

Increased perceptual reliability reduces membrane potential variability in cortical neurons

Ben von Hünenbein^{1*}, Jakob Jordan^{1,2}, Matthijs Oude Lohuis^{3,4}

Pietro Marchesi³, Umberto Olcese^{3,5}, Cyriel M.A. Pennartz^{3,5}, Walter Senn¹,

Mihai A. Petrovici¹

¹Department of Physiology, University of Bern, Bern, Switzerland

²Electrical Engineering, Yale University, New Haven, CT, United States

³Cognitive and Systems Neuroscience Group, Swammerdam Institute for Life Sciences,
University of Amsterdam, Amsterdam, The Netherlands

⁴Champalimaud Research, Champalimaud Foundation, Lisbon, Portugal

⁵Amsterdam Brain and Cognition, University of Amsterdam, Amsterdam, The Netherlands

March 13, 2024

Abstract

Uncertainty is omnipresent. While humans and other animals take uncertainty into account during decision making, it remains unclear how it is represented in cortex. To investigate the effect of stimulus reliability on uncertainty representation in cortical neurons, we analyzed single unit activity data recorded in mouse PPC, while animals performed a multisensory change detection task. We further used simulation-based inference (SBI) to infer membrane potential statistics underlying the spiking activity. Our analysis shows that stimulus changes increase spiking rate while decreasing its variability. The inferred membrane potential statistics suggest that PPC neurons decrease their membrane potential variability in response to task relevant stimuli. Furthermore, more perceptually reliable stimuli lead to a larger decrease in membrane potential variability than less reliable ones. These findings suggest that individual cortical neurons track uncertainty, providing Bayesian benefits for downstream computations.

Introduction

The brain needs to make sense of the world based on noisy, unreliable, and ambiguous information. We experience this when having a conversation in a noisy room, finding our way in the dark or trying to pick our bicycle from a crowded rack. To reduce uncertainty we automatically incorporate information from additional sources; we look at people's lips while they speak (Sumbly and Pollack, 1954), use our hands to feel our environment in the darkness (Ernst and Banks, 2002), and use our memory to recall where we left the bike. To achieve these feats, our brain must represent and compute with the relative uncertainty associated with sensory percepts and prior knowledge.

This idea is often formalized using Bayesian principles, suggesting that the brain represents and computes information using probability distributions, for example, computing a posterior probability of a given event (Jaynes, 2003). Indeed a large body of work suggests that humans and animals take stimulus- and prior uncertainty into account when performing tasks such as coincidence detection (Miyazaki et al., 2005), gaze direction perception (Tassinari et al., 2006; Landy et al., 2012) and dynamic sensorimotor tasks (Faisal and Wolpert, 2009; Fetsch et al., 2009; O'Reilly et al., 2013; Meijer et al., 2019). Moreover, maladaptive uncertainty weighting of sensory and prior information has been linked to neurological disorders such as autism spectrum disorder and schizophrenia (Lawson et al., 2014; Van de Cruys et al., 2014; Goris et al., 2018; Stevenson et al., 2014; Noel et al., 2018). Understanding how the brain represents and computes with the uncertainty linked to different sensory perceptions is key to understanding cortical computation.

The posterior parietal cortex (PPC) has been identified as a hub for integrating information from different sensory modalities. Neurons in PPC have been found to respond to visual, auditory and somatosensory

*Correspondence: ben.vonhuenbein@unibe.ch

inputs (Raposo et al., 2012; Olcese et al., 2013; Mohan et al., 2018; Nikbakht et al., 2018). Indeed, the response of multisensory neurons to cross-modal stimulation is enhanced compared to the most effective of the individual stimuli (Stein and Stanford, 2008). Moreover, the strength of this “multisensory enhancement” depends crucially on several factors such as the temporal congruency (Meijer et al., 2017) and the perceptual strength of the individual stimulus components (Ibrahim et al., 2016). Such effects are expected from a Bayesian perspective: temporally proximal cues are likely to be informative of each other, and weaker cues leave more room for gaining additional information than strong cues. However, the mechanisms underlying this integration of information at the level of single neurons, and how information from different sources are weighted are still not well understood.

Single deep-layer cortical neurons receive bottom-up (Song et al., 2017) input driven by different sensory modalities and top-down input from higher order areas (Larkum, 2013; Rindner et al., 2022). They thus make an excellent substrate for combining novel sensory information with prior expectations. A recent theory (Jordan et al., 2021) offers a novel view on the dynamics of individual cortical neurons and suggests that they may be optimally suited to perform Bayes-optimal integration. This theory proposes that conductances play an important role in neuronal computation by allowing microscopic dendritic compartments to represent distributions with local, biophysical quantities. It suggests that neurons represent likelihood functions and priors in basal and apical dendrites respectively, while the soma computes the corresponding posterior distribution (Figure 1a). Specifically, the local effective reversal potential, the potential to which the membrane potential is pulled by synaptic inputs, represents the mean, while the sum of local excitatory and inhibitory conductances represent the precision of a Gaussian likelihood in each basal compartment. Intuitively, increasing synaptic conductances, e.g., by upstream stimulation, increase the pull towards the effective reversal potential. As this effect propagates to the soma it reduces sensitivity to unspecific background input. Computationally, this can be interpreted as sharper likelihood functions resulting in sharper posteriors. Biophysically, the variability of the somatic membrane potential is proportional to the posterior uncertainty. In line with a Bayesian view of brain function, the theory predicts that since new information reduces uncertainty, it also decreases membrane potential variability; furthermore, it predicts that the magnitude of this decrease is proportional to stimulus reliability. Finding evidence for such Bayesian processing at the single neuron level represents a promising step towards understanding how the brain deals with perceptual uncertainty.

Here, we investigate neuronal responses in mouse PPC to uni- and multimodal changes of different magnitude in continuously presented auditory and visual stimuli (Figure 1b). First, we analyze how interspike interval statistics differ between pre and post stimulus change periods. Next, we infer and analyze membrane potential statistics underlying these spiking responses using simulation-based inference (SBI). Finally, we discuss our findings in the context of Bayesian computations in single neurons, and suggest future experiments.

Methods

Experimental Data

The data used in this study was collected to investigate the causal involvement of PPC in performing an audio-visual change detection task. For complete experimental details we refer to Oude Lohuis et al. (2022). All code used is openly accessible ¹.

Animals. Animal experiments were carried out in accordance with Dutch national legislation and institutional regulations. In total, 19 male mice were used from two transgenic mouse lines; PVcre (JAX 008069) and ai9-TdTomato cre-reported mice (JAX 007909). Mice were housed under a reversed day-night schedule (lights were on from 20:00PM until 8:00AM). As mice are nocturnal, experiments were performed during the dark period. Animals were aged at least 8 weeks at the beginning of the experiments.

