

A spiking neuron model of inferential decision making: Urgency, uncertainty, and the speed-accuracy tradeoff

Peter Duggins¹, Dominik Krzemiński², Chris Eliasmith³, and Szymon Wichary⁴

¹psipeter@gmail.com, Computational Neuroscience Research Group, University of Waterloo

²krzeminskidk@cardiff.ac.uk, Cardiff University Brain Research Imaging Centre

³celiasmith@uwaterloo.ca, Computational Neuroscience Research Group, University of Waterloo

⁴s.r.wichary@fsw.leidenuniv.nl, Leiden Institute for Brain and Cognition, Leiden University

Abstract

Decision making (DM) requires the coordination of anatomically and functionally distinct cortical and subcortical areas. While previous computational models have studied these subsystems in isolation, few models explore how DM holistically arises from their interaction. We propose a spiking neuron model that unifies various components of DM, then show that the model performs an inferential decision task in a human-like manner. The model (a) includes populations corresponding to dorsolateral prefrontal cortex, orbitofrontal cortex, right inferior frontal cortex, pre-supplementary motor area, and basal ganglia; (b) is constructed using 8000 leaky-integrate-and-fire neurons with 7 million connections; and (c) realizes dedicated cognitive operations such as weighted valuation of inputs, accumulation of evidence for multiple choice alternatives, competition between potential actions, dynamic thresholding of behavior, and urgency-mediated modulation. We show that the model reproduces reaction time distributions and speed-accuracy tradeoffs from humans performing the task. These results provide behavioral validation for tasks that involve slow dynamics and perceptual uncertainty; we conclude by discussing how additional tasks, constraints, and metrics may be incorporated into this initial framework.

Keywords: Neural Engineering Framework; decision making; computational model;

Introduction

In both natural and artificial environments, the dynamic and stochastic nature of decision making (DM) presents unique challenges for cognitive systems. Effectively performing DM tasks like appraising a potential mate or diagnosing an illness requires dealing with time and uncertainty, both when acquiring information and when comparing action alternatives. Although most humans perform these steps intuitively, the required cognitive operations are far from trivial: an agent must internally represent information sampled from the environment, judge the quality of that information, keep track of utilities for multiple actions, and choose when to make a decision. Furthermore, agents must flexibly adapt these operations to fit the demands of the current task: although DM would ideally proceed both quickly and accurately, either speed (time taken before making a choice) or accuracy (about the quality of the choice) must often be sacrificed in dynamic and noisy environments. We are interested in the neural and cognitive processes that underlie these operations, specifically (a) how brains flexibly manage tradeoffs between speed and accuracy, and (b) whether an integrated neural model can successfully reproduce human behavior.

Consider a DM task in which an interviewer must hire a new employee from a pool of candidates by asking them specific questions and numerically scoring their responses. Each question probes a specific *attribute*, such as intelligence, initiative, or creativity. Imperfect judgement implies the interviewer's score may not reflect the candidate's true *value* for that attribute; observations are noisy. Each attribute also has an associated *weight* indicating its importance in the interviewer's rubric. To select the best candidate, the interviewer may use the following algorithm: (1) choose a random candidate, (2) ask a random question and score the value, (3) multiply this score by the attribute's weight, and (4) add the result to the current score for that candidate. Once the interviewer has asked each candidate each question, she can select the individual with the greatest total score. Although this algorithm will produce the correct selection, it may be unnecessarily slow; for example, if one candidate receives a perfect score on each question, there is no need to continue interviewing. Various heuristics could similarly improve speed, but many come at the expense of no longer guaranteeing the correct selection.

The goal of this study is to introduce a framework for DM which is cognitively plausible, mathematically tractable, and biologically constrained. To do so, we present an anatomically mapped, spiking neuron model that performs the probabilistic inference task (PIT), a simplified version of the above interviewer task for which behavioral data is available. Model parameters corresponding to abstract variables such as uncertainty, urgency, and threshold are translated into connection weights and neural activities, bridging the gap between cognitive algorithm and biological implementation. This significantly extends previous works, which have either ignored neural plausibility or focused on implementing a single aspect of DM (e.g. evidence accumulation) in a spiking network. Our model covers multiple brain areas and diverse DM computations, producing agents that display heterogeneous behaviors. These behaviors align with strategies used by humans on the PIT. Specifically, we reproduce (a) the reaction time distributions of various individual participants, and (b) the speed-accuracy tradeoff across all participants in the study. The discussion compares our model with recent neural and computational models of DM. We conclude by discussing extensions of the model to expand its cognitive realism, to fine-tune its anatomical mapping, and to account for emotional state.

