Neural heterogeneity shapes the temporal structure of human working memory

Daria Kussovska (dnk2118@columbia.edu)
Department of Biomedical Engineering, Columbia University
New York, NY 10027, USA

Nuttida Rungratsameetaweemana (nr2869@columbia.edu)
Department of Biomedical Engineering, Columbia University
New York, NY 10027, USA

Abstract

The circuit computations that support memory maintenance in humans remain poorly understood. Persistent activity has long been thought to underlie the maintenance of information in working memory (WM). However, most supporting evidence has relied on trial-averaged neural responses, often overlooking the potential role of cellular heterogeneity. To address this, we analyzed single-trial spiking activity from intracranial recordings of neurosurgical patients performing a WM task. We developed a method for putative cell-type classification and examined the temporal dynamics of interneurons and pyramidal cells across encoding and maintenance task phases. Our findings reveal distinct cell-type-specific activity profiles during working memory maintenance and suggest that stimulus-specific information can be retained with minimal persistent firing, highlighting flexible and potentially energy-efficient circuit computations that support information storage in the human brain.

Keywords: working memory; circuit dynamics

Introduction

Working memory (WM) provides temporary storage of information, crucial for everyday tasks, such as learning, language comprehension, and decision-making. Persistent firing of individual neurons across distributed brain areas during retention intervals has been viewed as a primary mechanism for memory maintenance (Fuster & Alexander, 1971; Kamiński et al., 2017). However, sustained spiking is metabolically costly. Much supporting evidence relies on trial-averaged data and often assumes network homogeneity, with limited consideration of cell-type specificity or synaptic interactions shaping cortical dynamics (Lundqvist et al., 2018; English et al., 2017). Thus, it remains unclear how heterogeneous neuronal populations in the human brain maintain information over time.

To address this gap, we analyzed human intracranial recordings during a WM task, focusing on single-trial dynamics. We developed an approach for putative cell-type classification to examine how distinct neuronal populations contribute to memory-related dynamics. Our findings suggest that stimulus-specific information can be maintained without sustained spiking during the maintenance period, supporting alternative accounts of activity-silent mechanisms within recurrent cortical circuits (Stokes, 2015). Moreover, we show that cell types exhibit distinct temporal profiles across task phases, suggesting specialized roles in information maintenance. These findings advance our understanding of WM and underscore the value of cell-type-resolved, temporally agnostic analyses for uncovering the circuit dynamics that support human cognition.

Methods

To investigate the circuit computations supporting WM representations in the human brain, we used a dataset of 902 single cells recorded from neurosurgical patients with epilepsy during a Sternberg WM task (Fig. 1; Kyzar et al., 2024).

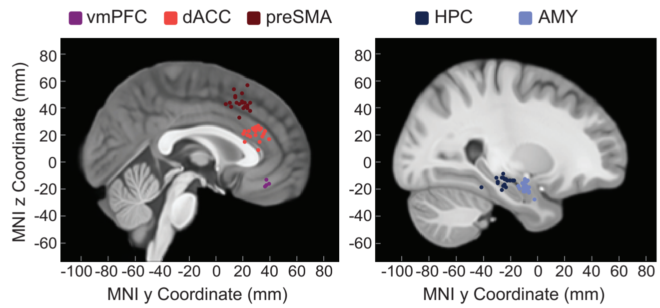

Figure 1: Location of recording sites in MNI space.

For putative cell-type classification, we first extracted autocorrelogram (ACG) features from each cell's spike train through a triple-exponential function. Cells were then clustered using spectral clustering on firing rate, mean ACG, and τ rise values, which have been found to exhibit robust differences between interneurons (IN) and pyramidal cells (PY) (Csicsvari et al., 1998; Dickey et al., 2021). To ensure reliability, we tested our method on a reference dataset (Petersen et al., 2021), yielding 97% accuracy. Firing characteristics were compared across cell types.

To identify bursts of spiking activity, we extended an approach introduced in a recent study (Xie et al., 2024). For each cell, spike times were concatenated across trials and analyzed separately for the encoding and maintenance periods. Using cells grouped by ACG features (mean ACG and τ decay), and by cell type (IN and PY), we binned spikes at 70 ms and smoothed with a Gaussian kernel ($\sigma = 40$ ms); bursts were defined as local peaks exceeding the top 10% in prominence. As a null model, we generated surrogate spike trains using a Poisson process matched to the firing rate of each cell, repeated 100 times. Burst counts from real and surrogate data were compared across participants using paired t-tests or Wilcoxon signed-rank tests, with the Benjamini/Hochberg false discovery rate (FDR) correction for multiple comparisons (Benjamini & Hochberg, 1995).

