HIGH SCHOOL EDITION
Journal of Student Research
Volume 12 Issue 3 (2023)

Distractor-Specific Single Neuron Activity Predicts Visual Working Memory Task Outcomes

Jia Lakhamraju[1] and Mare Stewart[#]

[1]Singapore American School
[#]Advisor

ABSTRACT

This paper explores the relationship between neural activity and behavioral performance in the form of visual working memory (VWM) task outcomes, by answering the question: Are there any significant differences in the firing rates of individual neurons during the distractor presentation period of a VWM task between success and error trials that can predict the outcome of a trial? Distractor-specific single neuron firing rates during a VWM task were analyzed to answer this question. A logistic regression was used to identify the predictive capability of neural firing rate on trial outcome with the neural activity of 51 cells from the lateral prefrontal cortex (LPFC) of a primate. This study found that a best-fit logistic model could predict the behavioral performance of the primate (success or error of the VWM task) with 63.01% accuracy, with additional machine learning techniques producing scores upwards of 68% accuracy. Moreover, greater firing rates in response to the distractor, indicating less efficient distractor suppression, accompanied the error trials of the VWM task. This suggests that stronger neural responses to task-specific distractors can hinder the attentional filtering required for efficient working memory, supporting previous research that found that distractor suppression is a mechanism that heavily influences WM efficiency. These findings indicate that people, particularly children, with disorders that affect WM capacity such as ADHD may experience stronger neural responses to distractors, and therefore inefficient distractor suppression, at the single neuron level when engaging in goal-oriented behaviors, which can significantly impact learning and other developmental processes.

Introduction

Visual working memory (VWM) has been well established in the field of neuroscience as a cognitive function essential to our daily lives, allowing us to process and temporarily maintain visual information in our working memory in order to focus on and complete tasks (Olivers et al., 2020). Decades of research have allowed us to try and identify the various functions of the brain, as well as the structures that contribute to certain processes, such as VWM. Although several brain structures contribute to the functioning of VWM, including the posterior sensory areas and basal ganglia, the prefrontal cortex (PFC) has been identified as the primary site of neural processing and filtering of VWM information (Lara et al., 2015).

The VWM system selects relevant visual information and suppresses distractors based on goal-oriented attentional controls dictated by the PFC. This suppression of distractors is simply the process of filtering out the task-irrelevant information that is obtained from visual stimuli, to ensure that it does not take up space in our limited WM storage, preventing what is known as task-irrelevant interference (Awh et al., 2008; Liesefeld et al., 2020, Lorenc et al., 2021). By doing so, VWM performs a crucial role for our optimal and efficient daily functioning in situations where distraction can lead to suboptimal or failed task completion, such as school, work, and even trying to follow a grocery list (Awh et al., 2008).

It should come as no surprise that inefficiency in VWM can have severe consequences in such situations. For example, if one needed to quickly buy apples at the grocery store, it would be necessary to be able to sort and filter task-irrelevant visual information such as items that are not apples, like other fruit and distracting advertisements (Geng, 2014).

Analyzing the subprocesses of WM is critical to the development of more effective treatments for neurological disorders that inhibit efficiency of VWM, such as attention deficit hyperactivity disorder (ADHD) (Kofler et al., 2008). As one of the most prevalent disorders in children, ADHD can have detrimental impacts on learning and school experiences. Children with ADHD have been shown to score lower on working memory (WM) tests, aligning with lower WM capacity overall (Kofler et al., 2008). Therefore, deepening our understanding of VWM and distractor suppression as a mechanism of cognitive filtering is crucial not only to grow our understanding of the brain and its neural processes, but also to better understand and create effective treatments for widespread disorders such as ADHD.

To better understand distractor suppression at the level of single neurons, this study addresses the following question: Are there any significant differences in the firing rates of individual LPFC neurons during the distractor presentation period of a VWM task between success and error trials that can predict the outcome of a trial? It aims to do this through statistical analysis of neural recordings from the lateral prefrontal cortex (LPFC) of a monkey during an oculomotor delayed response task. Since neural activity at the period of a VWM trial where the distractor is presented often includes information on whether or not single cells encode (internalize) the visual stimulus in WM, it is essential to analyze such activity when attempting to understand distractor suppression.

