

Modeling Human Information Acquisition Strategies

Annerieke Heuvelink (annerieke.heuvelink@tno.nl)

TNO Human Factors, Kampweg 5
3796 ZG Soesterberg, the Netherlands

Michel C. A. Klein, Rianne M. van Lambalgen (michel.klein@few.vu.nl, rm.van.lambalgen@few.vu.nl)

Vrije Universiteit Amsterdam, Department of Computer Science,
De Boelelaan 1081a, 1081 HV Amsterdam, the Netherlands

Abstract

The focus of this paper is the development of a computational model for intelligent agents that decides on whether to acquire required information by retrieving it from memory or by interacting with the world. First, we present a task for which such decisions have to be made. Next, we discuss an experiment that shows that humans do not apply rational expected utility analysis to make this decision, but instead adopt a simpler heuristic strategy. Then, we introduce a computational model that incorporates the rational as well as various heuristic task strategies. The human data is compared to the behavior of the model under several parameter settings. We were able to match the human actions with model actions for various task strategies, suggesting that humans differ in the task strategies they apply, and that our manner to deduce heuristic task strategies is feasible.

Keywords: Expected Utility Analysis, Interactive Behavior, Memory Retrieval, Strategies.

Introduction

For the execution of almost all tasks knowledge is required. For example, baking a cake requires explicit knowledge about its ingredients. When preparing for the task, a human will make an (often implicit) choice between retrieving the required knowledge from memory and looking it up. Intuitively, this choice is determined by the balance between the costs of looking up information on the one hand, and the costs of retrieval and the risk of mistakes on the other hand. In the baking example the choice could be to look up the recipe, as it is probably hard to retrieve it from memory and the cost of a mistake is quite high (distasteful cake).

Selecting actions based on their expected costs and benefits is often described as rational decision making. However, it is well known that humans do not always follow a rational process, but often depend on heuristic approaches to solve a problem (Tversky & Kahneman, 1974; Gigerenzer, Todd & the ABC research group, 1999). In addition, humans vary (between-subject) in the task-specific strategies they apply, and this choice is also influenced (within-subject) by the specific task circumstances (see, e.g., Beilock & DeCaro, 2007; Byrne, Kirlik & Fleetwood, 2008).

The overall aim of our work is to build intelligent agents that exhibit human-like behavior. In order to do so, we would like to build a computational model that can decide on whether to acquire information by retrieving it from memory or by interacting with the world.

In this paper, we first describe the experiment in which we analyzed the behavior of humans in a relative simple task that required them to choose between in-the-head information and in-the-world information under two cost conditions. We start with a description of the task. Then, we discuss how rational expected utility analysis could be applied to the task at hand, i.e., what the types of costs and benefits of its actions are. Subsequently, the behavioral experiment and its results are presented.

In the second part of the paper, we try to align the results of the experiment with a developed task model that takes both the rational-choice approach as heuristic-based approaches into account. We first discuss the possible heuristic strategies that people could apply for the task introduced. We then elaborate on the results of the technical experiment performed to find the model's parameter values that best fit the results of the behavioral experiment. Finally, the implications of the findings are discussed.

Task Description

The computer task we developed required participants to classify presented objects to specific bins. During the task, 9 objects were presented in a sequence of 36 trials. The objects were composed of a color (red, blue or yellow) and a shape (square, circle or triangle). Each object belonged to a specific bin, numbered 1 to 9, but initially the participants did not know the correct combinations. The goal of the task was to press the number of the correct bin upon presentation of the object. On each trial participants had the option to press the number of a bin first ('choose'), or to press a button to get more information about the bins ('sense'). Participants could choose one of three buttons: button 'j' revealed the bins of objects with the same color as the presented object; button 'k' revealed the bins with the same shape; and button 'l' revealed the bin of the specific object. After the information was shown, participants had to select a bin. After a bin was chosen, the correct bin was revealed.

Table 1: Costs of the two experiment conditions.

