Quantifying network behavior in the rat prefrontal cortex: a reproducibility crisis

Congzhou M. Sha[1,2], Jian Wang[2], Richard B. Mailman[2,3,*], Yang Yang[2,*], Nikolay V. Dokholyan[1,2,4,5,6,*]

[1] *Department of Engineering Science and Mechanics, Penn State University, University Park, PA USA*

[2] *Department of Pharmacology, Penn State College of Medicine, Hershey, PA USA*

[3] *Department of Neurology, Penn State College of Medicine, Hershey, PA USA*

[4] *Department of Biochemistry and Molecular Biology, Penn State College of Medicine, Hershey, PA USA*

[5] *Department of Chemistry, Penn State University, University Park, PA USA*

[6] *Department of Biomedical Engineering, Penn State University, University Park, PA USA*

*Corresponding authors
R.B.M. rbm17@psu.edu
Y.Y.: yyang1@pennstatehealth.psu.edu
N.V.D.: dokh@psu.edu

Abbreviations

Neuroelectrophysiology (NEP), dynamic time warping (DTW), medial prefrontal cortex (mPFC), Kolmogorov-Smirnov test (KS-test), receiver operating characteristic (ROC), area under the curve (AUC)

Abstract (122 words)

The question of how consciousness and behavior arise from neural activity is
fundamental to understanding the brain, and to improving the diagnosis and treatment
of neurological and psychiatric disorders. There is significant literature in murine and
primate species on how behavior is related to the electrophysiological activity of the
medial prefrontal cortex and its role in working memory processes such as planning and
decision-making. We have now applied dynamic time warping and statistical tests to
data from neuron spike trains and local field potentials, to quantify neural network
synchronicity and correlate neuroelectrophysiology with rat behavior. These analyses
have shown the important limitations of existing experimental designs and provide an
outline for improved experimental approaches that will yield more robust and
reproducible findings.

Introduction

Hodgkin and Huxley's model of neuron action potentials derived from experiments with the squid giant axon[1] was the seminal event in neurophysiology[2-4]. The central dogma of modern neuroscience is that neuron electrochemical activity and connectivity at the microscopic level can provide a clear understanding of complex behaviors. Thus, measuring electrical signals can be sufficient to integrate neural networks' chemical signaling and activity.

The visual cortex, and its relationship to visual processing, is one of the better-understood neural networks because there are significant spatial correlations in signals that result from local connections between neurons. These spatial autocorrelations result in a high signal-to-noise ratio due to the strong intensity of the electric field produced by local synchronized neuron firing. Additionally, the neuroelectrophysiology of individual cells in the visual pathway has allowed mathematical modeling of the performance of specific neurons[5,6]. The spatial autocorrelations of visual neural activity are sufficiently high that low-resolution, indirect measurements of neuron activity, such as local blood flow[7-9] and extracranial electrodes[10], are sufficient to reconstruct perceived images. The organization of the visual cortex may be explained by the fact that images projected on the retina are spatially correlated, and therefore biological neural networks have adapted to take advantage of these correlations.

From an evolutionary perspective, the visual cortex is an ancient structure, and therefore the neural networks involved in visual processing have optimized over time. In contrast, the prefrontal cortex as part of the neocortex is one of the least developed brain regions, especially through evolutionary analysis of its size in mammals[11,12]. Unlike the visual cortex, the prefrontal cortex does not seem to possess highly local and regular spatial organization[13]. Because of the lack of a regular organization of the prefrontal cortex and the consequent low intensity of electrical signals, electrophysiology is one of the only methods with adequate sensitivity to probe the fine structure of the prefrontal neural network. Unfortunately, such measurements are highly invasive, and experiments are only feasible on laboratory animals such as rats[14-16]. To find the relationship between neuron firing and behavior, methods are required to analyze these complex electrical measurements.

The characteristics of this neuron network constrain the firing patterns of neurons. Therefore, the methods are required to capture key network characteristics from recorded spike trains and local field potentials (LFPs). Our hypothesis is that these network characteristics are embodied by the synchronicity of neuron firing: neurons that work in concert will be highly (anti-)correlated, while neurons that process independent data will not be correlated. We chose to investigate neuron synchronicity using the dynamic time warping (DTW) method. DTW is an efficient, non-parametric approach to determine the best alignment of two time series, such that the overall shapes are matched. Compared to classical Fourier and wavelet analysis, DTW gives a global alignment of the data that is robust to local variations in timing. We predicted that DTW could discern if two neurons fire synchronously or asynchronously.

