

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Cognition in reach: continuous statistical inference
in optimal motor planning

Permalink

<https://escholarship.org/uc/item/4x88b75n>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 37(0)

Authors

Alonso-Diaz, Santiago

Cantlon, Jessica F

Piantadosi, Steven T

Publication Date

2015

Peer reviewed

Cognition in reach: continuous statistical inference in optimal motor planning

Santiago Alonso-Diaz, Jessica F. Cantlon, and Steven T. Piantadosi

Department of Brain and Cognitive Sciences, University of Rochester, 358 Meliora Hall, Rochester, NY, 14627
{sadiaz,spiantadosi}@bcs.rochester.edu; jcantlon@rcbi.rochester.edu

Abstract

We study the projection of cognitive representations into continuous motor (reaching) responses with a computational model that unifies three influential approaches: accumulation of evidence, statistical inference, and optimal feedback control. We modeled a number comparison task that asked participants to respond with a reaching gesture which of two side had more dots. The model successfully reproduced subjects' pattern of reach and performance across varying difficulties of numerical comparison. Our model parameterized several potentially relevant cognitive variables, including a threshold, memory decay, and mental sampling rate. Remarkably, a threshold for movement was not needed for modeling human behavior when statistical inference is combined with optimal motor planning. Overall, the model indicates that the motor-system positions the effectors optimally, both biomechanically through an optimal feedback controller, and cognitively by means of continuous statistical inference on the available evidence.

Keywords: Number Cognition; Threshold; Bayesian; Optimality; Continuous Responses

Introduction

Movement carries rich information about internal cognitive states (Song & Nakayama, 2009). Imagine a professional musician, with their movements and postures. The external motor signals that you observe are indicative of mental states and processes like feelings, moods, competence, and anxiousness. More quantitatively, recent research has been able to measure ongoing cognitive processes with continuous motor responses, mostly through reaching paradigms (e.g. Santens, Goossens, & Verguts, 2011).

One lingering question, though, is how motor structures integrate cognitive information into their plans. Here we present a computational model that uses developing cognitive states to update concurrent motor commands. Specifically, we model subjects' trajectories in a number comparison task (see Fig. 1) in which movements were modulated by the numerical ratio between the options while qualitatively maintaining an optimal profile. i.e. sigmoid positions, inverted u-shaped velocities, and sinusoidal-like accelerations (Fig. 3, first row). In brief, the model uses evidence from the approximate number system (ANS) to statistically infer the relative weighting of two parallel optimal reaching plans: one to the left and another to the right target (Fig. 2). It is parametrized with four cognitive variables, each of which has been independently proposed in this and other domains: threshold, numerical acuity, sampling rate, and sampling memory. The model reproduced human patterns in reaching,

including their overall trajectories, performance, reaction time, and changes of mind. Interestingly, thresholds were not critical for capturing human behavior. Rather, our results suggest that biomechanically optimal behavior likely employs continuous, online, and immediate statistical inference, rather than thresholded accumulation of evidence.

Experiment: methods and results

Methods

22 right handed participants (13 female; Mean age: 20.5 yrs, SD: 2.2 yrs) completed 420 trials of a number comparison task without feedback. They were asked to report the side with more dots by reaching to large clear targets on a screen with their index finger (Fig. 1). Recording was done with a Northern Digital Optotrak 3020, sampling at 200 Hz, and stopped 4 cm from the screen on the depth coordinate. Time was normalized to 101 points and only correct trials were considered. Movement started approximately 29 cm from the screen. Dots stayed on-screen for 200 ms and subjects were free to respond as soon as they appeared. Five numerical ratios were used (0.1, 0.25, 0.5, 0.75 and 0.9) equally distributed among the 420 trials. For example, a ratio of 0.1 is a trial with 1 dot vs. 10 dots (we used maximum 25 dots). The side with more dots was counterbalanced. In half of the trials both sides had equal cumulative dot area. Trials were separated into 4 blocks. Decision reaction time and movement time were defined as shown in (Fig. 1).




