JNeurosci
THE JOURNAL OF NEUROSCIENCE

Research Articles: Systems/Circuits

Stimulus-specific prediction error neurons in mouse auditory cortex

https://doi.org/10.1523/JNEUROSCI.0512-23.2023

Cite as: J. Neurosci 2023; 10.1523/JNEUROSCI.0512-23.2023

Received: 21 March 2023
Revised: 7 August 2023
Accepted: 4 September 2023

Stimulus-specific prediction error neurons in mouse auditory cortex

Nicholas J. Audette[1], David M. Schneider[1]

[1.] Center for Neural Science, New York University, 4 Washington Place, New York, NY 10003, USA

Corresponding author: david.schneider@nyu.edu

Abstract

Comparing expectation with experience is an important neural computation performed throughout the brain and is a hallmark of predictive processing. Experiments that alter the sensory outcome of an animal's behavior reveal enhanced neural responses to unexpected self-generated stimuli, indicating that populations of neurons in sensory cortex may reflect prediction errors – mismatches between expectation and experience. However, enhanced neural responses to self-generated stimuli could also arise through non-predictive mechanisms, such as the movement-based facilitation of a neuron's inherent sound responses. If sensory prediction error neurons exist in sensory cortex, it is unknown whether they manifest as general error responses, or respond with specificity to errors in distinct stimulus dimensions. To answer these questions, we trained mice of either sex to expect the outcome of a simple sound-generating behavior and recorded auditory cortex activity as mice heard either the expected sound or sounds that deviated from expectation in one of multiple distinct dimensions. Our data reveal that the auditory cortex learns to suppress responses to self-generated sounds along multiple acoustic dimensions simultaneously. We identify a distinct population of auditory cortex neurons that are not responsive to passive sounds or to the expected sound but that encode prediction errors. These prediction error neurons are abundant only in animals with a learned motor-sensory expectation, and encode one or two specific violations rather than a generic error signal. Together, these findings reveal that cortical predictions about self-generated sounds have specificity in multiple simultaneous dimensions and that cortical prediction error neurons encode specific violations from expectation.

Significance Statement: Audette et. al record neural activity in the auditory cortex while mice perform a sound-generating forelimb movement and measure neural responses to sounds that violate an animal's expectation in different ways. They find that predictions about self-generated sounds are highly specific across multiple stimulus dimensions and that a population of typically non-sound-responsive neurons respond to sounds that violate an animal's expectation in a specific way. These results identify specific

prediction error signals in the mouse auditory cortex and suggest that errors may be calculated early in sensory processing.

Introduction

Sensory responses in the cerebral cortex are influenced by an animal's behavior (Niell and Stryker, 2010; Polack et al., 2013; Zhou et al., 2014; McGinley et al., 2015; Kuchibhotla et al., 2017; Ayaz et al., 2019; Clancy et al., 2019; Musall et al., 2019; Steinmetz et al., 2019; Stringer et al., 2019) and can reflect an expectation for the sensory consequences of movement (Eliades and Wang, 2008; Flinker et al., 2010; Keller et al., 2012; Nelson et al., 2013; Zmarz and Keller, 2016; Rummell et al., 2016; Leinweber et al., 2017; Schneider et al., 2018; Knolle et al., 2019; Jordan and Keller, 2020; Reznik et al., 2021; Audette et al., 2022). This dynamism is consistent with the theory of predictive processing, which posits that cortical activity prioritizes representing deviations from expectation over directly representing features of the external world (Bastos et al., 2012; Keller and Mrsic-Flogel, 2018). Some expectations are purely sensory in nature, such as the repetition of a common stimulus. Following repeated presentation of a fixed stimulus, sensory cortical responses to the common stimulus become suppressed while responses to oddball stimuli are typically, though not always, enhanced (Ulanovsky et al., 2003, 2004; Khatri et al., 2009; Farley et al., 2010; Taaseh et al., 2011; Natan et al., 2015; Solomon et al., 2021). Experiments that alter the sensory outcomes of behavior have revealed that the sensory cortex is also modulated by learned expectations for the sensory outcome of specific movements (Eliades and Wang, 2008; Keller and Hahnloser, 2009; Keller et al., 2012; Mandelblat-Cerf et al., 2014; Rummell et al., 2016; Zmarz and Keller, 2016; Audette et al., 2022). As with sensory-only predictions, responses to the expected outcomes of movement are typically suppressed while responses to unexpected outcomes are unaffected or enhanced. The production of error signals during motor sensory behaviors may facilitate the encoding of sensory information, inform moment-by-moment behavior by influencing motor plans, and provide

learning signals that specify when motor-sensory predictions require updates (Schneider and Mooney, 2018; Schneider, 2020)

Sensory-motor error signals often manifest as modulations of a neuron's passive tuning curve. That is, expectation violation responses are heightened responses to stimuli that a neuron responds to even in a passive condition. In contrast, error signals elsewhere in the brain can be highly specific. For example, reward prediction error neurons in the midbrain explicitly encode violations from expectation but do not respond to predictable cues (Schultz et al., 1997; Glimcher, 2011; Eshel et al., 2016). Critical gaps remain in our understanding of cortical error signaling, including whether sensory cortex possesses neurons that explicitly encode sensory prediction errors akin to reward prediction error signals in midbrain dopamine neurons. It also remains unknown whether cortical error signals reflect a general error or whether they are specific to the nature of the expectation violation. Finally, it remains unresolved whether error signals could arise in a way that is unrelated to expectation, such as through the mixing of movement and sensation signals (Muzzu and Saleem, 2021).

Here, we employed a simple sound-generating forelimb behavior to generate a motor-sensory expectation in mice. We then recorded neural responses in the auditory cortex as mice experienced expected sounds and sounds that violated their expectation across multiple dimensions during behavior. To distinguish between the effects of motor-sensory predictions and other forms of modulation, we conducted an identical experiment in mice trained to perform the same forelimb behavior, but without prior motor-sound coupling. Our findings reveal that learned motor-sensory predictions have specificity across multiple feature dimensions simultaneously and that cortical prediction error neurons selectively encode specific violations from expectation with short latencies in the auditory cortex.

Materials and Methods

Animals

All experimental protocols were approved by New York University's Animal Use and Welfare Committee. Male and female wild-type (C57BL/6) mice were purchased from Jackson Laboratories and were subsequently housed and bred in an onsite vivarium. We used 2–4 month-old mice for our experiments that were kept on a reverse day-night cycle (12h day, 12 h night).

Surgeries

For all surgical procedures, mice of either sex were anaesthetized under isolfurane (1-2% in O_2) and placed in a stereotaxic holder (Kopf), skin was removed over the top of the head, and a Y-shaped titanium headpost (H.E. Parmer) was attached to the skull using a transparent adhesive (Metabond). Mice were treated with an analgesic (Meloxicam SR) and allowed to recover for 5 days prior to training. Following training and 24 to 48 hours prior to electrophysiology, a small craniotomy was made to expose the auditory cortex (~2mm diameter, -2.5mm posterior, 4.2 mm left from bregma). Another small craniotomy was made above the right sensory cortex and a silver-chloride reference electrode was positioned atop the surface of the brain for use as a ground electrode and covered (Metabond). Exposed craniotomies were covered with a silicone elastomer (Kwik-Sil) and the mouse was allowed to recover in its home cage, and an additional training session was performed prior to electrophysiology.