Background

Psychology

DM with multiple attributes requires many interrelated cognitive operations. Normative theories of choice rooted in economics postulate that before making a decision, humans should carefully process all available information; that is, multiply weights with attribute values and add the results. Descriptive theories of choice, on the other hand, postulate that people frequently use heuristics to simplify decision problems (Gigerenzer & Todd, 1999). Although generally quite successful, heuristics often trade off choice accuracy for speed, particularly in situations that require integrating many pieces of information. Given their simplicity, heuristics were proposed as plausible models of choice, particularly under time pressure (Rieskamp & Hoffrage, 2008) and emotional stress (Wichary, Mata, & Rieskamp, 2016). Research suggests that these factors drive individuals to make faster, less informed decisions in order to adapt to urgent or uncertain situations.

Cognitive Algorithms

Computational models have explored many aspects of DM, including the speed-accuracy tradeoff (SAT), reaction times (RT), and the roles of uncertainty, urgency, and threshold. A foundational model in the field is the drift-diffusion model (DDM), which states that evidence for alternative choices is accumulated in a dynamic, noisy manner until the process reaches a fixed decision threshold, at which point a choice is made. Although the DDM has been very successful in modelling human behavior across many domains (Ratcliff & McKoon, 2008), it is a purely mathematical model, and is therefore agnostic about neural implementation.

One important extension of DDM is the notion of time-varying decision thresholds, which may be realized by modulating the gain of accumulating sensory signals by an urgency variable that increases over time. The existence of an urgency signal is consistent with behavioral data (Ditterich, 2006), neural data (Cisek, Puskas, & El-Murr, 2009), and simple neural-network implementations (Standage, You, Wang, & Dorris, 2011). Anatomically-detailed computational models also have investigated variable decision thresholds (Standage et al., 2011; Frank, 2006): these models suggest that decision conflict (uncertainty about stimulus information) may compliment urgency signals in flexibly controlling the SAT.

Neuroanatomy

DM emerges from the interplay of distinct cognitive operations associated with activity in distinct brain structures. To understand how these operations align with the algorithms discussed above, and to anatomically map them onto a spiking neuron model, we draw from a growing literature documenting the neural correlates of DM. However, given the redundancy of neural computation and the diversity of operations required by DM tasks, these functional assignments remain a work-in-progress (see the Discussion).




Figure 1: Probabilistic Inference Task. See text for details.

Perceptual inputs on DM tasks are processed through various sensory hierarchies, producing high-dimensional cortical representations. To ascribe meaning to this information, an individual must draw on previous experience or abstract knowledge of task requirements: the retrieval and assessment of salience involves numerous brain areas that are outside of the current scope. However, it appears that the binding of sensory representations with this salience signal occurs especially in ventromedial prefrontal cortex and orbitofrontal cortex (OFC); these populations effectively weigh external stimuli by their relevance to the current task (Rangel & Clithero, 2014). Because perception and evaluation evolve over time, internal buffers are needed to track the accumulation of relevant information. Such a working memory retains the weighted evidence (utility) for various choice alternatives, and may be realized in dorsolateral prefrontal cortex (dlPFC) and sensorimotor areas (Thura & Cisek, 2014).

Before an action is selected, activity in certain regions may modulate DM to meet contextual demands. Areas such as the anterior cingulate cortex (ACC) and inferior frontal cortex (IFC) are thought to shift decision thresholds by gating inputs from presupplementary motor area (pSMA) to the basal ganglia (BG), effectively delaying decisions until more information is available (Forstmann et al., 2010). Given the role of ACC and IFC in monitoring cortical representations, this may realize an urgency- or uncertainty-based modulation of decision criteria (Aron, Robbins, & Poldrack, 2014). Furthermore, projections from locus coeruleus (LC) to these cortical regions may utilize norepinephrine to mediate arousal, further manipulating decision thresholds or evidence accumulation based on an agent's emotional state (Aston-Jones & Cohen, 2005); this would provide another mechanism whereby urgency (and affect) could mediate DM (Murphy, Boonstra, & Nieuwenhuis, 2016).