Head bar implantation. Animals were implanted with a head bar prior to experimentation. Animals were anesthetized using isoflurane (3% for induction, 1.5 – 2% for maintenance) and fixed in a stereotaxic apparatus. After exposing the animals skull, a circular titanium head bar was glued and cemented to it. Mice were given 2 – 7 days to recover following surgery. Prior to the training procedure mice were habituated to the head fixation and the handling of the experimenter. A subset of animals underwent additional surgery to virally infect individual areas for opsin expression for optogenetic manipulations. However, data collected during sessions with optogenetic manipulations were not used in this study.

Audiovisual change detection task. Mice were deprived of water and earned their daily ration by performing the behavioral task. Lick spouts were positioned within tongue reach symmetrically on the left and right side of the head-fixed animals. Auditory and visual stimuli were presented continuously throughout each behavioral session (Figure 1b).

¹<https://github.com/unibe-cns/reliable-membranes>




Figure 1: Perceptual reliability-weighted integration of sensory information in mouse PPC. (a) Conductance-based Bayes-optimal cue integration. According to Jordan et al. (2021), individual cortical neurons integrate novel (bottom-up) input with prior information (top-down) in a Bayes-optimal way at the level of somatic membrane potentials. Traces illustrate simulated membrane potentials of each compartment. The curves (right) represent the likelihood (basal dendrite, orange), prior (apical dendrite, grey), and posterior distribution (soma, green) pre- and post-stimulus. The presentation of novel information at 0 ms causes sharpening of the likelihood which reduces uncertainty in the posterior. (b) Behavioral task. A visual (drifting grating) and auditory (Shepard tone) stimulus were presented continuously. Mice learned to report changes in either grating orientation (“V” trial, lick right), tone frequency (“A” trial, lick left) or both (“AV” trial, lick either side). Stimulus changes occurred either at maximum or threshold magnitude. (c) Single-unit activity was measured in mouse PPC. The image represents a reference section (Paxinos and Franklin, 2019) overlaid on the histologic verification with DAPI staining (blue) and electrode tract stained with DiI (red). LPtA, Lateral parietal association area; V2L lateral secondary visual cortex.

The visual stimulus consisted of a drifting square-wave grating with a temporal frequency of 1.5 Hz and a spatial frequency of 0.08 cycles per degree at 70% contrast. The auditory stimulus was a stationary Shepard tone (Shepard, 1964), consisting of a center tone combined with its two lower and higher harmonics. Center tones used throughout the experiment ranged a one octave from 2^{13} Hz (8,372 Hz) to 2^{14} Hz (16,744 Hz). The weight of the center and harmonic tones were taken from a Gaussian distribution over weights, here centered at $2^{13.5}$ Hz (11,585 Hz). Tones were high-pass filtered (Beyma F100, Crossover frequency of 5 – 7 kHz) and presented at a sampling rate of 192 kHz using two bullet tweeters (300 W) positioned directly below the screen. The sound pressure level was calibrated at the position of the mouse, and the volume was adjusted per mouse to the minimum volume that maximized performance (on average around 70 dB).

In visual change trials (“V trials”), the orientation of the drifting grating was instantaneously changed while preserving the phase. Likewise, in auditory change trials (“A trials”) the stimulus was instantaneously changed from one Shepard tone to another, with different center frequency and associated harmonics. Animals were trained to respond to changes in a lateralised manner (A trial: lick left, V trial: lick right). Modality-side pairing was counterbalanced across mice. Thus, to successfully perform the task, the animal had to monitor both auditory and visual modalities to detect changes in either.

During training, trial types were ordered pseudo-randomly by block-shuffling; consecutive blocks of 10 trials each were comprised of randomly ordered trial types in a fixed proportion (8% catch, 46% V trials, 46% A trials). During testing, audio-visual change trials (“AV trials”) were introduced, where both auditory and visual stimuli changed (8% catch, 33.5% A trials, 33.5% V trials, 25% AV trials). AV trials were rewarded for the first lick on either side. Thus, in contrast to A and V trials, cues in AV trials were not informative about the response they should elicit.

Additionally, to manipulate their perceptual reliability, changes in either of the two modalities were delivered either at perceptual threshold or at maximum magnitude. The perceptual threshold for each modality was determined for each animal separately by fitting a psychometric function over the accuracy achieved across five levels of stimulus change magnitude. In trials with an auditory change, the center pitch changed by $1/10^{\text{th}}$ of an octave in threshold and a full octave on in maximum change trials. In trials with a visual change the movement orientation changed by 4 – 7 degrees in threshold and 90 degrees in maximum change trials. Consecutive trials were separated by inter-trial intervals sampled randomly from an exponential distribution with a mean of 6 s, a minimum of 2 s and maximum of 20 s. Following a stimulus change, mice were rewarded if the first lick after 100 ms and before 1500 ms was made to the correct side. Sessions were terminated after 20 trials of unresponsiveness, with these last 20 trials being discarded from all analyses. Mice achieved a median reaction time of 324 ms across auditory and 407 ms across visual hits.

Neural recordings. Mice were anesthetized (isoflurane, induction at 3%, maintenance at 1.5–2%) and small craniotomies (approx. 300 – 500 μm) over the areas of interest were made using a dental drill. Areas of interest were identified based on stereotaxic coordinates (V1: AP 0.0, ML ± 3.0 , PPC: AP 1.9, ML ± 1.6 , AC: AP 2.6, ML ± 4.3 , relative to lambda and bregma, respectively) (Goard et al., 2016; Song

et al., 2017; Le Merre et al., 2018). Extracellular recordings were performed on consecutive days with a maximum of four days to minimize cortical tissue damage. Microelectrode arrays (Figure 1c) with 32 or 64 channels (NeuroNexus, A1x32-Poly2-10 mm-50s-177, A1x64-Poly2-6 mm-23s-160) were slowly inserted either perpendicularly to the cortical surface (V1), or at an angle of approximately 30 degrees away from the midline (AC) After insertion, the probe was left in place for at least 15 min before recording to allow the tissue to stabilize. To identify the exact probe location post hoc, the probe was covered in Dil (Thermo Fisher Scientific) on the final day of recording. Neurophysiological signals were pre-amplified, bandpass-filtered (0.1 Hz to 9 kHz) and acquired continuously at 32 kHz using a Digital Lynx 128-channel system (Neuralynx). Post experiment, mice received an overdose of pentobarbital, were transcardially perfused (4% PFA in PBS) and their brains were extracted.

