REMOVING NEURAL CORRELATIONS IMPROVES POPULATION SENSITIVITY IN MT/MST IN RESPONSE TO RANDOM DOT STIMULI

Tristan A. Chaplin[1,2], Maureen A. Hagan[1,2], Benjamin J. Allitt[1], Marcello G.P. Rosa[1,2], Leo L. Lui[1,2]

1- Neuroscience Program, Biomedicine Discovery Institute and Department of Physiology, Monash University, Clayton, VIC 3800, Australia

2- ARC Centre of Excellence for Integrative Brain Function, Monash University Node, VIC 3800, Australia

Abbreviated title: Removing correlations improves population sensitivity

Correspondence: Leo L. Lui,

Department of Physiology, Monash University

Clayton, VIC 3800, Australia.

E-mail: Leo.Lui@monash.edu

Manuscript length: 32 pages

Key words: marmoset, population decoding, motion, random-dot, multi-electrode arrays.

Word count:

Abstract: 250

Introduction: 639

Discussion: 1445

Abstract

The effect of temporal correlations in the activity of sensory neurons remains a subject of debate, with some studies suggesting that correlations may be detrimental to population coding (by reducing the amount of information that can be extracted) or may actually enhance population coding. The study of neuronal responses to random-dot motion patterns has provided the some of the most valuable insights into how the activity of neurons is related to perception. However, it is currently unknown how changes in the strength of motion signals, through the reduction of coherence of random dot patterns, affects correlated activity and population decoding. To address this question, we recorded neuronal populations in the middle temporal (MT) and medial superior temporal (MST) areas of anaesthetized marmosets with electrode arrays, while varying the coherence of random dot patterns. We used the spike rates of simultaneously recorded neurons to decode the direction of motion at each level of coherence using linear classifiers. We found that reducing motion coherence increased neuronal correlations, yet the correlation structure was conserved. We also found that removing correlations with trial shuffling generally improved population decoding performance and ignoring correlations generally impaired decoding performance. Finally, we showed that decoders trained at 100% coherence have similar performance to decoders optimized for each level of coherence, demonstrating that the optimal linear readout is independent of coherence. These results have implications for how information is encoded by populations of neurons, as well as how they may be decoded by downstream areas in decision making tasks.

Introduction

Perception arises from the activity of neurons, and understanding the way in which stimulus features are represented by the activity of neurons is one of the key challenges in system neuroscience. One of the most effective paradigms for addressing this question has been the decoding the direction of visual motion from the activity of single neurons in the middle temporal area (MT) of the primate cerebral cortex. Decreasing the strength of the motion signal decreases both the behavioural performance and the amount of information neurons carry regarding direction in an opposite directions discrimination task (Newsome et al., 1989; Britten et al., 1992, 1996). While it was initially found that the activity of single neurons may account for behavioural performance (Newsome et al., 1989; Britten et al., 1992), it has since become clear that the activity of a pool of neurons must combined to form the perceptual decision (Britten et al., 1996; Shadlen et al., 1996; Law and Gold, 2008; Cohen and Newsome, 2009).

Combining the activity of a population of neurons can reduce the effects of single neuron variability (Tolhurst et al., 1983) and increase the reliability of the of the neural signal. However, combining the activity of single neurons recorded on different trials, even in response to the exact same stimulus, can only approximate the population responses, since it does not capture the trial to trial correlations of pairs of neurons (Zohary et al., 1994; Bair et al., 2001). Correlated activity has the potential to change the amount of information contained in populations of neurons. Initially, it was believed that these correlations would impair population decoding performance (Zohary et al., 1994), but it was found later that certain correlation structures can actually improve population decoding (Abbott and Dayan, 1999; Sompolinsky et al., 2001; Averbeck et al., 2006; Shamir and Sompolinsky, 2006; Graf et al., 2011; Zylberberg et al., 2016). Furthermore, it has recently been shown with modelling that neurons that are not informative individually, but are correlated with activity of informative neurons, could still contribute to improvements in population decoding (Goris et al., 2014; Zylberberg, 2017). Therefore, one has to consider population activity recorded concurrently on the same set of trials to fully account for the information in a neuronal population. In this study, we addressed the effects of correlations on population decoding in MT using random dot motion.

3

Previous studies have modelled large neuronal populations by assuming spiking correlation structures using data obtained by simultaneously recording pairs of neurons (Cohen and Newsome, 2009; Law and Gold, 2009), but these assumptions have not been fully substantiated in physiological experiments. For example, because decreasing motion signal strength causes changes firing rates (Britten et al., 1992, 1993; Chaplin et al., 2017a), and possibly correlations, decoding could require a different strategy depending on the strength of the motion signal. Alternatively, the optimal decoding method for high signal strength (supra-threshold) stimuli may be the same as the optimal decoding method for low signal strength (near threshold) stimuli, but this is yet to be tested with real neuronal population data.

We decoded population activity in area MT and the medial superior temporal area (MST, a higher order motion processing area: Van Essen et al., 1981; Tanaka et al., 1986; Celebrini and Newsome, 1994) for opposite directions while manipulating the strength of the motion signal through changes in the coherence of random dot patterns. We found that MT/MST cells were weakly correlated, and correlations increased as motion strength decreased for non-zero motion strengths. Correlations were detrimental to population decoding, however, using a decoder that accounted for the correlation structure mitigated some of the negative effects of correlated activity on population coding. Despite the changes in correlation with respect to coherence, the same decoding strategy can be generalized across all coherences.

Methods

Animals and surgical preparation

Single unit and multi-unit extracellular recordings in areas MT and MST were obtained from 5 marmoset monkeys; 2 males and 3 females, between 1.5 and 3 years of age, with no history of veterinary complications. These animals were also used for unrelated anatomical tracing and visual physiology experiments. Experiments were conducted in accordance with the Australian Code of Practice for the Care and Use of Animals for Scientific Purposes, and all procedures were approved by the Monash University Animal Ethics Experimentation Committee.

The preparation for electrophysiology studies of marmosets has been described previously (Bourne and Rosa, 2003; updated as in Yu and Rosa, 2010). Anesthesia was induced with alfaxalone (Alfaxan, 8 mg/kg), allowing a tracheotomy, vein cannulation and craniotomy to be performed. After all surgical procedures were completed, the animal was administered an intravenous infusion of pancuronium bromide (0.1 mg/kg/h) combined with sufentanil (6-8 μg/kg/h, adjusted to ensure no physiological responses to noxious stimuli) and dexamethasone (0.4 mg/kg/h), and was artificially ventilated with a gaseous mixture of nitrous oxide and oxygen (7:3). The electrocardiogram and level of cortical spontaneous activity were continuously monitored. Administration of atropine (1%) and phenylephrine hydrochloride (10%) eye drops was used to produce mydriasis and cycloplegia. Appropriate focus and protection of the corneas from desiccation were achieved by means of hard contact lenses selected by retinoscopy.

Electrophysiology, data acquisition and pre-processing

We recorded neural activity with single shaft linear arrays (NeuroNexus) consisting of 32 electrodes separated by 50 μm. MT and MST recording sites were identified during experiments using anatomical landmarks, receptive field progression and size (Rosa and Elston, 1998), and direction selectivity. The position of recording sites were confirmed post-mortem by histological examination.

Electrophysiological data were recorded using a Cereplex system (Blackrock Microsystems) with a sampling rate of 30 kHz. For offline analysis of spiking activity, each channel was high-pass filtered at 750 Hz and spikes were initially identified based on threshold crossings. Units were sorted using Offline Sorter (Plexon Inc.). Units were classified as single-units if they showed good separation on the (2 component) principal component analysis plot, and were confirmed by inspection of the inter-spike interval histogram and consistency of waveform over time. Any remaining threshold crossings were classified as multi-unit activity. We excluded five single units from adjacent channels since it was apparent they were duplicated across two channels, based on their sharp cross correlogram peak and high signal correlation (Bair et al., 2001).

