Title: Evidence for time division multiplexing of multiple simultaneous items in a sensory coding bottleneck

Short title: Fluctuations in neural activity

Classification*:* BIOLOGICAL SCIENCES, Neuroscience

Authors: V. C. Caruso[1,2,3,4]†, J. T. Mohl[1,2,3,4], C. Glynn[7,8], J. Lee[1,2,3,4]‡, S. Willett[1,2,3,4], A. Zaman[6], R. Estrada[5,9], S. Tokdar[1,6]†, J. M. Groh[1,2,3,4]†

Affiliations:

[1]Duke Institute for Brain Sciences, Duke University,
308 Research Drive, LSRC M051, Box 91003
Durham, NC 27708

[2]Center for Cognitive Neuroscience, Duke University,
308 Research Drive, LSRC M051, Box 91003
Durham, NC 27708

[3]Department of Psychology and Neuroscience, Duke University,
417 Chapel Drive, Box 90086
Durham, NC 27708-0086

[4]Department of Neurobiology, Duke University,
311 Research Drive, Box 3209
Durham, NC 27710

[5]Department of Computer Science, Duke University,
LSRC Building D101, 308 Research Drive, Box 90129
Durham, NC 27708-0129

[6]Department of Statistical Science, Duke University,
P.O. Box 90251
Durham, NC 27708-0251

[7]Department of Statistics, University of Washington,

Box 354322
Seattle, WA 98195-4322

[8]Department of Decision Sciences, University of New Hampshire,
Durham, NH 03824

[9]Teledyne Scientific Company, Research Triangle Park, NC,
5001 S. Miami Blvd., Suite 200, Durham, NC 27703

† Co-corresponding authors. Correspondence to
VCC at v.caruso@duke.edu, (919) 6846729
Duke Institute for Brain Sciences, Duke University,
308 Research Drive, LSRC M051, Box 91003
Durham, NC 27708

or JMG at jmgroh@duke.edu, (919) 681-6536
Duke Institute for Brain Sciences, Duke University,
308 Research Drive, LSRC M051, Box 91003
Durham, NC 27708

or ST at st118@stat.duke.edu, (919) 684 2152
Duke Institute for Brain Sciences, Duke University,
308 Research Drive, LSRC M051, Box 91003
Durham, NC 27708

‡ Current address: Department of Psychology, University of Canterbury, New Zealand,
Private Bag 4800, Christchurch

Keywords: neural coding, time division multiplexing, statistical methods, perception

ABSTRACT:

How the brain preserves information about multiple simultaneous items is poorly understood. Here, we provide evidence that the brain may accomplish this using time division multiplexing, or interleaving of different signals across time, to represent multiple items in a single neural channel. We evaluated single unit activity in an auditory coding "bottleneck", the inferior colliculus, while monkeys reported the location(s) of one or two simultaneous sounds. Using novel statistical methods to evaluate spiking activity on a variety of time scales, we found that on dual-sound trials, neurons sometimes alternated between firing rates similar to those observed for each single sound. These fluctuations could occur either across or within trials and appeared coordinated across pairs of simultaneously recorded neurons. Fluctuations could be predicted by the state of local field potentials prior to sound onset, and, in one monkey, predicted which sound the monkey would ultimately saccade to first. Alternation between activity patterns corresponding to each of multiple items may be a general strategy employed by the brain to enhance its processing capacity, suggesting a potential connection between such disparate phenomena as variable neural firing, neural oscillations, and limits in attentional or memory capacity.

SIGNIFICANCE STATEMENT:

In natural scenes, many things happen at once. Given that the neural populations activated by each stimulus overlap considerably, how does the brain preserve information about each item? We investigated whether the brain solves this problem using *time division multiplexing*, a telecommunications strategy for combining signals in a single channel. When two sounds were presented, we observed fluctuations in activity in the inferior colliculus at a variety of time scales. These fluctuations were not random but suggested that neurons switch back and forth between encoding different items. Such switching behavior provides a novel account for variability in neural firing, and suggests an explanation for limits in perception, attention, and working memory.

MAIN TEXT:

1. Introduction

In the natural world many stimuli or events occur at the same time, evoking activity in an overlapping population of neurons. When neurons are exposed to more than one stimulus to which they can respond, how might they preserve information about each stimulus? In this study we investigated whether spike trains contain interleaved signals corresponding to each stimulus, akin to time-division multiplexing used in telecommunications (Figure 1), and postulated to occur in some form in the brain (1-10).

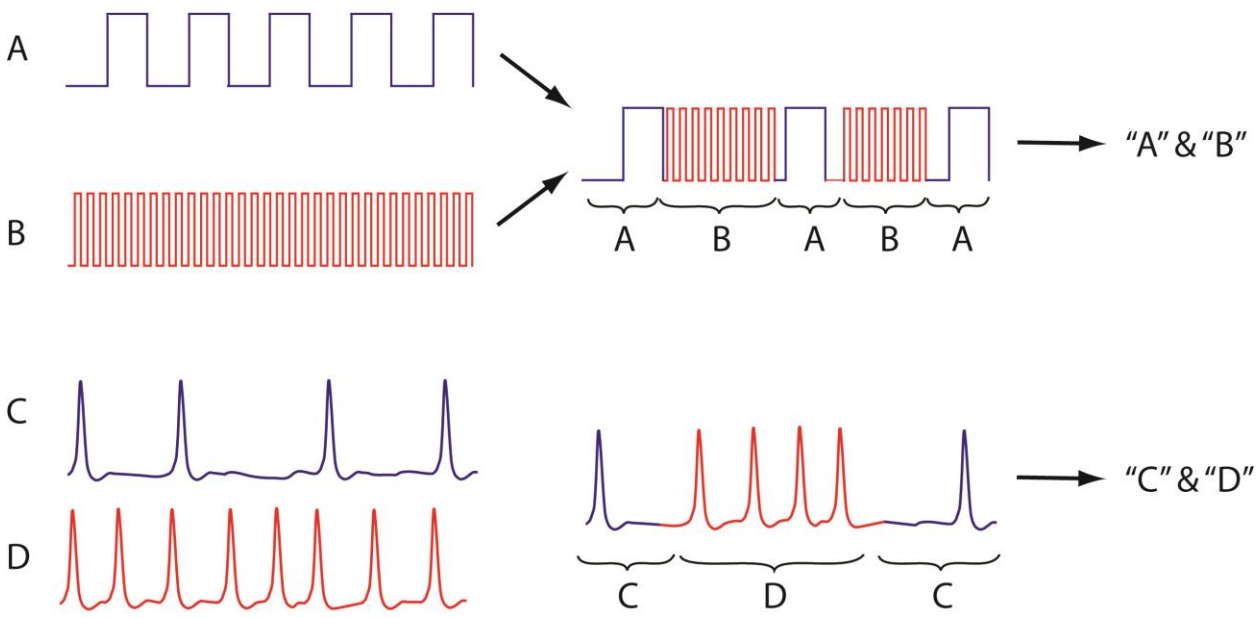

Figure 1. *In telecommunications, multiple signals can be conveyed along a single transmission line by interleaving samples (A and B). This process greatly increases the amount of information that can be transmitted by a single physical resource. In this study we investigated whether the brain might employ a similar strategy, i.e. do neurons encode multiple items using spike trains that alternate between the firing rates corresponding to each item, at some unknown time scale?*

Multiplexing is most likely to occur when there is an information-processing bottleneck. The coding of sound locations involves such a bottleneck. Sound waves stemming from two sources sum in the world and are sampled at only two locations, i.e. at each ear. In barn owls, multiple locations appear to be de-multiplexed from these signals and encoded as distinct peaks

in auditory space maps (11-14). But in primates (including humans) and several other mammalian species, the neural representations themselves involve a bottleneck (15-22). The inferior colliculus (IC) and other auditory structures encode sound location not in a map but in a "meter": a firing rate code in which neural activity is roughly proportional to the horizontal angle of the sound, reaching an apex (or nadir) at 90 degrees contralateral (or ipsilateral) along the axis of the ears, where the binaural timing and level differences reach their maximal (or minimal) values (Figure 2D,F) (15-22).

A strict meter/firing rate code would seem unable to represent more than one sound location *except* via multiplexing. The auditory pathway's maps for sound frequency can only partially ameliorate this situation. Such maps serve to separate the coding of sounds of different frequencies to somewhat different neural subpopulations. However, most natural sounds are spectrally rich and will activate overlapping "hills" of neural activity; even a single pure tone of a particular frequency can evoke activity in 40-80% of IC neurons (23). This raises the question of how a population consisting of such broadly-tuned neurons can preserve information about combinations of sounds, even when they differ in sound frequency. Alternating the coding of different sounds across time would potentially solve this problem.

In this study, we show that monkeys can localize both of two simultaneous sounds, but conventional time-and-trial pooled measures of IC activity do not seem to match this ability -- exhibiting instead an intermediate level apparently more consistent with the presence of a single sound between the two actual targets. We resolve this paradox using novel statistical methods that reveal IC activity actually fluctuates in time between the responses evoked by each sound when presented alone, potentially allowing both sounds to be represented simultaneously in the neural population. Such fluctuations occur at different time scales and have the following features: 1) they are (weakly) coordinated across the population of neurons, 2) they can be predicted from network state prior to sound onset (assessed by local field potentials), 3) they can predict which sound location the monkey would report first. We discuss how such a multiplexed code can be read out, as well as the broad implications of activity fluctuations for neural encoding.

2. Results

2.1. Monkeys can report the locations of both sounds, indicating that both are coded in brain

We first tested whether monkeys can perceptually preserve information about multiple sounds presented simultaneously. Monkeys performed a localization task in which they made eye movements to each of the sounds they heard: one saccade on single-sound trials and two saccades in sequence on dual-sound trials (Figure 2A). The sounds were separated horizontally by 30 degrees and consisted of band-limited noise with different center frequencies. They were thus physically distinguishable in principle, and humans can do so (24-26). The monkeys learned the task successfully (example session shown in Figure 2B), and, like humans, typically performed better when the frequency separation between the two sounds was larger (Figure 2C, ~72 vs. ~77% correct for frequency differences of 3.4 vs. 6.8 semitones).

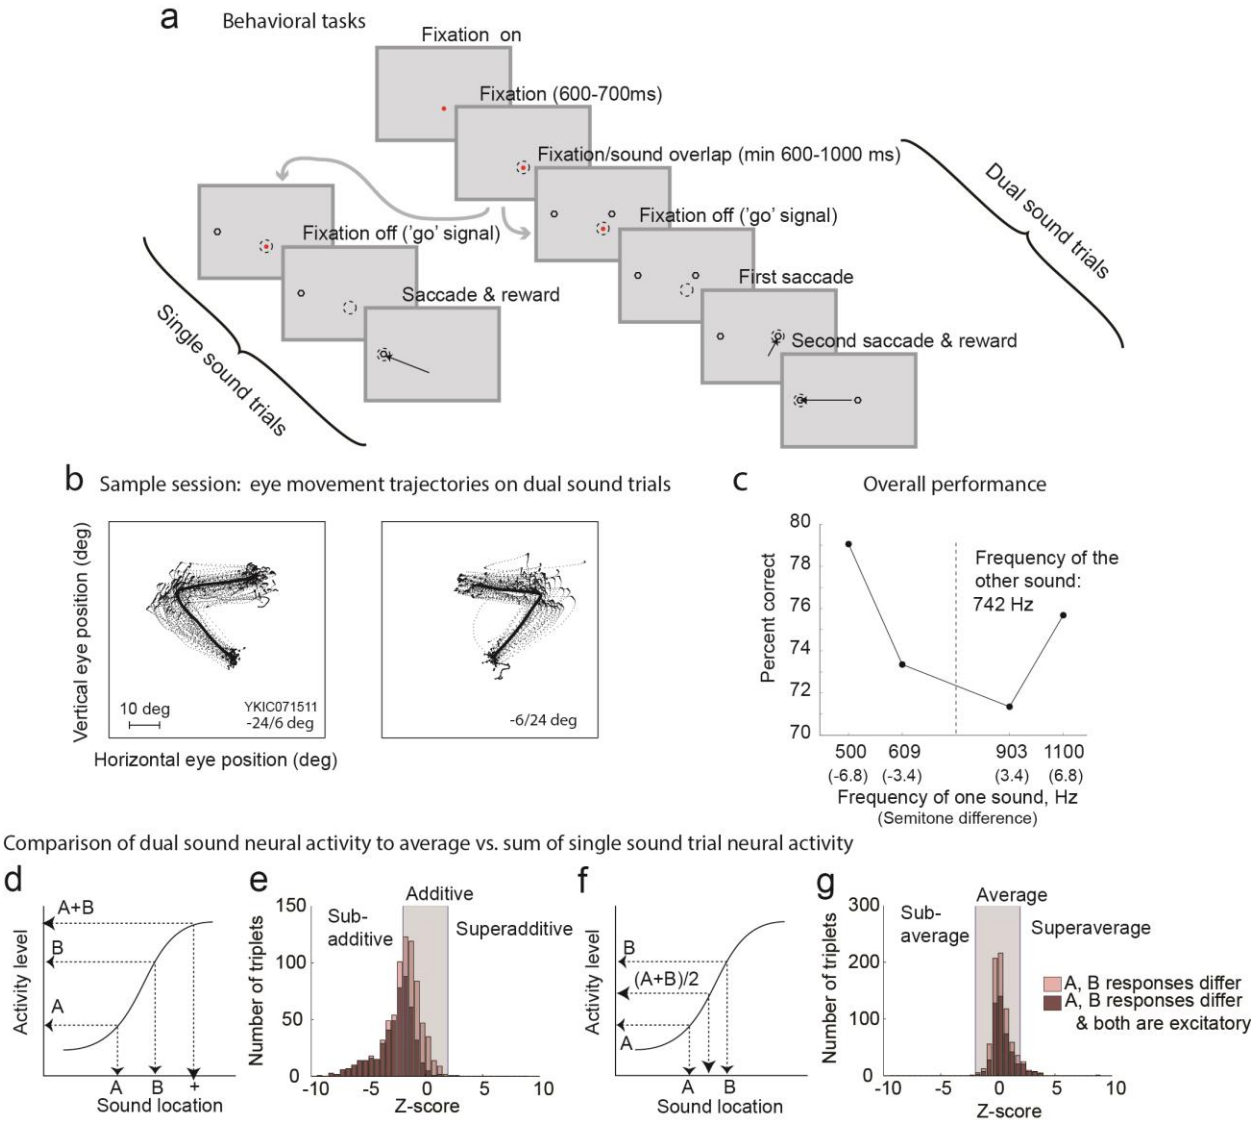

Figure 2. *Single- and dual-sound task, performance, and time-and-trial pooled neural activity. A. On dual-sound trials, monkeys made saccades to each of two simultaneous bandlimited noise sounds. Single-sound trials were similar but only required one saccade. B. Trajectories of eye movements on two sets of dual-sound trials for an example session (dash lines: individual correct trials; solid line: mean). Monkeys were permitted to look at the sounds in either order, but often showed stereotypical patterns based on their training history (see* Behavioral task and training*). C. Performance was better for larger frequency separations and was > 70% correct overall.* D-G. *Results of conventional analyses pooling across time and trials are inconsistent with summation (D-E), but are consistent with averaging (F-G). D, F. Schematic activity patterns of IC neurons in response to single sounds and predicted response for dual sounds if the neuron sums (D) or averages (F) inputs corresponding to individual sounds E, G. Observed Z-scores of activity on dual-sound trials differ from the sum (E), but correspond well to the average (G). The shaded areas indicate Z score values of +/- 1.96, or 95% confidence intervals. This analysis was conducted on "triplets" of single- and dual-sound trials with a given set of locations and frequencies, pooling across intensities. Triplets were included if the single-sound responses differed (light bars, two-tailed t-test, p<0.05, n=761); results were similar when single-sound responses were different and both excitatory (dark bars, one-tailed t-test, p<0.05, n=486). See Supplementary Figure 1 for a breakdown of conditions matched for the same signal levels on single- and dual-sound trials vs. signal levels adjusted to equate loudness on single- and dual-sound trials.*

If the monkeys can report the locations of two sounds presented simultaneously, it follows that their brains must preserve information about both sound items. To evaluate the neural basis of this, we focused on the IC because it lies comparatively early along the auditory pathway (a few synapses in from the periphery, and about two synapses prior to signals reaching auditory cortex) (27, 28) and because it is a nearly obligatory station along this pathway (29). Thus, preservation of information about both sound locations in the IC would appear to be required for performance of this task.