Finally, the brain must select a single action and inhibit any alternatives. This winner-take-all (WTA) competition between action representations is realized by recurrent connections between (and within) cortex and BG (Bogacz & Gurney, 2007; Frank, 2006). Once alternative actions have been suppressed, motor commands may be executed by projecting the cortical representation back to cerebellum, brainstem, and spinal cord.

Task

In the probabilistic inference task (PIT), participants begin by memorizing weights associated with six attributes. The

weights in this task are compensatory: they do not differ significantly from one another, encouraging participants to attend to each attribute rather than (heuristically) discard the low-weight attributes. During the task, participants are simultaneously shown the values (0 or 1) of objects *A* and *B* for one attribute. The participants respond by pressing one of three buttons, indicating their selection of *A*, selection of *B*, or a request for more information. If they choose the later, the current display is replaced by another pair of values for the next attribute. This is repeated until the participant makes a choice or until all attributes have been exhausted, at which point a choice is forced. Each participant performs the task 48 times; behavior on each trial is quantified by the number of attributes requested before the final decision (RT) and whether the choice corresponded to the highest-value option (accuracy). See Fig. 1 for a schematic and the original paper (Wichary, Magnuski, Oleksy, & Brzezicka, 2017) for further experimental details.

Neural Engineering Framework

The NEF (Eliasmith & Anderson, 2003) describes how spiking neural activity may represent a time-varying, vector-valued signal $\mathbf{x}(t)$ such as value, weight, or evidence. A neuron spikes most frequently when presented with its particular “preferred stimulus” and responds less strongly to increasingly dissimilar stimuli (i.e. values of $\mathbf{x}(t)$). In the NEF, each neuron *i* is accordingly assigned a preferred direction vector, or *encoder*, $\mathbf{e}_i$. To produce a variety of tuning curves that match electrophysiological variance within the brain, each neuron is also assigned a unique gain α_i and bias β_i . These quantities determine how strongly an incident vector $\mathbf{x}(t)$ drives the neuron:

$$I_{in}(t) = \alpha_i * (\mathbf{e}_i \cdot \mathbf{x}(t)) + \beta_i \quad (1)$$

where $I_{in}(t)$ is the current flowing into the neuron and $(\cdot)$ is the dot product between the encoder and input vector. So long as there is a well-defined relationship between input current and resulting firing rate, the neuron’s activity can be said to encode the vector $\mathbf{x}(t)$. A distributed encoding extends this notion: if $\mathbf{x}(t)$ is fed into multiple neurons, each with a unique tuning curve defined by $\mathbf{e}$, α , and β , then each neuron will respond with a unique spiking pattern $a_i(t)$, and the collection of all neural activities will robustly encode the signal.

For neural encoding to be meaningful, there must be methods to recover, or decode, the original vector from the neurons’ activities; together, encoding and decoding constitute neural *representation*. The NEF identifies neural *decoders* $\mathbf{d}_i$ that either perform this recovery or compute arbitrary functions, $f(\mathbf{x})$, of the represented vector. A functional decoding with $\mathbf{d}_i^f$ allows networks of neurons to *transform* the signal into a new state, which is essential for performing operations such as value-weight multiplication. To compute these transformations, a linear decoding is applied to the neural activities:

$$\hat{f}(\mathbf{x}(t)) = \sum_{i=0}^n a_i(t) * \mathbf{d}_i^f, \quad (2)$$

where a_i is the spiking activity of neuron *i* (smoothed by a

lowpass filter), *n* is the number of neurons, and the hat notation indicates that the computed function is an estimate. To find decoders $\mathbf{d}_i^f$ that compute the target function, we use least-squares optimization to minimize the error between the target value $f(\mathbf{x}(t))$ and the decoded estimate $\hat{f}(\mathbf{x}(t))$. The general-purpose Nengo neural simulator (Bekolay, Laubach, & Eliasmith, 2014) optimizes these decoders for the specified transformations. Connection *weights* between each presynaptic neuron *i* and each postsynaptic neuron *j* combine encoders and decoders into a single value used during simulation

$$w_{ij} = \alpha_j \mathbf{e}_j \cdot \mathbf{d}_i^f. \quad (3)$$

Finally, the NEF specifies methods to build neural networks that implement any dynamical system, including linear systems of the form $\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t)$. To do so, the matrices *A* and *B* must be modified to account for the dynamics that naturally occur when using neurons with non-instantaneous synapses. Nengo performs this optimization for the specified target dynamics; this is essential for constructing networks that include the recurrent connections required for working memory, evidence accumulation, and choice competition.