Behavioral Training and Data Collection

We adapted a custom head-restrained lever-based behavioral training paradigm where mice push a lever and hear closed-loop sounds (Audette et al., 2022). A custom-designed lever (7cm long, 3D printed using Formlabs Form2) was mounted to the post of a rotary encoder (US Digital) 5cm from the lever handle. A magnet (CMS magnetics) was mounted to the bottom of the lever, which was positioned 4cm above a larger static magnet which established the lever resting position and provided light and adjustable movement resistance. The lever handle (top) was positioned adjacent to a tube (Custom, 3D printed using Formlabs Form2) to hold mice directly below two plate clamps (Altechna) to secure the

mouse headpost. Lever and mouse apparatus was constructed with Thor-labs components. A water tube, controlled by a solenoid valve (The Lee Company), was positioned in front of the mouse. Digital signals for lever movement were collected by a data acquisition card (National Instruments) connected to a computer and logged by custom Matlab software (Mathworks, PsychToolBox) and sampled at 2Khz. Digital processing of lever movements received sufficient processing in real time to track important movement thresholds, which were used to trigger sound events based on user-defined closed-loop rules. Sound output was delivered from the computer to a sound card (RME Fireface UCX), the output of which was routed to an ultrasonic speaker (Tucker Davis Technologies) located lateral to the mouse, ~10cm from the mouse's right or left ear. We recorded sounds during test experiments using an ultrasonic microphone (Avisoft, Model # CM16/CMPA-P48) positioned 5 cm from the lever to confirm that the lever produced negligible noise (<1 dB SPL) and that experimenter-controlled sounds were delivered at a consistent volume of 50, 65, and 80 db depending on stimulus type. All training was performed in a sound-attenuating booth (Gretch-Ken) to minimize background sound and monitored in real-time via IR video.

During lever training, mice were water restricted and maintained greater than 80% of pre-restriction body weight and received all of their water (1-2ml) while performing the lever behavior. In practice, body weight was often above 90% since diminished body weight was not necessary to induce lever pressing once mice learned the task. During training, mice were head-fixed to the behavioral apparatus and presented with the lever and lick-port after ~10 minutes of quiet acclimation. Mice were then allowed to make outwards lever movements at will. For a movement to be considered valid, we required the lever to remain in the home position (~+/- 3mm from rest) for >200 ms prior to initiation. Valid movements that reached a reward threshold (~15mm from home position) elicited a small water reward (5-10uL) when the lever returned to home position. Auditory feedback in the form of a pure tone (50ms duration, 65dB, 12khz) was delivered on all trials when the lever crossed a set threshold 1/3

of the way between the home position and reward threshold for the first time in a trial. To ensure strong coupling between movement and sound, auditory feedback was provided on all trials, regardless of whether mice obeyed the home-position requirement and would subsequently receive a reward. Initially, 100% of successful trials produced a reward, but over the course of training that number was dropped to 25% to produce more lever movements per session. The reward rate was stable for at least 5 sessions prior to recording. Overall, mice received between 18 and 22 sessions of training over 10-12 days prior to electrophysiology, with either one or two sessions per day.

Electrophysiological Recording and Aggregate Neural Responses

Following training, we used stereotaxic coordinates and cranial landmarks to open a craniotomy above the auditory cortex. Following the experiment, onset latency of neural responses to passive sounds was used to confirm auditory cortex localization, though stereotaxic and latency data are insufficient to distinguish between subfields of the auditory cortex (Romero et al., 2020; Narayanan et al., 2023). After one subsequent training session, mice were positioned in the behavioral apparatus and a 128-channel electrode (128AxN, Masmanidis Lab) was lowered into the auditory cortex orthogonal to the pial surface (Yang et al., 2020). The electrode was connected to a digitizing head stage (Intan) and electrode signals were acquired at 30khz, monitored in real time, and stored for offline analysis (OpenEphys). The probe was allowed to settle for at least 20 minutes, at which point the lever and lick-port were introduced and mice were allowed to make lever movements at will as in any other training session. After performing at least 30 standard lever movements, we unexpectedly began a probe session in which mice heard several different sounds. 90% of sounds were as expected (‘Exp’, 12khz, 65db) while 1.4% each were a substituted frequency (‘Freq’, 5.6khz, 1.1 octave lower, 65db), both the unexpected and an unexpected frequency (‘Comp’, 5.6khz and 12khz, 65db), a higher intensity (‘Loud’, 12khz, 80db), a lower intensity (‘Quiet’, 12khz 50db), played from a different origin (‘Orig’, 12khz, 65db, played from a speaker on the left side of the mouse’s head), played during the return phase of the lever movement (‘Pos’, 12khz,

65db, half way between reward threshold and the return to the home position on trials reaching reward threshold), or omitted. The requirements for reward delivery were not influenced by the identify or timing of auditory feedback. Following probe sessions, the lever was removed and tone frequencies ranging from 3 to 32kHz (0.5 octave spacing) as well as all tones presented during the active phase of the task were presented with random inter-tone intervals drawn from a flat distribution with range 1 to 2 seconds.

After recording, electrical signals were processed and the action-potentials of individual neurons were sorted using Kilosort2.5 (Pachitariu et al., 2016), and manually reviewed in Phy2 based on reported contamination, waveform principal component analysis, and inter-spike interval histograms. Because the identification of prediction error neurons could be dramatically skewed by the loss of neural signals over the course of an experiment, we excluded any neuron that had a statistically significant difference ($p<0.05$) in baseline firing rate or the response rate to passively heard tones from the pre- and post-behavioral passive tone sessions. We analyzed neurons with non-fast-spiking waveforms, separated by plotting peak to valley ratio against action potential width. Tone-evoked average firing rate PSTHs were measured in 2 ms bins and aligned to sound onset for each neuron for each tone type (Fig. 1E). PSTHs and individual neuron modulation for a given tone type include all neurons that were responsive ($p<0.01$) to a given tone in either the active or passive condition measured as an increase in firing rate from baseline (60ms prior to stimulus onset) during the sound response window (0-60ms post stimulus onset) across trials using a paired rank sum test. To measure the movement-based modulation of each neuron's responses to the lever-associated or probe tones, we compared the neural sound response in our analysis window to the same sound in the active and passive condition using a radial modulation index (Audette et al., 2022). Radial modulation was calculated as the theta value resulting from a cartesian to polar transformation of the response strength in the active condition compared to the response strength in the passive condition. Theta values were converted to a scale of +/- 2 and rotated

such that a value of 0 corresponded to equal responses across the two conditions. The fraction of neuron overlap reported in the text measures the fraction of neurons responsive to the passively heard expected sound that also respond to each probe sound.

In a subset of animals, we performed electrophysiological recording of mice that had been trained on an identical version of the lever task but without sound feedback. On experiment day mice first performed silent lever pushes for 20-50 trials, then we delivered a range of sound frequencies (4-24khz, half octave intervals, 50ms duration, 65dB,) at the sound threshold during lever pushes, followed by presentation of the same sounds passively with the lever removed, as above.

