Learning time-varying categories

Daniel J. Navarro · Andrew Perfors ·
Wai Keen Vong

Published online: 20 April 2013
© Psychonomic Society, Inc. 2013

Abstract Many kinds of objects and events in our world have a strongly time-dependent quality. However, most theories about concepts and categories either are insensitive to variation over time or treat it as a nuisance factor that produces irrational order effects during learning. In this article, we present two category learning experiments in which we explored peoples' ability to learn categories whose structure is strongly time-dependent. We suggest that order effects in categorization may in part reflect a sensitivity to changing environments, and that understanding dynamically changing concepts is an important part of developing a full account of human categorization.

Keywords Categorization · Memory

At no two moments in time are we presented with the same world: Objects move, plants and animals are born and die, friends come and go, the sun rises and sets, and so on. More abstractly, while some of the rules that describe our world—like physical laws—are invariant over the course of our everyday experience, others—like legal rules—are not. Given some appropriate time scale, certain characteristics of an entity or class of entities can change; moreover, they may tend to change in systematic ways. For instance, the features that describe phones have changed considerably over recent decades: Not only do modern phones perform many new functions, they are also physically smaller, sleeker, and smoother. Not surprisingly, people's expectations about category members change to suit the environment as it stands: If asked to describe a phone in 2012, few

people would refer to a rotary dial, but in 1970, nearly everyone would.

In one sense, nothing is surprising about this observation. However, the changeable nature of many of the concepts and categories with which humans must interact has not been greatly emphasized in the categorization literature (but see Elliott & Anderson, 1995). In category-learning experiments, it is generally assumed that the underlying category is more or less static, and as such, the order in which one encounters category members should not matter to a rational learner. In statistics, this is referred to as the *assumption of exchangeability*, and for reasons of simplicity, it is generally the default assumption. Current probabilistic models of categorization make this assumption quite explicitly (e.g., Griffiths, Sanborn, Canini, & Navarro, 2008; Sanborn, Griffiths, & Navarro, 2010), and to the extent that standard exemplar and prototype models of categorization can be viewed as kinds of probabilistic models, they can also be seen to abide by this assumption (Ashby & Alfonso-Reese, 1995; Griffiths et al., 2008).

Perhaps because exchangeability is assumed in most real-world data analysis, it is generally taken to be a normative standard. However, human learners are also sensitive to the order in which category members are observed; this sensitivity appears to violate this normative standard. One way to account for order effects in cognitive models is to use learning rules that are sensitive to them. Some such rules can be viewed as modifications to standard probabilistic models. For instance, highlighting effects can be accounted for by assuming that people follow a “locally Bayesian” learning rule (Kruschke, 2006), whereas primacy effects can be captured by using a particle filtering learning rule (Sanborn et al., 2010). Another approach is to adopt connectionist, error-driven learning rules, which implicitly assume that recent items are more salient, and so are able to capture some kinds of recency effects, often better than a

D. J. Navarro (✉) · A. Perfors · W. K. Vong
School of Psychology, University of Adelaide, Adelaide, South
Australia 5005, Australia
e-mail: daniel.navarro@adelaide.edu.au

simple recency-weighting strategy would (e.g., Nosofsky, Kruschke, & McKinley, 1992; Sakamoto, Jones, & Love, 2008). A third approach is to alter the underlying stimulus representation: For instance, certain recency effects can be captured by the assumption that people track the differences between successive observations (Stewart, Brown, & Chater, 2002).

Although all of these approaches endeavor to account for order effects, it remains difficult to say whether it is rational to be sensitive to stimulus order. Many studies have avoided any explicit discussion of whether this sensitivity should be called normative (e.g., Stewart et al., 2002), others have argued that it reflects the cognitive limitations of the human learner (Sakamoto et al., 2008; Sanborn et al., 2010), and still others have suggested that order-sensitive learning rules are necessary if the learner is to be able to adapt to a changing world (Elliott & Anderson, 1995; Nosofsky et al., 1992). The latter perspective is mirrored rather explicitly in the literature on sequential effects (Yu & Cohen, 2009) and change detection (Brown & Steyvers, 2009).

Regardless of their views on the rationality of order effects, the literature shows a great deal of uniformity; in particular, the models that best capture human performance weight the observations by their recency. In most models, the weight assigned to a particular observation tends to decay approximately exponentially as a function of age (Brown & Steyvers, 2009; Nosofsky et al., 1992; Yu & Cohen, 2009), although in some cases the decay has been proposed to be a power function (Elliott & Anderson, 1995), which would be more in keeping with the literature on memory (Rubin, Hinton, & Wenzel, 1999; Wixted & Ebbesen, 1997). A similar outcome exists in the judgmental forecasting literature, which examines how people perceive and extrapolate time series data: Once again, an exponential weighting rule appears to account for human performance (see Goodwin & Wright, 1994, and Lawrence, Goodwin, O'Connor, & Onkal, 2006, for overviews). An exponential weighting scheme also emerges from the literature on stimulus generalization (Shepard, 1987).

