

Emergence of neural encoding of auditory objects while listening to competing speakers

Nai Ding^a and Jonathan Z. Simon^{a,b,1}

Departments of ^aElectrical and Computer Engineering and ^bBiology, University of Maryland, College Park, MD 20742

Edited by Thomas D. Albright, The Salk Institute for Biological Studies, La Jolla, CA, and approved June 6, 2012 (received for review April 5, 2012)

A visual scene is perceived in terms of visual objects. Similar ideas have been proposed for the analogous case of auditory scene analysis, although their hypothesized neural underpinnings have not yet been established. Here, we address this question by recording from subjects selectively listening to one of two competing speakers, either of different or the same sex, using magnetoencephalography. Individual neural representations are seen for the speech of the two speakers, with each being selectively phase locked to the rhythm of the corresponding speech stream and from which can be exclusively reconstructed the temporal envelope of that speech stream. The neural representation of the attended speech dominates responses (with latency near 100 ms) in posterior auditory cortex. Furthermore, when the intensity of the attended and background speakers is separately varied over an 8-dB range, the neural representation of the attended speech adapts only to the intensity of that speaker but not to the intensity of the background speaker, suggesting an object-level intensity gain control. In summary, these results indicate that concurrent auditory objects, even if spectrotemporally overlapping and not resolvable at the auditory periphery, are neurally encoded individually in auditory cortex and emerge as fundamental representational units for top-down attentional modulation and bottom-up neural adaptation.

spectrotemporal response function | reverse correlation | phase locking | selective attention

In a complex auditory scene, humans and other animal species can perceptually detect and recognize individual auditory objects (i.e., the sound arising from a single source), even if strongly overlapping acoustically with sounds from other sources. To accomplish this remarkably difficult task, it has been hypothesized that the auditory system first decomposes the complex auditory scene into separate acoustic features and then binds the features, as appropriate, into auditory objects (1–4). The neural representations of auditory objects, each the collective representation of all the features belonging to the same auditory object, have been hypothesized to emerge in auditory cortex to become fundamental units for high-level cognitive processing (5–7). The process of parsing an auditory scene into auditory objects is computationally complex and cannot as yet be emulated by computer algorithms (8), but it occurs reliably, and often effortlessly, in the human auditory system. For example, in the classic “cocktail party problem,” where multiple speakers are talking at the same time (9), human listeners can selectively attend to a chosen target speaker, even if the competing speakers are acoustically more salient (e.g., louder) or perceptually very similar (such as of the same sex) (10).

To demonstrate an object-based neural representation that could subserve the robust perception of an auditory object, several key pieces of evidence are needed. The first is to demonstrate neural activity that exclusively represents a single auditory object (4, 7). In particular, such an object-specific representation must be demonstrated in a range of auditory scenes with reliable perception of that auditory object, and especially in challenging scenarios in which the auditory object cannot be easily segregated by any basic acoustic features, such as frequency or binaural cues. For this reason, we investigate the existence of object-specific auditory representations by using an auditory scene consisting of a pair of

concurrent speech streams mixed into a single acoustic channel. In this scenario, the two speech streams each form a distinct perceptual auditory object but they overlap strongly in time and frequency, and are not separable using spatial cues. Therefore, any neural representation of an auditory object (i.e., in this case, a single stream of speech) would not emerge without complex segregation and grouping processes.

Second, the neural processing of an auditory object must also be adaptive and independent (2, 4). In particular, the neural processing of each auditory object should be modulated based on its own behavioral importance and acoustic properties, without being influenced by the properties of other auditory objects or the stimulus as a whole. Building on the well-established phenomena of feature-based top-down attentional modulation (11–14) and feature-based bottom-up neural adaptation to sound intensity (15), we investigate here whether such top-down and bottom-up modulations occur separately for individual auditory objects (i.e., in an object-based manner). Specifically, using this speech segregation paradigm, we ask the listeners to attend to one of the two speakers while manipulating separately the intensity of the attended and background speakers. If an observed neural representation is object-based, not only must it be enhanced by top-down attention but it must adapt to the intensity change of that speech stream alone, without being affected by the intensity change of the other stream or of the mixture as a whole.