Prediction Error Neuron Analysis

We defined prediction error neurons as having a significant response in the sound response window ($p<0.01$, 0-60ms post stimulus onset compared to 60ms prior to stimulus onset) for a given stimulus type, but not to the same stimulus heard passively ($p>0.1$), to the expected sound heard actively ($p<0.01$) or at the same position during movement on omission trials ($p>0.1$). Prediction error neurons were identified independently for each stimulus type. Prediction error neurons were identified in silent trained animals using the same functional definition comparing activity in the movement condition, passive sound condition, and active condition.

The fraction of prediction error neurons (Fig. 4A-C) was defined as the number of prediction error neurons for a stimulus type divided by the total number of sound-responsive neurons in an experiment (Active or passive) and are presented with data points representing one stimulus in one animal. For analyses involving individual animals (Fig. 2C, Fig. 4A-C), data was analyzed only for animals that had more than 40 sound responsive neurons in the population ($N = 4$). For regression comparisons, the neurometric difference between a probe stimulus and the expected stimulus was calculated by comparing average response responses across the two tones in the passive condition. The difference

between responses to the two tones for each neuron was summated across all neurons in an animal and used to represent the dissimilarity of neural response patterns between the probe sound an expected sound. These values were mean normalized within each animal to allow for comparison across animals. A similar process was used for passive response magnitude, but with average firing rates summated across all neurons in an animal instead of making a comparison to the expected sound. Onset latencies were defined for each neuron as the average of first post-stimulus spike times on each trial. Trials that did not produce an action potential in the sound response window were removed from the average. Histograms of onset latencies were created using 2ms bins.

Decoding Analysis

Decoding data were organized in a trials-by-neuron matrix within each animal, with each cell representing the response of an individual neuron on an individual trial. A consistent number of trials (20, randomly selected) was used for each stimulus type. Each trial, in sequence, was removed from the data set, and the remaining trials along with the ground truth identity of the experienced stimulus was used to train a multiclass error-correcting output codes model using support vector machine binary learners (Cristianini and Shawe-Taylor, 2000; Narsky and Porter, 2013). The trained model was then used to classify the withheld trial, which was then compared to the ground truth identify of the stimulus. This process was repeated for all trials in an animal, with the results visualized as a confusion matrix comparing the classification result to the ground truth identity of each trial. Each pixel represents the number of trials classified as a given stimulus type divided by the number of ground truth trials for a given stimulus type. The resultant confusion matrices were then averaged across animals.

Statistical Analysis

Throughout, animal values are denoted by a capital N while cell values are denoted by a lowercase n. Unless otherwise reported, all averages and error bars denote mean ± standard deviation. P values are

reported in text or on the relevant figure panels for all statistical comparisons. Statistical comparison of aggregate neural activity use a one-way ANOVA followed by two-sided, non-paired, non-parametric rank sum test and Bonferroni correction for multiple comparisons. The comparison of the number of 'active only' neurons for probe stimuli vs. the expected stimulus was performed by bootstrap resampling, with which we compared the observed counts for the two stimuli to 10,000 randomly generated distribution of counts created assuming equal probability. Statistical comparison of onset latency across groups was performed using a Kolmogorov-Smirnov test. The relationship between the number of prediction error neurons and neural response properties was measured using linear regression and correlation coefficient analysis with p and R values reported.

Results

Motor-sensory predictions are specific across multiple acoustic dimensions

The auditory cortex predicts the frequency of a self-generated sound and its expected position within an ongoing movement (Rummell et al., 2016; Schneider et al., 2018; Audette et al., 2022). But sounds have many features, including spatial location, intensity, and spectrum. We therefore aimed to determine whether movement-based predictions in the auditory cortex show specificity along multiple acoustic dimensions simultaneously. We trained head-fixed mice to produce a simple sound-generating behavior during which we could precisely control the acoustic outcome of each movement (Audette et al., 2022). Mice pushed a lever past a fixed threshold to trigger a water reward (on 25% of trials) when the lever was returned to the home position (Fig. 1A). During training, a pure tone (8 kHz) was presented at a consistent position early in each movement, and mice were free to initiate trials ad libitum. Mice rapidly learned to perform the task and averaged more than 2000 sound-generating trials per session. Lever movements in well-trained mice lasted about 275 ms on average (Fig. 1B,C) and mice experienced lever-evoked sounds roughly every second (Fig. 1D).

Following 10-12 days of training with the lever producing a predictable self-generated sound, we made large channel-count electrophysiological recordings from the auditory cortex while mice executed the learned lever behavior and heard either the expected sound (90% of trials) or a sound that unexpectedly varied in one of several different acoustic dimensions (probe trials, 1.4% each) (Fig. 1E). On these probe trials we did one of the following: substituted a sound shifted 1.1 octaves from the expected sound (Frequency), played an unexpected frequency simultaneously with the expected sound (Composite), changed the intensity of the expected sound by +/-15 dB (Quiet or Loud), changed the spatial origin of the sound (Origin), played the expected sound at the wrong lever position (Position), or omitted the sound altogether. Each of these sounds was also played in a passive listening context during which the lever was removed from the animal's reach. In total, we recorded from 1016 regular spiking neurons across five animals.

In the passive listening condition, we observed strong neural responses to each sound, including the expected sound (Fig. 1E). In the self-generated condition, neural responses to the expected sound were strongly suppressed (~50%) compared to the same sound heard passively (Audette et al., 2022). This strong suppression of neural responses to an expected self-generated sound provides a benchmark for comparing neural responses to unexpected self-generated sounds. If neural responses to an unexpected sound are less suppressed, unsuppressed, or enhanced, we can conclude that the auditory cortex recognizes that sound as a violation of its expectation.

We found that the auditory cortex did not display strong suppression of neural responses to any unexpected sound that we tested. Population-averaged neural responses to the unexpected probe sounds were not suppressed at all (Quiet, Loud, Origin), were mildly suppressed (Position), or were enhanced relative to the passive listening condition (Frequency). As a striking example, we found that neural responses to an unexpectedly quiet self-generated tone were significantly stronger than were responses to the self-generated tone heard at the expected volume ($p = 8*10^{-8}$). This is in direct contrast to the

passive listening condition, during which the expected intensity evoked stronger responses than the quieter intensity, as would be expected from typical mouse auditory cortex neurons (Joachimsthaler et al., 2014).

The acoustically selective suppression of neural responses to self-generated sounds was also recapitulated when we compared the sound responses of individual neurons across the passive and self-generated condition by computing a modulation index (see Methods)(Audette et al., 2022). The majority of neurons had weaker responses to the expected sound when it was self-generated compared to when it was heard passively (negative modulation values) (Fig. 2A). In contrast, neurons displayed less suppression to all unexpected sounds ($p<0.01$ for all), responding equally strongly on average to probe sounds when they were self-generated and heard passively, with some neurons enhanced, some suppressed, and many cells responding equally across the two conditions. The notable exception was the frequency probe, which generated enhanced neural responses relative to the passive condition, consistent with large population-level neural responses (see Fig. 1E).