The combined weight of this work suggests that weighting more recent items according to an exponential function is both theoretically and empirically justifiable. In fact, if the world changes at unpredictable times and in arbitrary ways, an exponential weighting scheme is close to optimal (e.g., Yu & Cohen, 2009). However, the world does not always change in an unpredictable fashion. For instance, the changes to the category of “phone” have been at least partially predictable: Newer phones tend to be faster, smaller, and more technologically capable. In this example, at least, it is clear that people have a strong expectation that the category will change systematically and in a particular direction. This ability to extrapolate the direction of future change cannot be explained by assuming that sensitivity to

order emerges simply from weighting more recent items more highly. A recency explanation would correctly predict that, for example, the iPad 4 would be more similar to the iPad 3 than to the iPad 2. What it does not predict, however, is that people expect the iPad 4 to be systematically different from (e.g., faster than) both. In order to capture this effect, we need to move beyond simple recency toward an explanation based on the idea that people detect trends and extend them.

The goal of this article is to address these open questions by investigating how people learn categories that change in a systematic fashion over time. We began by presenting an experiment that involves a simple “linear change” pattern, using a standard supervised classification task. We then followed this up with an experimental design that required participants to generate new category members. In both experiments, we found evidence that participants were sensitive to a systematic pattern of change in the observations that they were shown.

Experiment 1

In this section, we present a category-learning experiment in which people were presented with fairly obvious and systematic temporal changes, in order to investigate how well people would learn to anticipate those changes.

Method

Participants A group of 59 participants were recruited through a mailing list whose members consist primarily of current and former undergraduate psychology students. They were paid \$10/h for their time. The median age was 23, and the participants were predominantly (63%) female.

Materials and procedure The learning task was a standard supervised classification experiment, performed on a computer. The stimuli were little cartoon objects (“floaters”), which were displayed floating above a horizontal line (“the ground”). The height of the floaters was the only respect in which the stimuli varied from each other. An example of what a floater looked like is shown in Fig. 1.

Fig. 1 Example of the “floater” stimuli used in Experiment 1

On each of 100 trials, the participants were shown a single floater and asked to predict whether it would flash red or blue. After making their prediction, they would receive feedback for 2 seconds while the floater flashed the appropriate color. As the left panel of Fig. 2 illustrates, as the experiment progressed, all of the stimuli shown to people tended to rise, regardless of which category they belonged to. In the figure, black circles correspond to items that belonged to the “high” category, and white circles correspond to floaters that belonged to the “low” category. Assignment of the flash color (red or blue) to each category (high or low) was randomized across participants.

For the purposes of our analysis, we refer to the average rise (approximately 2 mm) as 1 unit. The reason for doing this is that it allows us to write the true classification rule in a very simple form. Specifically, the classification rule was such that, if x_t denotes the height of the stimulus on trial t , the optimal response is to select the response option corresponding to the high category if $x_t > t$. Such a rule, shown as the solid line in Fig. 2, achieves 100% accuracy on the task. However, because most stimuli tended to lie quite close to the classification boundary, the task was relatively difficult, even though the general trend was clear. Consistent with this, during informal discussions, the participants indicated that they detected the upward trend early in the experiment but still found the task to be quite challenging.

A model for the task Our data analysis relied on a simple categorization model. The model was inspired by decision bound models (e.g., Ashby & Gott, 1988; Ashby & Lee, 1991; Ashby & Maddox, 1993), although unlike most decision bound models, it was not explicitly derived from general recognition theory. In this section, we describe the structure of the model, because it is central to our analysis. It is worth noting, however, that we tried a variety of other categorization models, and the qualitative pattern of the results was not affected.¹

Recall from Fig. 2 that the stimuli were designed to be approximately normally distributed with mean μ_t , where μ_t increased linearly over time. Moreover, if a particular stimulus lay above the mean (i.e., $x_t > \mu_t$), it belonged to the high category, and otherwise it belonged to the low category. In other words, tracking the category boundary over time was equivalent to tracking the value of μ_t over trials. A simple model for estimating the value of μ_t is as follows. Suppose that the learner has some estimate $\hat{\mu}_{t-1}$ of the location of the

category boundary before the start of trial $t - 1$. When the stimulus x_{t-1} is observed, the boundary is shifted by some proportion w in the direction of that observation, yielding the following estimate for the location of the category boundary before the start of trial t :

$$\hat{\mu}_t = (1 - w)\hat{\mu}_{t-1} + wx_{t-1} \quad (1)$$

where w is a “twitchiness” parameter that indicates the extent to which the learner relies on the very last observation that he or she has seen.² Expanding the recursion in Eq. 1, we observe that this model produces an estimate $\hat{\mu}_t$ that is an exponentially weighted average of the previous items:

$$\hat{\mu}_t = w \sum_{k=1}^t (1 - w)^{k-1} x_{t-k} \quad (2)$$

In this equation, the fictitious “zero-th stimulus” x_0 corresponds to an initial value for the category boundary. The key thing to note in this equation is that recent trials will contribute more heavily to the estimate of μ_t ; large values of w imply that only a few observations are used, and small values of w allow multiple observations to be used. Having formed an estimate of where the category boundary lies, the learner is assumed to generate responses probabilistically, as a function of the deviation fit between the current stimulus and the estimated category boundary. This deviation is given by

$$\Delta_t = x_t - \hat{\mu}_t \quad (3)$$

and if $\Delta_t > 0$, then the item is more likely to be classified as a member of the high category. Specifically, we assume that the function relating distance to categorization probability is logistic (e.g., Navarro, 2007). If p_t denotes the probability of selecting the high category on trial t , then

$$p_t = \frac{1}{1 + \exp(-\lambda \Delta_t)} \quad (4)$$

where λ governs the rate at which the classification probability changes as a function of distance from the category

