Proceedings of the Annual Meeting of the Cognitive Science Society

The Role of Basal Ganglia Reinforcement Learning in Lexical Priming andAutomatic Semantic Ambiguity Resolution

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 41(0)

Authors
Ceballos, Jose M.
Stocco, Andrea
Prat, Chantel S.

Publication Date
2019

The Role of Basal Ganglia Reinforcement Learning in Lexical Priming and Automatic Semantic Ambiguity Resolution

Jose M. Ceballos, Andrea Stocco, Chantel S. Prat
{josemceb, stocco, csprat}@uw.edu

Department of Psychology and
Institute for Learning & Brain Sciences,
University of Washington
119A Guthrie Hall, Seattle, WA 98195 USA

Abstract

The current study aimed to elucidate the contributions of the subcortical basal ganglia to human language by adopting the view that these structures engage in a basic neurocomputation that may account for its involvement across a wide range of linguistic phenomena. Specifically, we tested the hypothesis that basal ganglia reinforcement learning mechanisms may account for variability in semantic selection processes necessary for ambiguity resolution. To test this, we used a biased homograph lexical ambiguity priming task that allowed us to measure automatic processes for resolving ambiguity towards high frequency word meanings. Individual differences in task performance were then related to indices of basal ganglia functioning and reinforcement learning, which were used to group subjects by learning style: primarily from choosing positive feedback (Choosers), primarily from avoiding negative feedback (Avoiders), and balanced participants who learned equally well from both (Balanced). The pattern of results suggests that balanced individuals, whom learn from both positive and negative reward equally well, had significantly lower access to the subordinate homograph word meaning. Choosers and Avoiders, on the other hand, had higher access to the subordinate word meaning even after a long delay between prime and target. Experimental findings were then tested using an ACT-R computational model of reinforcement learning that learns from both positive and negative feedback. Results from the computational model confirm and extend the pattern of behavioral findings, and provide a reinforcement learning account of lexical priming processes in human linguistic abilities, where a dual-path reinforcement learning system is necessary for precisely mapping out word co-occurrence probabilities.

Keywords: language; semantics; lexical selection; ambiguity resolution; priming; reinforcement learning; basal ganglia; dopamine; cognitive modeling; ACT-R

Introduction

The field of the neurobiology of language has traditionally focused on the contributions of cortical structures to linguistic processes (Tremblay & Dick, 2016). However, research from different sub-fields suggests that the subcortical basal ganglia are an essential part of the neurobiological bases of human linguistic abilities (Crosson, 1985; Booth, Wood, Lu, Houk, & Bitan, 2007; Seo, Stocco, & Prat, 2018). To date, no existing account of the neurobiology of language is able to systematically explain what the role of these subcortical structures is across the many levels of linguistic processing. Thus, in its current stage, the field suffers from a limited understanding of the neural processes that give rise to language. A detailed whole-brain understanding of this human ability is key to inform robust models of language neurobiology and also to advance our understanding of language disabilities for translational purposes. In an effort to contribute to a whole-brain model of language functioning, this work focuses on understanding the role of the basal ganglia in language.

Given that the basal ganglia are some of the most neurobiologically ancient structures (Lieberman, 2001), it is reasonable to assume that their role in human linguistic abilities is analogous to the more general motor or cognitive functions observed in other species. Indeed, many prominent theories and models of basal ganglia functioning stem from observations of motor control (Mink, 1996) and extend these functions to non-motor and abstract cognitive processes spanning from cognitive control (Graybiel, 1995; Stocco, Lebiere, & Anderson, 2010) to working memory capacity (Hazy, Frank, & O'Reilly, 2007). Thus, the current research aims to understand basal ganglia contributions to language in the context of the already well-understood and well-established theories of selection and reinforcement learning (RL). To test the hypothesis space of basal ganglia selection processes in language, we turned to semantic processing as a model system for competition between multiple viable alternatives. Specifically, this work is grounded on models of models of semantic activation spreading (Collins & Loftus, 1975).