In order to preserve an animal's expectation for the movement-associated sound for the duration of the experiment, animals heard the expected sound on 90% of movements with probe sounds occurring on just 10% of movements. Due to experimental time constraints, during passive playback all sounds were heard with equal probabilities and with an inter-sound-interval similar to that heard during the lever behavior. This unbalanced ratio of sounds between the two conditions could itself contribute to the observed pattern of neural responses to expected and unexpected sounds, through mechanisms such as stimulus-specific adaptation (Ulanovsky et al., 2004; Taaseh et al., 2011; Natan et al., 2015). To account for this possibility, we measured neural responses to lever-generated sounds using an identical experimental setup in mice that learned to make silent lever movements. We do observe some effects in silent-trained mice that could be attributed to stimulus-specific adaptation, specifically weak suppression of the expected sound compared to passive listening and compared to probe sounds that contained an

oddball frequency ('Frequency', 'Composite', Fig. 2B,C). However, the magnitude of this frequency-specific suppression was much smaller in than in sound-trained animals, and responses to expected sounds in silent-trained animals were statistically indistinguishable from other probe sounds that shared the same frequency (Fig. 2B,C). These findings demonstrate that while stimulus-specific adaptation contributes to the suppression of expected self-generated sounds, the magnitude and specificity of suppression measured in trained animals depends upon a learned motor-sensory prediction.

In addition to these population-level effects of motor-sensory expectation, we also observed highly specific suppression of the expected self-generated sound at the level of individual neurons. Measuring responses of each individual neuron to different self-generated sound types revealed that prediction-based suppression could diminish the magnitude of a neuron's response to the expected sound but have a small or minimum impact the neuron's ability to respond to other unexpected sounds (Fig. 2D). Indeed, neural responses to unexpected sounds that shared the same frequency as the expected sound largely escaped suppression despite substantial overlap in the neural population responsive to the sounds in the passive condition (60% ± 15%).

The different patterns of population-level activity evoked by passive and self-generated sounds were sufficient to decode the sound identify and behavioral context in which it was heard on individual trials from small groups of auditory cortex neurons (Fig. 2E). Taken together, these data are consistent with the auditory cortex simultaneously predicting the expected frequency, position, intensity, and spatial location of a self-generated sound and applying a highly selective mechanism of suppression.

Prediction error neurons respond to specific violations of a motor-sensory expectation

The single-neuron analyses outlined above reveal many neurons that respond more strongly to an unexpected self-generated sound than to the same sound heard passively (Keller et al., 2012; Jordan and Keller, 2020; Audette et al., 2022). While some of these neurons are likely responsive in both behavioral

conditions but with relatively larger responses in the active condition, the number of strongly enhanced neurons (i.e. neurons with MI close to 1 in Fig. 2A) for each unexpected sound raises the possibility that these sounds recruit a new group of cells that do not respond passively. We therefore quantified neurons that were activated by each sound in the passive condition, the active condition, or both. A relatively consistent number of neurons were responsive to each sound in the passive condition ('Passive only' and 'Shared', Fig. 3A). When mice heard the expected self-generated sound, only a small subset of passive-responsive neurons responded ('Shared'). In contrast, when mice heard any unexpected sound, a substantially larger number of neurons responded including many neurons that were unresponsive to these same sounds heard passively ('Active only', $p < 0.01$ for all).

Since these 'active only' neurons were abundantly recruited following unexpected, but not expected, self-generated sounds, we hypothesized that they may explicitly encode prediction errors. Enhanced neural responses following unexpected stimuli have been observed at the population and single-neuron level in prior experiments (Eliades and Wang, 2008; Keller et al., 2012; Rummell et al., 2016; Schneider et al., 2018; Audette et al., 2022), but it has not been conclusively established whether such responses depend upon a learned motor-sensory prediction. To determine whether prediction error neurons exist in the auditory cortex, we identify a subset of 'active only' neurons as putative prediction error neurons and measure their abundance following each sound in trained and untrained animals.

First, we established a stringent definition for putative prediction error (PE) neurons in the auditory cortex. We required that PE neurons respond to an unexpected self-generated sound ($p<0.01$) but not to the same sound heard passively ($p>0.1$), not to the expected self-generated sound ($p>0.1$), and not in the same window during silent movements ($p>0.1$, Fig. 3B). This ensures that our putative prediction error neurons respond to the presence of a sound that is self-generated and unexpected, and cannot arise due directly to movement, to the combination of movement and sound in a way that is not specific to an expectation violation, or to the enhancement of a neuron's passive response to the sound. Using these

criteria, we identified 85 PE neurons, corresponding to 8.4% of all recorded neurons and 29.8% of sound-responsive neurons (Fig. 3C). Neurons that fulfil these criteria could be highly selective for a single self-generated sound but could also respond to other sound types in either the active or passive condition. To determine the specificity of auditory cortex PE neurons, we visualized each neuron by displaying its responsiveness across active and passive stimuli, and the stimuli for which it signals a prediction error (Fig. 3D). Auditory cortex PE neurons fell into two general categories: Neurons that responded only to one or two unexpected self-generated sounds and no passive stimuli (Fig. 3E, Neuron 1), or neurons that responded to a different set of stimuli in the active and passive condition (Fig. 3E, Neuron 2).

Nearly half of auditory cortex PE neurons (45%) were unresponsive to any of the task sounds in the passive condition, with 70% responding to one or fewer, suggesting that many of these neurons would not classically be considered sound-responsive neurons (Fig. 3D). To further characterize the sound responsiveness of our PE neurons, we also presented pure tones at half octave intervals during passive listening following the playback of task sounds. Even across 14 unique stimuli, 28% of PE neurons did not respond, even weakly, to a single tone ($p < 0.1$), and 52% of neurons responded to two or fewer (Fig. 3F). Similarly, PE neurons had much weaker sound responses to passively heard sounds not present in the behavioral task than non-PE neurons (Mean 2.7 sp/s vs 8.1 sp/s, $p = 4*10^{-17}$).

In addition to responding weakly to passive sounds, PE neurons generally did not respond broadly to self-generated sounds. By definition, PE neurons cannot respond to the expected, self-generated sound, and most PE neurons signaled a prediction error for only one unexpected outcome (74%) and 97% of PE neurons signaled two or fewer outcomes, consistent with PE neurons signaling specific rather than generic errors (Fig. 3H). For the subset of PE neurons that responded to multiple unexpected self-generated sounds, we evaluated the specific sets of violation stimuli by which they were activated (Fig. 3I). The vast majority of these non-specific PE neurons were responsive to the frequency probe and composite probe

357 stimuli. This pairing makes sense since both the composite and frequency probe stimuli contained the
358 same unexpected frequency.

359 Stimulus-specific PE neurons could arise through computations of prediction errors at a higher cortical
360 level that are transmitted back to auditory cortex (Keller and Mrsic-Flogel, 2018). The computation of
361 prediction errors subcortically or in the auditory cortex should result in shorter latency error signals than
362 a mechanism that requires the feedback of a generic error signal. We therefore quantified the onset
363 latency of prediction error neurons measured as the time to first spike following stimulus onset (Fig. 4A).
364 Error responses to the frequency probe in PE neurons were as rapid as neural responses to passively heard
365 sounds and responses to self-generated sounds by non-PE neurons (Fig. 4B). Given the specificity and
366 early onset of prediction error signals following unexpected sounds, it is unlikely that these neurons are
367 driven by feedback of a general error signal calculated downstream from the auditory cortex. Together,
368 our criteria identify an abundant population of auditory cortex neurons that are selectively responsive to
369 a small number of sounds when they are heard as the violation of a motor-sensory prediction.