Thus, we analyzed electrophysiology data collected from rats performing the T-maze
task, which evaluates working memory. We use previously published spike train and
LFP recordings taken from the rat medial prefrontal cortex (mPFC)[14–16], along with new
data from three additional rats (rats A, B, and C) recorded using published methods[16].
We investigated if DTW would be useful in elucidating the connection between neural
activity and behavior.

Results

In the data we used from the three published studies (Table 1), rat neuron spike trains
and/or LFPs were recorded as the rats performed in T-maze (**Figure 1**-A). The exact
experimental details for the three studies differed slightly[14–16], but their combined
scientific goal was to analyze the link between neuron recordings and working memory
performance in the T-maze task. We focused on the four-second window of time
centered on the moment at which the rat leaves the T-intersection (decision box in
Figure 1-A), presumably making its choice. This window contained two seconds of pre-
decision neural activity (e.g., cognition and decision-making process) and two seconds
of post-decision neural activity (e.g., evaluation of reward or lack of reward). Our
hypothesis was that neural activity was significantly affected by the independent
variables of (1) the timing of the recording relative to the choice, (2) the correctness of
the choice, and (3) the sampled neurons (Table 2).

We used the DTW method on the neuron spike trains from all three studies and from
the local field potentials for Ito et al.[14] and Yang et al.[16] (Figure 1-B & 1-C). Since DTW
provides a numerical measurement of dissimilarity between firing of neurons, we
converted the resulting DTW matrices into undirected, unweighted graphs
(Supplemental Methods). For each set of neuron spike trains and/or LFPs, we
quantified the connectivity of the resulting DTW graph by a single number, d_{crit}. To
evaluate for the presence of differences among sets of experiments, we used non-
parametric statistical tests (Kruskal-Wallis, Kolmogorov-Smirnov, Mantel, and Boschloo
exact tests) to mitigate both our lack of knowledge of the true underlying probability
distributions and the small sample sizes. These statistical tests and the calculation of
d_{crit} for a given DTW matrix are described in detail in the Supplemental Methods. We
visualized d_{crit} as a function of our independent variables in Figure 2 (spike trains) and
Supplemental Figure 1 (LFPs).

We defined an ***event*** as a single traversal of the T-maze by a rat. We defined a ***run*** as
all events recorded on a single day for a single rat. In the Ito[14] and Stout[15] studies,
electrodes were moved between rats, and, therefore, a different set of neurons were
recorded between runs. We set the level of statistical significance to α =0.05, so that we
expect a false positive rate of α.

Neural network firing is consistent across a single run
Our first question was: when analyzing d_{crit}, can we pool events from a single run? The
null hypothesis was that during a single run, we may assume that the electrode
placement is unchanged and that the rat basal activity is also unchanged between
recordings in the same run, as measured by d_{crit}. The alternative hypothesis is that

specific factors affect d_{crit} between events in the run, such as the categorical variables in
Table 2.

First, we analyzed the normality of d_{crit} with the Shapiro-Wilk test (Supplemental Tables
1 and 2). For neuron spike trains, we found that in 96.7% (60 out of 62) of runs, the
samples were not normally-distributed with $p < 0.05$, and for LFPs, we found that 80.0%
(20 out of 25) were not normally-distributed. Comparing spike trains to LFP with
Boschloo's exact test, these fractions were not significantly different (p=0.615). We
performed subgroup analysis on the categorical variables (1) correctness and (2) timing
and we found that this proportion was unaffected by further stratification by these
variables. Therefore, we concluded that the Kruskal-Wallis non-parametric test was
more appropriate than classical one-way analysis of variance (ANOVA), due to violation
of the assumption of normally-distributed data and due to the limited number of events
per run.

Only 4.8% (3 out of 62) of the spike train runs displayed significant (Kruskal-Wallis $p <$
0.05) event-to-event variance of d_{crit}, regardless of stratification by correctness and/or
timing; the corresponding proportion for LFPs was 12% (3 out of 25). For completeness,
we also performed ANOVA and found that 11.2% (7 out of 62) of the spike train runs
displayed significant event-to-event variance, regardless of stratification; the
corresponding proportion for LFPs was 20% (5 out of 25). We tested if Kruskal-Wallis
gave different results as compared to ANOVA using Boschloo's exact test, and found
that neither the spike train data (p=0.238) nor the LFPs (p=0.378) were significantly
different.

