Hidden knobs: Representations for flexible goal-directed decision-making

Romy Frömer Department of Cognitive, Linguistic, and Psychological Sciences Brown University Providence, RI 02912 romy_fromer@brown.edu

Sebastian Gluth Department of Psychology

University of Hamburg Hamburg sebastian.gluth@uni-hamburg.de Amitai Shenhav Department of Cognitive, Linguistic, and Psychological Sciences Brown University Providence, RI 02912 amitai_shenhav@brown.edu

Abstract

Sequential sampling models have been tremendously successful in describing mechanisms of decision-making at the behavioral level, and at providing testable predictions at the neural level. What is missing to date is how these same mechanisms can flexibly give rise to the broad range of decisions humans are making every day. For instance, humans can choose the best item in a set, or they can assign a value to their option set as a whole. With rare exceptions, only the computational mechanisms underlying the former type of choice have been studied. More so, our understanding of value-based decisions is dominated by decisions that identify the most valuable item or how valuable it is. Our recent work has begun to uncover the necessary transformations to additionally afford least valuable item choices. Whether and how a single sequential sampling mechanism could flexibly accommodate all these, and more, types of decisions remains a gap in our understanding. To address this gaps, we developed a theoretical framework that makes explicit the necessary representations upon which sequential sampling operates, and outlines how these representations could adjust which information is used as evidence and how it is accumulated in support of one's current choice goals. We show that this framework can parsimoniously explain behavior across a range of different choice goals by implementing and simulating behavior from an extended leaky competing accumulator model. We also generate predictions for novel choice goals to test the generality of the framework. Our framework unifies mechanisms of cognitive control and mechanisms of decision-making, and in doing so provides a novel perspective on the dimensions along which choices can differ. By rendering visible the hidden knobs that afford qualitatively different decisions using similar mechanisms, we offer novel leverage for understanding why some decisions are hard while others are not, and how poor decisions may arise.

Keywords: decision-making, cognitive control, sequential sampling, evidence accumulation

Acknowledgements

We thank Xiamin Leng and Harrison Ritz for many valuable discussions about related work that have shaped our views and inspired the present work.

1 Introduction

Humans are remarkably flexible decision-makers and able to adjust how they make decisions to their myriad - even most arbitrary - goals (cf. Fig. 1). How do they achieve this flexibility? Both, value-based and perceptual decision-making have been successfully formalized using sequential sampling models [1]. These models capture a wide variety of decisions, including choosing the best among sets of multiple options [2], assigning a rating or monetary value to options [3], and selecting responses along a continuum (e.g., redness of the kettle on a color wheel) or even in 2D space (e.g., where is Waldo?) [4]. However, since this computational work models each type of decision in isolation, it does not explain how humans can flexibly make any of those decisions within the same cognitive architecture. Here, we link research into the computational mechanisms of decision-making with research into cognitive control to render visible the hidden knobs that afford the remarkable flexibility of human decision-making.

1.1 Framework

We start from the premise that a single, flexible architecture supports a variety of goaldirected decisions and outline how a typical sequential sampling model should be extended under this assumption. Our framework integrates a sequential sampling process into a cognitive control architecture that sets the parameters of the embedded decisionmaking process according to the decisionmaker's current choice goal and the characteristics of the current decision (Fig. 2). To flexibly make decisions and adjust decision parameters accordingly, a decision-maker needs to represent the following information:

a) What is the relevant feature dimension upon which I want to decide (e.g., size versus value)?

b) How does this property translate into evidence for my goal (e.g., finding largest vs smallest)?

c) What is the response structure (e.g., are all choice options equally mutually exclusive)?

In our flexible connectionist framework, these representations consist in weight changes between nodes, e.g., input feature nodes (perceptual vs value-based) and nodes of the hidden layer that transforms inputs to evidence; nodes of that hidden layer and response nodes; as well as between these response nodes, respectively. These weightchanges afford goal-dependent adjustments to information processing, paralleling similar mechanisms in cognitive control tasks [5], and parsimoniously enable a range of decisions within one architecture.



Figure 1. When faced with a given option set, people can engage in different types of decisions, including choosing among them or appraising the set as a whole. Current models of decision-making can explain either of these in isolation, but not both.



Figure 2. Flexible decision-making architecture. Control mechanisms set the parameters on the embedded sequential sampling process. These determine (a) which type of information is selected (attention goal), (b) how this information is transformed into evidence (transformation goal), and (c) how the information is integrated to generate a choice (integration goal).