2.2. Time-and-trial pooled neural activity in the IC is consistent with an "average", but an average is inconsistent with behavior

Conventional analysis of spike data typically involves two simplifications: spikes are counted within a fairly long window of time, such as a few hundred milliseconds, and activity is pooled across trials for statistical analysis. If IC neurons multiplex signals related to each of the two sounds (arbitrarily dubbed “A” or “B” for the single-sound trials), then they might appear to show "averaging" responses on dual (or “AB”) trials when activity is pooled across time and across trials. But they should not appear to show "summation" responses, i.e. in which the responses on dual-sound trials resemble the sum of the responses exhibited on single-sound trials involving the component sounds. Such summation has been observed in some neural populations in areas such as primary visual cortex (30, 31), the hippocampus (32), or the superior colliculus (33) when multiple stimuli are presented.

To investigate whether responses to two sounds are more similar to the sum or the average of the two single-sound responses, we considered matched combinations of a particular pair of stimuli A and B presented alone or in combination. The set of stimulus A alone, stimulus B alone, and stimuli A and B in combination is referred to as a “*triplet*”, a term we will use throughout. Using an analysis similar to that of (33), dual-sound responses were converted to Z-scores relative to either the sum or the average of the corresponding single-sound responses (see Methods). Figure 2D-G shows that such trial-and-time-pooled responses more closely resemble averaging than summation: 93% of Z scores (N=761) were consistent with averaging (gray zone indicating +/-1.96 units of standard deviation) whereas far fewer, 55%, were consistent with summation. This was true even when both sound A and sound B evoked excitatory responses (dark bars). Findings were similar regardless of whether the signals delivered to the audio speakers were identical on dual and single-sound trials vs. when the signals were adjusted to equate loudness across single- vs. dual-sound trials (See Methods and Supplementary Figure 1). Consequently, in subsequent analyses we pooled across sound level.

However, such apparent averaging response patterns are inconsistent with the behavioral results: if the neurons truly responded at an average firing rate, then presumably the monkeys should respond to dual sounds as if there were only a single sound at the midpoint of the two sources (Figure 2F). Since monkeys can indicate the locations of *both* sounds (Figure 2B, C), multiplexing might provide a better explanation for so-called averaging response patterns.

2.3. Fluctuations in responses to dual sounds appear consistent with multiplexing at various time scales

2.3.1 Visualization. To determine whether neural activity fluctuates within and/or between trials, creating an overall averaging response but retaining information about each sound at distinct moments, we first sought to visualize the activity on individual trials. Figure 3 shows the activity of two example neurons on dual-sound trials compared to single-sound trials. The colored backgrounds illustrate the median and 25-75% quantiles of the activity on single-sound trials, in 50ms time bins. Superimposed on these backgrounds is the activity on individual trials. Individual single-sound (A alone, B alone) trials align well with their corresponding 25-75% quantiles, by definition (Figure 3A-B; E-F). But on dual-sound (AB) trials, for any given trial or time bin, some individual traces correspond well to one of the component sound's 25-75% quantiles, and on other trials or time bins they correspond well to the 25-75% quantiles of the other component sound. For the neuron in Figure 3CD, there are whole trials in which the activity appears to match that evoked by sound "A" alone and others in which it better corresponds to that evoked by sound "B" alone. For the neuron in Figure 3G, the firing pattern on dual-sound trials appears to switch back and forth between the levels observed for sounds A and B as the trial unfolds. In short, for these two examples, the activity on dual-sound AB trials does not appear to occur at a consistent value intermediate between those evoked on single-sound A and B trials, but can fluctuate between those levels at a range of time scales.

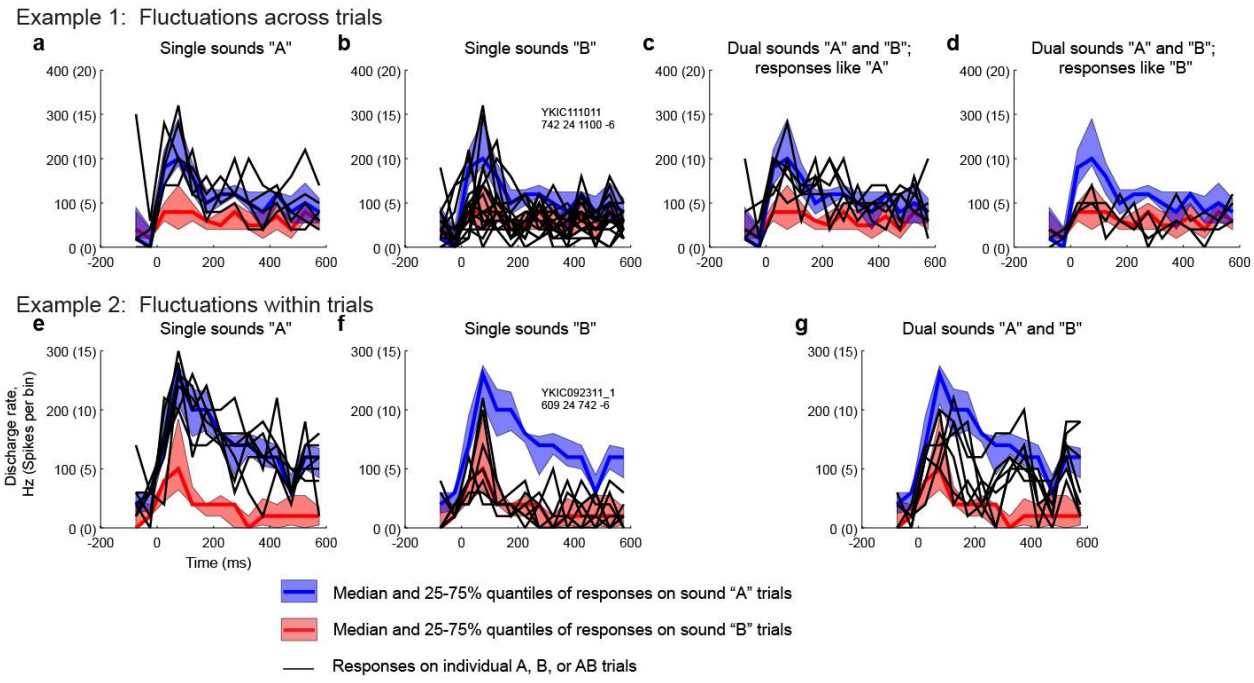

Figure 3. *Inspection of the individual trials of two IC neurons suggests that the overall intermediate firing rates observed on dual-sound trials may be due to alternating between firing rates corresponding to each component sound, fluctuating either across (A-D) or within trials (E-G). The red and blue shaded areas indicate the median and central 50% of the data on the single sound trials that make up the given triplet and are the same in panels A-D as well as E-G. The black traces superimposed on the shaded areas are the individual trials, for single sound and dual sound trials as indicated. For the neuron in A-D, individual traces on dual sound trials were classified based on the A vs. B assignment score (see Methods) and are plotted in two separate panels accordingly. For the neuron in E-G, the fluctuations occurred faster, within trials, and are plotted in the same panel (G) accordingly.*

We developed a series of statistical analyses to test for the presence of these various forms of alternation in firing rates. Several unknowns must be taken into consideration when testing for activity fluctuations. Specifically, the time scale, repeatability, and potential correlations across the neural population are uncertain. Accordingly, we sought to make minimal assumptions about the time scale at which neurons might alternate between encoding each stimulus, and we assumed that any such switching might vary from trial to trial and/or across time within a trial.

2.3.2. Statistical analysis of whole trial spike counts provides evidence consistent with multiplexing. If neurons alternate firing rates at the time scale of trials, as appears to be the case for the neuron in Figure 3A-D, then the spike counts from dual-sound responses should resemble a mixed bag of spike counts from each of the component single-sound responses. We statistically tested this hypothesis against other reasonable competing possibilities using the subset of triplets whose spike counts on single sound A and B trials could be well modeled by Poisson distributions with statistically different mean rates λ^A and λ^B (N=363, see methods for details).

The competing scenarios to describe the corresponding dual sound trials were:

(a) Mixture: The spike counts observed on individual trials are best described as having come from a weighted mixture of Poi(λ^A)and Poi(λ^B) (Figure 4A). This possibility is consistent with multiplexing across trials.

(b) Intermediate: A single Poisson distribution best describes the spike counts, and this Poisson has a rate λ^{AB} that is between λ^A and λ^B (Figure 4B). This possibility is consistent with either multiplexing at faster, sub-trial time scales or with true averaging/normalization.

(c) Outside: Again, a single Poisson, but the rate λ^{AB} is outside the range of λ^A and λ^B (i.e is greater than both or less than both; Figure 4C). Summation-type responses would be captured under this heading, as would inhibitory interactions.

(d) Single: A single Poisson describes the dual-sound trial spike counts, but the rate λ^{AB} is is equal to one of the single-sound rates λ^A or λ^B (Figure 4D). A winner- (or loser-)-take-all pattern would fit this category.

In summary, these four models capture the spectrum of possibilities at the whole-trial time scale. A Bayesian model comparison with default priors and intrinsic Bayes factor calculation was used to compute the posterior probabilities of the four hypotheses on a selected triplet given its neural data(34). Under the Bayesian inference paradigm, the hypothesis with the highest posterior probability (the winning hypothesis) can be interpreted as providing the best fit to the observed data for a given triplet.

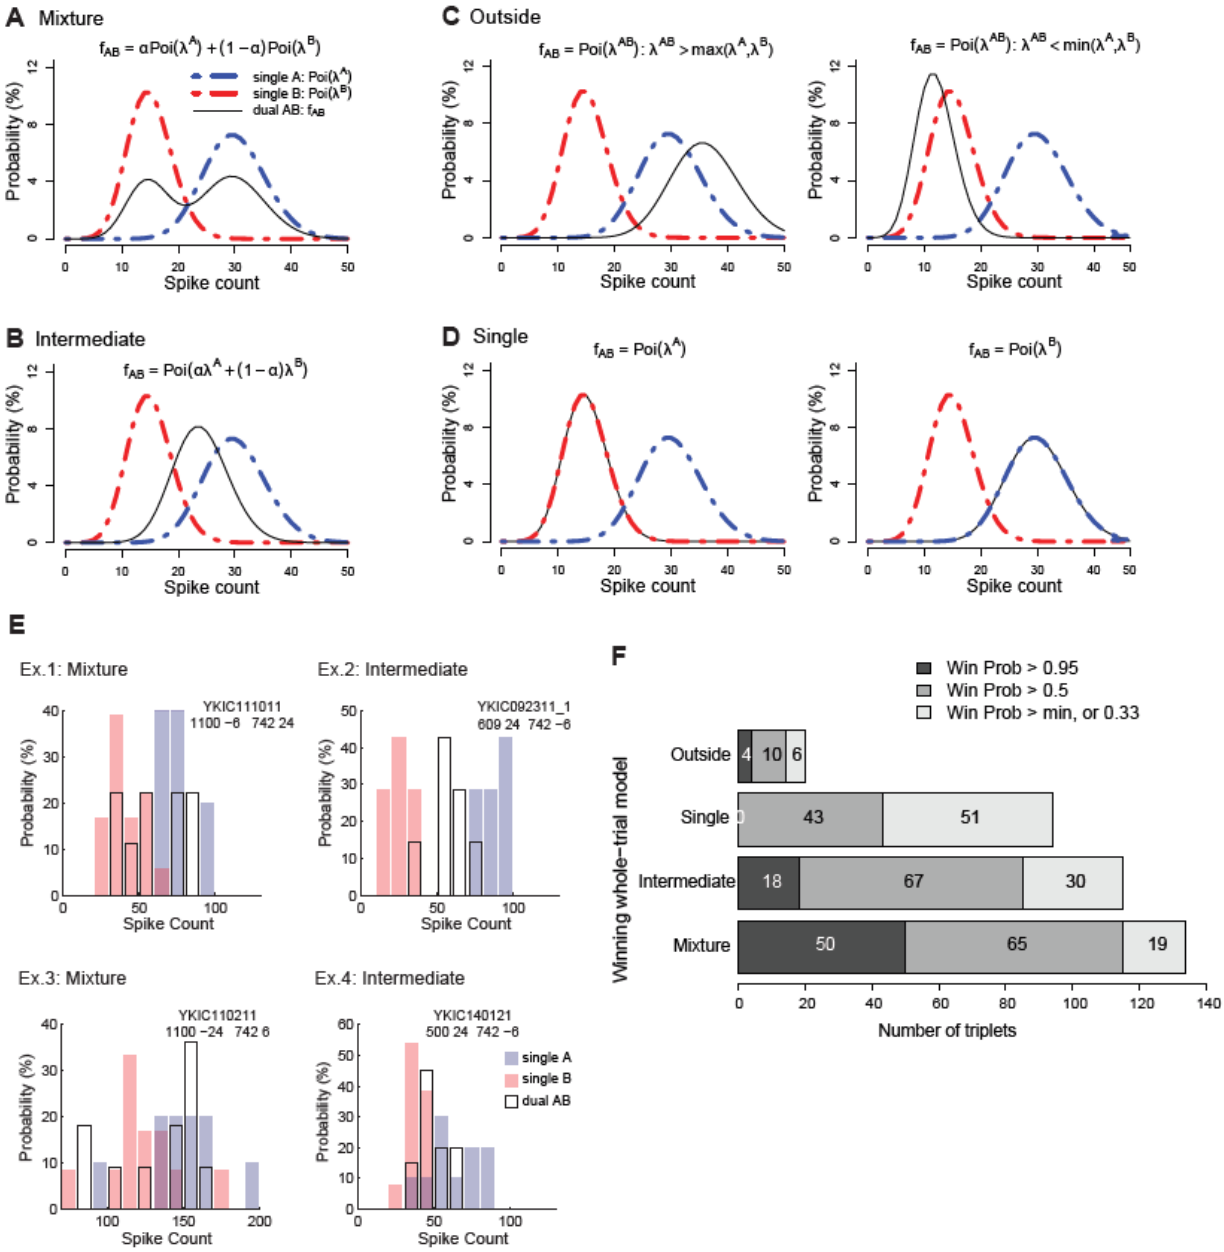

Figure 4. Whole-trial analysis. *Panels A to D show the four models that could describe the distribution of spike counts on individual dual-sound trials (0-600 or 0-1000 ms after sound onset, see Methods).* ***(A)*** mixture *of the Poisson distributions of spike counts for the component single-sound trials,* ***(B)*** intermediate *Poisson distribution, with rate between the rates of single sounds responses,* ***(C)*** outside, *Poisson distribution with rate larger or lower than the rates of single-sounds responses,* ***(D)*** single, *Poisson distribution with rate equal to one of the two single sound rates.* ***(E)*** *Four examples of spike count distributions for neurons/triplets classified as*