Literature Review

The prefrontal cortex (PFC) has long since been identified as the center for decision making, among other executive functions (Miyake & Friedman, 2012). The PFC and its role in WM has been studied fairly extensively, and although there has been debate regarding the exact structure that holds and maintains visual stimuli as neural representations in WM, it is agreed that the PFC plays a crucial role in VWM as the controller of the processing of relevant and irrelevant visual information through executive functions (Lara et al., 2015; Miyake & Friedman, 2012). This paper focuses on the PFC, specifically the lateral prefrontal cortex (LPFC), which is a key contributor to working memory mechanisms and capacity in primates, including humans.

Mechanisms and Capacity of Visual Working Memory

The process of VWM is thought to be composed of two primary components: target enhancement and distractor suppression (Awh et al., 2008; Liesefeld et al., 2020). Target enhancement is one part of attentional control that uses a top-down mechanism (signal from the PFC to process visual stimuli) that prioritizes task-relevant visual information to enter working memory (Liesefeld et al., 2020). Distractor suppression is the other attentional control mechanism that blocks task-irrelevant information from entering WM, preventing distraction and allowing goal-oriented behaviors to proceed (Geng, 2014).

Distractor suppression can be further split into proactive and reactive subprocesses. Geng (2014), a review examining these two subprocesses, established that proactive suppression is the suppressing of sensory information associated with distractors before they appear, which can be done by enhancing target priority or actively suppressing distractor features based on prior knowledge of these features. However, it is cognitively taxing and has metabolic limitations, making it impractical to rely on. Reactive suppression is the attentional rejection of unexpected distractors after they appear (Geng, 2014). The faster this suppression is, the more effectively we are able to reject distractors and focus our attention on task-relevant information. Geng (2014) suggests that it is more pertinent to study the process and mechanisms of reactive suppression, as it is inevitable to encounter unanticipated distractions throughout our daily lives and suppressing them is necessary for goal-oriented tasks to be completed. Reactive suppression can be tested in VWM tasks where the subject(s) is not informed of the visual and/or temporal details of a distractor's appearance during the task.

Working Memory and ADHD

Not only has working memory been identified as a cognitive function that can be severely impaired by ADHD and other cognitive disorders, but training to improve WM has been proven to lead to improvement on WM tasks in children with ADHD (Klingberg et al., 2010). Klingberg et al. (2010) aimed to evaluate the effect of WM training on children with ADHD when performing trained and non-trained WM tasks by engaging young children with and without ADHD through a "double blind, placebo controlled design". Alongside finding that WM training improved the performance of children with ADHD on both trained and untrained (new) WM tasks, this study reflected that in order to further advance WM training and develop other treatments for WM capacity deficits in children with ADHD, deeper exploration of WM mechanisms is required (Klingberg et al., 2010).

Further Study of Distractor Suppression

Studies exploring the mechanisms of VWM have researched its neural substrates using technologies such as fMRI, EEG, and single cell electrodes, amongst others. Each of these neural activity recording methods come with benefits and disadvantages, and all have been used to contribute to our understanding of VWM. Researchers have aimed to grow this understanding by analyzing sub processes of VWM through WM tasks with varying components, many focusing simply on WM and its representation in the brain (Pessoa et al., 2002). For example, single cell studies on primates have identified the significance of firing rate activity at different points during VWM tasks, while fMRI studies on humans have demonstrated how neural network activity attributed to VWM can predict the success of a trial (Pessoa et al., 2002).

Studies looking at the predictive capability of neural activity on a trial-by-trial basis, such as Pessoa et al. (2002), often utilize WM tasks without distractor components, and conclusions drawn by such studies cannot be applied to the distractor suppression subprocess of VWM. Studies focusing on distractor suppression have been conducted using various VWM tasks to grow our understanding of reactive distractor suppression. For example, Liesefeld et al. described the inverse correlation between the electrophysiological marker of distractor positivity (pD) and working memory capacity (Liesefeld et al. 2020). A 2012 study also exploring the neural mechanisms of distractor suppression noted that the LPFC, specifically the dorsolateral prefrontal cortex (dLPFC), significantly contributes to the regulation of the process by proving that reversible inactivation of dLPFC results in impaired distractor suppression (Suzuki & Gottlieb, 2012). However, when discussing the relationship between neural activity during distractor suppression-task trials and the trial outcome, this study simply found that the “distractor responses in the dLPFC were positively correlated with the monkey’s error rates”, but did not touch on the predictive capability of those distractor responses on the monkey’s behavior (Suzuki & Gottlieb, 2012). In other words, the predictive capability of electrical activity analyzed in studies such as Pessoa et al. has not yet been extended to VWM tasks involving distractors. This form of analysis not only identifies the relationship between neural activity and behavioral outcome, but also determines how significant the specific neural activity being analyzed is to the behavior based on prediction accuracy.

