

UC Davis

UC Davis Previously Published Works

Title

Feature-Selective Attention Adaptively Shifts Noise Correlations in Primary Auditory Cortex

Permalink

<https://escholarship.org/uc/item/9z6724mr>

Journal

Journal of Neuroscience, 37(21)

ISSN

0270-6474

Authors

Downer, Joshua D
Rapone, Brittany
Verhein, Jessica
[et al.](#)

Publication Date

2017-05-24

DOI

10.1523/jneurosci.3169-16.2017

Peer reviewed

Feature-Selective Attention Adaptively Shifts Noise Correlations in Primary Auditory Cortex

Joshua D. Downer, Brittany Rapone, Jessica Verhein, Kevin N. O'Connor, and Mitchell L. Sutter

Center for Neuroscience, University of California, Davis, California 95618

Sensory environments often contain an overwhelming amount of information, with both relevant and irrelevant information competing for neural resources. Feature attention mediates this competition by selecting the sensory features needed to form a coherent percept. How attention affects the activity of populations of neurons to support this process is poorly understood because population coding is typically studied through simulations in which one sensory feature is encoded without competition. Therefore, to study the effects of feature attention on population-based neural coding, investigations must be extended to include stimuli with both relevant and irrelevant features. We measured noise correlations (r_{noise}) within small neural populations in primary auditory cortex while rhesus macaques performed a novel feature-selective attention task. We found that the effect of feature-selective attention on r_{noise} depended not only on the population tuning to the attended feature, but also on the tuning to the distractor feature. To attempt to explain how these observed effects might support enhanced perceptual performance, we propose an extension of a simple and influential model in which shifts in r_{noise} can simultaneously enhance the representation of the attended feature while suppressing the distractor. These findings present a novel mechanism by which attention modulates neural populations to support sensory processing in cluttered environments.

Key words: amplitude modulation; attention; auditory cortex; feature-selective attention; noise correlation

Significance Statement

Although feature-selective attention constitutes one of the building blocks of listening in natural environments, its neural bases remain obscure. To address this, we developed a novel auditory feature-selective attention task and measured noise correlations (r_{noise}) in rhesus macaque A1 during task performance. Unlike previous studies showing that the effect of attention on r_{noise} depends on population tuning to the attended feature, we show that the effect of attention depends on the tuning to the distractor feature as well. We suggest that these effects represent an efficient process by which sensory cortex simultaneously enhances relevant information and suppresses irrelevant information.

Introduction

Listening is a challenging task. Multiple mixed sound sources arrive at the ear and the auditory system must then select relevant sound features from this mixture for further processing while suppressing irrelevant features. This is often called the “cocktail

party problem” and humans and other animals can solve this problem (Bee and Micheyl, 2008) despite its computational complexity (Haykin and Chen, 2005). Recent work has demonstrated that listeners do this in part by shifting attention between sound features (Woods and McDermott, 2015), although the mechanism of this adaptive feature selection is unknown (Shamma et al., 2011).

Recent work has highlighted attentional modulation of noise correlations (r_{noise}) as a potential mechanism to enhance perception (Cohen and Maunsell, 2009). r_{noise} arises between neurons when their response variability (“noise”) to repeated presentations of the same stimulus is correlated. These correlations are thought to affect the information bearing capacity of a neural population (Averbeck et al., 2006), although theoretical studies of the effect of r_{noise} on population coding suggest that the variables that mediate the impact of r_{noise} on coding are complex (Ecker et al., 2011; Hu et al., 2014; Moreno-Bote et al., 2014; Kanitscheider et al., 2015). One theory, the sign rule (SR) (Hu et al., 2014), posits that the effect of r_{noise} on coding accuracy can be

Received Oct. 11, 2016; revised Feb. 23, 2017; accepted Feb. 24, 2017.

Author contributions: J.D.D., J.V., K.N.O., and M.L.S. designed research; J.D.D., B.R., and J.V. performed research; J.D.D., K.N.O., and M.L.S. contributed unpublished reagents/analytic tools; J.D.D. and B.R. analyzed data; J.D.D., K.N.O., and M.L.S. wrote the paper.

This work was supported by the National Institute on Deafness and Other Communication Disorders—National Institutes of Health (Grant 02514 to M.L.S.), the National Science Foundation (Graduate Research Fellowship 1148897 to J.D.D.), and the ARCS Foundation (Fellowship to J.D.D.). We thank Felix Bacigalupo for comments on an earlier version of the manuscript and Jalina Graham for providing illustrations for figures.

The authors declare no competing financial interests.

B. Rapone's present address: Department of Ophthalmology, University of California, San Francisco, CA 94122.

J. Verhein's present address: School of Medicine, Stanford University, Stanford, CA 94305.

Correspondence should be addressed to Mitchell L. Sutter, Center for Neuroscience, University of California, 1544 Newton Court, Davis, CA 95618. E-mail: mlsutter@ucdavis.edu.

DOI:10.1523/JNEUROSCI.3169-16.2017

Copyright © 2017 the authors 0270-6474/17/375378-15\$15.00/0




Figure 1. Effect of r_{noise} in the presence of irrelevant (distractor) stimuli. **A, B**, Two graphic illustrations of the relationship between r_{noise} and r_{tuning} according to the SR. Within each panel, we represent the mean (points) and variance (dotted ellipses) of the joint response distributions of two neurons in response to a standard (gray) and target (red) stimulus. In **A**, this pair of neurons exhibits negative r_{tuning} ; that is, as the mean of neuron 1 increases, the mean of neuron 2 decreases. When r_{noise} is negative (i.e., the same sign as r_{tuning} ; **A**, left), the variance of the joint response distributions lies along approximately the same dimension as the signal, so this pair does not perfectly detect the target due to the overlap of the joint response distributions. However, when r_{noise} and r_{tuning} have opposite signs (**A**, right), target detection is enhanced, with no overlap in the joint response distributions. This rule can also be seen in **B** for a pair with positive r_{tuning} . **C–F**, Addition of a distractor stimulus (blue) necessitates an extension of the SR. Namely, when target and distractor r_{tuning} have the same sign (**D, E**), changes in r_{noise} will affect targets and distractors similarly (enhance or suppress detection of both) and thus nullify any benefit of modulating r_{noise} . However, when target and distractor r_{tuning} have opposite signs (**C, F**), modulating r_{noise} can simultaneously enhance target detection and suppress distractor detection.

determined based on the joint tuning or tuning correlation (r_{tuning}) between the pairs of neurons in the population: when r_{tuning} and r_{noise} have identical signs, sensory coding performance is impaired and, when they have opposite signs, performance is enhanced (Oram et al., 1998; Abbott and Dayan, 1999; Averbeck et al., 2006). However, despite theoretical disagreements regarding the effects of r_{noise} on coding, the effects of attention on r_{noise} have been remarkably reliable across studies; namely, studies have found that attention globally decreases r_{noise} (Cohen and Maunsell, 2009; Mitchell et al., 2009; Cohen and Maunsell, 2011; Herrero et al., 2013; Doiron et al., 2016; Nandy et al., 2017). Further, recent studies have shown that attention can modulate r_{noise} depending on r_{tuning} , consistent with the SR (Ruff and Cohen, 2014a; Downer et al., 2015). Importantly, the studies mentioned

above, both theoretical and empirical, have only dealt with the detection or discrimination of a single sensory variable in isolation (Fig. 1A, B). How r_{noise} contributes to coding in multifeature settings and how feature-selective attention may rely on dynamic changes in r_{noise} remains unexplored.

Modulation of r_{noise} presents a compelling potential mechanism for feature-selective attention. In cases in which a relevant (target) feature must be selected among competing irrelevant (distractor) features, attentional modulation of r_{noise} may allow for simultaneous enhancement of target detection alongside suppression of distractor detection. According to the SR, this would be possible between populations of neurons with opposite r_{tuning} signs for the target and distractor features (Fig. 1C, F); for example, positive r_{tuning} for the target feature and negative r_{tuning} for

the distractor feature (Fig. 1*F*). In such cases, a given shift in r_{noise} would have opposite effects on the detection of the target and distractor. For example, an increase in r_{noise} between pairs with negative target r_{tuning} and positive distractor r_{tuning} could simultaneously enhance target information and decrease distractor information (Fig. 1*C*).

We tested the effects of feature-selective attention on r_{noise} in a population of neurons with heterogeneous target and distractor tuning. To do so, we recorded from small populations of primary auditory cortical (A1) neurons from two rhesus macaques while they performed a feature-selective attention task. We found that attention's effects on r_{noise} depended upon both target r_{tuning} and distractor r_{tuning} in a manner consistent with simultaneous enhancement of target detection and suppression of distractor detection. Therefore, modulation of r_{noise} appears to support feature-selective attention efficiently. These results demonstrate a novel mechanism for listening in cluttered environments and contribute to an active debate on the role of r_{noise} in population coding.

Materials and Methods

Subjects. Our research subjects were two adult rhesus macaques, one male (Monkey U, 12 kg) and one female (Monkey W, 8 kg). All procedures were approved by the University of California–Davis animal care and use committee and met the requirements of the United States Public Health Service policy on experimental animal care. We implanted each animal with a head post centrally behind the brow ridge and a recording cylinder over an 18 mm craniotomy over the left parietal lobe using aseptic surgical techniques. Placement of the craniotomy was based on stereotactic coordinates of auditory cortex to allow us vertical access through parietal cortex to the superior temporal plane (Sallem and Logothetis, 2007).