Cond	Feature	Button Money	Button Time	Error Costs
1	Color	€0.10	1.0s = €0.02	€0.10
1	Shape	€0.10	1.0s = €0.02	€0.15
1	All	€0.15	1.5s = €0.03	€0.20
2	Color	€0.06	1.0s = €0.02	€0.12
2	Shape	€0.06	1.0s = €0.02	€0.18
2	All	€0.09	1.5s = €0.03	€0.24

Participants started the task with 10 euro. Money was subtracted when either a button was chosen, or an error was made; see Table 1 for the two specific cost-settings used. In addition, for every 500 ms 0.01 euro was subtracted. A typical trial started with presenting the object with below it 9 empty boxes. Furthermore, the three buttons were shown and in the upper right corner the amount of money left.

When participants choose to sense color or shape, they had to wait for 1.0 seconds until the requested information was shown. When participants choose to sense all, they had to wait for 1.5 seconds. Meanwhile, time costs were still subtracted. When the waiting time had passed, the object was presented again with below it the 9 bins of which the bins were revealed that matched the specific feature that was sensed: the three bins that matched the color of the object, the three bins that matched its shape, or the bin that matched the whole object. When a bin was chosen (immediately, or after sensing), the object and the 9 bins were presented again with the correct bin revealed. At the same time feedback was given on the choice of the participant.

The combination of 9 objects in 36 trials was determined previous to the experiment, to make sure that some objects would be often encountered so that over time it would be well known to which bin they belonged, while for others, less encountered, this could have been forgotten. See Table 2 for the number of specific objects presented over the trials.

Table 2: Overview of objects presented.

Feature	3 x Red	2 x Blue	1 x Yellow
3 x Circle	RC: 9x	BC: 6x	YC: 3x
2 x Square	RS: 6x	BS: 4x	YS: 2x
1 x Triangle	RT: 3x	BT: 2x	YT: 1x

Rational Expected Utility Analysis

The presented task requires interactive behavior: for its performance a mixture of elementary cognitive, perceptual, and motor operations are required. Gray and Boehm-Davis (2000) introduce interactive routines as the basis of interactive behavior. They envision interactive routines as dependency networks of low-level cognitive, perceptual, and motor operators that come together at a time span of about 1/3 to 3 seconds. Gray and Fu (2004) propose that at this time span, the human control system selects sequences of interactive routines that tend to minimize performance costs measured in time while achieving expected benefits.

For the presented task it is possible to rely to a smaller or larger degree on information in-the-world versus information in-the-head. In the first case more interaction with the world is required (button pressing), in the second case more intensive memory use (remembering the colors and shapes of the bins). Based on the specific task conditions it is expected that humans will adopt different interactive routines to minimize performance costs.

A rational strategy for performing the presented task would determine at each trial which of the four possible actions would be most optimal to execute: either directly choosing a bin, or first requesting which bins fit the color,

shape, or both these aspects of the presented object. For this, a cost-benefit analysis of each action needs to be made.

For the presented task four types of costs exist: 1) the money it costs when a certain mistake is made, 2) the money it costs to press a button, 3) the time it costs to do so, and 4) the time it costs to retrieve beliefs from memory. It is possible to express all the various types of costs in money, because time costs money. It could be debated that in addition to these money and time costs another type of costs exist, namely the cognitive and perceptual-motor effort involved in executing the actions. We do not separately distinguish these efforts but assume that time is a reasonable surrogate measure for them (Gray & Fu, 2004).

To determine the expected utility of each of the actions, the expected costs for each of the four types of costs need to be determined. The money and time it costs to press one or none of the buttons depends on the task condition, but apart from that can be determined in a straightforward way. It is more difficult to determine the expected costs of 1) making an error and of 2) retrieving beliefs from memory.

For the first aspect the chance that one of the three possible errors is made (color false, shape false, all false) is important together with their respective, task condition dependent, penalties. The chance that a specific error is made depends on what is remembered. When it is possible to retrieve the correct bin for a specific object, the chance on any error is zero. However, when this is not possible the chance on a specific error depends on the chance of correctly retrieving knowledge concerning bins with the to-be-classified object's color or shape, but also on the chance that knowledge is retrieved that exclude specific bins from selection, increasing the chance the correct bin is picked.

