The Rapid Emergence of Auditory Object Representations in Cortex Reflect Central Acoustic Attributes

Mattson Ogg^{a,c}, Thomas A. Carlson^b, L. Robert Slevc^{a,c}

^a Neuroscience and Cognitive Program, University of Maryland, 0219 Cole Student Activities Building, 4090 Union Drive, College Park, MD 20742

^b School of Psychology, The University of Sydney, Sydney, NSW, 2006, Australia

^c Department of Psychology, University of Maryland, Biology/Psychology Building, 4094 Campus Dr., College Park, MD 20742

Note: This is the authors' final version of an article accepted for publication in the Journal of Cognitive Neuroscience

Correspondence should be sent to Mattson Ogg, Neuroscience and Cognitive Science Program, Department of Psychology, University of Maryland, Biology-Psychology Building, Room 3150 4094 Campus Drive, College Park, MD 20742 or via email at: mogg@umd.edu

Abstract

Human listeners are bombarded by acoustic information that the brain rapidly organizes into coherent percepts of objects and events in the environment, which aids speech and music perception. The efficiency of auditory object recognition belies the critical constraint that acoustic stimuli necessarily require time to unfold. Using magnetoencephalography (MEG), we studied the time course of the neural processes that transform dynamic acoustic information into auditory object representations. Participants listened to a diverse set of 36 tokens comprising everyday sounds from a typical human environment. Multivariate pattern analysis was used to decode the sound tokens from the MEG recordings. We show that sound tokens can be decoded from brain activity beginning 90 milliseconds after stimulus onset with peak decoding performance occurring at 155 milliseconds post stimulus onset. Decoding performance was primarily driven by differences between category representations (e.g., environmental vs. instrument sounds), although within-category decoding was better than chance. Representational similarity analysis revealed that these emerging neural representations were related to harmonic and spectrotemporal differences among the stimuli, which correspond to canonical acoustic features processed by the auditory pathway. Our findings begin to link the processing of physical sound properties with the perception of auditory objects and events in cortex.

Keywords: Magnetoencephalography, Object Perception, Auditory, Multivariate Pattern Analysis, Representational Similarity Analysis

Introduction and Background

Successful navigation of the auditory environment requires quickly recognizing the stimuli or objects one encounters. This is true for the assessment of potential threats in the auditory scene as well as for the perception of speech and music (Bizley & Cohen, 2013, Griffiths & Warren, 2004; Ogg & Slevc, 2019b). However, the recognition of auditory objects is complicated by the fact that acoustic stimuli require time to develop. Thus, a listener's auditory system must rapidly transform acoustic information that is still evolving into a mental representation of an object or event to maintain behavioral relevance (for example, when hearing an approaching vehicle that poses a threat to the observer). Despite this complexity, listeners can make accurate judgments about the identity of a sound even when the sound's temporal development (and thus potential to relay useful acoustic detail) is limited in duration to just tens of milliseconds (Ogg, Slevc, & Idsardi, 2017) or less (Robinson & Patterson, 1995a; Robinson & Patterson, 1995b; Suied, Agus, Thorpe, Mesgarani, & Pressnitzer, 2014). In the current study, we asked how the auditory system uses incoming acoustic information from natural sound stimuli to help construct representations of auditory objects and events given these temporal constraints. Specifically, we investigated whether certain acoustic events or features are temporally prioritized in a listener's processing of natural sounds and how such prioritization relates to particular acoustic properties of the sounds.

One possibility is that the auditory system hierarchically processes certain kinds of behaviorally relevant information. Recordings in the ventral auditory pathway of nonhuman primates support this view. These studies find more selective responses for increasingly complex stimuli in anterior regions of the temporal lobes (Rauschecker &

Scott, 2009; Rauschecker & Tian, 2000; Tian, Reser, Durham, Kustov, & Rauschecker, 2001; Perrodin, Kayser, Logothetis, & Petkov, 2011) that correspond to different processing latencies (Kikuchi, Horwitz, & Mishkin, 2010). These findings have inspired human EEG work that suggests neural responses differ as a function of the sound source, with man-made sounds and human vocal sounds exhibiting stronger responses around 70-ms and 170-ms, respectively (Charest et al., 2009; De Lucia, Clarke, & Murray, 2010; Murray, Camen, Andino, Bovet, & Clarke, 2006). However, other results suggest stronger responses to instrument sounds compared to voices around 100-ms (Rigoulot, Pell, & Armony, 2015) or vice versa later around 320-ms (Levy, Granot, & Bentin, 2001; Levy, Granot, & Bentin, 2003).

Previous work has examined broad classes of sounds relative to one another. However, there are many inherent (and important) acoustic differences that exist between different categories of sound (or even among sound sources within the same category) that could drive these results (Ogg & Slevc, 2019a), which have not been explicitly analyzed in previous M/EEG studies of natural sounds. Additionally, work on classic auditory evoked responses (using simple synthesized stimuli) such as the mismatch negativity (Caclin et al., 2006; Giard et al., 1995; Rosburg, 2003), N1 (Näätänen & Picton, 1987) and M100 response (Roberts, Ferrari, Stufflebeam, & Poeppel, 2000) indicate that the auditory system is sensitive to specific acoustic qualities (e.g., spectral and temporal envelopes, noisiness, spectral variability). The manner by which acoustic features might influence early, time varying neural representations of natural sounds remains largely unexplored despite extensive examination via other methods such as fMRI (Spectral Centroid: Alluri et al., 2012; Giordano, McAdams, Zatorre, Kriegeskorte,

& Belin, 2012; Allen, Burton, Olman, & Oxenham, 2017; Spectral Flatness: Alluri et al., 2012, Lewis, Talkington, Tallaksen, & Frum, 2012; Spectral Envelope: Warren, Jennings, & Griffiths, 2005; Ogg, Moraczewski, Kuchinsky, & Slevc, 2019; Noisiness: Giordano et al., 2012; Lewis et al., 2009; Fundamental Frequency: Giordano et al., 2012, Allen et al., 2017; Patterson, Uppenkamp, Johnsrude, & Griffiths, 2002; Attack Time: Menon et al., 2002; Loudness: Giordano et al., 2012; Langers, van Dijk, Schoenmaker, & Backes, 2007; Spectrotemporal Modulation: Norman–Haignere, Kanwisher, & McDermott, 2015; Schönwiesner & Zatorre, 2009) and ECoG (Hullett, Hamilton, Mesgarani, Schreiner, & Chang, 2016).

The key to understanding the neural and auditory processing involved in the formation of object and event representations is not just in differentiating the magnitude of different neural responses. Rather, the goal is to understand how information pertaining to an object or stimulus feature is represented in the brain (Hebart & Baker 2018). Machine learning approaches have yielded such insights for vision researchers by using classifiers to "decode" (via multivariate pattern analysis, or MVPA; Tong & Pratte, 2012) the stimuli participants perceived based on time varying patterns of neural activity measured using MEG (Cichy, Pantazis, & Oliva, 2014; Carlson, Hogendoorn, Kanai, Mesik, & Turret, 2011; Carlson, Tovar, Alink, & Kriegeskorte, 2013; Grootswagers, Wardle, & Carlson, 2017; Figure 1a). Such an approach is often more sensitive than examining response magnitude or latency (Grootswagers et al., 2017; Formisano, De Martino, Bonte, & Goebel, 2008; Haynes & Rees, 2006; Hebart & Baker, 2018; Staeren, Renvall, De Martino, Goebel, & Formisano, 2009).