Neurophysiological data processing. Spike sorting was performed using Klusta software and manually curated using the Phy GUI (Rossant et al., 2016). Prior to spike sorting, common noise artifacts were eliminated by subtracting the median of the raw trace of surrounding channels (within 400 nm). Each potential single unit was inspected based on its waveform, autocorrelation function, and firing pattern across channels and time. Single units were only included if they met the following criteria: (1) isolation distance > 10 (Schmitzer-Torbert et al., 2005); (2) 0.1% of their spikes occurred inside the 1.5 ms refractory period (Vinck et al., 2016; Bos et al., 2017); and (3) stable presence throughout the session. To assess stability the session-wise activity was binned and single units were considered stable if they fired in at least 90% of bins.

Pre-processing. We split the trial-wise activity into a pre- and post-stimulus period. The pre-stimulus period spanned the 1500 ms before the stimulus change, while the post-stimulus period covered 200 ms to 1700 ms after the change. Due to transient responses in the initial 200 ms post-stimulus, we excluded this interval to better capture sustained within-trial statistics following stimulus onset.

Accordingly, we summarised the trial-wise spiking activity in each period by computing the mean and standard deviation of the inter-spike interval statistics (ISI, μ_s, σ_s). Note that this summary statistic assumes that the ISIs are unimodally distributed and thus would not adequately describe bimodal distributions. We thus quantified the amount of bursting within each spike train (Chen et al., 2009) and removed spike trains with excessive bursting (see also Supplementary Figure 2). Finally, as the inference procedure is less accurate for spike trains with fewer spikes (see also section Supplementary Figure 3), we only included trials with at least 10 Hz in both pre- and post-stimulus change periods in our analysis. Note that this may bias the sample towards neuron types associated with higher firing frequencies, namely interneurons and deep cortical pyramidal neurons (Contreras, 2004). After filtering, 4170 trials collected from 83 neurons in 6 animals remained for subsequent analysis ($V = 1674$, $A = 1663$, $AV = 524$, catch = 309; non-catch trials balanced across max and thresh change magnitudes).

Simulation-Based Inference (SBI)

SBI can be used to efficiently compute an approximate posterior distribution in cases where the likelihood function is intractable, e.g., stochastic simulations. This is achieved by leveraging universal function approximators to represent distributions (Tejero-Cantero et al., 2020). A typical application of SBI proceeds as follows: First, we choose an appropriate mechanistic model and a prior distribution over its parameters (see Figure 2). Next, we simulate the model with random samples from the prior distribution and record observations. We then train a density estimator to map recorded observations to the underlying parameters. Finally, we can use our trained density estimator to compute approximate posteriors over parameters for new observations, in our case spiking data obtained from experiments. We used SBI to infer membrane potential statistics of biological neurons by inverting the dynamics of leaky-integrate and fire (LIF) neurons using Sequential Neural Posterior estimation (SNPE; Greenberg et al., 2019). We implemented the SBI workflow using the `sbi` toolbox² in Python.

Mechanistic model. We chose the leaky-integrate-and-fire (LIF) neuron model due to its simplicity, ease of simulation, and common use in neuron modelling. Its dynamics is given by

$$\tau_m \frac{dV}{dt}(t) = -(V_m(t) - E_L) + RI(t) , \quad (1)$$

with the additional threshold condition

$$\begin{aligned} \text{if } V_m(t^*) \geq \Theta &\rightarrow V(t^* + \delta t) = V_{\text{reset}} , \\ &\text{and emit spike with timestamp } t^* . \end{aligned}$$

Here, V_m is the neuron’s membrane potential, τ_m the membrane time constant, E_L the reversal potential of the leak current, R the membrane resistance, Θ the spiking threshold and V_{reset} the reset potential. Due

²www.mackelab.org/sbi/




Figure 2: **Simulation-based inference (SBI)**. SBI is a method for recovering the hidden parameters underlying novel observations. (1) A single sample (red cross) from the prior distribution over the parameters (θ) of the mechanistic model. (2) Using this sample, observations x are generated by (stochastic) simulation. (3) The parameters and observation pairs are then used to train a density estimator to learn an approximate posterior distribution $p(\theta|x)$, i.e., to invert the simulation probabilistically. (4) The trained density estimator is used to infer a posterior distribution over the underlying parameters of a novel observation (x_o). Here, SBI was used to infer the parameters $\theta = (\mu_I, \sigma_I)$ of the input current I underlying the observation generated from a single spike train $x = (\mu_s^{\log}, \sigma_s^{\log}, N_s)$.

to the spiking mechanism a hard threshold on the membrane potential with subsequent reset, the membrane potential statistics of a LIF neuron are difficult to parameterize directly (Petrovici et al., 2016). Here we instead chose to parameterize an input current with Gaussian noise

$$I(t) \sim \mathcal{N}(\mu_I, \sigma_I), \quad (2)$$

with mean μ_I and standard deviation σ_I . This current subsumes all synaptic inputs to the neuron. All network simulations were carried out with NEST³ (Gewaltig and Diesmann, 2007; de Schepper et al., 2022).

Parameter prior. The prior defines the space of parameters used to simulate the mechanistic model. It thus implicitly defines the support (bounds) of the posterior distribution. We initialized the prior as a uniform distribution over the range $[-2 \text{ pA}, 2 \text{ pA}]$ for the mean μ_I of the injected Gaussian noise current, and $[0 \text{ pA}, 30 \text{ pA}]$ for its standard deviation σ_I . The prior was then iteratively restricted by removing parts of the parameter space that yielded no or excessive firing. Here we define “excessive firing” as firing rates larger than the maximum observed in the data.

Observables. Spike trains were the common quantity produced by simulations and experiments. Here, we summarized an individual spike train using the empirical mean, standard deviation and raw number of spikes ($\mu_s^{\log}, \sigma_s^{\log}, N_s$) of the log inter-spike interval (ISI) distribution (Figure 3a,b). We chose a summary statistic because we did not want to infer parameters that reproduce the exact spike times but only the spiking statistics. We log transform ISIs because it reduces the posterior uncertainty of our inference procedure.

Training. We trained a density estimator (Lueckmann et al., 2019; Papamakarios et al., 2019) to learn the mapping from observations to current parameters. To generate training data, i.e., parameter/observation pairs, we sampled parameters from the restricted prior distribution, performed a simulation of the neuron model and computed the respective observation from the simulation output. We simulated each neuron for 2000 ms, the first 500 ms of which were discarded to allow the neuron to reach its steady state. The spiking activity in the remaining 1500 ms was recorded. We generated and trained the mixed density estimator on 150,000 parameter/observable pairs.

