW, it has identical parameters. The only difference is that while EBRW assumes only a single representation (the exemplar representation), VABRW considers a range of representations (including the exemplar and prototype representation). These representations could be indexed by a parameter, that unlike more traditional parameters is discrete. The goal of applying VABRW to data, as we will do in the next section, is to estimate the value of this discrete parameter, thereby informing us about the level of abstraction people rely upon.

Application of VABRW

Experimental Procedure In an effort to test the EBRW model, Nosofsky and Palmeri (1997) administered an experiment (their Experiment 1) in which three participants were asked to perform a speeded classification task. The stimuli in the task were 12 Munsell colors. Nosofsky and Palmeri constructed two categories by dividing the stimuli in two categories of six members each: category A, represented by circles in Figure 2 and category B represented by squares in the same figure.

The speeded classification task was administered in five sessions of 30 blocks, each for a total of 150 blocks. Within a block, each color was presented only once, in a randomized order. Participants were asked to respond as quickly and accurately as possible. Corrective feedback was provided after each trial. After the speeded classification task, participants completed a similarity scaling task in which they were asked to rate the similarity of each pair of colors on a 10-point scale. These similarity ratings were then used to construct a multidimensional space for each participant.




Figure 2: The category structure used in Experiment 1 of Nosofsky and Palmeri (1997). Category A is represented by circles, category B by squares.

Data Nosofsky and Palmeri (1997) computed for each of the three participants, response times for individual stimuli averaged across blocks 31 to 150, yielding 12 data points, and response times for groups of 5 blocks averaged across individual stimuli, yielding 30 data points. Additionally, the accuracies for the individual stimuli were computed, but since these were close to ceiling they were not included in their model analyses. The response times for the individual stimuli showed a very regular pattern: stimuli lying far from the category boundary were responded to faster than stimuli lying close to the category boundary. The response times for each group of 5 blocks showed, as expected, a decrease in response times with practice. Nosofsky and Palmeri showed that the individual participant data could be well accounted for by the EBRW model.

We believed these data to be appropriate to apply the VABRW model. As mentioned above, the main focus of this application is investigating which category representations is most supported by the data. Furthermore fitting the model to the individual participant data allowed us to check whether there were individual differences in the participants category representations.

Fitting Procedure The most constrained version of the VABRW contained 36 different representations, resulting from the combination of 6 representations for each category which were identified using K-means clustering. All 36 models were fit separately following the same procedure as Nosofsky and Palmeri (1997). Each model was fitted to the data by searching for a single set of parameters to simultaneously predict both the mean response time for each stimulus and the speed-up curves (i.e. the mean response time as a function of grouped block), using the total sum of squared errors (SSE) as a goodness of fit criterion. Like Nosofsky and Palmeri, we further assumed that the amount of evidence that needed to be gathered in order to execute a categorization response was the same for the two target categories (A=B). In order to predict the speed-up curves,

Nosofsky and Palmeri assumed that on each block an additional token of the exemplar was stored in memory. This assumption was implemented by multiplying the similarity of the stimulus to the category with the number of blocks. We instantiated the same multiplication, formalizing the assumption that with each block the participants got more confident about the category representation they used. The only aspect in which our fitting procedure differed from Nosofsky and Palmeri's was that we restricted A to values between 1 and 10, to make the model fitting more feasible.

Results Consistent with earlier findings, the exemplar representation captured the data better than the prototype representation. The optimal representations for Participants 1, 2 and 3 are presented in Panels A, B and C of Figure 3. As in Figure 1, the subprototypes that make up the category representations are presented in black while the category members that have been clustered together to obtain these subprototypes are presented in white. When we consider the best fitting representation, the three participants differ in the type of representation that best accounted for their data. Participant 3 seemed to rely on an exemplar representation, a finding that fits well with the impressive fits that the EBRW model provided to the data. Participants 1 and 2 in contrast relied on less detailed representations. For Participant 1 the best fitting model consists of 5 subprototypes in Category A and 4 subprototypes in Category B and thus comes quite close to the fully detailed exemplar representation. Participant 2, in contrast, relied on a much more abstract representation consisting of 2 category A subprototypes and 2 category B subprototypes.

Table 1 shows the correlations between the observed and predicted values, as well as the best fitting parameter values for each participant. The relatively high correlations between the observed and predicted values indicate that the models did a fairly good job in accounting for the response time data.

Table 1: Best-fitting parameters and correlations.

	Participants		
	1	2	3
c	2.6831	0.6663	1.7910
A	10	9	4
α	0.1794	0.1964	0.0558
W_x	0.2871	0.3127	0.7469
k	224.4538	45.0123	661.5330
μ	293.7754	342.0962	447.1176
correlations			
individual	0.9339	0.9936	0.9570
speed-up	0.9281	0.6297	0.9334

Note. c= sensitivity parameter, A= the criterion parameter, α = time constant, W_x = weight for dimension x, k = slope and μ = intercept. Correlations on the row "individual" are the correlations with the individual mean response times and correlations on the row "speed-up" are the correlations with the speed-up curves.