Semantic ambiguity (also referred to as lexical ambiguity) occurs when a word refers to multiple different concepts (Vitello & Rodd, 2015). For example, the word "hot" can refer either to temperature or to food spiciness. Cases of semantic ambiguity may arise in conversational settings, and are also more commonly encountered in written form such as news headlines, puns, poetry, and novels. The ability to properly disambiguate an input into the contextually appropriate represented meaning is key for listening and reading comprehension. More importantly, this process provides details on a fundamental neurocognitive mechanisms, such as the contextual integration of information, statistical learning, inhibition, and selection processes used to manage simultaneous and competing neural representations that are at odds with the task goal of accurate transfer of information in communicative settings.

Semantic ambiguity can arise in a variety of different ways. The first class of ambiguity arises from words that have different unrelated meanings. For example, "bark" can refer to the sound a dog makes, or the outermost layer of a tree. In this case, both meanings of "bark" constitute a true homonym, but are also homographs and homophones (same spelling, and same sound, respectively). Furthermore, words can be encountered in contexts where only the written form is ambiguous (e.g., the homographs for "lead"), or only the spoken form is ambiguous (e.g., the homophones for "be/bee" or "seam/seem").

The cognitive mechanisms supporting the resolution of semantic ambiguities are best understood by exploring theories on the dynamics of semantic information and its representation. When a listener or reader first encounters a word with multiple meanings, all meanings are quickly activated and available in parallel. This refers to the automatic component in semantic processing. Furthermore, if encountered in isolation or in a highly ambiguous context, an ambiguous word will be automatically disambiguated towards the highest frequency meaning, reflecting another series of automatic selection processes. However, if an ambiguous word is encountered following a strong biasing context towards one specific meaning, only the contextually-relevant word meaning is available. This suggests that when ambiguous words are encountered, all meanings are initially activated, but this activation is modulated by multiple factors such as sentence context and meaning frequency.

While most research focused on understanding the neural mechanisms supporting lexico-semantic processing and ambiguity resolution has focused on cortical structures such as the left inferior frontal gyrus (for a review, see Vitello & Rodd, 2015), there is evidence suggesting a key involvement of subcortical structures in this process (e.g., Ketteler, Kastrau, Vohn, & Huber, 2008; Mason & Just, 2007). For example, a lexical priming investigation found that monolingual individuals experience abnormalities in the neurocognitive dynamics that shape lexical priming (Copland, Chenery, & Murdoch, 2001). Specifically, healthy participants show no traces of subordinate word activation following a long delay between prime and target, and thus reflect automatic semantic ambiguity resolution towards the dominant or highest frequency meaning. Parkinson's Disease (PD) patients, whom have decreased dopaminergic functioning resulting in a general hyperactivity of the basal ganglia indirect pathway, on the other hand, exhibit a longer-term activation of the multiple competing representations.

Although findings such as these have been traditionally framed under a selection and inhibition framework, we explore the hypothesis that the signature role of basal ganglia in RL may more accurately explain its role in semantic processing. In other words, the basal ganglia may be involved in statistical learning and predictive processing during language comprehension. Critical for the current investigation, the activity of the basal ganglia is often modeled as reflecting Temporal Difference (TD) learning. As it happens, TD-learning does not accurately reflect the computations of the basal ganglia, which are the result of the opposite contributions of two conflicting pathways. Their contribution have been modeled as the sum of competing RL systems (Frank, Seeberger, & O'reilly, 2004; Stocco, 2018). Individuals vary in the learning rates of the two pathways as a function of biological parameters (such as density of dopamine receptors: Frank, Moustafa, Haughey, Curran, & Hutchison, 2007) and external factors (administration of dopamine: Frank et al., 2004), and individual differences in the preponderance of each pathway can be indirectly measured through the PSS task (Frank et al., 2004; Stocco et al., 2017). Thus, the current investigation tests the hypothesis that individual differences in PSS task behavioral indices of basal ganglia pathways will be related to performance in a lexical prime style task. Furthermore, we make the prediction that a balance in functioning across both pathways is critical for optimal semantic processing and ambiguity resolution.

Methods

Participants

Informed consent was obtained from participants prior to the experiment, as outlined by the University of Washington Institutional Review Board. Participants were recruited using the Psychology Departments Participant Research Pool and all participants were compensated with course credit for their participation. Data were collected from 140 healthy monolingual participants (66 females, mean age = 19.4 years). Seven subjects were excluded from analyses due to low accuracy (≤ 0.50) in the primary experimental task, the Word-Pair Task (WPT).