2 Model implementation

We implemented this framework by extending a connectionist, biologically plausible sequential sampling model, the leaky competing accumulator model (LCA) [6]. In the LCA, evidence at each time step t is accumulated as

$$A_t = A_{t-1} + EI - kA_{t-1} - wWA_{t-1} + sN$$
⁽¹⁾

where *A* is the matrix of response activations, *I* is an input vector containing the evidence assigned uniquely to each response via identity matrix *E*, *k* is a leak parameter that scales how much evidence is "forgotten" from one sampling point to the next, *w* is a scalar on mutual inhibition *W*, suppressing activation for each response proportional to the current activation of the alternative responses, and s is a scalar on normally distributed noise (*N*) for each option. Note that this description already foreshadows critical aspects of our extension: 1) we formalize the components as matrices that can arbitrarily vary in size, adding *E*, 2) the inputs reflect evidence, not sensory or value inputs, and 3) the evidence is accumulated at the level of responses, which are separate from stimuli. In a typical value-based decision-making task, A and I would be vectors with one entry corresponding to one option. The inputs in *I* would correspond to the average value v_{ij} of the currently relevant feature dimension sampled for each option and be added to the accumulator for that respective option only via the excitation matrix *E*. All options are equally mutually exclusive, and inhibit each other with a constant weight (cf. Fig. 3 left). For the case of four options, we can therefore rewrite the equation above as follows:

$$\begin{bmatrix} a_{1t} \\ a_{2t} \\ a_{3t} \\ a_{4t} \end{bmatrix} = \begin{bmatrix} a_{1t-1} \\ a_{2t-1} \\ a_{3t-1} \\ a_{4t-1} \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \\ v_4 \end{bmatrix} - k \begin{bmatrix} a_{1t-1} \\ a_{2t-1} \\ a_{3t-1} \\ a_{4t-1} \end{bmatrix} - w \begin{bmatrix} 0 & -1 & -1 & -1 \\ -1 & 0 & -1 & -1 \\ -1 & -1 & 0 & -1 \\ -1 & -1 & -1 & 0 \end{bmatrix} \begin{bmatrix} a_{1t-1} \\ a_{2t-1} \\ a_{3t-1} \\ a_{4t-1} \end{bmatrix} + sN$$
(2)

In what follows we will unpack how changes to *I*, *E* and *W* can give rise to a qualitatively different decision based on the same options, the appraisal of their value as a set.

2.1 Integrating evidence in accord with relevant outputs

When appraising one or multiple options (e.g., rating their value or size), responses do not map onto concrete options (e.g., "choose the bottle"),but onto discrete levels of the relevant quantity (e.g., "choose the highest value level"). In contrast to options, these levels (e.g. ratings) are also not entirely independent, but neighboring levels are more similar than levels that are farther apart. For instance, when rating an option as a 2 (out of 5), 1 and 3 are more plausible alternatives than 4 or 5. To account for this structure we make the following changes to the accumulator (cf. Fig 3): Instead of mapping sampled option values 1:1 onto accumulator evidence, samples for each option are mapped onto the response space (e.g., 5 ordinal ratings) and integrated across options. Due to the ordinal response structure of ratings, instead of evidence for each option favoring that option only (e.g., the bottle), evidence for one response also provides partial evidence for other responses proportional to their distance (e.g., evidence for rating 2 also activates rating 1 and 3 via E, cf. Fig. 3). Likewise, instead of all options equally inhibiting each other, mutual inhibition increases with response distance (via W).

2.2 Transforming sampled values into suitable inputs

In each example so far, maximizing the relevant quantity is taken as a given. However, we frequently need to select the smallest item or one that has exactly the right size (or value). In our recent work we have shown that typical response speeding with increasing overall value reverses when participants aim to choose the worst instead of the



Figure 3. Reconfiguration of the accumulator structure flexibly affords choice and appraisal of the same options.

best item [7]. We could capture these findings through goal-dependent coding of reward as inputs to an LCA. To accommodate this, we introduced a hidden layer in which value information (v_i) is transformed into goal-dependent evidence (i_i) , suitable for accumulation. Here we expand on this work and generate predictions for how this goal manipulation should impact set appraisals (liking vs disliking). We will show below that introducing the hidden layer can parsimoniously reproduce behavior in best/worst choice, *and* - combined with the changes to integration above - generate liking/disliking appraisals. We will further generate behavioral predictions for novel transformation goals, mediocrity and extremity, respectively. We assume that all our simulations equally apply to perceptual decisions based on previous work comparing value-based and perceptual decisions, and modeling either by changing the relevant input variables[8].

3 Simulations

We simulate choices and appraisals across varying set sizes and transformation goals with threshold z = 2.243, non-decisiontime $t_0 = 0.2122$ per stimulus in a set, decay k = 0.153, and rectified activation at zero. For choices, we scaled noise by s = 0.587, and mutual inhibition by w = 0.671, and provided a constant input of 1 to all options essentially implementing a collapsing bound to assure timely choices when value difference and/or input values were zero. For appraisal, noise was scaled to s = 0, and inhibition was scaled by w = 1. We simulated appraisals with five discrete response options, as in our previous work [7].