Figure 1: Task

A general measure of number sensitivity ω (or Weber fraction) was computed by fitting average error rate

across participants to,

$$\frac{1}{2} \operatorname{erfc} \left(\frac{|n_1 - n_2|}{\omega \sqrt{2} \sqrt{n_1^2 + n_2^2}} \right) \quad (1)$$

erfc is the complementary error function, n_1 and n_2 are the presented numbers (Pica, Lemer, Izard, & Dehaene, 2004).

Results

Average performance was 90%; by numerical ratio: 0.1=99%, SD=0.002; 0.25=99%, SD=0.003; 0.5=98%, SD=0.010; 0.75=88%, SD=0.057; 0.9=65%, SD=0.065. The Weber fraction ω for the group was 0.17 (erfc goodness of fit: $R^2 = 0.99$). Decision reaction times also varied by numerical ratio ($F(4, 84) = 6.89$, $p < 0.001$, $\eta_g^2 = 0.091$) as well as movement time ($F(4, 84) = 21.50$, $p < 0.001$, $\eta_g^2 = 0.026$).

Critically, there was a ratio-based gradient in spatial positions, as revealed by differences in mean horizontal positions ($F(4, 84) = 67.22$, $p < 0.001$, $\eta_g^2 = 0.14$). There was no effect of side of response ($F(1, 21) = 0.13$, $p = 0.72$, $\eta_g^2 = 0.003$) (contact the first author for more detailed analysis). This means that subjects produced trajectories proximal to the midline in trials with close numerical distance between the dots (e.g. ratio 0.9, green trace in Fig. 3) and more lateral positions in trials with farther numerical distance (e.g. ratio 0.25, black trace in Fig. 3). Also, note how the overall kinematic profile resembles a movement that minimizes a cost, such as jerk (Hogan, 1984). (velocity and acceleration do not end in zero because movement recording stopped 4 cm from the screen).

The model: evidence and controller

The model accumulates evidence, executes statistical inference on the potential target of behavior given the evidence, and positions the motor effector with an optimal feedback controller. This set up allowed us to explore four plausible cognitive variables involved in a decision making task that requires a number judgment: an evidence threshold (α), sample rate (λ), sample memory (ζ), and numerical confidence/acuity (γ).

Evidence

The process of accumulation of evidence contains three elements: 1) Evidence, 2) Memory, 3) Threshold. First, evidence is assumed to be produced continuously by the Approximate Number System (ANS). Specifically, evidence is sampled from the internal representation of the numerical difference $|n_1 - n_2|$ with a linear number mapping and scalar variability (Whalen, Gallistel, & Gelman, 1999),

$$s(i) \sim N \left(|n_1 - n_2|, \omega \sqrt{n_1^2 + n_2^2} \right) \quad (2)$$

We assume that samples $s = (s_1, s_2, s_3, \dots, s_i, \dots, s_n)$ are taken every λ ms (a free parameter). The total number of

samples at a given time is denoted with n and the elapsed time with t_{thr} (the counter t_{thr} stops when the threshold is hit; the subscript is to differentiate it from movement time t). A negative sample from Eq. (2) is evidence favoring the wrong number.

Second, we introduce memory to the sample series with an exponential moving average (EMA)

$$EMA(i) = \zeta \cdot s(i) + (1 - \zeta) \cdot EMA(i - 1) \quad (3)$$

with $EMA(0) = s(1)$ and ζ a free parameter between 0 and 1. Intuitively, a ζ close to zero is related to a strong memory of *early* ANS samples, and if it is close to 1 this memory effect is non-present in the sample series. Thus, EMA is a generalized version of Eq. (2), in that with $\zeta = 1$ the samples from Eq. (2) are unmodified. Alternatively, an EMA close to 0 makes the series gravitate around the initial sample, akin to a system with primacy effects in memory (e.g. Shteingart, Neiman, & Loewenstein, 2013), or that use only a few samples (e.g. Vul, Goodman, Griffiths, & Tenenbaum, 2014).

Third, accumulated evidence (AE) towards a threshold α is defined as follow,

$$AE = \sum_{i=1}^n EMA(i) \quad (4)$$

if AE reaches $\pm\alpha$ (a free-parameter), motor positioning begins. Because subjects continued to accrue confidence, as suggested by the trajectory gradient (Fig. 3), the model also continues to accumulate evidence after the threshold is hit (as paced by λ). We also tried an accumulation based on the original, non-exponentiated, samples but the results were practically identical. Thus, rather than assuming different versions of evidence in the brain, one affected by memory and the other not, it was decided to keep the same exponentiated samples for the accumulation process and the statistical inference that follows.

