

Predicting the Effects of In-Task Instruction During Multi-cue Diagnosis

Andrew Halsey¹, Christopher W. Myers², Kevin Gluck², & Jack Harris²

phillip.halsey.1.ctr@us.af.mil, christopher.myers.29@us.af.mil, kevin.gluck@us.af.mil, jack.harris.3@us.af.mil

¹Oak Ridge Institute for Science and Education at AFRL ²Air Force Research Laboratory

Keywords: cognitive modeling; ACT-R; fast and frugal trees

Introduction

In this research we examine a specific type of human-machine teaming, decision support systems (DSS; Power, 2008), to facilitate decision-making in an uncertain environment (Eom, Lee, Kim, and Somarajan, 1998). We extended a previously reported instance-based learning cognitive model (Myers, Gluck, Harris, Veksler, Mielke, and Boyd, 2015) to receive decision support from a machine learning algorithm. To date, no models have integrated instance-based learning and decision support, though both are well represented individually in the literature (e.g., Power, 2008; Thomson, Lebiere, Anderson, and Staszewski, 2015). We are interested in examining the strengths and deficits of integrating these as a predictive model of human-machine teaming in the context of a multi-cue diagnosis decision task.

Multi-cue Diagnosis Task

The multi-cue diagnosis task is a two-alternative forced choice task where a response is made based on available cues. In the current task, individuals diagnose “patients” for heart attacks according to three binary cues available each trial. Each cue is associated with a different “symptom”—the presence of which is probabilistic—and this information may help determine whether the patient should go to the Coronary Care Unit or standard Nurse’s Bed (Green and Mehr, 1997; Marewski and Gigerenzer, 2012; Myers, et al., 2015). Feedback is provided based on the final decision. The next trial begins after delivery of feedback.

The learning difficulty within a particular multi-cue diagnosis environment is governed by the environment’s rule consistency and the symptom base rates. Rule consistency is the probability that using the underlying rule results in a correct diagnosis. The rule for the current experiment was: if and only if cue2 is true and cue3 is true, then it is a heart attack and the correct response is Coronary Care Unit. Given the current rule consistency, choosing Coronary Care Unit in the presence of these symptoms would result in a correct response 80% of the time. The presence of a positive “symptom” associated with each cue was: cue1=0.25, cue2=0.40, and cue3=0.75.

Instance-Based Learning Theory and Model

Gonzalez, Lerch, and Lebiere (2003) proposed Instance-based Learning Theory (IBLT) as a process account of human learning during repeated decision-making. IBLT posits how humans identify, store, and retrieve information for the explicit purpose of making decisions within a

dynamic, uncertain environment when performance feedback is provided (Gonzalez, et al., 2003).

IBLT has successfully accounted for human behavior in two-alternative forced-choice tasks (e.g., Gonzalez, and Dutt, 2011), classification tasks (Gagliardi, 2011), and dynamic tasks (Gluck, Stanley, Moore, Reitter, and Halbrugge, 2010; Reitter, 2010). We developed an IBLT model in ACT-R (Anderson, 2007; Thomson, et al., 2015) to make a decision based on a particular context (i.e., the presence of symptoms) and prior experience.

The model does not generate a response according to explicitly defined rules indicating the number and order of cues to check. Rather, the model generates its decision by using ACT-R’s blending mechanism to blend over chunks and to determine cue encoding order and a stopping rule according to a particular context. In the current paper, the IBLT model encoded decision support instruction similarly; rather than providing an explicit rule, decision support was represented as a collection of high feedback chunks that suggested a response given a particular context. Previous research has shown that the IBLT model is capable of predicting human behavior on a multi-cue diagnosis task (Myers, et al., 2015).

