1, 2, 3, More: An Accumulator Architecture for Anchoring and Adjustment

Kevin Burns (kburns@mitre.org)
The MITRE Corporation, 202 Burlington Road
Bedford, MA 01730-1420 USA

Abstract

We present a computational model of probabilistic reasoning in a synthetic task environment. The task is a solitaire game called TRACS in which a player must update odds as card are turned, like Blackjack. Initial experiments with this game showed that human subjects exhibit a "baseline bias" due to a heuristic strategy of "anchoring and adjustment". Here we develop a computational model of this behavior using fuzzy functions and simple summing to simulate the cognitive processes of memory activation and evidence accumulation. We compare the simulation results to experimental data and highlight the practical value of our approach.

Introduction

Computational models of human behavior are needed for a variety of practical applications (see National Research Council, 1998). Previous studies have demonstrated that human judgment and decision making are biased by simplified heuristics such as "anchoring and adjustment" (Kahneman, Slovic & Tversky, 1982), but most of these studies have employed static questions that do not reflect the dynamic conditions of "naturalistic" missions (Zsambok and Klein, 1997). In addition, most studies have produced theoretical descriptions of heuristics and biases rather than rather than computational simulations of judgment and decision making. Here we present an experimental study with more naturalistic context and we propose a computational model that simulates human behavior.

Our task is a solitaire game called Straight TRACS. TRACS (Burns, 2001) is a family of card games designed to measure and model thinking in tasks that are prototypical of "command and control", including probabilistic risk assessment and dynamic resource allocation. Card games, which involve "imperfect information", are widely regarded as micro models of business and warfare in the real world (McDonald, 1950). The TRACS games use two-sided cards that offer additional advantages of practical relevance and scientific rigor. To improve relevance, the back of each card gives some information about the front of the card, much like an x-ray or radar gives some information about a tissue or target. To improve rigor, the two-sided cards provide more constraint on the possible game states. This makes TRACS more tractable to mathematical analysis of normative solutions, which are needed to benchmark cognitive limitations.

TRACS Task

The backs of the cards show black shapes called "tracks" (triangle, circle or square) and the fronts of the cards show colored sets called "treads" (Red or Blue). There are multiple copies (2, 4 or 6) of each track/tread card type in the deck (see Table 1). The underlying analogy is that of a tire whose tread (set of shapes) leaves tracks (single shapes) in proportion to the number of shapes in the set. For example, a triangle is probably Red (6/8=75%), a square is probably Blue and a circle is 50-50. These are the baseline (full deck) odds at the start of the game, but the actual odds change as tracks are turned and their treads are revealed in play.

The basic game, called Straight TRACS, is played solitaire. Each turn is a forced choice between two tracks where the object is to turn the track that is most likely to match a given tread. The challenge is to update the odds as the deck is depleted in order to choose the best track on each turn.

To play Straight TRACS: The deck is held face down and three cards are dealt to a field. Two cards are dealt face down (tracks) and a third card is dealt face up (tread). The player indicates his subjective belief P(tread|track) for each track by clicking a button on a colored ruler (Figure 1). The player then turns one of the two tracks, trying to match the tread (color) of the third card. The result is scored as a "save" (match) or a "strike" (mismatch) and the pair (match or mismatch) is removed from the field. The track that remains on the field is turned to reveal its tread and this becomes the color to match on the next turn. Two new tracks are dealt from the deck to the field, the player turns a track, etc.

The game provides data on odds judgments and card choices, but here we are interested in odds judgments.

Table 1: Distribution of 24 cards in the TRACS deck.

# of Cards	6	4	2	2	4	6
Front (tread)	Red	Red	Red	Blue	Blue	Blue
Back (track)	▲	●	■	▲	●	■

Figure 1: Confidence meter used to measure subjective belief.

Experiment A: Random Decks

Experiment A tested 43 subjects, each playing 10 games of Straight TRACS on a personal computer with a mouse interface. All games used a random shuffle of the standard deck (Table 1) and all 43*10=430 shuffles were unique. For each track on each turn, we compared the reported odds to the actual odds.

Figure 2 shows that the average error (all games) for human subjects is similar to the average error for a baseline agent, i.e., a simulated player who always sets the confidence meter at the baseline odds. Figure 2 also shows that human subjects did somewhat better than the baseline agent near the ends of games, i.e., people were not just playing the baseline odds. These and other results (see Burns, 2002a; Burns, 2002b) suggest that subjects update the odds using a heuristic strategy of "anchoring and adjustment".