Mathematics aside, the NEF provides a framework for combining distinct cognitive subsystems into a coherent agent while respecting biological constraints. As explored in (Eliasmith, 2013) and numerous subsequent papers, models built with the NEF unify perception, cognition, and action using a standardized representational format. This can be contrasted with models that focus on detailed implementation of isolated cognitive processes, such as winner-take-all competition via lateral inhibition; although these models provide key insights into the neural basis of particular algorithms, we are interested in unifying several of these components into a stand-alone agent. As we show in this paper, such agents can be presented with inputs resembling the prompts humans receive during a task, perform internal operations that can be measured and compared with neural data, and produce outputs with well-defined action-space mappings for behavioral comparison. Furthermore, the integrated neural-, synaptic-, and network-level dynamics of NEF models are essential for studying temporally-extended DM. Finally, the complexity of the perceptual space, internal representations, and motor commands can be adjusted to model DM across many domains.

Model

The model architecture is summarized in Fig. 2 and outlined here, with detailed descriptions of key components in the following paragraphs. A two-dimensional vector representing the values of *A* and *B* for the currently displayed attribute is provided as external input to a population labelled **OFC**. This population also receives the “remembered” attribute weights; to reduce model complexity, we model this recall process as a noisy perturbation of externally-supplied weights. Neurons in **OFC** thus represent both the perceived attribute value and the remembered weight for the current attribute. Connection weights between **OFC** and **dIPFC** multiply values by weights




Figure 2: Model Schematic. Boxes are inputs and outputs, circles are spiking neuron populations. Red connections are inhibitory. See text for details on represented quantities and cognitive operations.

and send the result to the two-dimensional **dlPFC** population. Recurrent connections within **dlPFC** implement integration, leading to the accumulation of evidence from **OFC** as additional attributes are presented. When the difference between accumulated evidence for *A* and *B* exceeds a dynamic threshold, neurons in **rIFC** disinhibit the **pSMA** population. This allows information to flow from **dlPFC** through **pSMA** to the **BG** network. In the **BG**, mutual inhibitory competition selects the option with the greatest evidence as a final output. If **BG** does not output a selection after one second of input, then the next attribute is presented for one second, and so on.

The **dlPFC** population is a neural *integrator*, a system which maintains its currently represented value while additionally incorporating any inputs. This network has previously been used in neural models of working memory, where it has reproduced activity in PFC and behavior on several WM tasks (Eliasmith, 2013). The system is described by the target dynamics $\dot{\mathbf{x}} = B\mathbf{u}$; notice that changes in the represented value $\dot{\mathbf{x}}$ do not depend on the represented value $\mathbf{x}$ itself, but only on the input $\mathbf{u}$. This implies that the integrator will remember the current evidence perfectly, and steadily add any input evidence to arrive at a new value. However, because the integrator is implemented in noisy spiking neurons, the feedforward and recurrent connections do not perfectly implement these dynamics: evidence is slightly distorted and prone to drift.

The **rIFC** population determines when the accumulated evidence is sufficient to make a decision. Connections between **dlPFC** and **rIFC** compute the function

$$\mathbf{x}_{\text{unc}}(t) = T - |A(t) - B(t)|, \quad (4)$$

where the threshold T is a free parameter and the absolute difference is between accumulated evidence for *A* vs *B*. This signal can be interpreted as the current uncertainty about action selection, with larger values implying greater internal conflict. Neurons in **rIFC** have tuning curves with positive slope and an intercept at $\mathbf{x} = 0$, ensuring that neural activities remain positive for $\mathbf{x} > 0$ and go silent for $\mathbf{x} < 0$. **rIFC** connects to **pSMA** with strong inhibitory connections, such that any activity in **rIFC** dampens all activity in **pSMA**, restricting the flow of information to **BG** and delaying a decision.

Consequently, when $|A(t) - B(t)|$ exceeds the threshold T , disinhibition opens the gate, activating WTA competition (see below) and producing a decision with reaction time $\text{RT} = t$.

To simulate an urgency signal that modulates the SAT, we model an additional input $\mathbf{x}_{\text{urg}}$ to **rIFC**. This input grows linearly during the course of the trial and is additive with $\mathbf{x}_{\text{unc}}$,

$$\mathbf{x}_{\text{urg}}(t) = -R * t, \quad (5)$$

where R is a free parameter governing the magnitude of the urgency signal. Large R implies that time pressures cause uncertainty to shrink quickly during a trial. We also set $\mathbf{x}_{\text{urg}}(t)$ to a large negative value when $t > 5.8\text{s}$, which removes gating at the end of the trial and forces a selection.