Since few of the runs (within our 5% margin of error for false positives from $\alpha = 0.05$)
had a significant event-to-event variance in the spike trains, we concluded that we may
perform pooled analysis of d_{crit} for the spike trains of events across a single run.
Boschloo's exact test on the LFPs also demonstrated a non-significant deviation from
the 5% margin of error (p=0.413 for Kruskal-Wallis, p=0.140 for ANOVA).

Neural network structure is revealed by the DTW matrix
Instead of characterizing network connectivity with a single number (d_{crit}), we used all
the information available in the DTW matrix to determine if firing patterns were
consistent, using non-parametric Mantel tests. We performed all pairwise comparisons
for events in the same rat to ask the question: given a single run, is there a difference in
Pearson correlation coefficient of before vs after depending on if we compare (1) before
vs after of the same event to (2) before vs after of different events? We performed
pairwise Mantel test comparisons among events in the same run. We performed
Kruskal-Wallis and Kolmogorov-Smirnov tests to determine if these distributions of the
correlation coefficient and *p*-value were different (Supplemental Figures 2 and 3,
Supplemental Table 3). Overall, we found that there was no significant difference event-
to-event for the before vs after DTW matrix correlation.

When we stratified by correctness (Supplemental Table 3), we did observe an effect for
Kolmogorov-Smirnov tests on the Pearson correlation coefficient in 35.7% (5 out of 14)
runs. When we used d_{crit} alone, this proportion was 7.1% (1 out of 14). To evaluate if
the KS-test/Mantel test combination was more powerful than the Kruskal-Wallis test on

d_{crit}, we performed a Boschloo exact test on the contingency table, which failed to reach
statistical significance (p=0.159). We concluded that there was mixed evidence for the
efficacy of Mantel tests of before vs after DTW matrices in determining if a rat made the
correct or incorrect choice at the T-intersection.

Neural network differences exist across multiple runs for a single rat

Our second question was: when analyzing d_{crit}, can we pool all the events for a single
rat, regardless of the day the recordings were taken? Using the Kruskal-Wallis test, we
found that 70% (7 of 10) rats displayed significant ($p < 0.05$) day-to-day variance of d_{crit}
in their spike trains, regardless of additional stratification by (1) correctness and (2)
timing (Supplemental Table 1). The caveat was that we excluded six rats from the Yang
study [16] in this analysis, because we were unable to apply the Kruskal-Wallis test since
recordings for those rats were only performed on a single day (dof = 0). Additionally, the
three rats which did not display significant differences across runs were 17914 and
18471 from the Ito study[14] and rat A from Yang[16]. The same analysis was performed for
LFPs, in which 80% (4 out of 5) rats (all rats except rat A from Yang[16]) displayed
significant variance (Supplemental Table 2). For the LFPs, the 5 rats from Stout[15] were
excluded since the LFPs were not recorded for those rats. Boschloo's exact test
showed no difference between spike trains and LFPs ($p = 0.922$). We conclude that the
resting cognitive state for each rat may be subject to change depending on the day that
the recording is performed.

Significant neural network heterogeneity exists between rats and studies

Our third question was: does d_{crit} depend on the specific rat? Using the Kruskal-Wallis
test, we found that the Stout[15] and Yang[16] studies displayed significant ($p < 0.001$) rat-
to-rat differences in their spike trains, whereas Ito[14] was non-significant ($p = 0.206$). Ito[14],
however also had the fewest rats ($n = 3$) potentially leading to the insufficient statistical
power of the Kruskal-Wallis test. The LFPs showed significant ($p < 0.001$) rat-to-rat
variation in both Ito[14] and Yang[16]. Pooling all rats, we confirmed ($p < 0.001$) that the
specific study under consideration affects the value of d_{crit} for both spike trains and
LFPs. We conclude that significant heterogeneities exist among studies, animals, and
runs. The statistical tests supporting this conclusion are summarized in Supplemental
Tables 1 (spike trains) and 2 (LFPs).

DTW distance is inconsistently correlated with physical electrode distance

Do correlations in firing among neurons reflect the underlying spatial distance between
the neurons? We used both Welch's t-test for unequal variances and the non-
parametric Kolmogorov-Smirnov test (KS-test) to test for differences in the distribution
of the DTW distance when grouped by the spatial distance of the associated electrodes.

