Discontinuities in function learning

Roderick Yang Terng Seow (yseow@andrew.cmu.edu)
John R. Anderson (ja0s@andrew.cmu.edu)
Department of Psychology, Carnegie Mellon University, 5000 Forbes Ave
Pittsburgh, PA 15213 USA

Keywords: function learning, exemplar, rule, abrupt, discontinuity

Introduction

Function learning is the process by which humans acquire knowledge of functional relationships between continuous variables. For example, a frequent beachgoer might visit the beach on different nights and come to associate specific tide heights with specific moon phases. With experience, the beachgoer might then abstract an underlying functional relationship: the tide rises approaching the full moon, and lowers approaching the new moon.

Most theories of function learning largely focus on two types of models: exemplar-based and rule-based models. Exemplar-based models posit that humans learn to associate exemplar cues with their respective targets via error-driven updates of associative weights (Busemeyer et al., 1997). Rule-based models posit that humans instead begin with some parametric function and learn its coefficients through an error-driven update mechanism (e.g. polynomial rule model: Koh and Meyer (1991)). More recent studies have proposed hybrid models that combine associative learning with rules, and these models have been shown to better account for a wide range of function learning phenomena than their predecessors (e.g. EXtrapolation Association Model (EXAM): DeLosh et al. (1997); Population Of Linear Experts (POLE): Kalish et al. (2004)).

Despite their differences, existing process models mostly assume that function learning is a gradual and continuous process. In contrast, Brehmer (1974) proposed a two-staged hypothesis testing theory of function learning. The first stage involves discovering a suitable rule, and the second stage is concerned with learning the parameters of the rule. Although this theory has not been quantitatively formalized, it differs from the other theories by positing a discontinuity when the learner transitions from discovering a rule to applying a rule. In support of the role of rule discovery in human function learning, we present preliminary evidence of such discontinuities and demonstrate that existing process models do not adequately account for these observations.

Experiment

The experiment was a replication of study 1a of McDaniel et al. (2014), but with Amazon Mechanical Turkers instead of undergraduate students. 59 participants, 21 females, ages ranging from 20 to 53 years old (mean = 32.3), completed the experiment. Participants were paid \$4.50 for completion and an accuracy bonus up to \$0.02 on every training trial.

Participants completed 10 training blocks followed by 1 transfer block. Each training block consisted of the same 20 trials presented in random orders. For each training trial, the cue value was represented using the height of a colored bar, and participants made their predictions using arrow keys to adjust the height of a separate response bar. Feedback was presented in three forms: the response bar at the target height, an error score consisting of the numerical difference between the response and the target values, and an accuracy score computed as $100 - error^2$. Transfer trials consisted of novel cue values, both within (interpolation) and beyond (extrapolation) the range of training cue values. No feedback was provided during the transfer block.

Cues and targets were related through a V-shaped function. For $cue < 100$, $target = round(229.2 - 2.197 \cdot cue)$. For $cue \geq 100$, $target = round(2.197 \cdot cue - 210)$.

Detecting discontinuities

One potential behavioral correlate of rule discovery is an abrupt decrease in an individual's error rates as they proceed through the training phase. To detect such discontinuities if and when they occur, we fitted single- and double-function error curves for each participant. Error noise was assumed to be Poisson distributed and these functions specified how the error mean (λ) changed with trial number (t).

The set of single-function curves comprised a constant mean ($\lambda = c$), an exponentially decreasing mean ($\lambda = a \cdot e^{-b \cdot (t-1)} + c$), and a mean that decreased according to a power function with increasing number of trials ($\lambda = a \cdot t^{-b} + c$). The single-function curves were composed to create the set of double-function curves, with the restriction that the second function was a constant. All double-function curves required an additional change point parameter.

To determine if an individual demonstrated an abrupt decrease in error, we computed two measures. The first measure (ΔBIC) was the difference between the Bayesian Information Criterion of the best fitting single- and double-functions. A large and positive ΔBIC indicated that the error curve was much better fit by a double- than a single-function. To quantify abruptness if a transition exists, the second measure ($\Delta mean$) was the difference between the pre- and post-transition fitted means for the best fitting

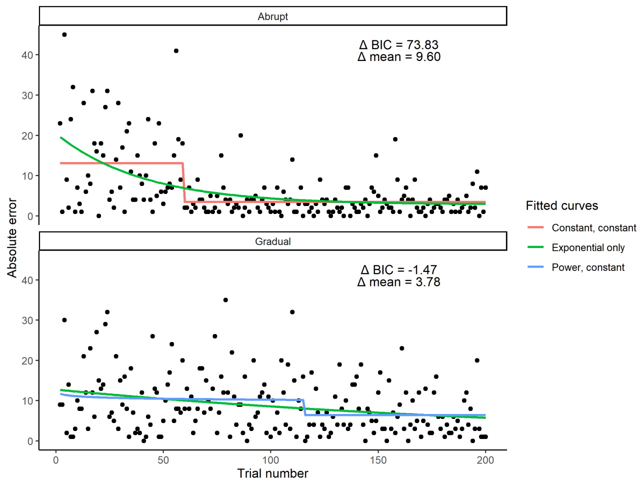

Figure 1: Error curves for two participants. Colored lines show the best fitting single- and double-function curves. The top panel displays an abrupt decrease in error, whereas the bottom panel displays a gradual decrease in error over time.