Validation. To validate our learned density estimator, we sampled 20,000 parameter pairs, not seen during training, from the prior. We simulated each parameter pair and recorded the observations. Next we performed the inference procedure for each simulated observation and computed the difference between the inferred and the original parameters. To determine the influence of the number of spikes on inference quality, we computed the errors separately for varying spike count. We found that the inference procedure was biased towards underestimating μ_I while overestimating σ_I at firing rates below 10 Hz (see also Supplementary Figure 3). This lower boundary of 10 Hz was used as an exclusion criterion for recorded trials.

³www.nest-simulator.org




Figure 3: **Inferring intracellular quantities from spike trains.** (a) Single PPC neuron spiking activity recorded during successive trials separated into pre- and post- stimulus change spiking activity. Colors denote trials from different conditions (see Figure 1). (b) Inter-spike interval (ISI) statistics. Each spike train is summarized by the mean and standard deviation of its log ISI distribution. (c) Inferring trans-membrane currents. The density estimator trained with SBI is used to infer an approximate posterior distribution over input current parameters from each observation (μ_s, σ_s). For subsequent analysis we collapse this distribution to its most likely values (grey cross).

Inference. To infer the parameters underlying single recorded spike trains we conditioned the posterior on a single observation. We collapse this distribution to its most likely value to obtain the parameters most likely underlying the observation (Figure 3c). In a final step we translated these current parameters to membrane potential statistics by simulating a neuron with the inferred parameters (μ_i, σ_i) and recording the membrane potential traces from which we can compute the relevant statistics (μ_u, σ_u , where u indicates that this is a quantity relating to membrane potentials).

Statistical Analysis

We used linear mixed modelling (LMM) in SPSS 29.0 to investigate the effect of change magnitude and modality both on the ISI and the inferred voltage statistics. We chose LMMs because they are robust to imbalanced designs and are able to account for the nested structure of the data (Aarts et al., 2014). LMMs describe the relationship between a response variable and multiple explanatory variables of two types; fixed effects are the independent variables (change magnitude, modality) whereas random effects are grouping variables that account for the nested structure of the data. Here, the nested structure arose due to repeated trials across conditions from the same animals. In our statistical model, we defined modality and change magnitude as our fixed effects and included animal identity as our random effect. We chose a covariance structure whereby the intercept for each animal was allowed to vary freely. We did not include the neuron identity from which trials were recorded as a nested random effect within animals as the sometimes limited number of trials recorded per neuron (> 5) meant intercepts could not be estimated reliably and the model would not converge. For our dependent measure, we computed the change both for the recorded ISI statistics ($\Delta\mu_s = \mu_s^{\text{post}} - \mu_s^{\text{pre}}, \Delta\sigma_s$) and the inferred membrane potential statistics ($\Delta\mu_u, \Delta\sigma_u$). The following statistical model was fitted for each of these dependent measures $y \in \{\Delta\mu_s, \Delta\sigma_s, \Delta\mu_u, \Delta\sigma_u\}$:

$$y_{ij} = \beta_0 + \beta_1 \text{magnitude}_{ij} + \beta_2 \text{modality}_{ij} + \beta_3 (\text{modality} \times \text{magnitude})_{ij} + \beta_{0i} + \epsilon_{ij} . \quad (3)$$

Here y_{ij} is the j^{th} observation within the i^{th} animal. The overall intercept (bias) is denoted by β_0 . The coefficients β_1, β_2 and β_3 correspond to the coefficients for magnitude, change modality and their interaction respectively. β_{0i} denotes the random intercept for each animal and ϵ_{ij} represents the residual error. For main effects we report the F-test as $F_{(df1, df2)} = F$ where $df1$ and $df2$ indicate the between- and within-group degrees of freedom respectively.

Results

ISI Statistics

We first analysed the effects of stimulus change and magnitude on the ISI statistics recorded during pre- and post-change period in mouse PPC (Figure 4a,b). Here, we expected stimulus changes to result in a decrease in within-trial spiking variability. Indeed, we found that across all conditions stimulus change led to a decrease in the mean inter-spike interval ($F_{(1,7.4)} = 31.72; p < 0.001$) and a decrease in the standard




Figure 4: **Effect of stimulus change on within trial ISI and inferred membrane potential statistics.** (a) The difference in post- and pre- stimulus mean (μ_s) and (b) standard deviation (σ_s) of the recorded inter-spike interval distributions across change modalities and for (T)hreshold and (M)ax change magnitudes. (c) The difference in post- and pre- membrane potential mean and (d) standard deviation (right) as inferred by SBI. Error bars indicates SEM and *, ** and *** indicate significance at the 0.05, 0.01 and 0.001 levels respectively.

deviation of the inter-spike interval ($F_{(1,3.9)} = 23.92$; $p = 0.008$). Further, we expected that larger change magnitudes would have a stronger effect on neural activity. While maximum change trials appeared to have a lower μ_s than threshold trials, this difference was not significant. However, max change trials showed a stronger decrease in σ_s than threshold change trials.

Further, we expected a difference between uni- and multi-modal stimulus changes as unimodal stimuli were matched for intensity. While a trend in this direction could be observed, AV trials did not differ significantly from unimodal A and V trials on the change in the mean or standard deviation of the ISI statistics. One reason for this could be that we are lacking statistical power due to the smaller number of AV trials compared to A and V trials. In correct rejection (catch) trials we did not observe a significant change in the ISI statistics between pre and post change (one sided t-test; $\Delta\mu_s$, $p = 0.282$, $\Delta\sigma_s$ $p = 0.257$).

Table 1: Effect of stimulus change on ISI statistics

	Conditions (a vs b)	Mean difference (a - b)	Standard Error	Significance	CI _{95%} [lower, upper]
$\Delta\mu_s$	Max vs Thresh	-0.830	0.695	0.233	[-2.193, -0.533]
	A vs V	0.398	0.647	1.000	[-1.151, 1.848]
	AV vs A	-1.312	0.945	0.496	[-3.576, 0.952]
	AV vs V	-0.913	0.944	1.000	[-3.174, 1.347]
$\Delta\sigma_s$	Max vs Thresh	-3.196***	0.991	0.001	[-5.140, -1.253]
	A vs V	0.502	0.922	1.000	[-1.708, 2.707]
	AV vs A	-2.914	1.354	0.094	[-6.156, 0.328]
	AV vs V	-2.412	1.351	0.223	[-5.648, 0.824]

Membrane potential statistics

Next, we analysed the membrane potential statistics underlying these changes in spiking statistics as inferred using SBI (Figure 4c,d). We expected, in line with the theoretical prediction, that a stimulus change would be associated with a decrease in the membrane potential variability. Indeed we found that stimulus change led to an increase in the mean membrane potential (i.e., a depolarization) ($F_{(1,4.3)} = 30.804$, $p = 0.003$) and a decrease in the membrane potential variability ($F_{(1,5.2)} = 30.798$, $p = 0.002$). Furthermore, we found that this effect of stimulus change associated increases in μ_u and decreases in σ_u was larger in maximum than in threshold change trials. Finally, we analysed the effect of change modality on membrane potential statistics. The comparison of uni- and multimodal change trial effects shows that AV trials led to a larger mean depolarization than A but not V trials.