Visual stimuli

Visual stimuli were presented on a VIEWPixx3D monitor (1920 x 1080 pixels; 520 x 295 mm; 120 Hz refresh rate, VPixx Technologies) positioned 0.35 to 0.45 m from the animal on an angle to accommodate the size and eccentricity of the receptive field(s), typically subtending 70° in azimuth, and 40° in elevation. All stimuli were generated with MATLAB using Psychtoolbox-3 (Brainard, 1997).

The main visual stimulus consisted of random dots presented full screen. White dots (106 cd/m^2) of $0.2°$ in diameter were displayed on a black (0.25 cd/m^2) background (full contrast). The density was such that there were on average 0.5 dots per $°^2$; this was chosen because these parameters have been shown to elicit good responses from marmoset MT when displayed on LCD monitors (Solomon et al., 2011; Zavitz et al., 2016). Dot coherence was controlled using the white noise method (i.e. Britten et al., 1992, 1996; see Pilly and Seitz 2009) by randomly choosing a subset of "noise" dots on each frame, which were displaced to random positions within the stimulus aperture. The remaining "signal" dots were moved in the same direction with a fixed displacement.

Determination of receptive fields and basic direction tuning: Visual receptive fields were quantitatively mapped using a grid of either static flashed squares or small apertures of briefly presented moving dots. Subsequent visual stimuli were presented full screen, so as to cover as many

6

neurons' receptive fields as possible. We also conducted visual direction tuning tests (12 directions, 100% coherence), which aided in identifying the location of MT and MST. Direction selectivity was determined using a Rayleigh test (p < 0.05) (Berens, 2009). The preferred direction was calculated with a vector sum (Ringach et al., 2002), and the strength of the direction selectivity was calculated with a direction index DI:

$$DI = 1 - \frac{rate_{null}}{rate_{pref}} \quad (1)$$

Where $rate_{pref}$ and $rate_{null}$ is the firing rate in response to preferred and null direction of motion respectively, at 100% motion coherence. DI values lie between 0 and 1, with 1 indicating a strongly direction selective neuron.

Stimulus protocol: We presented visual stimuli moving either leftwards or rightwards at 60°/s at different levels of motion coherence: 100, 82, 64, 46, 28, 10 and 0%. All stimuli were presented for 600 ms with 120 repeats per condition These data were obtained as part of a study to investigate the effects of auditory motion in MT and MST. Since we did find any effect of auditory stimuli in the responses of single neurons or populations of neurons (Chaplin et al., 2017b), we have grouped the two conditions (visual and audio-visual) into one dataset for these analyses.

Data Analysis

Time windows and inclusion criteria: Firing rates were calculated using a time window starting 10 ms after stimulus onset (to avoid a potential noise artifact caused by the speakers in the audio-visual condition, Chaplin et al., 2017) and finishing at the stimulus offset (600 ms). Units were deemed responsive if they passed a Wilcoxon rank sum test (p < 0.01) against the spontaneous firing rate, and if the firing rate in response to the stimulus that elicited the highest spike rate was at least 2 spikes/s above the spontaneous rate. Units were considered left-right selective if the firing rate to the best direction of motion (left or right) at 100% coherence was significantly greater than that of the other direction (Wilcoxon rank sum test, p < 0.05). Penetrations were included for analysis if they contained

at least 2 left-right selective units that were separated by at least 100 μm. We characterized the left-right selectivity of each unit with a left-right index (LRI):

$$LRI = \frac{rate_R - rate_L}{\max{(rate_R, rate_L)}} \quad (2)$$

Where $rate_R$ and $rate_L$ is the firing rate in response to rightwards and leftwards motion respectively, at 100% motion coherence. LRI values lie between -1 and +1, with -1 indicating a strongly leftwards preferring neuron, +1 indicating a strongly rightwards preferring neuron, and 0 indicating a neuron that is not selective for leftwards or rightwards motion.

Decoding: We used Linear Discriminant Analysis (Pesaran et al., 2002; Averbeck et al., 2003; Law and Gold, 2009; Adibi et al., 2014) to decode the direction of motion (leftwards or rightwards) at each coherence. Firing rates were first z-scored before being used for decoding. We used random subsampling cross validation by training on a randomly selected subset of 80% (96/120) of trials and testing on the remainder, and repeating this process 1000 times. For each iteration, we trained and tested the decoder at each level of motion coherence. Therefore, we obtained an estimate of the decoding accuracy (the mean percent correct across iterations) and the variability (95% interval across iterations). We also applied the same method to decode the direction of motion from each individual unit in order to compare them to the population performance. For the analysis of neuronal weights (the coefficients of the linear combination of neuronal responses, as determined by the decoder training), weights were normalized by dividing by the maximum absolute weight in the population.

Neurometric thresholds: To determine neurometric thresholds of single neurons and the population, the percent correct values of each cross-validation iteration were fitted using least squares regression with two variants of the Weibull function, resulting in a neurometric curve that described the neuron's performance with respect to coherence:

$$p = 1 - 0.5 \exp\left[-(c/\alpha)^\beta\right] \quad (3)$$

$$p = \delta - (\delta - 0.5) \exp\left[-(c/\alpha)^\beta\right] \quad (4)$$

where p was the probability of correctly discriminating the direction of motion at coherence c, α was

the coherence of threshold performance ($p = 82\%$, convention established by Britten et al., 1992), β

controlled the slope and δ was the asymptotic level of performance (less than 1). As Equation 4 has an

extra free parameter, we used an F-test to decide whether to reject the use of Equation 4 over Equation

3. The α was limited to between 0 and 3, β was limited to lie between 0 and 10, and δ was limited to

lie between 0 and 1. In some analyses, we used the "near-threshold" level of coherence, which was

the coherence that was closest to the exact threshold as determined by the curve fitting procedure. We

only analyzed penetrations in which the upper bound of the threshold 95% interval was less than

100% coherence, in order to ensure that threshold estimates were well constrained. Therefore, all

penetrations had a population threshold less than 100% coherence.

Decoding with and without correlations: To test the effects of correlations on population decoding, we

trained two types of decoder; the standard decoder, which was trained on the standard dataset (i.e.

contains correlations), and a "correlation blind" decoder, which was trained on trial shuffled datasets,

a process which removed all correlations. To test the effect of ignoring correlation structure, we

compared the performance of the blind decoder to the standard decoder on the standard dataset, i.e. a

data set that contained real correlations. To test the effect of removing correlations, we compared the

performance of the blind decoder on trial shuffled dataset to the standard decoder on the standard

dataset.

Spike count correlation: For each pair of units in each penetration, we calculated the spike count

correlation (r_{SC}) as the Pearson's correlation coefficient of the trial by trial spike counts for each level

of motion coherence. For each iteration of the decoding cross validation procedure, we calculated r_{SC}

and then averaged across all iterations to obtain the final estimate of r_{SC}.

9

Statistics: Measures of correlations were Spearman's rho (p), except for spike count correlations (see above). Tests between two groups were made with Wilcoxon's Rank Sign test (paired) or Wilcoxon's Rank Sum test (unpaired), and the α criterion was 0.05 unless otherwise specified. To test for differences in spike count correlations at different levels of motion coherences, we used a 2-way ANOVA with the Turkey-Cramer method for post hoc multiple comparisons, and we used an ANCOVA to account for differences in spike count correlations that might rise from differences in firing rate. We deemed differences in decoding thresholds to be statistically significant if the 95% interval of differences across cross-validation iterations did not overlap with zero.