370

371 **Prediction error neurons require a learned motor-sensory prediction**

372 Our strict criteria for PE neurons preclude the possibility that their responses to unexpected self-
373 generated sounds arise through a simple combination of suprathreshold sound and movement responses.
374 However, prediction error-like signals could potentially arise through subthreshold mechanisms that are
375 unrelated to expectation but instead reflect a simple convergence of subthreshold motor and auditory
376 inputs (Muzzu and Saleem, 2021). Prediction error signals could also emerge in response to the violation
377 of local probabilities through mechanisms like stimulus-specific adaptation (Näätänen et al., 2007; Taaseh
378 et al., 2011; Fishman and Steinschneider, 2012; Natan et al., 2015). To test whether the PE neurons
379 identified by our criteria truly reflect the violation of a learned motor-sensory prediction, we measured

the abundance of neurons that meet these criteria in an identical experiment performed in mice trained to make lever pushes in silence (See Fig. 2B). Unlike in mice that expected the lever to produce a sound, in silent-trained mice we observed a very small fraction of neurons responsive to a given sound that fulfilled our prediction error criteria (Fig. 5A). The comparative abundance of neurons responsive only to unexpected self-generated sounds in sound-trained mice demonstrates that the putative prediction error neurons identified by our criteria reflect the violation of a learned motor-sensory expectation rather than the mixing of subthreshold movement and sound signals or a response to local sound ratios.

A hallmark characteristic of prediction error neurons throughout the brain is the scaling of error responses with the magnitude of the perceived error (Tobler et al., 2005; Eshel et al., 2016). Given that different probe stimuli evoked different numbers of PE neurons (See Fig. 3C), we asked whether the number of PE neurons recruited by a stimulus was related to how different the stimulus was from "expected." We measured stimulus similarity using a population-level neurometric approach, computing the absolute difference between a neuron's response to the expected sound and a probe sound, summed across all non-PE neurons in an animal. This measure of response similarity relative to the expected sound varied across stimuli, providing a proxy for how strongly a stimulus violated expectation. We observed that the number of PE neurons responsive to an unexpected sound scaled with the magnitude of the estimated expectation violation in animals with a motor-sensory prediction, but not in animals trained in silence (Fig. 5B). The average response strength of the population of PE neurons activated by an unexpected stimulus was not significantly correlated with the estimated expectation violation ($r = 0.35$, $p = 0.11$). To ensure that this finding was not simply because some sounds activated the auditory cortex more strongly in general, we performed a similar analysis, comparing the number of PE neurons to the magnitude of a sound's response in the passive condition (Fig. 5C). We found no correlation between the number of PE neurons evoked by a sound and passive response strength regardless of animal training, supporting the conclusion that the number of PE neurons observed reflects the 'unexpectedness' of a movement's

sensory outcome. Together, these findings identify a substantial population of neurons in the auditory cortex whose responses signal the violation of a learned motor-sensory expectation.

Discussion

Our experiments show that movement-based predictions emerge with motor-sensory experience and result in sound-suppression that is specific across multiple feature dimensions. We also identify a population of auditory cortical neurons that signal specific violations of a learned, motor-sensory prediction.

Auditory cortex activity displayed prediction-based suppression that was specific for the frequency, intensity, timing, and spatial origin of an expected sound when measured at the population activity level and when measuring individual neuron modulation. These observations suggest that the auditory cortex learns a highly specific expectation for multiple simultaneous features of a sensation reliably caused by a movement, even when an animal is not explicitly tasked with learning these features. A simple circuit model in which somatic inhibition decreases the spiking of neurons tuned to the expected stimulus is likely inadequate to account for this multi-dimensional specificity, as it would lead to comparable inhibition in a given neuron in response to both expected and unexpected sounds (Wilson et al., 2012; Nelson et al., 2013; Schneider et al., 2014; Zhou et al., 2014; Singla et al., 2017). Instead, we observed many individual neurons that experienced strong suppression of responses to the expected sound during movement while experiencing weaker suppression or even enhancement in response to other self-generated sounds (Fig 2A). These data suggest a more subtle and targeted form of inhibition that can filter neural responses to an expected sensory outcome across multiple features simultaneously within a single neuron.

We identified an abundant population of PE neurons in the auditory cortex that are responsive only when a movement has an unexpected acoustic outcome. Prediction errors have been most commonly described

in midbrain dopamine neurons, which augment their firing rate in response to unexpected rewards (i.e. reward prediction errors)(Schultz et al., 1997; Glimcher, 2011). Reward prediction error signals in midbrain dopamine neurons are notable in that they only encode errors and not predictable outcomes and that their responses scale with the magnitude of an expectation (Schultz et al., 1997; Eshel et al., 2016). Here, we identify auditory cortical neurons that share these hallmarks with reward prediction error neurons, but instead of responding to unexpected rewards, respond to the unexpected acoustic consequences of an action. The PE neurons we identify in the auditory cortex also encode information about how a mouse's expectation was violated (e.g. which acoustic feature).

Our criteria for defining PE neurons exclude neurons that respond to a given sound heard passively, or that respond in the absence of sound on omission trials, eliminating the possibility that our PE neurons arise from a simple combination of sensory or motor tuning. Instead, we demonstrate that neurons with a prediction error phenotype are abundant only in animals that have a learned, motor-sensory expectation, and that the number of prediction error neurons recruited by an unexpected stimulus reflects how different the stimulus was from expectation. Individual prediction error neurons typically respond with short latency and to just one or two probe stimuli, indicating that these neurons do not reflect the feedback of a generic error signal calculated downstream from the auditory cortex. Although we cannot rule out that prediction errors are computed earlier in the ascending auditory pathway (Parras et al., 2017), previous work has shown that expectation violation signals are strongest in layer 2/3 and layer 5 but are largely absent in layer 4, which is a recipient of primary thalamic input. Precise characterization of prediction-related signals in subcortical auditory areas will be needed to confirm the hypothesis, that prediction error signals arise *de novo* within the auditory cortex. Similarly, experiments that measure prediction-related signals across subfields of the auditory cortex will be important for understanding precisely where and how predictive computations are implemented (Parras et al., 2021; Morandell et al., 2023).

As a population, PE neurons had much weaker responses to passive sounds than non-PE neurons and a large fraction of PE neurons were entirely unresponsive to passively heard sounds. The presence of such error-selective neurons that arise through learning clearly identifies neurons that functionally signal prediction errors, but it is less clear whether these neurons are a categorically distinct group that only signal expectation violations, or instead that they belong to a continuum of response phenotypes. Our analysis also identified PE neurons that encode prediction errors for one stimulus while encoding the passive playback of other sounds. Further, our battery of passive sounds was not exhaustive, and it is possible that there are other passive sounds that could reliably drive some of our identified PE cells. Such mixed functionality at the single-neuron level may serve important computational roles, especially when an animal can produce many actions, hear many different sounds, and must keep track of multiple different predictions, as is likely in more real-world contexts. Indeed, midbrain dopamine neurons are also implicated in computations beyond reward prediction errors, including movement vigor and temporal judgements, suggesting that prediction error neurons throughout the brain may play different functions depending on an animal's behavioral needs (Panigrahi et al., 2015; Soares et al., 2016). Auditory cortical prediction errors could be used to update an internal model when the sensory consequences of an action change (e.g. when transitioning from walking on leaves to walking on gravel) or could be routed to motor centers of the brain where they could be used to update subsequent motor plans (e.g. when learning how to play a musical instrument). Understanding whether motor-sensory prediction error signals map onto separable neural populations and how they are used across the brain to update internal models and behavior are important directions for future experiments.