That the above observed increase in firing (decrease in μ_s) is associated with a decrease in σ_u is unexpected in light of our mechanistic neuron model. This is because in the fluctuation driven regime (Kuhn et al., 2004) larger σ_u will make it more likely for the membrane potential to reach threshold, leading to an increase in firing. However, this finding aligns with the prediction that integration of information leads to decreases in

the membrane potential variability. These results support the predictions from Jordan et al. (2021) whereby individual neurons respond to more informative stimuli, i.e., larger cross-modal changes, by reducing their membrane potential variability.

Table 2: Effect of stimulus change on membrane potential statistics

	Conditions (a vs b)	Mean difference (a - b)	Standard Error	Significance	CI _{95%} [lower, upper]
$\Delta\mu_u$	Max vs Thresh	0.527*	0.244	0.031	[0.049, 1.006]
	A vs V	-0.315	0.227	0.495	[-0.859, 0.229]
	AV vs A	1.145**	0.333	0.002	[0.348, 1.943]
	AV vs V	0.830**	0.332	0.007	[0.034, 1.626]
$\Delta\sigma_u$	Max vs Thresh	-0.240*	0.118	0.042	[-0.471, -0.008]
	A vs V	0.098	0.89	0.810	[-0.115, 0.312]
	AV vs A	-0.520**	0.161	0.004	[-0.905, -0.134]
	AV vs V	-0.363	0.161	0.072	[-0.747, 0.022]

Discussion

Summary & interpretation

Here we investigated the effect of stimulus-associated reliability on the within-trial inter-spike interval and membrane potential statistics of single neurons in mouse PPC. The stimulus change associated reliability was manipulated (low vs. high) by adapting the magnitude (thresh vs. max) and modality (A, V vs. AV). We found that stimulus changes across conditions led to increased and more regular spiking. Further, our results demonstrate that the reduction in ISI variability was proportional to stimulus reliability, i.e., larger (max) stimulus changes led to a stronger reduction. However, we did not observe a multisensory enhancement (Stein and Stanford, 2008) at the level of spiking statistics. Apart from a lack of statistical power due to the imbalance in single and multimodal trials, this may be due to the stimulus components in AV trials not being mutually informative. As animals learned to respond to unimodal change stimuli in a lateralized way, multimodal stimuli may have led to ambiguity with respect to the desired response (see Methods). This stands in contrast to previous studies showing multi-modal enhancement where the integration of cross-modal stimuli was beneficial for task performance (Olcese et al., 2013; Nikbakht et al., 2018; Meijer et al., 2018). On the other hand, at the level of inferred membrane potential statistics, stimulus changes across all conditions lead to an increase in the mean and a decrease in the standard deviation of the membrane potential distribution. Moreover, this effect appeared to be sensitive to stimulus reliability as it was stronger in multi-modal and maximum change magnitude trials compared to unimodal and threshold magnitude changes respectively.

Our results suggest that neurons in PPC are sensitive to the stimulus associated perceptual uncertainty at the level of their membrane potential statistics. Previous work found a whisking-induced reduction of membrane potential variability in rat primary sensory cortex (Crochet et al., 2011; Yamashita et al., 2013; Poulet and Crochet, 2019). Although they did not investigate the effect of reliability, these findings are in line with ours, suggesting that this effect is a widespread cortical phenomenon. Secondly, we show that the effect is sensitive to the stimulus associated reliability, suggesting that individual neurons track the stimulus related uncertainty. In light of recent theoretical work (Jordan et al., 2021), this stimulus-driven reduction can be viewed as a signature of Bayesian computation at the single-neuron level. In a Bayesian framework, additional information leads to a reduction of uncertainty, explaining the post-change decrease of variability. Furthermore, more informative (e.g., more reliable or salient) stimuli lead to a larger reduction of uncertainty and thus to a larger decrease of membrane potential variability. In contrast to previous work, we thus suggest that the stimulus-driven reduction in membrane potential variability is a consequence of single-neuron computation, rather than an emergent property at the network level (Sussillo and Abbott, 2009; Deco and Hugues, 2012). Note that here we focused on within-trial variability, rather than trial-by-trial variability. While their relationship is complex, a reduction of within-trial variability can cause a reduction in trial-to-trial variability. We suggest that the well-documented reduction in stimulus-driven trial-by-trial variability (Finn et al., 2007; Churchland et al., 2010; Wright et al., 2017) could in part be attributed to the within-trial activity reductions observed here.

Probing the theory further

In the theory proposed by (Jordan et al., 2021), the stimulus-associated reliability is represented via local membrane conductances. Future work could apply SBI to a more complex neuron model including transmembrane conductances in addition to transmembrane currents. This would account for neuronal morphology by more accurately covering the effect of both distal (effectively current-based) and proximal (effectively conductance-based) synapses on the somatic membrane. Explicitly inferring a change in total input conductance following a change in the stimulus would lend further support to the concept of Bayesian processing at the single neuron level.

Furthermore, one could separately investigate the reliability-modulated contributions of presynaptic activity and synaptic coupling strength to the synaptic conductances. Whereas changes in the presynaptic rate lead to quick re-weighting across modalities (e.g., Fetsch et al., 2012), changes to the synaptic coupling reflect more long-term changes. These could be triggered by changes in reward probabilities or changes in sensory-sensory associations (Knöpfel et al., 2019). Future experiments could disentangle these components by using SBI together with a model where they are parameterized separately.

Alternatively, inferring intracellular parameters across time, rather than only their time-averaged statistics as we have done here, may allow for distinguishing between different models of biophysical computation. For example, probabilistic population codes (PPCs; Ma et al., 2006) consider a feedforward integration at the population level in which the respective efferent synaptic weights determine the relative weight of populations representing information from different modalities. Similarly, feedback normalization models (Ohshiro et al., 2011) assign effective “dominance weights” to afferents of multisensory neurons. These models thus exclusively rely on long-term plasticity to adapt to changes in relative reliability between modalities. In contrast, the Bayesian-dendrite theory considered here represents reliability via conductances. These are influenced by both synaptic weights and presynaptic rates and thus allow for rapid adaptation via presynaptic modulation of firing rates. Investigating how closely the neurally represented reliability follows rapid changes in stimulus-related reliability would allow us to differentiate between these model classes.

