Modeling Active and Passive Hypothesis Testing in Category Learning with ACT-R

Timon Dohnke (timon.dohnke@uni-luebeck.de)
Nele Russwinkel (nele.russwinkel@uni-luebeck.de)
Institute of Information Systems (IFIS), University of Luebeck, Luebeck, Germany

Abstract

Category learning constitutes a fundamental cognitive process in humans. While a considerable body of research has studied this through passive reception of stimuli, real-world scenarios often require learners to actively select the information they learn from. An approach to category learning involving active and passive hypothesis testing is employed to compare learning by reception (using a preselected sequence of stimuli) and selection (using self-chosen stimuli). For this, a cognitive plausible ACT-R model is presented with perception and motor abilities and without memorization of past stimuli. The ACT-R model is evaluated against a subset of a dataset from human participants and a Bayesian model on a rule-based and information-integration category learning task. In the rule-based learning task, the ACT-R model adequately captured the learning curves of human participants. However, in the information-integration task, the accuracy of the model was an average of 9.6 percentage points lower than human performance, a discrepancy primarily driven by the simplified hypothesis structure in the model. The ACT-R model demonstrates how a simple cognitive strategy can solve a category learning task without relying on the memorization of stimuli by employing hidden states, such as a confidence measure of hypotheses.

Keywords: Cognitive Models, ACT-R, Hypothesis Testing, Active Learning

Introduction

Category learning constitutes a fundamental cognitive process in humans. It is employed in numerous contexts, including the categorization of food as spoiled or good, or individuals as acquaintances or strangers. A large body of academic literature focuses on how these categories are learned and constructed (Ashby & Maddox, 2005, 2011; Richler & Palmeri, 2014). Most of this academic literature focuses on how human participants learn categories from a preselected sequence of observations. Although this setup facilitates a more straightforward comparison among participants. In real-world contexts, instances of categorization are frequently not acquired through passive reception alone, but individuals possess a degree of autonomy in the selection of examples from which they learn. This distinction between learning from passive reception and active selection is elaborated in Bruner, Goodnow, and Austin (1956). In other domains, such as hypothesis testing or diagnostic reasoning, the part of selection learning has been studied more extensively (Naghshvar & Javidi, 2013; Monteiro & Norman, 2013).

This paper builds on a study from the paper "Is it better to Select or to Receive? Learning via Active and Passive Hypothesis Testing" from Markant and Gureckis (2013) to further investigate the relationship between reception and selection in category learning. In the study, a rule-based and information-integration category learning task are used to capture the learning curves of selection and reception learners. This paper presents a cognitive model constructed in ACT-R and with a sequential hypothesis testing approach to model these category learning tasks. The ACT-R model is compared with the data from the study, as well as a Bayesian hypothesis testing model, which was also introduced in Markant and Gureckis (2013). The ACT-R model not only models the learning itself, but also the perception and motor skills associated with the tasks in the study. Furthermore, a simple framework for the adaptation of hypothesis without the memorized of past observations is introduced.

In the following a brief overview of the cognitive architecture ACT-R and it's application to category learning is given. Then the category learning tasks and the cognitive model are described. The ACT-R model is compared with previously collected data from human participants and the Bayesian hypothesis testing model. We conclude with a discussion of the results and limitations of the ACT-R model.

Background and Related Work

ACT-R (Adaptive Control of Thought—Rational) is a cognitive architecture that divides human cognition into multiple modules. These modules encompass a intentional (working memory), manual (motor control), visual, and declarative module (long-term memory). The modules can be associated with specific brain regions, thus providing a certain degree of biological grounding (J. Borst, Nijboer, Taatgen, Rijn, & Anderson, 2015). Each of the mentioned modules has a buffer that is used to provide information to other modules. A buffer can hold one chunk with multiple slots. An example of two chunks can be seen in Fig. 3. Another important part of ACT-R is the production system, which represents the procedural knowledge of the model with a collection of IF-THEN rules (i.e. productions). This allows the system to exchange, request, and process information between different modules (Anderson, Matessa, & Lebiere, 1997).