mixtures or intermediates. Responses to single sounds are in red and blue shades, to dual sounds are in black outline. Examples 1 and 3 were classified as having dual-sound spike counts drawn from a mixture of Poissons (winning probability >0.95, same as examples in Figure 3A-D and Figure 5B example 1 respectively). Examples 2 and 4 were classified as having dual-sound spike counts drawn from an intermediate Poisson (winning probability >0.95, example 2 is the same as examples in Fig3E-G and Figure 5B example 2; example 4 is the same as Figure 5B example 3). ***(F)*** *Population results of the whole-trial analysis. Shading indicates the confidence level of the assignment of an individual triplet to a winning model.*

For a sizeable portion of the triplets, the spike counts on dual-sound trials were better fit by a mixture of the single-sound Poisson distributions than by any single Poisson distribution (Figure 4F, bar labeled "mixture"). These response patterns are potentially consistent with time division multiplexing at the level of individual trials; the neurons illustrated in Figure 3A-D/Figure 4E example 1 and Figure 4E example 3 met these criteria. Of the 72 triplets in which one model had a winning probability >0.95, 50 or 69% were categorized as mixtures.

For the next largest category, the best fitting model involved a unique λ^{AB} between λ^{A} and λ^{B} (Figure 4F, bar labeled “intermediate”). These triplets are ambiguous: they could exhibit a true intermediate firing rate on the dual-sound trials, or they could simply show alternation at a time scale more rapid than individual trials (the neuron illustrated in Figure 3E-G/Figure 4E example 2 was classified as “intermediate”, as was Figure 4E example 4). Of the 72 triplets in which one model had a winning probability >0.95, 18 or ~25% were categorized this way.

The remaining triplets were categorized as “single”, or $\lambda^{AB} = \lambda^{A}$ or λ^{B} (a narrowly defined category that consequently did not produce any winning model probabilities >0.95) or “outside”, λ^{AB} greater or less than both λ^{A} and λ^{B}. “Single” can be thought of as a winner-take-all response pattern. “Outside” may be consistent with a modest degree of summation in the neural population, particularly as λ^{AB} was generally greater than both λ^{A} and λ^{B} in this subgroup.

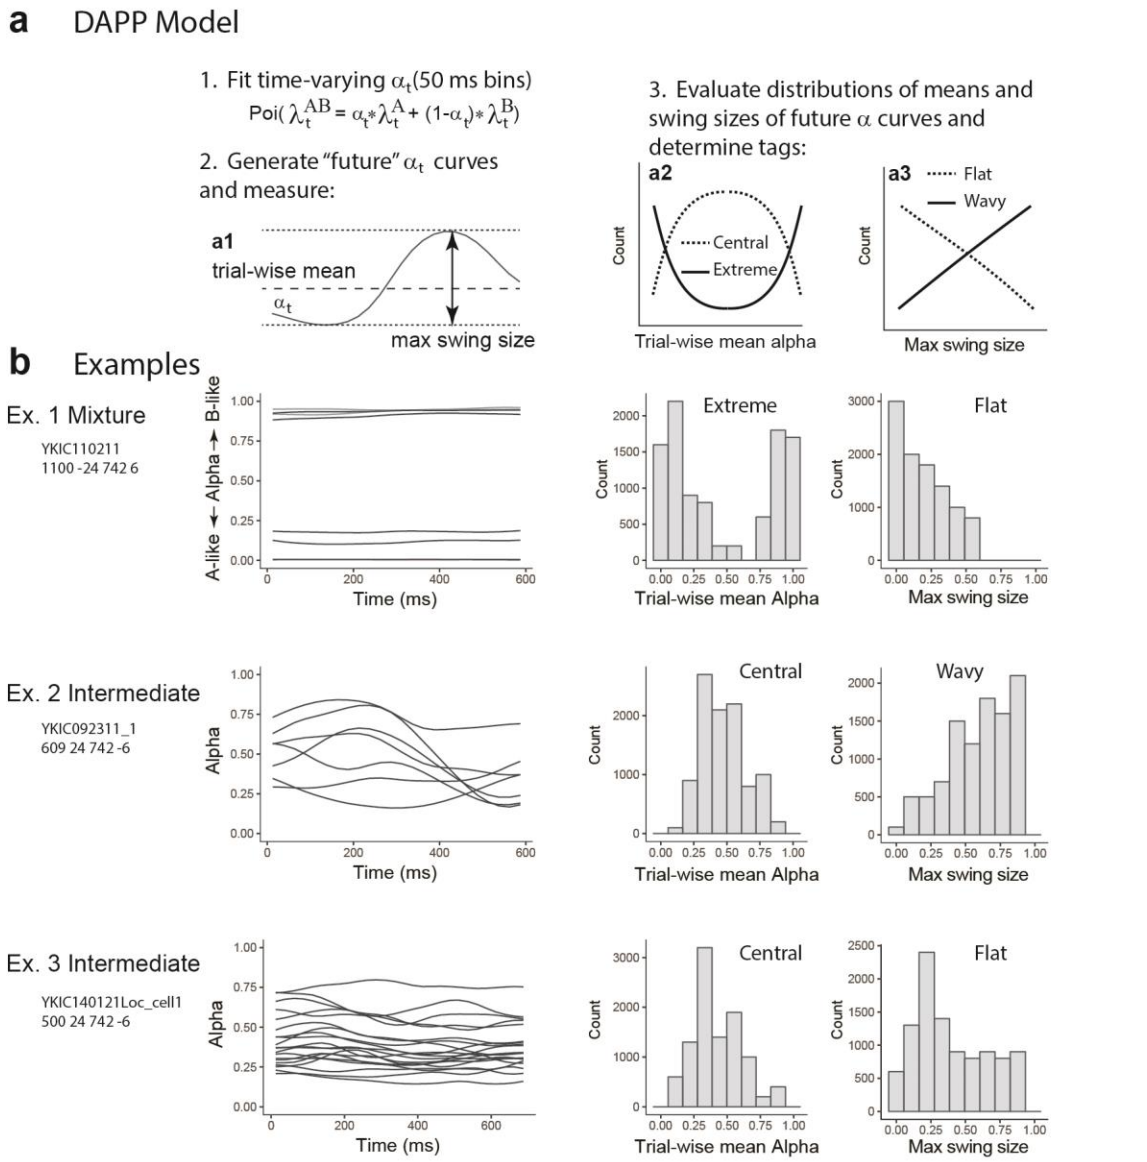
a DAPP Model
1. Fit time-varying α_t(50 ms bins)
Poi($\lambda_t^{AB} = \alpha_t * \lambda_t^A + (1-\alpha_t) * \lambda_t^B$)
2. Generate "future" α_t curves and measure:
a1
trial-wise mean
α_t
max swing size
3. Evaluate distributions of means and swing sizes of future α curves and determine tags:
a2
Central
Extreme
Count
Trial-wise mean alpha
a3
Flat
Wavy
Count
Max swing size
b Examples
Ex. 1 Mixture
YKIC110211
1100 -24 742 6
A-like ← Alpha → B-like
Time (ms)
Extreme
Flat
Count
Trial-wise mean Alpha
Max swing size
Ex. 2 Intermediate
YKIC092311_1
609 24 742 -6
Alpha
Time (ms)
Central
Wavy
Count
Trial-wise mean Alpha
Max swing size
Ex. 3 Intermediate
YKIC140121Loc_cell1
500 24 742 -6
Alpha
Time (ms)
Central
Flat
Count
Trial-wise mean Alpha
Max swing size

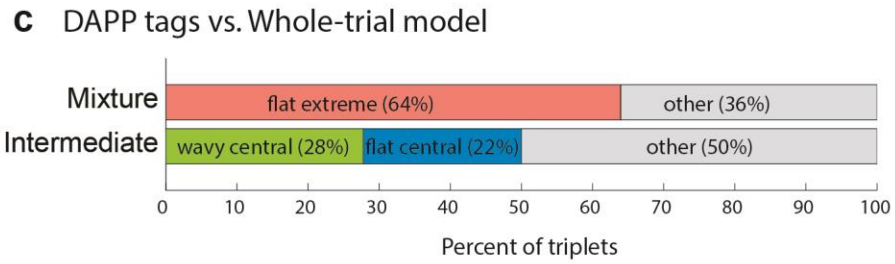
c DAPP tags vs. Whole-trial model
Mixture
flat extreme (64%)
other (36%)
Intermediate
wavy central (28%)
flat central (22%)
other (50%)
0 10 20 30 40 50 60 70 80 90 100
Percent of triplets

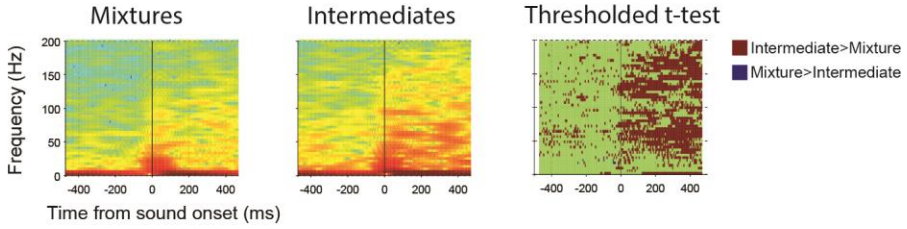
d LFP and Whole-trial model
Mixtures
Intermediates
Thresholded t-test
Intermediate>Mixture
Mixture>Intermediate
Frequency (Hz)
Time from sound onset (ms)

Figure 5. *Dynamic Admixture Point Process (DAPP) model rationale, results, and relationship to local field potential. A. The DAPP model fit smoothly time-varying weights capturing the relative contribution of A- and B-like response distributions to each AB dual sound trial. The dynamic tendencies of these curves were then used to generate projected alpha curves for hypothetical future draws from this distribution. The "waviness" and central tendencies were quantified by computing the max swing size and trial-wise mean for an individual trial draw from the distribution (panel A1). Low max swing sizes indicate "flat" curves and higher values indicate "wavy" ones (panel A2). Similarly, the distribution of trial-wise means could be bimodal ("extreme") or unimodal ("central")(panel A3). B. Fit alphas for three example neurons (Example 2 is the same as Figure 3E-G) and the distribution of trial-wise means and max swing sizes for future draws from the alpha curve generator. C. Triplets showed different patterns of tags, correlated with their classification at the whole-trial analysis. Specifically, triplets categorized as "Mixtures" (with a win probability > 0.95) tended to be tagged as "flat/extreme" (Example 1, same as Figure 4E example 3). Triplets categorized as "Intermediates" fell in two different main groups, "wavy-central" (Example 2, same as Fig 3E-G and Fig 4E example 2) and "flat-central" (Example 3, same as Fig 4E example 4). The "wavy-central" subgroup is consistent with multiplexing at the within-trial time scale. Triplets with a "skewed" subclassification (see Methods) were excluded from this graph; a complete listing of all the tag combinations is presented in Supplementary Table 1. See Supplementary Figures 2 and 3 for additional analyses. D. Average local field potentials also showed different patterns across the whole trial "mixture" vs "intermediate" classification, with the "intermediate" group tending to show greater power at a range of oscillatory frequencies. The third graph shows a thresholded (p<0.05) depiction of whether these differences were statistically significant according to a two-tailed t-test for each time point and frequency combination.*

<u>2.3.3 Statistical analysis of within-trial temporal patterns suggests presence of both fast multiplexing and stable averaging</u>. We next evaluated whether firing patterns fluctuated or remained stable across time within a trial. In particular, might triplets categorized as "intermediate" in the whole trial analysis show evidence of fluctuating activity on a faster time scale?

This is a more challenging statistical question, and required development of a novel statistical approach. We focused on the same triplets selected above, and analyzed the temporal patterns of the spike trains on dual sound trials under a novel Dynamic Admixture of Poisson Process model (Methods). For each triplet, spike trains from individual single sound trials were assumed to be independent realizations of a nonhomogeneous Poisson process with unknown time-dependent firing rates ($\lambda^A(t)$ for sound A and $\lambda^B(t)$ for sound B). To assess how individual time-varying dual sound responses related to single sound responses, each trial from the dual sound condition was assumed to be a realization of a Poisson process but with its own firing rate function $\lambda(t)$, modeled as an unknown weighted average of the two single sound firing rate functions $\lambda^{AB}(t) = \alpha(t)\lambda^A(t) + \big(1 - \alpha(t)\big)\lambda^B(t)$. The weight function $\alpha(t)$, unique to each dual sound trial, quantified the potentially time varying relative contribution of sound A on that trial at time *t*, while $1 - \alpha(t)$ quantified the complementary contribution of sound B (Figure 5A).