In all three studies, the d_{crit} between neurons recorded from the same electrode versus
from different electrodes were statistically significantly different ($p < 0.001$) by both the t-
test and KS-test (Supplemental Figures 4, 6-8); however, this classification was useless
as a predictor or regressor[17]. For example, using d_{cutoff} as a classifier of d_{crit} for
determining if two neurons were recorded from the same electrode vs from different
electrodes, the receiver operator characteristic (ROC) curve is essentially no different

from random guessing (Supplemental Figure 5), with area under the curve (AUC) of
0.51. However, if we stratify data by run (Supplemental Figures 6-8), the ROC curve
and AUC vary for specific rats/runs (Supplemental Figure 9), and in some cases
demonstrate high predictive value.

For the Yang study[16], we were able to obtain the exact geometry of the electrode.
Therefore, we were able to stratify DTW distances by the physical distance of the
neurons measured (**Figure 3**). For this stratification, we were unable to determine a
consistent trend. Rats B and C demonstrate decreased DTW distance (increased
synchronicity) in both neural spike trains and LFPs, whereas the remaining rats
demonstrated slightly increasing DTW distance (decreased synchronicity) in both,
except for rat S. For rat S, however, the large decrease in DTW distance seen in the
spike trains may not be captured in the LFPs, because there was no zero-distance
comparison for LFPs, whereas multiple spike trains recorded from the same electrode
could be thought of as having zero distance.

Discussion

There are two major issues with attempting to measure a specific population of neurons.
First, precise surgical implantation of electrodes is required; second, enough neurons
must be sampled to capture overall neural network dynamics. These two issues
synergize at our current level of intracranial electrode technology, complicating fine
measurements of neurons. While the rat brain is far smaller than a human brain (~1 cm),
it still contains an estimated 21 million neurons[18]. Reproducible surgical implantation in
a specific area of the rat brain such as the medial prefrontal cortex is therefore highly
dependent on the fine-motor skill of the researcher. Additionally, due to the lack of
spatial organization of the prefrontal cortex, we postulate that the surgical sampling of
neurons is essentially random. Therefore, the chances of measuring the same set of
neurons among different rats or after adjusting electrode positioning is extremely low,
an issue which can only be mitigated through the measurement of larger populations of
neurons. In the Methods, we discuss these issues further with a simple mathematical
illustration (*Reproducibility of electrode recordings*, **Figure 1**d).

In our analysis, we demonstrated that the DTW distance and the computed parameter
d_crit captured some of the mPFC neural network firing dynamics for both the spike trains
and LFPs. We found weak evidence that the correctness of rat choice influences the
firing dynamics (Supplemental Table 3). However, we expect that many more
experiments are needed to confirm the classification of the correctness of a decision
based on mPFC activity because our results indicated that firing events could be pooled
across the same run, but not across different runs or studies (Supplemental Tables 1
and 2). In terms of spatial dependence of firing synchronicity, we found that for specific
rats and runs, the classification of DTW distances for recordings taken from the same vs
different electrodes was both statistically significant (Supplemental Figures 4, 6-8) and
useful as a regressor (Supplemental Figure 9). For the Yang study[16], the spatial
dependence of DTW distances was highly dependent on the specific rat sampled
(**Figure 3**), providing evidence of the lack of spatial organization of the prefrontal cortex.

The difficulties we encountered in analyzing spike trains and LFPs can be addressed
through changes to the experimental design. The number of neurons sampled is an
important characteristic to consider, especially if only a few neurons were recorded and
some recordings represented only a single neuron (e.g., Stout[15]). If the mPFC is
composed of separate neural networks which perform specific, modular tasks, it is
unlikely that recording a low number of neurons will provide meaningful, interpretable
results. Since firing dynamics were consistent across a single run, we posit that multiple
recordings of a single neuron population are the best approach to characterizing
network firing behavior. The DTW spatial dependence remained consistent across
multiple runs for rats A and C (**Figure 3**), therefore we are reasonably confident that we
recorded the same population of neurons across the runs. In contrast, because the
electrodes were adjusted after each run, Ito[14] and Stout[15] only recorded a single run per
neuron population, which complicates the modelling we must perform. We hypothesize
that the additional variation across runs in Yang[16] may be due to changes in the resting
cognitive state of the rats. For the data analyzed in this work, it was not possible for us
to consistently distinguish between variance due to noisy measurements and variance
due to rat behavioral changes. It may be necessary to perform many control runs to
characterize in adequate detail all of the resting cognitive states of a specific rat.
Crucially, recording a single control run as Yang[16] did for the seven other rats (B, K, Q,
R, S, T, X) probably fails to detect within-rat variations in the resting state. To
complement the recording of many control runs, recording additional rat behaviors may
yield insight as to how the rat resting cognitive state influences the network firing
dynamics.