Tasks

All participants completed the following tasks in four pseudorandomized orders to control for possible fatigue effects induced by the WPT and PSS task length.

Word-Pair Task Measures of lexical priming were collected using the WPT. This task was designed to measure the availability of dominant and subordinate word meanings following the presentation of primes with multiple meanings. The primes used shared both phonetic and orthographic forms across both word meanings, making them true homographs (e.g., "Bat"). The prime and target words were presented in the center of the screen, one at a time, separated by an interstimulus interval (ISI) of either 150 ms (short) or 850 ms (long). Prior to starting the task, participants were asked to place their right index finger on the "P" key of the keyboard, and their left index finger on the "Q" key of the keyboard. Participants were then asked to respond with a button press if the target word was related or unrelated to the prime. Key mappings for related and unrelated were counterbalanced.

There were two conditions of interest (1 and 2) and two control conditions (3 and 4): (1) homograph prime / dominant target, (2) homograph prime / subordinate target, (3) prime / related target, and (4) prime / unrelated target. These four conditions will be referred to as dominant, subordinate, related, and unrelated (respectively) form here on for simplicity purposes. Participants completed 100 total prime-target pair trials, where 20 belonged to condition 1, 20 to condition 2, 30 to condition 3, and 30 to condition 4. Word frequency meanings were obtained from (Twilley, Dixon, Taylor, & Clark, 1994), and subordinate words were defined as having a relatedness frequency of less than 0.3 in a 0-1 scale, while dominant words had a relatedness frequency of greater than 0.7. Since each homograph prime is associated with two meanings, but each prime was presented once for each participant, two WPT versions were created. In one version, the dominant meaning of a homograph was used (e.g., version A contained "Bank" / "Money") while the other version used the subordinate meaning (e.g., version B contained "Bank" / "River"). The two lists were counter-balanced for word frequency, word length, and syllable length.

Probabilistic Stimulus Selection Task The PSS task is an iterative, two-alternative, forced-choice decision-making paradigm first introduced by Frank et al. (2007). In this task, participants are repeatedly asked to select one of two stimuli presented on the screen. Participants are also told that some of their choices would result in success, and some of them would result in failure, depending on which stimulus they choose. Feedback on the outcome of their decision is presented immediately after participants select a stimulus. To encourage participants to avoid explicit strategies (such as rote memorization of each stimulus history of successes), stimuli are implemented as complex shapes that are difficult to verbalize, typically Hiragana characters presented to non-Japanese speaking participants. Unbeknownst to participants, each stimulus has a predefined "success" probability. Six stimuli in total are used in the experiment, with success probabilities varying linearly from 80% to 20%. In the first phase, the stimuli are divided into fixed pairs, with the highest probability stimulus always paired with lowest probability one, then second higher stimulus paired with the second lowest one, and the third highest probability stimulus paired with the third lowest one.

Two values are calculated from the test phase of the PSS task: *Choose* accuracy, which represents the accuracy in choosing the most rewarding stimulus over others; and *Avoid* accuracy, that is, the proportion of times in which participants avoid the least rewarding stimulus. If we indicate the six stimuli with the letters $A, B \ldots F$, with A being the most rewarding stimulus and B the least rewarding one, then Choose and Avoid accuracies are calculated as the probability of choosing A when paired against C, D, E, and F, and the probability of choosing C, D, E, or F when they are paired with B, respectively.

Previous patients and genetic studies have demonstrated a functional connection between these two measures and the basal ganglia pathways. For example, Parkinson's patients, whose indirect pathway dominates over the direct one due to a loss of dopaminergic inputs from the substantia nigra pars compacta (SNc), have higher Avoid accuracy than Choose accuracy. Furthermore, this pattern is reversed when drugs are administered that overcompensate the direct pathway activity. Additionally, individuals with genetic alleles that cause a greater production of dopamine receptors in the direct pathway tend to be Choosers rather than Avoiders; conversely, individuals whose alleles cause greater number of dopamine receptors in the indirect pathway tend to be Avoiders (Frank et al., 2007; Frank & Hutchison, 2009).