The dynamics of the $\alpha(t)$ function characterize the dynamics of each dual-sound trial. A constant $\alpha(t)$ function indicates that the contributions of the two single sounds to the dual sound responses were constant across an individual trial. For example, a dual-sound response that matched (on a given trial) the response evoked by single sound A would be characterized by a constant $\alpha(t)$ at a value close to 0 or 1. Conversely, a dual-sound response encoding each of the single sounds at different times during the course of a single trial would be characterized by a wavy $\alpha(t)$ curve that fluctuated between 0 and 1.

For each selected triplet, a Bayesian inference technique was used to predict the $\alpha(t)$ curves the corresponding cell was likely to produce on future dual sound AB trials. Each predicted curve was summarized by two features: its time average over the response period of a given trial and its maximum swing size, that is, the difference between its highest peak and lowest trough on that trial (see Methods, Figure 5A-B). The triplet was then subjected to a two-way classification based on the distribution of these two features over the predicted curves (DAPP "tags", Figure 5A-B). The triplet was categorized as "wavy" vs. "flat" depending on whether the distribution of the maximum swing size peaked at high or low values, and as "central" vs "extreme" according to whether the distribution of the time average $\alpha(t)$ had a peak close to 0.5 or had one to two peaks at the extreme values of 0 and 1. In addition to this main

classification scheme, triplets were subcategorized as exhibiting "symmetric" or "skewed" response patterns, reflecting whether the $\alpha(t)$ curves reflected roughly equal contributions from the stimulus A and stimulus B response patterns or whether one or the other tended to dominate.

The DAPP tags for waviness/flatness and centrality/extremity confirmed and extended the results of the whole-trial analysis. Triplets categorized as "intermediate" in the whole trial analysis showed a different distribution of tags as compared to those categorized as "mixtures". "Mixture" triplets tended to be classified as showing "flat" single sound contributions, centering around "extreme" rather than "central" values of the time average of α(t) (Figure 5C), and the distribution of the these average α(t) values tended to be either symmetric, i.e. or unlabeled with regard to symmetry (Supplementary Table 1). In short, the dynamics of the "mixture" responses were consistent with fluctuations at the level of whole trials. In contrast, "intermediate" triplets showed a combination of two types of labelling patterns relevant to our hypothesis. Some showed flat firing at a central (and symmetric) intermediate value, indicating stable firing at roughly the average of the responses evoked by each sound separately. Such a firing pattern is consistent with some form of normalization occurring in this subpopulation. However, there were also triplets that showed wavy, i.e. fluctuating response patterns symmetric around a central value. This type of response pattern suggests that under some circumstances, neurons can "switch" relatively rapidly between a response pattern consistent with one stimulus vs the other on dual stimulus trials.

2.3.4. Local field potential shows greater oscillatory energy associated with "intermediate" classification. Consistent with this statistical evidence for activity fluctuations at the sub-trial timescale in the "intermediate" category, we also found that the local field potential (LFP) at such sites showed greater oscillatory activity. Figure 5D shows the average LFP power spectrum for dual trials of triplets categorized as "mixtures" vs. those categorized as "intermediates" and their statistical comparison (lower panel, two-tailed t-test between the LFP power spectrum of dual trials classified as Intermediate and that of dual trials classified as mixtures, for each time point and frequency combination). The LFP for intermediate sites showed higher energy across a range of frequencies, including frequencies well above the 20 Hz (50 ms) frequency range that we were able to evaluate at the spike-count single unit level

2.4 Activity fluctuations appear coordinated across the neural population and predict behavior

We next considered the question of whether and how activity fluctuations are coordinated across the neural population, in two ways: (1) by evaluating activity correlations across time within trials between pairs of simultaneously recorded neurons, and (2) by evaluating whether the state of the local field potential prior to sound onset predicts between-trial fluctuations in activity (e.g. 35, 36).

<u>2.4.1 Neural pairs show positive within trial correlations:</u> To evaluate correlations in *within-trial* switching patterns, we evaluated the neuron-to-neuron correlation between how "A-like" vs. how "B-like" the responses were on a time bin by time bin basis on individual trials, in a total of 91 pairs of triplet conditions from 34 pairs of neurons recorded simultaneously (from among the 363 triplets used for the previous analyses). For each 50 ms bin of a dual-sound trial in a given triplet, we assigned a probability score between 0 and 1 that the spike count in the bin was drawn from the Poisson distribution with rate equaling the bin's sound A rate, and the complementary probability to the same being drawn from the Poisson distribution with rate equaling the bin's sound B rate (Figure 6A; see Methods: A vs. B assignment scores). We normalized these probabilities by converting them to Z-scores within a given time bin but across trials, to minimize the contribution of shared correlations due to stimulus responsiveness or changes in motivational state across time (37). We then calculated the neuron-to-neuron correlation coefficients between the normalized assignment scores across the set of time bins within each trial, i.e. one correlation coefficient value estimated per trial. This analysis is conceptually similar to conventional cross-correlation analysis of spike trains in neural pairs, but does not focus on precise timing of spikes or the relative latency between them (38, 39).

Generally, the observed correlations were positive, indicating that the activity was coordinated within the neural population. Figure 6 illustrates analysis of the dual-sound trials for a particular triplet in an example pair of neurons (A), and the distribution of the mean neuron-to-neuron correlations in the population for all the triplets' dual-sound conditions (B). The distribution of mean correlation coefficients was skewed positive (t-test, $p = 6.8 \times 10^{-6}$). Similar results were obtained when the raw spike counts were analyzed rather than the assignment scores (Supplementary Figure 4). This was the case even though we included triplets that were not

categorized as showing "wavy" behavior in the DAPP analysis. It may be that coordinated activity fluctuations occur in more neurons than those that met our statistical criteria.

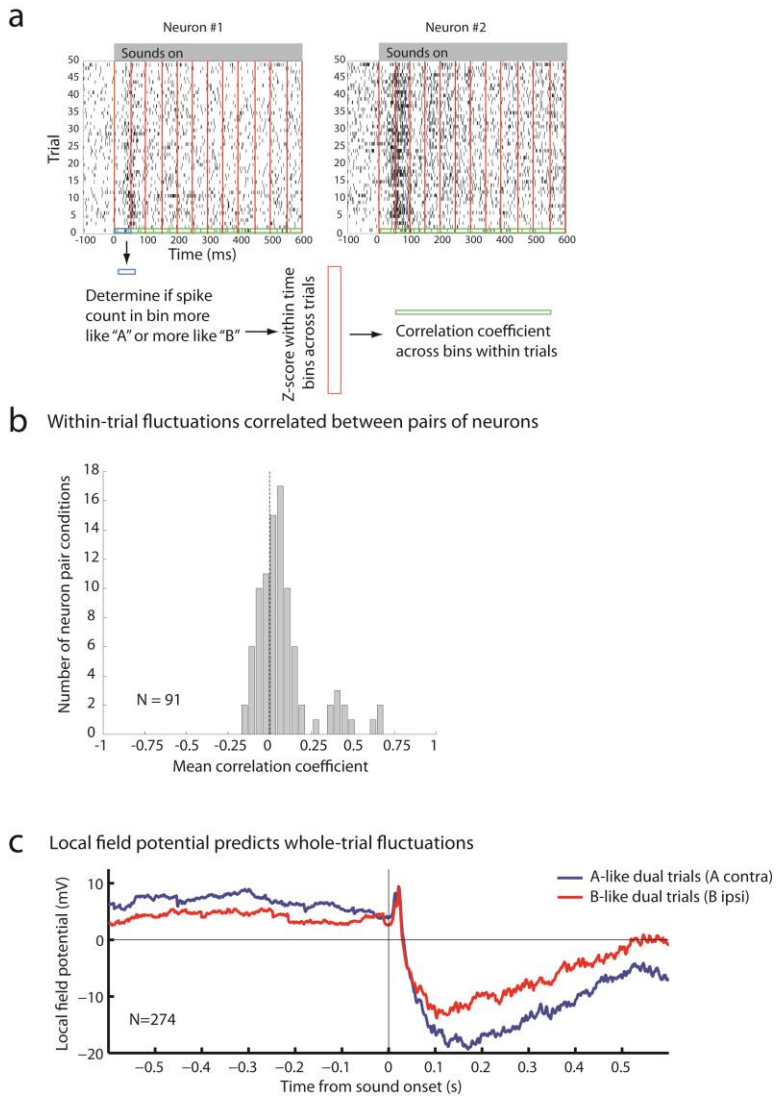

Figure 6. *Within-trial fluctuations are correlated between pairs of neurons and whole-trail fluctuations can be predicted by the state of the local field potential at sound onset. A. Pairs of neurons recorded simultaneously tended to show positive correlations with each other. Raster plots of two neurons recorded simultaneously; trials shown are for a particular set of dual-sound conditions. The spike count in a given 50 ms time bin, trial, and member of the neuron pair for a given set of dual-sound conditions was evaluated to determine if it was more similar to the spikes evoked during that bin on the corresponding sound "A" alone or "B" alone trials (blue box).*

These A vs. B assignment probabilities were then converted to a Z-score based on the mean and standard deviation of the assignment probabilities in that time bin on the other trials that involved the same stimulus conditions (red box). A correlation coefficient between the set of Z score values for a given trial between the pair of simultaneously recorded neurons was then calculated (green box). B. Across the population (N=91 conditions in the 34 pairs of neurons; triplets were included if their single sound response distributions were well-separated Poissons), the distribution of mean correlation coefficients tended to be positive (t-test comparing the mean correlation coefficients to zero; $p = 6.8 \times 10^{-6}$). C. The state of the local field potential prior to sound onset is predictive of whole-trial fluctuations in spiking activity. (Average of the LFP during dual-sound trials from 274 triplets at 87 sites; triplets were included if the single sound distributions of spike counts were well-separated Poissons). For this analysis, the single contralateral sound was dubbed "A" and the single ipsilateral "B". The LFP on each dual-sound trial was assigned to A-like or B-like categories based on the spike count of the corresponding single unit data during 0-600 ms after sound onset (see Methods: A vs. B assignment. N=1902 contra-like trials and N=1618 ipsi-like trials).

2.4.2 State of LFP prior to sound onset predicts between-trial fluctuations: To determine whether the state of the local field potential prior to sound onset predicts between-trial fluctuations in activity, we analyzed the LFP data recorded simultaneously with single unit spiking data. We combined data across triplets, creating two "bags" of trials based on whether the whole-trial spike count on a given dual-sound trial more closely resembled the responses evoked by sound A alone (where A is the contralateral sound) or sound B alone (see Methods: A vs. B assignment scores). Figure 6C shows the average LFP for the two groups of dual-sound trials. We quantified differences between these two groups with a t-test in the 600ms windows *before* and *after* sound onset (each trial contributed one mean LFP value in each time window). As expected, the LFP signals statistically differed after sound onset in these two trial groupings (red vs. blue traces, time period 0-600 ms, p-val = 1.0474×10^{-05}). But the LFP signals also differed *prior* to sound onset (p-val = 0.0064), suggesting that the state of activity in the local network surrounding an individual neuron at the time of sound onset is predictive of whether the neuron "encodes" the contra-lateral or the ipsi-lateral sound on that particular trial.

2.4.3 Neural fluctuations predict behavior. If fluctuations in neural activity are coordinated across the population, and if one particular stimulus dominates the representation at any given instant, it follows that there should be a relationship between trial-by-trial variability in neural activity and behavior. Accordingly, we investigated whether the activity on individual trials predicted whether the monkey would look first to sound “A” or sound “B” on that trial. As noted in the Methods, we trained the monkeys on sequential sounds first and this training strategy tended to promote performing the task in a stereotyped sequence. Partway through neural data collection, we provided monkey Y with additional training on the non-sequential task, after which that monkey began displaying less stereotypical behavior and sometimes saccaded first to A and sometimes first to B for a given AB dual sound combination (see Figure 7A for example). We then analyzed recording sessions after this training (N=73 triplets) and we found that at both the whole trial and sub-trial time scales, the activity of individual neurons was predictive of what saccade sequence the monkey would choose on that particular trial. Specifically, the average dual sound AB assignment score for a given triplet was computed separately for trials in which the first saccade was toward A vs. toward B. The average scores statistically differed between the two groups of dual-sound trials (t-test, pval = 5 x 10^-9, Figure 7B) and in the expected direction, with more A-like scores occurring on trials in which the monkey looked at A first. This relationship was also present when looking at finer, 50 ms bin time scales (Figure 7C).

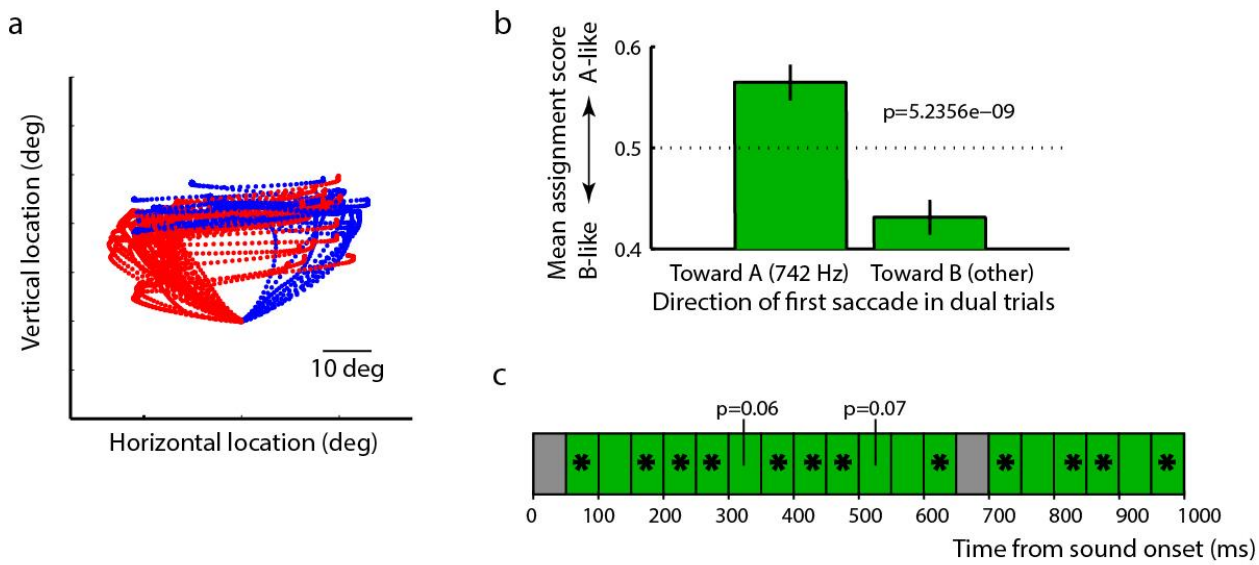

Figure 7. *The target of the first saccade on dual sound trials is predicted by the spike count before that saccade. (A) Eye trajectories during dual-sound trials to the same pair of single sounds (one triplet). The traces are color-coded based on which of the two sounds the monkey looked at first in the response sequence. For clarity, all traces are aligned on a common starting position despite some variation in fixation accuracy. (B) The average assignment score of trials in which the monkey looked at sound A first is more A-like than that of trials in which the monkey looked at sound B first. SEM are indicated with bars. (C). This relationship between assignment score and first saccade target was also evident at the scale of 50 ms bins (green = positive correlation; * indicate p<0.05 for t-test of assignment score on A-first vs. B-first trials).*

3. DISCUSSION

Our results show that the activity patterns of IC neurons fluctuate, and that these fluctuations may be consistent with encoding of multiple items in the same processing channels (i.e. the set of neural spike trains occurring in the IC). The time scale of these fluctuations ranges from the level of individual trials down to at least 50 ms bins within a trial. The fluctuations are positively correlated across pairs of neurons (at least, those recorded within the IC on a given side of the brain), are reflective of the state of local field potentials at the time of sound onset, and are predictive of the behavioral response to follow.