3.1 Integration Goal Simulations

Our simulations of choices and appraisals capture canonical behavioral findings (Fig. 4): Choices are faster and more consistent as the value difference between options increases, and response times further decrease as the overall value of the options increases [7]. Appraisal ratings increase with the overall value of the set [7, 9], and response times are faster for choices among more extreme compared to average overall values [10, 11].

In addition to replicating previous patterns, our model predicts interesting dissociations between appraisal and choice as set size varies. Consistent with Hick's Law and previous findings, choices are slower as the number of possible responses increases with set size. In the model this occurs due to increasing total inhibition across options [9]. However, since appraisals involve integrating across options to select among a constant set of responses, they demonstrate a different pattern. Our model predicts a non-linear effect of set size on RT due to two opposing effects. Increasing the number of options, increases the total input and non-decision time, the former speeding up decision time, the latter slowing overall RTs.

3.2 Transformation Goal Simulations



Figure 4. Flexible decision-making architecture. Control mechanisms set the parameters on the embedded sequential sampling process. These determine (a) which type of information is selected (attention goal), (b) how this information is transformed into evidence (transformation goal), and (c) how the information is integrated to generate a choice (integration goal).

Our simulations reproduce our previous findings for best/worst choice and show that our same architecture can generate goal-congruent appraisal ratings (Fig. 5). Since choice RTs speed up with increasing *magnitude* of the inputs [7], their relationship with overall value reverse as the goal changes from choosing the best to choosing the worst. However, appraisals are more sensitive to the *consistency* of the input [10]. Since inverting the appraisal goal does not change this consistency, our model predicts that this goal manipulation should not affect appraisal RTs. To show generality of this dissociation, we additionally simulated extremity/mediocrity choices and appraisals using the same option sets. Again, we find that the model can perform both tasks and that changing the choice goal strikingly changes choice RT, but much less affects appraisal RT patterns.

4 Conclusion

Our model integrates insights from cognitive control with a biologically inspired computational model of decisionmaking to offer insights into how humans flexibly decide based on their current goals. By accommodating a range



Figure 5. Transformation goals shape behavior. Our model predicts that transformation goals that symmetrically change the magnitude of the inputs (best vs worst or mediocrity vs extremity) affect choice but not appraisal RTs, which are only sensitive to changes to the consistency of the inputs (e.g., liking vs extremity lead to inverse U vs M-pattern).

of different decisions, our framework renders visible the necessary representations of one's response structure and how incoming information relates to one's goals. This change in perspective lays the ground for investigating how deviations in these representations, or poor maintenance, may give rise suboptimal decision-making. Our framework further extends beyond notions of reward and difficulty, and provides novel axes for critical properties which different kinds of decisions might share, or in which they might differ. These offer testable predictions and can give novel insights into decision-makers' performance and experience of their decisions.

References

- [1] J. R. Busemeyer, S. Gluth, J. Rieskamp, and B. M. Turner. Cognitive and neural bases of multi-attribute, multialternative, value-based decisions. *Trends Cogn Sci*, 23(3):251–263, 2019.
- [2] A. W. Thomas, F. Molter, and I. Krajbich. Uncovering the computational mechanisms underlying many-alternative choice. *Elife*, 10:e57012, 2021.
- [3] P. L. Smith. Diffusion theory of decision making in continuous report. *Psychol Rev*, 123(4):425–51, 2016.
- [4] R. Ratcliff. Decision making on spatially continuous scales. *Psychol Rev*, 125(6):888–935, 2018.
- [5] M. M. Botvinick, T. S. Braver, D. M. Barch, C. S. Carter, and J. D. Cohen. Conflict monitoring and cognitive control. *Psychological Review*, 108(3):624–652, 2001.
- [6] Marius Usher and James L. McClelland. The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3):550–592, 2001.
- [7] R. Frömer, C. K. Dean Wolf, and A. Shenhav. Goal congruency dominates reward value in accounting for behavioral and neural correlates of value-based decision-making. *Nature Communications*, 10(1):4926, 2019.
- [8] P. Sepulveda, M. Usher, N. Davies, A. A. Benson, P. Ortoleva, and B. De Martino. Visual attention modulates the integration of goal-relevant evidence and not value. *Elife*, 9:e60705, 2020.
- [9] A. Shenhav and R. L. Buckner. Neural correlates of dueling affective reactions to win–win choices. *Proceedings of the National Academy of Sciences*, 111(30):10978–10983, 2014.
- [10] M. Lebreton, R. Abitbol, J. Daunizeau, and M. Pessiglione. Automatic integration of confidence in the brain valuation signal. *Nat Neurosci*, 18(8):1159–67, 2015.
- [11] A. Shenhav and U. R. Karmarkar. Dissociable components of the reward circuit are involved in appraisal versus choice. *Sci Rep*, 9(1):1958, 2019.