To decode the encoded stimulus identity across task periods, we constructed cross-temporal decoding maps using a support vector machine (SVM) with leave-one-out cross validation (LOOCV) (Meyers, 2013). Binned activity was extracted from pyramidal and interneuron cells that had been previously classified as concept cells (Quiroga, 2012).

Results

While persistent dynamics in the maintenance period can be observed through averaging activity across population of cells (bootstrap test on paired differences, $p < 0.05$), this does not necessarily reflect how individual cells maintain information independently over time (two-sample t-test, $p < 0.05$).

After spectral clustering (Fig. 2C), from the 902 cells, 92 were identified as pyramidal (PY; 10.20%) and 202 as in-

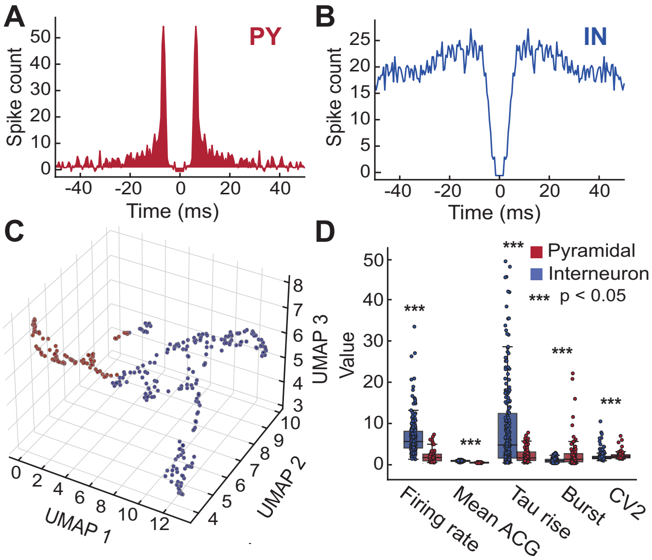

Figure 2: (A-B) Spike autocorrelogram (ACG) of an example pyramidal cell (PY) and an interneuron (IN). (C) PY and IN form separate clusters based on mean ACG, firing rate, and ACG tau rise. (D) Box plots comparing features of PY and IN.

terneurons (IN; 22.39%). Our clustering effectively distinguished bursty from non-bursty cells, as revealed by burst index and coefficient of variation squared (CV2) measures (Mann–Whitney U test, all $ps < 0.05$; Fig. 2D), reflecting distinct temporal dynamics of activity with possible contributions to WM.

Bursting analysis indicated that during encoding, only certain subpopulations, especially low ACG cells (W=0, $p = 0.004$) and PY cells (W=0, $p = 0.009$), exhibit temporally clustered activity that exceeds stochastic firing. Other cells appeared to follow a Poisson-like activity (Wilcoxon signed-rank test, $p > 0.01$), suggesting that encoding in these cells may rely more on rate coding or other mechanisms rather than burst firing. However, during maintenance, all subgroups of cells demonstrated structured bursting, implying coordinated, non-random dynamics (Wilcoxon signed-rank test, $p < 0.01$) (Fig. 3A). Decoding analyses using concept PY and IN cells revealed differences in how each cell type contributes to the representation and maintenance of stimulus information over time (Fig. 3B).

Discussion

Our findings demonstrate that while prominent persistent activity is observed when averaging across trials and neurons, these dynamics may not fully capture the mechanisms underlying working memory maintenance at the single-trial level. Our cell classification analysis yielded notable results, showing higher bursting and CV2 values for pyramidal cells, which are in line with empirical studies in rodents (Raus Balind et al., 2019). These findings highlight the utility of cell-type-specific analyses in generating improved understanding of circuit computations that underlie human cognitive processes.

Importantly, we observed distinct effects in bursting dynamics across neuronal subtypes during working memory. While encoding was characterized by relatively stochastic, Poisson-like firing patterns, the maintenance period showed a marked increase in structured, burst-like activity across the neural population. This shift suggests a dynamical reorganization in how information is carried during memory maintenance. Finally, decoding analyses revealed distinct temporal profiles between interneurons and pyramidal cells, suggesting complementary roles in the representation and maintenance of stimulus-specific information. In line with emerging theories, our findings provide cellular-level evidence for activity-silent schemes in WM. This emphasizes the role of more dynamic, computationally-efficient strategies for maintaining information over time.