To provide a guide for researchers investigating the connection between
neuroelectrophysiology and behavior, we asked what did each of the three studies do
right and what could be improved? There are three parameters we considered: $n_{neurons}$
(the number of neurons recorded); $n_{network\ repeat}$ (the number of times the same neural
network is recorded); and n_{rats} (the number of biological replicates). We contend that all
three of these parameters must be relatively high to make meaningful conclusions,
especially if one wishes to use more powerful parametric statistical tests over the non-
parametric tests we used in this study.

Stout[15] recorded too few neurons ($n_{neurons}$ ~1), and changed the neurons recorded
across runs ($n_{network\ repeat}$ = 1) for n_{rats} = 5. Ito[14] recorded a moderate number of neurons
($n_{neurons}$ ~20), but still changed the neurons recorded across runs ($n_{network\ repeat}$ = 1); n_{rats}
= 3. Yang[16] recorded a moderate number of neurons ($n_{neurons}$ ~20), and kept the same
neurons recorded across runs ($n_{network\ repeat}$ = 5), but only performed these replicates for
n_{rats} = 2; the other n_{rats} = 7 had ($n_{network\ repeat}$ =~1).

Recent advances provide hope that we may one day understand the relationship
between neuron firing and consciousness. Whole brain connectomes have been
mapped for simple organisms (*C. elegans*[19]*, Ciona intestinalis* larvae[20], and *Platynereis*
dumerilii[21]); and recently, all 3,000 neurons and 548,000 synapses of a fruit fly larva
(*Drosophila melanogaster*) have been mapped[22]. Although mammalian brains are
orders of magnitude larger than insect brains, whole-brain connectomes for rats may
not be out of reach in the coming decades. These connectomes are essential for

understanding cognition, but they provide only a static reference. This study highlights
the data requirements for the next steps in investigating the function and dynamics of
biological neural networks.

Methods and Materials

Dynamic time warping and statistical tests

297 Mathematical details on dynamic time warping, the Mantel test, and the Kolmogorov-
298 Smirnov test are included in the Supplemental Methods.

Estimated number of neurons in the rat medial prefrontal cortex

Our goal here is to provide a reasonable estimate for the number of neurons that one
must record to reconstruct the dynamics of the entire mPFC. We used existing
estimates of neuron and synaptic density in the rat prefrontal cortex as starting points
for our modeling. We set the total number of neurons in the rat brain[18] to $N_{brain} = 2.1 \times$
10^7 with volume[23] $V_{brain} = 2,500$ mm^3. We set the volume of the mPFC[23] to $V_{mPFC} = 20$
mm^3. Assuming a uniform distribution of neuron count throughout the brain, the
estimated number of neurons in the mPFC was $N_{mPFC} = N_{brain} \times \frac{V_{mPFC}}{V_{brain}} = 1.7 \times 10^5$.

Reproducibility of electrode recordings

Here, we discuss the problem of reproducibility in neuroelectrophysiology. Because of
the apparent lack of spatial organization in the mPFC, when electrodes are implanted,
they uniformly sample the estimated $N_{mPFC} = 50,000$ neurons in the mPFC. Depending
on how cooperative and synchronized neuron firing is, we assume that there is a certain
number of neural clusters $n_{cluster}$ for which the measurement of a single neuron
sufficiently captures the behavior of the entire cluster, and that measurement of these
clusters correlates with rat behavior. We also assumed that sampling is done with
replacement, i.e., $n_{cluster}$ << $N_{mPFC.}$ For reproducibility, one must (1) record the same
neural clusters across biological replicates and (2) identify or classify the neural clusters
to show that findings in one animal generalize to other animals.