In this study, we investigate whether a robust neural representation of an auditory object can be observed in the brain, and when and where it might emerge. In the experiment, the subjects selectively listened to one of two concurrent spoken narratives mixed into a single acoustic channel, answering comprehension questions about the attended spoken narrative after each 1-min stimulus. The neural recordings were obtained using magnetoencephalography (MEG), which is well suited to measure spatially coherent neural activity synchronized to speech rhythms (i.e., the slow temporal modulations that define the speech envelope) (16–19). Such spatially coherent phase-locked activity is strongly modulated by attention (20–22) and has been hypothesized to play a critical role in grouping acoustic features into auditory objects (3).

Specifically, we hypothesize that in cortical areas with an object-based representation, neural activity should phase lock to the rhythm of a single auditory object, whereas in cortical areas in which object-based representations are not yet formed, or formed only weakly, the neural response should phase lock to the envelope of the entire physical stimulus (i.e., the speech mixture) (both examples are illustrated in Fig. 1*A* and *B*). In other words, what is encoded in the neural response, whether a single speech stream or the mixture, can be easily distinguished by which sound's

Author contributions: N.D. and J.Z.S. designed research, performed research, contributed new reagents/analytic tools, analyzed data, and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. E-mail: jzsimon@umd.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1205381109/-DCSupplemental.




Fig. 1. Illustration of object-based neural representations. Here, the auditory scene is illustrated using a mixture of two concurrent speech streams. (A) If a complex auditory scene is not neurally parsed into separate auditory objects, cortical activity (*Upper, curve*) phase locks to the temporal envelope of the physical stimulus [i.e., the acoustic mixture (*Lower, waveform*)]. (B) In contrast, using the identical stimulus (but illustrated here with the unmixed instances of speech in different colors), for a hypothetical neural representation of an individual auditory object, neural activity would instead selectively phase lock to the temporal envelope only of that auditory object. (C) Neural representation of an auditory object should, furthermore, neurally adapt to an intensity change of its own object (*Upper*) but should remain insensitive to intensity changes in another auditory object (*Lower*). Neither of these modifications to the acoustic stimulus therefore significantly changes the neural representation (comparing A and C).

Results

Deciphering the Spatial-Temporal Code for Individual Speakers. In

the first experiment, listeners selectively listened to one of two competing speakers of different sex, mixed into a single acoustic channel with equal intensity. To probe object-specific neural representations, we reconstructed the temporal envelope of each of the two simultaneous speech streams by optimally integrating MEG activity over time and space (i.e., sensors). Such a reconstruction of the envelope of each speech stream, rather than the physical stimulus, can be successful only if the stimulus mixture is neurally segregated (“unmixed”) and the speech streams of the two speakers are represented differentially. We first reconstructed the temporal envelope of the attended speech. Fig. 24 shows representative segments of the different envelopes reconstructed by this decoder, from listeners hearing the identical speech mixture but attending to different speakers in it. Clearly, the reconstructed envelope depends strongly on the attentional focus of the listener and resembles the envelope of the attended speech. At the single-subject level and the single-trial level, the reconstructed envelope is

more strongly correlated with the envelope of the attended speaker than of the unattended speaker ($P < 0.001$, paired permutation test; Fig. 2*B*, *Left*). This attention-dependent neural reconstruction is seen in 92% of trials (Fig. S1).

We also reconstructed the temporal envelope of the background speech using a second decoder that integrates neural activity spatiotemporally in a different way. The result of this reconstruction is indeed more correlated with the envelope of the background speech rather than of the attended speech ($P < 0.005$, paired permutation test; Fig. 2B, Right). Therefore, by integrating the temporal and spatial neural responses in two distinct ways, the attended and background speech can be successfully decoded separately. On average, the reconstruction for the background speech is more correlated with the background speech in 73% of the trials from individual subjects (Fig. S1; significantly above chance level; $P < 0.002$, binomial test). In this experiment, the speakers are of opposite sex, but the neural representations of segregated speech streams can be similarly demonstrated even for

