Biologically Inspired Cognitive Architectures (2014) 10, 61–67

Available at www.sciencedirect.com

ScienceDirect

journal homepage: www.elsevier.com/locate/bica

A model of proactive and reactive cognitive control with anterior cingulate cortex and the neuromodulatory system

Matthias D. Ziegler [a,*], Suhas E. Chelian [a], James Benvenuto [a], Jeffrey L. Krichmar [b], Randall O'Reilly [c], Rajan Bhattacharyya [a]

[a] *HRL Laboratories, LLC, Malibu, CA, United States*
[b] *University of California Irvine, Irvine, CA, United States*
[c] *University of Colorado at Boulder, Boulder, CO, United States*

Received 7 November 2014; accepted 7 November 2014

KEYWORDS
Cognitive control;
Anterior cingulate cortex (ACC);
Neuromodulators

Abstract
Proactive and reactive cognitive control are often associated with anterior cingulate cortex (ACC). How ACC affects processing in other brain areas, however, is often not explicitly delineated. In this work, we describe a model of how ACC computes measures of conflict and surprise that are in turn relayed to the basal forebrain (BF) and locus coeruleus (LC) in that order. BF and LC signals then respectively sharpen posterior cortical processing and trigger the reframing of prefrontal cortical decision-making frames. We implemented this theory in a large-scale neurocognitive model that performs simulated geospatial intelligence tasks. Experiments demonstrate improved performance while minimizing additional processing. Alternate interpretations of neuromodulatory signals are also discussed.

Introduction

Anterior cingulate cortex (ACC) is implicated in a number of functions including proactive and reactive cognitive control (see Alexander & Brown, 2010 for a review). Proactive control is applied before feedback is given while reactive control is applied after feedback. For example, in taking a test, checking a calculation before submitting an answer

* Corresponding author.
E-mail addresses: mdziegler@hrl.com (M.D. Ziegler), sechelian@hrl.com (S.E. Chelian), jbenvenuto@hrl.com (J. Benvenuto), jkrichma@uci.edu (J.L. Krichmar), randy.oreilly@colorado.edu (R. O'Reilly), rbhattac@hrl.com (R. Bhattacharyya).

http://dx.doi.org/10.1016/j.bica.2014.11.008
2212-683X/© 2014 Published by Elsevier B.V.

and then after receiving results, changing the amount of studying one does for the next test are examples of proactive and reactive control in that order.

ACC models of cognitive control (e.g., Alexander & Brown, 2011; Shenhav, Botvinick, & Cohen, 2013) often do not explicitly state how control is exerted on other brain areas in terms of neuronal dynamics. One appealing candidate for the agent of this control is the neuromodulatory system. The neuromodulatory system is bidirectionally connected with ACC, prefrontal cortex and posterior cortical areas (Krichmar, 2008). Within the neuromodulatory system, the basal forebrain and locus coeruleus are of particular interest. The basal forebrain (BF) is the source of acetylcholine, and appears to modulate attention. For example, BF activity has been shown to increase attentional modulation in downstream cortical targets (Disney, Aoki, & Hawken, 2007; Goard & Dan, 2009; Herrero et al., 2008). BF inputs to posterior cortical neurons amplify cue detection and may also act to suppress irrelevant distractors (Broussard, Karelina, Sarter, & Givens, 2009). BF activity may be regulated by prefrontal cortex (PFC) to enable top-down control of attention (Nelson, Sarter, & Bruno, 2005). The locus coeruleus (LC) is the source of noradrenaline. Its effects are prevalent in cortical regions when there are dramatic environmental changes which cause large shifts in attention (Dalley et al., 2001) as opposed to the more gradual shifts in attention encoded by the BF. Thus, it has been proposed that the LC functions as a ''network reset'' to detect large changes, reject prior expectations, and enable the formation of new models of the environment (Bouret & Sara, 2005; Sara, 2009). Models of neuromodulatory control, however, often focus on perceptual or attentional processing (e.g., Avery, Dutt, & Krichmar, 2014; Avery, Nitz, Chiba, & Krichmar, 2012).