Histology

At the end of the recordings, the animals were given an intravenous overdose of sodium pentobarbitone and, following cardiac arrest, were perfused with 0.9% saline, followed by 4% paraformaldehyde in 0.1 M phosphate buffer pH, 7.4. The brain was post-fixed for approximately 24 hours in the same solution, and then cryoprotected with fixative solutions containing 10%, 20%, and 30% sucrose. The brains were then frozen and sectioned into 40 µm coronal slices. Alternate series were stained for Nissl substance and myelin (Gallyas, 1979). The location of recording sites was reconstructed by identifying electrode tracks and depth readings recorded during the experiment. Additionally, each electrode array was coated in DiI, allowing visualization under fluorescence microscopy prior to staining of the sections. In coronal sections, MT is clearly identifiably by heavy myelination in the granular and infragranular layers (Rosa and Elston, 1998), whereas MST is more lightly myelinated and lacks clear separation between layers (Palmer and Rosa, 2006). The majority of neurons reported here were histologically confirmed to be in MT or MST, but for some penetrations in which the histology was unclear (12% of units), neurons were included on the basis of their receptive field size and progression, and their direction tuning.

Results

Sample size

We made 27 electrode array penetrations in areas MT and MST, but restricted our analysis to 18 penetrations (see Methods for inclusion criteria, MT: n = 13; MST: n = 5) that were suitable for population decoding. Decoding performance and direction selectivity in MT was similar to that of MST, so we analyzed data from both areas together. The number of units per penetration varied from 6 to 35 (median = 12, 262 in total across all penetrations), comprising of both single (11%) and multi-units, but no distinctions were made between unit type for population decoding. We performed direction tuning tests using 100% coherence motion in 12 equally spaced directions (Figure 1A), and found that most units were directional selective (MT 76%, MST 60%, Rayleigh test, $p<0.05$), in agreement with previous reports (Zeki, 1974; Maunsell and Van Essen, 1983; Albright, 1984; Celebrini and Newsome, 1994; Born and Bradley, 2005; Lui and Rosa, 2015). For population decoding analyses, we presented motion in the left-right axis and at various levels of motion coherence (Figure 1B). We observed a range of left-right selectivity across all units, with most penetrations showing a mixture of left and right preferring units (Figure 1C).

Population thresholds are always lower than the best unit's threshold

We decoded the direction of motion (leftwards or rightwards) by training linear decoders at each level of motion coherence and calculating an overall population neurometric threshold for each penetration, defined as the level of coherence in which the decoder achieves 82% accuracy. In order to assess the improvement in decoding using a population of neurons over an individual neuron, we also decoded the direction of motion (and calculated thresholds) for each unit individually. Figure 2A shows the decoding performance of a representative penetration with the population decoding performance shown in blue and the best individual unit decoding performance in shown in red. In this penetration, the population threshold is lower than the best unit's threshold. In fact, across the full set of penetrations, the population threshold was always lower than the threshold of the best unit (Figure 1B, all points lying bellow the line of unity, median difference = -17%, $p<0.001$, Wilcoxon Rank Sign

11

test). Therefore, as expected, the populations as a whole contained more information about the direction of motion than any individual unit, and the decoders were able to better discriminate the direction of motion at lower levels of motion coherence using population activity.

MT/MST correlations depend on motion strength.

We measured the spike count correlations (r_{SC}) of all pairs of units (n = 2395) in the 18 penetrations. Confirming previous reports, we found that the activity of pairs MT and MST neurons was weakly correlated on a trial to trial basis (Zohary et al., 1994; Bair et al., 2001; Cohen and Maunsell, 2009; Solomon et al., 2015; Ruff and Cohen, 2016; Zavitz et al., 2016). However, we found that r_{SC} was dependent of the level of motion coherence, and increased significantly as motion strength decreased for non-zero motion strengths (Figure 3A and Table 1, repeated measures ANOVA F(6, 2394) = 314, p < 0.001). Interestingly, 0% coherence motion did not fit this trend, showing significantly lower values of r_{SC} compared to coherences from 10% through to 64% coherence, but significantly higher values of r_{SC} than 82% and 100% coherence motion (Table 1). Because motion coherence modulates the firing rate of MT/MST cells (e.g. Figure 1B) and firing rate is known to affect r_{SC} measurements (de la Rocha et al., 2007; Cohen and Kohn, 2011), we investigated the relationship between coherence and spike rate across coherences. We plotted the normalized r_{SC} by dividing by the minimum spike count of each pair (the principal firing rate metric that affects correlations: Cohen and Kohn, 2011) and found very similar results – the strength of the r_{SC} measurements decreased with coherence except for 0% coherence (Table 1, Figure 3B, repeated measures ANOVA F(6, 2394) = 178, p < 0.001).

We also found that the correlation structure was conserved between coherences, in that pairs of neurons that were highly correlated with each other at high motion strengths tended to have also high correlations at low motion strengths. This was exemplified by the fact that r_{SC} at 10% coherence and was highly correlated with r_{SC} at 100% coherence (Spearman's ρ =0.53, p < 0.001; Fig 3C), even though the r_{SC} was significantly higher at 10% coherence (median r_{SC} difference = 0.063, p < 0.001, Wilcoxon Rank Sum test). This relationship was significant for all pairs of coherences, with correlation coefficients (Spearman's ρ, all p < 0.001) ranging from 0.46 (0 vs 100% coherence) to 0.78

12

(82 vs 100% coherence). As before (Figure 1B), we tested if differences in r_{SC} for different coherences was due to differences in the minimum spike count of the pairs of units, and plotted the r_{SC} against the minimum spike count of the pair and fit linear models to the r_{SC} values for both the 10% and 100% coherence conditions (Figure 3D, green and cyan lines) and found that the fits were statistically significantly different (p < 0.001, ANCOVA). Similar relationships were found when comparing other combinations of high and low coherences.

Removing correlations improves decoding performance

In order to examine how the trial to trial correlations between neurons affect population decoding and therefore thresholds, we first tested the effects of removing correlations on population decoding. This was done by comparing the performance of the standard decoders to decoders that were trained and tested on trial shuffled data (i.e. no correlations present). Figure 4A shows the effect of removing correlations on decoding performance of a single example penetration (blue vs yellow lines), in which decoding performance was improved by removing correlations. In general, we found that removing correlations improved decoding performance across the full dataset, and resulted in a significant lowering of thresholds in 4 penetrations, and a significant decrease in median threshold across all penetrations (Figure 4B median difference = -2.7%, p=0.025, Wilcoxon Rank Sign test).

Because correlations had a detrimental effect on population decoding, we next asked if learning the correlation structure was actually advantageous for population decoding, or if similar performance could be achieved by ignoring correlations. To test this, we compared the performance of the standard decoders (which learnt correlations) and the decoders that ignored correlations on the standard dataset (that contains correlations). The example penetration in Figure 4A show a slight impairment in decoding when ignoring the presence of correlations (blue vs red lines). While no individual penetration showed a significant increase in threshold when ignoring correlations, there was a significant increase in median difference threshold across penetrations (Figure 4C, median difference = 1.6%, p = 0.007, Wilcoxon Rank Sign test). Therefore, a decoder which learnt the correlation structures usually performed better than one that ignored correlations.

13

Finally, we investigated if there was any relationship in the effect size of removing and ignoring correlations when decoding our dataset of neuronal populations. We plotted the effect size of ignoring correlations against the effect size of removing correlations for individual penetrations and found there was a significant correlation (Figure 4D, Spearman's $\rho = 0.647$, $p = 0.005$). Therefore, the penetrations whose decoding performance improved the most by removing correlations, were least affected by ignoring correlations. Conversely, penetrations with decoding performance which were minimally affected by the removal of correlations, were most affected by ignoring correlations.