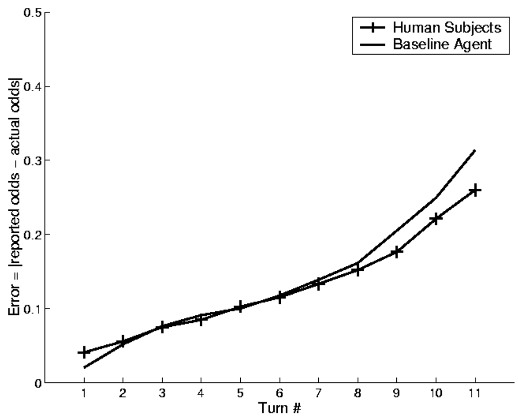

Figure 2: Average error versus turn in game.

Besides "anchoring and adjustment", Experiment A showed that subjects were quite limited in "concurrent counting" of cards. The figure of merit here is "hits" in setting the confidence meter at odds of 100% since odds of 100% arise only when all N of a given card type (N=2, 4 or 6, see Table 1) have been revealed. Our finding is that subjects can count up to 2 cards with reasonable reliability (~50%) but the hit rate drops sharply for counting 4 or more cards. That is, the limit of reliable performance is about 3±1 cards when counting multiple card types in parallel.

Mind Sets

The observed limit in concurrent counting is similar to the well-known limit on "subitizing" dots of a visual display (Dehaene, 1992). Numerous empirical studies describe a comparable limit and various theoretical studies suggest that the number 3 (or so) is a natural minimum and optimum (Cowan, 2001). Here we propose that subitizing and counting are both governed by a "mind set" (i.e., mental structure) that can only hold about three chunks at once.

In subitizing, three chunks let three dots be apprehended in parallel. In counting, three chunks {N, 1, N+1} let any number of cards be enumerated in series, i.e., as N+1 becomes N for the next step. In concurrent counting (Figure 3), we propose that a set of three chunks {L-2, L-1, L} is maintained in parallel for each card type and that attention must shift from set to set as different card types are seen in series. This reduces the reliability of counting each card type since the running total (L) for a given type must be recalled when resuming the count for that type.

Accumulator Architecture

Our computational model is based on an accumulator metaphor (Dehaene, 1997). We assume 6 distinct accumulators (one per card type), each with a fuzzy aperture for filling and viewing the level (Figure 3). This makes the apprehended (recalled) level L' a fuzzy function of the accumulated (encoded) level L. The fuzzy function F(L'|L) is defined by three numbers, High:Medium:Low, which govern the probability that L'=L (High), L'=L-1 (Medium) or L'=L-2 (Low).

This fuzzy function is designed to model the task (i.e., enumeration) and the brain (i.e., activation) in a simplified manner. The task requires links between adjacent levels (Figure 3b) so the brain can encode the new level and forget the old level as different accumulators are filled in "parallel distributed" fashion. The need to encode and forget makes the level fuzzy because residual activation will be maintained at level L-1 (via recurrent connections, not shown in Figure 3b) even after the level is incremented to L.

The residual activation at L-1 serves two purposes. First, it serves to excite (encode) L via the link from L-1 to L. Second, it serves to inhibit (forget) L-2 via the link from L-1 to L-2. Residual activation is also assumed at L-2 because decay is not instant. Subsequently, when the level is incremented to L+1, residual activation is maintained at L (to excite L+1) and residual activation must decay from L-1, etc.

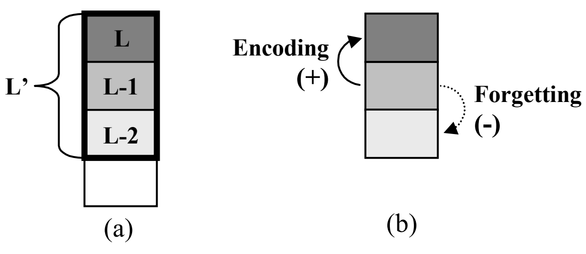

Figure 3: (a) A fluid accumulator with a fuzzy aperture. Gray depicts the probability (darker is higher) that the apprehended level L' will be equal to L or L-1 or L-2. (b) Residual activation at level L-1 allows the encoding of a new level L and the forgetting of an old level L-2.

Qualitatively, the average activation at L-1 should be lower than the activation at L but higher than the activation at L-2, otherwise the accumulator would not fill. Thus, we characterize the relative activation levels simply as High (L), Medium (L-1) and Low (L-2).

Quantitatively, the numerical values for High:Medium:Low are based on the results of Experiment A, which showed a hit rate of about 50% in concurrent counting up to 2. Since only High and Medium levels are involved in counting up to 2, and since counting to 2 involves two fuzzy viewing operations (first at L=1, then at L=2), we set High:Medum:Low to 70:25:5 (i.e., 0.7*0.7~0.50).

Adjustment Algorithm

Our algorithm for adjusting the accumulator level entails two distinct but related operations: inspecting (viewing) and incrementing (filling). The sequence is to inspect, inspect, increment, increment for each turn of the game because there are two tracks on the field before each turn. Each inspect operation (before cards are turned) or increment operation (after cards are turned) involves several steps (see below), most employing a fuzzy function.