¹ For instance, we also tried prototype models, exemplar models, and a range of possibilities in which the learner estimates a linear or nonlinear regression function to describe how the category changes over time. Most of these models are outlined in Navarro and Perfors (2012). The key point is that the substantive story is robust: As long as we use a model that can fit the empirical data (not all of them can), we find a systematic but weak effect in the hypothesized direction.

² It may not be obvious on the surface how the model incorporates feedback, since Eq. 1 does not explicitly specify a training signal. To illustrate how feedback is used, consider an error-driven learning model that adjusts its estimate of the category boundary every time the feedback indicates that an error was made. To learn from this feedback, the model needs a training signal. The true error here is $\mu_{t-1} - \hat{\mu}_{t-1}$, the deviation between the true and estimated category boundaries, but this is not observable to the learner. However, the fact that x_{t-1} was misclassified implies that the magnitude of the error is at least $x_{t-1} - \hat{\mu}_{t-1}$. This lower bound on the error is observable and can be used as a training signal. If the learner adjusts the estimated category boundary by some proportion w of this error, we obtain $\hat{\mu}_t = \hat{\mu}_{t-1} + w(x_{t-1} - \hat{\mu}_{t-1})$, which is identical to Eq. 1. In other words, if Eq. 1 is applied only on error trials, it is equivalent to a standard, error-driven learning model. We did consider using this restricted model, but felt that the full model (which learns from errors and from correct decisions) would be simpler.

Fig. 2 Experimental design. Black circles denote stimuli belonging to the “high” category, and white circles denote stimuli belonging to the “low” category. Although the classification rule changes over time—that is, the classification boundary is constantly rising, as is shown in the left panel—it does so in a regular fashion. The scale on the vertical axis is normalized so that the average “rise” from one trial to the next is

1 unit; onscreen, this corresponded to an average rise of approximately 2 mm per trial. Note that although it was logically possible to correctly classify all items, in practice the task was quite difficult, since most stimuli lay close to the boundary. The deviations of each stimulus from the true classification boundary are shown on the right

boundary. This model is closely related to the tracking model used by Brown and Steyvers (2009), but it has links to other models, too. For instance, in the extreme case in which $w = 1$, this heuristic corresponds to a relative judgment strategy in which each stimulus is compared only to the last stimulus in the experiment, and it becomes a slight simplification of the model used by Stewart et al. (2002). When $w = 1$ and λ is large, the model produces a very simple heuristic: Select the high category if and only if the current stimulus is higher than the previous one. Additionally, the model has links to prototype models of classification. In an equal-variance prototype model, the learner represents separate means (the prototypes) for each category, and the category boundary lies equidistant from the two prototypes. If the estimate for a category prototype takes recency into account by taking an exponentially weighted average (see, e.g., Navarro & Perfors, 2012), the classification probabilities will end up being almost identical to those produced by our model. The problem with using a heuristic such as this one is that it is very sensitive to random fluctuations in the data (when w is large), or else the estimate of the category boundary lags a long way behind the true one (when w is small). To see this, note that when w is large, the learner is strongly influenced by the most recent observation, and to the extent that this observation is noisy, or otherwise misleading as to the location of the category boundary, the learner will be unduly influenced by it. Decreasing the value of w allows the learner to avoid this mistake by aggregating information from multiple observations, but this comes at a price: Because the category is moving, setting w to too small a value means that the estimate $\hat{\mu}_t$ will always be “lagging” a long way behind the data.

This issue can be avoided to some extent if the learner is able to detect the pattern of change and anticipate the fact that the category boundary μ_t moves on each trial. For the sake of simplicity, we assume that this corresponds to the introduction of a simple correction factor b . This correction factor yields a slight modification of the model, in which the estimate of the category boundary on trial t is given by

$$\hat{\mu}_t = b + w \sum_{k=1}^t (1 - w)^{k-1} x_{t-k} \quad (5)$$

We now have a simple classification model with three parameters: twitchiness w that governs how reliant the learner is on the last stimulus; the slope parameter λ that describes the relationship between distance and generalization; and the bias parameter b that describes a correction factor, shifting the classification boundary to accommodate the fact that the task involves a clear trend over time. Although this model has not (to our knowledge) been used in the categorization literature previously, it has been used in the judgmental forecasting literature (see Goodwin & Wright, 1994) as a heuristic that is appropriate for modeling human extrapolation judgments for a trended time series.