Better or worse classification accuracy indicates that the brain's response patterns for two stimuli are more or less distinct and thus provides a difference metric for the neural representations of categories and object exemplars. These neural classification accuracy/dissimilarity rates can be contained in a matrix that is organized by item pairs (called a representational dissimilarity matrix or "RDM" see Figure 1a). Such classification results can then be used to understand how physical properties of the stimuli relate to participants' neural responses by correlating the observed patterns of neural classification accuracy with other matrices that pertain to differences among the stimuli (e.g., RDMs of differences among the same item-pairs along various acoustic feature dimensions; Figure 1b). This approach, called representational similarity analysis (Kriegeskorte & Kievit, 2013), provides a powerful tool that can reveal which feature dimensions the classifier uses to decode these neural representations as they emerge in time.

Relatively few studies have used decoding techniques to study neural responses to natural auditory stimuli over time in M/EEG (although see Khalighinejad, da Silva, & Mesgarani, 2017; Sankaran, Thompson, Carlile, & Carlson, 2018; Teng, Sommer, Pantazis, & Oliva, 2017). So far, the small body of work using these methods has focused only on specific subsets of stimuli (e.g., speech *or* musical pitches). This precludes a broader understanding of the acoustic dimensions that contribute to how human listeners and the auditory system distinguish auditory objects and events.

In the current study, we employed a wide array of natural sounds relevant to everyday functioning in human listeners (speech from different speakers, sounds from musical instruments, and sounds from everyday objects and materials) to probe the

acoustic dimensions that influence early neural responses and the formation of auditory object and event representations. As participants listened to these stimuli, we recorded their brain activity using MEG. Our approach leverages both the acoustic qualities that naturally differ among individual auditory stimuli and the temporal nature of sound. This allows us to glean an understanding of the acoustic qualities that subserve the emerging neural representations of auditory objects as both the stimuli and corresponding neural responses develop over time.

Methods

Participants

We recruited 20 participants (9 female) who were fluent English speakers, right handed and self-reported normal hearing and no neurological disorders. Participants were not selected for musical experience but, typical of a university population (Schellenberg, 2006), most had some musical training (mean = 6.6, SD = 3.9 years). This study was conducted with the approval of the University of Maryland Institutional Review Board and all participants gave their informed consent to participate.

Stimuli

Participants heard 36, 300-ms natural sound tokens evenly divided among three categories (see Figure 2). *Human-environmental* sounds were a selection of 12 everyday object sounds from a variety of media or events (e.g., fingers typing on a keyboard, toilet flushing, engine ignition; sampled from the BBC Sound Effects Library, 1997; and CMU sound events databases, 2008; Vettel, 2010). *Musical instrument* tokens were 12

individual notes played by 6 instruments (2 notes each) that spanned the range of orchestral timbres: bass clarinet, cello (arco), marimba, tenor trombone, acoustic guitar and piano (from the University of Iowa Musical Instrument Samples Database; 1997). *Speech* tokens were 12 excerpts of the utterances “bead,” “bod,” “heed,” and “hod” as well as an excerpted /a/ and /i/. These were spoken by a male and female raised in the Mid-Atlantic United States and recorded in-house (under conditions described in Ogg et al., 2017 and Ogg & Slevc, 2019a). These consonants were selected to match the variability of the other sound categories in terms of their noisiness and temporal envelope dynamics. A cat vocalization was also obtained (from the BBC sound effects library) for use as a catch trial stimulus.

Notes for each musical instrument were selected to match one of the male and female speakers’ utterances (median fundamental frequency calculated using the YIN algorithm; De Cheveigné & Kawahara, 2002) to control for large fundamental frequency differences between these categories. This required the use of lower register instruments (e.g., cello, trombone, etc.) to match the male speaker’s fundamental frequencies. Because the /h/ portions of the speech utterances were especially noisy, fundamental frequency matches for these stimuli were based on the last third of the token.

Stimuli were edited to begin at their onsets, defined here as 5-ms prior to where the absolute value of the stimulus’ amplitude reached 10% of its maximum. All stimuli were then truncated to a 300-ms duration, which is sufficient to facilitate sound source identification following onset (Ogg et al., 2017; Ogg & Slevc, 2019a; Robinson & Patterson, 1995; Suied et al., 2014). Cosine onset and offset ramps (both 5-ms) were then applied. Stimulus playback in the MEG scanner was equalized (Ultracurve Pro

DEQ2496, Behringer) to be approximately flat (± 4 dB) between 40 Hz and 5.5 kHz above which the sound level diminished. Thus, we low-pass filtered our stimuli (zero-phase 4th order Butterworth filter) at 5.5 kHz prior to RMS-level normalization. This achieved an average level of 72.5 dB-A as measured after playback by Presentation (NeuroBehavioral Systems), equalization and then delivery via a mixer (1202-VLZ Pro, LOUD Technologies Inc.), amplifier (SLA1, ART ProAudio), insert earphone tubing (E-A-RTONE Gold 3A Insert Earphone, 3M) and insert earphones (Etymotic Research).

Recording of Stimuli for Presentation and Analysis

In order to match our acoustic analyses to what participants heard in the MEG as closely as possible we derived acoustic features from high fidelity recordings of the stimuli from the scanner presentation system using a 2cc coupler and 1/2 inch condenser microphone amplified by a 824S/PRM902 sound level meter system (Larson Davis) into an Apogee Duet (Apogee Electronics). Recordings of the stimuli for behavioral screening were recorded using an approximate 2cc plastic coupler, a Sony ECM-144 Lavalier microphone (Sony Corporation) and a Zoom H1 (Zoom Video Communications) portable recorder.

Procedure

Screening and Familiarization

Participants attended an approximately half-hour screening session no more than 11 days prior to scanning (mean = 5.1 days), to ensure they could accurately identify the stimuli and to make sure any acoustic changes imposed by the scanner's playback

apparatus did not impair the identifiability of the tokens. Participants were first allowed to freely listen to each labeled sound file as many times as they liked. Screening then involved six blocks of trials run in PsychoPy (version 1.83.4; Peirce, 2007). On each trial a stimulus was repeated three times with an ISI of 700-ms (one stimulus trial per each of six blocks for a total of 18 presentations of each stimulus during screening). In the first block, participants simply listened to the stimuli and read each sound's label on-screen (no response required). In the second block, participants pushed a button to indicate the sound they heard on each trial. In the response key, each of the twelve speech and human-environmental sounds was assigned a unique response button (as was the cat stimulus) and each of the six instruments was assigned one button for use regardless of the note the instrument played (thus, 21 response buttons total). In this second block, participants were given feedback if they were right or corrective feedback (sound source and response button). The third block (the test block) was identical to the second but without feedback. The last three blocks of trials were the same as the first three except that these trials presented a recording of each stimulus that was made from the MEG sound playback system. Screening performance was assessed in terms of accuracy on the sixth block of trials (testing block with scanner-playback recordings). This ensured that participants could accurately identify the stimuli as they would hear them in the scanner. Participants were required to score at least 89% (33 of 37 correct) to move on to the scan (mean performance = 95%, SD = 3%).

Data Acquisition and Preprocessing

MEG responses were recorded using 157 axial gradiometers in the scanner (Kanazawa Institute of Technology), as participants listened to 100 repetitions of each stimulus presented in random order at a jittered rate of 1 per second ($\pm$ 150-ms drawn from a uniform distribution). MEG data were filtered online with a notch filter at 60 Hz, a low-pass filter at 200 Hz and were recorded at a sampling rate of 1000 Hz. Before scanning, head digitization was performed (POLHEMUS 3 Space Fast Track system), and electrodes were attached to two preauricular sites and three prefrontal sites to spatially coregister participants' head shapes with the MEG. To ensure participants paid attention to the stimuli, their task during the scan was to press a button upon hearing the cat vocalization (presented at an average rate of one catch trial for every 12 stimuli), while they watched a self-selected film on mute. Participants were very accurate on this task (mean accuracy = 93%, SD = 7%) indicating that they remained attentive to the sounds throughout the scan.