ACT-R has been used to model category learning tasks (Prezenski, Brechmann, Wolff, & Russwinkel, 2017; Taatgen, 1999; Rutledge-Taylor, Lebiere, Thomson, Staszewski, & Anderson, 2012). A common distinction in these modeling

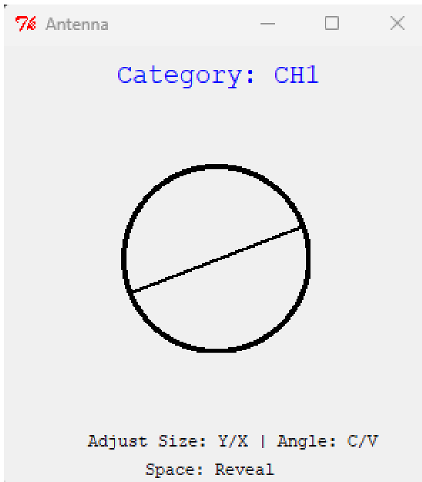

Figure 1: Experiment window of the category learning task during a training trial in the selection condition. The top shows the category of stimulus and in the center the stimulus is displayed. At the bottom the controls are outline to alter the stimuli or to reveal the category.

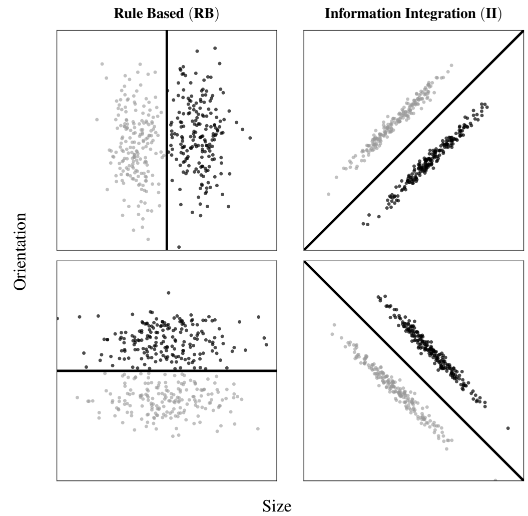

Figure 2: Visualization of the possible decision boundaries for the category learning task. On the horizontal axis the size of the stimulus is shown and on the vertical axis the orientation. Furthermore, it presents a sample of the distributions, which generate the stimuli for the reception condition. Figure recreated from Markant and Gureckis (2013, Figure 5A)

approaches is between exemplar-based and rule-based models. In the exemplar-based modeling approach, the categorization examples are typically saved in declarative memory. Through blending between chunks in declarative memory or partial matching of chunks, the model can learn specific categories. On the other hand, rule-based categorization relies on set predefined productions for classification (Anderson & Betz, 2001).

It is also feasible for models to acquire new productions or parameters for those productions to enable a more dynamic categorization. A prime example of this is the ACT-R model developed by Prezenski et al. (2017) for sound categorization. This type of categorization aligns the most with our approach. Furthermore, a hybrid approach between exemplar-based and rule-based is also possible, as demonstrated in (Anderson & Betz, 2001).

Category Learning Task

Before defining our ACT-R model, we outline the category learning task utilized for evaluation. The task is based on a well-established category learning task (e.g Garner and Felfoldy (1970); Ashby, Alfonso-Reese, Turken, and Waldron (1998); Nosofsky (1989)). We use a variation of the category learning task introduced by Markant and Gureckis (2013) in ”Is it better to Select or to Receive? Learning via Active and Passive Hypothesis Testing”. The goal of the task is to learn whether an ”antenna” picks up Channel 1 or 2. Thus, learning to classify it into two categories. The antenna is characterized by two stimuli, the (circle) size and the orientation of it. Figure 1 shows an example of such an antenna.

A rule-based (RB) and information-integration (II) category are part of the category learning task. In the RB category, the category of stimulus is only deepened on a single dimension of the stimulus, e.g the size or orientation. For the II category both dimensions need to be combined to determine the stimulus category, e.g only small and rotated stimuli or small and non-rotated stimuli. A visualization of the resulting decision boundaries in the state space is shown in Fig. 2 (Markant & Gureckis, 2013).