Decoders trained at 100% coherence generalize to all coherences

Because changes in motion coherence modulate spike rates and correlations, it is unclear if the optimal decoding weights would be the same across coherences. We therefore tested if the population readout method generalized across coherences. To do this, we used the decoders trained at 100% coherence to decode the direction of motion at every other level of coherence, and thereby obtained a new set of population thresholds. We then compared these thresholds to the thresholds obtained by training at each coherence level. We found there was no significant difference in population threshold between these two training methods (Figure 5A, median difference = -0.52%, $p = 0.879$, Wilcoxon Rank Sign test), suggesting that the readout method for 100% coherence can generalize to all other levels of motion coherence. To directly compare these decoding methods, we compared the weights of the decoders trained at 100% coherence with the decoders trained at the near threshold level of coherence (the closest level of coherence to the penetrations' threshold, therefore varying by penetration). We examined the normalized weights of these two decoders by plotting one against another and found they were highly correlated (Figure 5B, Spearman's $\rho = 0.711$, $p < 0.001$), again, suggesting weights are similar across coherences.

To further examine if the decoders that were optimized at 100% coherence is equivalent to the decoders optimized on a per coherence basis, we tested if the effect of removing or ignoring trial to trial correlations was the same for these two decoder types. We repeated the analysis used in Figure 4

14

to measure the increase in change when ignoring or removing correlations for the 100% coherence

trained decoder. We found that the magnitude of the change in threshold for the 100% trained

decoders was highly correlated with per coherence decoders both when correlations are removed

(Figure 5C, Spearman's $\rho = 0.938$, $p < 0.001$) and when correlations are ignored (Figure 5D,

Spearman's $\rho = 0.715$, $p = 0.001$), suggesting that the readout of the 100% coherence decoder is

similar to the per coherence decoders.

Discussion

In this paper, we have presented the first population recordings from areas MT and MST for random dot motion embedded in noise. We found that spike count correlations between neurons were dependent on motion coherence. Decreases in motion coherence generally led to increases in correlations in activity between pairs of neurons for non-zero motion strengths, but the correlation structures remained largely unchanged. When we decoded the direction of motion, one of two opposite directions of motion, we demonstrated the correlations generally impaired decoding performance, but this could be partly mitigated if the decoder to took into account the correlation structure. The optimal weights learned from high motion strengths can be applied to decoding weaker motion strengths without detrimentally affecting the performance of the decoder. Our results provide new insights into the information contained in populations of neurons in for opposite directions discrimination tasks including perceptual leaning, as discussed below.

Effects of correlations on decoding may be task-dependent

Recent work has demonstrated that, both theoretically and empirically, the presence of trial-to-trial correlations can improve population decoding (Abbott and Dayan, 1999; Sompolinsky et al., 2001; Averbeck et al., 2006; Shamir and Sompolinsky, 2006; Graf et al., 2011; Zylberberg et al., 2016). However, in our recordings, we found that correlations usually impaired population coding, and ignoring the correlation structure only had a relatively small impact on decoding performance. These results are in agreement with previous studies suggesting that correlations limit decoding performance (Zohary et al., 1994). Whether the correlation structure helps or hinders population decoding appears to be dependent on the type of task performed by the decoder (Averbeck et al., 2006; Ecker et al., 2011; Yarrow and Series, 2015). Studies that have found that population decoding performance was improved by the presence of correlations usually decode multiple directions (Zylberberg et al., 2016) or orientations (Graf et al., 2011) without stimulus noise. These studies are comparable with a fine discrimination task, in which subjects make judgements between small differences in stimulus attributes. Therefore, the correlation structure of neurons in MT/MST may be beneficial for other

16

types of stimulus feature decoding (Zavitz et al., 2017), even though they impair population decoding for a 2AFC opposite direction task.

Our results are compatible with studies of attention, which show that attention decreases neuronal correlations and improves stimulus feature decoding (Cohen and Maunsell, 2009, 2011; Mitchell et al., 2009). In the context of a 2AFC opposite directions task, the reduction in neuronal correlations caused by attention could improve decoding performance, since our study showed that removing correlations improved population decoding performance. However, it should be noted though that small decreases in correlations due to attention are not equivalent to the artificial removal of correlations in this study, and to date, no study has examined the effects of attention on correlations and decoding in MT for a 2AFC opposite directions task.

Stronger correlations in response to weaker motion signals

We found that the mean spike count correlation (r_{SC}) in MT and MST in response to random dot stimuli ranged from 0.17 (100% coherence) to 0.26 (10% coherence). The strength of the r_{SC} measurements at 100% coherence was similar to previous studies that recorded pairs of nearby neurons (0.12: Zohary et al., 1994; 0.2: Bair et al., 2001; 0.13: Cohen and Newsome, 2008; 0.1: Huang and Lisberger, 2009). In addition, we found that r_{SC} was generally higher in response to weaker motion signals compared to stronger motion signals, which had not been reported before. Zohary et al. (1994) did measure responses to different levels of motion coherence as part of their r_{SC} measurement, but reported that there was no statistically significant difference in r_{SC} between coherence conditions on and individual pair basis, and therefore combined their r_{SC} measurements across coherence.

The finding that lower coherence stimuli produced higher values for r_{SC} was somewhat surprising given that stronger motion signals elicit high spike rates in the preferred direction of motion (Britten et al., 1992, 1993; Chaplin et al., 2017a), which would be expected to increase correlations (de la Rocha et al., 2007; Cohen and Kohn, 2011). We demonstrated that correlations were higher at lower levels of motion coherence even when taking spike rates into consideration. This is in agreement with the

17

effects of changes in contrast of sinusoidal gratings on r_{SC} measurements in the primary visual cortex of macaques (Smith and Kohn, 2008), in which low contrast gratings (which elicit lower spikes rates) produced higher correlations. These findings support the notion that correlations are higher when sensory signals are weaker.

Correlations in spiking activity are thought to arise from both shared inputs and local horizontal connections (Zohary et al., 1994; Bair and Movshon, 2004; Kohn and Smith, 2005; Smith and Kohn, 2008). Our measurements of r_{SC} suggest that the strength of the inputs that produce correlations in activity are reduced as motion strength increases. This could be explained by an inhibitory normalization process, in which high motion coherences elicit greater overall spike rates, which in turn act to suppress cortical activity, presumably though inhibitory interneurons, and reduce the strength of the inputs that cause correlations. An alternative explanation is that since the low coherence stimuli also contain motion signals in many different directions (due to the random displacement of the noise dots), they may activate many more direction selective neurons at once, both within MT/MST, and in their input regions. Therefore, low coherence stimuli could potentially activate larger pool of neurons, which could produce more correlated activity between pairs of neurons within in MT and MST.

Surprisingly, we also found that the 0% coherence did not fit the general trend of lower coherence stimuli producing higher spike count correlations, since it elicited significantly lower spike count correlations than stimuli of low motion strength (Table 1, Figure 3A), implying that the effects of the 0% coherence stimuli on MT/MST neurons is fundamentally different to that of low coherence stimuli. It is likely that the 10% coherence stimulus (the lowest non-zero coherence used in this study) is near or above the perceptual thresholds of marmosets, since perceptual behavioral thresholds of macaques ranges from 5 to 20% coherence (Britten et al., 1992; Roitman and Shadlen, 2002; Law and Gold, 2008; Cohen and Newsome, 2009), and therefore the 10% coherence stimulus may be sufficient to drive direction selective responses in MT/MST neurons and result in motion percepts. In contrast, the 0% coherence stimulus would not produce a reliable percept of motion, presumable because it does not produce direction selective responses, but perhaps also because it results in weaker

18

correlations in spiking activity. This could be tested by presenting very weak but non-zero coherences (for example 1%) which are well below perceptual thresholds (i.e. the animal cannot perceive the direction of motion at all) to see it results in higher correlations than the zero coherence condition.