Controller

The model uses an optimal feedback control, a framework that has been successful at explaining a wide array of motor behavior (Todorov, 2004). In general, this framework proposes that motor commands try to minimize a cost function along the movement, while concurrently taking into account the present state of the system and task demands. This flexibility is essential for our purposes of generating movements that are affected by continuous number processing.

Specifically, we tackle the kinematic problem with a criterion of jerk (third derivative of position) minimization (Hogan, 1984) and an optimal feedback controller derived by Hoff and Arbib (1993). At each time step t , changes to the state q of the motor effector are deter-

mined by,

$$\dot{q} = \begin{bmatrix} \dot{x} \\ \ddot{x} \\ \ddot{x} \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ \frac{-60}{D^3} & \frac{-36}{D^2} & \frac{-9}{D} \end{bmatrix} \begin{bmatrix} x \\ \dot{x} \\ \ddot{x} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \frac{60}{D^3} \end{bmatrix} x_f \quad (5)$$

The controller uses the current position x , velocity \dot{x} , acceleration \ddot{x} , the remaining time $D = t_f - t$, and the end target x_f , to update the effector state. Here t refers to current movement time and t_f to total movement time (Fig. 1). As in the data, there will be different t_f depending on the numerical ratio. They will be sampled from a log-normal distribution with parameters determined by the data of the appropriate ratio. The kinematic state q is initialized at 0.

The motor system is capable of producing parallel motor plans (Cisek, 2007), and the model simultaneously generates two distinct plans at each time t : one to the left (\dot{q}_L ; $x_{fL} = -1$), and one to the right (\dot{q}_R ; $x_{fR} = 1$). Both are generated with Eq. (5), using the current position, velocity and acceleration of the effector but with different end targets. The overall kinematic state q_O is updated with a weighted sum of the left and right updates (Fig. 2),

$$q_O(t+1) = q_O(t) + w_L \dot{q}_L + w_R \dot{q}_R \quad (6)$$

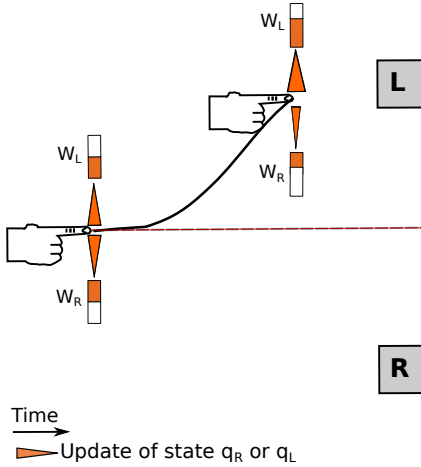


Figure 2: Model diagram

The weights w_L and w_R can be interpreted as measures of confidence of where to point to and evolve according to the sampling history. More concretely, we assume that motor plans are trying to position the effector with the best possible estimate of the numerical difference $\mu = n_1 - n_2$, given all the samples e in EMA taken up to time $t_{thr} + t$. The numerical difference will be inferred with an uniform prior $p(\mu)$, and normal likelihood $p(e|\mu, \gamma, n_1, n_2, t)$; with γ as a free parameter representing the likelihood noise. The resulting posterior $p(\mu|\gamma, n_1, n_2, e)$ is normally distributed (details in Gel-

man, Carlin, Stern, and Rubin, 2004),

$$\sim N\left(\bar{e}, \frac{\gamma}{n}\right) \quad (7)$$

where, \bar{e} is the mean of the EMA samples. The free parameter γ scales with the magnitude of the numerical distance (Harvey, Klein, Petridou, & Dumoulin, 2013; Whalen et al., 1999). This follows from the intuition that large numerical distances enhance confidence, while close numerical distances reduce it; thus γ can be interpreted as an overall confidence/acuity parameter.

To establish the weights in Eq. (6), we determined the overall probability that $\mu \geq 0$,

$$w_R = p(\mu > 0 | \gamma, n_1, n_2, e), \quad (8)$$

$$w_L = 1 - w_R \quad (9)$$

If the right side has the larger number, Eq. (8) represents the probability of being correct.