Analyses

Behavioral Data Cleaning Target words in the WPT were cleaned on a by-participant basis for RT outliers, defined as trial RTs greater than or lower than three standard deviations from the mean.

Participant Groups Participant groups were created using PSS Choose and Avoid scores. Since one of the guiding assumptions of this investigation was that one's ability to learn from *both* positive and negative feedback, groups were created using a relative score where Avoid was subtracted from Choose, which resulted in scores between 100 and -100. Third-group splits were then used to separate individuals into participant groups. Thus, high values (approximately 33 to 100) reflected participants who learned primarily from positive feedback (Choosers), low values (approximately -100 to -33) reflected participants who learned primarily from negative feedback (Avoiders), and values around zero (-33 to 33) reflected individuals who learned equally as well from positive and negative feedback (Balanced). This resulted in 44 Choosers, 38 Avoiders, and 52 Balanced participants.

Analysis with Linear Mixed Effects Models The data were analyzed using linear mixed effects (LME) models, as this method has been previously shown to outperform the traditional procedures such as ANOVA (Kristensen & Hansen, 2004), and can adequately handle imbalances in group sizes. However, for validation purposes, the same results were reproduced using ANOVA (although not reported herein). LME models were specified using the R lme4 package (Bates, Mächler, Bolker, & Walker, 2015). The model was specified using the following formula:

$$\text{Target Accuracy} \sim \text{ISI} \times \text{Condition} \times \text{PSS Group} \\ + (1 + \text{Condition} \mid \text{Participant})$$

where the dependent variable is Target Word accuracy, the fixed-effects term is the factors for ISI (short or long) × Condition (dominant or subordinate) × PSS Group (Choosers, Balanced, or Avoiders), and the random effects term allows for each participant to have a different slope (or effect) for Condition, while intercepts and slopes for each participant by Condition are allowed to be correlated (e.g., higher intercepts may also have steeper slopes). Finally, a type III ANOVA with Satterthwaite's method was used to test for significance between the factors of interest in the LME model.

Computational Model

A theoretical model was implemented to examine predictions on the relationship between reward learning and lexical retrieval[1]. The model was developed in the ACT-R cognitive architecture (Anderson, Fincham, Qin, & Stocco, 2008; Anderson et al., 2004), a general theory of cognition that enables the development of complete models capable of end-to-end simulations of a complete task while maintaining a high degree of psychological plausibility. The model described herein is based on a previously published model of the role of the basal ganglia in the PSS task (Stocco, 2018). According to this model, the conflict between the two pathways can be simulated in ACT-R as a conflict between the selection of opposite and symmetric *productions*, that is, state-action pairs that implement minimal cognitive steps. Productions representing the direct pathway implements "Go" actions, while those representing the indirect pathway represent opposite "No Go" actions. For example, the choice between two options in the PSS task, *A* and *B*, can be represented as the competition between two alternative pairs of productions, "Choose A" and "Avoid A" and "Choose B" and "Avoid A". In ACT-R, the competition between productions is resolved through a softmax algorithm that preferentially selects the actions with the highest estimated *utility*, a scalar quantity that depends on the history of previous successful uses of the production and is learned through a reinforcement-learning algorithm. Importantly, Stocco, 2018 noticed that both individual differences due to differential expressions of dopamine genes (Frank et al., 2007) and the effects of basal ganglia pathologies (Frank et al., 2004) can be successfully captured by differentially altering the learning rates of "Choose" and "Avoid" productions. The different learning rates will be indicated as α_C and α_A, respectively.

An ambiguity resolution experiment can also be understood as as a two-alternative forced choice (2AFC) task in the context of lexical retrieval. In essence, two homographs are competing for access to semantic retrieval. Consequently, for each choice, two competing selections are performed. Thus, if the two homographs are a dominant and a subordinate interpretation of the same written word, each of them will have two production rules associated with them, "Choose Dominant" and "Avoid Dominant", and "Choose Subordinate" and "Avoid Subordinate".

Contrary to traditional 2AFCs, in lexical access the two options are not equivalent in terms of response times. Selection of the dominant meaning is usually associated with much shorter retrieval times than selection of the non-dominant meaning. In our model, this was captured by forcing those production rules that select the subordinate meaning ("Avoid Dominant" and "Choose Subordinate") to have a longer execution time. As a consequence, under short ISI, the subordinate meaning is never successfully selected. Under longer ISIs, however, participants *do* have a chance to select these meanings, so that the eventual firing of productions that select the subordinate interpretation could result in the successful retrieval of the least common meaning of the homograph.