(1) The first issue of capturing a certain percentage of the neural clusters in the rat brain
corresponds to a classical problem in combinatorial probability known as the coupon
collection problem. The mathematical question is: given $c = n_{cluster}$ categories, what is
the expected number of samples S which needs to be drawn from those categories so
that all categories are represented at least once? In the case of equal probabilities for
drawing each category,

$$S = c \cdot \sum_{i=1}^{c} \frac{1}{i} = c H_c \ ,$$

324   (1)

325 where H_c is the c^{th} harmonic number. Alternatively, we may require that a certain
326 number $1 < k < c$ out of all the categories is represented[24], so that Eq. (1) is a special
327 case of

$$S(k) = c(H_c - H_{c-k}) = c \sum_{i=0}^{k-1} \frac{1}{c-i} \ .$$

328   (2)

In both cases, $S \sim c \log c$ asymptotically. We have plotted the expected number of
clusters we may identify depending on the maximum number of neurons recorded from
each of the studies (**Figure 1**-D). In practice, it is necessary to first determine the
number of neural clusters that exist and to choose a number of neurons sampled that
can adequately sample all the clusters with high probability.

(2) The second issue of identifying neural clusters across different animals relies on the
researcher's ability to determine which of the $n_{cluster}$! possible matchings is the correct
matching between any two rats. To reduce the number of possibilities, one must rely
either on prior information (such as typical firing rates or patterns of certain clusters) or
on other measurements of network behavior that are invariant under permutations of
nodes in the graph.

Computational analysis and plotting
All computational analysis was carried out in Python 3.10 and R 4.1, on an M1 Max
MacBook Pro with 64 GB RAM. For statistical tests: Shapiro-Wilk, Kruskal-Wallis, and
ANOVA were carried out using pingouin 0.5.3[25], Mantel tests[26,27] were carried out using
scikit-bio 0.5.7[26], and Kolmogorov-Smirnov and Boschloo's exact tests were carried out
using SciPy 1.8.1[28]. We implemented dynamic time warping[29] using Cython 0.29.32[30].
We used the DataFrame structure from Pandas 1.5.1[31] to organize our data, and we
used seaborn 0.12.1[32] and Matplotlib 3.5.2[33] for plotting. We used NumPy 1.23.3[34,35] for
numerical array operations. We used ggpubr 0.6.0[36] for the balloon plots.

Acknowledgements

We thank Hiroshi T. Ito, May-Britt Moser, John J. Stout, and Amy L. Griffin for supplying
rat neuroelectrophysiology recordings and precise experimental details for our data
analysis.
354
355   We acknowledge support from the National Institutes of Health (NIH) 1R35 GM134864
356   (N.V.D.), 1RF1 AG071675 (Y.Y., N.V.D., R.B.M.), 1R01 AT012053 (N.V.D.), the
357   National Science Foundation 2210963 (N.V.D.), and the Passan Foundation (N.V.D.).
358

Author contributions (CRediT statement)

C.M.S.: Conceptualization, methodology, software, validation, formal analysis,
investigation, data curation, writing – original draft, writing – review & editing,
visualization.
J.W.: Conceptualization, methodology, software, formal analysis, writing – review &
editing.

R.B.M.: Conceptualization, methodology, validation, resources, writing – review &
editing, supervision, project administration, funding acquisition.
Y.Y.: Conceptualization, methodology, validation, investigation, resources, data curation,
writing – review & editing, supervision, project administration, funding acquisition.
N.V.D.: Conceptualization, methodology, software, validation, formal analysis,
resources, writing – review & editing, supervision, project administration, funding
acquisition.

Competing interests

We confirm that we have no financial or other interests related to this work which could
affect or have the perception of affecting our objectivity or the content of this article.
375

376　**Code and Data Availability**

377　All code and data which are necessary to reproduce the results in this work are posted
378　on Zenodo (https://doi.org/10.5281/zenodo.7933346).

Tables

Table 1: Characteristics of the three studies analyzed in this work.

	Ito[14]	Stout[15]	Yang[16]
Year of publication	2018	2020	2022
Number of rats	3	5	9
Total number of runs (including runs in which only one neuron was recorded)	13	45	13
Runs excluded (runs in which only a single neuron was recorded)	0	9	0
Number of events (excluding single-neuron recordings)	790	1508	472
Mean neurons recorded (s.d.)	17.56 (4.24)	3.31 (2.21)	31.8 (28.55)
Range of neurons recorded	[10, 26]	[1, 10]	[5, 76]
Fixed electrode	No	No	Yes
Spatial metadata	Yes	Yes	Yes

Table 2: Listing of categorical variables used for stratification/grouping in our analyses.