This study uses a similar predictive analysis strategy to assess the relationship between single neuron distractor specific firing rates and VWM trial outcome. Because visual working memory is a crucial cognitive process that is often negatively impacted by several cognitive disorders such as ADHD, understanding the mechanisms that contribute to individual differences in WM capacity, such as the process and efficiency of distractor suppression, is necessary for growing our understanding of the disorders themselves and their potential treatments (Miyake & Friedman, 2012).

Methodology

Dataset

This study analyzed data from a previous experiment run by Professor Camilo Libedinksy in his National University of Singapore (NUS) Department of Psychology neuroscience research laboratory. Professor Libedinksy provided access to the raw data and consent for its use. In the experiment, the electrical activity of 202 single neurons from the LPFC of a monkey was recorded using electrodes while the monkey engaged in a VWM oculomotor delayed response task involving a distractor. The monkey was presented with a visual stimulus on a screen, and trained to understand which stimulus represented the task-specific target that it was supposed to remember the location of, and which stimulus represented the distractor, which it was supposed to ignore. The target and distractor were differentiated only by color. A dot on the screen, known as a fixation point, was used as a go-cue, where the monkey was trained to give its task response after the disappearance of the dot.

The target was presented first, for 300 milliseconds (ms), and following a 1 second delay period, the distractor was presented for 300 ms as well. After another 1 second delay, the fixation point disappeared from the screen, prompting the monkey to move its eyes toward the remembered location of the target. An eye movement, or microsaccade, to the correct location was categorized as a success trial, while eye movement to any other location was categorized as an error trial. The data from this experiment has not been published, and the firing rates between error and success trials have not been analyzed.

Analysis Methodology

This study first analyzed the primate LPFC neuronal firing rate data using two statistical analysis methods: the t-test and logistic regression model. Additional classification machine learning techniques, including the decision tree, random forest, and neural network were then used to further assess the predictive capability of distractor-specific firing rate on VWM trial outcome.

An independent two-tailed t-test was run for each cell to compare the mean firing rates during the distractor presentation period for all success trials and all error trials of the experiment ($n = 202$). The t-test was chosen as the simplest method to determine which cells would be appropriate to use when identifying a potential relationship between firing rate and trial outcome. For the purposes of this study, the error trials were defined as those in which the monkey incorrectly reported the location of the target, while success trials were those in which the monkey correctly reported the target location. The distractor presentation period consisted of the 300ms that the distractor was presented for, as well as a subsequent 400ms to account for any delay in firing rate in response to the distractor. This analysis operated at a significance level of $\alpha < 0.01$, and a t-test result with this p-value determines that there is a significant difference between the two firing rate distributions (success trials vs error trials) for a particular cell. This significance indicates that the firing rates of the cell at the distractor presentation period hold information about the success of the trial.

A logistic regression was used to quantify the relationship between single cell firing rate at the distractor presentation period and trial success for only the cells that have a significant difference in firing rates between success and error trials (as per the t-test results). This analysis was chosen because of the dichotomous and categorical nature of the dependent variable, trial outcome (success/error). It is important to note that this analysis does not determine any form of causation, that is, the reason behind trial success based on firing rate. It only determines whether there is any relationship between firing rate (independent variable) and trial success. This method of analysis was used by Pessoa et al. (2002) for a similar purpose: to quantify the contingency between fMRI amplitude and subject performance, where a logistic regression “revealed that fMRI signal amplitude during the delay interval predicted successful performance on a trial-by-trial basis” (Pessoa et al., 2002).

Similarly, this study attempted to fit a logistic model to the firing rates of cells with a significant t-test result, where the "slope of the best-fitting logistic function measures the strength of the predictive effect" (Pessoa et al., 2002). In other words, the results of the logistic regression determined whether the firing rate at the distractor presentation period of a cell can predict the outcome of the trial and to what degree of accuracy it can do so.