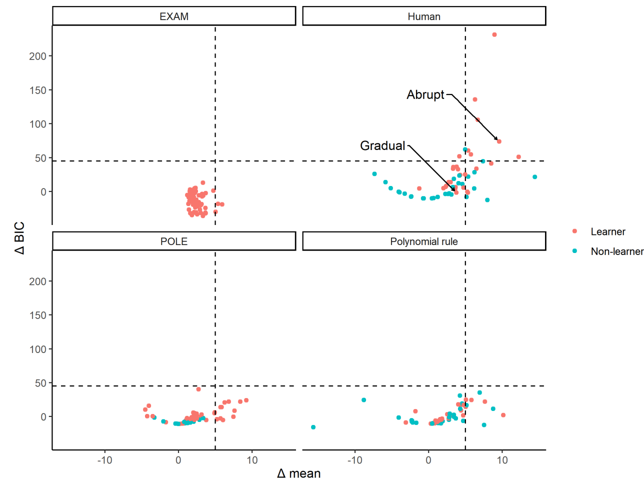

Figure 2: Individual participants and model simulations on abruptness measures. The two participants from Fig. 1 are labeled. Thresholds for the two measures are represented by the dashed lines. Individuals in the upper right quadrants of each panel are classified as abrupt learners.

double-function. A large and positive $\Delta mean$ indicated an abrupt decrease in error around the estimated change point.

Using these two measures, we classified participants into those who did and did not show abrupt learning. The 59 participants were first separated into 28 learners and 31 non-learners according to the criterion in McDaniel et al. (2014): learners were those who attained an average absolute error of less than 10 on the last training block. We then determined a combined threshold on our two measures. The threshold was chosen to be as inclusive as possible with the constraint that only learners could be classified as abrupt learners. This yielded a threshold of $\Delta BIC > 45$ and $\Delta mean > 5$. Based on the threshold, 7 out of 59 participants were classified as abrupt learners (Fig. 2).

Model comparisons

To generate individual model simulations, we found the best fitting set of parameters per participant for three process models (polynomial rule, EXAM, and POLE) by maximizing the log-likelihood with respect to the participant's responses. Applying the same classification procedure as above, none of the model simulations were classified as abrupt learners.

Conclusion

In this study, we identified a subset of participants that demonstrated abrupt decreases in error over the course of a function learning task. Our simulations of the existing process models confirmed that gradual update mechanisms cannot reproduce the observed discontinuities, which is consistent with the hypothesis that these discontinuities correspond to moments of rule discovery. To test this hypothesis, we are currently investigating the nature of rules and the role of rule discovery in human function learning.

Acknowledgments

This research was supported by ONR Grant N00014-15-1-2151 and AFOSR/AFRL award FA9550-18-1-0251.

References

Brehmer, B. (1974). Hypotheses about relations between scaled variables in the learning of probabilistic inference tasks. *Organizational Behavior and Human Performance*, *11*(1), 1–27.

Busemeyer, J. R., Byun, E., Delosh, E. L., & McDaniel, M. A. (1997). Learning functional relations based on experience with input-output pairs by humans and artificial neural networks. In K. Lamberts & D. R. Shanks (Eds.), *Knowledge, concepts and categories* (p. 408—437). The MIT Press.

DeLosh, E. L., Busemeyer, J. R., & McDaniel, M. A. (1997). Extrapolation: the sine qua non for abstraction in function learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*(4), 968–986.

Kalish, M. L., Lewandowsky, S., & Kruschke, J. K. (2004). Population of linear experts: knowledge partitioning and function learning. *Psychological review*, *111*(4), 1072–1099.

Koh, K., & Meyer, D. E. (1991). Function learning: induction of continuous stimulus-response relations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*(5), 811–836.

McDaniel, M. A., Cahill, M. J., Robbins, M., & Wiener, C. (2014). Individual differences in learning and transfer: stable tendencies for learning exemplars versus abstracting rules. *Journal of Experimental Psychology: General*, *143*(2), 668–693.