Inspect Operation

(1) Attention: See the track type τ (triangle, circle or square) with probability P=1.

(2) Activation: Get the Red and Blue accumulators $A_{R\tau}$ and $A_{B\tau}$ for card types Rτ and Bτ with probability $P=F(A_{R\tau}|R\tau)$ and $P=F(A_{B\tau}|B\tau)$.

(3) Apprehension: View the levels $L'_{R\tau}$ and $L'_{B\tau}$ of the accumulators $A_{R\tau}$ and $A_{B\tau}$ with probability $P=F(L'_{R\tau}|L_{R\tau})$ and $P=F(L'_{B\tau}|L_{B\tau})$.

(4) Aggregation: Use the levels $L'_{R\tau}$ and $L'_{B\tau}$ from step 3 to compute the change in odds, relative to the baseline odds, for track type τ. The change is computed as: $C_{R\tau}=L'_{B\tau}*\delta_{B\tau} - L'_{R\tau}*\delta_{R\tau}$ or $C_{B\tau}=L'_{R\tau}*\delta_{R\tau} - L'_{B\tau}*\delta_{B\tau}$ where $\delta_{R\tau}$ (or $\delta_{B\tau}$) is a constant equal to the change in odds for a single card of type Rτ (or Bτ).

(5) Action: Use the change in odds C_τ from step 4 to compute the odds O_τ that are C_τ away from the baseline odds for track type τ. Set the confidence meter at odds of O'_τ with probability $P=F(O'_\tau|O_\tau)$.

Increment Operation

(6) Attention: See the card type $\Gamma\tau$ (either Rτ or Bτ) with probability P=1.

(7) Activation: Get the accumulator $A_{\Gamma\tau}$ for card type $\Gamma\tau$ with probability $P=F(A_{\Gamma\tau}|\Gamma\tau)$.

(8) Apprehension: View the level $L'_{\Gamma\tau}$ of accumulator $A_{\Gamma\tau}$ with probability $P=F(L'_{\Gamma\tau}|L_{\Gamma\tau})$ where $L_{\Gamma\tau}$ is the last level set by step 9 ($L_{\Gamma\tau}=0$ at the start).

(9) Addition: Add 1 to the apprehended level $L'_{\Gamma\tau}$ to get the newly accumulated level $L_{\Gamma\tau}(new)=L'_{\Gamma\tau}+1$.

The details of each step are outlined below.

(1) Attention

The model assumes that a track is seen with probability 1 because there are only two tracks on each field and because the confidence meter must be set for each track.

(2) Activation

The model assumes that there is one accumulator $A_{\Gamma\tau}$ for each card type $\Gamma\tau$. Both accumulators (Red and Blue) for a track τ on the field are activated in order to inspect their accumulated levels. The activation is accomplished by a fuzzy function $F(A_{\Gamma\tau}|\Gamma\tau)$ that accounts for both static and dynamic associations between card types.

The static associations are assumed to be High, Medium or Low based on the similarity of card types with respect to frequency (baseline odds), track type and tread type. A match of all features is High; a match of frequency and track type is Medium; a match of frequency alone is Low. Using order of magnitude estimates, we assume 90:9:1 for High:Medium:Low. For the dynamic associations, the accumulators that were incremented on the previous turn are assumed to have residual (base-level) activation and their static association strengths are increased by a factor of 2.

(3) Apprehension

The model uses the fuzzy function F(L'|L) discussed above (see "Accumulator Architecture") to compute the apprehended level L' from the accumulated level L.

(4) Aggregation

If $L'_{R\tau}$ or $L'_{B\tau}$ is equal to its known maximum (2, 4 or 6) then the odds are set to 100% Red or Blue. Otherwise the model uses a simple summing heuristic to compute the change in odds relative to the baseline odds.

The heuristic assumes $\delta_{Rt}=\delta_{Bs}=0.04$, $\delta_{Rc}=\delta_{Bc}=0.08$ and $\delta_{Rs}=\delta_{Bt}=0.12$, where Rt denotes Red triangle, Bs denotes Blue square, Rc denotes Red circle, etc. These values are in a simple ratio of 1:2:3, and are approximately equal to the exact change in odds when the first card of a given type is removed from the initial (baseline) deck, i.e., 6/8–5/7=0.04, 4/8-3/7=0.07 and 2/8-1/7=0.11. The 1:2:3 ratio reflects the relative importance of card types judged by experimental subjects, as measured in a post-game questionnaire.

The assumed heuristic avoids the need for division to compute the absolute odds. Instead, it allows the relative change (from baseline odds) to be computed more simply as the difference between $L_{R\tau}'*\delta_{R\tau}$ and $L_{B\tau}'*\delta_{B\tau}$, where each term $L'*\delta$ can be further simplified to L' additions (adjustments) of a constant value δ.