Finally, to derive predictions from the model, we conducted an extensive set of simulations of the utility values associated to production rules under different reward conditions, corresponding to different situations in which the selection of the dominant or subordinate meaning are correct. Specifically, we examined a hypothetical situation in which the dominant meaning is contextually correct 80% of the time and the subordinate 20% of the time. To simulate the large amount of experience with the occurrence statistics of different lexical items that is associated with adult native speakers, the model was let to learn the corresponding utility values until they reached asymptotic values.

Importantly, these simulations of language experience were conducted under different learning rate parameters. The parameter values were chosen to reflect the values that were found to best capture genetic variance of dopamine receptors in healthy adults in Stocco (2018). Specifically, we simulated three groups of individuals, exhibiting a preference to learn from positive feedback ($\alpha_C = 1.5, \alpha_A = 1.0$), a preference to learn from negative feedback ($\alpha_C = 1.0, \alpha_A = 1.5$), or no preference between the two ($\alpha_C = 1.5, \alpha_A = 1.5$). These parameters are associated with different performance profiles in the PSS task, corresponding to a preference for "Choose A", for "Avoid B", or for a balance between the two (Stocco, 2018).

Results

General Word-Pair Task Results

Mean accuracy for dominant trials (M = 0.90, SD = 0.14) was significantly higher than for subordinate trials (M = 0.55, SD = 0.15, t(138) = 22.17, $p < 0.0001$). Differences in mean RTs were also observed, faster for dominant trials (M = 832.04, SD = 194.71) than subordinate trials (M = 996.25, SD = 238.26, t(138) = -13.77, $p < 0.0001$).

General Probabilistic Stimulus Selection Task Results

Subjects performed similarly across Choose (M = 69.78, SD = 22.24) and Avoid (M = 67.99, SD = 22.22) trials. Furthermore, as in previous studies using the PSS Task (Stocco et al., 2017; Frank et al., 2007; Frank & Hutchison, 2009), Choose and Avoid trials were not correlated (r(138) = -0.12, p = 0.14).

Linear Mixed Effects Model Results: Relating WPT Performance and PSS Groups

The LME model predicting Target accuracy had a total explanatory power (conditional R^2) of 90.62%, in which the fixed effects explained 68.43% of the variance (marginal R^2). The model revealed a significant main effect of Condition (F(1, 131) = 1096.33, $p < 0.0001$). A significant two way interaction between Condition × ISI was also observed (F(1, 262) = 6.47, p = 0.012), alongside a significant three-way interaction between Condition × ISI × PSS Group (F(2, 262) = 3.86, p = 0.022). Marginal two-way interactions were observed for Condition × PSS Group (F(2, 131) = 2.45, p =

---

[1] Code for the model is available on our laboratory's GitHub repository: *http://github.com/UWCCDL/BAGELS_ACTR*

0.087) and also ISI × PSS Group (F(2, 262) = 3.00, p = 0.051). For details, see Figure 1.

A follow-up analysis using the orthogonal contrasts extracted from the LME model suggest that the three-way interaction between Condition × ISI × PSS Group is explained by higher accuracy to Target Words during the Subordinate condition observed in PSS Choosers (difference = 0.083, t(166.85) = 2.41, p = 0.028) and Avoiders (difference = 0.086, t(166.95) = 2.60, p = 0.017), relative to the Balanced group, during the long ISI.