The study from Markant and Gureckis (2013) consisted of eight blocks and each block comprised of 16 training trials and followed by 32 test trials. Four different conditions were investigated in the study, for which only two are relevant for this paper: Selection (S) and Random-Reception (R). These conditions only influence the training trials. In the selection condition, a randomly generated stimulus is presented at the beginning of each training trail, and the participants are able to adjust it and reveal the category of the altered stimulus. In the reception condition, participants received a random stimulus and the corresponding category. The distribution for these random stimuli depends on the specific decision boundary of the task, samples from these distributions can be seen in Fig. 2. In the test trial, a specific stimulus is shown to participants, which must be classified into one of the two categories. Special consideration is taken when generating test stimuli to avoid bias for certain kind of stimulus. Participants in this study were not subject to any time constraints. The study was conducted with a sample size of 240 participants. However, since the present paper exclusively focuses on two of the four conditions, the effective sample size is reduced to 120 (Markant & Gureckis, 2013).

The ACT-R model in this paper will not only be compared to the study data, but also a Bayesian hypothesis testing model, which was also introduced in Markant and Gureckis

(2013). The Bayesian model represents a hypothesis as a linear decision boundary in the state space. At each training trial, the current hypothesis is compared with a new hypothesis drawn from a prior. The posteriors for both of these hypotheses are calculated based on the last eight stimuli. If the new hypothesis has a greater posterior, it is accepted as the new current hypothesis. Otherwise, the current hypothesis adapted in proportion to the acceptance ratio, i.e the quotient between the posterior of the new and current hypothesis.

Due to space constraints, a comprehensive description of the experimental setup and the Bayesian hypothesis testing model is not possible here; for a full account, the reader is referred to Markant and Gureckis (2013).

Cognitive Model

The following section explains the implementation of the category learning task in ACT-R. Followed by the key principles of the cognitive model and a detailed description of it.

The category learning task was implemented with the help of the ACT-R GUI Interface, as seen in Fig. 1, to allow for a seamless integration with the visual and manual module of ACT-R. Due to the limitations of ACT-R, the adjustment of antenna dimensions was not implemented using a sliding mouse movement as in the original study. Instead, the size or orientation can be altered by predefined amounts using key presses. For the perception of the stimulus, two hidden text fields were used to gather size and orientation values. As in the Bayesian model, the ACT-R model was provided with normalized values of size and orientation ranging from 0 to 600. Furthermore, for the ACT-R model, Gaussian noise $\varepsilon \sim \mathcal{N}(0,9)$ is added to the presented stimulus values to model the uncertainty in their perception. The remaining information was retrieved from the other text fields in the GUI. The timings and transitions were implemented as described in the original paper.

The subsequent section will focus on the key principles of the cognitive model. A sequential hypothesis testing approach is utilized for the model based on a large body of research (Nosofsky & Palmeri, 1998; Noh, Yan, Bjork, & Maddox, 2016) to model category learning as such. Specifically, a hypothesis is maintained as long as the current stimulus does not falsify it. Furthermore, rather than relying on the memorization of past stimuli, only a few hidden states are used, such as confidence in a hypothesis. Lastly, to create novel stimuli in the selection condition hypothesis-dependent sampling bias, a key thesis of Markant and Gureckis (2013), is incorporated, where the novel stimulus is on average close to the current subjective decision boundary.

Figure 4 shows the schematic control flow of the ACT-R model. Each block in the diagram consists of multiple ACT-R productions. The goal buffer contains a control chunk that manages the model's flow of execution. Specifically, the "state" and "step" slots keep track of the current block in the control flow graph and the production in that block, respectively. The control chunk is illustrated in Fig. 3.

