Evaluating Instance-based Learning in Multi-cue Diagnosis

Christopher W. Myers[1], Kevin A. Gluck[1], Jack Harris[1], Vladislav Veksler[2], Thomas Mielke[3], & Rachel Boyd[4]
christopher.myers.29@us.af.mil, kevin.gluck@us.af.mil, jack.harris@us.af.mil, vdv718@gmail.com,
thomas.r.mielke@gmail.com, rachel.boyd.rb1@gmail.com

[1]*Air Force Research Laboratory* [2]*Army Research Laboratory* [3]*L3 Link Simulation at AFRL*
[4] *Oak Ridge Institute for Science & Education at AFRL*

Keywords: instance-based learning; fast and frugal trees

Introduction

Decision heuristics are often described as fast and frugal, meaning that they take little time and require relatively few computations to make a decision when compared to optimal decision systems (Gigerenzer & Todd, 1999). Fast & Frugal Trees are one heuristic that are a special case of decision trees in which there is a possible exit out of the decision process at every cue considered in the tree (Luan, Schooler, & Gigerenzer, 2011).

There is currently no computational account of how humans learn heuristics like F&FT-based decision processes. This is a significant gap in our scientific understanding, and we aim to begin addressing that gap in this effort. In this abstract we report results from a pilot study assessing Instance-based Learning Theory (IBLT) as an account of human learning from experience in domains where F&FTs may be good decision heuristics, such as diagnostic tasks.

Instance-based Learning Theory

Instance-based Learning Theory (IBLT) is a theory of how humans acquire and apply new knowledge given performance feedback and a particular context. It was developed to explain and understand human decision processes in dynamic task environments (Gonzalez, Lerch, & Lebiere, 2003) The four components of any IBLT model are (1) episodic memory elements (i.e., *instances*), (2) retrieving the instance from memory, (3) contextual similarity, and (4) integrating feedback across multiple, contextually similar events. In essence, an instance provides the utility of a particular action given a specified context in a way similar to expected utility theory

As far as we can tell, IBLT has not been applied to decision tasks where a set of different cues can be discriminately sampled for improving decision making. In the following section we introduce a multi-cue diagnosis task.

Multi-cue Diagnosis Task

The multi-cue diagnosis task is an extension of 2AFC tasks, where a decision-maker is provided with two alternative responses and a set of cues with which to inform the decision. Cues are binary (i.e., present or absent) and may be related to particular responses; part of decision-makers' task is to learn which cue(s) is (are) important.

In the task, there are three cues that a decision-maker can choose to use for determining a response. Cue information is not immediately visually available and requires clicking on a cue button to reveal its presence or absence. The decision maker is free to use any number of the cues in any order for informing their decision, and the only cost with accessing cue information was behavioral (i.e., moving to, clicking, etc.). Further, decision-makers were not speeded in their response and no penalty was issued based on trial response time. Given this basic task, we derived two environments: an easy environment (EZ) and data recreated from real-world CCU diagnoses (GnM; Green & Mehr, 1997). These environments were selected to provide approximate ceiling and floor performance in not only response accuracy, but also the adoption of prescribed F&FTs.

There were two payoff regimes: balanced (BAL) and heavy-miss (HM). In balanced, hits and correct-rejections received 10 points whereas misses and false alarms were penalized -10. The HM regime was the same as BAL except misses received -50.

In the pilot study reported here we ran a 2 (environment difficulty) x 2(payoff regime) between subjects design. We ran five participants through each of the four conditions. Each participant performed nine blocks of 30 trials. For each subject, on each trial, we captured their accuracy, the symptoms they revealed, and the order in which they were revealed. For each block the proportion of correct responses (i.e., accuracy), response time (RT), the proportion of selected responses (i.e., response selection), and the adherence to the prescribed F&FT (i.e., rule adherence) was calculated. Subjects' performance improved with experience in each condition and the EZ environment was easier than the GnM environment. Further subjects' RTs decreased as their acquisition and adherence to the prescribed rule increased (see Figure 1).

ACT-R Instance-based Learning Model

We developed an IBLT model in the ACT-R architecture (Anderson, 2007). We used ACT-R's declarative memory system to instantiate IBLT components one and two. We did not vary the degree of similarity between instances, instead opting for identity. This is justified as there was no hypothetical relationship between the binary cues. Finally, we used the ACT-R *blending* mechanism to instantiate IBLT component four.

The model used IBLT to determine the order of cues to check, when to stop checking, and which response to make. We believe this to be a novel use of IBLT, and the model represents complete adherence to the theory for execution.