Each participant's MEG data was first cleaned with a time-shift PCA algorithm (De Cheveigné & Simon, 2007). The data were then epoched around the presentation of each stimulus from -200 to 600-ms, and down-sampled to 200 Hz (5-ms temporal resolution). The extra 100-ms on each end of the -100 to 500-ms epoch was eliminated to remove edge artifacts. Each SQUID sensor was then baseline corrected to the average of its -100-ms baseline.

Decoding Analysis

Decoding was performed to distinguish the neural responses to each pair of sound tokens at the individual subject level using CoSMoMVPA (Oosterhof, Connolly &

Haxby, 2016; see Figure 1). At each 5-ms time point throughout the epoch, a given pair of sounds was selected and a linear discriminant classifier was trained (using 90% of the data) to associate each of the two item labels with their corresponding MEG responses using the amplitudes of the SQUID sensors (all 157 sensors for that time point) as feature vectors for classification. The classifier was then tested on its ability to correctly assign item labels to the responses in the left-out trials (the last 10% of the data). This repeated such that each cross-validation fold served as both a training and test case and the analysis iterated through each unique item-pair (630 pairs of tokens, i.e., one half of the symmetrical dissimilarity matrix excluding the diagonal) and through each 5-ms time point in the epoch (see Figure 1a). Note that essentially the same results were obtained based on a PCA of the MEG data prior to decoding (thus eliminating any dependence among feature vectors in training and testing, see Grootswagers et al., 2017, for a discussion).

The decoding analyses yielded a matrix of classification accuracies for each item-pair and subject at each time point (averaged over cross-validation folds). Classification results within the RDMs were averaged for each subject to assess when: 1) overall decoding accuracy across all item-pairs exceeded chance (i.e., averaged across the entire RDM, Figure 2B top), 2) decoding accuracy for items within or between different categories exceeded chance or differed from one another (i.e., averaged among quadrants on or off the diagonal in the RDM; quadrants in Figure 2C vs. 2D resulting in Figure 2B bottom), or 3) decoding accuracy for individual within or between-category comparisons exceeded chance or differed from one another (i.e., averaged within quadrants in the RDM, Figure 2C or 2D). Because we always decoded individual pairs of stimuli from

one another, chance performance for all classification analyses (i.e., all item and category-level comparisons) was 50%. Statistical significance with multiple comparison corrections was assessed via nonparametric (Maris & Oostenveld, 2007) bootstrapped (10,000 iterations) sign-permutation tests (vs. chance), t -tests (between vs. within-category) or ANOVAs (comparing among the three within-category or between-category decoding accuracies in Figure 2C and 2D, respectively) followed by threshold-free cluster-enhancement (Smith & Nichols, 2009) as implemented by CoSMoMVPA (Oosterhof et al., 2016).

Acoustic Feature Extraction and Analyses

To ensure the acoustic analyses were the most accurate representation of the stimuli heard in the scanner, acoustic features were derived from high fidelity recordings of each stimulus from the scanner's playback system. From these recordings we extracted acoustic features known to influence the neural processing of sound in humans (Alluri et al., 2012; Giordano et al., 2012; Allen et al., 2017; Lewis et al., 2012; Warren et al., 2005; Lewis et al., 2009; Patterson et al., 2002; Ogg & Slevc, 2019a; 2019b Menon et al., 2002; Langers et al., 2007; Norman-Haignere et al., 2015; Schönwiesner & Zatorre, 2009), via widely used algorithms (*YIN*: De Cheveigné & Kawahara, 2002; *Timbre Toolbox*: Peeters, Giordano, Susini, Misdariis, & McAdams, 2011; Kazazis, Esterer, Depalle, & McAdams, 2017; *Modulation Power Spectra*: Elliott & Theunissen, 2009). Temporal dynamics of the stimuli (regarding the speed of sound onset) were characterized via the log-attack-time and temporal centroid of the energy envelope (Peeters et al., 2011, Kazazis et al., 2017). Spectral qualities were assessed via spectral

centroid (similar to overall brightness), spectral flatness (similar to overall noisiness of the spectrum) and spectral variability (index of spectral change over time) extracted from an ERB (equivalent rectangular-bandwidth) gammatone representation (i.e., a cochleagram) of the sounds (Peeters et al., 2011, Kazazis et al., 2017) in 5-ms windows/increments. We included a proxy for loudness (“frame energy” feature from Peeters et al., 2011, which we denote as “ERB Energy”) to examine the influence of local changes in dynamics in the stimuli, and also to account for any differences imparted by the playback system (because perceived loudness is difficult to characterize even after standard RMS normalization, especially for complex natural stimuli that vary in their temporal and frequency characteristics; Giordano et al., 2012; Langers et al., 2007; Moore, 2012). We also included the raw ERB cochleagram that these features were based on (Peeters et al., 2011, Kazazis et al., 2017).

Modulation power spectra were used to capture different spectral and temporal rates of change in the stimuli (Elliott & Theunissen, 2009; Thoret, Depalle, & McAdams, 2017; Gaussian windowed spectrogram, 50 dB dynamic range, 16 Hz frequency and 10-ms temporal resolution yielding maximum modulation of 31.22 cyc/kHz and 48.39 Hz, respectively). Finally, we extracted aperiodicity and fundamental frequency estimates (De Cheveigné & Kawahara, 2002) from each stimulus in 5-ms increments (minimum of the fundamental frequency search range set to 100 Hz). Missing values were returned for the incomplete analysis windows in the very first and last increments, which were replaced with the nearest non-missing values.

Acoustic analyses examined both global (representing the stimulus as a whole) and local (each 5-ms increment) stimulus qualities. Global stimulus features comprised

any attribute that aggregated over the entire duration of the stimulus (log-attack-time, temporal centroid, modulation power spectrum), as well as the median across windowed features. Local stimulus features comprised the 5-ms increments of the windowed measures: spectral (i.e., ERB-representation-based) features, aperiodicity, and fundamental frequency features. Acoustic RDMs were created by taking the absolute value of the difference of each feature between stimulus pairs (single-valued features) or by calculating Euclidean distances between stimulus pairs (for features comprised of multiple values: modulation power spectra and ERB cochleagram).

Acoustic correlation analyses were first carried out for each subject individually via a Kendall's tau correlation between the neural RDMs for that subject at each time point and the acoustic RDMs for each feature (either the global acoustic RDMs, Fig. 3a, or the local acoustic RDMs at each time point, Fig. 3b). These individual subject correlation statistics were assessed for each feature at the group level at each time point using a Wilcoxon signed-rank test against 0. Group-level statistics for each acoustic feature were then false discovery rate corrected across time points at a threshold of FDR-corrected $p < 0.001$ (Benjamini & Hochberg, 1995).

Figure 1. Schematic of the decoding and acoustic analyses. (A) Decoding algorithm for item-pairs used to construct a neural RDM. (B) Acoustic correlation analysis between the neural RDMs throughout the epoch and the global or local acoustic feature RDMs. The red box in B around the acoustic waveform (middle) indicates that the data correspond to a global attribute (aggregated across all time points in the stimulus). Red vertical bars on the MEG and acoustic waveforms indicate that the data relate to an individual 5-ms time point. Note, only one instance of each of the unique 630 stimulus pairs was decoded and the data from this half of the RDM (e.g., item 1 vs. item 2) was copied to the other half (e.g., item 2 vs. item 1) only for illustrative purposes here and in Figure 2.

Results

We first determined how accurately our sound tokens could be decoded from one another based on participants' MEG responses throughout the time series. For this, we analyzed the average pairwise decoding accuracy rates for all possible pairs of sounds (Figure 2A) for each subject throughout the time series. As seen in Fig. 2B (top), these grand averaged decoding accuracy rates exceeded chance in a cluster of time points beginning 90-ms after sound onset. This cluster remained statistically significant throughout the epoch (500-ms), notably even after sound offset at 300-ms (assessed against chance using bootstrapped sign-permutation testing and threshold-free cluster enhancement across time points, all $p < 0.05$ corrected). Peak grand-average decoding accuracy was 58.4%, which occurred 155-ms after sound onset. Thus, information related to the neural representations of these auditory objects and events can be reliably distinguished very early on in the listeners' MEG response.