In this work, we present a model of cognitive control over more deliberate processing such as probability inference calculations and assessing the utility of various factors in spatial reasoning tasks. We embed our model in a large-scale neurocognitive model implemented in Emergent (Aisa, Mingus, & O'Reilly, 2008) that includes nine brain areas (parietal cortex, temporal cortex, hippocampus, prefrontal cortex, basal ganglia, anterior cingulate cortex, orbitofrontal cortex, neuromodulatory system, and thalamus). The multi-region model performs several simulated geospatial tasks each with multiple steps for inference and decision making. For proactive control, we use a measure of conflict to consider a ''look—relook'' decision. For reactive control, we use a measure of surprise, derived from error, to control the use or disuse of features in a spatial reasoning task. Proactive and reactive control increased our ability to model human performance on the simulated intelligence tasks while minimizing extraneous processing.

Materials and methods

Tasks

In this work, we consider two tasks—tasks 2 and 3—that are a subset of six simulated geospatial intelligence tasks more fully described in Lebiere et al. (2013). (We do consider task 1 because it was an introductory task with limited complexity to introduce humans to subsequent tasks. Tasks 4—6, on the other hand, emphasized updating probability estimates as opposed to spatial processing.) One hundred-three subjects were recruited from intelligence analysis graduate studies programs to complete the tasks and we modeled their inference and decision making using a large-scale neurocognitive model. Our goal was to produce quantitatively similar distributions of responses to human response. Tasks were not specifically designed to elicit conflict or surprise, but we found modeling these mechanisms increased our fidelity in modeling human behavior.

Tasks 2 and 3 are very similar: given a map that displays a history of previous attacks from four groups, subjects attempt to assess which group is most likely to attack at a new probe location. Subjects were told to judge the probability of a group attacking based on three features: (1) distance from each group's center to the probe location (the closer the probe location is to a group's centroid, the more likely that group is to attack), (2) the radius associated with each group's attacks (if a probe location is equidistant from two groups' centroids, the group with a larger radius is more likely to attack—that group's region of interest is larger than the other group), and (3) the base rate of each group's activity (with equal distance and radius, the group with the higher base rate is more likely to attack—that group attacks more frequently). In task 2, subjects were asked to consider distances ''as the crow flies'' (Fig. 1a) while in task 3, a network of roads meant subjects had to consider distances ''as the cow walks'' (Fig. 1b). Each subject performs five trials where each trial consists of a set of twenty attack histories that appear sequentially as icons on the screen. Each trial builds on the last and events are not erased. Each attack location is determined by a 2D Gaussian distribution with mean (or center) and variance (or spread) that does not change for the duration of the five trials. Neither the process to generate attack locations nor its parameters were known to the subjects.

Model

Our model is composed of posterior cortical areas, frontal cortical areas, and neuromodulatory areas as illustrated in Fig. 2. Proactive and reactive control emerges through the interaction of these brain areas.

Posterior cortex: probability calculations

Within posterior cortex, Parietal Cortex (PC) was responsible for calculating the probability that each group was likely to attack at the target location; full details can be found in Sun and Wang (2013). This calculation took place using either a subset or all of the possible features—distance, radius, and base rate—provided. If a single feature was used, the marginal attack probability estimate was used (e.g., $P(\text{attack}|\text{distance})$). If more than one feature was used, an average across marginal attack probability estimates was performed (e.g., $0.5 * P(\text{attack}|\text{distance}) + 0.5 * P(\text{attack}|\text{base rate})$). This simulates the blending mechanism that approximates average human behavior when combing two or more sources of information (Lebiere, 1999). Attack probability estimates were computed sequentially for all four groups and then held in