Stimuli and task. We presented sounds that varied along two feature dimensions, spectral and temporal, as well as an unmodified broadband noise sound (Fig. 2*B*). The unmodified sound was broadband (white) Gaussian noise with a 9 octave (40 to 20480 Hz) range. We used four different seeds to create noise signals and these were frozen across recording sessions. To introduce variance along each feature dimension, this sound was narrowed in terms of the spectral bandwidth (ΔBW) or sinusoidally amplitude modulated (AM). We also manipulated the degree (value) of variation in each dimension to measure behavioral and neural responses near the perceptual threshold for detecting each feature.

Single-feature threshold determination. We first assessed the behavioral thresholds for detecting each feature in isolation using a yes/no

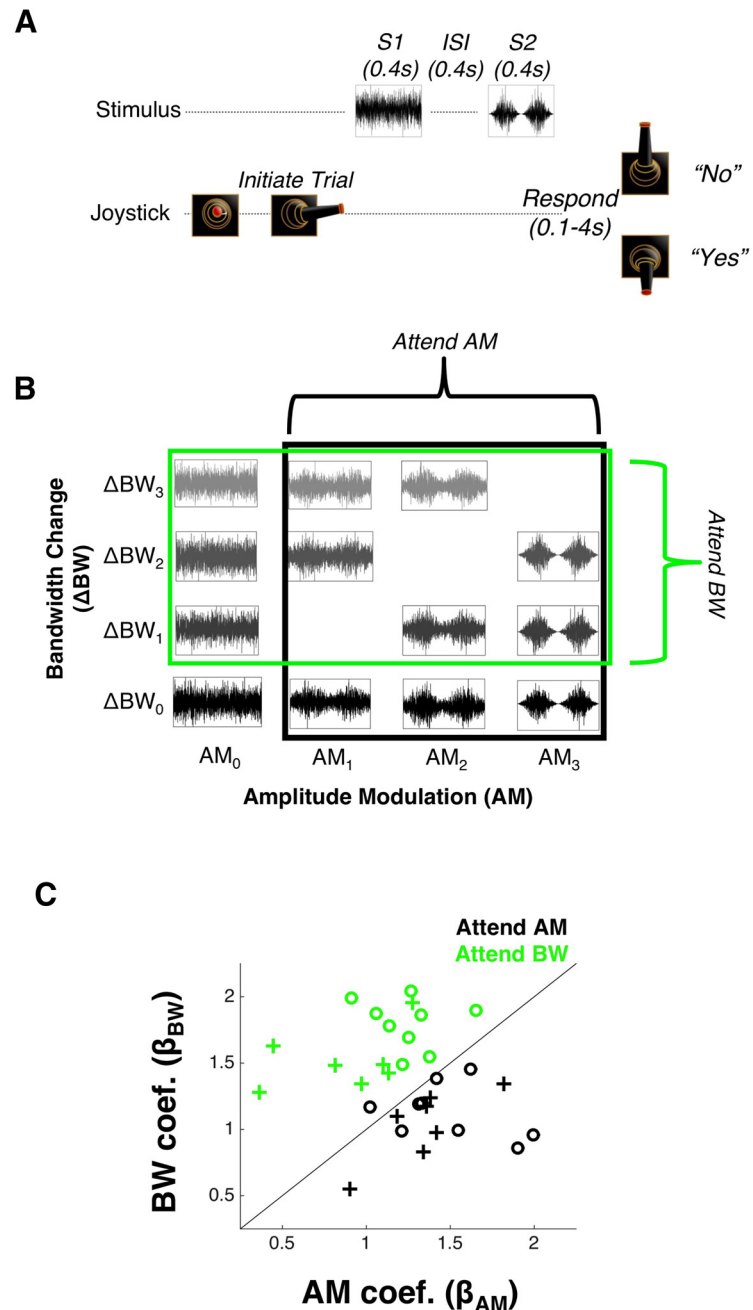


Figure 2. Rhesus macaques perform a novel auditory feature-selective attention task. **A**, Subjects use a joystick to initiate trials and to respond. We present two sequential sounds after subjects initiate a trial. The first sound (S1) is always a broadband (9 octave wide) noise burst. The second sound (S2) could be any of the stimuli in the set (**B**). After S2 offset, the subject was required to make either a “yes” or a “no” response (to indicate target detection or rejection, respectively) with a vertical joystick movement (up or down). Correct responses were rewarded with liquid and incorrect responses were penalized with a timeout. **B**, Broadband noise burst (“unmodified sound”), as well as sounds that deviate from the unmodified sound along the temporal and/or spectral sound feature dimensions (x and y axes, respectively). The temporal feature is amplitude modulation (AM); sinusoidal modulation of the sound envelope. We parametrically increase the value of AM by increasing its depth. To introduce spectral variation, we change (decrease) the bandwidth of the unmodified sound (ΔBW). We parametrically increase the value of ΔBW by narrowing the width of the band-pass filter used to synthesize these sounds. Changes in spectral BW are schematized in **B** with increasing lightness relative to the unmodified sound. When subjects attend to AM, the sounds within the black box are targets; that is, AM depth > 0 . Likewise, when subjects attend BW, all the sounds in the green box are targets: $\Delta\text{BW} > 0$. Here, and throughout the text, whenever feature values are collapsed across subjects, AM and ΔBW values are given as ranks. The exact values presented to each subject were determined separately for each subject and this process is described in detail in the main text. **C**, Performance of both subjects (+ for Monkey W and \circ for Monkey U) in each condition represented as regression coefficients that correspond to the influence of each feature on subjects’ behavioral responses. An increased value of a given coefficient corresponds to an increased probability that the subject will report “yes” as the value of that feature increases. Subjects’ behavioral responses are more strongly influenced by the target feature than the distractor feature, a hallmark of feature-selective attention.

task (Fig. 2A). Subjects moved a joystick laterally to initiate a trial, after which the 400 ms S1 sound (always the 9-octave wide unmodified sound) was presented, followed by a 400 ms interstimulus interval (ISI), followed by the 400 ms S2 sound. Subjects then reported whether the S2 contained the target feature by either moving the joystick down (up) to report detection or up (down) to report no detection. The response contingencies were counterbalanced between subjects. For determining ΔBW thresholds, the S2 was either the unmodified sound or one of six sounds that differed in spectral BW from the unmodified sound. We varied the value of ΔBW from a minimum of 0.25 octaves (8.75 octave spectral BW; 440–18,780 Hz) to a maximum of 2 octaves (7 octave spectral BW; 800–10,240 Hz), with 4 other intermediate ΔBW values in between (0.5, 0.75, 1.0, 1.5 ΔBW ; 480–17,222, 520–15,792, 570–14,482, 670–12,177 Hz). Both the high- and low-frequency changes of the ΔBW sounds fall in sensitive areas of the rhesus macaque audiogram (Pfungst et al., 1978). After S2 offset, the subject used the joystick to report whether s/he detected ΔBW . Hits and correct rejections were rewarded with water or juice and misses and false alarms resulted in a penalty (5–10 s timeout). Animals were rewarded (penalized) equally for all correct (incorrect) trials regardless of trial difficulty. Threshold was determined as the ΔBW at which subjects' average sensitivity (d') was 1. Subjects' thresholds for detecting ΔBW value were similar between each other (~ 0.4 octaves from Monkey U, ~ 0.5 octaves for Monkey W). It is important to note that, with some methods, changing the spectral BW of a sound induces variations in that sound's envelope. However, we have implemented a noise synthesis method that constructs noise using a single-frequency additive technique and thereby avoids introducing envelope variations that may serve as cues for ΔBW (Strickland and Viemeister, 1997).

We used an identical procedure for determining thresholds for detecting AM: the S2 was either the unmodified sound, or 1 of 6 sounds with AM depth $>0\%$, with values from 6% to 100%. For this pilot experiment, we presented AM at a fixed frequency of 30 Hz. For the full experiment, in which we presented AM at other frequencies, we only used frequencies for which the established average detection thresholds for rhesus macaques were similar to those found for 30 Hz AM noise (O'Connor et al., 2011).

Feature attention task. For feature attention (e.g., during recording sessions), we presented only three values of each feature because we presented sounds that varied along both feature dimensions and we needed to keep the size of our stimulus set small enough to obtain reasonable power for analyzing neural data (Fig. 2B). We also reduced the stimulus space by presenting only a subset of the possible co-modulated stimuli (stimuli with modulation in each feature dimension). Within each recording session, we presented 13 total stimuli (Fig. 2B; the blank diagonals indicate co-modulated sounds that were not used in the task). To equilibrate difficulty between the two features, we presented values of each feature so that one was near threshold, one was slightly above, and one far above threshold. For Monkey U, ΔBW values were 0.375, 0.5, and 1 octave and AM depth values were 28%, 40%, and 100%. For Monkey W, ΔBW values were 0.5, 0.75, and 1.5 octaves and AM depth values were 40%, 60%, and 100%. For all analyses in which data are collapsed across subjects, ΔBW values and AM values are presented as ranks (ΔBW_{0-3} and AM_{0-3}) (e.g., Fig. 2B). Within a given session, we presented AM sounds at only a single frequency. Across sessions, we used a small range of frequencies, for which behavioral detection thresholds were similar to 30 Hz (15, 22, 30, 48, and 60 Hz). The AM frequency was selected randomly each day.