Results

Human performance was significantly above chance for both categories: 76% of the high-category items and 61% of the low-category items were classified correctly, as is shown in the left panel of Fig. 3. Note that, while performance in both categories was significantly above chance, people performed better in the high category. This is not surprising: Because both categories were rising throughout

the experiment, novel items from the low category tended to be close to the previous high-category items. By contrast, new items from the high category were much less confusable, since they were not close to previous low-category items.

When we plot the performance of all 59 participants separately, as in the right panel of Fig. 3, it is clear that the improved performance on the high-category items holds at the individual-participant level, as well. Nine of the 59 participants were not significantly above chance (the $p < .05$ significance threshold is plotted as a dashed line). Of the remaining 50 participants, 47 classified the high-category items more accurately than the low-category items while performing above 50% correct for both categories. This is illustrated visually by the fact that the vast majority of the dots in Fig. 3 fall within the solid black triangle.

Model fits were performed at the individual-participant level, obtaining separate values of b , w , and λ for all 59 participants by using maximum likelihood estimation to estimate the parameter values. The first five trials were excluded for the purpose of the model fitting.³ As is shown in Fig. 4, the model produces classification behavior very similar to that of the human participants. Across participants, the correlations between the human data and the model predictions were $r = .94$ ($p < .001$), for the probability of correctly classifying low-category items, and $r = .92$ ($p < .001$), for the high-category items. At a within-participant level, the model fit was significantly better than chance (as assessed via a likelihood ratio test at the $p < .05$ level) for 52 of the 59 participants. Of the seven participants whose data were not well-fit by the model, five were among those who did not classify the stimuli above chance. It is no surprise that the model cannot account for the performance of those participants, as the empirical data from those participants are extremely noisy.

The key empirical test here related to the parameter estimates, most notably of the bias parameter b . Given this, we restricted the analysis to those participants whose data were well-fit by the model. Our inclusion criterion here was that the model needed to explain at least 25% of the variance in a participant's choices. This criterion was met for 43 of the participants. Among these participants, the model explained 43% of the variance on average ($SD = 12\%$), corresponding to a 71% probability that the model would make the same response as the participant on any given trial ($SD = 6\%$).

³ This decision was motivated by the intuition that the first few trials would be “special” and not easily modeled. The intuition was also backed up by additional analyses using generalized linear mixed models: Including a random-effect term for trial number revealed that this was true for Trial 2 in particular. Given this, it seemed prudent to exclude the first few trials.

Fig. 3 Overall performance (left) with 95% confidence intervals, as well as the individual participants' data (right). Both plots incorporate data from all participants and all trials. Note that nine participants did not perform significantly above chance (the dashed line corresponds to the $p < .05$ significance threshold), and for this reason were excluded from the subsequent analyses

Descriptive statistics for the parameter estimates are provided in Table 1. From a theoretical perspective, the first analysis to consider was a test of whether the b parameter is in fact necessary within the model. To that end, for all 43 participants, we fit a restricted model with the bias parameter fixed at $b = 0$; a likelihood ratio test rejected this null model (at $p < .05$) in all 43 cases, favoring the model that included bias. A second analysis to consider was to look at the magnitude of the bias parameter. The mean value of b was 1.07, with a 95% confidence interval of [0.28, 1.87], implying that participants did shift their classification boundaries upward.

Closer examination revealed a somewhat more complicated story. There was a moderately strong negative correlation of $r = -.77$ ($p < .001$) between the bias parameter b and the twitchiness parameter w (this was the only significant correlation). This was to be expected: As noted earlier, smaller values of w mean that participants are aggregating information across more trials, which in turn implies that the uncorrected estimate of the category boundary will lag farther behind the true boundary. In other words, the optimal value of b should be higher for smaller values of w . This is illustrated in Fig. 5, which plots b against w for all participants (black dots), along with a regression line (solid line) that depicts the relationship between the two. The fact that the confidence bands for the regression line (gray region) sit above zero (dotted line) for most values of w indicates that, in general, participants were extrapolating.

The fact that people extrapolate does not imply that the extent of the extrapolation is optimal. Indeed, Fig. 5 also

Fig. 4 Model performance at an individual-participant level. As in Fig. 3, each participant is characterized in terms of the probabilities of correctly classifying low-category items (left panel) and high-category items (right panel). Each panel plots each participant as a circle, showing the model probability correct against the human probability correct

shows the optimal value of b for all values of w (dashed line). It is clear that the participant responses lie below the optimal value, indicating that they did not extrapolate far enough. This is consistent with the raw data in Fig. 3: The simple fact that people are more accurate for high-category items than for low-category items provides strong evidence that they do not entirely anticipate the extent of the changes across trials. Moreover, when we plot the mean category boundary (taken across participants) extracted from the model, as per Fig. 6, it is clear that the model reproduces this behavior.

Discussion

The results from Experiment 1 provided some indication that people are capable of detecting changes to a category over time and are able extrapolate a trend that they have observed when making decisions about new items. However, the effect is subtle and relies on the assumption that, if the categories had not been changing over time, participants would have used a recency-weighted average to estimate the category representation (i.e., the model without b). This assumption seems sensible, insofar as the task was a standard supervised categorization task, and models of that kind have been highly successful in explaining human behavior in such tasks. Nevertheless, it is desirable to show that the effect can be detected without requiring detailed modeling, if we allow the task to be modified to capture the effect more directly.