Our experiments focused specifically on expectation violations when a movement produces an unexpected sound, consistent with mismatch-negativity signals that have been observed in humans and other animals during vocalizations and other sound-generating behaviors (Näätänen et al., 2007; Ylinen et al., 2016). The auditory cortex is involved in other forms of predictive processing as well, including

adaptive responses to repeated sounds, known as stimulus-specific adaptation (SSA)(Ulanovsky et al., 2003, 2004; Farley et al., 2010; Taaseh et al., 2011; Natan et al., 2015). In SSA, neural responses to a commonly occurring sound become weaker and neurons produce larger responses to uncommon sounds. During behavior, mice in our experiments heard the expected sound on 90% of trials to preserve the mouse's motor-sensory prediction throughout the experiment, which raises the question of whether stimulus-specific-adaptation or other forms of auditory-only prediction contribute to our results. By performing identically structured experiments in mice trained in the absence of a motor-sensory prediction, we were able to compare the impact of motor-sensory prediction to the combined effect of other forms of modulation including SSA, general task engagement, and general modulation related to movement. While our task is not designed to delineate the relative contribution of each form of modulation, we observed net suppression of lever-generated sounds in mice lacking a motor-sensory prediction. However, this suppression was smaller and less selective than in sound-trained mice, aligning with a previous study that observed prediction-based suppression in a paradigm that excludes any confounds caused by SSA (Audette et al., 2022). Importantly, in animals that experienced identical auditory consequences of movement but without a prior motor-sensory prediction, we did not observe enhancement of responses to uncommon sounds relative to passive listening and our prediction error criteria were met by very few neurons. At a mechanistic level, SSA and motor-sensory predictions likely involve at least partially different neural circuits. Models of SSA involve computations that are local to the auditory cortex, whereas motor-sensory predictions likely require the integration of long-range signals from motor regions with local auditory cortical circuitry (Farley et al., 2010; Natan et al., 2015; Leinweber et al., 2017; Schneider et al., 2018; Park and Geffen, 2020).

Figure Legends

Figure 1: Specific suppression of expected sounds across multiple acoustic dimensions. (A) Schematic of head-fixed lever press training paradigm (Left) and stimulus and reward timing for lever movements (Right). Grey area indicates home position. (B) Global average lever movement trace (Black) and individual animal average lever movement traces (Gray) with position measured as a fraction of the reward threshold. All movements were included, even those that did not reach the reward threshold. (C) Histogram of global lever movement durations (Black, mean 0.77s, median 0.274s) averaged across individual animal histograms (Gray). (D) Histogram of global inter-movement interval (Black, mean 2.9s, median 0.86s) averaged across individual animal histograms (Gray). (E) Schematic of multi-array recording sessions in trained mice (Left) and aggregate neural responses to expected and multiple unexpected sounds in the passive (Darker) and movement-evoked (Lighter) context (Right). Of the 1016 regular-spiking neurons we recorded (N = 5 Animals), a subset of neurons are analyzed for each sound type if they respond to that sound in either context ($p < 0.01$, 0 - 60ms post sound onset). Cell counts are listed below each PSTH. Color differences represent sound frequency, and the likelihood of each lever press producing a given sound type during the recording session is displayed in black bar.

Figure 2: Precise suppression of expected sound responses in individual neurons. (A) Modulation (See Methods) of individual neurons comparing responses to sounds heard in the active and passive condition to each tone type. Negative values indicate weaker responses in the active condition, i.e. suppression. A one-way ANOVA detected differences amongst the groups (F-statistic $p = 2*10^{-17}$), with Exp and Freq being significantly different from all other groups (Exp, $p < 0.005$; Freq, $p < 0.005$). Neuron values and inclusion are the same as (Fig. 1E). (B) Identical experimental setup and analysis as (A), but performed in a subset of mice trained to perform the lever task in the absence of sound (F-statistic p = 0.01). The expected sound was assigned as the sound heard on 90% of trials, though mice had no prior experience with the sound. (C) Matrix (left) representing the absolute difference in neural modulation for stimuli heard in sound trained (top) and silent trained (right) mice. Comparison (bottom) of the absolute difference between the modulation of probe sounds and the expected sound (corresponding to the top row in heatmaps) in sound trained and silent trained mice. Error bars represent standard error in each dimension. (D) Average responses across trials of three individual neurons to each tone type, showing suppression that is specific for the expected sound at the individual neuron level. (E) Confusion matrix showing how accurately sounds could be decoded from auditory cortex neural responses on individual trials. Matrix shows decoding performance averaged across 4 animals.

Figure 3: Abundant prediction error neurons in mouse auditory cortex. (A) Number of neurons responsive ($p < 0.01$) to a given sound in the active context (light), passive context (dark), or both (white). (B) Example neuron depicting the identification of putative prediction error neurons, defined as neurons which respond to a given stimulus type in the active context, but not in the passive context, not at the time of expected sound on omission trials, and not to the expected self-generated sound. Stimulus window of 0 - 60ms post sound onset compared to the 60ms prior to sound onset. (C) Number of neurons that fulfil our putative prediction error criteria for each unexpected trial type. (D) Visual representation of each prediction error neuron's responsiveness (white) to task tones heard in the active condition (left), responsiveness in passive condition (middle), and whether a neuron obeyed our prediction error criteria for a given stimulus (right, see Fig. 3B). To match our strict prediction error criteria, a probability value of 0.01 was used for actively heard unexpected sounds, while a cutoff of $p = 0.1$ was used for the expected, self-generated sound and all passive sounds. Rows with color represent example neurons in (E). (E) Responses of two example neurons to sounds heard actively (top) and passively (bottom). Black PSTHs show significant responses using the p values described in (A). Scale Bar: 25ms, 50 Sp/s (F) Histogram representing the number of passively heard stimuli to which each neuron responded ($p>0.1$) for prediction error neurons (green) and non prediction error neurons (black). This includes all task sounds heard passively as well as passively heard pure tones at half octave intervals between 4khz and 32khz. (G) Histogram representing the neural response to passive tones, averaged across all passive sounds, including half octave separated pure tones described in (F). (H) Quantification of the number of different stimuli for which a neuron signals prediction error. (I) Color-coded matrix showing the number of prediction error neurons that are shared across pairs of stimuli.

Figure 4: Prediction error responses in auditory cortex are short-latency. (A) Raster of example neuron showing action potential timing following frequency probe sounds, with the first spike on a given trial (orange) used to calculate an average onset latency. (B) Histogram of average onset latency following frequency probe trials for prediction error neurons (Green), all neurons responsive to the frequency probe (Orange), and latency of neurons responsive to the passive frequency probe following passive presentation. No difference between prediction error neuron latencies and general latencies in the active condition ($p = 0.97$) or to passive sound responses ($p = 0.97$, $p = 0.87$, KS Test).