A strength of our approach is that we do not simply assess unimodality/bimodality (e.g. (10)) but anchor our consideration of the dual stimulus response distributions to the observed distributions for single stimulus trials. However, there are several limitations to the present statistical approach. First, the analyses could only be conducted on a subset of the data, requiring a good fit of a Poisson distribution to the single-sound trials and adequate separation of the responses on those trials. For the moment, it is unknown whether any of the excluded data exhibit meaningful response fluctuations. In principle, the modeling approach can be extended to other types of response distributions which should reduce the amount of data that is excluded. Second, the range of time scales at which fluctuations occur is unknown. Fluctuations that occur faster than the 50 ms bin time scale used for the DAPP model would likely have been (erroneously) categorized as flat-central. Third, our statistical approach based on the DAPP model involves a categorization step that summarizes the dominant features of a triplet. If a neuron sometimes behaves as a "flat-extreme" type and sometimes as a "wavy-central" type for a given triplet of conditions, it would likely be categorized as ambiguous. In other words, even though the DAPP model can pick up composite response patterns, the results we present ignore the existence of any such patterns.

The observed fluctuations have broad implications because they provide a novel account linking a number of other well-known aspects of brain function under a common explanation. First, it is widely recognized that neural firing patterns are highly variable. This variability is often thought to reflect some fundamental inability of neurons to code information accurately. Here, we suggest that some of this variability may actually reflect interleaved periods of potentially quite accurate coding of different items. What else individual neurons may commonly be coding for in experiments involving presentation of only one stimulus at a time is not known, but possibilities include stimuli not deliberately presented by the experimenter, memories of previous stimuli, or mental imagery as suggested by the theory of embodied cognition (40). In the present study, we were able to demonstrate signal in these fluctuations by virtue of statistical tests comparing each of the trial types in A-B-AB triplets, but it may be the case that fluctuations were occurring in the single stimulus trials as well. We could not test this because our analysis required having as benchmarks the response distributions corresponding to the potentially encoded items.

Second, as a concept, multiplexing provides insight into why limitations in certain types of cognition exist. Working memory capacity is limited; attention filters stimuli to allow in depth processing of a selected set of items. These limitations may stem from using the same population of neurons for each attended or remembered item. If this is the case, then the puzzle becomes why these limits are often greater than one. Multiplexing suggests that cycling between different items across time allows evading what might otherwise be a one-item limit (2). Here, we investigated only two time scales, 50 ms and whole trials. Future work will be needed to more fully explore the time scales on which this occurs and to tie the resulting information on duty cycle to perceptual capacity.

Third, brain oscillations are ubiquitous, have been linked specifically to attentional and memory processes (36, e.g. 41, see also 42), and have been suggested as indicating multiplexing (2-8). Oscillations indicate that neural activity fluctuates, although they capture only the portion of such fluctuation that is coordinated across the underlying neural population and is regular in time. It remains to be determined to what degree oscillations in field potentials reflect the activity of neural circuits that control such temporal coordination in other neural populations vs. the activity of the neural circuits subject to the effects of such coordination. In a highly interconnected system such as the brain, both are likely to occur.

In the case of our particular experimental paradigm, several additional questions arise. How do signals related to different items come to be multiplexed? Are they later de-multiplexed? If so, how?

To some degree, sounds are multiplexed in the world. That is, the sound waves from multiple sources sum in the world and are never purely distinct from one another. The air pressure waves arriving at each ear reflect the combined contribution of all sound sources. However, if the IC's neural fluctuations were driven by the sound signals arriving at the ears, then individual neurons should always respond the same way on every trial, and they do not. Instead, it seems likely that the externally-multiplexed sound waves interact with neural circuit states at the time that the incoming signal arrives to govern how individual neurons respond on a moment by moment basis.

Where and how signals may be de-multiplexed critically depends on the nature of the representation to which a de-multiplexed output could be written. In barn owls, which have maps of auditory space, the coding bottleneck intrinsic to meter/rate coding does not occur, and two

sounds produce two separate active populations (11-14). Such distinct peaks suggest that the multiplexed-in-the-air signals have been de-multiplexed and segregated into two hills of activity.

In primates and several other mammals, neural representations of space employ meters (rate codes) rather than maps throughout the pathway from sound input to eye movement output, as far as is currently known (15-22, 43). This is the case even at the level of the superior colliculus (44), which has a well-deserved reputation for mapping when activity is evoked by non-auditory stimuli (45, 46).

Given that different types of codes exist in different species, and given that coding format is not known in all the circumstances in which multiplexing might apply (e.g. attention, working memory), we developed two different models to illustrate a range of different de-multiplexing possibilities (Figure 8) based on the nature of the recipient representation. In the first (Figure 8A), a multiplexed signal in a meter is converted into two hills of activity in a map, using a basic architecture involving graded thresholds and inhibitory interneurons suggested previously (47). Adding an integration mechanism such as local positive feedback loops would then serve to latch activity "on" at the appropriate locations in the map, producing a more sustained firing pattern. No clock signal is necessary for this model.

In the second model (Figure 8B), there are multiple output channels, each capable of encoding one item. An oscillating circuit that knows about the timing of the input gates signals to each output channel at the appropriate moments. As in the first model, a local positive feedback mechanism acts to sustain the activity during the gaps in the input. This model thus retains the efficient coding format of a meter but requires a controlling signal with knowledge of when to latch input flow through to each output channel. In our data, it is possible that within-trial fluctuating units lie at the input stage of such a circuit, and that between-trial fluctuating units actually lie at the output stage. A given unit might be allocated to either the "A" or the "B" pools based on state of the network (as detected by the LFP measurements) on different trials.

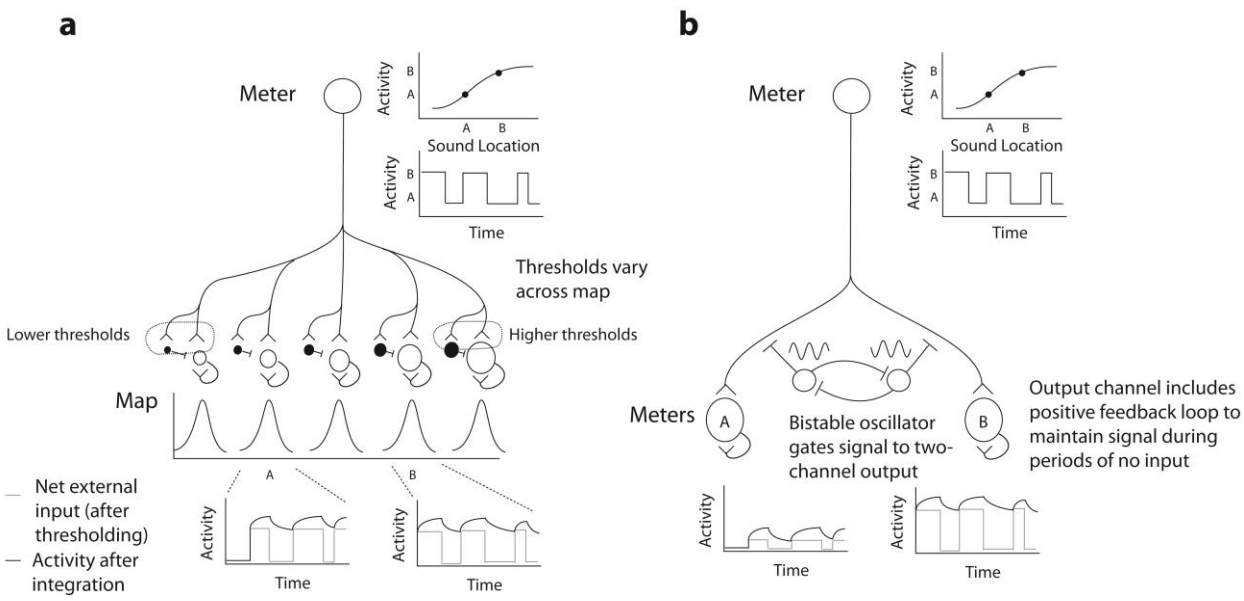

Figure 8. *Two possible mechanisms for de-multiplexing a fluctuating signal. A clock signal that knows about coding transitions is not necessarily needed if signals are read out into a map, but is required if signals are retained in a meter or rate-coded format.*

An important unresolved question posed by our study is whether multiplexing may be a general mechanism that is commonly at play to enhance the total processing power of the brain. The statistical tools developed here can be applied to any "triplet" data. Additional studies with both single stimulus conditions, to define the distributions of signals, and dual stimulus conditions, to evaluate fluctuations between membership in those distributions, will be important for delineating the extent of this phenomenon. Digging under the hood of the time-and-trial pooled activity to look at activity patterns on a moment by moment basis will be essential to advancing our understanding of how the brain operates dynamically to maximize its processing power.

ACKNOWLEDGMENTS:

We are grateful for expert technical assistance from Jessi Cruger, Karen Waterstradt, Christie Holmes, Stephanie Schlebusch, Tom Heil, and Eddie Ryklin. We have benefitted from thoughtful discussions with Michael Lindon, Winrich Freiwald, Liz Romanski, Stephen Lisberger, Marty Woldorff, Daniel Pages, David Bulkin, Kurtis Gruters, Bryce Gessell, Luke

Farrell, David Murphy, and Akinori Ebihara. We thank Bao Tran-Phu, Will Hyung, Stephen Spear, Francesca Tomasi, and Ashley Wilson for assistance with animal training and/or recordings. Financial support for the research was provided by the National Science Foundation (0924750) to JMG and the National Institutes of Health (5R01DC013906-02) to ST and JMG.

REFERENCES AND NOTES

1. Meister M (1996) Multineuronal codes in retinal signaling. *Proc Natl Acad Sci U S A* 93(2):609-614.
2. Lisman JE & Idiart MA (1995) Storage of 7 +/- 2 short-term memories in oscillatory subcycles. *Science* 267(5203):1512-1515.
3. Hoppensteadt FC & Izhikevich EM (1998) Thalamo-cortical interactions modeled by weakly connected oscillators: could the brain use FM radio principles? *Bio Systems* 48(1-3):85-94.
4. Akam T & Kullmann DM (2014) Oscillatory multiplexing of population codes for selective communication in the mammalian brain. *Nat Rev Neurosci* 15(2):111-122.
5. Tallon-Baudry C (2009) The roles of gamma-band oscillatory synchrony in human visual cognition. *Front Biosci (Landmark Ed)* 14:321-332.
6. Lisman JE & Jensen O (2013) The theta-gamma neural code. *Neuron* 77(6):1002-1016.
7. Siegel M, Warden MR, & Miller EK (2009) Phase-dependent neuronal coding of objects in short-term memory. *Proc Natl Acad Sci U S A* 106(50):21341-21346.
8. McLelland D & VanRullen R (2016) Theta-Gamma Coding Meets Communication-through-Coherence: Neuronal Oscillatory Multiplexing Theories Reconciled. *PLoS computational biology* 12(10):e1005162.
9. Jezek K, Henriksen EJ, Treves A, Moser EI, & Moser MB (2011) Theta-paced flickering between place-cell maps in the hippocampus. *Nature* 478(7368):246-249.
10. Li K, *et al.* (2016) Neurons in Primate Visual Cortex Alternate between Responses to Multiple Stimuli in Their Receptive Field. *Front Comput Neurosci* 10:141.
11. Baxter CS, Nelson BS, & Takahashi TT (2013) The role of envelope shape in the localization of multiple sound sources and echoes in the barn owl. *J Neurophysiol* 109(4):924-931.
12. Takahashi TT, *et al.* (2008) Object localization in cluttered acoustical environments. *Biol Cybern* 98(6):579-586.
13. Keller CH & Takahashi TT (2005) Localization and identification of concurrent sounds in the owl's auditory space map. *J Neurosci* 25(45):10446-10461.
14. Takahashi TT & Keller CH (1994) Representation of multiple sound sources in the owl's auditory space map. *Journal of Neuroscience.* 14(8):4780-4793.
15. Groh JM, Kelly KA, & Underhill AM (2003) A monotonic code for sound azimuth in primate inferior colliculus. *Journal of Cognitive Neuroscience* 15:1217-1231.
16. McAlpine D & Grothe B (2003) Sound localization and delay lines--do mammals fit the model? *Trends Neurosci* 26(7):347-350.
17. Grothe B, Pecka M, & McAlpine D (2010) Mechanisms of sound localization in mammals. *Physiol Rev* 90(3):983-1012.
18. Salminen NH, May PJ, Alku P, & Tiitinen H (2009) A population rate code of auditory space in the human cortex. *PLoS One* 4(10):e7600.
19. Salminen NH, Tiitinen H, Yrttiaho S, & May PJ (2010) The neural code for interaural time difference in human auditory cortex. *J Acoust Soc Am* 127(2):EL60-65.
20. Groh JM (2014) Moving with maps and meters. *Making space: How the brain knows where things are*, (Harvard University Press), pp 143-159.
21. Werner-Reiss U & Groh JM (2008) A rate code for sound azimuth in monkey auditory cortex: implications for human neuroimaging studies. *J. Neurosci.* 28(14):3747-3758.
22. Woods TM, Lopez SE, Long JH, Rahman JE, & Recanzone GH (2006) Effects of stimulus azimuth and intensity on the single-neuron activity in the auditory cortex of the alert macaque monkey. *J Neurophysiol* 96(6):3323-3337.
23. Bulkin DA & Groh JM (2011) Systematic mapping of the monkey inferior colliculus reveals enhanced low frequency sound representation. *J. Neurophysiol.* 105(4):1785-1797.