Categorical variable	Values	Interpretation
timing	before, after	Data from before vs after the rat left the T-intersection
correctness	true, false	True if rat chose the correct (i.e. alternate) branch of the T-maze during the task
study	Ito, Stout, Yang	The study to which the rat belongs
rat	17 total animals	The rat identity
run	62 total runs (with more than 1 neuron recorded)	The rat identity and the date of recording

Figures

388 **Figure 1: Outline of experiments, our working model, and data analysis.** (a) Rats
389 complete the T-maze task, in which the reward switches locations after each completion
390 of the task. This task tests the working memory and decision-making capacity of the rat.
391 (b) Neurons in the medial prefrontal cortex are recorded using intracranial electrodes,
392 resulting in the measurement of specific clusters of neurons. (c) Dynamic time warping
393 computes an optimal alignment of neuron spike trains and local field potentials,
394 producing a single number characterizing the synchronicity of the neurons. (d)
395 Depending on the number of distinct neuron clusters in the region of interest, a certain
396 number of neurons must be sampled to produce a representative sample of all the
397 clusters; in combinatorics, this is known as the coupon collector problem. We plot the
398 number of clusters that the maximum number of neurons recorded in each of the three
399 studies can theoretically discern.

400

Figure 2: Balloon plot of neuron spike train d_{crit}, grouped by rat, correctness of T-arm choice, and timing with respect to T-arm choice. Each box represents a single rat. The x-axis represents timing (before vs after the rat visits the T-maze intersection) and the y-axis represents correctness of T-arm taking (True vs False). The size and color of each marker represents $\log_{10} d_{\text{crit}}$. For rats with multiple trials, we took the mean of $\log_{10} d_{\text{crit}}$. Rats 17913, 17914, and 18471 are from Ito[14] (2018); single letter rats are from Yang[16] (2022); the remaining five rats are from Stout[15] (2020). The most visually striking difference is in the Ito[14] rats, in which there is a stark contrast for spike trains recorded before vs after making a choice.

Figure 3: DTW distance for neuron spike trains (left) and LFPs (right) as a function of electrode distance (from Yang[16]). The median DTW distance is plotted for each electrode distance, across all runs for the indicated rat (rat A had 4 runs, rat C had 3 runs, and the remaining rats had a single run). 95% confidence intervals were estimated using 1,000 bootstrap samples. There is no consistent correlation between the DTW and electrode distance for either dataset, supporting the idea that there is no discernible spatial organization of neurons in the mPFC. Note that the DTW metric is different for the two datasets, with the distance for neuron spike trains reported as a time (ms) and the distance for LFPs reported as a voltage (V). Physical distance is reported in multiples of 0.25 mm.

References

1. Hodgkin, A. L. & Huxley, A. F. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol* **117**, 500–544 (1952).

2. Coull, J. T. Neural correlates of attention and arousal: insights from electrophysiology, functional neuroimaging and psychopharmacology. *Prog Neurobiol* **55**, 343–361 (1998).

3. Citron, F. M. M. Neural correlates of written emotion word processing: A review of recent electrophysiological and hemodynamic neuroimaging studies. *Brain Lang* **122**, 211–226 (2012).

4. Kucyi, A. & Davis, K. D. The Neural Code for Pain: From Single-Cell Electrophysiology to the Dynamic Pain Connectome. *The Neuroscientist* **23**, 397–414 (2017).

5. Hartline, H. K. Visual Receptors and Retinal Interaction. *Science (1979)* **164**, 270–278 (1969).

6. Gilbert, C. D. & Wiesel, T. N. Receptive field dynamics in adult primary visual cortex. *Nature* **356**, 150–152 (1992).

7. Miyawaki, Y. *et al.* Visual Image Reconstruction from Human Brain Activity using a Combination of Multiscale Local Image Decoders. *Neuron* **60**, 915–929 (2008).

8. Shen, G., Dwivedi, K., Majima, K., Horikawa, T. & Kamitani, Y. End-to-End Deep Image Reconstruction From Human Brain Activity. *Front Comput Neurosci* **13**, (2019).

9. Ren, Z. *et al.* Reconstructing seen image from brain activity by visually-guided cognitive representation and adversarial learning. *Neuroimage* **228**, 117602 (2021).

10. Wakita, S., Orima, T. & Motoyoshi, I. Photorealistic Reconstruction of Visual Texture From EEG Signals. *Front Comput Neurosci* **15**, (2021).