559 **Figure 5: Prediction error neurons reflect the violation of a learned expectation.** (A) Quantification of
560 the number of putative prediction error neurons in trained animals, and in an identical experiment and
561 analysis in animals trained to make lever presses in silence. Each dot represents the fraction of neurons
562 responsive in any context to a given sound in a recording session that met the criteria for prediction
563 error neurons. (B) Comparison between the number of prediction error neurons for a stimulus (as in A)
564 and how 'different' a stimulus was from the expected sound. Differences were quantified between
565 neural responses to each probe sound and the expected sound in the passive condition (See Methods).
566 Each dot represents one unexpected stimulus in a sound trained animal (left, N = 4) and difference
567 values were mean-normalized within animal to enable a comparison across animals. Linear regression is
568 shown with shaded standard error. P values and correlation coefficients are listed. Identical analysis but
569 comparing the fraction of PE neurons to a stimulus and the absolute magnitude of an animal's
570 population response to that stimulus in the passive condition (Right). (C) Same as (B) but for mice
571 trained in a silent version of the lever task (N = 3).

572

Acknowledgements

We thank Alessandro La Chioma, Ralph Peterson, and Grant Zempolich for their thoughtful comments on the manuscript. We thank members of the Schneider lab for fruitful discussions. We thank Jessica A Guevara for expert animal care and technical support. This research was supported by the National Institutes of Health (1R01-DC018802 to DMS); a Career Award at the Scientific Interface from the Burroughs Wellcome Fund (D.M.S); fellowships from the Searle Scholars Program, the Alfred P. Sloan Foundation, and the McKnight Foundation (D.M.S.); and an investigator award from the New York Stem Cell Foundation (D.M.S). D.M.S. is a New York Stem Cell Foundation - Robertson Neuroscience Investigator.

Declaration of Interests

The authors declare no competing interests.

585 **References**

586 Audette NJ, Zhou W, Chioma A La, Schneider DM (2022) Precise movement-based predictions in the
587 mouse auditory cortex. Curr Biol:1–16 Available at: https://doi.org/10.1016/j.cub.2022.09.064.

588 Ayaz A, Stäuble A, Hamada M, Wulf MA, Saleem AB, Helmchen F (2019) Layer-specific integration of
589 locomotion and sensory information in mouse barrel cortex. Nat Commun 10 Available at:
590 https://doi.org/10.1038/s41467-019-10564-8.

591 Bastos AM, Usrey WM, Adams RA, Mangun GR, Fries P, Friston KJ (2012) Canonical Microcircuits for
592 Predictive Coding. Neuron 76:695–711 Available at:
593 http://dx.doi.org/10.1016/j.neuron.2012.10.038.

594 Clancy KB, Orsolic I, Mrsic-Flogel TD (2019) Locomotion-dependent remapping of distributed cortical
595 networks. Nat Neurosci 22:778–786 Available at: www.nature.com/reprints.

596 Cristianini N, Shawe-Taylor J (2000) An Introduction to Support Vector Machines and Other Kernel-based
597 Learning Methods. Cambridge University Press. Available at:
598 https://www.cambridge.org/core/product/identifier/9780511801389/type/book [Accessed
599 January 6, 2023].

600 Eliades SJ, Wang X (2008) Neural substrates of vocalization feedback monitoring in primate auditory
601 cortex. Nature 453:1102–1106.

602 Eshel N, Tian J, Bukwich M, Uchida N (2016) Dopamine neurons share common response function for
603 reward prediction error. Nat Neurosci 19:479–486.

604 Farley BJ, Quirk MC, Doherty JJ, Christian EP (2010) Stimulus-specific adaptation in auditory cortex is an
605 NMDA-independent process distinct from the sensory novelty encoded by the mismatch
606 negativity. J Neurosci 30:16475–16484.

607 Fishman YI, Steinschneider M (2012) Searching for the mismatch negativity in primary auditory cortex of
608 the awake monkey: Deviance detection or stimulus specific adaptation? J Neurosci 32:15747–
609 15758.

610 Flinker A, Chang EF, Kirsch HE, Barbaro NM, Crone NE, Knight RT (2010) Single-trial speech suppression
611 of auditory cortex activity in humans. J Neurosci 30:16643–16650.

612 Glimcher PW (2011) Understanding dopamine and reinforcement learning: The dopamine reward
613 prediction error hypothesis. Proc Natl Acad Sci U S A 108:15647–15654.

614 Joachimsthaler B, Uhlmann M, Miller F, Ehret G, Kurt S (2014) Quantitative analysis of neuronal
615 response properties in primary and higher-order auditory cortical fields of awake house mice (Mus
616 musculus). Eur J Neurosci 39:904–918.

617 Jordan R, Keller GB (2020) Opposing Influence of Top-down and Bottom-up Input on Excitatory Layer 2/3
618 Neurons in Mouse Primary Visual Cortex. Neuron 108:1194-1206.e5 Available at:
619 https://doi.org/10.1016/j.neuron.2020.09.024.

620 Keller GB, Bonhoeffer T, Hübener M (2012) Sensorimotor Mismatch Signals in Primary Visual Cortex of
621 the Behaving Mouse. Neuron 74:809–815.

622 Keller GB, Hahnloser RHR (2009) Neural processing of auditory feedback during vocal practice in a
623 songbird. Nature 457:187–190.

624 Keller GB, Mrsic-Flogel TD (2018) Predictive Processing: A Canonical Cortical Computation. Neuron
625 100:424–435 Available at: https://doi.org/10.1016/j.neuron.2018.10.003.

626 Khatri V, Bruno RM, Simons DJ (2009) Stimulus-specific and stimulus-nonspecific firing synchrony and its
627 modulation by sensory adaptation in the whisker-to-barrel pathway. J Neurophysiol 101:2328–
628 2338.

629 Knolle F, Schwartze M, Schröger E, Kotz SA (2019) Auditory Predictions and Prediction Errors in
630 Response to Self-Initiated Vowels. Front Neurosci 13:1–11.

631 Kuchibhotla K V., Gill J V., Lindsay GW, Papadoyannis ES, Field RE, Sten TAH, Miller KD, Froemke RC
632 (2017) Parallel processing by cortical inhibition enables context-dependent behavior. Nat Neurosci
633 20:62–71.

634 Leinweber M, Ward DR, Sobczak JM, Attinger A, Keller GB (2017) A Sensorimotor Circuit in Mouse Cortex
635 for Visual Flow Predictions. Neuron 95:1420-1432.e5 Available at:
636 https://doi.org/10.1016/j.neuron.2017.08.036.

637 Mandelblat-Cerf Y, Las L, Denisenko N, Fee MS (2014) A role for descending auditory cortical projections
638 in songbird vocal learning. Elife 3:1–23.

639 McGinley MJ, Vinck M, Reimer J, Batista-Brito R, Zagha E, Cadwell CR, Tolias AS, Cardin JA, McCormick
640 DA (2015) Waking State: Rapid Variations Modulate Neural and Behavioral Responses. Neuron
641 87:1143–1161 Available at: http://dx.doi.org/10.1016/j.neuron.2015.09.012.

642 Morandell K, Yin A, Triana Del Rio R, Schneider DM (2023) Movement-related modulation in mouse
643 auditory cortex is widespread yet locally diverse. bioRxi.