24. Perrott DR (1984) Discrimination of the spatial distribution of concurrently active sound sources: some experiments with stereophonic arrays. *J Acoust Soc Am* 76(6):1704-1712.
25. Perrott DR (1984) Concurrent minimum audible angle: a re-examination of the concept of auditory spatial acuity. *J Acoust Soc Am* 75(4):1201-1206.
26. Best V, van Schaik A, & Carlile S (2004) Separation of concurrent broadband sound sources by human listeners. *J Acoust Soc Am* 115(1):324-336.
27. Adams JC (1979) Ascending projections to the inferior colliculus. *J Comp Neurol* 183(3):519-538.
28. Moore RY & Goldberg JM (1963) Ascending projections of the inferior colliculus in the cat. *J Comp Neurol* 121:109-135.
29. Aitkin LM & Phillips SC (1984) Is the inferior colliculus an obligatory relay in the cat auditory system? *Neuroscience Letters* 44(3):259-264.
30. Carandini M, Heeger DJ, & Movshon JA (1997) Linearity and normalization in simple cells of the macaque primary visual cortex. *J Neurosci* 17(21):8621-8644.
31. Movshon JA, Thompson ID, & Tolhurst DJ (1978) Spatial summation in the receptive fields of simple cells in the cat's striate cortex. *J Physiol* 283:53-77.
32. Cash S & Yuste R (1999) Linear summation of excitatory inputs by CA1 pyramidal neurons. *Neuron* 22(2):383-394.
33. Alvarado JC, Vaughan JW, Stanford TR, & Stein BE (2007) Multisensory versus unisensory integration: contrasting modes in the superior colliculus. *J Neurophysiol* 97(5):3193-3205.
34. Berger JO & Pericchi LR (1996) The Intrinsic Bayes Factor for Model Selection and Prediction. *J Am Stat Assoc* 91(433):109-122.
35. Haslinger R, Ulbert I, Moore CI, Brown EN, & Devor A (2006) Analysis of LFP phase predicts sensory response of barrel cortex. *J Neurophysiol* 96(3):1658-1663.
36. Vanrullen R, Busch NA, Drewes J, & Dubois J (2011) Ongoing EEG Phase as a Trial-by-Trial Predictor of Perceptual and Attentional Variability. *Front Psychol* 2:60.
37. Metzger RR, Greene NT, Porter KK, & Groh JM (2006) Effects of reward and behavioral context on neural activity in the primate inferior colliculus. *J. Neurosci.* 26(28):7468-7476.
38. Gerstein GL, Perkel DH, & Subramian KN (1978) Identification of functionally related neural assemblies. *Brain Res.* 140:43-62.
39. Alonso JM, Usrey WM, & Reid RC (1996) Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature* 383(6603):815-819.
40. Barsalou LW (2008) Grounded cognition. *Annu Rev Psychol* 59:617-645.
41. Raghavachari S, *et al.* (2001) Gating of human theta oscillations by a working memory task. *J Neurosci* 21(9):3175-3183.
42. Engel TA, *et al.* (2016) Selective modulation of cortical state during spatial attention. *Science* 354(6316):1140-1144.
43. Ashida G & Carr CE (2011) Sound localization: Jeffress and beyond. *Curr Opin Neurobiol* 21(5):745-751.
44. Lee J & Groh JM (2014) Different stimuli, different spatial codes: a visual map and an auditory rate code for oculomotor space in the primate superior colliculus. *PLoS One* 9(1):e85017.
45. Robinson DA (1972) Eye movements evoked by collicular stimulation in the alert monkey. *Vision Res* 12:1795-1807.
46. Mohler CW & Wurtz RH (1976) Organization of monkey superior colliculus: intermediate layer cells discharging before eye movements. *J. Neurophysiol.* 39(4):722-744.
47. Groh JM & Sparks DL (1992) Two models for transforming auditory signals from head-centered to eye-centered coordinates. *Biological Cybernetics* 67(4):291-302.
48. Barton EJ & Sparks DL (2001) Saccades to remembered targets exhibit enhanced orbital position effects in monkeys. *Vision Res* 41(18):2393-2406.
49. Metzger RR, Mullette-Gillman OA, Underhill AM, Cohen YE, & Groh JM (2004) Auditory saccades from different eye positions in the monkey: implications for coordinate transformations. *J. Neurophysiol.* 92(4):2622-2627.
50. Berger J (2006) The case for objective Bayesian analysis. *Bayesian Anal* 1.3:385-402.
51. Berger JO & Pericchi LR (1996) The intrinsic Bayes factor for model selection and prediction. *J Am Stat Assoc* 91.433:109-122.
52. Escobar MD & West M (1995) Bayesian Density-Estimation and Inference Using Mixtures. *J Am Stat Assoc* 90(430):577-588.

SUPPLEMENTARY MATERIALS:

MATERIALS AND METHODS

General procedures

All procedures conformed to the guidelines of the National Institutes of Health (NIH Pub. No. 86-23, Revised 1985) and were approved by the Institutional Animal Care and Use Committee of the university. Two adult rhesus monkeys (Macaca mulatta) participated (monkey P, and monkey Y, both female). Under general anesthesia and in sterile surgery we first implanted a head post holder to restrain the head and a scleral search coil to track eye movements (Robinson 1963; Judge et al. 1980). After recovery with suitable analgesics and veterinary care, we trained the monkeys in the experimental task. In a second surgery, we implanted a recording cylinder (2 cm diameter) over the right (monkey Y) or left (monkey Y, P) IC respectively. We determined the location of the cylinder with stereotactic coordinates and verified it with MRI scans (e.g. 23).

Behavioral task and training

<u>Events of task and performance criteria</u>

The monkeys performed a single- or dual-sound localization task (Figure 2A) by making saccades toward one or two simultaneously-presented auditory targets with one or two saccades as appropriate. All sound targets were located in front of the monkey at eye level; the horizontal location, frequency and intensity were varied pseudorandomly as described below (Recording Procedures). Each trial began with 600-700ms of fixation of a visual stimulus (light emitting-diode, LED, located straight ahead and 10-14° below the speakers). During fixation we presented one sound (single-sound trials) or two simultaneous sounds (dual-sound trials). After a fixation time of either 600-800 (Data Set I, some of Data Set II) or 1000-1100 (remainder of Data Set II), the fixation light was extinguished and the monkey was required to make a single saccade on single-sound trials or a sequence of two saccades (in either order) on dual-sound trials. Trials were considered correct if each saccade was directed within 10-17.5 degrees horizontally and 20-40 degrees vertically of a target (due to vertical inaccuracies in localizing non-visual targets in primates, 48)

and if the gaze was maintained on the final target for 100 – 200 ms. On correct trials monkeys were rewarded with juice drops.

Training

Training was accomplished in three stages. We initially trained the monkeys to report the location of single visual targets by saccading to them. We then introduced single auditory targets. As these were novel and unexpected in the silent experimental booth monkeys readily saccaded to them (49). To help the monkeys calibrate their auditory saccades, a visual feedback was added on trials where the auditory saccade was not initiated correctly within 700 ms. The feedback was presented only at the most peripheral target locations (+/- 24 degree) and only for a few initial days of training. Finally, we trained monkey to localize dual-sound targets. Initially we presented the two sounds sequentially in a specific order, then we gradually reduced the temporal gap between them until the sounds were simultaneous.

In the final version of the task, monkeys were allowed to look at the targets in either order, as noted above. However, due to the initial training with sequential sounds, they retained stereotyped patterns of saccades in which they tended to look first to whichever sound location had been presented first during the sequential and partial overlap stages of training. Monkey P was trained with more central target locations (e.g. -6 or 6 degree targets) initially occurring first and more peripheral targets (e.g. -24 or 24 degree targets) occurring second, and monkey Y was trained with sounds initially occurring in the opposite sequence. Midway through neural data collection, we provided additional training to monkey Y to encourage free choice of which sound to look at first. This allowed us to investigate the relationship between each behavioral response and the neural representation at that moment.

Recording procedure and strategy

General procedure

Recordings were made with one or two tungsten electrodes (FHC, impedance between 1 and 3 MΩ at 1 kHz). Each electrode was lodged in a stainless-steel guide tube (manually advanced through the dura) and controlled independently with an oil hydraulic pulse micropositioner (Narishige International USA, Inc. and NAN INSTRUMENTS LTD, Israel). First, we localized the IC (and isolated single neurons) while the monkey listened passively to

sounds of different frequencies. We then collected single unit spiking activity and local field potential while the monkey performed the single- and dual-sound localization tasks. We used a Multichannel Acquisition Processor (MAP system, Plexon Inc., Dallas, TX) and Sort Client software. The single unit spiking activity was filtered between 150 Hz and 8 kHz and sampled at 20 kHz, while the LFP signal was filtered between 0.7 and 300 Hz and sampled at either 20kHz or 1kHz (see Local Field Potential). Data were collected as long as the neurons were well isolated and the monkey performed the tasks

Neural signals were recorded primarily from two functionally-defined subregions of the IC, the low frequency area and the tonotopic area (23). Neurons in the low frequency tuned area generally respond best to low frequencies and there is little heterogeneity in tuning, whereas neurons recorded in the tonotopic area had best frequencies that could be either low or high depending on the position of the recording electrode.

Data Sets and Auditory Stimuli: Locations, Frequencies, and Levels

The spiking activity of 166 single neurons was recorded, in two datasets involving the same task but differing in which sound levels and frequencies were included. A total of 68 of these neurons were recorded as pairs from separate electrodes positioned in the IC on the same side of the brain at a minimum spatial separation of 2 mm. Local field potentials (LFP) were also recorded from 87 of these recording sites.

In both datasets, the sounds consisted of bandpass noise with a bandwidth of +/- 200 Hz. On dual-sound trials, the sounds were delivered from pairs of locations (24 degrees and -6 degrees), and (-24 and +6 degrees) i.e. 30 degrees apart. The two sounds differed in frequency, with one of the two sounds having a 742 Hz center frequency and the other differing by at least 0.285 octaves or multiples of this distance. Single-sound trials involved the same set of locations and frequencies as on dual-sound trials, but with only a single sound presented at a time. All sounds were “frozen” within an individual session; that is, all trials with a given set of auditory parameters involved the same time series signal delivered to the relevant speaker.

In data set I (N=98 neurons), the sounds presented on dual-sound trials were 742 Hz and a sound from the set (500, 609, 903, 1100 Hz); these frequencies were ±0.285 octave or ± 0.57 octaves above or below 742 Hz, or ±3.4 and 6.8 semitones. Combining two sounds will produce a combination that is louder than either component. Sound levels were therefore calibrated to

provide two sets of conditions: dual sounds for which the component sounds involve the *same signals* to the audio speakers as on single-sound trials, producing a *louder* dual sound, and dual sounds for which the level of the component sounds was *reduced* so that the overall *loudness* was the *same* on dual as on single trials. The levels used for the components were 51 and 55 dB, producing sound levels of minimum 55 or maximum 60 dB on dual-sound trials. The same-signal comparison involved using the 55 dB component levels, singly and on dual-sound trials. The same-loudness comparison involved using the 55 dB levels on single-sound trials and the 51 dB levels for the components of dual-sound trials. Calibrations were performed using a microphone (Bruel and Kjaer 2237 sound level meter) placed at the position normally occupied by the animal's head.

Because results did not differ substantively when comparisons were made between same-signal and same-loudness conditions (Figure 2 vs. Supplementary Figure 1), we pooled across sound levels for subsequent analyses, and we dispensed with the multiple sound levels for data set II (monkey Y only, N=68 neurons), using either 50 or 55 dB levels for all components. We also incorporated additional sound frequencies, [1340 1632 1988 Hz], to improve the odds that responses to each of the component sounds differed significantly. Again, one of the two sounds on dual-sound trials was 742 Hz; the other sound frequency was either from the original list of [500 609 903 1100] or from the new frequencies. Most of the neurons in this data set were tested with [500 742 1632].

<u>Cell Inclusion/Exclusion criteria and trial counts</u>

The N=166 neurons (N=98 from Data Set I and N=68 from Data Set II) included for analysis were drawn from a larger set of 325 neurons. Neurons were excluded from analysis if the neuron proved unresponsive to sound (Student's t-test, spike counts during the 600 ms after sound onset compared to the same period immediately prior to sound onset, one-tailed, $p>0.05$), or if there were too few correct trials (mininum of five correct trials for each of the components [A, B, and AB trials] that formed a given "triplet" of conditions or if there were technical problems during data collection (e.g. problems with random interleaving of conditions or with computer crashes). The average number of correct trials for a given set of stimulus conditions in the included dataset (N=166) was 10.5 trials. The total number of included triplets was 1484.

Data Analysis

All analyses concerned correctly performed trials.

<u>Analysis of activity pooled across time and/or trials: Summation and Averaging</u>

To evaluate IC activity using conventional analysis methods that pool across time and/or across trials, we counted action potentials during two standard time periods. The baseline period (*Base*) was the 600ms period before target onset, and the sensory-related target period (*Resp*) was the 600ms period after target onset (i.e. ending before, or at the time of, the offset of the fixation light. Figure 2A).

Summation/Averaging Indices: We quantified the activity on dual-sound trials in comparison to the sum and the average of the activity on single-sound trials, expressed in units of standard deviation (Z-scores), similar to a method used by (33). Specifically, we calculated,

$$PredictedSum_{A,B,} = mean(Resp_A) + mean(Resp_B) - mean(Base_{A,B}) \quad (1)$$

and

$$PredictedAvg_{A,B,} = (mean(Resp_A) + mean(Resp_B))/2 \quad (2)$$

where $Resp_A$ and $Resp_B$ were the number of spikes of a given neuron for a given set of single-sound conditions *A* and *B* (location, frequency, and intensity) that matched the component sounds of the dual-sound trials being evaluated. As the "response" may actually include a contribution from spontaneous baseline activity, we subtracted the mean of the baseline activity for the single sounds ($Base_{A,B}$). Without this subtraction, the predicted sum would be artificially high because two "copies" of baseline activity are included under the guise of the response activity.