Even though the model was designed for movement trajectories, ideally it should also reproduce reaction times, percentage of correct numerical discriminations, and changes of mind. Reaction times were defined as *time to threshold* (t_{thr}) + *non-decision times*. Non-decision times refer to stimulus encoding and output generation (Ratcliff, 2014). For simplicity, we assumed encoding time to be zero and we made a proxy of output generation with lift-off time, defined as the time t when positions were greater than a random value between the accuracy (0.1mm) and resolution (0.01mm) of the recording device (Optotrak 3020). This is a naïve approach for reaction times, but it will be shown that it qualitatively replicates data patterns. As for percentage of correct discriminations, it was computed with all the trajectories that had at least 50% of points on the appropriate side, including the end position (we only plot correct trajectories). Finally, changes of mind were trajectories that fulfilled these conditions: 1) at least one point penetrated more than 1 cm of the other side; and 2) the end point was on the opposite side of 1). The model can produce changes of mind without a second post lift-off threshold (Resulaj, Kiani, Wolpert, & Shadlen, 2009) because its initial trajectory depends on the sampling history that can be modified as new samples arrive.

The model: fit and results

The model has four parameters: threshold (α), sampling rate (λ), confidence/acuity (γ), and memory (ζ). These may be viewed as formalizing *potentially* relevant processes. We fit the model to humans' behavioral data in order to determine which of these factors influence measured behavior.