Table 1 Means and standard deviations for model parameters, estimated from the data of all 43 participants for which the model fit was deemed adequate

	Mean	SD
Bias, b	1.07	2.58
Twitchiness, w	0.27	0.11
Slope, λ	0.56	0.30

Experiment 2

One of the problems with the supervised classification task used in Experiment 1 is that the participants provided only a very limited amount of data (a binary choice) on any particular trial. This made it a little difficult to measure very subtle effects related to how category representations changed at a trial-by-trial level. To redress this, in Experiment 2 we employed a very different dependent measure: The participants were asked to generate new category members on each trial. After training participants on a category that changes over time, we asked them to generate a sequence of future category members over multiple time points. If participants were genuinely able to extrapolate their category knowledge, this sequence of responses should extend the trend that had appeared in the training data.⁴

Method

Participants A total of 110 participants (34 female, 76 male) were recruited via Amazon Mechanical Turk, an online service that coordinates people to complete tasks requiring human intelligence. The participants were located in 17 different countries, but the vast majority were in India (71) or the United States (17). They were required to be English speakers (assessed via a series of simple test questions) and were paid US\$0.25 for their participation in the study. The structure of the experiment allowed a fairly straightforward way to check whether participants had understood the task and were making a genuine attempt to complete it: examination of their responses. Exclusions are thus discussed together with the data analysis.

⁴ This experiment formed part of a larger study that made up the third author's Honours thesis.

Fig. 5 Scatterplot (black dots) showing the relation of the twitchiness parameter w and the bias parameter b , along with the linear regression (solid line) and corresponding 95% confidence band. The fact that the regression line is above zero (dotted line) for most values of w indicates that, in general, participants were extrapolating. However, the regression line also lies below the optimal value of b (dashed line) for most values of w , indicating that participants did not extrapolate far enough

Materials and procedure Participants completed the study online through the Amazon Mechanical Turk website. After accepting the task, they were asked demographic questions and then presented with the experiment instructions. They were then quizzed on the experimental procedure, to check that they had understood the experimental instructions (this implicitly served as a check that they understood English). If they did not answer all questions correctly, they were redirected back to the instructions and required to repeat the instruction checks until all of the questions were answered correctly. To verify that any effect observed did not depend on particular details of the stimulus or task, four versions of the task were employed. Participants were randomly allocated to different groups for the Stimulus Type and Prediction Type factors, yielding a 2×2 design. In the stimulus type

Fig. 6 Estimated category boundary on each trial (solid line), expressed in terms of the deviation from the true category boundary (dashed line), and averaged across all 43 participants whose data were well-fit by the model. The fact that the solid line lies below the dashed line most of the time indicates that participants' implied category boundaries still lagged behind the true rule

manipulation, we altered the surface representation of the stimuli. The prediction type manipulation focused on the kinds of responses that people were asked to provide.

In the cover story for the task, the participants were asked to imagine that they were helping a confectionery company to determine the future direction of some of their products. The experiment involved two phases for all conditions. Phase 1 consisted of ten trials, each of which corresponded to a different year: On each trial, participants were shown a set of ten stimuli, intended to represent different examples of confectionery on sale that year. All ten items were displayed onscreen together. The visual representation of the confectionery depended on the stimulus type: In the candy cane condition, the stimuli were displayed as striped horizontal bars of different lengths. The length of the bar was the relevant stimulus dimension. In the chocolate condition, the stimuli were boxes of chocolates in which each chocolate could be either white or dark. Each box showed a 4×5 grid of 20 chocolates, and the number of dark chocolates was the relevant stimulus dimension. Example stimuli are shown in Fig. 7.

As in Experiment 1, the stimuli changed systematically across trials, as illustrated by the dashed lines in Fig. 8. In the candy cane condition, the mean length of the stimuli was 300 pixels on Trial 1, which rose linearly to 525 pixels by Trial 10, corresponding to an average increase in length of 25 pixels per trial. In the chocolates condition, the mean number of dark chocolates in the boxes started at three on Trial 1, rising by one chocolate per trial until the average number of chocolates on Trial 10 was 12. On any given trial, however, participants were shown ten stimuli that varied around this mean. In the candy cane condition, the length of any one stimulus was uniformly sampled from a range of ± 25 pixels around the true mean. In the chocolate condition, the number of dark chocolates in any one box was equal to the true mean plus -1 , 0 , or 1 , with each outcome being equally likely. In our data analysis, however, all stimulus magnitudes were rescaled so that the mean magnitude on Trial 1 was 0 , and on Trial 10 it was 1 .

After being presented with a set of stimuli, people were asked to generate three new examples of confectionery.