Proxies for reliability

In this study we manipulated the perceptual reliability of the stimuli by varying the magnitude of the change. The task-relevant cue consisted of a change in one of two continuously presented stimuli. The reason we chose stimulus change as our cue (as opposed to stimulus onset) was to manipulate the stimulus reliability separately from its perceptual intensity. Thus, while larger stimulus changes were more reliably informative towards the availability of reward, the perceptual intensity, i.e., sound volume or visual contrast, of the stimuli did not change. However, given that neurons adapt to continuous stimulation (Kohn, 2007; Ulanovsky et al., 2003) it may be that stimulus change was associated with a perceived change in intensity. It would be interesting to explore whether the effects observed here are maintained under experimental manipulations where the stimulus reliability is fully decoupled from, or even inversely related to, perceptual intensity. Based on Bayesian principles we would expect that, independent of the way reliability is modulated, increasing reliability should decrease membrane potential variability. Here we only probed two levels of change magnitude per modality. It would be interesting to investigate whether neuronal responses are sensitive to more finely graded reliabilities as Bayesian principles would suggest a monotonic decrease of variability with reliability.

Validation with ground-truth data

Here we used simulation-based inference together with simulations of a LIF model parameterized by a Gaussian current input to obtain single neuron membrane potential statistics underlying spiking activity. We validated our approach in simulation by inferring known parameters, both to assess inference quality and determine the amount of information needed to correctly infer model parameters. To further strengthen the inference procedure one would ideally compare the quantities inferred from recordings to a ground truth. This was not feasible in the current work as it would require, simultaneously recorded intra- and extracellular data which are challenging to obtain, especially over longer periods. In addition to providing a ground truth, these parameter and observation pairs recorded in vivo could be used to fine tune the density estimator used for the inference procedure using transfer learning (e.g., Finn et al., 2017). Establishing such ground truth benchmarks would allow researchers to test model predictions about computationally relevant but difficult to access quantities in a standardized way.

Leveraging SBI to test theoretical predictions

Computational theories are often formulated with normative principles and mathematical convenience in mind. Thus many theoretical predictions pertain to abstract and difficult-to-access quantities. Testing these predictions in experiments is often not feasible due to cost and time constraints. As demonstrated here, SBI provides a way to bridge this gap by fitting the relevant quantities in mechanistic neuron models (Gonçalves et al., 2020). These inferred parameters can then be interpreted to gain valuable insights about the mechanisms underlying otherwise obscure phenomena. SBI could be particularly useful for testing the implicit predictions of microcircuit models which describe cortical computation as the product of canonically arranged groups of neurons (Sacramento et al., 2018; Wilmes and Clopath, 2019; Hertäg and Clopath, 2022; Haider et al., 2021; Granier et al., 2023).

With SBI, these predictions could be more readily tested based on existing data. For example, a recent model for how cortical microcircuits may implement efficient learning by approximating the backpropagation of errors algorithm (Sacramento et al., 2018) suggests that apical compartments receive lateral and top-down information to compute local errors (Fişek et al., 2023). This prediction is challenging to test as it requires in vivo recordings from apical dendrites. However, synaptic inputs to cortical pyramidal cells form a considerable contribution to local field potentials (LFPs) (Einevoll et al., 2013). One could use SBI together with existing simulation tools for microcircuits and their associated LFPs (Skaar et al., 2020; Rimehaug et al., 2023) to infer net currents flowing into apical dendrites during task learning. If apical compartments represent errors, these currents should shrink over the course of learning. Likewise one could test the prediction by more recent work (Haider et al., 2021) that apical and somatic compartments receive strongly correlated inputs.

Conclusion

In this work we investigated how cortical neurons respond to uni- and multi-modal stimuli associated with varying perceptual reliabilities. We demonstrated that task relevant information decreases variability both on the spiking as well as the inferred membrane potential level. As predicted by a recent Bayesian model of single neuron computation we found a reduction in variability in proportion to the stimulus associated reliability. Our results shine new light on the representation of uncertainty in cortex and suggest that individual cortical neurons keep track and compute with uncertainty.

Acknowledgements

This work has received funding from the European Union 7th Framework Programme under grant agreement 604102 (HBP), the Horizon 2020 Framework Programme under grant agreements 720270, 785907 and 945539 (HBP), the Swiss National Science Foundation (Sinergia grant CRSII5-180316), the Manfred Störk Foundation, FLAG-ERA JTC projects CANON (2015) and DOMINO (2019), both cofinanced by the Netherlands Organization for Scientific Research. We also thank Sietse Kiewiet for helpful contributions to the manuscript.

References

- Aarts, E., Verhage, M., Veenliet, J. V., Dolan, C. V., and Van Der Sluis, S. (2014). A solution to dependency: using multilevel analysis to accommodate nested data. *Nature Neuroscience*, 17(4):491–496.
- Bos, J. J., Vinck, M., van Mourik-Donga, L. A., Jackson, J. C., Witter, M. P., and Pennartz, C. M. (2017). Perirhinal firing patterns are sustained across large spatial segments of the task environment. *Nature Communications*, 8(1):15602.
- Chen, L., Deng, Y., Luo, W., Wang, Z., and Zeng, S. (2009). Detection of bursts in neuronal spike trains by the mean inter-spike interval method. *Progress in Natural Science*, 19(2):229–235.
- Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., Newsome, W. T., Clark, A. M., Hosseini, P., Scott, B. B., et al. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience*, 13(3):369–378.
- Contreras, D. (2004). Electrophysiological classes of neocortical neurons. *Neural Networks*, 17(5-6):633–646.
- Crochet, S., Poulet, J. F., Kremer, Y., and Petersen, C. C. (2011). Synaptic mechanisms underlying sparse coding of active touch. *Neuron*, 69(6):1160–1175.