All sounds were 400 ms in duration (5 ms cosine ramped at onset and offset). We have reported our sound generation methods previously (O'Connor et al., 2011). Briefly, sound signals were produced using an in-house MATLAB program and a digital-to-analog converter (Cambridge Electronic Design [CED] model 1401). Signals were attenuated (TDT Systems PA5, Leader LAT-45), amplified (RadioShack MPA-200), and presented from a single speaker (RadioShack PA-110 or Optimus Pro-7 AV) positioned approximately 1 m in front of the subject centered at the interaural midpoint. Sounds were generated at a 100 kHz sampling rate. Intensity was calibrated across all sounds (Bruel Kjaer model 2231) to 65 dB at the outer ear.

The task structure for feature attention was identical to that for the single-feature task described in the section "Single-feature threshold determination" except that the S2 could be any of the 13 sounds in the set. The subject was cued visually via an LED above the speaker as to which feature to attend (green or red light, counterbalanced between subjects). Moreover, each block began with 60 "instruction" trials in which only the unmodified sound and sounds only containing the target feature were presented (i.e., sounds containing the distractor feature were not presented). Subjects were to respond with a "yes" (up or down joystick movement) on any trial in which the attended feature was presented (Fig. 2B: sounds in black brackets during Attend AM and sounds in green brackets during Attend BW). Otherwise, the correct response was "no." The S2 was the unmodified sound on 25% of trials, co-modulated on 25% of trials, contained only ΔBW on 25% of trials, and only AM on 25% of trials. Sounds in the set were presented pseudorandomly such that, over sets of 96 trials, the stimulus set was presented exhaustively (including all four random noise seeds). Block length was variable, based in part on subjects' performance, to ensure sufficient correct trials for each stimulus. Not including instruction trials, block length was at least 180 trials and at most 360 trials to ensure that subjects performed in each attention condition at least once during the experiment. Subjects could perform each attention condition multiple times within a session. Only sessions in which subjects completed at least 180 trials per condition (excluding instruction trials) were considered for analysis in this study. We analyze neural data from both correct and incorrect trials in this study; excluding error trials makes no qualitative difference in the results because they make up a relatively small fraction of the overall trials.

Recording procedures. All recordings took place within a sound-attenuated, foam-lined booth (IAC: 9.5 \times 10.5 \times 6.5 feet) while subjects sat in an acoustically transparent chair (Crist Instruments). We advanced three quartz-coated tungsten microelectrodes (Thomas Recording, 1–2 M Ω ; 0.35 mm horizontal spacing; variable, independently manipulated vertical spacing) vertically to the superior surface of the temporal lobe. During electrode advancement, we isolated sound-responding neurons as follows. We presented sounds to the subject while the experimenter monitored neural responses to various sounds, including the tested stimuli (described above) and other natural and synthetic sounds. When neural signals from the electrodes exhibited responsiveness to sound, we then attempted to isolate single neurons. When at least one single neuron was well isolated from at least one of the electrodes, we measured spike count responses for the two sound features while the animal sat passively awake. We presented at least 10 repetitions each of the following stimuli: the unmodified sound, each of the 3 ΔBW stimuli, as well as 100% depth AM across the range of frequencies (15, 22, 30, 48, and 60 Hz). After we measured the responses of the isolated neurons, we cued the subject via cue light to begin the task and continued recording throughout task performance. When possible, we also measured responses to the tested stimuli after task performance in part to ensure stability of electrodes throughout the recording. Only recordings in which neurons were well isolated for at least 180 trials within each condition (excluding instruction trials) are included in this report.

We estimated our recording locations using established measures of neural responses to pure tones (Merzenich and Brugge, 1973; Petkov et al., 2006). Across recordings, we mapped neurons' best pure tone frequency and drew boundaries around A1 based on the tonotopic frequency gradient and reversal (caudal–rostral axis), the width of neurons' frequency response areas (medial–lateral axis), and response latency. Recordings were assigned to their putative cortical fields *post hoc*. Here, we present data from 92 single neurons assigned to A1.

Extracellular signals were amplified (AM Systems model 1800), band-pass filtered between 0.3 Hz and 10 kHz (Krohn-Hite 3382), and then converted to a digital signal at a 50 kHz sampling rate (CED model 1401). Contributions of single units to the signal were determined offline using principal components analysis-based spike sorting tools from Spike2 (CED). Spiking activity was at least 4–5 times the background noise level. Fewer than 0.1% of spike events assigned to single unit clusters fell within a 1 ms refractory period window. In the present study, we report data from 92 individual neurons (57 from Monkey W, 35 from Monkey U)

and 434 simultaneously recorded pairs of neurons (274 from Monkey W, 160 from Monkey U).

Analysis of single neuron feature tuning. Spike counts (SCs) were calculated over the entirety of each 400 ms stimulus. Each neuron was categorized as “increasing” or “decreasing” for each feature based on its SC in response to the presence of a feature (ΔBW or depth of AM). To characterize a neuron’s response function, we calculated a feature-selectivity index (FSI) for each feature. This index characterizes a neuron’s change in spike count in the presence of a feature, normalized over its full spike count distribution across both features. It is calculated as follows:

$$FSI_{AM} = (\overline{SC}_{AM>0\Delta BW_0} - \overline{SC}_{AM_0\Delta BW_0}) / (\overline{SC}_{AM>0\Delta BW_0} + \overline{SC}_{AM_0\Delta BW_0})$$

$$FSI_{\Delta BW} = (\overline{SC}_{AM_0\Delta BW>0} - \overline{SC}_{AM_0\Delta BW_0}) / (\overline{SC}_{AM_0\Delta BW>0} + \overline{SC}_{AM_0\Delta BW_0})$$

where \overline{SC}_x is the mean SC in response to the class of stimuli designated by the subscript. We performed a Kruskal–Wallis rank-sum test between the distribution of SCs in response to sounds with a feature value above 0 (e.g., $AM_{>0}$ for AM) and the distribution of SCs in response to sounds with a feature value of 0 (e.g., AM_0 for AM) to determine the significance of each FSI for each neuron.

Analysis of r_{noise} and r_{tuning} . We measured AM r_{tuning} , ΔBW r_{tuning} and r_{noise} for each pair in each condition. We calculated r_{tuning} as the Pearson correlation between the mean spike count of each neuron in a pair to the set of four stimuli along each feature axis (each axis includes the unmodified sound). It is worth noting that Pearson correlation provides a noisy measure of the overall fit between the mean responses, it is not sensitive to relationships between means that are not linear, and it is affected by outliers. For the purposes of this study, we do not intend this correlation metric to provide an exact account of the relationship between mean responses; rather, we use it to approximate the relationship between neurons’ selectivity to a given feature. For instance, if a given neuron tends to increase its spike count in the presence of AM depth $>0\%$, then it will have a positive AM r_{tuning} with other neurons that increase their spike count in the presence of AM, and a negative AM r_{tuning} with neurons that decrease their spike count in the presence of AM. We also calculated r_{tuning} using two other methods, Spearman correlation and the weighted least-squares coefficient on the normalized mean spike counts, both of which are more robust to outliers than Pearson correlation. The results obtained with those metrics (not reported here) are qualitatively quite similar to those obtained using Pearson correlation and do not alter our major findings or conclusions.

We calculated r_{noise} across all presentations of the unmodified stimulus separately within each condition to allow for analyses of r_{noise} due to changes in attention. We calculated r_{noise} using spike counts calculated over the entire 400 ms S2 presentation epoch within each trial. We selected the unmodified stimulus (i.e., S2 = S1) as the stimulus upon which to calculate r_{noise} due to the fact that its behavioral “meaning” and associated behavioral response were consistent across contexts. These factors have been shown to modulate nonsensory variability in A1 neurons and this nonsensory variability may modulate r_{noise} because it is, on average, shared across neurons (Nienborg and Cumming, 2010; Niwa et al., 2012). Moreover, r_{noise} is also often affected by stimulus drive, so restricting our analyses to responses to a single stimulus type should increase the reliability of our r_{noise} estimate. We calculated r_{noise} over an average of 85 trials of spike count responses to the unmodified sound (range 48–168). When we calculated r_{noise} by collapsing across all stimuli, the results that we observed were qualitatively similar and statistically indistinguishable from the analyses conducted when only the unmodified sound was used to calculate r_{noise} . This implies that the effect of attention on r_{noise} is, on average, consistent across all stimuli. Those analyses are not presented here. To calculate r_{noise} collapsed across stimuli, spike counts within each stimulus were z-scored, combined into a single vector of normalized spike counts within each condition, and then the Pearson correlation between these vectors was calculated. Low spike counts reduce the reliability of r_{noise} estimates. We therefore excluded 16 pairs due to low average pairwise spike counts (<1 spike/s).