Fig. 7 Examples of the stimuli used in Experiment 2. The image on the left shows a display containing three candy canes, and the image on the right shows a display containing two chocolate boxes. In both cases, the actual displays contained ten stimuli

Fig. 8 Mean responses on each trial during the supervised phase. In the current-prediction condition, participants on trial t generated predictions for additional examples at trial t . In the future-prediction condition, participants on trial t generated predictions for examples

that they would expect to see at trial $t + 1$. One would therefore expect that, in most cases, the future predictions would be higher than the current predictions. This effect is evident for the candy cane but not for the chocolate stimuli

The task differed depending on the prediction type: In the current-prediction condition, participants were asked to generate three more examples that would be representative of the confectionery for the current year; in the future-prediction condition, the examples were supposed to represent predictions for the next year's confectionery. Participants who saw candy canes made their predictions by dragging a slider bar to change the length of the candy canes, and the participants who saw chocolate boxes were able to click on each individual chocolate piece to change its color from white chocolate to dark chocolate, or vice versa.

The critical part of the experiment was Phase 2, which was identical, regardless of what prediction type participants had been assigned to. During this phase, participants were asked to generate three examples of candy canes or chocolate boxes for each of the next three years (i.e., generating nine examples in total). During this phase, no information was given as to what kinds of candies or chocolate blocks would be popular for those years. Note that, because these unsupervised trials followed immediately from the supervised ones in Phase 1, we did not ask participants in the future-prediction condition to make any responses for the last trial of Phase 1. The reason for this was that the last trial of Phase 1 was essentially equivalent to the first trial of Phase 2 in the future-prediction condition, and we felt it more important to preserve the comparability of judgments in Phase 2.

Results

Because the stimuli increased across trials in a linear fashion, one would expect that the predictions generated by participants would also rise across trials. At a minimum,

one would expect the magnitude of the responses generated by participants in Phase 1 to correlate with the stimulus magnitude. This suggests a natural exclusion criterion: whenever the Pearson correlation between the stimuli and responses was less than or equal to zero, we had evidence that the participant either did not understand the task or failed to make a genuine attempt to do so. In total, 16 participants were excluded on this basis, including two participants who produced identical responses on every trial. After these exclusions were made, we ended up with 24 participants in the candy stimulus current-prediction condition, 29 in the chocolate stimulus current-prediction condition, 34 in the candy stimulus future-prediction condition, and 23 in the chocolate stimulus future-prediction condition. The mean responses on each supervised trial are plotted in Fig. 8, broken down by conditions (error bars denote 95% confidence intervals). These plots provide some evidence that participant responses differed between the conditions during Phase 1, but the important thing for our purposes was to note that in all conditions, the participants understood the task, correctly realizing that the stimulus magnitude was increasing over trials. Given that the primary interest lay in how participants would respond in Phase 2, this is the topic to which we now turn.

Mean responses and confidence intervals⁵ for the Phase 2 trials (i.e., extrapolations over a three-year range) are plotted in Fig. 9: On the left the data are shown aggregated across conditions, and on the right each condition is shown

⁵ The confidence intervals in the right panel of Fig. 9 are the intervals associated with each group mean, constructed in the usual way. Because the left panel collapses across groups with different means, the error bars in this case show the confidence interval associated with the grand mean in the corresponding 2×2 ANOVA (i.e., the confidence interval for the intercept term when Helmert contrasts are used).

Fig. 9 Data from Phase 2 of Experiment 2, showing future predictions made over three years, without further examples shown. The plots display the mean responses and 95% confidence intervals for participant responses in Phase 2 of the experiment. The left panel shows the mean response averaged across conditions, whereas the right panel

shows the responses broken down across all four conditions. As this figure indicates, participants' generalizations in Phase 2 reflected the linearly increasing trend that they had experienced in Phase 1 of the experiment

separately. Given that the conditions involved different stimuli, response methods, and task instructions, it is not entirely surprising to see that the baseline response in Phase 2 (i.e., the response on the first trial of Phase 2) differed across conditions. The key question of interest with respect to Phase 2, however, was not whether the four conditions would differ, but whether the three trials would. That is, would the responses show a significant rise from trial to trial? Would people extend the linear trend that they had observed in Phase 1 into the future?

To address this question, we fit a linear mixed effect model to the Phase 2 data using the lme4 package in R (Bates, Maechler, & Bolker, 2011). The model included fixed effects of prediction type, stimulus type, and trial number (i.e., year). Individual variation was captured by including a random intercept and a random effect of trial number (i.e., random slope) for each participant. The Bayesian information criterion (BIC) was used for model checking, testing each effect by comparing the full model to a restricted model in which the terms corresponding to that effect were removed. The BIC deteriorated substantially if we removed the random intercepts (BIC difference = 100.0) or the random slopes (BIC difference = 14.3). Removing the fixed effect of prediction condition improved model performance slightly (BIC difference = −5.0), whereas removing the fixed effect of stimulus condition caused it to decline slightly (BIC difference = 6.1). Most importantly, the effect of removing the fixed effect of trial caused a large deterioration (BIC difference = 22.4). When we converted the BIC difference of 22.4 to approximate posterior odds, we observed that this corresponded to odds of 73,000:1 in favor of the model that included the effect of trial number. This suggests that participants were extrapolating sensibly over

the next three years, capturing the linear trend in the data they had been shown.