Generalizability of decoding

Our earlier work demonstrated that sensitivity of individual neurons to noisy random dot stimuli can be reliably predicted from its responses to the 100% coherence stimulus (Chaplin et al., 2017). Here we showed that the neurons that contained the most information at 100% coherence will more than likely be the neurons carry the most information at lower coherences. These results suggest, in the context of population decoding, that weights learnt using the 100% coherence can also be applied to lower, and even near threshold, coherences with little loss of information in comparison to weights optimized at each individual coherence. The coherence invariant read-out was aided by the relatively stability of neural correlations at different coherences (Figure 3). The fact that pairwise neuronal correlations were dependent on motion coherence meant that the coherence invariant readout may not be possible, however, this did not prove to be the case.

While there is evidence to suggest that perceptual leaning involves an improvement in the sensory representation of stimulus features (Schoups et al., 2001; Yang and Maunsell, 2004; Raiguel et al., 2006), the improvements appear to be minimal, particularly for area MT in a 2AFC opposite directions of motion task (Law and Gold, 2008). Improvements can be best accounted for by optimizing weights via changing feedforward connectivity (Law and Gold, 2009; Bejjanki et al., 2011). Essentially, the process will enable the neurons which carry the most task-relevant information to contribute the most to the decision, which is the same process as training the decoder to optimize weights in the present study. The results of the present study imply that the system can be trained to perform the task at 100% coherence, and then apply same decoding strategy at lower coherences, and still perform relatively well. This makes perceptual learning a simpler process than if the weights have to be refined substantially with respect to changes in coherences.

References

Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. Neural Comput 11:91–101.

Adibi M, McDonald JS, Clifford CWG, Arabzadeh E (2014) Population decoding in rat barrel cortex: optimizing the linear readout of correlated population responses. PLoS Comput Biol 10:e1003415.

Albright TD (1984) Direction and orientation selectivity of neurons in visual area MT of the macaque. J Neurophysiol 52:1106–1130.

Averbeck BB, Crowe DA, Chafee M V., Georgopoulos AP (2003) Neural activity in prefrontal cortex during copying geometrical shapes. Exp Brain Res 150:142–153.

Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. Nat Rev Neurosci 7:358–66.

Bair W, Movshon JA (2004) Adaptive temporal integration of motion in direction-selective neurons in macaque visual cortex. J Neurosci 24:7305–23.

Bair W, Zohary E, Newsome WT (2001) Correlated firing in macaque visual area MT: time scales and relationship to behavior. J Neurosci 21:1676–1697.

Bejjanki VR, Beck JM, Lu ZL, Pouget A (2011) Perceptual learning as improved probabilistic inference in early sensory areas. Nat Neurosci 14:642–650.

Berens P (2009) CircStat⊡: A MATLAB toolbox for circular statistics. J Stat Softw 31:1–21.

Born RT, Bradley DC (2005) Structure and function of visual area MT. Annu Rev Neurosci 28:157–189.

Bourne JA, Rosa MGP (2003) Preparation for the in vivo recording of neuronal responses in the visual cortex of anaesthetised marmosets (Callithrix jacchus). Brain Res Protoc 11:168–177.

Brainard DH (1997) The psychophysics toolbox. Spat Vis 10:433–6.

Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis Neurosci 13:87–100.

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. J Neurosci 12:4745–4765.

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. Vis Neurosci 10:1157–1169.

Celebrini S, Newsome WT (1994) Neuronal and psychophysical sensitivity to motion signals in extrastriate area MST of the macaque monkey. J Neurosci 14:4109–4124.

Chaplin TA, Allitt BJ, Hagan MA, Price NSC, Rajan R, Rosa MGP, Lui LL (2017a) Sensitivity of neurons in the middle temporal area of marmoset monkeys to random dot motion. J Neurophysiol 118:1567–1580.

Chaplin TA, Allitt BJ, Hagan MA, Rajan R, Rosa MGP, Lui LL (2017b) Auditory motion does not modulate spiking activity in visual motion processing areas MT and MST. bioRxiv.

Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. Nat Neurosci 14:811–819.

Cohen MR, Maunsell JHR (2009) Attention improves performance primarily by reducing interneuronal correlations. Nat Neurosci 12:1594–1600.

Cohen MR, Maunsell JHR (2011) Using neuronal populations to study the mechanisms underlying spatial and feature attention. Neuron 70:1192–204.

Cohen MR, Newsome WT (2008) Context-dependent changes in functional circuitry in visual area MT. Neuron 60:162–73.

Cohen MR, Newsome WT (2009) Estimates of the contribution of single neurons to perception depend on timescale and noise correlation. J Neurosci 29:6635–48.

de la Rocha J, Doiron B, Shea-Brown E, Josić K, Reyes A (2007) Correlation between neural spike trains increases with firing rate. Nature 448:802–6.

Ecker AS, Berens P, Tolias AS, Bethge M (2011) The effect of noise correlations in populations of diversely tuned neurons. J Neurosci 31:14272–83.

Gallyas F (1979) Silver staining of myelin by means of physical development. Neurol Res 1:203–209.

Goris RLT, Movshon JA, Simoncelli EP (2014) Partitioning neuronal variability. Nat Neurosci

17:858–865.

Graf AB, Kohn A, Jazayeri M, Movshon JA (2011) Decoding the activity of neuronal populations in macaque primary visual cortex. Nat Neurosci 14:239–245.

Huang X, Lisberger SG (2009) Noise correlations in cortical area MT and their potential impact on trial-by-trial variation in the direction and speed of smooth-pursuit eye movements. J Neurophysiol 101:3012–30.

Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. J Neurosci 25:3661–73.

Law C-T, Gold JI (2008) Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. Nat Neurosci 11:505–513.

Law C-T, Gold JI (2009) Reinforcement learning can account for associative and perceptual learning on a visual-decision task. Nat Neurosci 12:655–663.

Lui LL, Rosa MGP (2015) Structure and function of the middle temporal visual area (MT) in the marmoset: Comparisons with the macaque monkey. Neurosci Res 93:62–71.

Maunsell JHR, Van Essen DC (1983) Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. J Neurophysiol 49:1127–1147.

Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. Neuron 63:879–88.

Newsome WT, Britten KH, Movshon JA (1989) Neuronal correlates of a perceptual decision. Nature 341:52–54.

Palmer SM, Rosa MGP (2006) A distinct anatomical network of cortical areas for analysis of motion in far peripheral vision. Eur J Neurosci 24:2389–2405.

Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. Nat Neurosci 5:805–811.

Raiguel S, Vogels R, Mysore SG, Orban GA (2006) Learning to see the difference specifically alters the most informative V4 neurons. J Neurosci 26:6589–602.

Ringach DL, Shapley RM, Hawken MJ (2002) Orientation selectivity in macaque V1: diversity and

laminar dependence. J Neurosci 22:5639–51.

Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J Neurosci 22:9475–89.

Rosa MGP, Elston GN (1998) Visuotopic organisation and neuronal response selectivity for direction of motion in visual areas of the caudal temporal lobe of the marmoset monkey (Callithrix jacchus): middle temporal area, middle temporal crescent, and surrounding cortex. J Comp Neurol 393:505–527.

Ruff DA, Cohen MR (2016) Attention increases spike count correlations between visual cortical areas. J Neurosci 36:7523–7534.

Schoups A, Vogels R, Qian N, Orban GA (2001) Practising orientation identification improves orientation coding in V1 neurons. Nature 412:549–553.

Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996) A computational analysis of the relationship between neuronal and behavioral responses to visual motion. J Neurosci 16:1486–1510.

Shamir M, Sompolinsky H (2006) Implications of neuronal diversity on population coding. Neural Comput 18:1951–86.

Smith M a, Kohn A (2008) Spatial and temporal scales of neuronal correlation in primary visual cortex. J Neurosci 28:12591–603.

Solomon SS, Chen SC, Morley JW, Solomon SG (2015) Local and global correlations between neurons in the middle temporal area of primate visual cortex. Cereb Cortex 25:3182–3196.

Solomon SS, Tailby C, Gharaei S, Camp AJA, Bourne JA, Solomon SG (2011) Visual motion integration by neurons in the middle temporal area of a New World monkey, the marmoset. J Physiol 23:5741–5758.