Model Evaluation

To simulate human-machine teaming, the IBLT model received decision support from a machine learner using a constrained version of the A* algorithm that constructed a decision rule with maximum expected reward. Three decision support types were tested across a single rule consistency: correct DSS (optimal rule), incorrect DSS (non-optimal rule), and no DSS. Each decision support condition completed 20 runs of 267 trials. Decision support was delivered at trial 60 to allow the machine learner to settle on an environmentally consistent optimal rule and to ensure the IBLT model had not reached asymptotic response accuracy performance.

The IBLT model goes through three different stages across the experiment: exploration, instruction, and exploitation. The focus in the current paper was on the exploration and exploitation phase, each with unique questions concerning model performance. During the exploration stage, we were interested in model behavior according to rule acquisition, rule adherence, and accuracy. Specifically, because the model generates a rule based on learning and experience, does the model find the underlying rule or does it generate an alternative rule? We examined optimal rule-adherence with respect to accuracy for insights into these questions.

After the machine learner delivers decision support instruction, it would be unsurprising that rule adherence and accuracy change. However, of interest is how the model behaves during the exploitation stage when the decision support rule is encoded in declarative memory. Two primary questions need to be addressed during this stage. First, to what extent does the model appropriately use the provided rule? Second, does the decision support coerce or force the model into a pattern of inflexible responding by suppressing exploratory behavior? In other words, if the environment transitioned to a new rule, would the model continue responding according to the provided decision support rule, or would the model still be able to detect environmental changes and subsequently adjust its strategy?

Model Predictions

Trial data from each run was binned to 9 blocks for ease of interpretation. Decision support instruction occurred at the beginning of block 3 (trial 60; indicated in Figure 1 as a dashed vertical line). Data from model accuracy and reward across blocks were identical and therefore we used model accuracy to examine model performance.

During the exploratory stage of blocks 1 and 2, the IBLT model began to learn and respond according to a self-generated rule that, according to the model, appeared to best explain the environment. Given accuracy and rule adherence, the rule generated by the model during these blocks was sub-optimal and as a result accuracy was unable to match the probability associated with the environment rule consistency.

After receiving decision support from the machine learner at block 3, the model was responsive in incorporating the provided instruction into its rule strategy (Figure 1). Correct DSS instruction resulted in increased model accuracy and rule adherence by following the underlying environmental rule. Thus, not only was the model making appropriate responses, but these were a result of adherence to the decision support rather than a decision strategy or rule self-generated by the model. Incorrect DSS and no DSS also responded expectedly according to the type of rule (or lack thereof) provided.

During the exploitation phase, after delivery of decision support, model behavior remains—to a degree—flexible and exploratory. For example, in the correct DSS condition, rule adherence of the model increases to nearly 100% then begins to drift lower in subsequent blocks. The reason for this behavioral variability is a result of the model forgetting the rule over time. Rule forgetting across blocks is gradual; rule adherence and accuracy both remain higher relative to the same metrics during the exploratory stage. However, some degree of forgetting behavior can be advantageous. It allows the model to continue responding dynamically according to the environment, and should an environmental change occur (i.e., rule change or change in probabilities), the model can detect these changes and adjust the response strategy accordingly.

Response time predictions corroborated documented deficits of the IBLT model (Myers, et al., 2015).

Infrequently used chunks in declarative memory increased response time due to lower probability of recall. Response time decreased at block 7 because the chunks used thereafter included only frequently or recently used chunks. Based on the current results and previous findings (Myers et al., 2015), the IBLT model may not be capable of accurately accounting for human response times in a two-alternative forced choice task.