In order to conduct the logistic regression, it was necessary to format the firing rates into input and output values, respectively, that the model could understand. The original dataset stored the single neuron firing rates in a three dimensional array composed of cell number, trial number, and "bin" number, respectively. Each "bin" was simply the firing rate of a particular 50 millisecond period in the graph. The outcome of each trial of each cell, either success (1) or error (0) was stored separately, which became "y" values of the logistic regression.

The 700 millisecond distractor presentation period being analyzed in this study included 14 bins in total, which were labeled as bins 32 to 45 in the dataset. These firing rates needed to be stored into a data frame that the logistic regression could use as an input, where each row contained the firing rates across the 14 bins for one trial. However, in the original experiment, not every cell participated in the same number of trials, which meant that the firing rates for the cells with fewer trials had to be zero-padded before converting to a single data frame. The logistic regression was first run with these zero-padded trials included, and then again after removing them.

Decision tree and random forest classifiers were then used to test the relationship between LPFC distractor-specific firing rate and VWM task outcome. These classifiers were chosen for their relative simplicity and ability to map nonlinear relationships. The decision tree also allows for a clear visual representation of the most influential features, or bins, towards success or error outcomes. Both techniques used 14 features, with each feature representing one of the 14 bins of the distractor presentation period.

Lastly, a neural network was used to obtain a final accuracy score, to indicate predictive capability, for comparison. This method of machine learning is a computing system with "neurons" that can "learn" and improve performance on a certain task. The neural network structure was composed of 14 input nodes with 9 hidden layers arriving at a single output node. The rectified linear activation function (ReLU) was used for all hidden layers, the "Swish" activation function was used for the input layer and the standard sigmoid function was used for the output layer.

Results

An independent-samples t-test was conducted to compare the mean firing rates of success and error trials at the distractor presentation period for each individual cell out of the 202 total cells. The null hypothesis was that the two mean firing rates would be the same for each cell. The t-test was used to determine that 51 out of the 202 cells recorded had a statistically significant ($\alpha < .01$) difference in firing rates between success and error trials at the distractor presentation period. Table 1 describes the results of the t-test for one of these 51 cells.

Table 1. Firing Rates Across 700ms for all Success and Error Trials for Cell 20 (t-test statistic = -4.154)

	M	SD	n	p
Success Trials	3.122	0.160	107	3.345e-5
Error Trials	3.485	0.202	108	

Across the 700ms distractor period, cell 20 had a mean firing rate of 3.122 per 50ms and a standard deviation of 0.160 for 107 success trials, compared to a mean firing rate of 3.485 per 50ms and a standard deviation of 0.202 for 108 error trials. The mean firing rate for error trials is statistically higher than the mean firing rate for success trials, indicating that this cell has a greater response to the distractor in trials when the monkey does not correctly report the location of the VWM task-specific target.

The 51 cells with statistically different firing rates were then used to conduct a logistic regression to determine whether the model could predict the outcome of a trial based on the firing rate at the distractor presentation period. Figure 1 presents the true trial outcomes of the dataset alongside the predicted trial outcomes classified by the logistic regression.