Figure 2. Decoding accuracy. (A) Group-averaged neural dissimilarity matrices (RDMs) at 0, 150, and 300ms. (B) Grand-average decoding accuracy (top; $p < 0.05$ corrected vs. chance) and average between vs. within-category decoding accuracy (bottom; $p < 0.05$ corrected between vs. within). (C) Within-category decoding accuracy for each sound category ($p < 0.05$ corrected vs. chance). (D) Between-category decoding accuracy for each pair of sound categories ($p < 0.05$ corrected vs. chance). Square insets depict the region of the RDM that was averaged for each comparison. The top of B encompasses all the insets in C and D, and the bottom of B represents the insets of C vs. D. Error bands represent SEM calculated over subject-level average decoding accuracies for each comparison.

Decoding Categories of Sound

An examination of the neural RDMs (Figure 2A) suggests that neural responses to pairs of stimuli from different categories (i.e., instrument vs. human-environmental) were better decoded than pairs of stimuli from the same category. To quantify this, we averaged accuracy rates for the regions of the RDM corresponding to between and within-category comparisons (Figure 2B, bottom). Between- and within-category decoding followed a very similar time course: both exceeded chance in clusters of time points beginning 90-ms after onset and continuing to the end of the epoch (all $p < 0.05$ corrected). Decoding accuracy peaked for both comparisons at similar times as well (155-ms for between and 160-ms for within). However, between-category classification was more accurate (peak = 60.0%) than within-category decoding (peak = 54.8%) throughout a cluster of time points starting at 95-ms and lasting through the end of the epoch (assessed using nonparametric bootstrapped t -tests and threshold-free cluster enhancement across time points comparing between vs. within-category accuracy, all $p < 0.05$ corrected).

We next examined the decoding accuracy for specific within- or between-category comparisons of individual sound categories (Figure 2, C-D insets). Within-category decoding accuracy rates were lower overall (compare Figures 2C and 2D, see also Figure 2B), but sounds within each category could nevertheless be decoded from one another better than chance among clusters of time points spanning most of the epoch (Figure 2C; cluster start times for within-speech: 90-ms; within-human-environment: 95-ms; within-instrument: 90-ms, all lasting through 500-ms, all $p < 0.05$ corrected). However, the time courses and peaks of the within-category decoding varied within a

cluster of time points beginning at 115-ms, lasting until 470-ms (assessed using nonparametric bootstrapped ANOVA and threshold-free cluster enhancement over time points comparing within-category accuracy rates, all $p < 0.05$ corrected). Peak within-instrument decoding (54.1%) occurred at 145-ms, and peak within-human-environmental sound decoding (54.7%) occurred later at 215-ms. Interestingly, within-speech decoding accuracy was the highest among the within-category comparisons (56.7%) and peaked at 160-ms.

Significant clusters of time points for all combinations of between-category decoding began at 90-ms and lasted throughout the epoch (Figure 2D; all $p < 0.05$ corrected). Accuracy rates for these comparisons also differed significantly from one another in a cluster of time points throughout the epoch beginning at 120-ms (all $p < 0.05$ corrected). Decoding between human-environmental and instrument sounds was the highest overall (peak: 62.5% at 155-ms) followed by human-environmental vs. speech decoding (peak: 59.3% at 165-ms) and instrument vs. speech token decoding (peak: 58.9% at 150-ms).

In sum, shortly after overall decoding first exceeded chance at 90-ms and prior to the peak in overall decoding at 155-ms, neural representations already began to differ across sound categories. Note that neither the subjects' level of musical training or behavioral performance (at screening or during the scan) were correlated with subject-level averages of decoding accuracy at any time point (overall, in Figure 2B top, or for any combination of categories in Figure 2C or 2D, all FDR corrected $p > 0.05$ across time points for each measure).

Processing of Global Acoustic Features

Speech, instrument and human-environmental sounds contain different acoustic regularities that likely relate to the neural representations we decoded. To investigate this possibility, we computed Kendall's tau correlations between the pairwise neural classification accuracy rates (each subject's neural RDMs throughout the epoch) and the pairwise acoustic differences among our stimuli along a variety of dimensions (acoustic RDMs; Figure 1B, center; see Methods for details). Correlations between neural decoding rates and global acoustic differences among the stimuli (i.e., for acoustic measures that operated over the whole sound token or the median across time-varying features) are shown in Figure 3A. These features were strongly associated with decoding rates before and around the peak in decoding accuracy that was observed around 155-ms (all FDR-corrected $p < 0.001$). Differences among stimuli in terms of their modulation power spectra exhibited the earliest significant correlations with the neural RDM (onset $r_\tau = 0.03$ at 100-ms) followed a few milliseconds later by overall spectral variability (onset $r_\tau = 0.03$ at 115-ms), and aperiodicity (onset $r_\tau = 0.06$ at 120-ms). Other significant correlations emerged later for temporal centroid (onset $r_\tau = 0.05$ at 135-ms), overall spectral envelope (assessed by an ERB cochleagram; onset $r_\tau = 0.03$ at 135-ms), spectral centroid (onset $r_\tau = 0.05$ at 145-ms) and the sum of the squared ERB amplitudes across cochlear filter channels (which we used as a proxy for loudness to account for any changes from the sound playback apparatus in the scanner, 125 to 135-ms, peak $r_\tau = 0.04$). These features mostly remained correlated with the neural RDMs and increased in their correlation strength around the decoding peak at 155-ms except for ERB energy, and temporal centroid which were non-significant after 135-ms. The strongest

correlations with the neural RDMs around the decoding peak were observed for differences in aperiodicity ($r_t = 0.18$ at 165-ms) followed by modulation power spectra ($r_t = 0.14$ at 140-ms) and spectral variability ($r_t = 0.13$ at 165-ms). Interestingly, many of the features that exhibited a significant association with decoding rates earliest in the epoch are complex acoustic representations known to be associated with some of the earliest processes in auditory cortex: aperiodicity, spectral envelope and spectrotemporal variability (Lewis et al., 2009; Norman–Haignere et al., 2015; Chi, Ru, & Shamma, 2005; Theunissen & Elie, 2014). Later in the epoch, particularly after sound offset, significant correlations pertained to spectral centroid (peak $r_t = 0.09$ at 445-ms), spectral flatness (peak $r_t = 0.06$ at 460-ms), and temporal centroid (peak $r_t = 0.11$ at 250-ms). Overall, spectral variability (peak $r_t = 0.16$ at 435-ms), and aperiodicity (peak $r_t = 0.19$ at 350-ms), exhibited the strongest correlations throughout the epoch.

We conducted an additional analysis to examine the influence of fundamental frequency. This was possible because we allowed our speech stimuli to vary naturally in their fundamental frequencies and matched the notes of the instrument tokens along this dimension. However, a global acoustic analysis of fundamental frequency among just the instrument and speech stimulus pairs (excluding the pairs with /h/ consonant tokens, because these were predominantly noise) did not reveal any significant correlation with the neural decoding results.

Figure 3. The correlation between neural decoding accuracy (neural RDMs) and acoustic feature differences (acoustic RDMs) throughout the epoch plotted as averages of subject-level correlations. (A) Global acoustic analysis results. Mean correlation statistics and standard errors across participants are plotted for each feature throughout the neural epoch (* FDR-corrected $p < 0.001$; error bands represent SEM calculated over subject-level correlations). (B) Local acoustic analysis results. Mean correlation statistics across participants at the neural and acoustic time points that were significant at the group level (FDR-corrected $p < 0.001$) are shown.