Having verified that participants' responses did show a significant rising trend in Phase 2 of the experiment, it was useful to consider how large that trend was. For simplicity, we considered the aggregated data plotted on the left-hand side of Fig. 9. The slope of this plot is 0.077. The linear trend in Phase 1 involved the stimuli rising (in rescaled units) from an average value of 0 on Trial 1 to an average of 1 on Trial 10. This corresponded to a rise of 1 unit in nine steps—that is, to a slope of 0.111. In other words, although participants did extrapolate a linear trend in Phase 2, the slope of the trend was shallower than the slope of the original trend in the data, with a slope that was only 70% as large as the empirical slope in Phase 1.

Discussion

The results of Experiment 2 are in agreement with those of the earlier experiment. In Experiment 1, we found evidence that people extrapolated their knowledge about categories in a manner that was consistent with the pattern of change in the task. However, the effect was subtle and relied on a detailed model-fitting exercise. By asking people to generate new category members that they would expect to observe several years into the future, we were able to observe the effect more directly. As is illustrated in Fig. 9, across all four conditions, participants' judgments showed a linearly rising trend when they were asked to make predictions for the next three years, even in the absence of any explicit feedback. Similarly, just as in Experiment 1, we found that the extent of the extrapolation was suboptimal: People did the right thing qualitatively, but not to the extent required to

perfectly match the trend in the raw data. Taken together, the two experiments tell a fairly consistent story about how people make sense of categories that change over time.

General discussion

The world is not a static place: The rules and regularities that characterize our environments do not remain constant over time. Sensitivity to change is a necessary characteristic of any intelligent learner operating in a changing world. As trivial as this sounds, it has important implications: Any model that is invariant to stimulus ordering, such as the original generalized context model (GCM; Nosofsky, 1984), standard decision bound models (Ashby & Lee, 1991), or the statistical model underpinning the “rational” model of categorization (Sanborn et al., 2010), is fundamentally the wrong model for some kinds of categorization problems that humans need to solve in real life—namely, those that involve change over time. As such, models that are sensitive to order information, such as the GCM with recency weighting (Nosofsky et al., 1992) or models that rely on connectionist learning rules (e.g., Kruschke, 1992; Sakamoto et al., 2008), are better suited to this sort of category-learning problem. Similarly, to the extent that learning about change is a key task facing the learner, directly encoding stimulus differences (as per Stewart et al., 2002) may be viewed as a rational thing to do. In some respects, these models are closer to the correct rational analysis of this learning problem than are most “rational” categorization models.⁶

This work shows that once we start to think of category learning as a problem that applies in a changing environment, there is much more to the problem than simply assigning more importance to recent observations. When the pattern of change is systematic rather than arbitrary, it is not sufficient merely to detect changes (e.g., Brown & Steyvers, 2009): An ideal learner is able to anticipate them by forming sensible expectations about when the world changes and how that change occurs. Although our experiments are very simple examples of this, involving regular linear change in one dimension, they provide evidence that people do form these sorts of expectations during category learning.

This research also opens up questions as to the mechanisms by which people actually form these expectations. The model that we used in Experiment 1 does not provide an answer to this question. We do not advocate any theory of

how the bias parameter is learned, although it seems likely that, since $b > 0$ produces superior classification performance in this task, almost any sensible learning rule (e.g., error-driven learning or Bayesian updating) would infer a positive value. Indeed, it seems likely that such models would infer the *optimal* value for b , which would open up the question of why the human extrapolations were suboptimal in both experiments. More generally, we are not convinced that a simple “bias” is the right answer to the problem in general: The model from Experiment 1 would be highly inappropriate for a category that changed in a nonlinear fashion, for instance. These questions are left open for future work.

Although the model in Experiment 1 was used primarily as an aid to the data analysis, the bias parameter b within that model is perhaps the most explicit attempt to formalize the notion of extrapolation in this study. As such, it is worth giving some consideration to what b actually does. Viewed purely as a mathematical device, the effect that b has on the estimated category boundary is identical to the effect of a category base rate: It moves the boundary closer to one 684 category or the other. Given this, one might wonder whether the bias learned by participants in Experiment 1 was acquired via a base-rate learning mechanism. If so, one might then question whether any genuine extrapolation has occurred. However, this concern is unfounded: It conflates the learning mechanism with the substance of what is learned. In the experiments presented in this article, the base rates for both categories were always 50%. If the learned bias b were equivalent to base-rate learning, one would expect an optimal value of $b = 0$ in any experiment in which the categories did not differ in base rate. Yet, as Fig. 5 shows, the optimal value of b is decidedly nonzero. Even if the mechanism used to learn b were formally equivalent to a base-rate learning rule, the thing that was learned would clearly not be a base rate. Rather, it would be a bias term whose substantive effect would be to systematically shift the learners’ expectations about future events away from their experience of past events. “Extrapolation” seems as good a name as any for this inference.

To sum up, this article illustrates that people are capable of tracking changes over time during category learning. This work also opens up a range of additional issues. Regular linear change is only one pattern of change that could characterize real-world categories. In our initial work on this topic (Navarro & Perfors, 2009), we also considered the possibility that changes could be discrete jumps, as in Brown and Steyvers (2009), or sinusoidal patterns, as in Sakamoto et al. (2008). However, these possibilities are not exhaustive, either. Understanding how people adapt to systematic changes in categories will require a broader investigation across a wider range of possible change patterns.