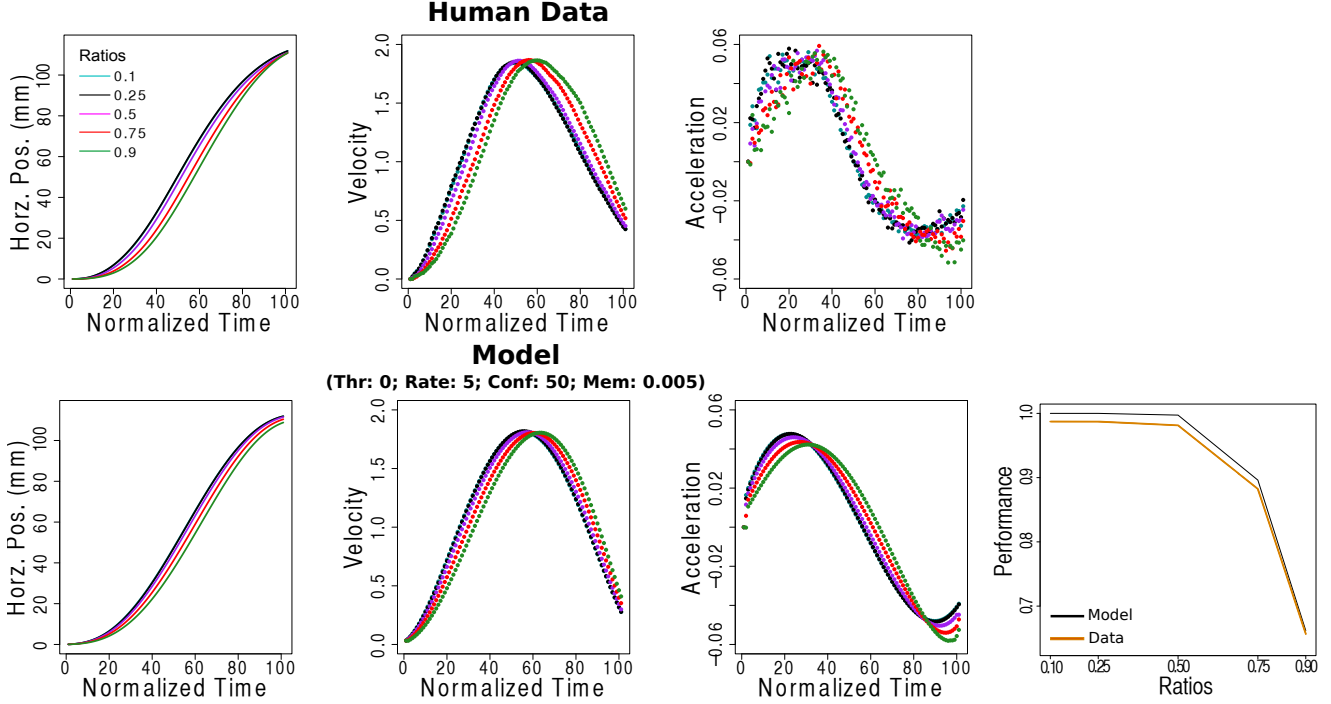


Figure 3: Kinematics and performance. First row has human data, second row model with best parameters: threshold (α), sampling rate (λ), confidence/acuity (γ), and memory (ζ). Final plot in row 2 is proportions of correct responses

Fit

Fitting was done with a grid search. The range of values for each parameter was selected heuristically based on preliminary runs of the model: $\alpha = [0, 5, 10, 50, 100, 250, 500]$; $\lambda = [1, 5, 10, 50, 150]$; $\zeta = [0, 0.005, 0.25, 0.5, 1]$; $\gamma = [30, 40, 45, 50, 100]$. For each parameter combination, we ran 240 trajectories per numerical ratio (1200 in total) and computed an average trajectory. We normalized time to 101 points and use positions up to 99% of the end target to account for the 4 cm gap in the depth axis that was not recorded and that truncated accelerations and velocities in the data (see methods and Fig. 3). A mean-squared error (MSE) was calculated for each average trajectory and summed. The best parameter combination was determined as the one that minimized this overall sum. Intuitively, this procedure just selects the best parameter combination that worked best for all numerical ratios.

The grid procedure was taken as a first manual approximation. We followed it by a finer sensitivity analysis by fixing 3 out of 4 of the best fitted parameters and ran the model with 300 values of the non-fixed parameter. We did this for all parameters with the following ranges: $\alpha = [0, 1000]$; $\lambda = [1, 1000]$; $\zeta = [0, 1]$; $\gamma = [1, 1000]$. The vectors were more dense around the best fitted values. Due to space constraints and its conceptual relevance we will report the analysis on α (threshold).

Results

The first main result is that the model was successful at reproducing the observed gradient in positions (and the velocity and acceleration profiles, but that necessarily followed given the selected controller; they are only shown for a complete visual comparison) (Fig. 3). Even though the model was only fitted to position data, the proportion of correct numerical comparisons is also remarkably similar (Fig. 3). Thus, the motor-system was rationally positioning its effector based on online cognitive information produced by the approximate number system (ANS).

The second main result is the zero value for the threshold parameter (Fig. 3). The amount of information or evidence that the model required in order to emulate the observed data was null. This means that as soon as the first sample was received the model started to position the effector.

To confirm that near-zero threshold values matched human data, we computed the sensitivity of the model to a large range of values, holding other parameters constant. As the threshold value became larger the model rapidly failed to reproduce the ratio-based gradient of the positions (Fig. 4; right panel). This is explainable because higher thresholds increase the amount of samples at lift-off, which in turn reduces the uncertainty of the estimated numerical difference (Eq. (7)), and the movement is clearly weighted in favor of one of the targets. It is worth to emphasize that the zero value for the threshold parameter is not the central argument. In fact, there

were other good thresholds, the ones close to the blue dot in Fig. 4. The critical finding here is that because the evidence was in units of numerical difference (Eq. (4)) the low threshold values in Fig. 4, including zero, indicate that movement started as soon as the first evidence arrived. This means that thresholds were irrelevant cognitive constructs in our modeling approach. In fact, removing the threshold parameter and running the grid search again produced similar results.

The model with the best parameters also captured, at a qualitative level, two other relevant behavioral effects. First, the mean and standard deviation of reaction times increased as a function of numerical ratio (Fig. 5, first row, but notice the scale difference). Though, for the hardest ratio (0.9) the mean reaction time was slightly faster than for the other hard ratio (0.75). This dip disappears with simple tweaks to the model (e.g. small changes to the best parameters or accumulating the original samples rather than EMA samples). What is interesting is that it confirms the presence of a central feature of models that accumulate evidence: RTs depend on the quality of the evidence. Due to space constraints we can not present a formal analysis of this feature but for the hardest ratio, given the best parameters, initial sampling was really favorable in correct trials and this reduced RTs for this ratio.