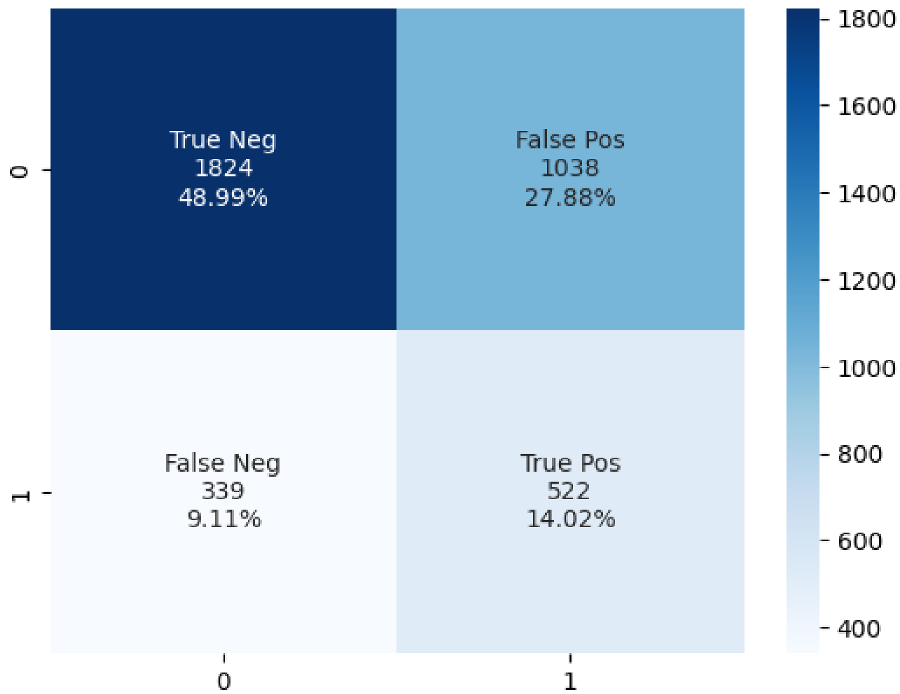

Fig. 1. Heatmap of Logistic Regression Confusion Matrix

Figure 1 visualizes the data used to test the logistic regression through a heatmap, with 3723 total trials in the test. The four boxes represent the predicted and true trial outcomes. True Negative refers to the error trials that were predicted as error, False Positive refers to the error trials that were predicted as success, False Negative refers to the success trials that were predicted as error, and True Positive refers to the success trials that were predicted as success. The scale shows the distribution of trials, with the test data having a much larger proportion of error trials compared to success trials, as is described by the exact number of trials in each box as well as their percentage of the total trials.

A best-fit logistic regression model was trained with 75% of the total data and tested with the remaining 25%, with the portions selected at random. A logistic regression confusion matrix (Figure 1) is used to calculate additional relevant statistics. For example, the model produced an accuracy score of 0.6301[1]. In other words, with every 100 sets of firing rates at the distractor presentation period given, the logistic regression model would be able to accurately predict the outcome of 63.01% of them. This accuracy score indicates that distractor-specific firing rate can predict VWM trial outcome to some extent.

However, to better quantify how well the logistic regression fits the dataset, it is necessary to consider the additional statistics of precision, recall and F1-score. Precision is the fraction of true positives among the true positive

[1] Accuracy score: (True Pos + True Neg)/Total = 0.6301

and false positives, while recall (also known as sensitivity) is the fraction of true positives among the true positives and false negatives. A higher precision value means that the model returns more "relevant" results, or success trials, and higher recall means the model returns most of the total relevant results. The precision and recall scores of the logistic regression model are 33.46%[2] and 60.63%[3] respectively. These values indicate that about a third of the success trials predicted by the model are accurate, but that it is able to correctly predict the majority of the true success trials.

The F1-score refers to the harmonic mean of the precision and recall, and is not sensitive to outliers. It provides a metric that takes into account both the quality and quantity of the relevant results. The F1-score of the logistic regression is 42.58%[4]. To increase this value, both the precision and recall of the model would need to be higher.

The accuracy score produced by the logistic regression was compared to those produced by the decision tree, random forest, and neural network, all of which indicated stronger predictive capability of distractor-specific LPFC neural firing rate on VWM trial outcome. Figure 2 presents a visual representation of the decision tree, which was set with a depth of 3 layers.

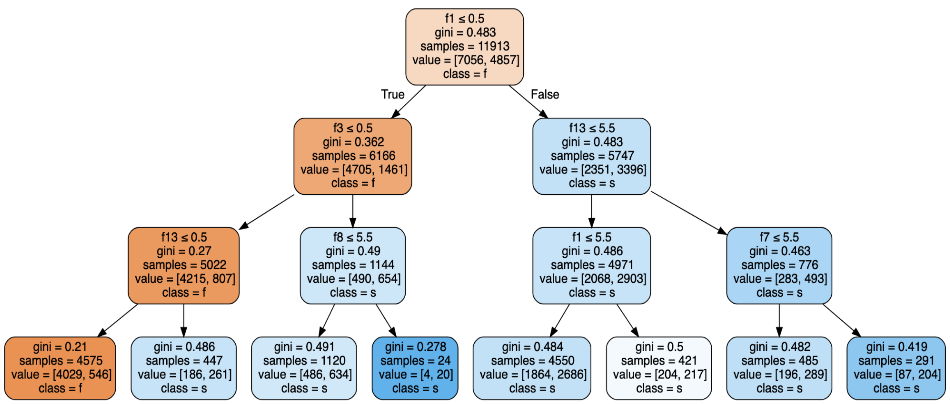

Fig. 2. Diagram of Decision Tree Split

Figure 2 identifies the features, or time bins, that played the largest role in determining the outcome of the VWM task by the decision tree. Features 1, 3, and 13 were identified as particularly crucial, which indicates that the firing rates when the distractor is presented and towards the end of the distractor presentation period may hold information on the outcome of the VWM task.