- de Schepper, R., Eppler, J. M., Kurth, A., Nagendra Babu, P., Deepu, R., Spreizer, S., Trench, G., Pronold, J., Vennemo, S. B., Graber, S., Morales-Gregorio, A., Linssen, C., Benelhedi, M. A., Mørk, H., Morrison, A., Terhorst, D., Mitchell, J., Diaz, S., Kitayama, I., Enan, M., Kamiji, N. L., and Plesser, H. E. (2022). NEST 3.2.
- Deco, G. and Hugues, E. (2012). Neural network mechanisms underlying stimulus driven variability reduction. *PLoS Computational Biology*, 8(3):e1002395.
- Einevoll, G. T., Kayser, C., Logothetis, N. K., and Panzeri, S. (2013). Modelling and analysis of local field potentials for studying the function of cortical circuits. *Nature Reviews Neuroscience*, 14(11):770–785.
- Ernst, M. O. and Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429–433.
- Faisal, A. A. and Wolpert, D. M. (2009). Near optimal combination of sensory and motor uncertainty in time during a naturalistic perception-action task. *Journal of Neurophysiology*, 101(4):1901–1912.
- Fetsch, C. R., Pouget, A., DeAngelis, G. C., and Angelaki, D. E. (2012). Neural correlates of reliability-based cue weighting during multisensory integration. *Nature Neuroscience*, 15(1):146–154.
- Fetsch, C. R., Turner, A. H., DeAngelis, G. C., and Angelaki, D. E. (2009). Dynamic reweighting of visual and vestibular cues during self-motion perception. *Journal of Neuroscience*, 29(49):15601–15612.
- Finn, C., Abbeel, P., and Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135. PMLR.
- Finn, I. M., Priebe, N. J., and Ferster, D. (2007). The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. *Neuron*, 54(1):137–152.
- Fişek, M., Herrmann, D., Egea-Weiss, A., Cloves, M., Bauer, L., Lee, T.-Y., Russell, L. E., and Häusser, M. (2023). Cortico-cortical feedback engages active dendrites in visual cortex. *Nature*, pages 1–8.
- Gewaltig, M.-O. and Diesmann, M. (2007). NEST (NEural Simulation Tool). *Scholarpedia*, 2(4):1430.
- Goard, M. J., Pho, G. N., Woodson, J., and Sur, M. (2016). Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. *eLife*, 5:e13764.
- Gonçalves, P. J., Lueckmann, J.-M., Deistler, M., Nonnenmacher, M., Öcal, K., Bassetto, G., Chintaluri, C., Podlaski, W. F., Haddad, S. A., Vogels, T. P., et al. (2020). Training deep neural density estimators to identify mechanistic models of neural dynamics. *eLife*, 9:e56261.
- Goris, J., Braem, S., Nijhof, A. D., Rigoni, D., Deschrijver, E., Van de Cruys, S., Wiersema, J. R., and Brass, M. (2018). Sensory prediction errors are less modulated by global context in autism spectrum disorder. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 3(8):667–674.
- Granier, A., Petrovici, M. A., Senn, W., and Wilmes, K. A. (2023). Precision estimation and second-order prediction errors in cortical circuits. *arXiv preprint arXiv:2309.16046*.
- Greenberg, D., Nonnenmacher, M., and Macke, J. (2019). Automatic posterior transformation for likelihood-free inference. In *International Conference on Machine Learning*, pages 2404–2414. PMLR.
- Haider, P., Ellenberger, B., Kriener, L., Jordan, J., Senn, W., and Petrovici, M. A. (2021). Latent equilibrium: A unified learning theory for arbitrarily fast computation with arbitrarily slow neurons. *Advances in Neural Information Processing Systems*, 34:17839–17851.
- Hertäg, L. and Clopath, C. (2022). Prediction-error neurons in circuits with multiple neuron types: Formation, refinement, and functional implications. *Proceedings of the National Academy of Sciences*, 119(13):e2115699119.
- Ibrahim, L. A., Mesik, L., Ji, X.-y., Fang, Q., Li, H.-f., Li, Y.-t., Zingg, B., Zhang, L. I., and Tao, H. W. (2016). Cross-modality sharpening of visual cortical processing through layer-1-mediated inhibition and disinhibition. *Neuron*, 89(5):1031–1045.
- Jaynes, E. T. (2003). *Probability theory: The logic of science*. Cambridge University Press.
- Jordan, J., Sacramento, J., Wybo, W. A., Petrovici, M. A., and Senn, W. (2021). Learning Bayes-optimal dendritic opinion pooling. *arXiv preprint arXiv:2104.13238*.

- Knöpfel, T., Sweeney, Y., Radulescu, C. I., Zabouri, N., Doostdar, N., Clopath, C., and Barnes, S. J. (2019). Audio-visual experience strengthens multisensory assemblies in adult mouse visual cortex. *Nature Communications*, 10(1):5684.
- Kohn, A. (2007). Visual adaptation: Physiology, mechanisms, and functional benefits. *Journal of Neurophysiology*, 97(5):3155–3164.
- Kuhn, A., Aertsen, A., and Rotter, S. (2004). Neuronal integration of synaptic input in the fluctuation-driven regime. *Journal of Neuroscience*, 24(10):2345–2356.
- Landy, M. S., Trommershäuser, J., and Daw, N. D. (2012). Dynamic estimation of task-relevant variance in movement under risk. *Journal of Neuroscience*, 32(37):12702–12711.
- Larkum, M. (2013). A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. *Trends in Neurosciences*, 36(3):141–151.
- Lawson, R. P., Rees, G., and Friston, K. J. (2014). An aberrant precision account of autism. *Frontiers in Human Neuroscience*, 8:302.
- Le Merre, P., Esmaeili, V., Charrière, E., Galan, K., Salin, P.-A., Petersen, C. C., and Crochet, S. (2018). Reward-based learning drives rapid sensory signals in medial prefrontal cortex and dorsal hippocampus necessary for goal-directed behavior. *Neuron*, 97(1):83–91.
- Lueckmann, J.-M., Bassetto, G., Karaletsos, T., and Macke, J. H. (2019). Likelihood-free inference with emulator networks. In *Symposium on Advances in Approximate Bayesian Inference*, pages 32–53. PMLR.
- Ma, W. J., Beck, J. M., Latham, P. E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9(11):1432–1438.
- Meijer, G. T., Mertens, P. E., Pennartz, C. M., Olcese, U., and Lansink, C. S. (2019). The circuit architecture of cortical multisensory processing: Distinct functions jointly operating within a common anatomical network. *Progress in Neurobiology*, 174:1–15.
- Meijer, G. T., Montijn, J. S., Pennartz, C. M., and Lansink, C. S. (2017). Audiovisual modulation in mouse primary visual cortex depends on cross-modal stimulus configuration and congruency. *Journal of Neuroscience*, 37(36):8783–8796.
- Meijer, G. T., Pie, J. L., Dolman, T. L., Pennartz, C. M., and Lansink, C. S. (2018). Audiovisual integration enhances stimulus detection performance in mice. *Frontiers in Behavioral Neuroscience*, 12:231.
- Miyazaki, M., Nozaki, D., and Nakajima, Y. (2005). Testing Bayesian models of human coincidence timing. *Journal of Neurophysiology*, 94(1):395–399.
- Mohan, H., Gallero-Salas, Y., Carta, S., Sacramento, J., Laurenczy, B., Sumanovski, L. T., De Kock, C. P., Helmchen, F., and Sachidhanandam, S. (2018). Sensory representation of an auditory cued tactile stimulus in the posterior parietal cortex of the mouse. *Scientific Reports*, 8(1):7739.
- Nikbakht, N., Tafreshiha, A., Zoccolan, D., and Diamond, M. E. (2018). Supralinear and supramodal integration of visual and tactile signals in rats: psychophysics and neuronal mechanisms. *Neuron*, 97(3):626–639.
- Noel, J.-P., Stevenson, R. A., and Wallace, M. T. (2018). Atypical audiovisual temporal function in autism and schizophrenia: similar phenotype, different cause. *European Journal of Neuroscience*, 47(10):1230–1241.
- Ohshiro, T., Angelaki, D. E., and DeAngelis, G. C. (2011). A normalization model of multisensory integration. *Nature Neuroscience*, 14(6):775–782.
- Olcese, U., Iurilli, G., and Medini, P. (2013). Cellular and synaptic architecture of multisensory integration in the mouse neocortex. *Neuron*, 79(3):579–593.
- O’Reilly, J. X., Jbabdi, S., Rushworth, M. F., and Behrens, T. E. (2013). Brain systems for probabilistic and dynamic prediction: computational specificity and integration. *PLOS Biology*, 11(9):e1001662.
- Oude Lohuis, M. N., Marchesi, P., Pennartz, C. M., and Olcese, U. (2022). Functional (ir) relevance of posterior parietal cortex during audiovisual change detection. *Journal of Neuroscience*, 42(26):5229–5245.