Control chunk
+ modi: Train or Test
+ condition: Selection or Reception
+ state: block in control graph
+ step: step in the block
+ decision: CH1 or CH2
+ target-value: number (float)

Hypothesis
+ dimension: Size or Orientation
+ pred-low: CH1 or CH2
+ pred-high: CH1 or CH2
+ switch-marker: nil or Marked
+ uncer-direction: Aboce or Below
+ threshold: number (float)
+ confidence: number (float)

Figure 3: Schematic visualization of the control and hypothesis chunk. "nil" indicates that the variable has no value.

The model starts each new trial in the "New trial and mode detection" block. In this block, the visual module is utilized to discern between training and test trials and whether the task is in the selection or reception condition. This can be determined by referring to the instructions provided in the ACT-R GUI. The results of this are stored in the "modi" and "condition" slots of the control chunk. In the first trial of the experiment, an initial hypothesis is created on the basis of the initial stimulus. For this, the most extreme dimension of the stimulus is taken as a uni-dimensional threshold to divide the state space.

Hypothesis The current hypothesis of the model is saved in the imaginal buffer of the ACT-R model. The hypotheses of the model can only represent uni-dimensional rules. The chunk for the hypothesis consists of six slots. The first slot "dimension" determines whether the size or orientation of the stimulus is tested. The "threshold" slot controls the position of the decision boundary, where as "pred-low" and "pred-high" decide whether above or under the threshold "Channel 1" is predicted. This information is divided into two slots to simplify the modeling. The "switch-marker" is used to tell whether the values of "pred-high" and "pred-low" were already swapped. The switch-marker is introduced to enable the model to remember whether the predicted category above the threshold has already been switched, thereby improving the model's control flow. The "confidence" slot gives a numeric representation of how confident the model is in this particular hypothesis. The starting confidence of a hypothesis is 2. Lastly, "uncer-direction" indicates on which side of the threshold there is greater uncertainty about the category. The full chunk is illustrated in Fig. 3.

Train trial The sequence of events for training trials is marked in blue in the control graph (Fig. 4). For the selection condition, the stimulus is first adjusted based on the current hypothesis. For this, the "uncer-direction" determines whether the altered stimulus is above or below the current threshold value of the hypothesis. Then the "target-value" in the control chunk is set to the sum (or difference) of the threshold value and the absolute value of random number x drawn from the normal distribution $\mathcal{N}(0,60^2)$. This follows from the aforementioned principle of hypothesis-dependent sampling bias. Through a feedback mechanism between the

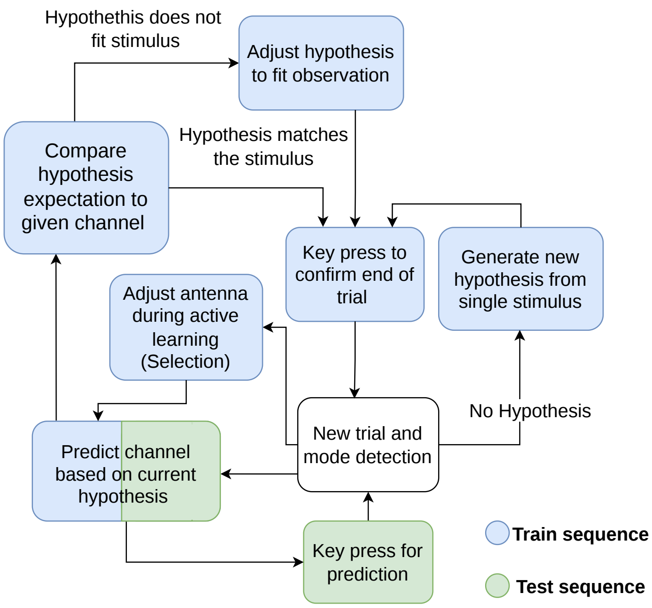

Figure 4: Schematic control flow graph of the ACT-R model. The sequence of events of a training trial are marked in blue and for a test trial in green.