Processing of Local Acoustic Features

Sound stimuli necessarily unfold over time. Thus, it is possible that local changes in the acoustic signal could have consequences for participants' unfolding neural responses. To understand how local acoustic developments precipitated corresponding changes in neural representations of these sounds, we expanded our global acoustic

analysis approach in a way that leveraged the dynamic information in our stimuli. In this local acoustic analysis (Figure 1B, right) the pairwise acoustic feature differences among the sounds at every 5-ms time point in the stimuli (local acoustic RDMs) were correlated with every subsequent 5-ms RDM throughout the MEG epoch (neural RDMs; see Methods for details).

The local acoustic analysis results are depicted in Figure 3B. The earliest significant correlation between a local acoustic RDM and the neural RDM was for spectral flatness (onset $r_t = 0.05$ at 100-ms in the neural epoch correlated with the 5-ms local acoustic RDM) and spectral variability (onset $r_t = 0.03$ at 100-ms in the neural epoch correlated with the 40-ms local acoustic RDM). Again, aperiodicity, spectral centroid/envelope and spectral variability (across many stimulus time points) were strongly and significantly correlated with the decoding results around the decoding peak (155-ms). We also see a small influence of differences in local ERB energy from approximately the first 100-ms of the stimuli on subsequent neural representations. These local acoustic features were more strongly associated with neural decoding accuracy rates than the global features, with aperiodicity (peak $r_t = 0.21$ at 360-ms in the neural epoch correlated with the 55-ms local acoustic RDM and $r_t = 0.20$ at 165-ms in the neural epoch correlated with the 105-ms local acoustic RDM) and spectral variability (peak $r_t = 0.20$ at 160-ms in the neural epoch correlated with the 110-ms local acoustic RDM) exhibiting some of the strongest correlations. Aperiodicity and spectral variability were also significantly associated with decoding across neural and acoustic time points. Meanwhile, spectral differences (including spectral centroid, flatness and ERB energy) from the earlier portions of the stimuli appear to exert the most influence.

Discussion

The questions of “when” and “how” a listener recognizes what they hear are fundamentally linked due to the temporal nature of sound. We examined these questions using MVPA analyses of MEG time series data in conjunction with representational similarity analyses of acoustic differences among a set of natural sound stimuli that are behaviorally relevant to human listeners. Using this approach, we gained a better understanding of the acoustic qualities associated with the formation of auditory object and event representations as they unfold in time.

We found that neural activity measured using MEG contained information sufficient to distinguish individual sounds beginning 90-ms after sound onset, reaching a peak at 155-ms, and continuing after sound offset at 300-ms until the end of the neural epoch at 500-ms. This pattern of results was driven by especially accurate decoding of stimuli from different sound categories, although individual instruments, speech utterances and everyday object sounds could all be decoded from one another better than chance. Acoustic analyses indicated that these emergent representations were associated with differences among the stimuli pertaining to spectrotemporal variability, aperiodicity (or harmonic content), and the sounds’ spectral and temporal envelopes. The importance of these features fluctuated over time with spectral envelope and modulation spectra exerting their greatest influence soon after onset, while spectral centroid and temporal centroid had a greater influence later in the epoch. Spectral variability and aperiodicity were significantly associated with decoding throughout much of the epoch. These are fundamental acoustic qualities that have been shown to influence neural responses and perception, but our results go further to suggest that these cues are specifically associated

with the dynamic formation of mental representations of what a listener hears early after sound onset.

Our analyses needed to manage dynamic information within our stimuli due to the physical (temporal) constraints of sound. By leveraging this information, we obtained a picture of how the brain rapidly integrates and organizes sensory information from natural sounds. However, processing dynamic stimulus information is also an important part of everyday object and event perception beyond the auditory domain. Thus, the approach we describe here could provide valuable insight into other areas of cognitive neuroscience and object perception as well. For example, this could be applied to dynamic visual stimuli or multi-sensory integration paradigms to better understand how the brain integrates and represents specific stimulus information from both modalities to achieve percepts of complex objects and events.

Relative to the decoding of visual objects from MEG, the decoding onsets and latencies that we observed were slower than some findings (Cichy et al., 2014; Isik, Meyers, Leibo, & Poggio, 2014) but comparable to others (Carlson et al., 2011; Carlson et al., 2013). This difference might be due to the inherent temporal nature of our sound stimuli relative to the pictures used in previous work on visual objects. Unlike sounds, picture stimuli present all the information necessary for object recognition at once. Auditory stimuli on the other hand require time to unfold, thus incrementally relaying the information that supports object recognition. Also, while overall decoding accuracy differed for within and between-category comparisons among our auditory stimuli, the emergence of category and individual object representations followed a roughly similar time course. In most cases these exceeded chance around 90-ms and peaked around 155-

ms. Studies of visual object decoding suggest somewhat wider variability in the latency of decoding individual objects or categories (Cichy et al., 2014; Carlson et al., 2011; Carlson et al., 2013).

A related study (Teng et al., 2017) decoded neural representations of different impact sounds and reverberant spaces from MEG, finding earlier onsets and latencies in their sound source decoding than what we observe here. The slower latencies in our results are likely a function of the more varied stimulus onsets (amplitude envelopes) we employed in our set of sounds. This follows from the diverse set of natural sound sources and events that we examined, compared to the constrained set of impact sounds used in previous work (Teng et al., 2017). Thus, our results provide a useful expansion upon this prior work given that the heterogeneous stimulus set examined here is more representative of the diversity of sounds humans encounter in everyday life (which can involve different amplitude envelopes).

Exploring the mechanisms of sound source decoding further, our results suggest that the auditory system prioritizes cues related to a sound's spectrotemporal variability and aperiodicity as well as its spectral and temporal envelopes (at least among this set of sounds). These features can effectively distinguish instrument and human-environmental sounds (especially early after onset), which occupy opposing extremes on these dimensions (Ogg et al., 2017; Ogg & Slevc, 2019a). Speech utterances, on the other hand (an especially important sound category for humans), exploit all of these acoustic cues and their extremes from moment to moment to relay linguistic information (Smith & Lewicki, 2006; Stilp & Kluender, 2010). Speech sounds also exhibited the highest within-category decoding indicating that individual speech sounds are

more distinguishable from one another in listeners' neural responses relative to the instrument and human-environmental sounds. Thus, the acoustic fluctuation among speech tokens might account for this increased within-category decoding.

While these acoustic qualities have previously been shown to influence the strength of neural responses on different time scales (Caclin et al., 2006; Giard et al., 1995; Rosburg, 2003; Näätänen & Picton, 1987; Roberts et al., 2000; Alluri et al., 2012; Giordano et al., 2012; Allen et al., 2017; Lewis et al., 2012; Warren et al., 2005; Lewis et al., 2009; Patterson et al., 2002; Menon et al., 2002; Langers et al., 2007; Norman–Haignere et al., 2015; Schönwiesner & Zatorre, 2009; Hullett et al., 2016), our results go further by beginning to link these features to representations of auditory objects and events early in perception. Additionally, we show that these features, which have typically been examined in fMRI (Alluri et al., 2012; Giordano et al., 2012; Allen et al., 2017; Lewis et al., 2012; Warren et al., 2005; Lewis et al., 2009; Patterson et al., 2002; Menon et al., 2002; Langers et al., 2007; Ogg et al., 2019; Norman–Haignere et al., 2015; Schönwiesner & Zatorre, 2009), vary in their influence over time, with more primary acoustic processes in cortex associated with decoding accuracy early in the response such as spectral envelope and modulation power spectra (Norman–Haignere et al., 2015, Chi et al., 2005, Theunissen & Elie, 2014). Later in the response, however, neural representations were more related to spectral and temporal centroid features (with spectral variability and aperiodicity correlating strongly throughout). Finally, the time course of these neural and acoustic processes aligns with behavioral identification and categorization results for duration-gated stimuli (e.g., Ogg et al., 2017; Suied et al., 2014). Similar acoustic qualities (from the first tens of milliseconds in the stimuli) appear

to be associated with the differentiation of sounds in early MEG decoding time windows (starting around 100-200 milliseconds) and in subsequent behavioral responses (later around 300-600 milliseconds; Agus et al., 2012; Ogg et al., 2017; Suied et al., 2014).