The expected cost of retrieving beliefs from memory is equal to the time to do so or to the time to failure. These times, as well as the chance that knowledge can be retrieved in the first place, are important to know for calculating the expected utilities. Insight in these aspects can come from models of human memory. A well known model of memory retrieval is embedded in the cognitive theory ACT-R (Anderson et al., 2004). In ACT-R declarative knowledge is presented by chunks, whose activation values determine their chance and speed of retrieval, the latter according to this formula:

$$RT = Fe^{-Ai}$$

RT: The time to retrieve the chunk in seconds.

A_i: The activation of the chunk *i* which is being retrieved.

F: The latency factor parameter.

The latency factor parameter depends on the retrieval threshold, *T*, which varies substantially between ACT-R models. However, the following general relationship has been discovered: $F = 0.35 e^T$ which means that the retrieval latency at threshold (when $A_i = T$) is approximately 0.35 seconds (Anderson et al., 2004). The full equation used by ACT-R to determine a chunk's activation takes into account several aspects, but its basis is the chunk's base-level activation. The base level activation *B_i* reflects the recency and frequency of use of the chunk, and is calculated by:

Table 3: Results of Regression Analysis.

Dependent Variables	Independent Variables			First_Choice			RT_First			RT_Bin			Sense_Feature			Correct_Bin		
				<i>p</i>	<i>R</i> ²	<i>r</i>	<i>p</i>	<i>R</i> ²	<i>r</i>	<i>p</i>	<i>R</i> ²	<i>r</i>	<i>p</i>	<i>R</i> ²	<i>r</i>	<i>p</i>	<i>R</i> ²	<i>r</i>
	Act-Color	0.002	0.27	0.52	0.000	0.48	-0.69	0.000	0.46	-0.68	0.002	0.28	-0.53	0.004	0.24	0.49		
Act-Shape	0.000	0.38	0.62	0.000	0.36	-0.60	0.000	0.35	-0.59	0.000	0.40	-0.63	0.006	0.22	0.46			
Act-All	0.001	0.35	0.59	0.000	0.53	-0.73	0.000	0.59	-0.73	0.000	0.43	-0.65	0.002	0.32	0.57			

$$B_i = \ln\left(\sum_{j=1}^n t_j^{-d}\right) + \beta_i$$

n: The number of presentations for chunk *i*.

t_j: The time since the *j_{th}* presentation.

d: The decay parameter. Standard this one is set at 0.5.

β_i : A constant offset.

When we assume that people can unconsciously employ a kind of utility analysis (which includes having some kind of implicit knowledge about what they can remember, see Gray et al. (2006)) and adopt these interactive routines that minimize performance costs, we expect to find differences in behavior between the two cost conditions introduced.

Experiment

Sixteen first year AI students, aged between 17 - 24 years, participated in the experiment. The experiment's duration was approximate 30 minutes and participants received 1 to 10 euro for participation, depending on their performance.

In the experiment a 2-factor, between subjects design was used, with *costs* varied between participants. In condition 1, the costs of pressing a button were relatively high compared to the costs of an error, while in condition 2 the opposite was the case. For an overview of exact costs, see Table 1.

Participants started by reading a written instruction on how to perform the experiment and the costs of errors, time and sensing. Next, a practice task was given to familiarize them with the task and the costs. This task was similar to the main task, but in order to keep a low interference, color and shapes of objects were altered. Furthermore, the bin in which often or rarely encountered objects belonged and the order in which the objects were presented was altered.

Data Analysis

For data analysis we first calculated for each bin and at each trial the expected activation value of the participant's knowledge concerning the color, shape and the whole object (all) that would fit in the bin. For this we used the ACT-R formula with a standard decay of 0.5 and an offset of 0. As 'presentations' we counted the display of bin information due to button use, and the display of the correct bin at the end of each trial. Next, these activation values were used for regression analysis across participants for each trial. Trials where the activation was 0 (e.g. the object had not been presented before) were excluded from analysis.

Univariate variance analysis was used to check for differences between the two conditions. For the difference between color and shape, a repeated measure ANOVA was

conducted, using the Huyn-Feldt correction. For all analysis, trials with a RT exceeding 8000ms were excluded.

Experimental Results

Over all the participants, the percentage correctly classified objects ranged from 30 to 97 percent; the average percentage correct was 61 (*SD*=21). The number of times a participant chose a bin immediately ranged from 5 to 34; the average was 24.44 (*SD*=7.87). So overall, there was a wide variety in the participant's behavior.