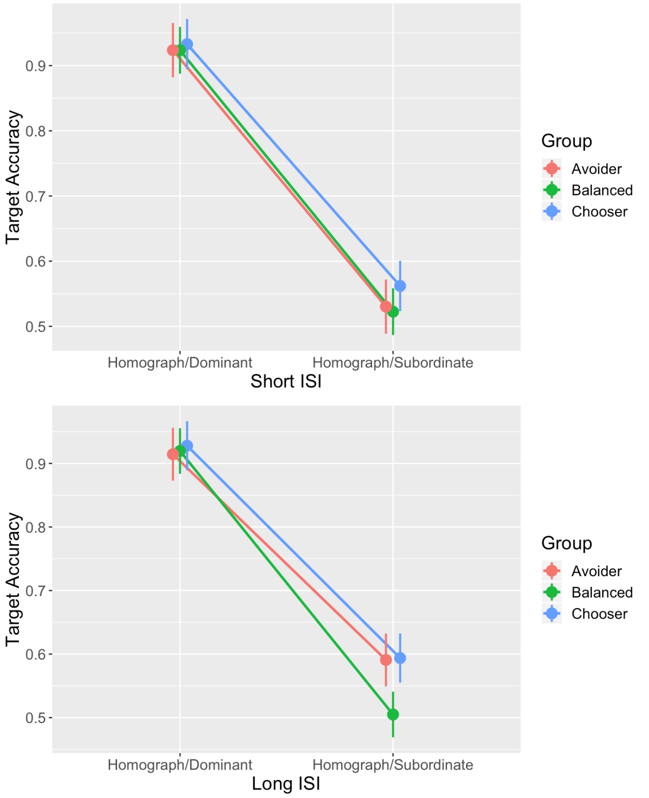

Figure 1: Top: Accuracy Dominant and Subordinate conditions for the short ISI. Bottom: Accuracy for Dominant and Subordinate conditions for the long ISI.

PSS Groups & Reading Experience Control Measure

The Author Recognition Test (Stanovich & West, 1989) score differences computed in order to ensure that differences in sensitivity to the subordinate word meaning observed was not driven by reading experience. There were no differences between Choosers and Balanced participants (t(94) = -0.44, p = 0.66), nor between Avoiders and Balanced (t(88) = -0.06, p = 0.95) that could account for the effect observed in the LME model results reported previously.

Computational Model Results

To generate predictions, the model was run for 1,000 times under the different values of α_C and α_A associated with Choosers, Avoiders, or Balanced individuals. The model predicts that, under short ISI, all three groups should perform at chance for the subordinate meaning, with no difference in performance. Under long ISI, however, the model predicts that Avoiders and Choosers should have greater than chance performance for the subordinate condition (62% and 63%, respectively), while Balanced individuals should still perform essentially at chance (55% accuracy). Note that these predictions are parameter-free, and come remarkably close to the actual results of our experiment. In the model, this asymmetry in behavior is due to the fact that different initial learning rates α_C and α_A result in biased estimates of success when selecting dominant and subordinate meanings, respectively. In particular, the model predicts that Choosers would tend to overestimate the probability of the subordinate meaning, while Avoiders would tend to underestimate the probability of the dominant meaning, with both cases resulting in a tendency to favor the selection of the subordinate meaning. Under balanced learning rates, instead, the model correctly estimates the rarity of the subordinate meaning and tends to select it significantly less often.

Discussion

The current project explored the hypothesis that human linguistic ability, and specifically semantic processing, is dependent on core basal ganglia RL mechanisms. The results provide evidence for the proposed hypothesis, and more specifically, suggest that individual differences in learning from positive or negative feedback are predictive of automatic semantic ambiguity resolution in context-free lexical ambiguity priming paradigm. Specifically, task performance was in line with behavioral predictions by the computational cognitive model, which predicted that action-selection in the basal ganglia for dominant and subordinate meanings would happen in line with an individual's estimate of success for choosing either meaning. To illustrate this, when a Balanced participant reads the word "bank" they co-activate the associated "money" and "river" meanings. Selection happens in line with their learned estimate that"river" rarely occurs following "bank," and this subordinate meaning is unavailable for the semantic relatedness judgment, resulting in poor task performance (for this condition, only). Thus, the signal generated by the basal ganglia during semantic selection can be seen as reflecting an individual's estimate of that word-meaning co-occurrence, or in other words, the individual's representation of relative frequency of a meaning associated with a lexical form.