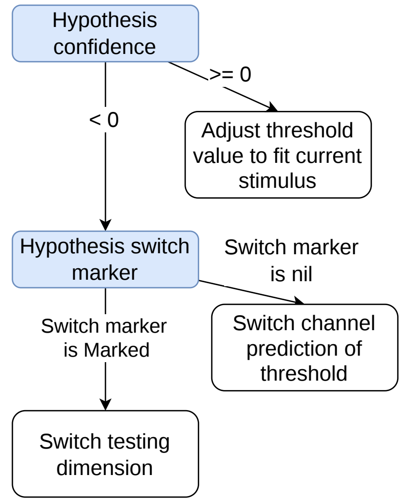

Figure 5: Schematic graph of the adjustment of the hypothesis chunk following a false prediction in a training trial.

visual and motor module the stimulus is iteratively adjusted until the observed stimulus is less than 20 units away from the target-value. In a following step, the category of the stimulus is predicted based on the current hypothesis and compared to the given category. This also marks the first step in the reception condition. If the predicted category matches the category shown, the confidence in the hypothesis increases by 0.5, and the corresponding key is pressed to advance to the next trial.

However, if the predicted category does not match the given category, the confidence in the hypothesis is reduced by 2 and the hypothesis is adjusted. If the confidence of hypothesis is still positive, only the threshold-value is adjusted to fit to the current stimulus. In contrast, if confidence is negative and the switch-marker is not set, the direction of the threshold where Channel 1 and Channel 2 are predicted is switched and the switch-marker is set. If the switch-marker is already set, the testing dimension of the hypothesis, e.g size or orientation, is altered. In both of these cases, the confidence is set back to 2. A schematic view of this algorithm can be seen in Fig. 5. Furthermore, at every training trial the 'uncerdirection' is adjusted to point toward the side of the threshold opposite to the current stimulus.

Test trial The test trial procedure is identical across all conditions and is marked in green in Fig. 4. With the help of the test stimulus and the current hypothesis, a prediction about the category is made. The corresponding key to the predicted channel is pressed and advanced to the next trial.

A subset of the hyperparameters for the model was determined via an exhaustive grid search and selected on the basis of the mean absolute error between the learning curves (accuracies over test blocks) between the model and the study data. The hyperparameters determined by the grid search included the starting confidence of a new hypothesis $(\mathbf{2}, 6, 10)$, the confidence penalty for a wrong prediction $(0.5, 1, \mathbf{2})$, the confidence reward for a correct prediction $(0.2, \mathbf{0.5}$ and the standard deviation, when adjusting the hypothesis in the selection condition $(1, 10, 20, 40, \mathbf{60}, 80)$. The hyperparameter values selected for the ACT-R model are marked in bold. Although these hyperparameters provide a good fit for the study data, this study does not assess the cognitive validity of the parameters employed.

Results

In the results, we compare the ACT-R model with the previously collected study data and with the aforementioned Bayesian hypothesis testing model, both obtained from Markant and Gureckis (2013). The category learning task is analyzed for the selection and reception condition and the rule-based and information-integration category. For each of the four resulting configurations, 30 experiment runs are compared. Figure 6 shows the overall accuracy and the average accuracy of each test block. Additionally, the Standard Error of Mean is also plotted. The test blocks consist of 32 trials and are each separated by 16 training trials. To begin with, in the selection condition with the rule-based category task, the ACT-R model matches the overall accuracy and the learning curve across the test blocks well. The mean absolute error (MAE) between the accuracies of the test blocks between the ACT-R model and the study data is only 2.5 percentage points. For the reception condition with the rule-based category task, the ACT-R model performed considerably better than human participants, an increase of 12.3 percentage points compared to the study data. However, the shape of the learning curve on the test blocks is quite similar, the MAE of the relative change between the test blocks is only 1.4 percentage points.