The Z scores for the dual-sound trials were computed by subtracting these predicted values from the mean of the dual-sound trials (mean($Resp_{AB}$)) and dividing by the mean of the standard deviations of the responses on single-sound trials:

$$Zsum_{AB} = \frac{mean(Resp_{AB}) - PredictedSum_{A,B}}{mean(std(Resp_A), std(Resp_B))} \quad (3)$$

and

$$ZAvg_{AB} = \frac{mean(Resp_{AB}) - PredictedAvg_{A,B}}{mean(std(Resp_A), std(Resp_B))} \quad (4)$$

If the dual response was within +/- 1.96 of the predicted sum or predicted average, we could say the actual dual response was within the 95% confidence intervals for addition or averaging of two single responses, respectively.

<u>Analyses of fluctuations in neural firing across and within-trials, and inclusion criteria</u>

Our statistical tests for fluctuations in neural firing were conducted on triplets, or related sets of single and dual-sound trials (A, B, AB trials). To evaluate whether neural activity fluctuates *across* trials in a fashion consistent with switching between firing patterns representing the component sounds, we evaluated the Poisson characteristics of the spike trains on matching dual and single-sound trials (triplets: AB, A and B). Spike train data from each trial was summarized by the total spike count between 0-600ms or 0-1000 ms from sound onset (i.e. whatever the minimum duration of the overlap between fixation and sound presentation was for that recorded neuron, see section Events of Task). We modeled the distribution of spike counts in response to single sounds A and B as Poisson distributions with unknown rates λ^A , denoted $Poi(\lambda^A)$, and λ^B , denoted $Poi(\lambda^B)$. Four hypotheses were considered for the distribution of sound AB spike counts:

1. a mixture distribution $\alpha \cdot Poi(\lambda^A) + (1 - \alpha) \cdot Poi(\lambda^B)$ with an unknown mixing weight α ("mixture")
2. a single $Poi(\lambda^{AB})$ with some λ^{AB} in between λ^A and λ^B ("intermediate")
3. a single $Poi(\lambda^{AB})$ where λ^{AB} is either larger or smaller than both λ^A and λ^B ("outside")
4. a single $Poi(\lambda^{AB})$ where λ^{AB} exactly equals one of λ^A and λ^B ("single")

Relative plausibility of these competing hypotheses was assessed by computing their posterior probabilities with equal prior weights (1/4) assigned to the models, and with default Jeffreys' prior (50) on model specific Poisson rate parameters, and a uniform prior on the mixing weight parameter α. Posterior model probabilities were calculated by computation of relevant intrinsic Bayes factors (51).

Triplets were excluded if either of the following applied: 1) the Poisson assumption on A and B trial counts was not supported by data; or 2) λ^A and λ^B were not well separated. To test the Poisson assumption on single-sound trials A and B of a given triplet, we used an approximate chi-square goodness of fit test with Monte Carlo p-value calculation. For each sound type, we estimated the Poisson rate by averaging counts across trials. Equal probability bins were constructed from the quantiles of this estimated Poisson distribution, with number of bins determined by expected count of 5 trials in each bin or at least 3 bins -- whichever resulted in more bins. A lack-of-fit statistic was calculated by summing across all bins the ratio of the square of the difference between observed and expected bin counts to the expected bin count. Ten thousand Monte Carlo samples of Poisson counts, with sample size given by the observed number of trials, were generated from the estimated Poisson distribution and the lack-of-fit statistic was calculated from each one of these samples. P-value was calculated as the proportion of these Monte Carlo samples with lack-of-fit statistic larger than the statistic value from the observed data. Poisson assumption was considered invalid if the resulting Monte Carlo p-value < 0.1.

For triplets with valid Poisson assumption on sound A and B spike counts, we tested for substantial separation between λ^A and λ^B, by calculating the intrinsic Bayes factor of the model $\lambda^A \neq \lambda^B$ against $\lambda^A = \lambda^B$ with the non-informative Jeffreys' prior on the λ parameters: λ^A, λ^B or their common value. The triplet was considered well separated in its single sounds if the logarithm of the intrinsic Bayes factor equaled 3 or more, which is the same as saying the posterior probability of $\lambda^B \neq \lambda^A$ exceeded 95% when a-priori the two models were given 50-50 chance.

Dynamic Admixture Point Process Model

To evaluate whether neural activity fluctuates *within* trials, we developed a novel analysis method we call a Dynamic Admixture Point Process model (DAPP) which characterized the

dynamics of spike trains on dual sound trials as an admixture of those occurring on single sound trials. The analysis was carried out by binning time into moderately small time intervals. Given a predetermined bin-width $w = T/C$ for some integer C, we divided the response period into contiguous time intervals $I_1 = [0;w)$; $I_2 = [w; 2w)\ldots$ $I_C = [(C-1)w,Cw)$ and reduced each trial to a C-dimensional vector of bin counts $(X^e_{j1},\ldots,X^e_{jC})$ for $e \in \{A;B;AB\}$ and $j = 1,\ldots, n_e$. Mathematically, $X^e_{jC} = N^e_j(Ic)$. The results reported in Section 2 were based on $w = 50$ (with time measured in ms and T = 600 or 1000), but we also repeated the analyses with $w = 25$ and noticed little difference.

Our model for the bin counts was the following. Below we denote by t_c^* the mid-point $(c - 1/2)w$ of sub-interval I_c.

1. $X_{jc}^e \sim Poi\big(w \cdot \lambda^e(t_c^*)\big)$, $e \in \{A, B\}$, $c \in \{1, \ldots, C\}$, $j \in \{1, \ldots, n_e\}$. We assume both $\lambda^A(t)$ and $\lambda^B(t)$are smooth functions over $t \in [0, T]$.

2. $X_{jc}^{AB} \sim Poi(w \cdot \lambda_j(t_c^*))$, where $\lambda_j(t) = \lambda_j(t) = \alpha_j(t)\lambda^A(t) + \{1 - \alpha_j(t)\}\lambda^B(t)$ with $\alpha_j: [O, T] \to (0,1)$ being unknown smooth functions.

We modeled $\alpha_j(t) = S(\eta_j(t))$, where $S(t) = 1/(1 + e^{-t})$, and, each $\eta_j(t)$ was taken to be a (smooth) Gaussian process with $E\{\eta_j(t)\} \equiv \phi_j$, $Var\{\eta_j(t)\} \equiv \psi_j$, and, $Cor\{\eta_j(t), \eta_j(t')\} = \exp\{-0.5(t - t')^2/\ell_j^2\}$. The three parameters (ϕ_j, ψ_j, ℓ_j) respectively encoded the *long-term average value*, the *total swing magnitude* and the *waviness* of the $\alpha_j(t)$ curve. While the temporal imprint carried by each α_j was allowed to be distinct, we enforced the dual trials to share dynamic patterns by assuming $(\phi_j, \psi_j, \ell_j), j = 1, \ldots, n_{AB}$, were drawn from a common, unknown probability distribution P, which we called a dynamic pattern generator and viewed as a characteristic of the triplet to be estimated from the data.

To facilitate estimation of P, we assumed it decomposed as $P = P_{\phi\psi} \times P_\ell$, where $P_{\phi\psi}$ was an unknown distribution on $(-\infty, \infty) \times (0, \infty)$ generating (ϕ_j, ψ_j), and, P_ℓ was an unknown distribution on $(0, \infty)$ generating ℓ_j. To simplify computation, we restricted ℓ_j to take only a finitely many positive values, representative of the waviness range we are interested in (in our

analyses, we took these representative values to be {75, 125, 200, 300, 500}, all in ms). This restricted P_ℓ to be a finite dimensional probability vector.

We performed an approximate Bayesian estimation of model parameters. Note that only $\lambda^A(t)$ and $\lambda^B(t)$ were informed by the single sound trial data. All other model parameters were informed only by the dual sound trial data conditionally on the knowledge of $\lambda^A(t)$ and $\lambda^B(t)$. To take advantage of this partial factorization of information sources, we first smoothed each set of single sound trial data to construct a conditional gamma prior for the corresponding $\lambda^e(t_c^*), e \in \{A, B\}, c = 1, \dots, C$, where the gamma distribution's mean and standard deviation were matched with the estimate and standard error of $\lambda^e(t_c^*)$. A formal Bayesian estimation was then carried out on all model parameters jointly by (a) using only the dual sound trial data, (b) utilizing the conditional gamma priors on $\lambda^A(t)$ and $\lambda^B(t)$, and, (c) assuming a Dirichlet process prior (52) on $P_{\phi\psi}$ and an ordinary Dirichlet prior on P_ℓ. This final step involved a Markov chain Monte Carlo computation whose details will be reported in a separate paper.

Next, the estimate of the generator P was utilized to repeatedly simulate $\alpha(t)$ functions for hypothetical, new dual trials for the triplet. For each simulated $\alpha(t)$ curve, we computed its maximum swing size $|\alpha| = max_t\, \alpha(t) - min_t \alpha(t)$, and, time aggregated average value $\bar{\alpha} = \int_0^T \alpha(t)dt\, /T$. The waviness index of the triplet was computed as the odds of seeing an $\alpha(t)$ function exhibiting a swing of at least 50% between its peak and trough:

$$r_w = \frac{P(|\alpha|>0.5)}{P(|\alpha|<0.5)}$$

where P denotes the sampling proportion of the simulated α draws. The triplet's extremeness index was computed as the odds of seeing an $\alpha(t)$ function with its long-term average $\bar{\alpha}$ being closer to the mid-way mark of 50% than the extremes:

$$r_c = \frac{P(\bar{\alpha} \in (0.25, 0.75))}{P(\bar{\alpha} \notin (0.25, 0.75))}.$$

The two indices were then thresholded to generate a 2-way classification of all triplets. On waviness, a triplet was categorized as "wavy", "flat" or "ambiguous" according to whether $r_w > 1.3$, $r_w < 0.77$, or, $0.77 \leq r_w \leq 1.3$, respectively On extremeness, the categories were "central", "extreme", or, "ambiguous" according to whether $r_c > 3.24$, $r_c < 1.68$, or, $1.68 \leq r_c \leq 3.24$, respectively.

In addition to the flat/wavy and extreme/central classification, a third parameter was evaluated for each triplet: the degree of skewness in the distribution of $\overline{\alpha_*}$:

$$r_S = max\left(\left\{\frac{P(\overline{\alpha_*}<0.5)}{P(\overline{\alpha_*}>0.5)}, \frac{P(\overline{\alpha_*}>0.5)}{P(\overline{\alpha_*}<0.5)}\right\}\right)$$

which ranges in (1,∞). Each triplet's flat/wavy/central/extreme tag could then be subcategorized as either "skewed" or "symmetric" depending on whether $r_S > 4$ or $r_S < 2$ (with no label in the middle). This subcategorization step was useful for distinguishing the dynamic admixtures associated with the whole-trial categorizations of "single" and "outside" from "intermediate" and "mixture", with "single" and "outside" tending to be classified as "skewed". Supplementary Table 1 and Supplementary Figures 2 and 3 give the full results of the main 2-way classification together with the symmetry/skewness subclassification, cross tabulated with the classification done under the whole trial spike count analysis.

<u>A vs. B assignment scores: individual neurons, pairs of neurons, local field potential, and behavioral prediction</u>

A vs. B assignment scores were computed for several analyses (the example shown in Figure 3A-D; pairs of recorded neurons; the relationship between spiking activity and local field potential; and the relationship between saccade sequences and spiking activity). For each triplet, every dual-sound trial received an "A-like" score and a "B-like" score, either for the entire response window (600-1000 ms after sound onset) or for 50 ms time bins. The scores were computed as the posterior probability that the spike count in each dual-sound trial was drawn from the Poisson distribution of single-sound spike counts,

For the pairs analysis, the A vs. B assignment scores were computed within each 50 ms time bin independently for each pair of neurons recorded simultaneously. The scores were normalized across trials by subtracting the mean score and dividing by the standard deviation of scores for that bin (a Z-score in units of standard deviation). Only conditions for which both recorded neurons exhibited reasonably different responses to the "A" vs. the "B" sound and for which there were at least 5 correct trials for A, B, and AB trials were included (t-test, $p < 0.05$). A total of 206 conditions were included in this analysis.

<u>Local field potential analysis</u>

We analyzed the local field potential from 87 sites in both monkeys (30 sites from monkey P's left IC, 31 sites from monkey Y's right IC and 26 sites from monkey Y's left IC). The LFP acquisition was either recorded in discrete temporal epochs encompassing behavioral trials (roughly 1.2 to 2 seconds long) and at a sampling rate of 20 kHz (Dataset I, part of Dataset II), or as a continuous LFP signal during each session, at a sampling rate of 20 kHz or 1kHz (rest of Dataset II). We standardized the LFP signals by trimming the continuous LFP into single trial intervals and down-sampling all signals to 1 kHz. The MAP system filters LFP signals between 0.7 and 300 Hz; no additional filtering was applied. For each site we subtracted the overall mean LFP value calculated over the entire session, to remove any DC shifts, and we excluded trials that exceeded 500mV. For each triplet, we assigned individual dual-sound trials to two groups based on the total spike count in a 600 ms response window (see Methods: A vs. B assignment scores). The average LFP was then compared across the two groups in two 600 ms windows before and after sound onset (baseline and response periods). The results reported here refer to these mean-normalized LFP signals. We obtained similar results when the amplitude of each trial's LFP was scaled as a proportion of the maximum response within the session.