⁶ In fact, although we do not discuss the topic in this article, it is clear that models that lack sensitivity to change will perform extremely poorly when applied to experiments such as the one here. This issue was discussed in the conference paper that this article extends (Navarro & Perfors, 2012).

Author note Portions of this work were presented at the 2009 and 2012 annual conferences of the Cognitive Science Society. Early work on this project, including salary support for D.J.N., was supported through ARC Grant No. DP0773974. Later work was supported through ARC Grant No. DP110104949, with salary support for D.J.N. provided by ARC Grant No. FT110100431, and for A.P. through ARC Grant No. DE120102378. We thank Natalie May and Yiyun Shou for helping run the experiments.

References

- Ashby, F. G., & Alfonso-Reese, L. A. (1995). Categorization as probability density estimation. *Journal of Mathematical Psychology*, 39, 216–233.
- 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. doi:10.1037/0278-7393.14.1.33
- Ashby, F. G., & Lee, W. W. (1991). Predicting similarity and categorization from identification. *Journal of Experimental Psychology: General*, 120, 150–172.
- Ashby, F. G., & Maddox, W. T. (1993). Relations between prototype, exemplar, and decision bound models of categorization. *Journal of Mathematical Psychology*, 37, 372–400. doi:10.1006/jmps.1993.1023
- Bates, D., Maechler, M., & Bolker, B. (2011). *lme4: Linear mixed-effects models using s4 classes* (R package version 0.999375-42) [Computer software manual]. Available from <http://CRAN.R-project.org/package=lme4>
- Brown, S. D., & Steyvers, M. (2009). Detecting and predicting changes. *Cognitive Psychology*, 58, 49–67.
- Elliott, S., & Anderson, J. R. (1995). Effect of memory decay on predictions from changing categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 815–836.
- Goodwin, P., & Wright, G. (1994). Heuristics, biases and improvement strategies in judgmental time series forecasting. *International Journal of Management Science*, 22, 553–568.
- Griffiths, T. L., Sanborn, A. N., Canini, K. R., & Navarro, D. J. (2008). Categorization as nonparametric Bayesian density estimation. In M. Oaksford & N. Chater (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 303–328). Oxford, UK: Oxford University Press.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99, 22–44. doi:10.1037/0033-295X.99.1.22
- Kruschke, J. K. (2006). Locally Bayesian learning with applications to retrospective revaluation and highlighting. *Psychological Review*, 113, 677–699. doi:10.1037/0033-295X.113.4.677
- Lawrence, M., Goodwin, M., O'Connor, M., & Onkal, D. (2006). Judgmental forecasting: A review of progress over the last 25 years. *International Journal of Forecasting*, 22, 493–518.
- Navarro, D. J. (2007). Similarity, distance and categorization: A discussion of Smith's (2006) warning about "colliding parameters". *Psychonomic Bulletin & Review*, 14, 823–833.
- Navarro, D. J., & Perfors, A. (2009). Learning time-varying categories. In N. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 412–424). Austin, TX: Cognitive Science Society.
- Navarro, D. J., & Perfors, A. (2012). Anticipating changes: Adaptation and extrapolation in category learning. In N. Miyake, D. Peebles, & R. P. Cooper (Eds.), *Building bridges across cognitive sciences around the world: Proceedings of the 34th Annual Conference of the Cognitive Science Society* (pp. 809–814). Austin, TX: Cognitive Science Society.
- Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 104–114. doi:10.1037/0278-7393.10.1.104
- Nosofsky, R. M., Kruschke, J. K., & McKinley, S. C. (1992). Combining exemplar-based category representations and connectionist learning rules. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 211–233.
- Rubin, D. C., Hinton, S., & Wenzel, A. (1999). The precise time course of retention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 1161–1176.
- Sakamoto, Y., Jones, M., & Love, B. C. (2008). Putting the psychology back into psychological models: Mechanistic versus rational approaches. *Memory & Cognition*, 36, 1057–1065. doi:10.3758/MC.36.6.1057
- Sanborn, A. N., Griffiths, T. L., & Navarro, D. J. (2010). Rational approximations to rational models: Alternative algorithms for category learning. *Psychological Review*, 117, 1144–1167.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, 237, 1317–1323. doi:10.1126/science.3629243
- Stewart, N., Brown, G. D. A., & Chater, N. (2002). Sequence effects in categorization of simple perceptual stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 3–11.
- Wixted, J. T., & Ebbesen, E. B. (1997). Genuine power curves in forgetting: A quantitative analysis of individual subject forgetting functions. *Memory & Cognition*, 25, 731–739. doi:10.3758/BF03211316
- Yu, A. J., & Cohen, J. D. (2009). Sequential effects: Superstition or rational behavior? In D. Koller, D. Schuurmans, Y. Bengio, & L. Bottou (Eds.), *Advances in neural information processing systems 21* (pp. 1873–1880). Cambridge, MA: MIT Press.