644 Musall S, Kaufman MT, Juavinett AL, Gluf S, Churchland AK (2019) Single-trial neural dynamics are
645 dominated by richly varied movements. Nat Neurosci 22:1677–1686 Available at:
646 http://dx.doi.org/10.1038/s41593-019-0502-4.

647 Muzzu T, Saleem AB (2021) Feature selectivity can explain mismatch signals in mouse visual cortex
648 Graphical abstract. CellReports 37:109772 Available at:
649 https://doi.org/10.1016/j.celrep.2021.109772 [Accessed December 9, 2021].

650 Näätänen R, Paavilainen P, Rinne T, Alho K (2007) The mismatch negativity (MMN) in basic research of

central auditory processing: A review. Clin Neurophysiol 118:2544–2590.

Narayanan DP, Tsukano H, Kline AM, Onodera K, Kato HK (2023) Biological constraints on stereotaxic targeting of functionally-defined cortical areas. Cereb Cortex 33:3293–3310.

Narsky I, Porter FC (2013) Reducing Multiclass to Binary. Stat Anal Tech Part Phys 1:371–379.

Natan RG, Briguglio JJ, Mwilambwe-Tshilobo L, Jones SI, Aizenberg M, Goldberg EM, Geffen MN (2015) Complementary control of sensory adaptation by two types of cortical interneurons. Elife 4:1–27.

Nelson A, Schneider DM, Takatoh J, Sakurai K, Wang F, Mooney R (2013) A circuit for motor cortical modulation of auditory cortical activity. J Neurosci 33:14342–14353.

Niell CM, Stryker MP (2010) Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex. Neuron 65:472–479 Available at: http://dx.doi.org/10.1016/j.neuron.2010.01.033.

Pachitariu M, Steinmetz N, Kadir S, Carandini M, Kenneth D. H (2016) Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. bioRxiv:061481.

Panigrahi B, Martin KA, Li Y, Graves AR, Vollmer A, Olson L, Mensh BD, Karpova AY, Dudman JT (2015) Dopamine Is Required for the Neural Representation and Control of Movement Vigor. Cell 162:1418–1430.

Park Y, Geffen MN (2020) A circuit model of auditory cortex. PLoS Comput Biol 16:1–29 Available at: http://dx.doi.org/10.1371/journal.pcbi.1008016.

Parras GG, Casado-Román L, Schröger E, Malmierca MS (2021) The posterior auditory field is the chief generator of prediction error signals in the auditory cortex. Neuroimage 242:118446 Available at: https://doi.org/10.1016/j.neuroimage.2021.118446.

Parras GG, Nieto-Diego J, Carbajal G V., Valdés-Baizabal C, Escera C, Malmierca MS (2017) Neurons along

672 the auditory pathway exhibit a hierarchical organization of prediction error. Nat Commun 8
673 Available at: http://dx.doi.org/10.1038/s41467-017-02038-6.

674 Polack P-O, Friedman J, Golshani P (2013) Cellular mechanisms of brain state-dependent gain
675 modulation in visual cortex. 16.

676 Reznik D, Guttman N, Buaron B, Zion-Golumbic E, Mukamel R (2021) Action-locked Neural Responses in
677 Auditory Cortex to Self-generated Sounds. Cereb Cortex 31:5560–5569.

678 Romero S, Hight AE, Clayton KK, Resnik J, Williamson RS, Hancock KE, Polley DB (2020) Cellular and
679 Widefield Imaging of Sound Frequency Organization in Primary and Higher Order Fields of the
680 Mouse Auditory Cortex. Cereb Cortex 30:1603–1622.

681 Rummell BP, Klee JL, Sigurdsson T (2016) Attenuation of responses to self-generated sounds in auditory
682 cortical neurons. J Neurosci 36:12010–12026.

683 Schneider DM (2020) Reflections of action in sensory cortex. Curr Opin Neurobiol 64:53–59 Available at:
684 https://doi.org/10.1016/j.conb.2020.02.004.

685 Schneider DM, Mooney R (2018) How movement modulates hearing. Annu Rev Neurosci 41:553–572.

686 Schneider DM, Nelson A, Mooney R (2014) A synaptic and circuit basis for corollary discharge in the
687 auditory cortex. Nature 513:189–194 Available at: http://dx.doi.org/10.1038/nature13724.

688 Schneider DM, Sundararajan J, Mooney R (2018) A cortical filter that learns to suppress the acoustic
689 consequences of movement. Nature 561:391–395 Available at: http://dx.doi.org/10.1038/s41586-
690 018-0520-5.

691 Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. Science (80-)
692 275:1593–1599.

693 Singla S, Dempsey C, Warren R, Enikolopov AG, Sawtell NB (2017) A cerebellum-like circuit in the
694 auditory system cancels responses to self-generated sounds. Nat Neurosci 20:943–950.

695 Soares S, Atallah B V., Paton JJ (2016) Midbrain dopamine neurons control judgment of time. Science
696 (80-) 354:1273–1277.

697 Solomon SS, Tang H, Sussman E, Kohn A (2021) Limited Evidence for Sensory Prediction Error Responses
698 in Visual Cortex of Macaques and Humans. Cereb Cortex 31:3136–3152.

699 Steinmetz NA, Zatka-Haas P, Carandini M, Harris KD (2019) Distributed coding of choice, action and
700 engagement across the mouse brain. Nature 576:266–273 Available at:
701 http://dx.doi.org/10.1038/s41586-019-1787-x.

702 Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, Harris KD (2019) Spontaneous behaviors
703 drive multidimensional, brainwide activity. Science (80-) 364.

704 Taaseh N, Yaron A, Nelken I (2011) Stimulus-specific adaptation and deviance detection in the rat
705 auditory cortex. PLoS One 6.

706 Tobler PN, Fiorillo CD, Schultz W (2005) Adaptive Coding of Reward Value by Dopamine Neurons.
707 Science (80-) 203:1642–1645.

708 Ulanovsky N, Las L, Farkas D, Nelken I (2004) Multiple time scales of adaptation in auditory cortex
709 neurons. J Neurosci 24:10440–10453.

710 Ulanovsky N, Las L, Nelken I (2003) Processing of low-probability sounds by cortical neurons. Nat
711 Neurosci 6:391–398.

712 Wilson NR, Runyan CA, Wang FL, Sur M (2012) Division and subtraction by distinct cortical inhibitory
713 networks in vivo. Nature 488:343–348.

714 Yang L, Lee K, Villagracia J, Masmanidis SC (2020) Open source silicon microprobes for high throughput
715 neural recording. J Neural Eng 17.

716 Ylinen S, Huuskonen M, Mikkola K, Saure E, Sinkkonen T, Paavilainen P (2016) Predictive coding of
717 phonological rules in auditory cortex: A mismatch negativity study. Brain Lang 162:72–80 Available
718 at: http://dx.doi.org/10.1016/j.bandl.2016.08.007.

719 Zhou M, Liang F, Xiong XR, Li L, Li H, Xiao Z, Tao HW, Zhang LI (2014) Scaling down of balanced excitation
720 and inhibition by active behavioral states in auditory cortex. Nat Neurosci 17:841–850.

721 Zmarz P, Keller GB (2016) Mismatch Receptive Fields in Mouse Visual Cortex. Neuron 92:766–772
722 Available at: http://dx.doi.org/10.1016/j.neuron.2016.09.057.

723