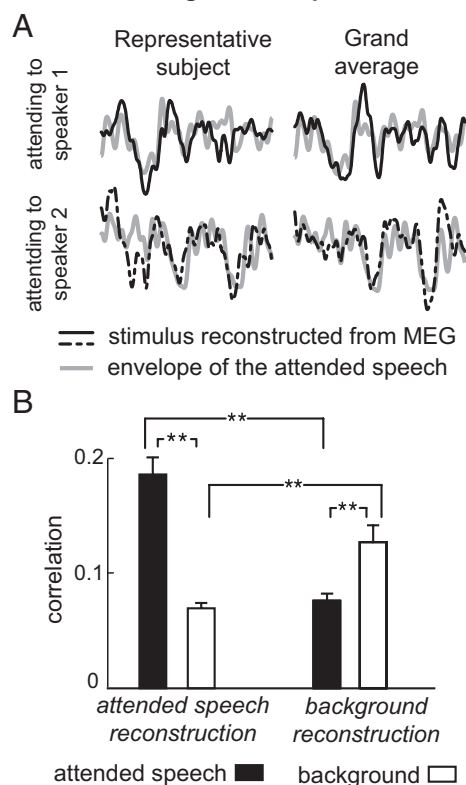


Fig. 2. Decoding the cortical representation specific to each speech stream. (A) Examples of the envelope reconstructed from neural activity (black), superimposed on the actual envelope of the attended speech when presented in isolation (gray). (*Upper and Lower*) Different envelopes are decoded from neural responses to identical stimuli, depending on whether the listener attends to one or the other speaker in the speech mixture, with each resembling the envelope of the attended speech. Here, the signals, 5 s in duration, are averaged over three trials for illustrative purposes, but all results in the study are based on single-trial analysis. (B) Two separate decoders reconstruct the envelope of the attended and background speech, respectively, from their separate spatial-temporal neural responses to the speech mixture. The correlation between the decoded envelope and the actual envelope of each speech stream is shown in the bar graph (averaged over trials and speakers), with each error bar denoting 1 SEM across subjects (** $P < 0.005$, paired permutation test). The separate envelopes reconstructed by the two decoders selectively resemble that of attended and background speech, demonstrating a separate neural code for each speech stream.

the more challenging scenario where the two speakers are of the same sex (Fig. S2). Consequently, these results demonstrate that the neural representation in auditory cortex goes beyond encoding just the physically presented stimulus (the speech mixture) and shows selective phase locking to auditory objects (Fig. S3).

Robustness to the Intensity of Either Speaker. When the intensity of either of the two competing speaker changes, up to 10 dB, human listeners can still understand either speaker with more than 50% intelligibility (10). Intensity gain control may contribute to this robustness in speech perception. Here, we address whether intensity gain control occurs globally for an auditory scene or separately for each auditory object. In a second Varying-Loudness experiment, the intensity level of one speech stream, either the attended or the background, is kept constant, whereas the other is reduced (up to 8 dB). Under this manipulation, the intensity ratio between the attended and background speakers [i.e., the target-to-masker ratio (TMR)] ranges between -8 dB and 8 dB.

To distinguish how different intensity gain control mechanisms would affect the neural representation of each speech stream, we simulate possible decoding outcomes (SI Methods). The MEG activity is simulated by the sum of activity precisely phase locked to each speech stream and interfering stimulus-irrelevant background activity. The strength of the phase-locked activity is normalized by either the strength of whole stimulus, for a global gain control mechanism, or the strength of the encoding auditory object, for an object-based gain control mechanism. The simulated decoding outcomes under different gain control mechanisms are shown in Fig. 3A.

The neural decoding from actual MEG measurements is shown in Fig. 3B. For the decoding of the attended speech, the decoded envelope is significantly more correlated with the envelope of the attended speech [$P < 0.004$, $F(1,71) = 25.8$, attentional focus \times TMR two-way repeated measures ANOVA], and this correlation is not affected by TMR. The result is consistent with the object-based gain control model rather than with the global gain control model. Similarly, the neural decoding of the background speech is also affected by the attentional focus [$P < 0.02$, $F(1,71) = 14.65$, higher correlation with the background speech, two-way ANOVA], without interaction between attention and TMR. Consequently, the neural representation of a speech stream is stable both against the intensity change of that stream and against the intensity change of the other stream, consistent with the hypothesized object-specific gain control (compare the examples shown in Fig. 1C).

Spatial Spectrotemporal Response Function and Neural Sources. The decoding analysis above integrates neural activity, spatiotemporally, to reveal an object-specific neural representation optimally. To characterize the neural code that the decoder extracts information from, we analyze the neural encoding process via the spectrotemporal response function (STRF) for each MEG sensor (26, 27). The linear STRF and the linear decoder are, respectively, the forward and backward models describing the same relationship between the stimulus and neural response. Nevertheless, only the forward STRF model can reveal the timing and spatial information of the neural encoding process.