One important caveat to these findings and a potential avenue for future research is that it is not clear how (or when) these decoded representations correspond to more abstract qualities or semantic labels associated with these sounds compared to their acoustic features. Separating semantic labels from acoustic processing in the neural responses obtained in the present study is difficult because these were not orthogonal in this natural and diverse stimulus set. Indeed, much of our decoding results could follow from differences in acoustic or stimulus-level processing rather than object-level representations per se. This potentially aligns with the timing of the decoding peak we observed, which also coincides with components of the auditory evoked response (particularly N1-P2, between 100 to 200 ms). Evoked responses in this time window are known to be influenced by acoustic differences among stimuli (Caclin et al., 2006; Giard et al., 1995; Rosburg, 2003; Näätänen & Picton, 1987; Roberts et al., 2000) and appear to support representations of speech (Bidelman, & Walker, 2017; Chang et al., 2010; Näätänen & Picton, 1987; Poeppel et al., 1997; Roberts et al., 2000). Note, however, that this decoding peak could also be related to the overall improved signal to noise ratio in the neural response that follows from these prominent evoked components. Note also that decoding remained significantly above chance after these time points as well. In either case, we go beyond much of the previous work examining natural sounds in M/EEG by describing the acoustic features that might underpin processing differences between sound categories.

The ostensibly modest absolute decoding accuracy rates we observed (peaking at around 8.4% above chance) might raise some concerns regarding the reliability of these results. However, as Hebart and Baker (2018) point out, while high decoding rates are clearly desirable for brain-computer-interface applications, for understanding *brain function* (which is our present goal), above chance decoding is all that is necessary to inform how the brain processes information (especially given the resolution we have access to with current imaging technologies). Moreover, they make the important point that decoding accuracy is not equivalent to an effect size, and that, again, the key is to understand where (or, in our case, “when”) information is contained or processed in the system (i.e., when accuracy exceeds chance). Analogous points can be made regarding the strength of correlations made using these decoding accuracy rates (such as in our acoustic analyses) since decoding accuracy places an upper limit on the strength of the correlations that can be observed. Finally, it is worth keeping in mind that essentially the same decoding results were obtained based on a principal components analysis of the MEG data, which indicates that our results are not due to any deleterious effects of interdependence among feature vectors (MEG data) in training and testing (Grootswagers et al., 2017).

It is interesting that we did not observe an influence of fundamental frequency or musical training in these neural decoding results. However, this might be due to aspects of our design or task. For example, the non-influence of fundamental frequency could reflect our use of isolated sound tokens (rather than longer concurrent streams) or a range restriction in the fundamental frequencies of the speaker’s utterances (which spanned approximately one octave). Similarly, it is possible that fundamental frequency requires

larger time scales on which to exert an influence (Robinson & Patterson, 1995a; Robinson & Patterson, 1995b; Walker, Bizley, King & Schnupp, 2011), beyond the time windows examined here. And while musical training has sometimes been found to influence auditory cortical responses (Bidelman, Weiss, Moreno, & Alain, 2014), these effects might arise more robustly during demanding behavioral tasks (see also Alho et al., 2016; Bidelman & Walker, 2017). Indeed, other findings based on passive tasks with diverse, readily distinguishable stimuli have not found an influence of musical training on neural representations (Ogg et al., 2019).

More generally, attention is known to strengthen neural responses and neural representations of auditory stimuli (Ding & Simon 2012), particularly during categorization tasks (Alho et al., 2016; Bidelman, & Walker, 2017). However, the current study used a mostly passive paradigm (non-target odd-ball detection task). Thus, it is likely that stronger attentional demands to specific target stimuli (or categories) could boost their decoding accuracy, and could potentially modify the influence of certain target-specific acoustic features. Some findings also indicate that attention plays a particularly acute role in auditory scene analysis, or streaming, tasks around 200-ms (O’Sullivan et al., 2015; Snyder, Alain, & Picton, 2006). Viewed in conjunction with our results, this suggests that, after object identification occurs most acutely around 150-ms, attention might act to tune responses to a given stimulus and track a specific sound source through an auditory scene. This is thought to occur via feedback mechanisms from frontal areas (Fritz, David, Radtke-Schuller, Yin, & Shamma, 2010). Viewed alongside our significant acoustic correlations with aperiodicity, modulation spectra and the spectral envelope, this suggests that these features could be avenues through which auditory

cortex is tuned to a given sound via the modulation of attention (Fritz, Elhilali, David, & Shamma, 2007; Petkov et al., 2004). Future work examining how semantic or acoustic cues factor into attentional modulation will also help further illuminate this issue. These exciting new directions along with our results will help paint a fuller picture of how sounds are transformed from acoustic phenomena into the events and objects we identify and subsequently track in the environment in service of more complex behaviors and individual goals.

Acknowledgements

We are grateful for the discussions, feedback and assistance we received in the course of this work from Natalia Lapinskaya, Christopher Neufeld, Jonathan Simon, Christian Brodbeck, Ed Smith, Tijl Grootswagers, Amanda Robinson, and Stefanie Kuchinsky. A portion of the human-environmental sounds were obtained courtesy of the Sound Events Database (<http://www.auditorylab.org/>; Copyright 2008, Laurie M. Heller. Funding provided by NSF award 0446955). Thanks also to Carnegie Mellon's Center for the Neural Basis of Cognition regarding the human-environmental sounds, and to Lawrence Fritts for managing the University of Iowa Musical Instrument Samples database.

References

- Alho, J., Green, B. M., May, P. J., Sams, M., Tiitinen, H., Rauschecker, J. P., & Jääskeläinen, I. P. (2016). Early-latency categorical speech sound representations in the left inferior frontal gyrus. *NeuroImage*, *129*, 214–223.
- Allen, E. J., Burton, P. C., Olman, C. A., & Oxenham, A. J. (2017). Representations of pitch and timbre variation in human auditory cortex. *Journal of Neuroscience*, *37*, 1284–1293.
- Alluri, V., Toiviainen, P., Jääskeläinen, I. P., Glerean, E., Sams, M., & Brattico, E. (2012). Large-scale brain networks emerge from dynamic processing of musical timbre, key and rhythm. *NeuroImage*, *59*, 3677–3689.
- BBC Sound Effects Library (1997) BBC Worldwide, London, United Kingdom.
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, *57*, 289–300.
- Bidelman, G. M., & Walker, B. S. (2017). Attentional modulation and domain-specificity underlying the neural organization of auditory categorical perception. *European Journal of Neuroscience*, *45*, 690–699.
- Bizley, J. K., & Cohen, Y. E. (2013). The what, where and how of auditory–object perception. *Nature Reviews Neuroscience*, *14*, 693–707.
- Caclin, A., Brattico, E., Tervaniemi, M., Näätänen, R., Morlet, D., Giard, M. H., & McAdams, S. (2006). Separate neural processing of timbre dimensions in auditory sensory memory. *Journal of Cognitive Neuroscience*, *18*, 1959–1972.
- Carlson, T. A., Hogendoorn, H., Kanai, R., Mesik, J., & Turret, J. (2011). High temporal resolution decoding of object position and category. *Journal of Vision*, *11*.
- Carlson, T. A., Tovar, D. A., Alink, A., & Kriegeskorte, N. (2013). Representational dynamics of object vision: The first 1000 ms. *Journal of Vision*, *13*.
- Charest, I., Pernet, C. R., Rousselet, G. A., Quiñones, I., Latinus, M., Fillion-Bilodeau, S., Chartrand, J. P., & Belin, P. (2009). Electrophysiological evidence for an early processing of human voices. *BMC Neuroscience*, *10*, 127.
- Chi, T., Ru, P., & Shamma, S. A. (2005). Multiresolution spectrotemporal analysis of complex sounds. *Journal of the Acoustical Society of America*, *118*, 887–906.
- Cichy, R. M., Pantazis, D., & Oliva, A. (2014). Resolving human object recognition in space and time. *Nature Neuroscience*, *17*, 455–462.