- Papamakarios, G., Sterratt, D., and Murray, I. (2019). Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 837–848. PMLR.
- Paxinos, G. and Franklin, K. B. (2019). *Paxinos and Franklin’s the mouse brain in stereotaxic coordinates*. Academic Press.
- Petrovici, M. A., Bill, J., Bytschok, I., Schemmel, J., and Meier, K. (2016). Stochastic inference with spiking neurons in the high-conductance state. *Physical Review E*, 94(4):042312.
- Poulet, J. F. and Crochet, S. (2019). The cortical states of wakefulness. *Frontiers in systems neuroscience*, 12:64.
- Raposo, D., Sheppard, J. P., Schrater, P. R., and Churchland, A. K. (2012). Multisensory decision-making in rats and humans. *Journal of Neuroscience*, 32(11):3726–3735.
- Rimehaug, A. E., Stasik, A. J., Hagen, E., Billeh, Y. N., Siegle, J. H., Dai, K., Olsen, S. R., Koch, C., Einevoll, G. T., and Arkhipov, A. (2023). Uncovering circuit mechanisms of current sinks and sources with biophysical simulations of primary visual cortex. *eLife*, 12:e87169.
- Rindner, D. J., Proddutur, A., and Lur, G. (2022). Cell-type-specific integration of feedforward and feedback synaptic inputs in the posterior parietal cortex. *Neuron*, 110(22):3760–3773.
- Rossant, C., Kadir, S. N., Goodman, D. F., Schulman, J., Hunter, M. L., Saleem, A. B., Grosmark, A., Belluscio, M., Denfield, G. H., Ecker, A. S., et al. (2016). Spike sorting for large, dense electrode arrays. *Nature Neuroscience*, 19(4):634–641.
- Sacramento, J., Ponte Costa, R., Bengio, Y., and Senn, W. (2018). Dendritic cortical microcircuits approximate the backpropagation algorithm. *Advances in Neural Information Processing Systems*, 31.
- Schmitzer-Torbert, N. ., Jackson, J., Henze, D., Harris, K., and Redish, A. (2005). Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience*, 131(1):1–11.
- Shepard, R. N. (1964). Circularity in judgments of relative pitch. *The Journal of the Acoustical Society of America*, 36(12):2346–2353.
- Skaar, J.-E. W., Stasik, A. J., Hagen, E., Ness, T. V., and Einevoll, G. T. (2020). Estimation of neural network model parameters from local field potentials (LFPs). *PLOS Computational Biology*, 16(3):e1007725.
- Song, Y.-H., Kim, J.-H., Jeong, H.-W., Choi, I., Jeong, D., Kim, K., and Lee, S.-H. (2017). A neural circuit for auditory dominance over visual perception. *Neuron*, 93(4):940–954.
- Stein, B. E. and Stanford, T. R. (2008). Multisensory integration: current issues from the perspective of the single neuron. *Nature Reviews Neuroscience*, 9(4):255–266.
- Stevenson, R. A., Siemann, J. K., Woynaroski, T. G., Schneider, B. C., Eberly, H. E., Camarata, S. M., and Wallace, M. T. (2014). Evidence for diminished multisensory integration in autism spectrum disorders. *Journal of Autism and Developmental Disorders*, 44:3161–3167.
- Sumby, W. H. and Pollack, I. (1954). Visual contribution to speech intelligibility in noise. *The Journal of the Acoustical Society of America*, 26(2):212–215.
- Sussillo, D. and Abbott, L. F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63(4):544–557.
- Tassinari, H., Hudson, T. E., and Landy, M. S. (2006). Combining priors and noisy visual cues in a rapid pointing task. *Journal of Neuroscience*, 26(40):10154–10163.
- Tejero-Cantero, A., Boelts, J., Deistler, M., Lueckmann, J.-M., Durkan, C., Gonçalves, P. J., Greenberg, D. S., and Macke, J. H. (2020). Sbi—a toolkit for simulation-based inference. *arXiv preprint arXiv:2007.09114*.
- Ulanovsky, N., Las, L., and Nelken, I. (2003). Processing of low-probability sounds by cortical neurons. *Nature Neuroscience*, 6(4):391–398.
- Van de Cruys, S., Evers, K., Van der Hallen, R., Van Eylen, L., Boets, B., De-Wit, L., and Wagemans, J. (2014). Precise minds in uncertain worlds: predictive coding in autism. *Psychological Review*, 121(4):649.

- Vinck, M., Bos, J. J., Van Mourik-Donga, L. A., Oplaat, K. T., Klein, G. A., Jackson, J. C., Gentet, L. J., and Pennartz, C. M. (2016). Cell-type and state-dependent synchronization among rodent somatosensory, visual, perirhinal cortex, and hippocampus ca1. *Frontiers in Systems Neuroscience*, 9:187.
- Wilmes, K. A. and Clopath, C. (2019). Inhibitory microcircuits for top-down plasticity of sensory representations. *Nature Communications*, 10(1):5055.
- Wright, N. C., Hoseini, M. S., and Wessel, R. (2017). Adaptation modulates correlated subthreshold response variability in visual cortex. *Journal of Neurophysiology*, 118(2):1257–1269.
- Yamashita, T., Pala, A., Pedrido, L., Kremer, Y., Welker, E., and Petersen, C. C. (2013). Membrane potential dynamics of neocortical projection neurons driving target-specific signals. *Neuron*, 80(6):1477–1490.