Results

Rhesus macaques perform an auditory feature-selective attention task

We developed a feature-selective attention task in which subjects must listen for the occurrence of a target feature while ignoring a concurrent distractor feature (Fig. 2A,B). The task simulates feature-based tracking of relevant information in a crowded environment (Shamma et al., 2011; Mesgarani and Chang, 2012; Woods and McDermott, 2015). We assessed subjects’ performance by calculating, within each attention condition within each session, the influence of each of the two features on subjects’ responses. To do so, we use a binomial logistic regression to calculate coefficients that quantify the degree to which the value of a given feature affects the probability of a “yes” response, according to the following logistic function:

$$P(\text{“Yes”}) = 1 / (1 + e^{-reg})$$

$$reg = \beta_{AM}(VAL_{AM}) + \beta_{BW}(VAL_{BW})$$

$$+ \beta_{AM*BW}(VAL_{AM} * VAL_{BW}) + \alpha$$

where VAL_{AM} and VAL_{BW} are the ranked values of AM and ΔBW (from 0–3), respectively, β is the coefficient for the value term, and α is an offset term to capture response bias. Intuitively, as the influence of a given feature on the subject’s probability of responding “yes” increases, the value of the coefficient will increase; when a given feature has no impact on the behavioral response, the value of the coefficient will be ~ 0 . We calculated β_{AM} and β_{BW} within each condition for each recording session [$n = 16$ sessions; 9 Monkey U (circles, Fig. 2C) and 7 Monkey W (crosses, Fig. 2C)] and then compared these coefficient values collapsed across sessions (Fig. 2C). We found a significant increase in the average coefficient value for the target versus the distractor feature ($t_{(31)} = 7.89$, $p = 6.59e^{-9}$) revealing that animals were using feature-selective attention, although imperfectly (as evidenced by nonzero coefficients for the distractor feature). Specifically, this nonzero distractor coefficient arises because animals increase their likelihood to make a “yes” response in the presence of the distractor feature for both single feature stimuli (Fig. 1B, marginals) and co-modulated stimuli. We emphasize that this analysis is meant to estimate the strength of selective attention, rather than assess task accuracy. Both the analysis and behavioral results are similar to those from a feature-selective attention study in the visual system (Sasaki and Uka, 2009).

Single neurons exhibit heterogeneous selectivity for the two sound features

We measured single neurons’ selectivity to each feature by analyzing whether the presence of the feature significantly changed a given neuron’s spike count. We assessed this by performing a Mann–Whitney rank-sum test comparing spike count distributions in response to the unmodified sound with spike count distributions in response to each of the single-feature AM and ΔBW sounds (Fig. 2B, marginals of the stimulus matrix) in each of the two conditions. For each neuron, we performed 12 tests and used the false discovery rate method to correct for multiple comparisons at an α level of 0.05 (Benjamini and Yekutieli, 2001).

We found that 73/92 of recorded neurons exhibited selectivity (either increasing or decreasing spike count) to at least one feature. Only one of these 73 neurons exhibited increasing spiking to at least one value of a feature and decreasing spiking to another value of the same feature; therefore, we classify selectivity as either “increasing” or “decreasing.” In addition, although neurons

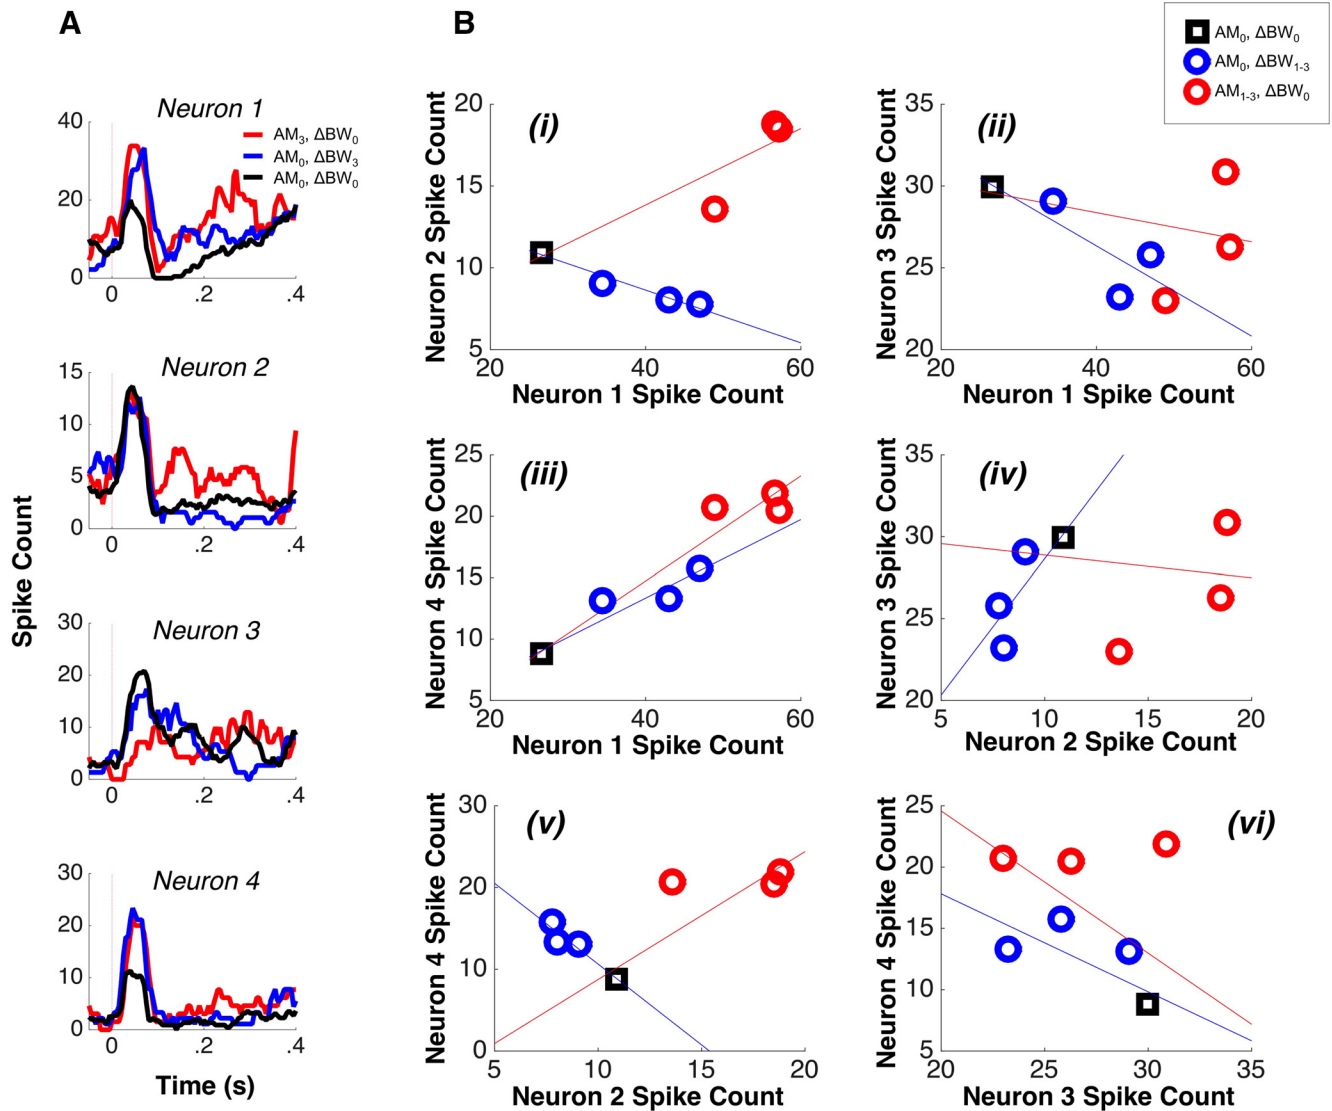


Figure 3. Heterogeneous selectivity for AM and ΔBW leads to heterogeneous feature-specific tuning correlations. **A**, Peristimulus time histograms of four single neurons in response to three stimuli: (1) the largest value AM-only stimulus (red; $AM_3, \Delta BW_0$), (2) the largest value ΔBW only stimulus (blue; $AM_0, \Delta BW_3$), and (3) the unmodified sound (black; $AM_0, \Delta BW_0$). We display only the largest value stimulus for each feature for clarity. Neuron 1 (top) exhibits an increased spike count for both AM and ΔBW ; neuron 2 increases for AM, but decreases for ΔBW ; neuron 3 decreases both AM and ΔBW ; neuron 4 increases for both (similar to neuron 1). Whereas in **A**, we only used the largest AM and ΔBW values, in **B**, we show responses to all single-feature stimuli because these were the stimuli used to calculate feature-specific r_{tuning} . In each panel is plotted the mean spike count for each of the three single feature sounds (red and blue circles), as well as the unmodified sound (black square). The lines correspond to weighted least-squares regression fits and illustrate the sign of r_{tuning} (the upper and lower bounds of each fit’s 95% confidence interval have the same sign). There is one panel for each of the six possible pairings of the four neurons in **A**. Pairs can have positive r_{tuning} for both features (**Bii**; similar to Fig. 1*D*), negative r_{tuning} for both features (**Bij** and **Bvi**; similar to Fig. 1*E*), or positive r_{tuning} for one and negative r_{tuning} for the other (**Bi**, **Biv**, and **Bv**; similar to Fig. 1*C,F*).

did show some change in spike count between conditions, the valence of selectivity was constant across conditions. The responses of four example neurons are shown in Figure 3*A*. As illustrated, neurons may exhibit increased spike count to both AM and ΔBW (Fig. 3*A*, Neurons 1, 4), they may exhibit decreased spike count to both (Fig. 3*A*, Neuron 3), or they may increase for one feature and decrease for the other (Fig. 3*A*, Neuron 2).