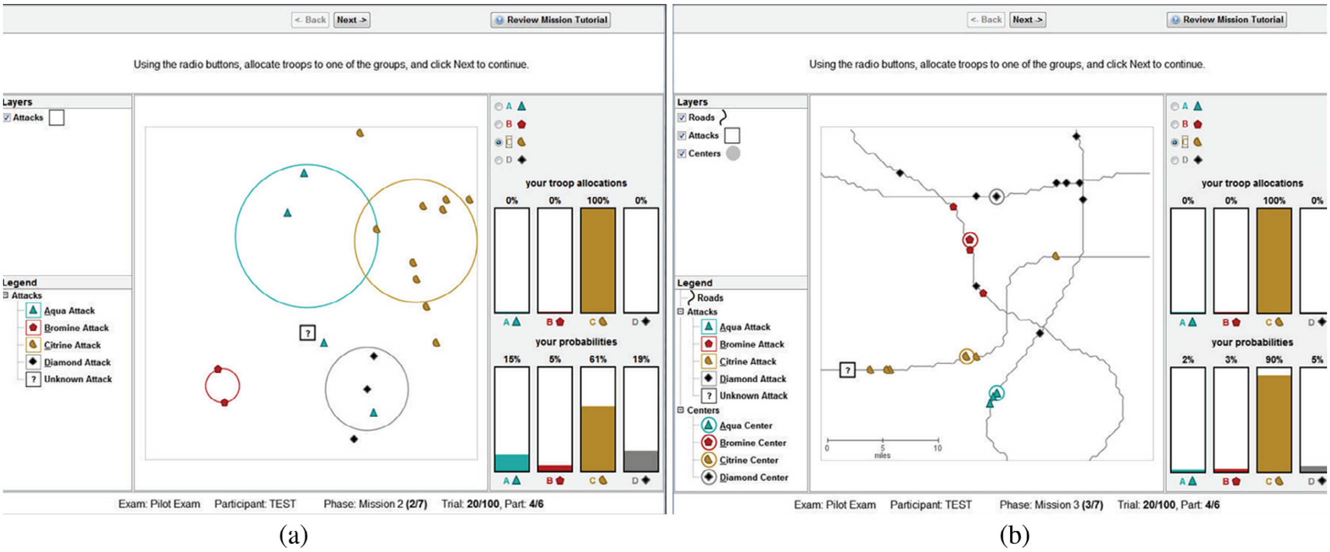

Fig. 1 Screenshots of (a) task 2 and (b) 3 performed by human subjects. In both cases, subjects had to determine the likelihood of a group attacking at a probe location based on distance, radius, and base rate.