Furthermore, findings from this investigation are compatible with the widely accepted view that prefrontal cortex (PFC) regions and specifically the left inferior frontal gyrus (LIFG), are involved in semantic selection processes (Vitello & Rodd, 2015). While the LIFG may very well be the primary driver of semantic selection, it is known to make use of biasing signals to rule out multiple competing representations (Schnur et al., 2009). This biasing signal is posited to stem from the basal ganglia, as research on the functional and anatomical properties of the PFC-basal ganglia network has shown that two of the five main cortico-striatal-thalamo-cortical loops project directly to lateral prefrontal regions, including dorsolateral PFC and lateral orbitofrontal cortex (Alexander, Crutcher, & DeLong, 1991). Thus, the basal ganglia posses the functional, anatomical, and computational properties necessary to provide biasing signals to LIFG during semantic ambiguity resolution.

Interestingly, these results reproduce and extend, by artificially segmenting a continuum of basal ganglia-mediated Choose and Avoid learning in a healthy population, findings observed in clinical groups. As mentioned previously herein, PD patients show abnormal lexical priming effects, with disrupted automatic semantic ambiguity resolution and sustained multiple competing representations. Additionally, literature focusing on the cognitive effects of Huntingtons Disease (HD) a basal ganglia dysfunction characterized by hyper-dopaminergic signaling and thus a hyper-active direct pathway, reveals that HD patients also have an increased susceptibility to semantic priming (Randolph, 1991). Taken together, these findings highlight the importance of a competitive dual-path RL system that gives rise to learning from both positive and negative feedback.

Possible alternative explanations exist for the current set of experimental results. Many theoretical and computational models of basal ganglia functioning focus on its role as "gates" that modulate prefrontal cortex functioning through selection (or Choose) and inhibition (or Avoid) mechanisms. Thus, under this framework, we would anticipate to find that Choosers would manage conflict in multiple competing representations by selecting the relevant or dominant word meaning, while Avoiders would inhibit the subordinate meaning. This is, however, not what is observed in the behavioral results, where both Choosers and Avoiders show identical performance in the subordinate condition after the long delay. This pattern of results is most compatible with a RL explanation of statistical learning, where a one-path mechanism (akin to traditional TD-learning) would over-estimate the utility of the lower frequency meaning. In other words, it is possible that Choosers are overly sensitive to low frequency reward probabilities, while Avoiders are less sensitive to high frequency reward probabilities. This results in a misrepresentation of the relative frequency effect observed between the dominant and subordinate word meanings.

This proposed role of the basal ganglia in RL through statistical mapping of the rich and dynamic linguistic environment, and engaging in live predictive processing may ultimately account for its involvement across multiple language processing modalities. In fact, a great deal of work exists that discusses evidence of basal ganglia involvement in language through the lens of a pacemaker-like, live, temporal processing machine that synchronizes internal states with external inputs (Kotz, Schwartze, & Schmidt-Kassow, 2009). While this research has focused mostly on morphosyntactic processing, its framework is both compatible with the one proposed herein and can be extended to multiple processing domains, including those beyond linguistic processing (e.g., non-linguistic cognitive functioning and motor processing). We consider these exciting areas for future research.

Acknowledgments

This research was supported by a National Science Foundation Graduate Research Fellowship (DGE-1256082) awarded to Jose M. Ceballos and by an award from the Office of Naval Research (ONRBAA13-003) to Chantel S. Prat and Andrea Stocco.

References

Alexander, G. E., Crutcher, M. D., & DeLong, M. R. (1991). Basal ganglia-thalamocortical circuits: parallel substrates for motor, oculomotor,"prefrontal" and "limbic" functions. In *Progress in brain research* (pp. 119–146). Elsevier.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). *An integrated theory of the mind* (Vol. 111) (No. 4).

Anderson, J. R., Fincham, J. M., Qin, Y., & Stocco, A. (2008). A central circuit of the mind. *Trends in Cognitive Sciences*, *12*(4), 136–143.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 1–48.

Booth, J. R., Wood, L., Lu, D., Houk, J. C., & Bitan, T. (2007). The role of the basal ganglia and cerebellum in language processing. *Brain research*, *1133*, 136–144.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological review*, *82*(6), 407.

Copland, D. A., Chenery, H. J., & Murdoch, B. E. (2001, August). Discourse Priming of Homophones in Individuals With Dominant Nonthalamic Subcortical Lesions, Cortical Lesions and Parkinsons Disease. *Journal of Clinical and Experimental Neuropsychology*, *23*(4), 538–556.