- De Cheveigné, A., & Kawahara, H. (2002). YIN, a fundamental frequency estimator for speech and music. *Journal of the Acoustical Society of America*, *111*, 1917–1930.
- De Cheveigné, A., & Simon, J. Z. (2007). Denoising based on time-shift PCA. *Journal of Neuroscience Methods*, *165*, 297–305.
- De Lucia, M., Clarke, S., & Murray, M. M. (2010). A temporal hierarchy for conspecific vocalization discrimination in humans. *Journal of Neuroscience*, *30*, 11210–11221.
- Ding, N., & Simon, J. Z. (2012). Emergence of neural encoding of auditory objects while listening to competing speakers. *Proceedings of the National Academy of Sciences*, *109*, 11854–11859.
- DeWitt, I., & Rauschecker, J. P. (2012). Phoneme and word recognition in the auditory ventral stream. *Proceedings of the National Academy of Sciences, USA*, *109*, E505–E514.
- Elliott, T. M., & Theunissen, F. E. (2009). The modulation transfer function for speech intelligibility. *PLoS Computational Biology*, *5*, e1000302.
- Formisano, E., De Martino, F., Bonte, M., & Goebel, R. (2008). "Who" is saying "what"? Brain-based decoding of human voice and speech. *Science*, *322*, 970–973.
- Fritz, J. B., David, S. V., Radtke-Schuller, S., Yin, P., & Shamma, S. A. (2010). Adaptive, behaviorally gated, persistent encoding of task-relevant auditory information in ferret frontal cortex. *Nature Neuroscience*, *13*, 1011–1019.
- Fritz, J. B., Elhilali, M., David, S. V., & Shamma, S. A. (2007). Auditory attention—focusing the searchlight on sound. *Current Opinion in Neurobiology*, *17*, 437–455.
- Giard, M. H., Lavikainen, J., Reinikainen, K., Perrin, F., Bertrand, O., Pernier, J., & Näätänen, R. (1995). Separate representation of stimulus frequency, intensity, and duration in auditory sensory memory: An event-related potential and dipole-model analysis. *Journal of Cognitive Neuroscience*, *7*, 133–143.
- Giordano, B. L., McAdams, S., Zatorre, R. J., Kriegeskorte, N., & Belin, P. (2012). Abstract encoding of auditory objects in cortical activity patterns. *Cerebral Cortex*, *23*, 2025–2037.
- Griffiths, T. D., & Warren, J. D. (2004). What is an auditory object?. *Nature Reviews Neuroscience*, *5*, 887–892.
- Groetswagers, T., Wardle, S. G., & Carlson, T. A. (2017). Decoding dynamic brain patterns from evoked responses: A tutorial on multivariate pattern analysis applied to time series neuroimaging data. *Journal of Cognitive Neuroscience*, *29*, 677–697.

- Haynes, J. D., & Rees, G. (2006). Neuroimaging: Decoding mental states from brain activity in humans. *Nature Reviews Neuroscience*, 7, 523–534.
- Hebart, M. N., & Baker, C. I. (2018). Deconstructing multivariate decoding for the study of brain function. *NeuroImage*, 180, 4–18.
- Hullett, P. W., Hamilton, L. S., Mesgarani, N., Schreiner, C. E., & Chang, E. F. (2016). Human superior temporal gyrus organization of spectrotemporal modulation tuning derived from speech stimuli. *Journal of Neuroscience*, 36, 2014–2026.
- Isik, L., Meyers, E. M., Leibo, J. Z., & Poggio, T. (2014). The dynamics of invariant object recognition in the human visual system. *Journal of Neurophysiology*, 111, 91–102.
- Kazazis, S., Esterer, N., Depalle, P., & McAdams, S. (2017). A performance evaluation of the timbre toolbox and the MIRtoolbox on calibrated test sounds. *Proceedings of the 2017 International Symposium on Musical Acoustics*, 144–147.
- Khalighinejad, B., da Silva, G. C., & Mesgarani, N. (2017). Dynamic encoding of acoustic features in neural responses to continuous speech. *Journal of Neuroscience*, 37, 2176–2185.
- Kikuchi, Y., Horwitz, B., & Mishkin, M. (2010). Hierarchical auditory processing directed rostrally along the monkey's supratemporal plane. *Journal of Neuroscience*, 30, 13021–13030.
- Kriegeskorte, N., & Kievit, R. A. (2013). Representational geometry: Integrating cognition, computation, and the brain. *Trends in Cognitive Science*, 17, 401–412.
- Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S., & Baker, C. I. (2009). Circular analysis in systems neuroscience: The dangers of double dipping. *Nature Neuroscience*, 12, 535–540.
- Langers, D. R., van Dijk, P., Schoenmaker, E. S., & Backes, W. H. (2007). fMRI activation in relation to sound intensity and loudness. *NeuroImage*, 35, 709–18.
- Levy, D. A., Granot, R., & Bentin, S. (2001). Processing specificity for human voice stimuli: Electrophysiological evidence. *Neuroreport*, 12, 2653–2657.
- Levy, D. A., Granot, R., & Bentin, S. (2003). Neural sensitivity to human voices: ERP evidence of task and attentional influences. *Psychophysiology*, 40, 291–305.
- Lewis, J. W., Talkington, W. J., Tallaksen, K. C., & Frum, C. A. (2012). Auditory object salience: Human cortical processing of non-biological action sounds and their acoustic signal attributes. *Frontiers in Systems Neuroscience*, 6, 27.
- Lewis, J. W., Talkington, W. J., Walker, N. A., Spirou, G. A., Jajosky, A., Frum, C., & Brefczynski-Lewis, J. A. (2009). Human cortical organization for processing

- vocalizations indicates representation of harmonic structure as a signal attribute. *Journal of Neuroscience*, 29, 2283–2296.
- Liegeois-Chauvel, C., Musolino, A., Badier, J. M., Marquis, P., & Chauvel, P. (1994). Evoked potentials recorded from the auditory cortex in man: Evaluation and topography of the middle latency components. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section*, 92, 204–214.
- Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, 164, 177–190.
- Menon, V., Levitin, D. J., Smith, B. K., Lembke, A., Krasnow, B. D., Glazer, D., Glover, G. H., & McAdams, S. (2002). Neural correlates of timbre change in harmonic sounds. *NeuroImage*, 17, 1742–1754.
- Moore, B. C. J. (2012). *An Introduction to the Psychology of Hearing* (Emerald, Bingley, UK), 6th Ed.
- Murray, M. M., Camen, C., Andino, S. L. G., Bovet, P., & Clarke, S. (2006). Rapid brain discrimination of sounds of objects. *Journal of Neuroscience*, 26, 1293–1302.
- Näätänen, R., & Picton, T. (1987). The N1 wave of the human electric and magnetic response to sound: a review and an analysis of the component structure. *Psychophysiology*, 24, 375–425.
- Norman-Haignere, S., Kanwisher, N. G., McDermott, J. H. (2015). Distinct cortical pathways for music and speech revealed by hypothesis-free voxel decomposition. *Neuron*, 88, 1281–1296.
- O’Sullivan, J. A., Power, A. J., Mesgarani, N., Rajaram, S., Foxe, J. J., Shinn-Cunningham, B. G., Slaney, M., Shamma, S. A., & Lalor, E. C. (2015). Attentional selection in a cocktail party environment can be decoded from single-trial EEG. *Cerebral Cortex*, 25, 1697–1706.
- Ogg, M., Moraczewski, D., Kuchinsky, S. E., & Slevc, L. R. (2019). Separable neural representations of sound sources: Speaker identity and musical timbre. *NeuroImage*, 191, 116–126.
- Ogg, M., Slevc, L. R., & Idsardi, W. J. (2017). The time course of sound category identification: Insights from acoustic features. *Journal of the Acoustical Society of America*, 142, 3459–3473.
- Ogg, M., & Slevc, L. R. (2019a). Acoustic Correlates of Auditory Object and Event Perception: Speakers, Musical Timbres, and Environmental Sounds. *Frontiers in Psychology*, 10.