The STRF functionally describes how the spectrotemporal acoustic features of speech are transformed into cortical responses. It deconvolves the neural activity evoked by the continuous envelope of speech. In this STRF model, the encoding of each speech stream is modeled using the auditory spectrogram (28) of the unmixed speech signal with unit intensity. The STRF shows the neural response to sound features at different acoustic (i.e., carrier) frequencies (as labeled by the y axis). At any given frequency, the horizontal cross-section of the STRF (e.g., Fig. 4A) characterizes the time course of the neural response evoked by a unit power increase of the stimulus at that frequency for one MEG sensor.

Neural Adaptation to the Intensity of Each Speech Stream

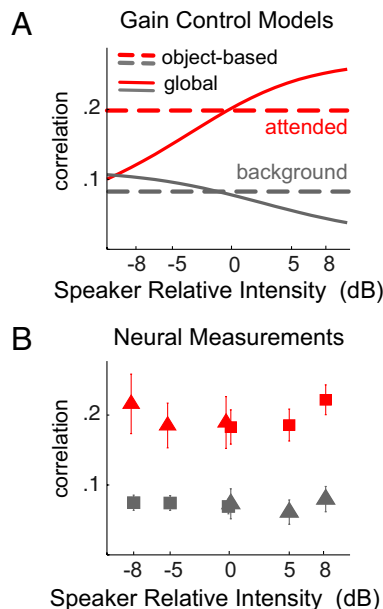


Fig. 3. Decoding the attended speech over a wide range of relative intensity between speakers. (A) Decoding results simulated using different gain control models. The x axis shows the intensity of the attended speaker relative to the intensity of the background speaker. The red and gray curves show the simulated decoding results for the attended and background speakers, respectively. Object-based intensity gain control predicts a speaker intensity invariant neural representation, whereas the global gain control mechanism does not. (B) Neural decoding results in the Varying-Loudness experiment. The cortical representation of the target speaker (red symbols) is insensitive to the relative intensity of the target speaker. The acoustic envelope reconstructed from cortical activity is much more correlated with the attended speech (red symbols) than the background speech (gray symbols). Triangles and squares are results from the two speakers, respectively.

The MEG STRF contains two major response components (Fig. S4): one with latency near 50 ms, here called the $M50_{STRF}$, and the other with latency near 100 ms, here called the $M100_{STRF}$. This indicates that two major neural response components continuously follow the temporal envelope of speech, with delays of 50 ms and 100 ms, respectively. Because an STRF is derived for each MEG sensor, the neural source locations of the $M50_{STRF}$ and $M100_{STRF}$ can be estimated based on the distribution over all sensors of the strength of each component (i.e., the topography of the magnetic fields at each latency). An equivalent current dipole (ECD)-based neural source analysis reveals that the $M50_{STRF}$ and $M100_{STRF}$ responses arise from different neural sources (SI Methods), with the source of $M100_{STRF}$ being 5.5 ± 1.5 mm and 7.1 ± 2.0 mm more posterior in the left and right hemispheres, respectively (Fig. 4B; $P < 0.005$ for both hemispheres, paired t test). The ECD location of the neural source of the $M100_{STRF}$ peak is consistent with that observed for the M100 response to tone pips, localized to the superior temporal gyrus and roughly in the planum temporale (29).

The amplitudes of the $M50_{STRF}$ and $M100_{STRF}$ are further analyzed in the neural source space, based on the STRF at the ECD location of each component. The amplitude of the $M100_{STRF}$ is much stronger for the attended speech than for the background speech [Fig. 4A; $P < 0.007$, $F(1,87) = 11.85$, attentional focus \times hemisphere \times speaker, three-way repeated-measures ANOVA]. The amplitude of the $M50_{STRF}$ is, in contrast, not significantly modulated by either attention or TMR. The latency of the $M50_{STRF}$