Sompolinsky H, Yoon H, Kang K, Shamir M (2001) Population coding in neuronal systems with correlated noise. Phys Rev E 64:51904.

Tanaka K, Hikosaka K, Saito H, Yukie M, Fukada Y, Iwai E (1986) Analysis of local and wide-field movements in the superior temporal visual areas of the macaque monkey. J Neurosci 6:134–44.

Tolhurst DJ, Movshon JA, Dean AF (1983) The statistical reliability of signals in single neurons in cat and monkey visual cortex. Vision Res 23:775–85.

Van Essen DC, Maunsell JHR, Bixby JL (1981) The middle temporal visual area in the macaque: myeloarchitecture, connections, functional properties and topographic organization. J Comp Neurol 199:293–326.

Yang T, Maunsell JHR (2004) The effect of perceptual learning on neuronal responses in monkey visual area V4. J Neurosci 24:1617–26.

Yarrow S, Series P (2015) The influence of population size, noise strength and behavioral task on best-encoded stimulus for neurons with unimodal or monotonic tuning curves. Front Comput Neurosci 9:1–20.

Yu H-H, Rosa MGP (2010) A simple method for creating wide-field visual stimulus for electrophysiology⍰: Mapping and analyzing receptive fields using a hemispheric display. J Vis 10:1–16.

Zavitz E, Yu H-H, Rosa MGP, Price NSC (2017) Correlated variability in the neurons with the strongest tuning improves direction coding. Cereb Cortex 1–12.

Zavitz E, Yu H-H, Row EG, Rosa MGP, Price, Nicholas SC (2016) Rapid adaptation induces persistent biases in population codes. J Neurosci 36:4579–4590.

Zeki S (1974) Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. J Physiol 236:549–573.

Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. Nature.

Zylberberg J (2017) Untuned but not irrelevant: A role for untuned neurons in sensory information coding. bioRxiv 1–18.

Zylberberg J, Cafaro J, Turner MH, Shea-Brown E, Rieke F (2016) Direction-selective circuits shape noise to ensure a precise population code. Neuron 89:369–383.

Table 1: Results of post hoc comparisons of differences in spike count correlation (r_{SC}) across

coherences.