The decision tree produced an accuracy score of 68.043%. A random forest classifier, using the standard 100 decision trees, similarly formatted but with a max depth of 5 layers, produced a slightly higher accuracy score of 69.45%. In a similar vein, the neural network, for which the structure was described in section 3.2, produced a test accuracy score of 70.26% with a train accuracy 55.09%. The model was built with a decay rate of 0.001 and three dropout layers to prevent overtraining.

Discussion

Summary

The present study explored the relationship between single neuron firing rate and VWM task outcome based on the neural activity in response to a task-specific distractor. This relationship was first quantified using a logistic regression,

[2] Precision = True Positives / (True Positives + False Positives) = 522/(522 + 1038) = 0.3346
[3] Recall = True Positives / (True Positives + False Negatives) = 522/(522 + 339) = 0.6063
[4] F1-Score = 2*((Precision*Recall)/(Precision+Recall)) = 0.4258

where the model was trained and tested with the firing rates and corresponding trial outcomes of 51 cells with significant differences in firing rate between success and error trials. The logistic regression was able to predict trial outcome with 63.66% accuracy, a score moderately above chance, which indicates that there is a relationship between LPFC single neuron activity and behavioral outcome. However, the F1-score of 42.58% indicated that the fit of the logistic model was not as ideally aligned with the data, and further machine learning classification techniques were required to better assess the predictive capability of distractor-specific firing rates for VWM trial outcome. Decision tree, random forest, and neural network methods produced better accuracy scores of 68.043%, 69.45%, and 70.26% respectively. These accuracy scores, which are fairly high above chance, indicate that the response of single neurons in the LPFC to the distractor does predict the monkey's likelihood of correctly remembering the task-specific target location and reporting the answer accordingly.

Impact of Distractor-Specific Firing Rate on Trial Outcome

Understanding that cells fire in response to the visual stimulus being presented is crucial to understanding the significance of the distractor-specific firing rate as it relates to trial outcome. Because this study analyzed the firing rate during the period of distractor presentation and slightly afterward, the firing rates of the cells quantify the degree to which the distractor is encoded and internalized by each individual cell, which can be used to predict the behavioral outcome of a VWM task with fairly high accuracy. Table 1 reveals that the mean firing rate for all error trials of one particular cell is significantly different, and more specifically, greater than the mean firing rate for all success trials. A greater firing rate indicates a stronger response to the distractor, meaning that this particular cell encoded the distractor more in error trials than in success trials, most likely as a result of less efficient distractor suppression. The logistic regression extends this implication to all cells with significant firing rate differences between success and error trials, revealing that the greater the firing rate of a trial, or the less efficient the distractor suppression mechanism is, the more likely the outcome is to be an error.

These findings are in line with previous research about VWM and filtering ability. The efficiency of distractor suppression is based on the degree of neural response to a distractor, supporting the results of this study that show that greater firing rate is correlated with an increase in likelihood for an error, where the monkey was distracted and unable to correctly complete the VWM task. These findings demonstrate the importance of efficient distractor suppression, that is, the mechanism that controls the neural response to distractors when engaging in goal-oriented behavior, to the successful completion of tasks. Following this line of thought, it may be the case that those with lower WM capacity, and less efficient distractor suppression leading to less efficient filtering ability, have a greater single neuron response to task-specific distractors as well.

Conclusion

Our ability to filter through task-specific targets and distractors as well as suppress the latter is crucial even for the simplest of tasks. This study proves that millisecond differences in the timing of single neuron firing are associated with opposing behavioral outcomes. This study ascertained that trial outcome for a VWM task could be predicted by distractor-specific firing rate using a logistic regression, which demonstrated that stronger neural response to the task-specific distractor was associated with trial failure. Because distractor suppression efficiency is crucial to the successful completion of goal-oriented behaviors, specifically, WM tasks, these findings support previous research that people with lowered WM capacity might experience this decrease as a result of lower distractor suppression efficiency (Kofler et al., 2008). By demonstrating that behavioral outcomes can be predicted by the firing of single neurons in the LPFC of the brain, this study emphasizes the importance of addressing disorders such as ADHD not only at the level of behavior, but at the level of neural activity.