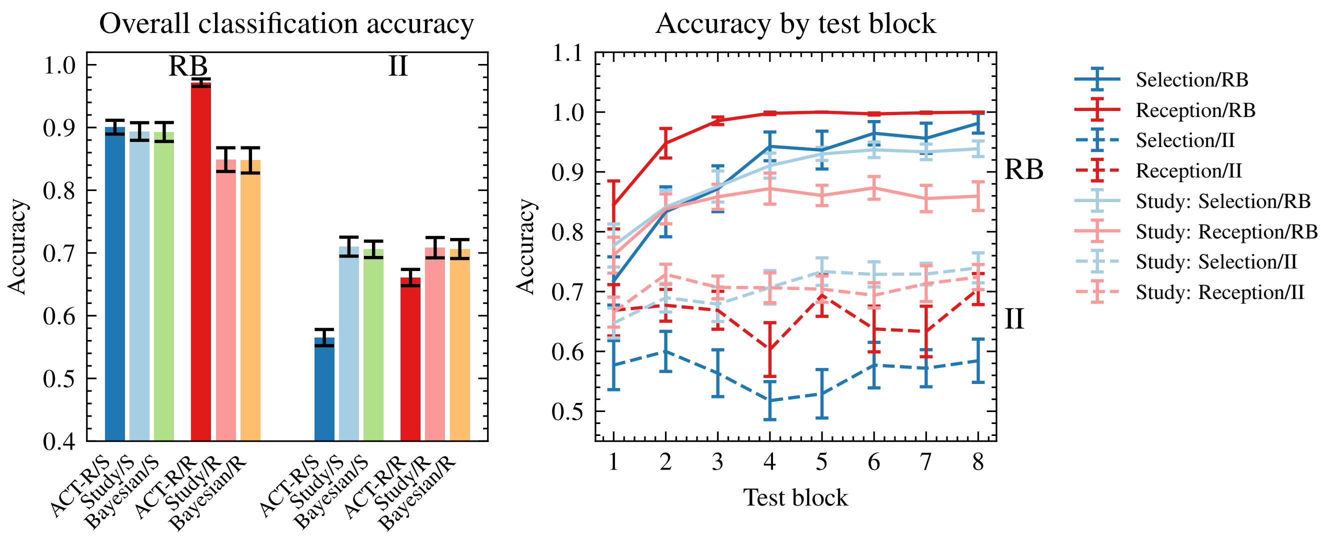

Figure 6: Left: Overall classification accuracy from the ACT-R model, Bayesian model and human participants over the eight test blocks. Right: Accuracy of the ACT-R model and human participants across the eight test blocks. The diagrams show the data for the selection (S) and reception (R) conditions, as well as the rule-based (RB) and information integration (II) categories. The sample size for each of the data points is 30. The error bars are showing the Standard Error of Mean (SEM). Figure adapted from (Markant and Gureckis (2013), Fig. 7)

Continuing with the information-integration category, the ACT-R model performed in both conditions worse than human participants; in the selection case by 14.5 percentage points and in the reception case by 4.8. Overall, the reception condition outperforms the selection condition in both rule categories in the ACT-R model, where, as in the study, the selection group achieves a better accuracy than the reception group in the rule-based category and similar performance in the information-integration category. The comparison with the Bayesian models yields a very similar pattern, as the Bayesian approach is able to closely reproduce the empirical data from the study.

Discussion

In this paper, a cognitive plausible ACT-R model for active and passive hypothesis testing was introduced and compared with human participant data and a Bayesian model on a rule-based and information-integration category learning task. The ACT-R model simulates a specific strategy to solve these category learning tasks. It is reasonable to assume that the human participants in the study used a wide variety of strategies to solve these category learning tasks. Thus, the aim of the model is not to perfectly reproduce the data, but to give some intuition of the cognitive processes that may be involved in solving these tasks.

The ACT-R model demonstrated a level of accuracy that was either equivalent (selection condition) to or greater (reception condition) than that of the Bayesian model in the rule-based category, although the Bayesian model had access to the previous 8 stimuli and compared two competing hypotheses in each trial. This illustrates how a simple strategy without relying on the memorization past stimuli of as in the ACT-R can exhibit similar or better performance as a more complicated Bayesian model. Additionally, some advantages are gained with the shift from a computational Bayesian approach to a more mechanistic approach in ACT-R. For example, ACT-R models offer the possibility of comparing and predicting response times and brain activity of human participants, as described in J. Borst et al. (2015); J. P. Borst and Anderson (2015); Atashfeshan and Razavi (2017). Thus, ACT-R's bottom-up approach allows for greater grounding with sufficient experimental data.