Second, the model produced changes of mind, but at a lower proportion than in the data (Fig. 5, bar plot). They were concentrated in the hardest ratios (0.75 and 0.9), still present in the intermediate (0.5), and non-existent for the easiest (0.1 and 0.25). Importantly, the trajectories for changes of mind were similar in the model and the participants. We hand-picked two similar ones to show that the model was able to produce changes of mind when the effector was relatively deep or shallow in the opposite side (Fig. 5, lower right panels).

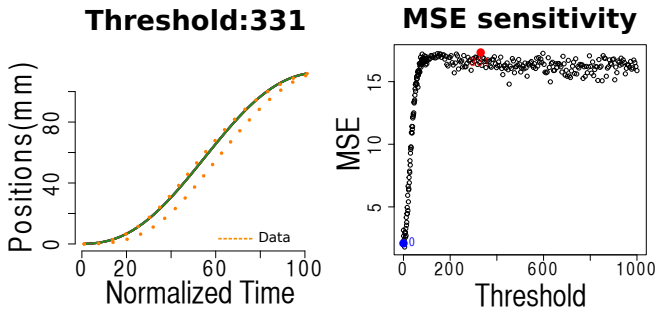


Figure 4: Sensitivity to threshold. Left panel has the position plot for the maximum MSE. Blue and red dot in the right panel are the MSE for Fig. 3 and left panel, respectively

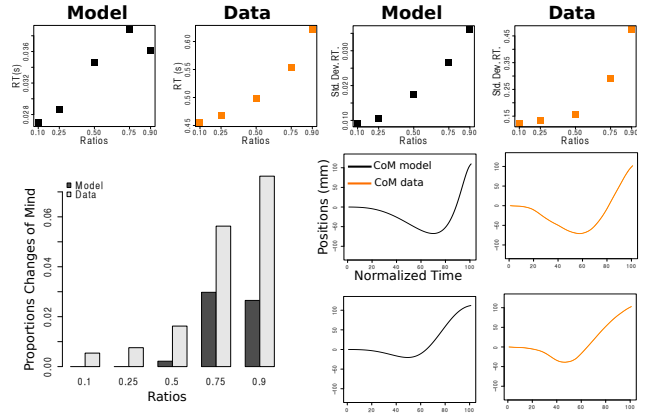


Figure 5: Reaction times and changes of mind. First row has RT (mean and standard deviations) in the data and the model (notice the scale differences). Lower left panel has proportions of changes of mind. Lower right panels have hand-picked examples with noticeable resemblance in the data and the model

Discussion

We have layout an optimal model of reach that can be modulated by ongoing cognitive processing. We now discuss several more general implications of this work.

The need of thresholds

A dominant way of thinking about decisions in the literature (perceptual or cognitive) is that the system generates evidence for the main purpose of hitting a threshold (Shadlen & Kiani, 2013). Our results pose a challenge to the framework because trajectories overlap with sufficiently large thresholds (Fig. 4). Rather, in the model thresholds are not critical. Practically-null thresholds assure that during movement time there is sufficient uncertainty in regards to which is the appropriate target to hit, and this produces the observed spatial gradient in reach. This finding may have more general implications for movement and decision-making models where thresholds are commonly employed. Behavior may appear as though subjects accumulate evidence until a threshold is passed, but our work shows that this behavior could also result from a unitary mechanism: a statistically-optimal motor planner that moves only very slowly at first due to statistical uncertainty.

Also, the model addresses a fairly underrated challenge to the threshold-framework: changes of mind. The prominent way to tackle the problem is to assume a second, post lift-off, threshold (Resulaj et al., 2009). A problem of adding additional thresholds is that there is not a principle argument against adding more. For example, if someone has a double change of mind, does this suggest a third threshold? The model presented here produced changes of mind naturally, as a result of the quality of evidence Eq. (8).