The **BG** population is based off an anatomical reconstruction of the BG and implements WTA competition between action alternatives (Stewart, Choo, & Eliasmith, 2010). As with the neural integrator, this network has been used in numerous functional brain models as part of the action selection system (Eliasmith, 2013). Here, it is used to select the action with the greatest evidence value by engaging a WTA circuit that produces a single motor output (choose *A* or choose *B*).

Each model agent is initialized with unique T , R , and random seed for generating neuron parameters $\mathbf{e}$, α , and β . As with human participants, each agent repeats the task 48 times, and its RT and accuracy on each trial are recorded.

Results

To clarify model dynamics, we begin by looking at time series for the state variables represented in spiking neural populations. Fig. 3 shows an easy trial from an agent with moderate urgency and threshold ($R = 0.4, T = 2.5$). On this trial, choice *B* has positive values for the first three attributes, while choice *A* has zero value. The difference in accumulated evidence (**dlPFC**) steadily grows (blue vs. red line), diminishing the activity of the “hold” population (**rIFC**, black line) until its neurons cease firing around $t = 2.3\text{s}$. With the inhibitory hold on the **pSMA** gate removed, information flows from **dlPFC** to **BG**, where WTA dynamics quickly selects the dominant choice (*B*), leading to nonzero motor output by $t = 2.4\text{s}$.




Figure 3: An easy trial with a typical agent ($R = 0.4, T = 2.5$) leads to an accurate selection with moderate RT (2.4s).

Fig. 4 shows a hard trial from an agent with low urgency and high threshold ($R = 0.3, T = 3.0$). On this trial, attributes for both A and B have numerous positive values, making an accurate choice difficult. As a result, the difference in accumulated evidence does not remove the inhibitory hold. At the end of the trial ($t = 5.8s$), x_{urg} is externally boosted, opening the gate and forcing a decision. By this point, noise-induced errors have accumulated in the **dIPFC** representation, resulting in a poor estimate of the total evidence for each option (solid vs. dashed lines). This drift favors B , leading to the correct selection; in other trials, such errors produce incorrect responses.




Figure 4: A hard trial with a patient agent ($R = 0.3, T = 3.0$) leads to stochastic selection with large RT (5.9s).

Fig. 5 shows a trial from an agent with high urgency and low threshold ($R = 0.5, T = 2.0$). In the third trial, the evidence initially favors B ; because this agent is impatient, its decision criteria is met by $t = 3.4s$, leading to selection of B . However, the evidence from subsequent attributes favors A , making the agent's early choice ultimately incorrect.

To explore the SAT predicted by the model, we first examined individuals (both humans and agents) who strongly favored either speed or accuracy. Fig. 6 compares the RTs of fast, inaccurate decision makers and slow, accurate decision makers across all 48 experimental trials. As expected, an




Figure 5: A trial from an impatient agent ($R = 0.5, T = 2.0$) leads to a hasty (RT = 3.4s) and inaccurate selection.




Figure 6: RT distributions of agents and humans with opposing DM strategies. Legend indicates mean accuracy.

agent with high urgency and low threshold ($R = 0.5, T = 1.8$) typically makes selections after 2 – 3 attributes but has low accuracy (73%). Conversely, an agent with low urgency and high threshold ($R = 0.32, T = 3.0$) typically views 4 – 6 attributes before making an accurate choice (92%). The RTs and accuracies of both agents are closely aligned with the behavior of individual human participants performing the PIT.

We also generated a population of agents with random parameters in the range $R \in [0.3, 0.5]$, $T \in [2.0, 3.0]$, then had each agent perform the 48 experimental trials. Fig. 7 plots the mean RT of each agent against its mean accuracy (across trials). The SAT is readily apparent: larger RTs correlate with more accurate decisions, and the slopes of the best-fit trendlines are remarkably similar between humans and agents.