The distributions of selectivity profiles for each feature are summarized in Figure 4. For AM, neurons overwhelmingly tend toward increasing selectivity (median $FSI_{AM} = 0.095$; sign-rank test, $p = 1.21e-9$) and, for ΔBW , neurons are approximately evenly distributed between increasing and decreasing selectivity (median $FSI_{\Delta BW} = 0.010$; sign rank test, $p = 0.327$) (Fig. 4*A,B*). There is no significant relationship between a neuron’s AM selectivity and ΔBW selectivity (Fig. 4*A*; $r_{(91)} = 0.17$, $p = 0.102$).

Figure 4*B* illustrates the frequency with which we observed each of nine possible classes of neurons (increasing, decreasing, or nonselective for each feature). In our sample, we found multiple neurons within each of these classes. Such a diverse population allows for the possibility of heterogeneous feature-specific joint tuning, which is crucial to the type of simultaneous modulation of target and distractor detection illustrated in Figure 1.

Neural pairs exhibit heterogeneous feature-specific r_{tuning} in A1

We calculated feature-specific r_{tuning} for both AM and ΔBW for each of the 418 pairs, as described in the Materials Methods. We found that approximately equal numbers of pairs exhibited opposite (188/418) versus the same (230/418) r_{tuning} sign for AM




Figure 4. Distributions of feature-specific single neuron tuning. In **A**, we plot each neuron's AM selectivity against each neuron's Δ BW selectivity and show the marginal distributions for single feature selectivity. For AM, the majority of neurons exhibit AM selectivity above 0 (i.e., primarily Increasing AM functions), whereas Δ BW selectivity is approximately equally distributed about 0 (i.e., \sim equal occurrence of increasing and decreasing Δ BW functions). We found no significant relationship between a neuron's AM selectivity and Δ BW selectivity, although across the population, there was a preponderance of neurons with increasing selectivity for both AM and Δ BW. **B**, We classify neurons' selectivity for each feature as increasing (Inc.), decreasing (Dec.), or nonselective (n.s.), for both AM and Δ BW. There is quite some heterogeneity in feature-specific single neuron selectivity, with approximately half of the neurons (45/92) exhibiting differential selectivity for the two features.

and Δ BW. The pairs with opposite AM and Δ BW r_{tuning} are essential for allowing the neural population to exploit a mechanism by which selective shifts in r_{noise} can simultaneously enhance target detection and suppress distractor detection (Fig. 1C–F). The feature-specific r_{tuning} between the four example neurons in Figure 3A is shown in Figure 3B. Each point corresponds to the joint mean for a given stimulus and the lines correspond to weighted least-squares fits, which illustrate the r_{tuning} sign. Among this small population of four neurons with diverse feature tuning, three of the six resulting pairs exhibit feature-specific r_{tuning} profiles that would allow for efficient attentional modulation of r_{noise} to support feature-selective attention.

The frequency distribution of feature specific r_{tuning} across the 418 pairs is displayed in Figure 5. As we have reported previously, AM r_{tuning} tends to be positive in A1 (Downer et al., 2015). Conversely, Δ BW r_{tuning} appears to be approximately evenly distributed between positive and negative. There is no significant relationship between AM and Δ BW r_{tuning} ($r_{(417)} = -4.8e-4$, $p = 0.99$; Fig. 5A). We observe approximately equal numbers of pairs within each of the four r_{tuning} classes (positive or negative for each feature; Fig. 5B). Figure 5C shows the distribution of r_{tuning} differences (AM $r_{\text{tuning}} - \Delta$ BW r_{tuning}). We found a large number of pairs with very large differences between AM and Δ BW r_{tuning} .

Feature-selective attention modulates noise correlations between pairs with opposite feature-specific r_{tuning} signs

We hypothesized that feature-selective attention may operate by modulating r_{noise} between pairs with opposite feature-specific r_{tuning} signs, as in Figure 1, C and F, because a single shift in r_{noise} between these pairs can efficiently enhance target detection and suppress distractor detection. Namely, the SR predicts that attention should decrease r_{noise} when the r_{tuning} for the target feature is positive and increase r_{noise} when the r_{tuning} for the target feature is negative (Fig. 1C). In the absence of distractors, such modulations have been argued to enhance coding (Romo et al., 2003; Averbeck et al., 2006; Jeanne et al., 2013; Hu et al., 2014; Ruff and Cohen, 2014a). We therefore compared the average r_{noise} between attention conditions and tested for influences of both

feature-specific r_{tuning} similarity (whether the pair had the same or opposite r_{tuning} sign for AM and Δ BW) and the sign of r_{tuning} of the target feature. We used an independent-means ANOVA for these analyses, with two levels each of three factors (factor 1: AM r_{tuning} sign, positive or negative; factor 2: Δ BW r_{tuning} sign, positive or negative; and factor 3: attention condition, AM or BW).

Our results support adaptive modulation of r_{noise} to enhance representation of the attended feature and suppress representation of the distractor feature. We found no main effect of attention on r_{noise} across the population ($F_{(1,1,835)} = 0.01$, $p = 0.92$), but we did find a significant interaction between attention condition and target r_{tuning} when collapsing across distractor r_{tuning} ($F_{(1,1,835)} = 5.01$, $p = 0.025$). When we analyzed the nature of this interaction more deeply, we found that feature-selective attention shifts r_{noise} based not only on the target r_{tuning} , but also on the distractor r_{tuning} , as our illustrations in Figure 1, C and F, would predict to enhance performance. Figure 6 shows r_{noise} values in each attention condition across four groups (classified based on both AM and Δ BW r_{tuning}). Among pairs with opposite feature-specific r_{tuning} signs (Fig. 6A, D), we observed that r_{noise} is lower when the r_{tuning} of the target feature is positive and higher when r_{tuning} of the target feature is negative ($F_{(1,1,375)} = 8.64$, $p = 0.003$; one-sided pairwise Tukey's HSD tests, $p < 0.05$). Among pairs with the same feature-specific r_{tuning} signs (Fig. 6B, C), we observed no effect of attention on r_{noise} ($F_{(1,1,459)} = 0.09$, $p = 0.76$).

We also analyzed these data using multivariate linear regression, treating r_{tuning} as a continuous variable, to ensure that our findings are robust across the range of r_{tuning} values. Similar to the ANOVA, the regression analysis revealed a significant interaction between attention and r_{tuning} . Namely, attention modulates the relationship between r_{tuning} and r_{noise} depending on both target and distractor r_{tuning} ($t_{(417)} = 3.65$, $p = 0.0003$). To better visualize this interaction, we show the effect of attention on r_{noise} (attend AM – attend BW) across the range of AM and Δ BW r_{tuning} values (Fig. 6E). We calculated the r_{noise} difference (attend AM – attend BW) between attention conditions for each pair, calculated the average r_{noise} difference value within bins of width 0.33 for both AM and Δ BW r_{tuning} , and then smoothed these




Figure 5. Distributions of feature-specific r_{tuning} . In **A**, we plot each pair's AM r_{tuning} against each neuron's Δ BW r_{tuning} and show the marginal distributions for single feature r_{tuning} . Similarly to single neuron tuning, we find that most pairs exhibit positive AM r_{tuning} , whereas Δ BW is approximately evenly distributed between positive and negative r_{tuning} . **B**, Pairs classified according to their AM r_{tuning} sign and Δ BW r_{tuning} sign, yielding four groups. We show the frequency of occurrence of each of those groups because our hypothesized potential mechanism for modulating r_{noise} to support feature-selective attention requires pairs with opposite AM r_{tuning} and Δ BW r_{tuning} signs. We found that approximately equal numbers of pairs exhibit opposite AM r_{tuning} and Δ BW r_{tuning} signs as exhibit identical AM r_{tuning} and Δ BW r_{tuning} signs. In **C**, we show the distribution of r_{tuning} differences, calculated as AM r_{tuning} - Δ BW r_{tuning} . We see a small peak near 0, indicating many pairs for which AM and Δ BW r_{tuning} is very similar, as well as a large proportion of pairs at each tail, indicating many pairs for which AM and Δ BW r_{tuning} are very different. The large peaks on the right and left of the distribution correspond to pairs for which r_{tuning} sign differs between AM and Δ BW.

averages with a Gaussian kernel of width 0.5. The heat map in Figure 6E shows increased r_{noise} difference values in the upper left quadrant and decreased r_{noise} difference values in the lower right quadrant, with average r_{noise} difference values near 0 in the upper right and lower left quadrants. However, it is worth noting that, in the lower left quadrant, we found a local increase in r_{noise} difference in pairs with slightly negative AM r_{tuning} and strongly negative Δ BW r_{tuning} . This local anomaly arises due to effects in ~ 20 pairs and therefore has little impact when folded into the coarser category shown in Figure 6C.

These results suggest that feature-selective attention can operate by selectively modulating r_{noise} in specific subnetworks within sensory cortex. The observed effects are consistent with feature-selective attention operating by simultaneous enhanced detection of the target feature and decreased detection of the distractor via a single shift in r_{noise} .