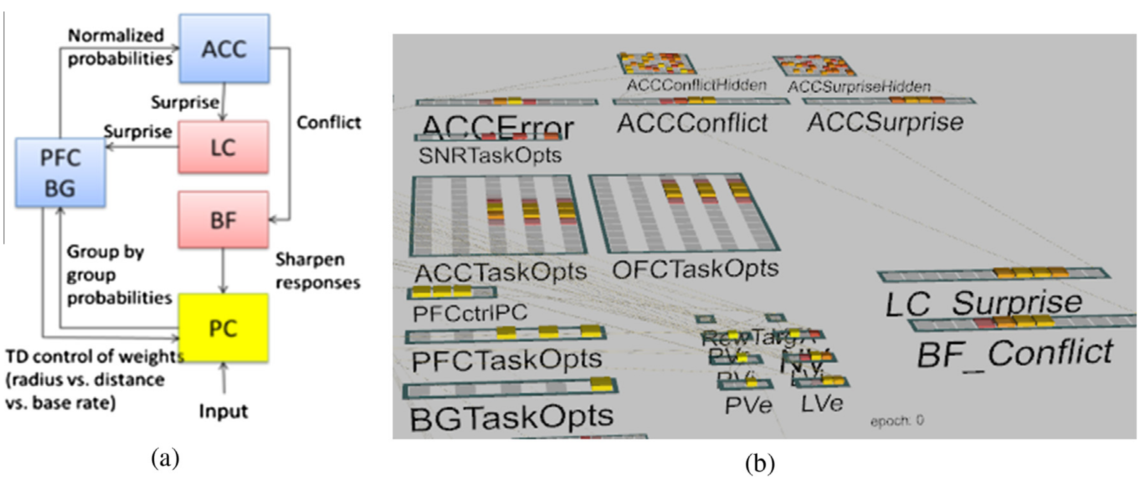

Fig. 2 (a) Block diagram of our model which includes frontal cortical areas (blue), neuromodulatory areas (pink) and posterior cortical areas (yellow). (b) Emergent implementation of our model focusing on frontal cortical areas and neuromodulatory areas. The Primary Value, Learned Value (PVLV) subsystem of the Prefrontal cortex, Basal Ganglia Decision Making (PBDM) model with layers such as PVe and LVe is also shown (Herd, Krueger, Kriete, Huang, & O'Reilly, 2013).

prefrontal cortex, an area known to be involved in working memory since the 1930s (e.g., Jacobsen, 1935). Control of which features were used in assessing a group's probability of attack was dictated by other regions of prefrontal cortex described in the next section.

Prefrontal cortex: decision making for proactive and reactive control

Within prefrontal cortex (PFC), we consider orbitofrontal cortex (OFC), anterior cingulate (ACC), and dorsolateral prefrontal cortex (DLPFC) to represent respectively the reward, effort and utility (reward − effort) of strategies in a Prefrontal cortex, Basal ganglia Decision Making or PBDM network (Herd, Krueger, Kriete, Huang, & O'Reilly, 2013) and a program code abstraction of its dynamics. In the programmatic encapsulation of PBDM, the effort of strategies in ACCTaskOpts of Fig. 2b is subtracted from the reward of strategies in OFCTaskOpts of Fig. 2b. This competition is believed to be carried out by indirect (or No Go) and direct (or Go) pathways in the striatum with the winner—i.e., the option with the highest utility—being selected in the basal ganglia, which is BGTaskOpts in Fig. 2b. PBDM and its program code approximation try to balance between strategies with higher reward but often higher effort.

Proactive and reactive control from the other regions of ACC triggered utility computations in PBDM-like dynamics as listed in Table 1. In the case of proactive control, PFC had two options. The first option is ''Look'' and output the first pass of group attack probability estimates. The second option is to ''Relook'' by increasing attention in PC computation, reconsider each group in turn, and then

Table 1 PFC options for proactive and reactive control.

	Option 1	Option 2	Option 3	Utility computations
Proactive control	''Look'': output group attack probability estimates after 1st pass	''Relook'': increase attention in PC computation, redo calculations, then output updated group attack probability estimates	—	Reward is constant across the options. Effort is proportional to conflict for option 1 and fixed for the 2nd strategy As conflict increases, the utility of the ''Look'' option decreases and the ''Relook'' option is more likely to be chosen
Reactive control	''Maintain'': Maintain current level of PC engagement	''Increase'': Increase level of PC engagement	''Randomize'': Randomize factors used in PC computation	Reward is constant across the options. Effort is proportional to surprise for option 1 and fixed for the 2nd and 3rd options As surprise increases, the utility of the ''Maintain'' option decreases and the ''Increase'' option is more likely to be chosen. If all PC features are being used in option 2, then ''Randomize'' is selected

output updated group attack probability estimates. Using a relook is more effortful but more accurate. This ''look—relook'' effect aligns with speed-accuracy trade-offs in psychological experiments (e.g., Laming, 1979; Rabbitt, 1966). In the case of reactive control, PFC controls the use or disuse of features in PC. There are three options. The first option, ''Maintain,'' retains the current level of PC engagement. The second option, ''Increase,'' raises the level of PC engagement by adding features used in assessing a group's probability of attack. The third option, ''Randomize,'' is only used when all PC features are being used and yet more control is needed; it randomly selects which factors are used in PC computation. Using more of features entails greater effort but greater accuracy. However, if no more control can be applied, a random strategy might be tried out of frustration. In Fig. 2b, PFCCtrlPC represents which features to use in PC probability computation.