Discussion

Our model was designed to recreate the anatomy and cognitive function of the human DM system using populations of spiking neurons. Although DM is too complex to capture in a simple network, we believe that our computational model distills many of its core features into a functionally clear circuit which aligns with the central tenants of the DDM. **OFC**




Figure 7: The speed-accuracy trade off. Mean accuracy increases as a function of mean RT among humans and agents.

receives inputs from sensory and memory systems and passes weighted values to **dIPFC**. These inputs are incorporated into actively maintained **dIPFC** representations of the accumulated evidence for each choice alternative. Connections between **dIPFC** and **BG** are responsible for sending the utility of potential actions to an action selection system, but these connections are gated by intermediate structures like **pSMA**, ensuring that individuals wait for sufficient information before executing a behavior. Only once an inhibitory hold, represented in areas like **rIFC**, has been released will **BG** be freed to select the best action and engage a motor response. This delay is modulated by signals representing uncertainty and urgency, which are estimated from the current evidence and the elapsed time, and together determine the SAT implemented by the agent.

We showed that our model agents reproduced human RTs and SATs on this task: agents with small thresholds and high urgency chose quickly but inaccurately, while agents with large thresholds and low urgency chose slowly but accurately. The distributions of model RTs matched both extremes of human behavior (Fig. 6), and the SAT evident across a population of agents matched the trend from a small human dataset (Fig. 7). These behavioral results provide further support for the urgency-gating hypothesis while demonstrating that the proposed cognitive algorithms can be realized in an anatomically-mapped spiking neural system. In particular, they highlight the role of modulatory signals in dynamically controlling a neural implementation of the DDM, and indicate that individual variability may be explained by two parameters: the slope of the urgency signal R and the decision threshold T .

Dynamic decision thresholds have been explored in previous neural and computational models, with similar findings regarding the SAT. Our model stands apart in several respects. First, our agents model DM from start-to-finish: they receive inputs that resemble the attributes presented to humans, process that information using spiking neural activities and weighted synaptic connections, and produce outputs that are directly comparable to human behavioral data. While

other models explore the mathematical basis for urgency-gated DDM (Ditterich, 2006), the neural activities associated with various populations (Cisek et al., 2009), or the detailed neuroanatomy of particular cognitive algorithms (Frank, 2006), our framework unifies these perspectives into a single model. This unification provides clarity about how neural activities relate to cognitive operations, mathematical representations, and behavior, providing numerous means for model validation. In this paper, we chose to focus on behavioral validation; future work should compare the neural activities predicted by our model with measured activities from the relevant areas.

Our model also has a well-defined anatomical mapping that includes the principal areas associated with human DM. By including both cortical and subcortical structures, we capture the dynamic interplay of weighted evidence accumulation and thresholded action selection. While some models attribute these process exclusively to competition within cortex (Standage et al., 2011) or BG (Frank, 2006), we believe that both anatomic divisions are active when the human brain engages in high-level, abstract DM. Although simpler tasks (e.g., classifying random dot motion) may be solved by isolated neural systems, we hypothesize that complex, temporally-extended tasks (e.g., interviewing job candidates) will involve both cortex and BG. Our model proposes an initial framework for this coupling, one that can be expanded to include more anatomical and cognitive detail, as discussed below.

Although our model strives for neural and anatomical detail, it remains biologically unrealistic in several respects. With respect to the model’s fundamental units, we use LIF neurons with mixed-sign connection weights and current-based synapses. This makes our model more biologically-detailed than (Standage et al., 2011), whose network is built from rate-approximated cortical columns and connected using weights drawn from mathematical distributions, but less biologically-detailed than (Frank, 2006), whose network is built from biophysical neuron models and connected using a reconstruction of BG circuitry that permits STDP learning. However, both these models emphasise validation with neural activities, whereas we emphasise behavioral validation. With respect to cognitive algorithms, we implement urgency-induced gain by adding a linearly-ramping signal to the decision variable. Many alternative forms of this function are possible: urgency may ramp according to a sigmoid function, and the effects of urgency may be multiplicative rather than additive, if this signal is e.g. modulated by noradrenergic inputs from LC (Aston-Jones & Cohen, 2005). With regards to anatomy, the exact targets of urgency modulation are unclear: while our model posits that urgency causes **rIFC** to disinhibit **pSMA**, some evidence suggests that urgency increases the rate of evidence accumulation in sensorimotor areas (Thura & Cisek, 2014). Similarly, the decision threshold may be realized in any of several areas, including **rIFC**, **pSMA**, lateral intraparietal area (Standage et al., 2011), striatum, or subthalamic nucleus (Frank, 2006). Further empirical and modelling work is needed to clarify the contribution of each area.