The results of the linear regression analysis are shown in Table 3. The *R*² (explained variance), *r* (correlation) and *p*-values are given for each analysis. The results show that the activation value of color, shape and the whole object was successful in predicting a number of variables, confirming that the ACT-R theory correctly captures how human memory operates.

First_Choice (the number of participants who chose a bin immediately) is positively dependent on activation value: as activation increased, First_Choice increased. Furthermore **RT** (reaction time) was examined: RT when the object is shown for the first time ('**RT_First**') and the time from the presentation of the object to the moment the bin was chosen ('**RT_Bin**'). Both RT's are dependent on the activations: RT decreased when activation value increased.

In addition, the percentage of correct classifications concerning color, shape and all was found to be positively dependent on the activation of color, shape and all, see Table 3. When the activation increased, the percentage correct increased as well. The number of times a specific feature was sensed ('**Sense_Feature**') for color, shape or all decreased as the activation value of that feature increased.

Figure 1 shows the results of the ANOVA on the interaction between condition and feature. A trend is revealed when looking at the percentage incorrect. In condition 1 participants' percentage incorrect of shape (*M*=0.32, *SD*=0.15) was higher than that of color (*M*=0.21, *SD*=0.15; *F*(1,7)=6.81, *p*<0.04). For participants in condition 2 no such difference was found. An interaction is found between condition and feature on the number of times participants sensed a feature. For participants in condition 2 a trend was revealed, which showed that the percentage of sensed shape (*M*=0.30, *SD*=0.27) was higher than the percentage of sensed color (*M*=0.19, *SD*=0.28; *F*(1,7)=4.37, *p*<0.1). For participants in condition 1 no significant difference was found between the percentage of sensed color and the percentage of sensed shape.

Other than these interactions, no differences were found between the two conditions. This indicates that participants did not always make a rational decision, otherwise we would have expected to find more variety, e.g., in the total number of times features were sensed. Support for the thesis that humans instead rely on a prefixed strategy is found in the data, e.g., of two participants in the same condition, one always chose to acquire unknown information from the world (by pressing the *all* button ‘1’), whereas the other always attempted to retrieve it from memory (never pressed any button). Additional support can be found in the description of their approach by the subjects themselves. At least 5 participants described an approach that is different from rational decision making, e.g. “I choose for shape and color (“1”) if unsure or unknown, else I answered”.

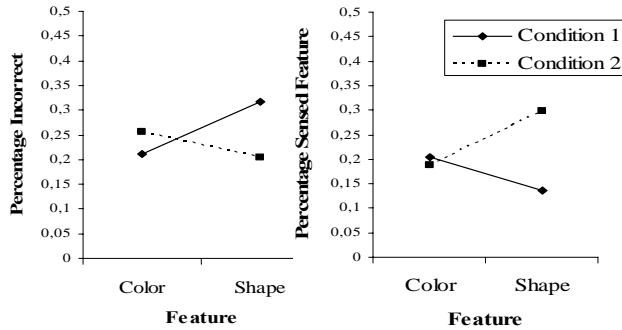


Figure 1: Interaction between feature and condition on percentage incorrect and percentage sensed feature.

Discussion of Experimental Results

Overall, the results show that people’s decision to acquire information from the world or from memory correlates with the activation of that information in memory following ACT-R’s base-level activation formula, and is thus dependent on the frequency and recency of that information.

A difference is found between color and shape, in that shape appears more difficult to retrieve from memory than color. This is shown by the fact that when people retrieve information from memory, the chance of making a mistake concerning shape is higher than the chance of making a mistake concerning color, see Figure 1. When the costs of acquiring information from the world are relatively low, this difference disappears as in such a situation people request shape (button ‘k’) more than color (button ‘j’).

No other differences are found between condition 1 and 2 when looking at the participants’ reaction times or actions (sense or choose bin). This indicates that the decision to rely on information in-the-world versus information in-the-head is not influenced by the specific costs of acquiring that information. Rather it seems that people make a decision based on their own (pre-)specified strategy.