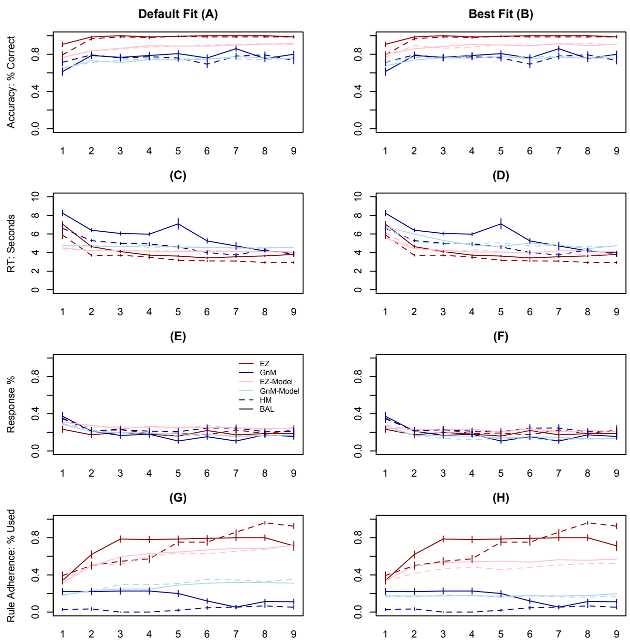

Figure 1: Human data and model results.

The model did not use production compilation nor production utility learning for acquiring any skill in the task environment.

Model Evaluation

We performed three evaluations of the model. For each evaluation, the model was run five times across 19 blocks of 30 trials, resetting after each run. The first evaluation (i.e., *default fit*) used parameters that were either default, taken from the central tendency of parameters in the Max Planck ACT-R parameter database (Wong, Cokely, & Schooler, 2010), or hypothesized where no guidance was available (see Table 1). The second and third evaluations (*best fit-all but RT* and *best fit-RT*, respectively) varied the retrieval threshold, blending temperature, activation noise, and decay parameters in a full combinatorial design producing 37,632 combinations of values. We only modified declarative memory and blending parameters as the investigation was on the adequacy of IBLT to account for multi-cue diagnosis tasks. Further, we report two different RMSEs for each evaluation: one for *RT* and another for the rest of the dependent variables (i.e., *Other*). We did this because the RT and the other dependent variables are on quite different scales.

The *best fit-all but RT* and *best fit-RT* surprisingly resulted in the same model parameters, and thus is referred to as *best-fit* (see Table 1, Best-fit column). The model performed quite well in the *default fit* evaluation, with an RT $RMSE = 1.162; R^2 = 0.552$ and an Other $RMSE = 0.682; R^2 = 0.914$. The model also performed well in the *best-fit* evaluation, with an RT $RMSE = 0.786; R^2 = 0.747$ and an Other $RMSE = 0.466; R^2 = 0.906$. Interestingly, with an improved fit in RT with best-fitting parameters (panel D, Figure 1) over the default (panel C), there is a reduction in rule adherence fitness (see panel H & G). Further, the :rt and :ans parameters are quite different from the central tendency of those reported by the community (see Table 1).

Parameter Name	Best-fit	Default	Source
decay (:bll)	0.1	0.4	MPIB-DB
base level constant (:blc)	0	1	Free
retrieval threshold (:rt)	-50	-0.4	MPIB-DB
activation noise (:ans)	0.75	0.4	MPIB-DB
blending temp	1	1	Free
imaginal-activation	1	1	Free

Table 1: Parameter values for the *default* and *best-fit* models. *MPIB-DB* refers to the Max Planck ACT-R Database and *Free* refers to hypothesized values due to no guidance on its setting. The source refers only to the *default* model parameters as the *best-fit* were derived by iterating over the large parameter space and minimizing RMSE. All other parameters were default values.

Conclusions

Generally, IBLT seems well suited for multi-cue diagnosis tasks. However, there appears to be a tradeoff between accounting for rule adherence and response times. Specifically, when fitting the model RTs, rule adherence decreased relative to the *default-fit* parameters. Consequently, we conclude that IBLT may not be sufficient to account for both RTs and rule adherence in this environment. Finally, the default fit model performed quite well, highlighting the value of making model parameter databases available to the community.

Acknowledgments

This work 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?* Oxford University Press.

Gigerenzer, G., & Todd, P. (1999). *Simple Heuristics That Make Us Smart*. Oxford, England: Oxford University Press.

Gonzalez, C., Lerch, F. J., & 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.

Luan, S., Schooler, L. J., & Gigerenzer, G. (2011). A signal-detection analysis of fast-and-frugal trees. *Psychological Review*, *118*(2), 316–338.

Wong, T. J., Cokely, E. T., & Schooler, L. J. (2010). An Online Database of ACT-R Parameters : Towards a Transparent Community-based Approach to Model Development Managing Parameters for ACT-R Models. In *International conference of cognitive modeling* (Vol. 1, pp. 282–286).