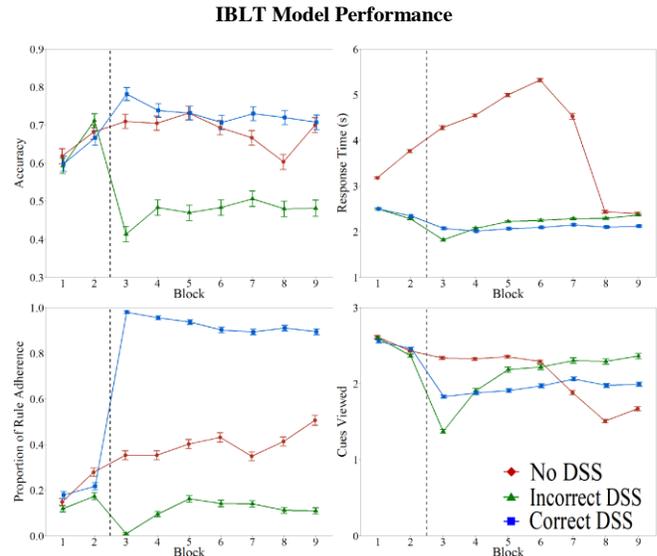


Figure 1. Performance data for the three models (+/- 1 SEM). The dashed line represents decision support.

Conclusions

The integration of the IBLT model with a machine learning decision support system demonstrated several strengths and weaknesses. The IBLT model is capable of taking instruction and incorporating it within its rule discovery strategy, as evidenced by accuracy and rule adherence changes between the exploratory and exploitation stages. The incorporation of the rule strategy does not suppress future learning. In fact, the model resumes some amount of exploratory behavior after instruction, thereby allowing the model to remain flexible and adaptive to possible environmental changes. These core strengths demonstrate a model with explanatory potential when validated against human behavior. The main weakness of the model relates to the IBLT model's inability to model response times, such as those demonstrated by humans engaged in a similar task (Myers, et al., 2005). Future research will tackle issues such as direct human to model comparisons according to instruction-taking, and determining the timing and frequency of instruction delivery.

Acknowledgments

This research was supported by the Air Force Office of Scientific Research, grant 13RH06COR.

References

- Anderson, J. R. (2007). *How can the human mind exist in the physical universe?* New York, NY: Oxford University Press.
- Eom, S. B., Lee, S. M., Kim, E. B., & Somarajan, C. (1998). A survey of decision support system applications (1988-1994). *Journal of the Operational Research Society*, 49, 109-120.
- Gagliardi, F. (2011). Instance-based classifiers applied to medical databases: Diagnosis and knowledge extraction. *Artificial Intelligence in Medicine*, 52, 123-139.
- Gluck, K., Stanley, C., Moore, L., Reitter, D., & Halbrügge, M. (2010). Exploration for understanding in cognitive modeling. *Journal of Artificial General Intelligence*, 2, 88-107.
- Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, 118, 1-29.
- Gonzalez, C., Lerch, J. F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, 27, 591-635.
- Green, L., & Mehr, D. R. (1997). What alters physicians' decisions to admit to the coronary care unit?. *Journal of Family Practice*, 45, 219-226.
- Kaplan, B. (2001). Evaluating informatics applications—clinical decision support systems literature review. *International Journal of Medical Informatics*, 64, 15-37.
- Marewski, J. N., & Gigerenzer, G. (2012). Heuristic decision making in medicine. *Dialogues in Clinical Neuroscience*, 14, 77-89.
- Myers, C. W., Gluck, K. A., Harris, J., Veksler, V., Mielke, T., & Boyd, R. (2015). Evaluating instance-based learning in multi-cue diagnosis. In N. A. Taatgen, M. K., van Vugt, J. P Borst, & K. Mehlhorn. *Proceedings of ICCM 2015*. Paper presented at International Conference on Cognitive Science, Groningen, Netherlands (198-199).
- Power, D. J. (2008). Decision support systems: A historical overview. In F. Burnstein & C. W. Holsapple (Eds.), *Handbook on decision support systems 1: Basic themes* (pp. 121-140). Springer.
- Reitter, D. (2010). Metacognition and multiple strategies in a cognitive model of online control. *Journal of Artificial General Intelligence*, 2, 20-37.
- Thomson, R., Lebiere, C., Anderson, J. R., & Staszewski, J. (2015). A general instance-based learning framework for studying intuitive decision-making in a cognitive architecture. *Journal of Applied Research in Memory and Cognition*, 4, 180-190.