| Coherences (%)	Mean r_{SC} difference (CI)	p value	Mean r_{SC} / spike difference (CI)	p value
0 vs 10	-0.043 (-0.050, -0.036)	< 0.001	-0.046 (-0.058, -0.034)	< 0.001
0 vs 28	-0.047 (-0.055, -0.040)	< 0.001	-0.041 (-0.053, -0.029)	< 0.001
0 vs 46	-0.025 (-0.032, -0.018)	< 0.001	-0.027 (-0.039, -0.014)	< 0.001
0 vs 64	-0.036 (-0.043, -0.029)	< 0.001	-0.020 (-0.033, -0.008)	< 0.001
0 vs 82	0.012 (0.005, 0.019)	< 0.001	0.037 (0.025, 0.050)	< 0.001
0 vs 100	0.039 (0.032, 0.046)	< 0.001	0.057 (0.045, 0.070)	< 0.001
10 vs 28	-0.004 (-0.012, 0.003)	0.548	0.005 (-0.007, 0.017)	0.906
10 vs 46	0.018 (0.011, 0.025)	< 0.001	0.019 (0.007, 0.032)	< 0.001
10 vs 64	0.007 (0.000, 0.014)	0.077	0.026 (0.013, 0.038)	< 0.001
10 vs 82	0.055 (0.048, 0.062)	< 0.001	0.083 (0.071, 0.096)	< 0.001
10 vs 100	0.082 (0.075, 0.089)	< 0.001	0.103 (0.091, 0.116)	< 0.001
28 vs 46	0.022 (0.015, 0.030)	< 0.001	0.014 (0.002, 0.027)	0.011
28 vs 64	0.011 (0.004, 0.019)	< 0.001	0.021 (0.008, 0.033)	< 0.001
28 vs 82	0.059 (0.052, 0.067)	< 0.001	0.078 (0.066, 0.091)	< 0.001
28 vs 100	0.086 (0.079, 0.094)	< 0.001	0.098 (0.086, 0.111)	< 0.001
46 vs 64	-0.011 (-0.018, -0.004)	< 0.001	0.006 (-0.006, 0.019)	0.731
46 vs 82	0.037 (0.030, 0.044)	< 0.001	0.064 (0.052, 0.077	< 0.001
46 vs 100	0.064 (0.057, 0.071)	< 0.001	0.084 (0.072, 0.097)	< 0.001
64 vs 82	0.048 (0.041, 0.055)	< 0.001	0.058 (0.045, 0.070)	< 0.001
64 vs 100	0.075 (0.068, 0.082)	< 0.001	0.078 (0.065, 0.090)	< 0.001
82 vs 100	0.027 (0.020, 0.034)	< 0.001	0.020 (0.008, 0.032)	< 0.001
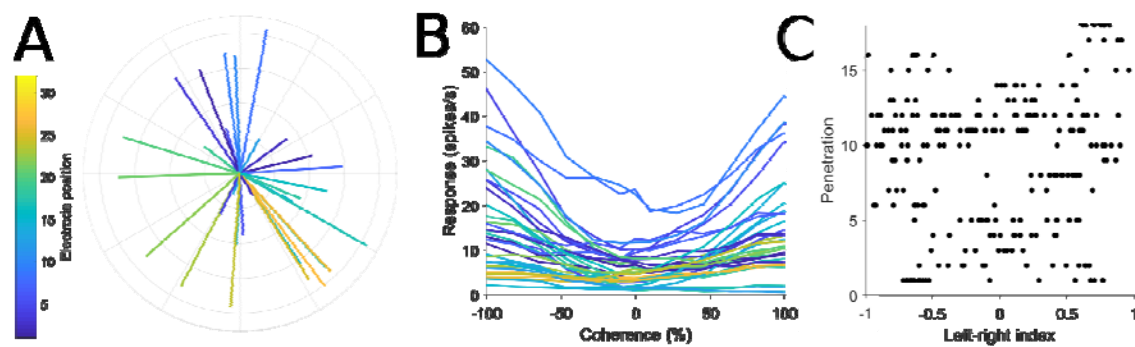
Figure 1: Direction preference and left-right selectivity of neural populations. A: Polar plot showing the direction tuning of each unit in an example penetration. The direction of each line represents the preferred direction, the lengths represent the strength of direction selectivity (direction index, DI) and the color shows the position of the unit on the electrode shank. B: The same penetration from panel A, showing the responses to motion on the left-right axis of motion at different levels of motion coherence, which were used for decoding. Each line shows the coherence response function for a unit from the example penetration shown in A, using the same color convention. Leftwards motion is represented by negative coherence values, rightwards motion by positive coherence values. C: Distribution of left-right selectivity across all penetrations in this study. Each row of dots represents a penetration and the position of each dot on the x-axis represents the left-right index (LRI). Dots positioned on the left (negative) side of the axis prefer leftwards motion, dots positioned on the right (positive) side of the axis prefer rightwards motion.
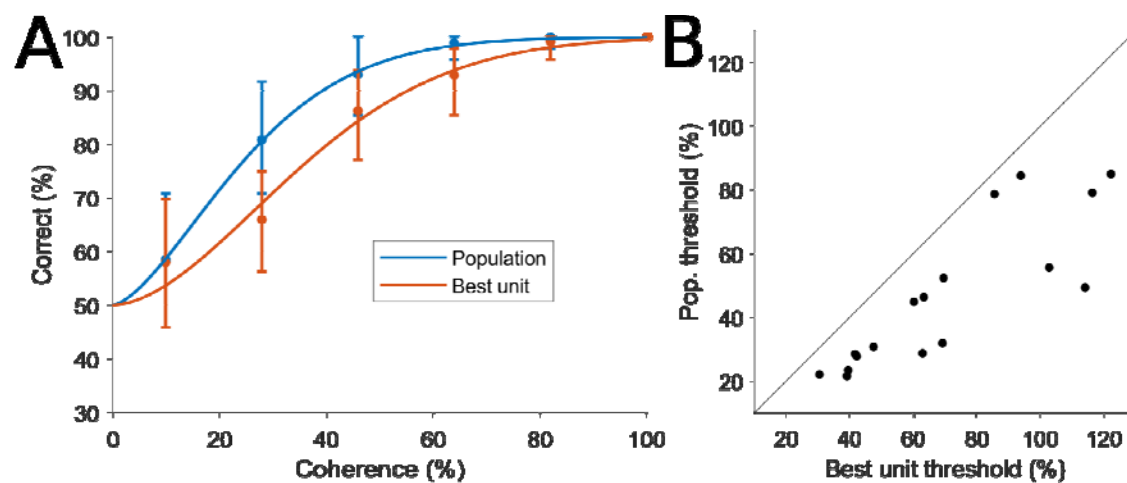
Figure 2: Improvement of population decoding over individual units. A: An example penetration showing the decoding performance of the population (blue) and the best individual unit (red) plotted against coherence. The two data sets were both fit with a Weibull curve to determine the threshold, defined as the level of coherence that achieves 82% correct. B: Population thresholds plotted against the threshold of the best unit for all populations. All penetrations had population thresholds lower than the threshold of the best individual unit.
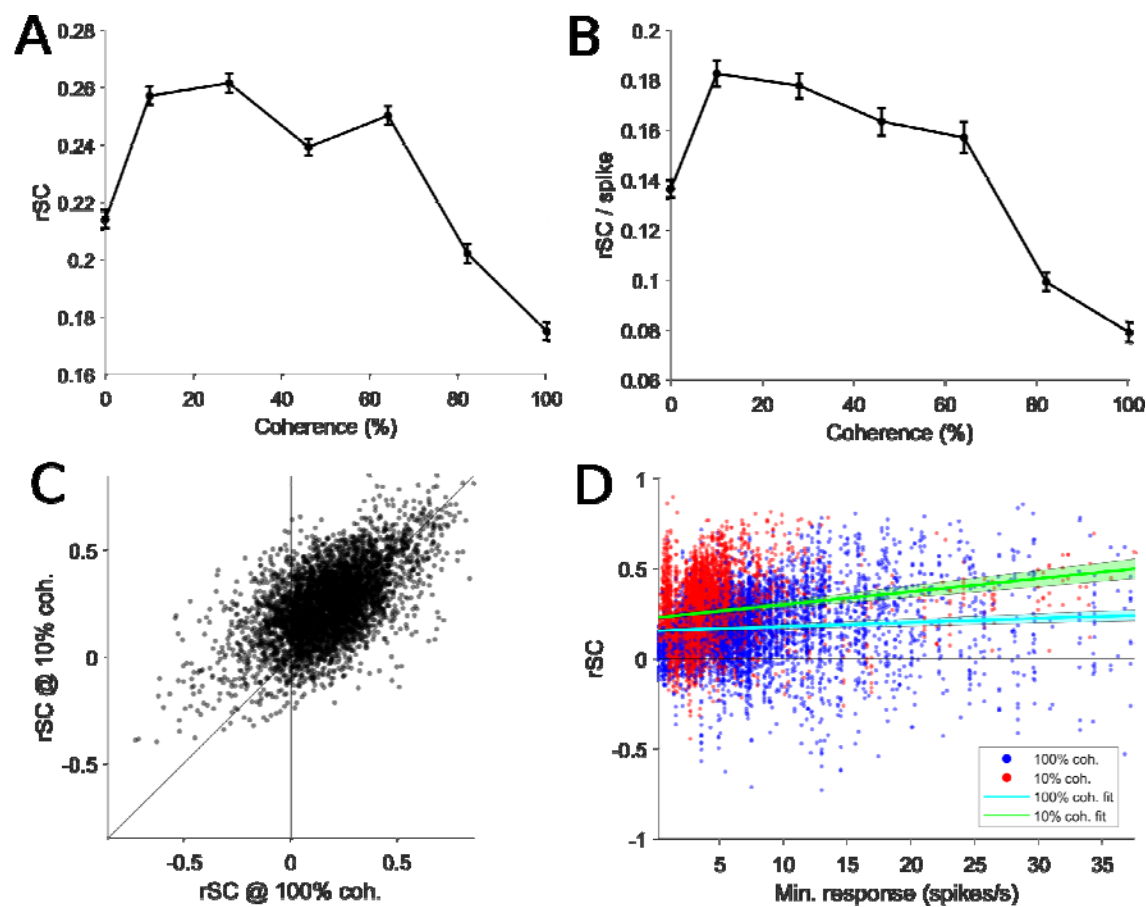
Figure 3: Spike count correlation (r_{SC}) at different levels of motion coherence. A: The mean r_{SC} plotted against each level of coherence, error bars show the standard error of the mean, and r_{SC} varied with coherence (repeated measures ANOVA $p < 0.001$). B: The mean r_{SC} normalized by dividing by the minimum spike count of the neuronal pair, plotted against each level of coherence, error bars show the standard deviation. As in C, r_{SC} varied with coherence (repeated measures ANOVA $p < 0.001$). C: r_{SC} at 100% coherence plotted against the r_{SC} at 10% coherence. Each point represents the correlation of pair of units for a particular direction of motion (left or right), showing a statistically significant correlation (Spearman's $\rho = 0.529$, $p < 0.001$). D: r_{SC} values from C plotted against the minimum spike count of the two units in the pair. Blue points represent data from 100% coherence, red points represent the data from 10% coherence, and the lines of best fit are plotted in cyan and green respectively, with the 95% confidence intervals of the fit shown by the light shading.
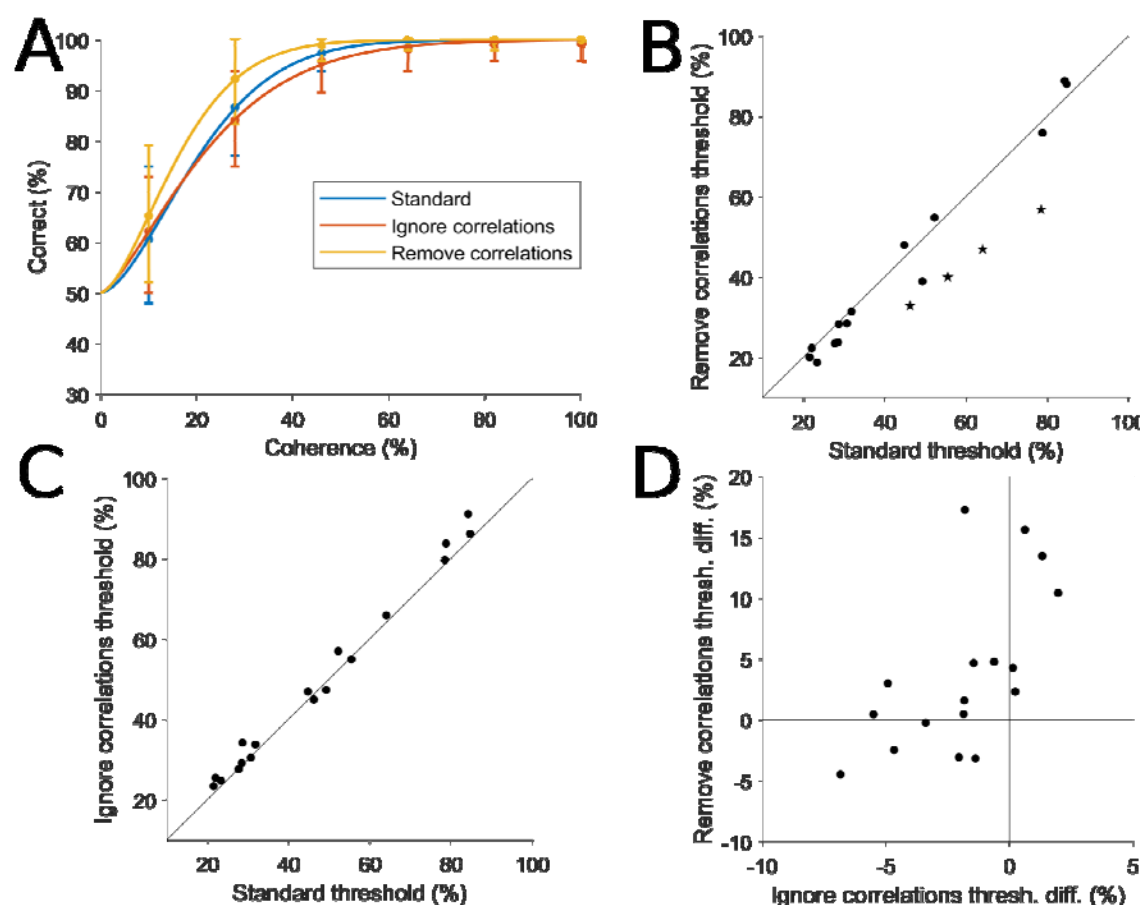
Figure 4: Effects of neuronal correlations on population decoding. A: An example penetration, showing the performance of the standard decoding procedure (blue), performance when correlations were removed (yellow) and performance when correlations were ignored (red). Each data sets were fit with a Weibull curve to determine the threshold, defined as the level of coherence in which the decoding accuracy reaches 82% correct. B: Effects of removing correlations. The thresholds from the standard decoding procedure are plotted against the thresholds obtained when correlations were removed, showing a statistically significant decrease in the median threshold (p = 0.025, Wilcoxon Rank Sign test). Star symbols represent penetrations in which the difference is statistically significant. C: Effects of ignoring correlations. The thresholds from the standard decoding procedure are plotted against the thresholds obtained when correlations were ignored, showing a small but statistically significant increase in the median threshold (p = 0.007, Wilcoxon Rank Sign test). D: The change in thresholds when correlations were ignored was significantly positively correlated with change in thresholds when correlations were removed (Spearman's ρ = 0.647, p = 0.005). The change in