This finding does not necessary conflict the hypothesis that humans optimize their interactive routines to minimize performance costs. Gray and Fu (2004) and Gray et al. (2006) only consider performance costs measured in time, and argue that humans are evolved to conserve the resource of time. For the task presented in this paper performance

costs are a combination of time and money costs, and it is conceivable that humans are not good in taking into account the money costs of actions. Since the time costs of actions do not alter between the two conditions, this might explain that no more differences can be found between them. On the other hand, people do attempt to optimize their performance based on time and money costs: if they would only optimize the time costs, they would never press a button.

Task Model

As mentioned in the introduction, our research goal is the development of methods and techniques that will enable intelligent agents to display human-like behavior which might be rational, but often is not. For this goal we previously developed a memory model enabling rational as well as biased reasoning (Heuvelink, Klein & Treur, 2008). This model was implemented in SWI-Prolog (Wielemaker, 2003), and incorporates ACT-R’s base-level activation formula for declarative knowledge in memory. In this paper we take that model as basis for the development of a task specific model capable of executing the task previously introduced: <http://human-ambience.few.vu.nl/docs/CogSci-IIAModel.pl>

Heuristic Strategies

A task specific extension of the original model is the specification of heuristic strategies suited for this task. Gray and Fu (2004) state that the cost-benefit considerations for interactive routines only provide a soft constraint on their selection as they may be overridden by deliberately adopted top-down strategies. The statistical analysis and the description of the approach by participants are indications that this might have happened in our task.

Based on logical reasoning and inspired by the participants’ answers, we came up with 37 possible heuristic strategies participants could follow. The strategies mainly differ in the number of retrieval actions humans are willing to execute (1, 2 or 3), and the order in which they do so. There is also the possibility of an extra security check, to see whether the bin selected to be chosen is not in conflict with the given object (e.g., when checked, it turns out that the shape of the selected bin can be retrieved and conflicts that of the object). Possible actions that can be taken after one of the retrieval steps are:

- choose a random bin (a)
- choose a random bin with security check (b)
- press show color/shape button, then choose random one of the three presented bins with security check. (c/d)
- press show all button, then choose that bin. (e)

Figure 2 summarizes all strategies. In the **first retrieval step** it is tried to retrieve the bin that matches the whole object which is presented. When retrieval is unsuccessful, any one of the actions a, b, c, d and e can be taken, which results respectively in strategies 1, 2, 3, 4, 5.

Instead of directly choosing an action after unsuccessful object retrieval, a participant can make a **second retrieval step** to retrieve a bin of which either the color or the shape

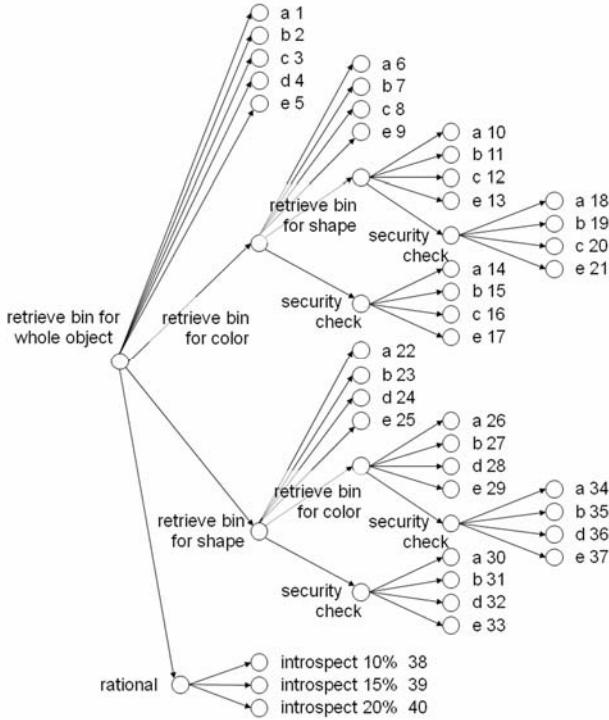


Figure 2: Schematic overview of all strategies.

fits that of the object. If it is possible to retrieve the specific feature, that bin will be chosen. If it is not possible to retrieve it, again a specific action will be taken. For strategy 6 to 9 and 14 to 17, action a, b, c and e will be taken directly after an unsuccessful attempt to retrieve *color*. The difference between strategies 6 to 9 and 14 to 17 is that the latter perform a security check when color can be retrieved. Strategy 22 to 25 and 30 to 33 are the same, but attempt to retrieve shape instead of color and take actions a, b, d and e.