11. Teffer, K. & Semendeferi, K. Human prefrontal cortex. in 191–218 (2012). doi:10.1016/B978-0-444-53860-4.00009-X.

12. Preuss, T. M. & Wise, S. P. Evolution of prefrontal cortex. *Neuropsychopharmacology* **47**, 3–19 (2022).

13. O'Reilly, R. C. The What and How of prefrontal cortical organization. *Trends Neurosci* **33**, 355–361 (2010).

14. Ito, H. T., Moser, E. I. & Moser, M.-B. Supramammillary Nucleus Modulates Spike-Time Coordination in the Prefrontal-Thalamo-Hippocampal Circuit during Navigation. *Neuron* **99**, 576-587.e5 (2018).

15. Stout, J. J. & Griffin, A. L. Representations of On-Going Behavior and Future Actions During a Spatial Working Memory Task by a High Firing-Rate Population of Medial Prefrontal Cortex Neurons. *Front Behav Neurosci* **14**, (2020).

16. Yang, Y., Kocher, S. D., Lewis, M. M. & Mailman, R. B. Dose-Dependent Regulation on Prefrontal Neuronal Working Memory by Dopamine D1 Agonists: Evidence of Receptor Functional Selectivity-Related Mechanisms. *Front Neurosci* **16**, (2022).

17. Lo, A., Chernoff, H., Zheng, T. & Lo, S.-H. Why significant variables aren't automatically good predictors. *Proceedings of the National Academy of Sciences* **112**, 13892–13897 (2015).

18.   Korbo, L. *et al.* An efficient method for estimating the total number of neurons in rat brain cortex. *J Neurosci Methods* **31**, 93–100 (1990).
19.   Cook, S. J. *et al.* Whole-animal connectomes of both Caenorhabditis elegans sexes. *Nature* **571**, 63–71 (2019).
20.   Ryan, K., Lu, Z. & Meinertzhagen, I. A. The CNS connectome of a tadpole larva of Ciona intestinalis (L.) highlights sidedness in the brain of a chordate sibling. *Elife* **5**, (2016).
21.   Verasztó, C. *et al.* Whole-animal connectome and cell-type complement of the three-segmented Platynereis dumerilii larva. *bioRxiv* 2020.08.21.260984 (2020) doi:10.1101/2020.08.21.260984.
22.   Winding, M. *et al.* The connectome of an insect brain. *Science (1979)* **379**, (2023).
23.   Hamezah, H. S. *et al.* Volumetric changes in the aging rat brain and its impact on cognitive and locomotor functions. *Exp Gerontol* **99**, 69–79 (2017).
24.   Ferrante, M. & Frigo, N. On the expected number of different records in a random sample. *ArXiv* (2012).
25.   Vallat, R. Pingouin: statistics in Python. *J Open Source Softw* **3**, 1026 (2018).
26.   The scikit-bio development team. scikit-bio: A Bioinformatics Library for Data Scientists, Students, and Developers. Preprint at (2023).
27.   Mantel, N. The Detection of Disease Clustering and a Generalized Regression Approach. *Cancer Res* **27**, 209–220 (1967).
28.   Virtanen, P. *et al.* SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat Methods* **17**, 261–272 (2020).
29.   Vintsyuk, T. K. Speech discrimination by dynamic programming. *Cybernetics* **4**, 52–57 (1972).
30.   Behnel, S. *et al.* Cython: The Best of Both Worlds. *Comput Sci Eng* **13**, 31–39 (2011).
31.   The pandas development team. pandas-dev/pandas: Pandas. Preprint at (2022).
32.   Waskom, M. seaborn: statistical data visualization. *J Open Source Softw* **6**, (2021).
33.   Hunter, J. D. Matplotlib: A 2D Graphics Environment. *Comput Sci Eng* **9**, 90–95 (2007).
34.   Van Der Walt, S., Colbert, S. C. & Varoquaux, G. The NumPy array: A structure for efficient numerical computation. *Comput Sci Eng* **13**, 22–30 (2011).
35.   Harris, C. R. *et al.* Array programming with NumPy. *Nature* vol. 585 357–362 Preprint at https://doi.org/10.1038/s41586-020-2649-2 (2020).
36.   Kassambara, A. ggpubr: 'ggplot2' Based Publication Ready Plots.
507