In the reception condition and rule-based category, the ACT-R model demonstrates considerably higher accuracy compared to the human participant. One potential explanation for this phenomenon is that the initial hypothesis in the ACT-R model is highly dependent on the presented stimulus. As shown in Figure 2, the average stimulus shown in the reception condition is already on average close to the decision boundary. Therefore, it is plausible that the model will initially start with a strong hypothesis.

Due to this strong initial hypothesis the ACT-R model fails to capture the positive relationship between the selection condition and human accuracy in the rule based category, while the Bayesian model manages to do so. The superior alignment of the Bayesian model with the study data is attributable to its access to past stimuli, as well as its superior computational flexibility, in contrast to the more mechanistic ACT-R model.

The inferior performance of the ACT-R model in the information-integration category is likely attributable to the ability of hypothesis chunk to only represent uni-dimensional decision boundaries. The assumption that human participants continue to employ uni-dimensional decision boundaries even when the decision boundary is more complex is not implausible. The original paper indicated that, in 34.5% of the instances, a uni-dimensional decision boundary was identified as the best fit model for the responses within the information integration category (Markant & Gureckis, 2013). Incorporating a more complex hypothesis chunk would likely result in enhanced alignment with the study data.

Outlook

Although the ACT-R model demonstrates a high degree of correlation with the study data in certain scenarios, it also fails to reproduce certain key findings from the study, such as superior performance of the selection condition over reception one in the rule-based category. An option to improve the modeling performance of the ACT-R model would be to restructure the hypothesis chunk in ACT-R to allow for a more complex decision boundary. This would in particular improve the correlation with the study data in the information integration category. Additionally, the adjustment of hyperparameters for each study participant would presumably allow for a superior fit of the individual learning curves and on collective as a whole. Furthermore, for a more complete understanding of the mechanisms underlying category learning, further research is required on the factors that influence the selection of learning stimuli and the subsequent acquisition of categories.

Conclusion

This paper presents a cognitive model for active and passive hypothesis testing in category learning. The cognitive model is implemented in ACT-R and integrates perception as well as motor abilities. The model is evaluated on a category learning task with a selection condition, where stimuli for learning could be actively selected, and a reception condition, where stimuli for learning are preselected. Furthermore, the category learning task offers a simpler rule-based decision category and a more complex information integration one. The model is compared to existing data from human participants and data from a Bayesian hypothesis testing model, both of which were obtained from Markant and Gureckis (2013).

The Bayesian model demonstrates superior alignment with the data recorded from human participants. However, the ACT-R model also displays the capacity to adequately model study data in the selection condition with the rule-based category. For the reception condition in this category, the accuracy of the ACT-R model is 12.3 percentage points higher, but it has a similar relative learning curve. This is particularly interesting because the ACT-R model does not compare competing hypotheses or rely on the memorization of the previous eight stimuli, as with the Bayesian model. Achieving this mainly involves introducing simple hidden states, such as confidence in the current hypothesis and a strategy for adapting the hypothesis when it fails. The ACT-R model is rooted in a more mechanistic bottom-up approach compared to the Bayesian model. The ACT-R approach offers the additional advantage of allowing the prediction of reaction times or the activity of brain regions by using a suitable model, which can subsequently be compared with empirical data (J. P. Borst & Anderson, 2015; Atashfeshan & Razavi, 2017). In summary, further research is necessary to determine how ACT-R models can achieve behavior similar to that of Bayesian models while utilizing more verifiable cognitive processes. In addition, category learning with active selection of stimuli remains an area of promising research.

Acknowledgments

The authors thank Dr. Todd Gureckis and Dr. Douglas Markant for collecting and providing the dataset.

References

Anderson, J. R., & Betz, J. (2001, December). A hybrid model of categorization. *Psychon Bull Rev*, *8*(4), 629–647.

Anderson, J. R., Matessa, M., & Lebiere, C. (1997). Act-r: A theory of higher level cognition and its relation to visual attention. *Human–Computer Interaction*, *12*(4), 439–462. Retrieved from https://doi.org/10.1207/s15327051hci1204_5

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*(3), 442–481. doi: 10.1037/0033-295X.105.3.442

Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annu Rev Psychol*, *56*, 149–178.