Anterior cingulate cortex: conflict and surprise

Several subregions of ACC with different functions have been identified including those for conflict, error and surprise or unexpectedness (Nee, Kastner, & Brown, 2011). An influential subset of performance monitoring models include conflict models (Botvinick, Braver, Carter, Barch, & Cohen, 2001). Typically, conflict is a measure of the incompatibility of opposing responses or actions. Conflict is often measured by energy or the sum across response options. However, for probability distributions this sum would be constant (viz., 1). Thus, we chose normalized entropy to measure conflict instead (Chelian, Oros, Zaldivar, Krichmar, & Bhattacharyya, 2012). As examples, the normalized entropy of probabilities (.25.25.25.25) is 1; for probabilities (1000) it is 0. A 3-layer network is trained to map probabilities within the working memory portion of prefrontal cortex to normalized entropy values in ACCConflict (Fig. 2b).

Surprise was defined by the following formula:

$$surprise = \begin{cases} \frac{error}{2} & if\ error < 0.2 \\ \frac{error}{conflict} & otherwise \end{cases}$$

with limits to prevent division by zero or exceeding the value of one. When error is small, surprise is also small. However when errors are larger, surprise is the ratio of error and conflict. Conflict acts as a proxy for uncertainty or lack of confidence in a response. For example, with probabilities (.25.25.25.25), conflict is high because no response is differentiated from others and confidence is low. On the other hand, with probabilities of (1000), conflict is low so confidence is high—the neurocognitive model is sure that the first option is the true attacker. (Other measures of confidence might sum the spread or confidence interval within the probability estimate of each group.) Error is the mean of absolute differences (MAD) between ground truth and predictions of the true attacker shown as ACCError in Fig. 2b, while conflict is defined above. For the same error value (e.g., 0.4), surprise scales with confidence. When error is smaller than conflict, surprise is low (e.g., 0.4/0.8 = 0.5)—although an error was made, there was little confidence in the answer. Conversely, when error exceeds or is the same as conflict (e.g., 0.4/0.4 = 1), surprise is high—an error was made and the neurocognitive model was confident in its original response. A 3-layer network is trained to perform the calculation in ACCSurprise (Fig. 2b).

Neuromodulatory system: basal forebrain and locus coeruleus

ACC Conflict values are relayed to the basal forebrain (BF). If the activity of the BF crosses a critical threshold, it will trigger the programmatic encapsulation of PBDM to analyze the utility of the two possible options for proactive control as in Table 1. When conflict is low, the neurocognitive model provides a direct response; however when conflict is high, the neurocognitive model reconsiders, or relooks at input data before providing a response. So long as conflict

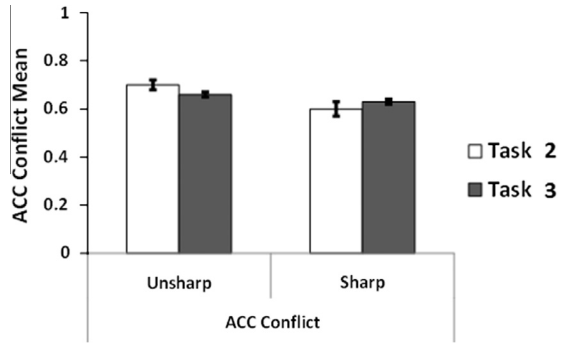

Fig. 3 ACC Conflict decreases when PC calculation is changed from an Unsharp to a Sharp state in both tasks 2 and 3. Unsharp is before the relook while Sharp is after the relook.