Potential impact of changes in spike count on r_{noise}

Previous studies have shown both that r_{noise} and spike count covary (de la Rocha et al., 2007; Cohen and Maunsell, 2009; Hu et al., 2014; Zylberberg et al., 2016) and that attention affects spike counts in A1 (Sutter and Shamma, 2011; Osmanski and Wang, 2015). Therefore, we analyzed the relationship between the effect of attention on spike count and the effect of attention on r_{noise} to determine whether attention uniquely affects r_{noise} or, alternatively, if the observed effects can be attributed to attention-mediated changes in spike count. We calculated the spike count of a pair separately for each stimulus within each condition as the geometric mean between each neuron's mean spike count. Figure 7A shows the relationship between the attention-related changes in spike count and attention-related changes in r_{noise} across all pairs (attend AM - attend BW). Using a Pearson correlation, we found a nonsignificant negative relationship between the two




Figure 6. Attention-related effects on r_{noise} depend on both target and distractor r_{tuning} . **A**, Average r_{noise} (bars represent means, whiskers represent \pm SEM here and throughout the figure) between attention conditions for pairs with negative $\text{AM } r_{\text{tuning}}$ and positive $\Delta\text{BW } r_{\text{tuning}}$ (Figure 1, **C** and **F**, would predict attention to enhance task performance by increasing r_{noise} between these pairs during AM attention (because this may concurrently enhance AM detection and suppress ΔBW detection) and decreasing r_{noise} during BW attention (because this may concurrently suppress AM detection and enhance ΔBW detection). We found that r_{noise} was significantly higher during AM attention (black) and lower during BW attention (green), consistent with our predictions ($p < 0.05$). For the pairs in **B** and **C**, $\text{AM } r_{\text{tuning}}$ and $\Delta\text{BW } r_{\text{tuning}}$ are the same sign, either both positive (**B**) or both negative (**C**). We predicted that, for these pairs, shifts in r_{noise} will have the same effect on both target and distractor detection and therefore may provide no benefit to task performance. We found no effect of attention on r_{noise} in either group ($p > 0.05$). In **D**, these pairs are similar to those in **A** in that their AM and $\Delta\text{BW } r_{\text{tuning}}$ have opposite signs. Our finding here is similar to what we observed in **A**: when the attended feature (AM) has positive r_{tuning} , r_{noise} is lower than when the attended feature has negative r_{tuning} (BW) ($p < 0.05$). In **E**, we used smaller r_{tuning} bins to show how the effect of attention on r_{noise} depends on r_{tuning} (Figure legend continues.)

variables ($r_{(417)} = -0.047, p = 0.342$). We also conducted this analysis separately for each of the feature-specific r_{tuning} groups (Fig. 7B–E) and found no significant relationship between attention-related changes in spike count and attention-related changes in r_{noise} in any of the 4 groups (Fig. 7B: $r_{(79)} = 0.026, p = 0.817$; Fig. 7C: $r_{(159)} = -0.062, p = 0.444$; Fig. 7D: $r_{(69)} = -0.094, p = 0.41$; Fig. 7E: $r_{(107)} = -0.036, p = 0.738$). Therefore, we could find no evidence that the observed effects of attention on r_{noise} could be explained by changes in spike count.

Attention-related r_{noise} changes during the ISI

Differences between stimulus-driven and non-driven r_{noise} have been reported in multiple studies across a range of cortical structures (Kohn and Smith, 2005; Cohen and Maunsell, 2009; Churchland et al., 2010; Doiron et al., 2016; Nandy et al., 2017). Stimulus-related changes in r_{noise} provide insights into the neural mechanisms that give rise to r_{noise} . In the present study, we have observed highly selective attention-related changes in r_{noise} that may be due to changes to feedforward inputs to A1, changes in connections among neurons within A1, changes in the structure of top-down inputs, or some combination thereof. To begin to address the source of the observed stimulus-driven r_{noise} effects, we analyzed the effect of attention on r_{noise} during the ISI between S1 and S2 (Fig. 1A, ISI).

The results of these analyses are shown in Figure 8. Interestingly, the effect of attention on nondriven r_{noise} is quite similar to the effects of attention on driven r_{noise} (Fig. 6). Factorial ANOVA reveals no significant main effect of attention ($F_{(1,1,835)} = 0.12, p = 0.73$), but a trending nonsignificant interaction effect between attention and target r_{tuning} when collapsing across distractor r_{tuning} ($F_{(1,1,835)} = 3.47, p = 0.0625$). Moreover, we also find that for pairs with opposite feature-specific r_{tuning} signs (Fig. 8A, D), attention inverts the relationship between r_{tuning} and r_{noise} , similar to the observed effects on driven activity ($F_{(1,1,375)} = 8.14, p = 0.0044$) and elicits no interaction between target r_{tuning} and attention on r_{noise} in pairs with the same feature-specific r_{tuning} signs (Fig. 8B, C; $F_{(1,1,459)} = 0.02, p = 0.897$). We also represent these results in a more continuous manner in Figure 8E (similar to Fig. 6E). Although these results are consistent with what we found during driven neural activity, we did observe some differences. For instance, there was an interaction between AM r_{tuning} and attention wherein attending to AM decreases average r_{noise} when AM r_{tuning} is positive (Fig. 8B, D) and exhibits a nonsignificant increase in r_{noise} when AM r_{tuning} is negative (Fig. 8A, C) ($F_{(1,1,459)} = 4.6, p = 0.02$). Moreover, the pairwise test of the effect of attention between pairs with positive AM and ΔBW r_{tuning} reveals a reduction in r_{noise} during the attend AM condition (Fig. 8B). Finally, we find a significant main effect of stimulus on r_{noise} : consistent with previous reports, stimulus-driven r_{noise} is significantly lower than nondriven r_{noise} ($F_{(1,1,835)} = 32.34, p = 1.52 \times 10^{-8}$) regardless of r_{tuning} or attention condition, consistent with previous reports (Kohn and Smith, 2005; Churchland et al., 2010; Snyder et al., 2014).

←

(Figure legend continued.) across the range of values. We calculated r_{noise} difference values (attend AM – attend BW), calculated the average r_{noise} difference within bins of width 0.33 for both AM and ΔBW r_{tuning} , and then smoothed the binned averages. We observed increased r_{noise} difference values (r_{noise} higher during attend AM) in the upper left quadrant and decreased r_{noise} difference values (r_{noise} lower during attend AM) in the lower right quadrant, consistent with our findings using coarse categorization using r_{tuning} sign. In A–D, inset histograms show the distribution of r_{noise} in each condition. All statistical analyses performed are described in the main text.

These results show that highly selective attention-related effects of r_{noise} in A1 are mainly consistent throughout a trial and are distinct from nonselective stimulus-related decreases in r_{noise} . Together, these findings support distinct sources of r_{noise} in A1 and suggest that attention can target distinct network mechanisms selectively to modulate r_{noise} according to specific task demands. Frontal cortex inputs to A1 appear particularly poised to modulating r_{noise} selectively relative to r_{tuning} because these inputs have been shown to shift the relationship dynamically between r_{tuning} and r_{noise} in A1 pairs (Winkowski et al., 2013).

Discussion

Perception is a fundamentally constructive process in which attention plays a critical role (Rensink et al., 1997). One proposed role for attention is in the selection of the sensory features that compose perceptual objects (Treisman, 1998). Although the study of feature-selective attention has a rich history, its mechanisms remain obscure. We present a novel mechanism by which different features of the same object can be adaptively selected or discarded based on the ongoing sensory needs of the organism. Importantly, this mechanism extends the current framework for studying how attention affects neural populations in two ways. First, we present an extension of the SR that is of potential value in understanding how populations may encode sensory information when more than one feature or object is present, an area of study that is currently lacking (Orhan and Ma, 2015). Second, although effects of feature attention on r_{noise} have been reported previously (Cohen and Maunsell, 2011), we present the first known mechanism for how attention may shift r_{noise} to affect both target and distractor sensitivity simultaneously, rather than simply enhancing detection of a single object or feature in isolation.

Across modalities, early sensory neurons represent constituent features of stimuli (Hubel and Wiesel, 1962; Merzenich and Brugge, 1973; Fuss and Korsching, 2001), whereas later stages tend to integrate over multiple features (Barlow, 1972). Researchers debate whether the features that constitute perceived objects must be either selected or rejected at early stages or if both relevant and irrelevant features are passed from early to late stages (Chun et al., 2011). There is strong evidence that multiple mechanisms may be at play and that the brain can change its strategy based on task demands (Chen et al., 2012; David et al., 2012; Mante et al., 2013; Bacigalupo and Luck, 2015; Woods and McDermott, 2015). In the auditory domain, feature attention supports “cocktail party” listening, in which sound features such as the fundamental frequency of a speaker’s voice can be used to listen selectively to one speaker over another (Cherry, 1953; Shamma et al., 2011; Mesgarani and Chang, 2012). Feature attention has long been understood to constitute one of the primary skills required to solve the cocktail party problem, but has received less attention than other skills such as preattentive sound segregation (Shinn-Cunningham, 2008; McDermott, 2009). Our results suggest an efficient mechanism by which auditory neurons can prioritize relevant features to facilitate tracking an attended sound source to solve the cocktail party problem. Moreover, our findings contribute to a growing literature suggesting that heterogeneous joint tuning in a population allows for efficient coding of complex signals (Pagan et al., 2013).