There is also the possibility of a **third retrieval step** after retrieving color or shape. That is, if color can not be retrieved, in such strategies people will first try to retrieve shape before taking an action. Strategy 10 to 13 first try to retrieve color, then try to retrieve shape. Strategies 18 to 21 do the same, but with an extra security check. Actions a, b, c and e are taken when retrieving is unsuccessful. Strategy 26 to 29 first try to retrieve shape, then try to retrieve color (strategy 34 to 37 with an extra security check). Actions a, b, d and e are taken with unsuccessful retrieval.

In addition to the 37 strategies just introduced, we also implemented the rational strategy and included it as strategy 38-40. These strategies were equal in their determination of the expected costs of each action, but varied in the time it took them to introspect the activation values of the beliefs. This took them respectively 10, 15 and 20% of the time that it would take to actually retrieve the belief inspected.

Parameter Fitting

The developed model contains a large number of parameters. Each specific parameter setting will result in

different behavior of the model. To answer the question to what extent the model can correctly describe human behavior, we performed a technical experiment to find parameter settings for which the model displays behavior close to that of a participant. Due to the large number of parameters, we were unable to fit them all, so we focused on fitting the strategy parameter as well as the parameters that influence the storage and retrieval of beliefs. For each of the selected parameter settings we ran the model and gave it the 36 objects to classify. Subsequently, we compared each participant with the simulation results. To do this in a structured way, we developed a distance measure that calculates for each trial a distance between the model data and the data of the participant.

Results Parameter Fitting

The results of the parameter fitting are not yet a thorough validation of the model, but do still provide evidence for the feasibility of the model. The fitted subjects, two for each condition, were selected based on typical behavior patterns: participant 2 (condition 2) almost always requested information, participant 7 (condition 1) almost never did. Participant 9 (condition 1) and 10 (condition 2) were chosen because they seemed to perform rational behavior (more sensing in the beginning, less sensing at the end). For all participants the settings with distance values that lie within 1% of the lowest distance value were analyzed. For these settings we found that per participant the parameters for *strategy* and *retrieval_threshold* were equal.

The strategy parameter that fits participant 2 is strategy 5, with a retrieval threshold of 0.5. This strategy entails that when an object can not be retrieved from memory, its position will be requested. Analysis of the best matching setting showed that action of subject 2 indeed correlates with action of the model ($r=0.47, p<0.01$). Reaction time of subject 2 does not correlate with reaction time of the model.

Strategy 30 and a retrieval threshold of 0.5 fit best with participant 7. This strategy often results in directly choosing a bin as when shape can not be retrieved, a random bin is chosen. This is apparent in participant 7, who only pressed a button at the first two trials.

Participant 9 fits best with strategy 39 and a retrieval threshold of -0.1. Strategy 39 is a rational strategy taking the costs of acquiring information from the world and from memory into account. Further analysis revealed a significant correlation between human action and model action ($r=0.68, p<0.01$), but also between human RT and model RT ($r=0.40, p<0.02$).

Strategy 36 and a retrieval threshold of 0.2 fit best with participant 10. Strategy 36 is, contrary to our expectations, not a rational strategy. The strategy either results in choosing a bin (when either shape or color is known), or in sensing the shape (when shape and color are both unknown or one of them conflicts). Indeed there is a significant correlation between human action and model action ($r=0.61, p<0.01$). In addition, a trend in correlation was found between human RT and model RT ($r=0.31, p<0.1$).

Discussion & Conclusion

The results show that it was possible to find parameter settings that match reasonably well with the four investigated participants, especially on the executed actions. Reaction time proved to be a less optimal measurement for parameter fitting. This could be due to the fact that we set a fixed time to observe information, and to press a bin or a button for all participants. As reaction time is personal, such parameters need to be fitted as well.

We can also conclude that people adopt different strategies to decide whether to acquire information in-the-world versus information in-the-head. At this moment we think that many of our participants already decided on how to act beforehand. The descriptions of the strategies by the subjects themselves support this hypothesis.