One of the strengths of threshold models is their ability to fit reaction time data (Ratcliff, 2014). The model here qualitatively reproduced the pattern of reaction times with a fairly naïve approach that did not include any encoding phase. Future work should be able to reproduce reaction times by explicitly including an encoding phase, in which harder ratios take longer to encode. This should improve the quantitative fit, but more importantly, it questions again the need for thresholds. If reaction times can be accounted with encoding and output stages, then thresholds seem unnecessary. The main advantage of continuous motor-responses is that they made more transparent output generation, hidden from discrete button responses. Thus, future work should continue to explore how cognition is reflected in continuous responses and produce a threshold-less quantitative account of RT.

The need of more complex evidence accumulation models

It is possible that more complex accumulation models can solve the "anomaly" of (practically) null thresholds, such as collapsing the thresholds over time or accumulating likelihood ratios to determine degrees of uncertainty at lift-off. While this is a potential future direction of our work, a central implication of the model here presented is that evidence can be used in richer ways beyond threshold purposes. The motor system needs a mechanism to efficiently account for evolving cognitive states. Waiting, for example, for a threshold to recompute a new course of action given a changing context might not be an efficient computational strategy. Thus, in addition to more complex diffusion models, there are further interesting questions. For instance, we proposed a linear combination of two parallel motor plans but is unclear if this is the best approach. Also, even though we sampled a random movement time for each trial, movement time was fixed. Future work should make movement times a natural consequence of the computations. Finally, at the neural level, the Bayesian inference could be done with probabilistic population codes (Ma, Beck, Latham, & Pouget, 2006) and implemented in neural networks. Those, and related questions, are critical areas for future research.

Conclusion

The model produced movements modulated by cognitive processing through weighting of optimal controllers and statistical inference. Also, we found that accumulation of evidence serves the purpose of uncertainty reduction rather than threshold arrival.

Acknowledgments

We dedicate this paper to the memory of Dave Knill, whose input and advice were invaluable to the project. We would also like to thank the members of the CAOs lab at the University of Rochester for support and feedback.

References

- Cisek, P. (2007). Cortical mechanisms of action selection: the affordance competition hypothesis. *Philosophical Transactions of the Royal Society of London*, 362(1485), 1585–99.
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2004). *Bayesian Data Analysis* (2nd Edition). Boca Raton, Florida: Chapman & Hall/CRC.
- Harvey, B. M., Klein, B. P., Petridou, N., & Dumoulin, S. O. (2013). Topographic representation of numerosity in the human parietal cortex. *Science*, 341(6150), 1123–6.
- Hoff, B. & Arbib, M. A. (1993). Models of Trajectory Formation and Temporal Interaction of Reach and Grasp. *Journal of Motor Behavior*, 25(3), 175–192.
- Hogan, N. (1984). An Organizing Principle for a Class of Voluntary Movements. *The Journal of Neuroscience*, 4(11), 2745–2754.
- Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9(11), 1432–8.
- Pica, P., Lemer, C., Izard, V., & Dehaene, S. (2004). Exact and approximate arithmetic in an Amazonian indigene group. *Science*, 306(5695), 499–503.
- Ratcliff, R. (2014). Measuring psychometric functions with the diffusion model. *Journal of Experimental Psychology: HPP*, 40(2), 870–88.
- Resulaj, A., Kiani, R., Wolpert, D. M., & Shadlen, M. N. (2009). Changes of mind in decision-making. *Nature*, 461(7261), 263–6.
- Santens, S., Goossens, S., & Verguts, T. (2011). Distance in motion: response trajectories reveal the dynamics of number comparison. *PloS one*, 6(9), e25429.
- Shadlen, M. N. & Kiani, R. (2013). Decision making as a window on cognition. *Neuron*, 80(3), 791–806.
- Shteingart, H., Neiman, T., & Loewenstein, Y. (2013). The role of first impression in operant learning. *Journal of Experimental Psychology: General*, 142(2), 476.
- Song, J.-H. & Nakayama, K. (2009). Hidden cognitive states revealed in choice reaching tasks. *Trends in Cognitive Sciences*, 13(8), 360–6.
- Spivey, M., Grosjean, M., & Knoblich, G. (2005). Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences*, 102(29), 10393–10398.
- Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience*, 7(9), 907–15.
- Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive science*, 38(4), 599–637.
- Whalen, J., Gallistel, C., & Gelman, R. (1999). Nonverbal counting in humans: the psychophysics of number representation. *Psychological Science*, 10(2), 130–137.