Relationship to previous studies of the effect of behavior

on r_{noise}

Our results support adaptive modulation of r_{noise} among specific neural networks as a property of effective listening. The dynamic




Figure 7. Attention-related effects on r_{noise} are independent of attention-related changes in spike count. **A**, For each pair in our sample, we plotted the spike count difference against the r_{noise} difference (attend AM — attend BW). We found no significant relationship between the effect of attention on spike count and the effect of attention on r_{noise} ($p > 0.05$). **B–E**, Same analysis as in **A**, but for pairs from each of the four r_{tuning} categories. For none of these groups did we find any significant relationship between attention-related r_{noise} effects and attention-related spike count effects.

selection and suppression of sound features is a crucial aspect of this problem, the neural implementation of which has been elusive. Our results suggest that network noise can be reshaped by the ongoing sensory demands of the listener to simultane-

ously enhance relevant and suppress irrelevant representations. Importantly, our results suggest that this can be done efficiently and without changes in the level of overall neural activity.




Figure 8. Attention-related effects on r_{noise} during spontaneous spiking. Conventions in this figure are same as for Figure 6. Similarly to r_{noise} measured during stimulus-driven activity (Fig. 6), we found a significant interaction among attention, target r_{tuning} , and distractor r_{tuning} . Namely, in pairs with opposite r_{tuning} signs (**A**, **D**), there is higher r_{noise} when attending the feature with negative r_{tuning} and lower r_{noise} when attending the feature with positive r_{tuning} . We observe one compelling difference between the effect of r_{noise} during spontaneous spiking and driven spiking: in **B** (pairs with positive AM and $\Delta\text{BW } r_{\text{tuning}}$), we found that attending to AM reduces r_{noise} during spontaneous spiking, but not during stimulus driven activity. However, the (null) effect observed in **C** mirrors that observed during driven activity. In **E**, we reiterate these results using narrower bins. We calculated the average r_{noise} difference value (attend AM — attend BW) within bins of width 0.33 and then smoothed these binned averages. The matrix in **E** exhibits similar structure to that in Figure 6E, illustrating similar effects of attention on r_{noise} during driven and spontaneous spiking. In **A–D**, inset histograms show the distribution of r_{noise} in each condition.

There has been recent interest in the optimal relationship between r_{tuning} and r_{noise} . That r_{noise} could be modulated selectively according to r_{tuning} has only recently been shown (Jeanne et al., 2013; Ruff and Cohen, 2014a; Downer et al., 2015). Previously, many studies have revealed that, in situations of increased sensory demand, r_{noise} would globally decrease in a population (Cohen and Maunsell, 2009; Mitchell et al., 2009; Gu et al., 2011; Herrero et al., 2013; Issa and Wang, 2013; Ruff and Cohen, 2014b; Nandy et al., 2017). According to the SR, the effect of r_{noise} on coding accuracy depends on the sign of r_{noise} relative to the sign of r_{tuning} , namely that coding is enhanced when they have opposite signs but impaired when they have the same sign (Oram et al., 1998; Abbott and Dayan, 1999; Averbeck et al., 2006). Therefore, according to the SR, the global reduction of r_{noise} across a population of neurons could enhance coding only if a large majority of neurons exhibited similar tuning (yielding mainly positive r_{tuning}). Widespread similar tuning may very well have been the case in many studies. Recently, Ruff and Cohen (2014a) presented evidence that a more demanding attention task, in which the activity of distinct neural populations code for different perceptual decisions, can increase or decrease r_{noise} in a manner consistent with the SR. Similarly, we show that attention can either increase or decrease r_{noise} depending on pairs' r_{tuning} , as the SR predicts. Moreover, our results provide the first evidence that we know of that r_{noise} can be modulated, not only based on target r_{tuning} , but also distractor r_{tuning} .

Multiple theoretical studies have pointed out that the effect of r_{noise} on neural coding may depend on other factors in addition to those described by the SR, including heterogeneity of tuning functions (Ecker et al., 2011), stimulus dependence of r_{noise} (Zylberberg et al., 2016), r_{noise} with a very large magnitude (Hu et al., 2014), and the slopes of neurons' tuning curves (Moreno-Bote et al., 2014). These studies tend to assume different population codes than those that might be at work coding AM and Δ BW in A1. Namely, AM depth and Δ BW are largely encoded by monotonic changes in spike count in A1, whereas most studies of the effect of r_{noise} on population coding assume bell-shaped feature coding. In addition, these studies tend to measure population coding accuracy in terms of discrimination whereas we focus on feature detection. We find that neurons exhibit monotonic, as opposed to bell-shaped, feature tuning for the features we present in the current study; this difference in single neuron coding strategy can significantly change how a population's activity is decoded. In addition, as we mentioned earlier, most studies of population coding and the effect of r_{noise} thereupon have not considered cases in which a relevant feature must be detected in the presence of irrelevant features, so our study presents a step toward understanding the role of r_{noise} in more naturalistic sensory settings. Moreover, because r_{noise} cannot be manipulated precisely in a behaving animal, evidence for a role of r_{noise} in perception is only correlative. Although we believe that our findings support the SR as a valid model of the effect of r_{noise} on coding accuracy, future work is needed to test these ideas directly both theoretically and experimentally.

Feature attention and the cocktail party problem

Feature attention likely plays a special role in listening for two primary reasons. First, the spatial acuity of the auditory system is poor relative to the visual system (Alais and Burr, 2004), so the ability to attend to nonspatial features is likely more important in sound perception than, say, visual or tactile perception. Second, the soundscape undergoes constant change such that a given feature may uniquely define a target sound in one instance but not

the next (Woods and McDermott, 2015). Therefore, experiments in which listeners alternate attention between different nonspatial sound features provide insights into one of the most challenging aspects of listening in crowded environments.

A few studies examining the neural basis of nonspatial feature attention during cocktail party listening have been performed in humans. In one prominent study, Mesgarani and Chang (2012) presented monaural mixtures of two talkers, one of which subjects attended to. Using electrocorticography (ECoG) to measure brain activity, they found that local field potential activity in nonprimary (parabelt) auditory cortex selectively represented the attended speaker relative to the unattended speaker. Moreover, they provide evidence that the auditory cortex selectively represents the spectral feature values that may best segregate the attended from the unattended speaker. Other studies in humans have examined shifts in attention between nonspatial features such as duration, pitch and loudness (for review, see Caporello Bluvas and Gentner, 2013). These studies used methods coarse in spatial or temporal resolution due to limitations inherent in measuring human brain activity. The primary finding has been that brain areas that putatively code the attended feature exhibit greater activity. Our results would not be detectable using methods without very fine spatial and temporal resolution (such as that available with microelectrode recording) and thus would not be apparent in these human studies. We argue that our results provide evidence of a novel and independent property of auditory cortical function in enhancing listening in crowded environments.

Future work should focus on bridging the gap between mesoscopic measures of human neural activity (e.g., ECoG high-gamma signals) and microscopic measures obtained in animals (e.g., single neurons) in the context of r_{noise} and attention. Although there is evidence that the high-gamma local field potential signal indexes multiunit neuron activity (Steinschneider et al., 2008; Ray and Maunsell, 2011; Chang, 2015), the direct relationship between ECoG high-gamma signals recorded at the pial surface and correlated single neuron activity recorded across the depth of cortex remains unknown. A promising and feasible avenue for future studies is the simultaneous acquisition of ECoG and single neuron data during complex attention tasks in monkeys. Heretofore, simultaneous acquisition of ECoG and single neuron data in rhesus macaque auditory cortex has been hampered by difficulties in recording from the exposed superior temporal gyrus corresponding to parabelt, where ECoG recordings are most feasible. Recent advances in the ability to record from the surface of the superior temporal gyrus in rhesus macaques provide an excellent opportunity to both address the relationship between ECoG and single neuron activity in the auditory cortex and understand hierarchical changes in the effects of attention in the auditory cortex (Kajikawa et al., 2015).

References

- Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11:91–101. [CrossRef Medline](#)
- Alais D, Burr D (2004) The ventriloquist effect results from near-optimal bimodal integration. *Curr Biol* 14:257–262. [CrossRef Medline](#)
- Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. *Nat Rev Neurosci* 7:358–366. [CrossRef Medline](#)
- Bacigalupo F, Luck SJ (2015) The allocation of attention and working memory in visual crowding. *J Cogn Neurosci* 27:1180–1193. [CrossRef Medline](#)
- Barlow HB (1972) Single units and sensation: a neuron doctrine for perceptual psychology? *Perception* 1:371–394. [CrossRef Medline](#)
- Bee MA, Micheyl C (2008) The cocktail party problem: What is it? How can