The level of specificity in this study of both the analysis and the VWM experiment itself opens the possibility for various manipulations and exploration in future research. For example, changing the nature of the task, the properties of the task-specific target and/or distractor, or the ratio of targets to distractors may reveal more about the WM subprocesses of target enhancement and distractor suppression, as well as their importance to trial outcome. Conducting experiments on humans with VWM tasks involving distractors and using alternative brain imaging techniques, such as EEG, would allow for the analysis of neural activity for much more complex tasks, as well as a comparison of results between single neuron activity and EEG recordings. In terms of analysis, using more sophisticated data analysis techniques may produce more accurate and possibly stronger predictive scores between firing rate and trial outcome. In future study of the neural processes associated with lower WM capacity, it is necessary to address the micro level of neural activity alongside the macro level when considering treatments, particularly for children, for whom WM capacity can play a large role in learning and development.

Acknowledgements

I would like to thank Professor Camilo Libedinsky and Dr. Roger Herikstad from the Department of Psychology at the National University of Singapore for introducing me to this field of research, and for their guidance and support. I would like to thank my teacher, Ms. Stewart, for her encouragement and support throughout the project. I would also like to thank Dr. D. Narayana for helping me understand the statistical methods and machine learning techniques mentioned in this paper.

References

Awh, E., & Vogel, E. K. (2008). The bouncer in the brain. *Nature Neuroscience*, *11*(1), 5–6. https://doi.org/10.1038/nn0108-5

Christensen, L. B., Turner, L. A., & Johnson, B. (2014). *Research methods, design, and analysis* (12th ed.). Pearson.

Geng, J. J. (2014). Attentional mechanisms of distractor suppression. *Current Directions in Psychological Science*, *23*(2), 147–153. https://doi.org/10.1177/0963721414525780

Klingberg, T., Forssberg, H., & Westerberg, H. (2010). Training of working memory in children with ADHD. *Journal of Clinical and Experimental Neuropsychology*, *24*(6), 781–791. https://doi.org/10.1076/jcen.24.6.781.8395

Kofler, M. J., Rapport, M. D., Bolden, J., & Altro, T. A. (2008). Working memory as a core deficit in ADHD: Preliminary findings and implications. *The ADHD Report*, *16*(6), 8–14. https://doi.org/10.1521/adhd.2008.16.6.8

Lara, A. H., & Wallis, J. D. (2015). The role of prefrontal cortex in working memory: A mini review. *Frontiers in Systems Neuroscience*, *9*. https://doi.org/10.3389/fnsys.2015.00173

Liesefeld, H. R., Liesefeld, A. M., Sauseng, P., Jacob, S. N., & Müller, H. J. (2020). How visual working memory handles distraction: Cognitive mechanisms and electrophysiological correlates. *Visual Cognition*, *28*(5-8), 372–387. https://doi.org/10.1080/13506285.2020.1773594

Lorenc, E. S., Mallett, R., & Lewis-Peacock, J. A. (2021). Distraction in visual working memory: Resistance is not futile. Trends in Cognitive Sciences, 25(3), 228–239. https://doi.org/10.1016/j.tics.2020.12.004

Miyake, A., & Friedman, N. P. (2012). The nature and organization of individual differences in executive functions. *Current Directions in Psychological Science*, *21*(1), 8–14. https://doi.org/10.1177/0963721411429458

Olivers, C. N., & Van der Stigchel, S. (2020). Future steps in visual working memory research. *Visual Cognition*, *28*(5-8), 325–329. https://doi.org/10.1080/13506285.2020.1833478

Pessoa, L., Gutierrez, E., Bandettini, P. A., & Ungerleider, L. G. (2002). Neural Correlates of Visual Working Memory: fMRI Amplitude Predicts Task Performance. *Neuron*, *35*, 975–987. https://doi.org/10.1016/S0896-6273(02)00817-6

Suzuki, M., & Gottlieb, J. (2012). Distinct neural mechanisms of distractor suppression in the frontal and parietal lobe. *Nature Neuroscience*, *16*(1), 98–104. https://doi.org/10.1038/nn.3282