Crosson, B. (1985). Subcortical functions in language: a working model. *Brain and language*, *25*(2), 257–292.

Frank, M. J., & Hutchison, K. (2009). Genetic contributions to avoidance-based decisions: striatal d2 receptor polymorphisms. *Neuroscience*, *164*(1), 131–140.

Frank, M. J., Moustafa, A. A., Haughey, H. M., Curran, T., & Hutchison, K. E. (2007). Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. *Proceedings of the National Academy of Sciences*, *104*(41), 16311–16316.

Frank, M. J., Seeberger, L. C., & O'reilly, R. C. (2004). By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science*, *306*(5703), 1940–1943.

Graybiel, A. M. (1995). Building action repertoires: memory and learning functions of the basal ganglia. *Current opinion in neurobiology*, *5*(6), 733–741.

Hazy, T. E., Frank, M. J., & O'Reilly, R. C. (2007). Towards an executive without a homunculus: computational models of the prefrontal cortex/basal ganglia system. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *362*(1485), 1601–1613.

Ketteler, D., Kastrau, F., Vohn, R., & Huber, W. (2008, February). The subcortical role of language processing. High level linguistic features such as ambiguity-resolution and the human brain; an fMRI study. *NeuroImage*, *39*(4), 2002–2009.

Kotz, S. A., Schwartze, M., & Schmidt-Kassow, M. (2009, September). Non-motor basal ganglia functions: A review and proposal for a model of sensory predictability in auditory language perception. *Cortex*, *45*(8), 982–990.

Kristensen, M., & Hansen, T. (2004). Statistical analyses of repeated measures in physiological research: a tutorial. *Advances in physiology education*, *28*(1), 2–14.

Lieberman, P. (2001). Human language and our reptilian brain: The subcortical bases of speech, syntax, and thought. *Perspectives in Biology and Medicine*, *44*(1), 32–51.

Mason, R. A., & Just, M. A. (2007, May). Lexical ambiguity in sentence comprehension. *Brain Research*, *1146*, 115–127.

Randolph, C. (1991). Implicit, explicit, and semantic memory functions in alzheimer's disease and huntington's disease. *Journal of Clinical and Experimental Neuropsychology*, *13*(4), 479–494.

Schnur, T. T., Schwartz, M. F., Kimberg, D. Y., Hirshorn, E., Coslett, H. B., & Thompson-Schill, S. L. (2009). Localizing interference during naming: Convergent neuroimaging and neuropsychological evidence for the function of broca's area. *Proceedings of the National Academy of Sciences*, *106*(1), 322–327.

Seo, R., Stocco, A., & Prat, C. S. (2018). The bilingual language network: Differential involvement of anterior cingulate, basal ganglia and prefrontal cortex in preparation, monitoring, and execution. *NeuroImage*, *174*, 44–56.

Stanovich, K. E., & West, R. F. (1989). Exposure to print and orthographic processing. *Reading Research Quarterly*, 402–433.

Stocco, A. (2018). A biologically plausible action selection system for cognitive architectures: Implications of basal ganglia anatomy for learning and decision-making models. *Cognitive science*, *42*(2), 457–490.

Stocco, A., Lebiere, C., & Anderson, J. R. (2010). Conditional routing of information to the cortex: A model of the basal ganglias role in cognitive coordination. *Psychological review*, *117*(2), 541.

Stocco, A., Murray, N. L., Yamasaki, B. L., Renno, T. J., Nguyen, J., & Prat, C. S. (2017). Individual differences in the simon effect are underpinned by differences in the competitive dynamics in the basal ganglia: An experimental verification and a computational model. *Cognition*, *164*, 31–45.

Tremblay, P., & Dick, A. S. (2016). Broca and wernicke are dead, or moving past the classic model of language neurobiology. *Brain and language*, *162*, 60–71.

Twilley, L. C., Dixon, P., Taylor, D., & Clark, K. (1994). University of alberta norms of relative meaning frequency for 566 homographs. *Memory & Cognition*, *22*(1), 111–126.

Vitello, S., & Rodd, J. M. (2015). Resolving semantic ambiguities in sentences: Cognitive processes and brain mechanisms. *Language and Linguistics Compass*, *9*(10), 391–405.