Ashby, F. G., & Maddox, W. T. (2011). Human category learning 2.0. *Annals of the New York Academy of Sciences*, *1224*(1), 147–161.

Atashfeshan, N., & Razavi, H. (2017). Determination of the proper rest time for a cyclic mental task using act-r architecture. *Human Factors*, *59*(2), 299-313. Retrieved from https://doi.org/10.1177/0018720816670767 (PMID: 27738278) doi: 10.1177/0018720816670767

Borst, J., Nijboer, M., Taatgen, N., Rijn, H., & Anderson, J. (2015, 03). Using data-driven model-brain mappings to constrain formal models of cognition. *PloS one*, *10*, e0119673. doi: 10.1371/journal.pone.0119673

Borst, J. P., & Anderson, J. R. (2015). Using the act-r cognitive architecture in combination with fmri data. In B. U. Forstmann & E.-J. Wagenmakers (Eds.), *An introduction to model-based cognitive neuroscience* (pp. 339–352). New York, NY: Springer New York. Retrieved from https://doi.org/10.1007/978-1-4939-2236-9_17 doi: 10.1007/978-1-4939-2236-9_17

Bruner, J. S., Goodnow, J. J., & Austin, G. A. (1956). *A study of thinking*. John Wiley and Sons. doi: 10.4324/9781315083223

Garner, W., & Felfoldy, G. L. (1970). Integrality of stimulus dimensions in various types of information processing. *Cognitive Psychology*, *1*(3), 225-241. doi: https://doi.org/10.1016/0010-0285(70)90016-2

Markant, D. B., & Gureckis, T. M. (2013, March). Is it better to select or to receive? learning via active and passive hypothesis testing. *J Exp Psychol Gen*, *143*(1), 94–122.

Monteiro, S. M., & Norman, G. (2013). Diagnostic reasoning: Where we've been, where we're going. *Teaching and Learning in Medicine*, *25*(sup1), S26–S32. Retrieved from `https://doi.org/10.1080/10401334.2013.842911`

Naghshvar, M., & Javidi, T. (2013). Active sequential hypothesis testing. *The Annals of Statistics*, *41*(6), 2703 – 2738. Retrieved from `https://doi.org/10.1214/13-AOS1144`

Noh, S. M., Yan, V. X., Bjork, R. A., & Maddox, W. T. (2016). Optimal sequencing during category learning: Testing a dual-learning systems perspective. *Cognition*, *155*, 23-29. doi: https://doi.org/10.1016/j.cognition.2016.06.007

Nosofsky, R. M. (1989, Jul 01). Further tests of an exemplar-similarity approach to relating identification and categorization. *Perception & Psychophysics*, *45*(4), 279-290. Retrieved from `https://doi.org/10.3758/BF03204942`

Nosofsky, R. M., & Palmeri, T. J. (1998, Sep 01). A rule-plus-exception model for classifying objects in continuous-dimension spaces. *Psychonomic Bulletin & Review*, *5*(3), 345-369. Retrieved from `https://doi.org/10.3758/BF03208813`

Prezenski, S., Brechmann, A., Wolff, S., & Russwinkel, N. (2017, August). A Cognitive Modeling Approach to Strategy Formation in Dynamic Decision Making. *Frontiers in Psychology*, *8*. doi: 10.3389/fpsyg.2017.01335

Richler, J. J., & Palmeri, T. J. (2014). Visual category learning. *Wiley Interdisciplinary Reviews: Cognitive Science*, *5*(1), 75–94.

Rutledge-Taylor, M., Lebiere, C., Thomson, R., Staszewski, J., & Anderson, J. (2012, 03). A comparison of rule-based versus exemplar-based categorization using the act-r architecture..

Taatgen, N. A. (1999). Explicit learning in act-r. In *Mind modelling: A cognitive science approach to reasoning, learning and discovery* (pp. 233–252). Pabst, Lengerich.