With hindsight knowledge, we can make a few critical remarks about our experimental setup and our model. First, the task that was given to the subjects was too complex, in the sense that it contained too many cost parameters. This made it difficult for the participants to do an accurate cost-benefit analysis, shown by the fact that we were not able to clearly distinguish an effect of the different cost conditions.

Second, it became clear that the setup of the task made it possible to choose a strategy that optimizes the utility *over different trials*. Some participants preferred to sense ‘color’ or ‘shape’ over ‘all’ because the first two options revealed information about objects in three bins instead of information about an object in one bin. As the rational strategies in our model do not take this into account, such (actual rational) strategies did not fit the rational strategy.

Third, we can conclude that we made a suboptimal choice in selecting the parameters to be fitted. Major parameter settings were fixed (time to observe information and time to execute actions) while it was attempted to fit others that were of much less importance to task execution.

Fourth, it is a question whether our ‘meta-model’ for deriving the 37 strategies is correct, i.e., the idea that the heuristic strategies vary in the number (and order) of retrieval actions humans make. On the other hand, modeling different strategies is in line with the work of Dickison and Taatgen (2007), who state that for complex tasks it may become impossible to model individual differences by parameter tuning. Instead, they propose that people differ in the control strategies they employ, and that these manifest themselves as different problem-solving strategies. The control strategies supposedly differ in the amount of top-down control exerted on behavior, opposed to this behavior being driven by bottom-up processes.

It could well be that people differ in the type of control they exert (with top-down control leading to more rational behavior) based on other individual differences, e.g. the capacity of their working memory (WM). Differences in WM capacity have been used to explain the differences between the task strategies selected by different humans under the same task circumstances, as by the same human under varying circumstances (Beilock & DeCaro, 2007). Given these findings, we think that our approach to capture

variations in human decision-making by modeling (heuristic) strategies that differ in the number of retrieval actions humans are willing to make, is a feasible one.

In future work, we would like to redo the experiments using the insights that are described above, i.e., using a simpler task. In addition, we want to vary and fit on more model parameters, and we would like to extend the model so it does not execute a pre-determined strategy, but online selects one, e.g., based on the available WM capacity. Furthermore, it might be interesting to further investigate the difference we found between color and shape retrieval.

References

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y (2004). An Integrated Theory of the Mind. *Psychological Review*, 111(4), 1036-1060.

Beilock, S.L. & DeCaro, M.S. (2007). From poor performance to success under stress: Working memory, strategy selection, and mathematical problem solving under pressure. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 33(6), 983-998.

Byrne, M. D., Kirlik, A., & Fleetwood, M. D. (2008). An ACT-R approach to closing the loop on computational cognitive modeling: Describing the dynamics of interactive decision making and attention allocation. In D. C. Foyle & B. L. Hooey (Eds.) *Human performance modeling in aviation*. Boca Raton, FL: CRC Press.

Dickison, D & Taatgen, N.A. (2007). ACT-R Models of Cognitive Control in the Abstract Decision Making Task. In *Proceedings of the Eighth International Conference on Cognitive Modeling*. New York, NY: Psychology Press.

Heuvelink, A., Klein, M. C. A., & Treur, J. (2008). An agent memory model enabling rational and biased reasoning. In *Proceedings of the 2008 IEEE/WIC/ACM International Conference on Intelligent Agent Technology* (pp. 193–199). IEEE Computer Society Press.

Gigerenzer, G., Todd, P. M., & ABC Research Group (1999). Simple heuristics that make us smart. New York: Oxford University Press.

Gray, W. D., & Boehm-Davis, D. A. (2000). Milliseconds Matter: An introduction to microstrategies and to their use in describing and predicting interactive behavior. *Journal of Experimental Psychology: Applied*, 6(4), 322-335.

Gray, W. D., & Fu, W.-T. (2004). Soft constraints in interactive behavior: The case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science*, 28(3), 359-382.

Gray, W. D., Sims, C. R., Fu, W.-T., & Schoelles, M. J. (2006). The soft constraints hypothesis: A rational analysis approach to resource allocation for interactive behavior. *Psychological Review*, 113(3), 461-482.

Tversky, A. & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124-1131.

Wielemaker, J. (2003). An Overview of the SWI-Prolog Programming Environment. In *Proceedings of the 13th International Workshop on Logic Programming Environments* (pp. 1-16).