Discussion

Generally, different computational models of category learning are compared on their ability to account for the choice proportion in a category learning task. Though useful to discriminate among competing models, choice proportions provide only a limited window on categorization behavior. Also other variables like typicality and classification response times are informative about the underlying category representations and processes. Several researchers attempted to extend categorization models to be able to predict classification response times in category learning tasks. (Ashby, Boynton & Lee, 1994; Lamberts, 2000; Nosofsky & Palmeri, 1997; Nosofsky & Stanton, 2005).

In this paper we extended the VAM, designed to predict choice proportions, to enable it to account for response times by combining the set of category representations provided by the VAM with the process assumption of the EBRW model. Based on response time data from Nosofsky and Palmeri (1997), we showed that the best-fitting representations of two of the three participants relied on some amount of partial abstraction. This finding adds to the growing body of evidence that categorization should not always be exemplar based, but can rely on intermediate representations as well (e.g., Griffiths et. al. 2007, Vanpaemel & Storms, 2008).

Although response times are not the primary variable of interest in categorization research, we believe that they can provide additional information about the nature of category representations and should be used more often in the evaluation of computational models of category learning. One problem, for example, with evaluating models using choice proportions is that it becomes difficult to study participants that have received extensive training. In long learning tasks, like the one used by Nosofsky and Palmeri (1997), participants will often reach expertise and have, by the end of training, classification proportions that are close to ceiling for the training examples. A clear advantage of working with response times is that, even after extensive training, they show variation that is likely to be linked to the categorization processes and representations. Therefore, response times can offer insights in the category representation of the expert engaging in the task.

Being able to investigate abstraction after extensive learning is of particular theoretical importance, since one key factor that is thought to influence abstraction seems to be the time point in learning (Smith and Minda, 1998). Further, when models of category learning are applied outside a lab context to semantic concepts, such as *fruits* or *vehicles*, relying on choice proportions is very difficult, since most people are experts for these categories. Models such as VABRW escape this conundrum, because they can account for response times. In this capacity, VABRW has the promise to provide fruitful insights in the representation of semantic concepts.




Figure 3: The best category representations for the three participants.

Acknowledgments

This work was partly supported by research grant G.0513.08 funded by the Research Foundation – Flanders (FWO), awarded to Gert Storms.

References

Ashby, F. G., Boynton, G., & Lee, W. W. (1994). Categorization response time with multidimensional stimuli. *Perception & Psychophysics*, 55, 11-27.

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98, 409-429.

Griffiths, T. L., Canini, K. R., Sanborn, A. N., & Navarro, D. J. (2007). Unifying rational models of categorization via the hierarchical Dirichlet process. In D. S. McNamara & J. G. Traffon (Eds.), *Proceedings of the 29th Annual Conference of the Cognitive Science Society* (pp. 323-328). MahWah, NJ: Erlbaum.

Hastie, T., Tibshirani, R., & Friedman, J. (2001). *The elements of statistical learning: Data mining, inference, and prediction*. New York: Springer.

Lamberts, K. (2000). Information-accumulation theory of speeded categorization. *Psychological Review*, 107, 227-260.

Logan, G.D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95, 492-527.

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

Minda, J. P., & Smith, J. D. (2001). Prototypes in category learning: the effects of category size, category structure, and stimulus complexity. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 27, 775-799.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115(1), 39-57.

Nosofsky, R. M. (1992). Exemplars, prototypes, and similarity rules. In A. F. Healy, & S. M. Kosslyn (Eds.), *Essays in honor of William K. Estes, vol. 1: From learning theory to connectionist theory; vol. 2: From learning processes to cognitive processes* (pp. 149-167). Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc

Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, 104, 266-300.

Nosofsky, R. M., & Stanton, R. D. (2005). Speeded classification in a probabilistic category structure: contrasting exemplar-retrieval, decision-boundary, and prototype models. *Journal of Experimental Psychology: Human Perception and Performance*, 31, 608-629.

Rosseel, Y. (2002). Mixture models of categorization. *Journal of Mathematical Psychology*, 46, 178-210.

Smith, E. E. & Medin, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.

Smith, J.D., & Minda, J. P. (1998). Prototypes in the mist: the early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 24, 1411-1436.

Vanpaemel, W. (2011). Constructing informative model priors using hierarchical methods. *Journal of Mathematical Psychology*, 55, 106-117.

Vanpaemel, W. & Storms, G. (2008). In search of abstraction: the varying abstraction model of categorization. *Psychonomic Bulletin and Review*, 15, 732-749.

Vanpaemel, W., & Storms, G. (2010). Abstraction and model evaluation in category learning. *Behavior Research Methods*, 42, 421-437.

Verbeemen, T., Vanpaemel, W., Pattyn, S., Storms, G., & Verguts, T. (2007). Beyond exemplars and prototypes as memory representations of natural concepts: a clustering approach. *Journal of Memory and Language*, 56, 537-554.