Future work can profitably proceed in several directions. In the current model, attribute weights were provided as inputs, whereas humans performing the task were required to memorize the weights beforehand. To increase the model's cognitive realism, we plan to add an associative memory system that, when presented with a high-dimensional vectors representing the current attribute, recalls the associated one-dimensional weight value. Online learning rules would be used to train these associations during an initial phase of the task. When presented with an attribute at test time, this system would return a noisy recall of the associated weight, which would then be routed to **OFC** as in the current model. We are also interested in extending the model to more complex DM tasks, including the interview task described in the introduction. This task involves more choice alternatives, a wider range of attribute weights, and a less strict schedule: participants must choose which option to investigate next, rather than being presented values and weights in a fixed order. This last feature introduces an extra dimension of exploration, in which the decision about which "questions to ask" may interact in interesting ways with the current evidence for various options. Finally, experimental data is available for a variant of the PIT in which participants are shown aversive images before testing, priming them with an emotional state that affects DM (Wichary et al., 2016). We plan to investigate whether arousal-mediated changes in the urgency signal can capture such emotional biases.

Conclusion

We leveraged anatomy, a cortical integrator, a detailed basal ganglia model, the drift-diffusion model, and the Neural Engineering Framework to build a spiking neuron model of urgency- and uncertainty-gated decision making. We showed that model agents with varying degrees of urgency and decision threshold reproduced the reaction times and the speed-accuracy tradeoff of humans performing a probabilistic inference task.

Acknowledgments

This work was funded by Marie Skłodowska-Curie Fellowship (Project no. 791181), an EPSRC PhD studentship (EP/N509449/1), CFI and OIT infrastructure funding, the Canada Research Chairs program, an NSERC Discovery grant (261453), and an AFOSR grant (FA9550-17-1-0026).

References

Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2014). Inhibition and the right inferior frontal cortex: one decade on. *Trends in cognitive sciences*, 18(4).

Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.*, 28.

Bekolay, T., Laubach, M., & Eliasmith, C. (2014). A spiking neural integrator model of the adaptive control of action by the medial prefrontal cortex. *The Journal of Neuroscience*, 34(5).

Bogacz, R., & Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural computation*, 19(2).

Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: the urgency-gating model. *Journal of Neuroscience*, 29(37).

Ditterich, J. (2006). Evidence for time-variant decision making. *European Journal of Neuroscience*, 24(12).

Eliasmith, C. (2013). *How to build a brain: A neural architecture for biological cognition*. Oxford University Press.

Eliasmith, C., & Anderson, C. H. (2003). *Neural engineering: Computation, representation, and dynamics in neurobiological systems*. MIT press.

Forstmann, B. U., Anwander, A., Schäfer, A., Neumann, J., Brown, S., Wagenmakers, E.-J., ... Turner, R. (2010). Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *Proceedings of the National Academy of Sciences*, 107(36).

Frank, M. J. (2006). Hold your horses: a dynamic computational role for the subthalamic nucleus in decision making. *Neural Networks*, 19(8).

Gigerenzer, G., & Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In *Simple heuristics that make us smart*. Oxford University Press.

Murphy, P. R., Boonstra, E., & Nieuwenhuis, S. (2016). Global gain modulation generates time-dependent urgency during perceptual choice in humans. *Nature communications*, 7(1).

Rangel, A., & Clithero, J. A. (2014). The computation of stimulus values in simple choice. In *Neuroeconomics*. Elsevier.

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural computation*, 20(4).

Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: How opportunity costs affect strategy selection. *Acta psychologica*, 127(2).

Standage, D., You, H., Wang, D., & Dorris, M. C. (2011). Gain modulation by an urgency signal controls the speed-accuracy trade-off in a network model of a cortical decision circuit. *Frontiers in computational neuroscience*, 5.

Stewart, T. C., Choo, X., & Eliasmith, C. (2010). Dynamic behaviour of a spiking model of action selection in the basal ganglia. In *Proceedings of the 10th international conference on cognitive modeling*.

Thura, D., & Cisek, P. (2014). Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron*, 81(6).

Wichary, S., Magnuski, M., Oleksy, T., & Brzezicka, A. (2017). Neural signatures of rational and heuristic choice strategies: a single trial erp analysis. *Frontiers in human neuroscience*, 11.

Wichary, S., Mata, R., & Rieskamp, J. (2016). Probabilistic inferences under emotional stress: how arousal affects decision processes. *Journal of Behavioral Decision Making*, 29(5).