(5) Action

If the odds are not 100% then the model also accounts for random error in estimating the odds and positioning the mouse. A fuzzy function $F(O'_\tau|O_\tau)$ sets $O'_\tau=O_\tau$ with High probability; $O'_\tau=O_\tau\pm1$ button on the confidence meter with Medium probability; and $O'_\tau=O_\tau\pm2$ buttons with Low probability. Using order of magnitude estimates, we assume 90:9:1 for High:Medium:Low.

(6), (7), (8) Attention, Activation, Apprehension

These steps of the increment operation are similar to steps 1, 2 and 3 of the inspect operation.

(9) Addition

The model assumes no potential for error at this step. Any apparent accumulation errors are instead attributed to (7) activation errors and (8) apprehension errors.

Comparisons

Experiment B: Stacked Decks

Experiment B tested 25 subjects on 10 deliberately ordered decks played in balanced design. For comparison, the model was run 25 times on the same 10 decks since the fuzzy functions generate stochastic variability in the results. Figure 4 compares model to data for one deck. The values on the y-axis correspond to buttons on the confidence meter (100% Red at y=1). Results were similar for the other decks, all showing a good match between model and data, especially compared to the actual odds computed from perfect card counting and odds updating (dotted line). Sensitivity studies showed that the model performance was robust with respect to minor variations in model parameters.

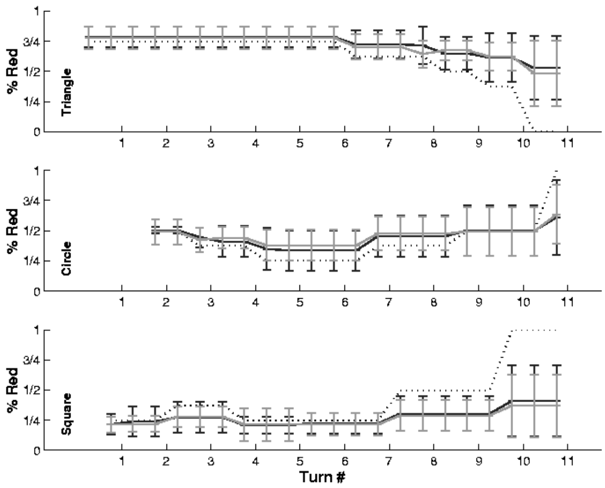

Figure 4: Odds versus turn for each track type; deck #2. Model (gray) against data (black). Lines show mean; bars show standard deviation; dotted line is actual odds.

Applications

We believe that our modeling approach strikes a useful balance between cognitive theory and practical applications. With respect to cognitive theory, our fuzzy functions capture the effects of both decay and interference (Altmann & Gray, 2002). With respect to practical applications, our fuzzy functions are simple to express and easy to program so domain experts who are not cognitive scientists can estimate the parameters and implement the algorithm. This is important because expert opinion is the primary source of model input to many applications that simulate human and organizational behavior (National Research Council, 1998). Engineering applications often use fuzzy logic to simulate human performance, and our research seeks to strengthen the scientific foundation for this practical approach through psychological modeling and empirical testing.

Acknowledgements

This research was supported by the MITRE Technology Program. Thanks to Craig Bonaceto for his help in performing the experiments and simulations.

References

Altmann, E. M. & Gray, W. D. (2002). Forgetting to remember: The functional relationship of decay and interference. *Psychological Science* 13(1), 27-33.

Burns, K. (2001). TRACS: Tool for Research on Adaptive Cognitive Strategies. www.tracsgame.com

Burns, K. (2002 a). On Straight TRACS: A baseline bias from mental models. *24th Annual Meeting of the Cognitive Science Society*. Fairfax, Virginia.

Burns, K. (2002 b). Dealing with TRACS: The game of confidence and consequence. *AAAI Fall Symposium on Chance Discovery and Management*. Falmouth, Massachusetts.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences* 24, 87-185.

Dehaene, S. (1992). Varieties of numerical abilities. *Cognition* 44, 1-42.

Dehaene, S. (1997). *The Number Sense: How the Mind Creates Mathematics* (New York: Oxford).

Kahneman, D., Slovic, P. & Tversky A., Eds. (1982). *Judgment under Uncertainty: Heuristics and Biases* (New York, NY: Cambridge University Press).

McDonald, J. (1950). *Strategy in Poker, Business and War* (New York: Norton).

National Research Council (1998). *Modeling Human and Organizational Behavior: Application to Military Simulations* (Washington, DC: National Academy Press).

Zsambok, C. E., & Klein, G. (1997). *Naturalistic Decision Making*. Mahwah, NJ: Lawrence Elrbaum.