- Ogg, M., & Slevc, L. R. (2019b). Neural mechanisms of music and language. In G. Zubizaray, and N. Schiller (Eds.), *Oxford Handbook of Neurolinguistics* (pp. 907–952). New York, NY: Oxford University Press.
- Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J. M. (2011). FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational Intelligence and Neuroscience*, 2011, 1.
- Oosterhof, N. N., Connolly, A. C., & Haxby, J. V. (2016). CoSMoMVPA: Multi-modal multivariate pattern analysis of neuroimaging data in Matlab/GNU Octave. *Frontiers in Neuroinformatics*, 10, 27.
- Patterson, R. D., Uppenkamp, S., Johnsrude, I. S., & Griffiths, T. D. (2002). The processing of temporal pitch and melody information in auditory cortex. *Neuron*, 36, 767–776.
- Peeters G, Giordano BL, Susini P, Misdariis N, McAdams S (2011) The timbre toolbox: Extracting audio descriptors from musical signals. *Journal of the Acoustical Society of America*, 130, 2902–2916.
- Peirce, J. W. (2007). PsychoPy—psychophysics software in Python. *Journal of Neuroscience Methods*, 162, 8–13.
- Perrodin, C., Kayser, C., Logothetis, N. K., & Petkov, C. I. (2011). Voice cells in the primate temporal lobe. *Current Biology*, 21, 1408–1415.
- Petkov, C. I., Kang, X., Alho, K., Bertrand, O., Yund, E. W., & Woods, D. L. (2004) *Nature Neuroscience*, 7, 658–663.
- Poeppel, D., Phillips, C., Yellin, E., Rowley, H. A., Roberts, T. P., & Marantz, A. (1997). Processing of vowels in supratemporal auditory cortex. *Neuroscience Letters*, 221, 145–148.
- Raij, T., Ahveninen, J., Lin, F. H., Witzel, T., Jääskeläinen, I. P., Letham, B., Israeli, E., Sahyoun, C., Vasios, C., Stufflebeam, S., Hämäläinen, M., & Belliveau, J. W. (2010). Onset timing of cross-sensory activations and multisensory interactions in auditory and visual sensory cortices. *European Journal of Neuroscience*, 31, 1772–1782.
- Rauschecker, J. P., & Scott, S. K. (2009). Maps and streams in the auditory cortex: Nonhuman primates illuminate human speech processing. *Nature Neuroscience*, 12, 718–724.
- Rauschecker, J. P., & Tian, B. (2000). Mechanisms and streams for processing of “what” and “where” in auditory cortex. *Proceedings of the National Academy of Sciences, USA*, 97, 11800–11806.

- Rigoulot, S., Pell, M. D., & Armony, J. L. (2015). Time course of the influence of musical expertise on the processing of vocal and musical sounds. *Neuroscience*, 290, 175–184.
- Roberts, T. P., Ferrari, P., Stufflebeam, S. M., & Poeppel, D. (2000). Latency of the auditory evoked neuromagnetic field components: Stimulus dependence and insights toward perception. *Journal of Clinical Neurophysiology*, 17, 114–129.
- Robinson, K., & Patterson, R. D. (1995a). The duration required to identify the instrument, the octave, or the pitch chroma of a musical note. *Music Perception*, 13, 1–15.
- Robinson, K., & Patterson, R. D. (1995b). The stimulus duration required to identify vowels, their octave, and their pitch chroma. *Journal of the Acoustical Society of America*, 98, 1858–1865.
- Rosburg, T. (2003). Left hemispheric dipole locations of the neuromagnetic mismatch negativity to frequency, intensity and duration deviants. *Cognitive Brain Research*, 16, 83–90.
- Sankaran, N., Thompson, W. F., Carlile, S., & Carlson, T. A. (2018). Decoding the dynamic representation of musical pitch from human brain activity. *Scientific Reports*, 8, 839.
- Schellenberg, E. G. (2006). Long-term positive associations between music lessons and IQ. *Journal of Educational Psychology*, 98, 457–468.
- Schönwiesner, M., & Zatorre, R. J. (2009). Spectro-temporal modulation transfer function of single voxels in the human auditory cortex measured with high-resolution fMRI. *Proceedings of the National Academy of Sciences, USA*, 106, 14611–14616.
- Smith, E. C., & Lewicki, M. S. (2006). Efficient auditory coding. *Nature*, 439, 978–982.
- Smith, S. M., & Nichols, T. E. (2009). Threshold-free cluster enhancement: Addressing problems of smoothing, threshold dependence and localisation in cluster inference. *NeuroImage*, 44, 83–98.
- Snyder, J. S., Alain, C., & Picton, T. W. (2006). Effects of attention on neuroelectric correlates of auditory stream segregation. *Journal of Cognitive Neuroscience*, 18(1), 1–13.
- Sound Events and Real World Events Databases (2008) Carnegie Mellon University, Pittsburgh, PA.
- Staeren, N., Renvall, H., De Martino, F., Goebel, R., & Formisano, E. (2009). Sound categories are represented as distributed patterns in the human auditory cortex. *Current Biology*, 19, 498–502.

- Stilp, C. E., & Kluender, K. R. (2010). Cochlea-scaled entropy, not consonants, vowels, or time, best predicts speech intelligibility. *Proceedings of the National Academy of Sciences, USA*, 107, 12387–12392.
- Suied, C., Agus, T. R., Thorpe, S. J., Mesgarani, N., & Pressnitzer, D. (2014). Auditory gist: Recognition of very short sounds from timbre cues. *Journal of the Acoustical Society of America*, 135, 1380–1391.
- Teng, S., Sommer, V. R., Pantazis, D., & Oliva, A. (2017). Hearing scenes: A neuromagnetic signature of auditory source and reverberant space separation. *eNeuro*, 4, ENEURO-0007.
- Theunissen, F. E., & Elie, J. E. (2014). Neural processing of natural sounds. *Nature Reviews Neuroscience*, 15, 355–366.
- Thoret, E., Depalle, P., & McAdams, S. (2017). Perceptually salient regions of the modulation power spectrum for musical instrument identification. *Frontiers in Psychology*, 8, 587.
- Tian, B., Reser, D., Durham, A., Kustov, A., & Rauschecker, J. P. (2001). Functional specialization in rhesus monkey auditory cortex. *Science* 292, 290–293.
- Tong, F., & Pratte, M. S. (2012). Decoding patterns of human brain activity. *Annual Review of Psychology*, 63, 483–509.
- Tsunada, J., & Cohen, Y. E. (2014). Neural mechanisms of auditory categorization: From across brain areas to within local microcircuits. *Frontiers in Neuroscience*, 8, 161.
- University of Iowa (1997) Musical Instrument Samples Database
- Vettel, J. M. (2010). Neural integration of multimodal events. Doctoral dissertation, Brown University, Providence, RI.
- Walker, K. M., Bizley, J. K., King, A. J., & Schnupp, J. W. (2011). Multiplexed and robust representations of sound features in auditory cortex. *Journal of Neuroscience*, 31, 14565–14576.
- Warren, J. D., Jennings, A. R., & Griffiths, T. D. (2005). Analysis of the spectral envelope of sounds by the human brain. *NeuroImage*, 24, 1052–1057.