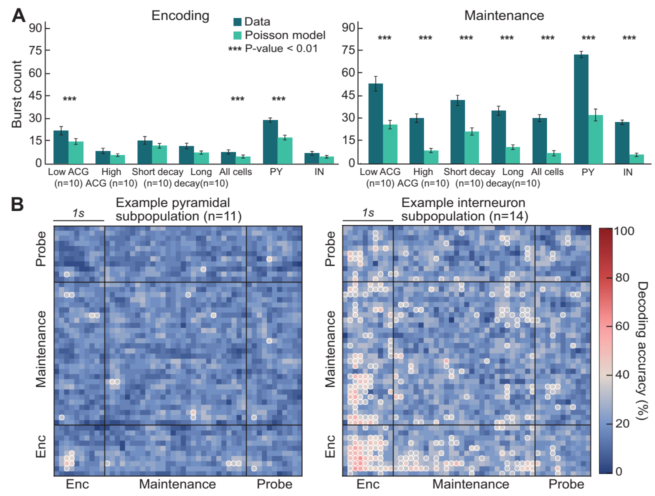

Figure 3: (A) Burst counts for different neuronal subgroups during encoding and maintenance. Error bars represent standard error of the mean (SEM). (B) Decoding analysis on PY and IN concept cells. White circles indicate decodability higher than chance ($p < 0.05$).

Limitations and Broader Impacts

One limitation of the current work is the absence of ground truth labels for human cells as pyramidal cells or interneurons, which naturally constrains the interpretation of our cell-type-specific results. While we cannot definitively confirm the identities we assigned to these cells, our findings suggest that distinct populations may contribute differently to working memory maintenance. These functional distinctions, shaped by properties like connectivity patterns, firing dynamics, and temporal profiles, provide a useful lens for thinking about how the brain efficiently maintains information beyond models that assume homogeneity across cells. More broadly, our findings highlight the importance of moving beyond homogeneous assumptions to consider diverse cellular contributions in studying the circuit dynamics underlying cognition.

References

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, *57*(1), 289–300.

Csicsvari, J., Hirase, H., Czurko, A., & Buzsáki, G. (1998). Reliability and state dependence of pyramidal cell–interneuron synapses in the hippocampus: An ensemble approach in the behaving rat. *Neuron*, *21*(1), 179–189. doi: 10.1016/S0896-6273(00)80525-5

Dickey, C. W., Sargsyan, A., Madsen, J. R., Eskandar, E. N., Cash, S. S., & Halgren, E. (2021). Travelling spindles create necessary conditions for spike-timing-dependent plasticity in humans. *Nature Communications*, *12*(1). doi: 10.1038/s41467-021-21298-x

English, D. F., McKenzie, S., Evans, T., Kim, K., Yoon, E., & Buzsáki, G. (2017). Pyramidal cell-interneuron circuit architecture and dynamics in hippocampal networks. *Neuron*, *96*(2), 505–520.e7. doi: 10.1016/j.neuron.2017.09.033

Fuster, J. M., & Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science*, *173*(3997), 652–654. doi: 10.1126/science.173.3997.652

Kamiński, J., Sullivan, S., Chung, J. M., Ross, I. B., Mamelak, A. N., & Rutishauser, U. (2017). Persistently active neurons in human medial frontal and medial temporal lobe support working memory. *Nature Neuroscience*, *20*(4), 590–601. doi: 10.1038/nn.4509

Kyzar, M., Kamiński, J., Brzezicka, A., Reed, C. M., Chung, J. M., Mamelak, A. N., & Rutishauser, U. (2024). Dataset of human-single neuron activity during a sternberg working memory task. *Scientific Data*, *11*(1), 89. doi: 10.1038/s41597-024-02943-8

Lundqvist, M., Herman, P., & Miller, E. K. (2018). Working memory: Delay activity, yes! persistent activity? maybe not. *The Journal of Neuroscience*, *38*(32), 7013–7019. doi: 10.1523/JNEUROSCI.2485-17.2018

Meyers, E. M. (2013). The neural decoding toolbox. *Frontiers in neuroinformatics*, *7*, 8.

Petersen, P. C., Siegle, J. H., Steinmetz, N. A., Mahallati, S., & Buzsáki, G. (2021). Cellexplorer: A framework for visualizing and characterizing single neurons. *Neuron*, *109*(22), 3594–3608.

Quiroga, R. Q. (2012). Concept cells: the building blocks of declarative memory functions. *Nature Reviews Neuroscience*, *13*(8), 587–597.

Raus Balind, S., Magó, Á., Ahmadi, M., Kis, N., Varga-Németh, Z., Lőrincz, A., & Makara, J. K. (2019). Diverse synaptic and dendritic mechanisms of complex spike burst generation in hippocampal ca3 pyramidal cells. *Nature communications*, *10*(1), 1859.

Stokes, M. G. (2015). 'activity-silent' working memory in prefrontal cortex: a dynamic coding framework. *Trends in Cognitive Sciences*, *19*(7), 394–405. doi: 10.1016/j.tics.2015.05.004

Xie, W., Jr, J. H. W., Chapeton, J. I., El-Kalliny, M., Jackson, S. N., Inati, S. K., & Zaghloul, K. A. (2024). Neuronal sequences in population bursts encode information in human cortex. *Nature*, *635*(8040), 935–942. doi: 10.1038/s41586-024-08075-8