- it be solved? And why should animal behaviorists study it? *J Comp Psychol* 122:235–251. [CrossRef Medline](#)
- Benjamini Y, Yekutieli D (2001) The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics* 29:1165–1188. [CrossRef](#)
- Caporello Bluvas EC, Gentner TQ (2013) Attention to natural auditory signals. *Hear Res* 305:10–18. [Medline](#)
- Chang EF (2015) Towards large-scale, human-based mesoscopic neurotechnologies. *Neuron* 86:68–78. [CrossRef Medline](#)
- Chen X, Hoffmann KP, Albright TD, Thiele A (2012) Effect of feature-selective attention on neuronal responses in macaque area MT. *J Neurophysiol* 107:1530–1543. [CrossRef Medline](#)
- Cherry EC (1953) Some experiments on the recognition of speech, with one and with two ears. *J Acoust Soc Am* 25:975–979. [CrossRef](#)
- Chun MM, Golomb JD, Turk-Browne NB (2011) A taxonomy of external and internal attention. *Annu Rev Psychol* 62:73–101. [CrossRef Medline](#)
- Churchland MM et al. (2010) Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci* 13:369–378. [CrossRef Medline](#)
- Cohen MR, Maunsell JH (2009) Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci* 12:1594–1600. [CrossRef Medline](#)
- Cohen MR, Maunsell JH (2011) Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron* 70:1192–1204. [CrossRef Medline](#)
- David SV, Fritz JB, Shamma SA (2012) Task reward structure shapes rapid receptive field plasticity in auditory cortex. *Proc Natl Acad Sci U S A* 109:2144–2149. [CrossRef Medline](#)
- Doiron B, Litwin-Kumar A, Rosenbaum R, Ocker GK, Josić K (2016) The mechanics of state-dependent neural correlations. *Nat Neurosci* 19:383–393. [CrossRef Medline](#)
- Downer JD, Niwa M, Sutter ML (2015) Task engagement selectively modulates neural correlations in primary auditory cortex. *J Neurosci* 35:7565–7574. [CrossRef Medline](#)
- Ecker AS, Berens P, Tolias AS, Bethge M (2011) The effect of noise correlations in populations of diversely tuned neurons. *J Neurosci* 31:14272–14283. [CrossRef Medline](#)
- Fuss SH, Korsching SI (2001) Odorant feature detection: activity mapping of structure response relationships in the zebrafish olfactory bulb. *J Neurosci* 21:8396–8407. [Medline](#)
- Gu Y, Liu S, Fetsch CR, Yang Y, Fok S, Sunkara A, DeAngelis GC, Angelaki DE (2011) Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron* 71:750–761. [CrossRef Medline](#)
- Haykin S, Chen Z (2005) The cocktail party problem. *Neural Comput* 17:1875–1902. [CrossRef Medline](#)
- Herrero JL, Gieselmann MA, Sanayei M, Thiele A (2013) Attention-induced variance and noise correlation reduction in macaque V1 is mediated by NMDA receptors. *Neuron* 78:729–739. [CrossRef Medline](#)
- Hu Y, Zylberberg J, Shea-Brown E (2014) The sign rule and beyond: boundary effects, flexibility, and noise correlations in neural population codes. *PLoS Comput Biol* 10:e1003469. [CrossRef Medline](#)
- Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol* 160:106–154. [CrossRef Medline](#)
- Issa EB, Wang X (2013) Increased neural correlations in primate auditory cortex during slow-wave sleep. *J Neurophysiol* 109:2732–2738. [CrossRef Medline](#)
- Jeanne JM, Sharpee TO, Gentner TQ (2013) Associative learning enhances population coding by inverting interneuronal correlation patterns. *Neuron* 78:352–363. [CrossRef Medline](#)
- Kajikawa Y, Frey S, Ross D, Falchier A, Hackett TA, Schroeder CE (2015) Auditory properties in the parabelt regions of the superior temporal gyrus in the awake macaque monkey: an initial survey. *J Neurosci* 35:4140–4150. [CrossRef Medline](#)
- Kanitscheider I, Coen-Cagli R, Pouget A (2015) Origin of information-limiting noise correlations. *Proc Natl Acad Sci U S A* 112:E6973–E6982. [CrossRef Medline](#)
- Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J Neurosci* 25:3661–3673. [CrossRef Medline](#)
- de la Rocha J, Doiron B, Shea-Brown E, Josić K, Reyes A (2007) Correlation between neural spike trains increases with firing rate. *Nature* 448:802–806. [CrossRef Medline](#)
- Mante V, Sussillo D, Shenoy KV, Newsome WT (2013) Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503:78–84. [CrossRef Medline](#)
- McDermott JH (2009) The cocktail party problem. *Curr Biol* 19:R1024–R1027. [CrossRef Medline](#)
- Merzenich MM, Brugge JF (1973) Representation of the cochlear partition of the superior temporal plane of the macaque monkey. *Brain Res* 50:275–296. [CrossRef Medline](#)
- Mesgarani N, Chang EF (2012) Selective cortical representation of attended speaker in multi-talker speech perception. *Nature* 485:233–236. [CrossRef Medline](#)
- Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63:879–888. [CrossRef Medline](#)
- Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, Pouget A (2014) Information-limiting correlations. *Nat Neurosci* 17:1410–1417. [CrossRef Medline](#)
- Nandy AS, Nassi JJ, Reynolds JH (2017) Laminar organization of attentional modulation in macaque visual area V4. *Neuron* 93:235–246. [CrossRef Medline](#)
- Nienborg H, Cumming B (2010) Correlations between the activity of sensory neurons and behavior: how much do they tell us about a neuron's causality? *Curr Opin Neurobiol* 20:376–381. [CrossRef Medline](#)
- Niwa M, Johnson JS, O'Connor KN, Sutter ML (2012) Activity related to perceptual judgment and action in primary auditory cortex. *J Neurosci* 32:3193–3210. [CrossRef Medline](#)
- O'Connor KN, Johnson JS, Niwa M, Noriega NC, Marshall EA, Sutter ML (2011) Amplitude modulation detection as a function of modulation frequency and stimulus duration: comparisons between macaques and humans. *Hear Res* 277:37–43. [CrossRef Medline](#)
- Oram MW, Földiák P, Perrett DI, Sengpiel F (1998) The “Ideal Homunculus”: decoding neural population signals. *Trends Neurosci* 21:259–265. [CrossRef Medline](#)
- Orhan AE, Ma WJ (2015) Neural population coding of multiple stimuli. *J Neurosci* 35:3825–3841. [CrossRef Medline](#)
- Osmanski MS, Wang X (2015) Behavioral dependence of auditory cortical responses. *Brain Topogr* 28:365–378. [CrossRef Medline](#)
- Pagan M, Urban LS, Wohl MP, Rust NC (2013) Signals in inferotemporal and perirhinal cortex suggest an untangling of visual target information. *Nat Neurosci* 16:1132–1139. [CrossRef Medline](#)
- Petkov CI, Kayser C, Augath M, Logothetis NK (2006) Functional imaging reveals numerous fields in the monkey auditory cortex. *PLoS Biol* 4:e215. [CrossRef Medline](#)
- Pfingst BE, Laycock J, Flammino F, Lonsbury-Martin B, Martin G (1978) Pure tone thresholds for the rhesus monkey. *Hear Res* 1:43–47. [CrossRef Medline](#)
- Ray S, Maunsell JH (2011) Different origins of gamma rhythm and high-gamma activity in macaque visual cortex. *PLoS Biol* 9:e1000610. [CrossRef Medline](#)
- Rensink RA, O'Regan JK, Clark JJ (1997) To see or not to see: The need for attention to perceive changes in scenes. *Psychol Sci* 8:368–373. [CrossRef](#)
- Romo R, Hernández A, Zainos A, Salinas E (2003) Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. *Neuron* 38:649–657. [CrossRef Medline](#)
- Ruff DA, Cohen MR (2014a) Attention can either increase or decrease spike count correlations in visual cortex. *Nat Neurosci* 17:1591–1597. [CrossRef Medline](#)
- Ruff DA, Cohen MR (2014b) Global cognitive factors modulate correlated response variability between V4 neurons. *J Neurosci* 34:16408–16416. [CrossRef Medline](#)
- Saleem KS, Logothetis NK (2007) A combined MRI and histology atlas of the rhesus monkey brain in stereotaxic coordinates. Burlington, MA: Academic.
- Sasaki R, Uka T (2009) Dynamic readout of behaviorally relevant signals from area MT during task switching. *Neuron* 62:147–157. [CrossRef Medline](#)
- Shamma SA, Elhilali M, Micheyl C (2011) Temporal coherence and attention in auditory scene analysis. *Trends Neurosci* 34:114–123. [CrossRef Medline](#)

- Shinn-Cunningham BG (2008) Object-based auditory and visual attention. *Trends Cogn Sci* 12:182–186. [CrossRef Medline](#)
- Snyder AC, Morais MJ, Kohn A, Smith MA (2014) Correlations in V1 are reduced by stimulation outside the receptive field. *J Neurosci* 34:11222–11227. [CrossRef Medline](#)
- Steinschneider M, Fishman YI, Arezzo JC (2008) Spectrotemporal analysis of evoked and induced electroencephalographic responses in primary auditory cortex (A1) of the awake monkey. *Cereb Cortex* 18:610–625. [CrossRef Medline](#)
- Strickland EA, Viemeister NF (1997) The effects of frequency region and bandwidth on the temporal modulation transfer function. *J Acoust Soc Am* 102:1799–1810. [Medline](#)
- Sutter ML, Shamma SA (2011) The relationship of auditory cortical activity to perception and behavior. *The Auditory Cortex* 617–641.
- Treisman A (1998) Feature binding, attention and object perception. *Philos Trans R Soc Lond B Biol Sci* 353:1295–1306. [CrossRef Medline](#)
- Winkowski DE, Bandyopadhyay S, Shamma SA, Kanold PO (2013) Frontal cortex activation causes rapid plasticity of auditory cortical processing. *J Neurosci* 33:18134–18148. [CrossRef Medline](#)
- Woods KJ, McDermott JH (2015) Attentive tracking of sound sources. *Curr Biol* 25:2238–2246. [CrossRef Medline](#)
- 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. [CrossRef Medline](#)