SUPPLEMENTARY FIGURES AND TABLES

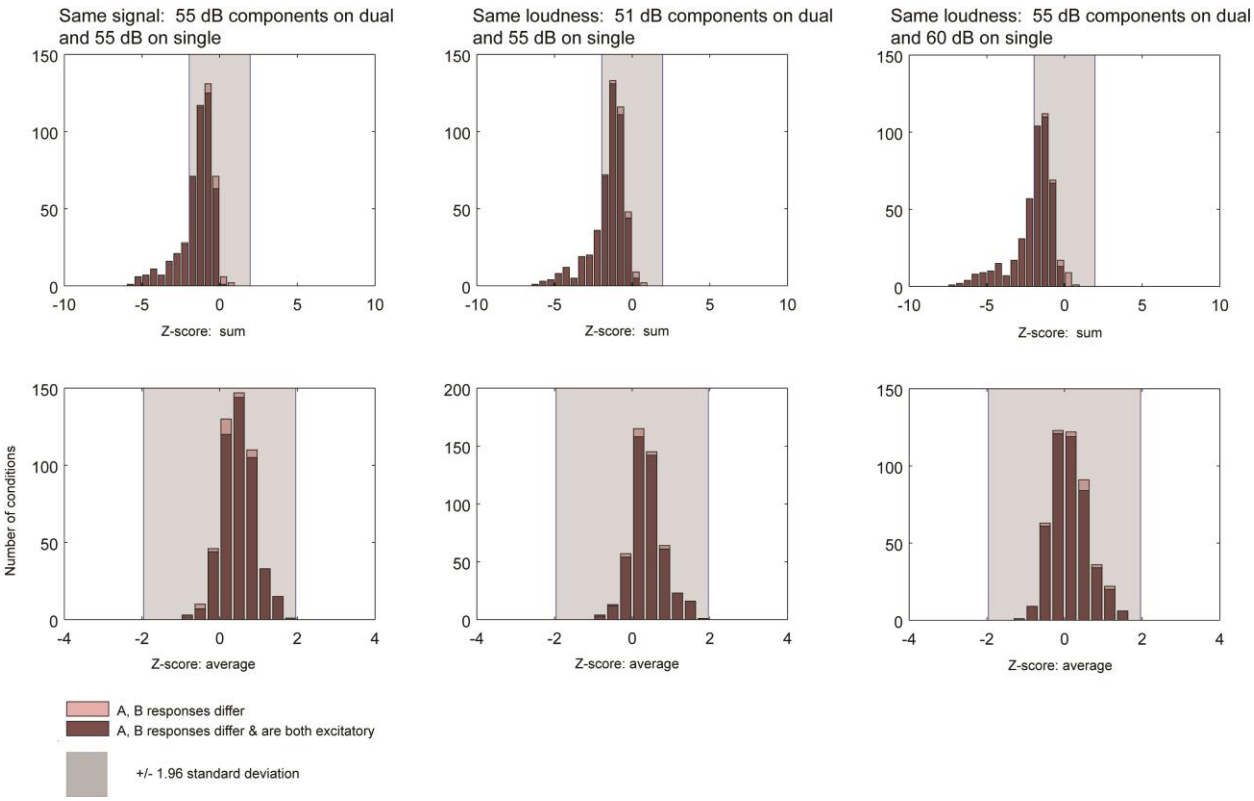

Supplementary Figure 1. *Same analysis as figure 2, but comparing dual-sound trials to single-sound trials that used the same signal sent to the speakers (left column), or for which the signal was amplified on single-sound trials to match the dual sounds in loudness (middle and right columns). The results are essentially identical to each other. Accordingly, the remainder of the analyses in the paper ignored sound intensity as a factor.*

	Intermediate	Mixture	Outside	Single	Total
Wavy-central+symmetric	13 (5)	2			15 (5)
Wavy-central	1				1
Wavy+symmetric	4	2			6
Flat-central+symmetric	7 (3)				7 (3)
Flat-central	2 (1)	1			3 (1)
Flat-extreme+skewed	8 (1)	12 (3)	19 (4)	46	85 (8)
Flat-extreme+symmetric	2	22 (19)		4	28 (19)
Flat-extreme	3	31 (13)	1	16	51 (13)

Flat+skewed	4 (2)				4 (2)
Flat+symmetric	6	11 (4)		4	21 (4)
Flat	13 (2)	1		2	16 (2)
Extreme+symmetric	1	4 (4)		1	6 (4)
Extreme		3 (2)		1	4 (2)
Central+symmetric	9 (2)	1		2	12 (2)
Symmetric	36 (1)	42 (5)		14	92 (6)
Skewed	1 (1)				1 (1)
Ambiguous	5	2		3	10
Total	115 (18)	134 (50)	20 (4)	93	362 (72)

Supplementary Table 1. *Complete listing of "tags" applied in the DAPP analysis. Categories not shown (e.g. wavy-central+skewed) did not receive any members.___The numbers in parentheses are for triplets that produced a winning model in the whole trial analysis with a probablility > 95%. The total is 362, since for one "single" triplet, the DAPP analysis failed to compute the necessary metrics.*

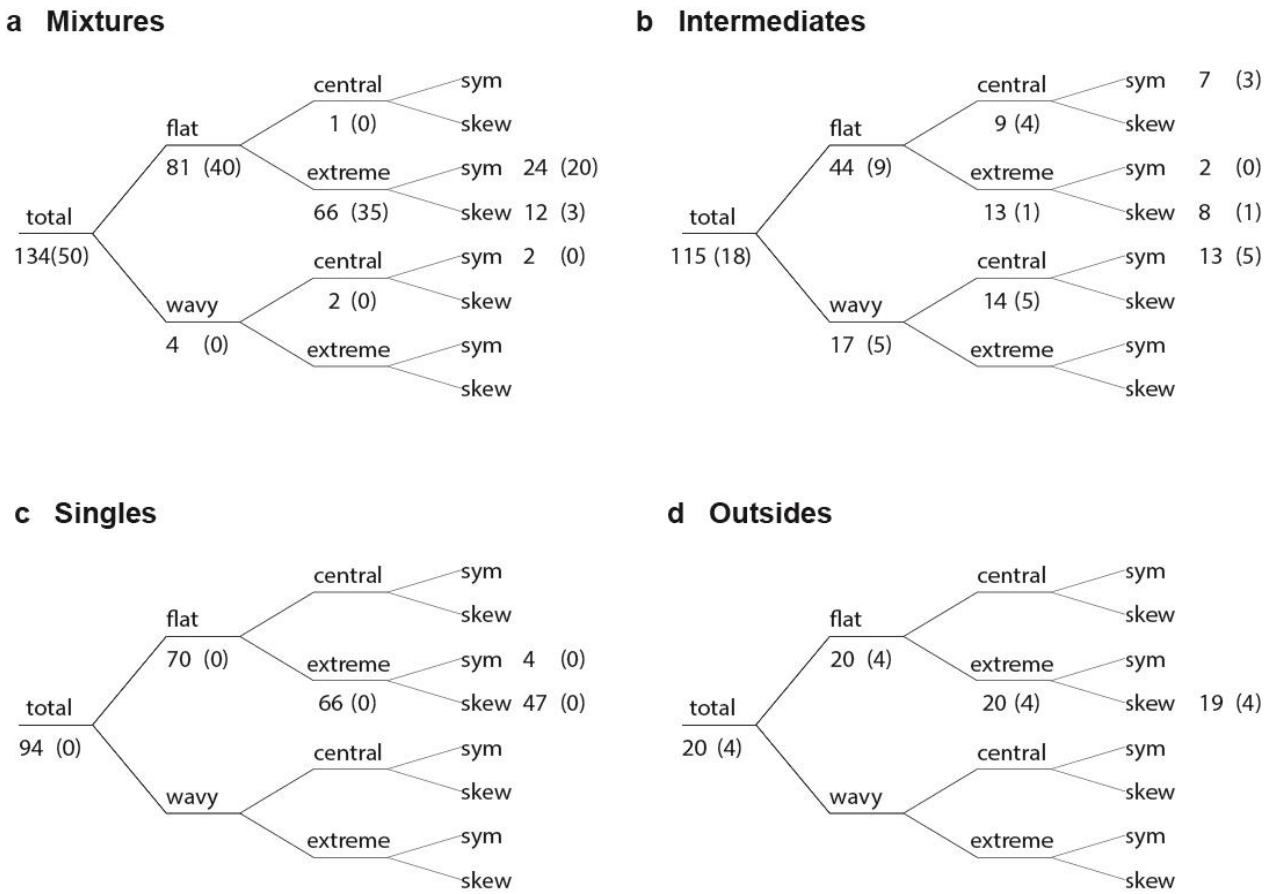

Supplementary Figure 2. *Hierarchical depiction of the DAPP tags, considering first flat vs. wavy, then central vs. extreme, then symmetric vs. skewed, as a function of whole trial model classification. As in Supplementary Table 1, numbers in parentheses are for triplets that produced a winning model in the whole trial analysis with a probability > 95% and the total included here is 362. For clarity, labels "0(0)" are not shown.*

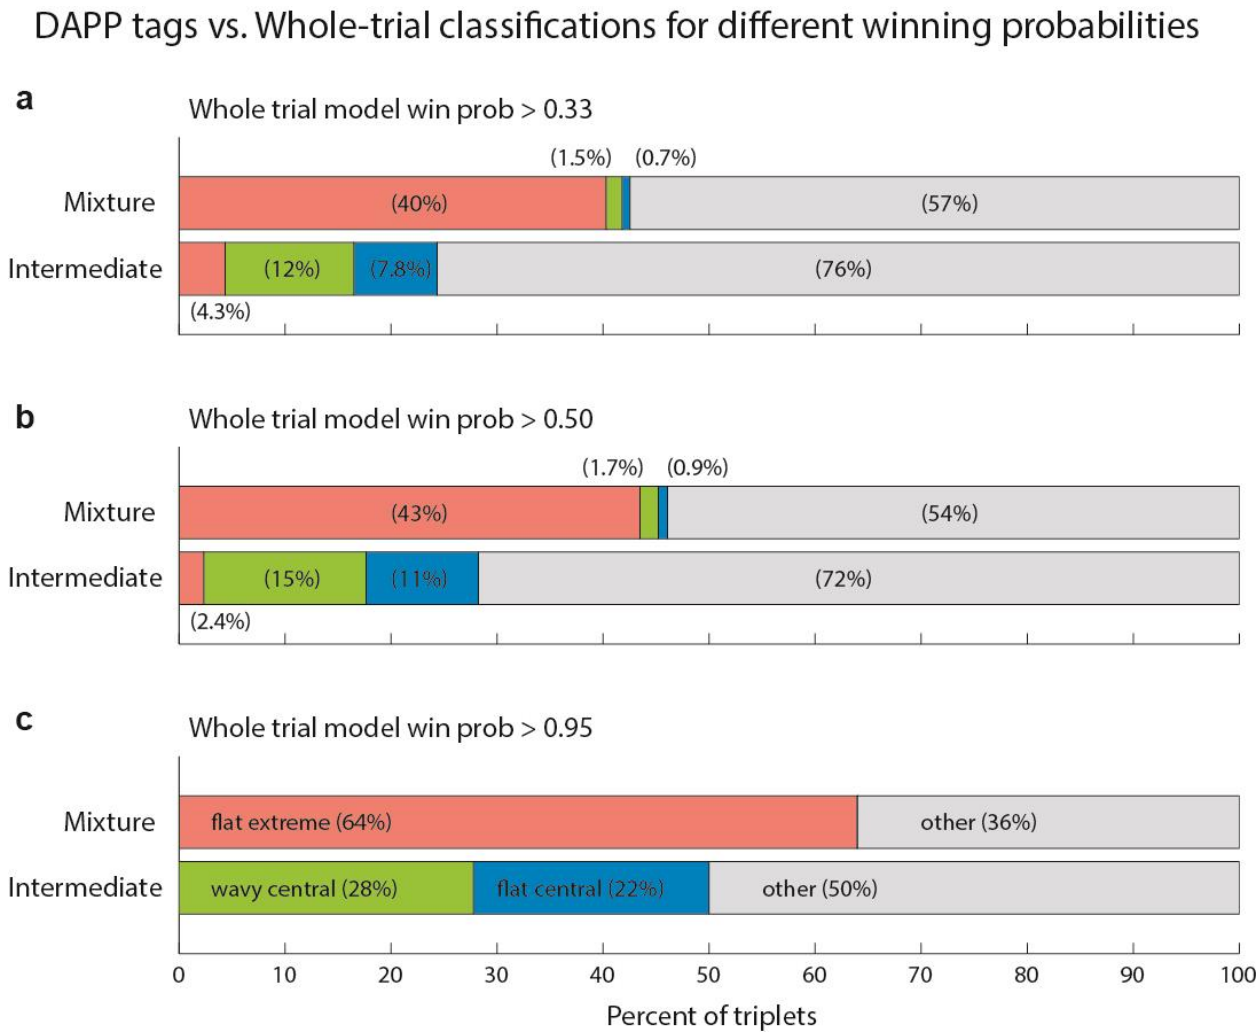

Supplementary Figure 3. *Relationship between the three relevant DAPP tag combinations (flat-extreme, wavy-central, and flat-central) and whole-trial Poisson classification, for different levels of winning probability for the whole-trial analysis. Panel C is identical to Figure 5C in the main text, and triplets that received a "skewed" subclassification are excluded.*

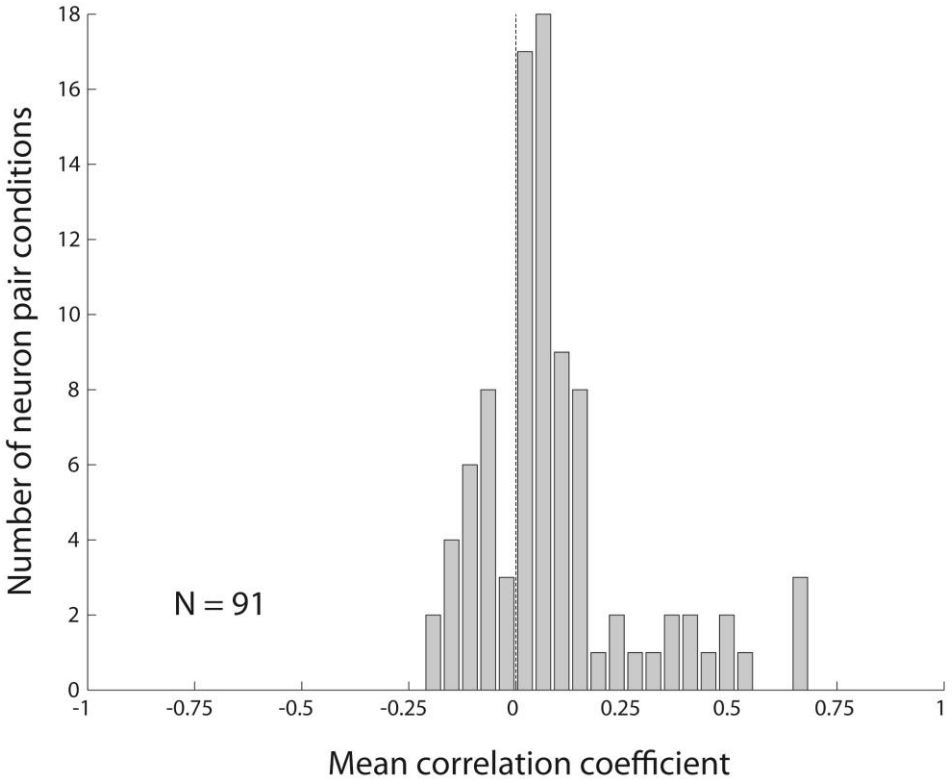

Supplementary Figure 4. *Same analysis as in Figure 6B, but using spike counts in each bin instead of A vs. B assignment scores.*