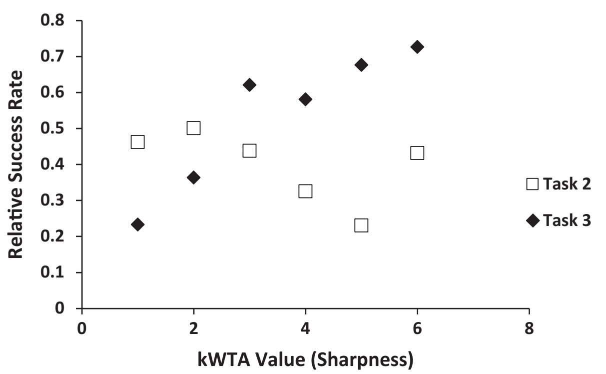

Fig. 4 The sharpness of PC layers (sharper for smaller kWTA value) results in either an increased relative success rate—a better fit to human responses—for task 3 or a response that fluctuates depending on the sharpness for task 2.

is low, utility will be high for the ''Look'' strategy and it will continue to be used. However, when conflict is high, utility is low for the ''Look'' strategy and the ''Relook'' strategy has a chance of winning. BF activity increases attention in PC by increasing competition between neurons effectively increasing the network's signal-to-noise ratio. In Emergent, increasing competition corresponds to reducing the number of winners in the kWTA mechanism.

ACC Surprise values are relayed to the locus coeruleus (LC). If the activity of the LC crosses a critical threshold, it will trigger PFC to analyze the utility of three possible options for reactive control as in Table 1. As surprise increases, the utility of the ''Maintain'' option decreases and the ''Increase'' option is more likely to be chosen. If all PC features are being used in option 2, then ''Randomize'' is selected. This is meant to model frustration—no more control can be applied and perhaps a random strategy might be worth trying.

Results

Conflict

In Fig. 3, when ACC computed high conflict (>0.6), PC recalculated the probability of each group with increased attention. When conflict is calculated again after sharpening of attention in PC, there is less conflict (<0.6) and a decision can be made. We tested multiple different levels of sharpening of PC and found it does decrease conflict in both tasks presented, but it only increases accuracy of PC calculation in task 3 as shown in Fig. 4. In task 2, it has little or even detrimental changes to the accuracy of the calculation. In Fig. 4, accuracy is measured in relative success rate, a measure of how well models fit human responses described in Lebiere et al. (2013). (The detrimental changes occur in this task because the differences in group probability estimates between the groups was below the range of error for PC calculation even in the sharpened condition, therefore even though conflict decreases, the calculation actually worsens.)

Surprise

In addition to conflict, when the result of the previous trial results in surprise, ACC mean error reaches a threshold and additional resources are used within PC to better calculate future trials. Fig. 5 shows that ACC Mean Error decreases in both task 2 and 3 when either base rate or radius is added to the distance only calculation to provide more information on each group. Similarly as with conflict, the addition of additional information may end up helping or hurting the final solution as shown in Fig. 6. In task 2, when the neurocognitive model is already using Distance and Radius, a surprising result causes PBDM to instruct PC to use additional information, i.e. Base rate, which highly increases the accuracy of the calculation with respect to human responses.

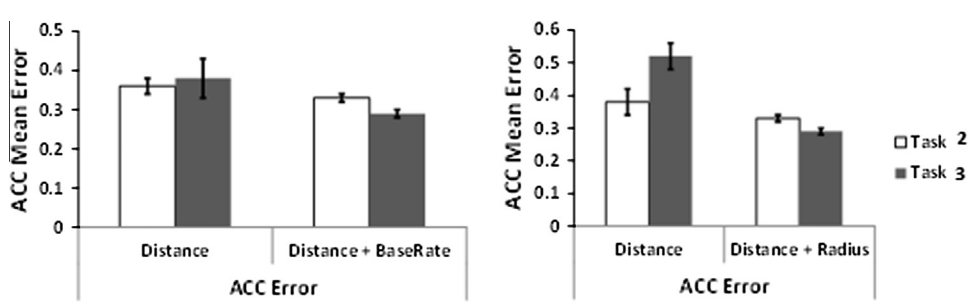

Fig. 5 After high error and surprise levels, the model selects more features to be used in future trails. On average this causes a decrease in mean error to occur in both tasks.