threshold was calculated as the standard threshold minus the remove or ignore correlations threshold.

Therefore, for the x-axis, positive values represent penetrations which had higher (worse performance)

thresholds in the standard decoding compared to decoding when correlations where ignored. These

were the units that were below the line of unity in C, and were the minority of penetrations. For the y-

axis, positive values indicate penetrations that had higher thresholds (worse performance) for the

standard decoder in comparison to decoding when correlations where removed. These were the units

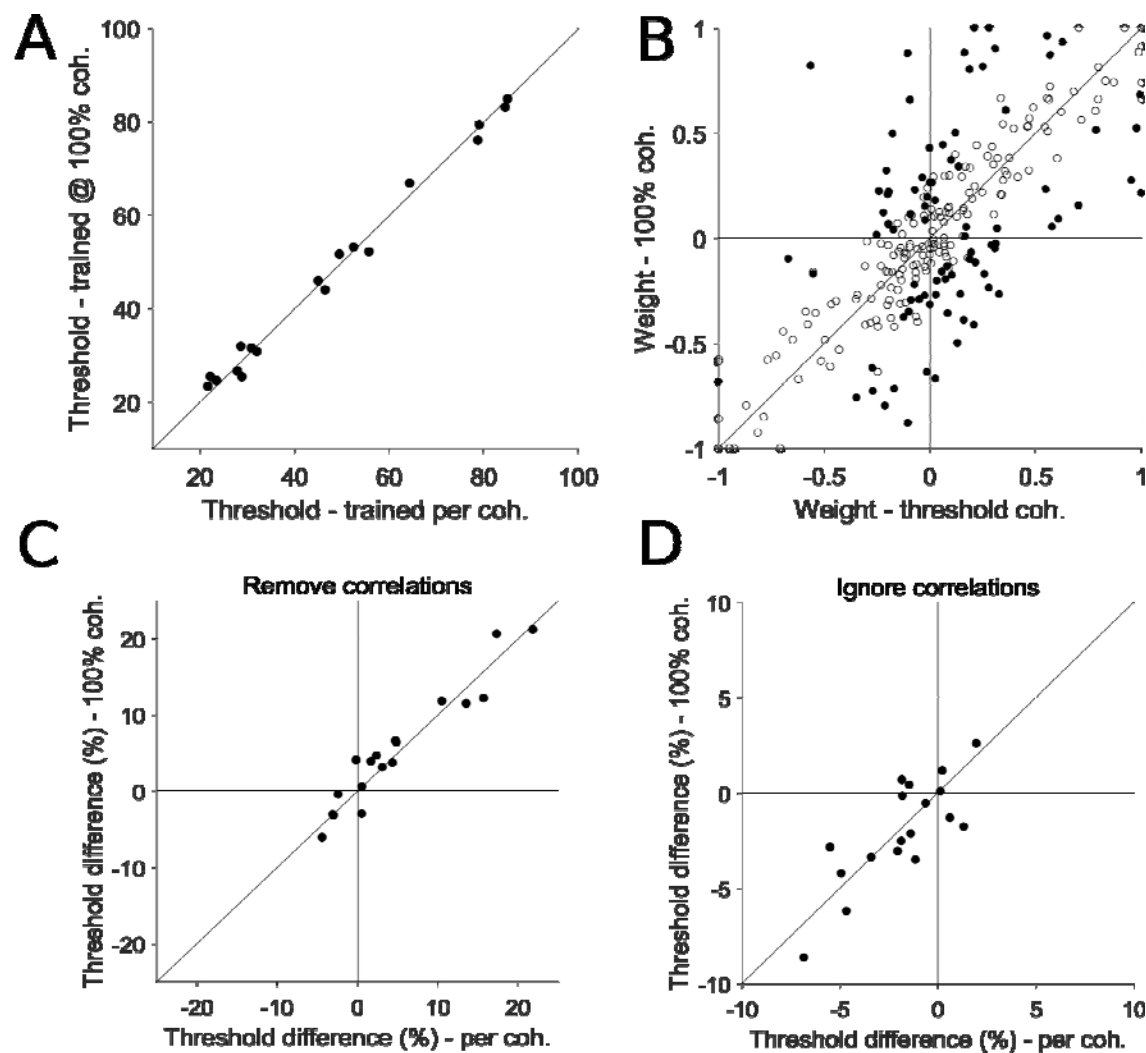above the line of unit in B and were the majority of the penetrations.

Figure 5: Generalizability of decoders across motion coherence. A: Population thresholds of the standard decoder (trained and tested for each coherence) are plotted against the thresholds of the decoders trained only at 100% coherence (and tested at every other coherence). The median threshold difference was not statistically significantly different to zero (p = 0.826, Wilcoxon Rank Sign test). B: Normalized decoder weights of the per coherence decoders plotted against the 100% coherence decoders. Each point represents the weight of an individual unit, filled circles represent weights that are statistically different at the 100% and near threshold conditions. The weights of the two decoder types were strongly correlated (Spearman's ρ = 0.711, p < 0.001). C: The changes in thresholds for the two types of decoders was very similar when removing correlations (Spearman's ρ = 0.938, p < 0.001). As in Figure 4D, the change in threshold was calculated as the standard threshold minus the remove correlations threshold. D: The changes in threshold for the two types of decoders was also

very similar when ignoring correlations (Spearman's $\rho = 0.715$, p = 0.001). Similar to panel C, the

change in threshold was calculated as the standard threshold minus the ignore correlations threshold.