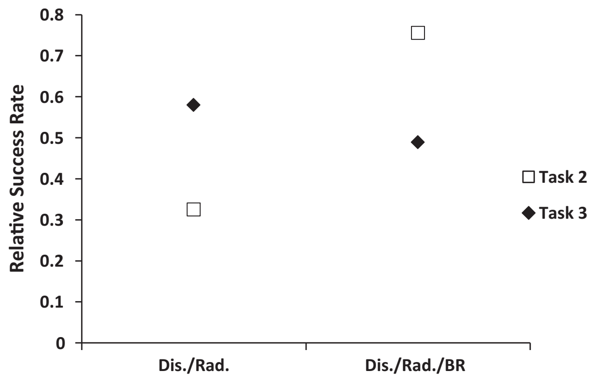

Fig. 6 Depending on the task, more information may increase the relative success rate as in task 2 or be detrimental as in task 3. However, in task 3, subsequent trials would subsequently reduce the amount of control.

However, in task 3 if the neurocognitive model is surprised and PBDM requests additional information the results worsen. This result would cause additional surprise resulting in PFC resetting its expectations and no longer requesting the use of Base rate in future calculations.

Discussion and conclusion

Typically models of cognitive control are demonstrated on simple tasks such as the Stroop task (1935). However in this work, we found that proactive and reactive control produces better fits to human data in relatively complicated inference and decision-making tasks. Furthermore, in our model, cognitive control was coordinated across several brain areas such as posterior cortex, frontal cortical areas and neuromodulatory areas. When group attack estimates were relatively undifferentiated, conflict was high. High conflict was relayed from ACC to the BF to increase attention in PC. Increased competition in PC, in turn, led to more distinction between response options and greater overall success in matching human performance of predicting attackers. When group attack estimates were incorrect, surprise was high. This was relayed from ACC to LC to trigger a re-evaluation of the spatial strategies used by PC. More terms in the spatial strategy in turn drove down error.

With respect to ACC, our model used conflict and surprise to initiative proactive and reactive control respectively. Broader theories of ACC such as PRO (Alexander & Brown, 2011) and EVC (Shenhav et al., 2013) attempt to unify several cognitive control functions of ACC, including conflict and surprise, with various objective functions. These models are also tied to neuroimaging and neurophysiological studies but do not discuss the effects of the neuromodulatory system.

With respect to neuromodulators, we primarily focused on acetylcholine and noradrenaline. Other models of decision making often include dopamine and serotonin (e.g., Chelian et al., 2012; Herd, Krueger, Kriete, Huang, & O'Reilly, 2013). Doya (2002) has also presented alternate interpretations of the neuromodulatory functions. In his work, dopamine signals reward prediction errors, serotonin controls the time scale of reward prediction (discounting), acetylcholine controls the speed of memory update, and noradrenaline controls the randomness in action selection. However, his work primarily considers reactive control in a reinforcement learning framework whereas in this work we also consider proactive control. Acetylcholinergic and noradrenergic effects for modulating attention in visual processing were also modeled by Avery et al. (2012, 2014). The latter work also uses detailed anatomical models of the primary visual area or V1. To our knowledge a similarly detailed reconstruction of parietal cortex used in this work has not been performed.

Acknowledgements

Supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of the Interior (DOI) contract number D10PC20021. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained hereon are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI, or the U.S. Government.

References

Aisa, B., Mingus, B., & O'Reilly, R. (2008). The emergent neural modeling system. *Neural Networks, 21*, 1146–1152.

Alexander, W. H., & Brown, J. W. (2010). Computational models of performance monitoring and cognitive control. *Topics in Cognitive Science, 2*, 658–677.

Alexander, W. H., & Brown, J. W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nature Neuroscience, 14*, 1338–1344.

Avery, M. C., Dutt, N., & Krichmar, J. L. (2014). Mechanisms underlying the basal forebrain enhancement of top-down and bottom-up attention. *The European Journal of Neuroscience, 39*, 852–865.

Avery, M. C., Nitz, D. A., Chiba, A. A., & Krichmar, J. L. (2012). Simulation of cholinergic and noradrenergic modulation of behavior in uncertain environments. *Frontiers in Computational Neuroscience, 6*, 1–16.

Botvinick, M. M., Braver, T. S., Carter, C. S., Barch, D. M., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review, 108*, 624–652.

Bouret, S., & Sara, S. J. (2005). Network reset: A simplified overarching theory of locus coeruleus noradrenaline function. *Trends in Neuroscience, 28*, 574–582.

Broussard, J. I., Karelina, K., Sarter, M., & Givens, B. (2009). Cholinergic optimization of cue-evoked parietal activity during challenged attentional performance. *European Journal of Neuroscience, 29*, 1711–1722.

Chelian, S.E., Oros, N., Zaldivar, A., Krichmar, J., & Bhattacharyya, R. (2012). Model of the interactions between neuromodulators and prefrontal cortex during a resource allocation task. In *Proceedings of the IEEE international conference on development and learning and epigenetic robotics*, San Diego, USA.

Dalley, J. W., McGaughy, J., O'Connell, M. T., Cardinal, R. N., Levita, L., & Robbins, T. W. (2001). Distinct changes in cortical acetylcholine and noradrenaline efflux during contingent and noncontingent performance of a visual attentional task. *Journal of Neuroscience, 21*, 4908–4914.

Disney, A. A., Aoki, C., & Hawken, M. J. (2007). Gain modulation by nicotine in macaque V1. *Neuron, 56*, 701–713.

Doya, K. (2002). Metalearning and neuromodulation. *Neural Networks, 15*, 495—506.
Goard, M., & Dan, Y. (2009). Basal forebrain activation enhances cortical coding of natural scenes. *Nature Neuroscience, 12*, 1444—1449.
Herd, S. A., Krueger, K. A., Kriete, T. E., Huang, T., & O'Reilly, R. C. (2013). Strategic cognitive sequencing: A computational cognitive neuroscience approach. *Computational Intelligence and Neuroscience*, 149329.
Herrero, J. L., Roberts, M. J., Delicato, L. S., Gieselmann, M. A., Dayan, P., & Thiele, A. (2008). Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. *Nature, 454*, 1110—1114.
Jacobsen, C. F. (1935). Functions of frontal association area in primates. *Archives of Neurology and Psychiatry, 33*, 558—569.
Krichmar, J. L. (2008). The neuromodulatory system — A framework for survival and adaptive behavior in a challenging world. *Adaptive Behavior, 16*, 385—399.
Laming, D. (1979). Choice reaction performance following an error. *Acta Psychologica, 43*, 199—224.
Lebiere, C. (1999). The dynamics of cognition: An ACT-R model of cognitive arithmetic. *Kognitionswissenschaft, 8*, 5—19.
Lebiere, C., Pirolli, P., Thomson, R., Paik, J., Rutledge-Taylor, M., Staszewski, J., et al (2013). A functional model of sensemaking in a neurocognitive architecture. *Computational Intelligence and Neuroscience*, 921695.
Nee, D. E., Kastner, S., & Brown, J. W. (2011). Functional heterogeneity of conflict, error, task-switching, and unexpectedness effects within medial prefrontal cortex. *NeuroImage, 54*, 528—540.
Nelson, C. L., Sarter, M., & Bruno, J. P. (2005). Prefrontal cortical modulation of acetylcholine release in posterior parietal cortex. *Neuroscience, 132*, 347—359.
Rabbitt, P. M. (1966). Errors and error correction in choice-response tasks. *Journal of Experimental Psychology, 71*, 264—272.
Sara, S. J. (2009). The locus coeruleus and noradrenergic modulation of cognition. *Nature Reviews Neuroscience, 10*, 211—223.
Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron, 79*, 217—240.
Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology, 18*, 643—662.
Sun, Y., & Wang, H. (2013). The parietal cortex in sensemaking: The dissociation of multiple types of spatial information. *Computational Intelligence and Neuroscience*, 152073.