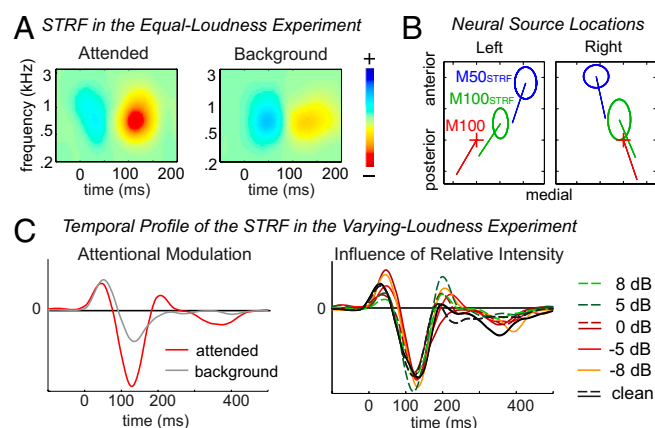


Fig. 4. Cortical encoding of the spectral-temporal features of different speech streams. (A) STRFs for the attended and background speech, at the neural source location of the M100_{STRF}. Attention strongly enhances the response with latency near 100 ms. (B) Neural source locations for the M50_{STRF} and M100_{STRF} in each hemisphere, as estimated by dipole fitting. The location of the neural source of the M50_{STRF} is anterior and medial to that of the M100_{STRF} and M100. The source location for each subject is aligned based on the source of the M100 response to tone pips, shown by the cross. The span of each ellipse is 2 SEM across subjects. The line from each dipole location illustrates the grand averaged orientation of each dipole. Each tick represents 5 mm. (C) Temporal profile of the STRF in the Varying-Loudness experiment for the attended speech. The M100_{STRF} (averaged over TMR) is strongly modulated by attention, whereas the M50_{STRF} is not (Left). Neither response peak is affected by the intensity change of the two speakers (Right).

and M100_{STRF} is also modulated by attention [$P < 0.03$, $F(1,87) > 7$ for both peaks, three-way repeated-measures ANOVA] and is, on average, 11 and 13 ms shorter, respectively, when attended.

The temporal profile of the STRF in the Varying-Loudness experiment is shown in Fig. 4C, which is extracted by applying singular value decomposition to the STRF. The M100_{STRF} is significantly modulated by attention [$P < 0.03$, $F(1,143) = 9.4$, attentional focus \times hemisphere \times speaker \times TMR, four-way repeated-measures ANOVA], whereas the M50_{STRF} is not. Neither response component is affected by TMR (compare the examples shown in Fig. 1C). The invariance of the M50_{STRF} and M100_{STRF} to the intensity of both the attended and background speech streams provides further evidence for the hypothesized object-specific gain control.

Discussion

This study investigates whether a multisource auditory scene, perceptually represented in terms of auditory objects, is neurally represented in terms of auditory objects as well. From subjects selectively listening to one of two spectrotemporally overlapping speech streams, we do observe neural activity selectively synchronized to the speech of a single speaker (as is illustrated in Fig. 1B). Furthermore, in an ecologically valid listening setting, this selective representation of an individual speech stream is both modulated by top-down attention and normalized by the intensity of that sound stream alone (as is illustrated in Fig. 1C). In sum, this meets all the criteria of an object-based representation, [e.g., those specified by Griffiths and Warren (4)]: The observed neural representation is selective to the sound from a single physical source, is minimally affected by competing sound sources, and is insensitive to perceptually unimportant acoustic variations of the stimulus (e.g., changes in intensity).

Temporal Coherence, Attention, and Object-Based Representations. The object-specific representations seen here are precisely synchronized to the temporal envelope of speech. In speech and

natural sounds in general, the temporal envelope synchronizes various acoustic features, including pitch and formant structures. Therefore, they provide important cues for perceptual auditory grouping (30) and are critical for robust speech recognition. For example, major speech segregation cues, such as pitch, are not sufficient for speech recognition, whereas acoustic features necessary for speech recognition (e.g., the spectrotemporal envelope) are not easily distinguishable between speakers. A solution to this dilemma would be to group acoustic features belonging to the same auditory object, both speech segregation and intelligibility-relevant cues, through temporal coherence analysis, and then to process selectively the attended auditory object as a whole (3). In other words, the auditory cortex selects the attended speech stream by amplifying neural activity synchronized to the coherent acoustic variations of speech (i.e., the envelope). This idea is highly consistent with the large-scale synchronized and object-specific activity seen in this study.

At the neuronal mechanistic level, it is plausible that the low-frequency, phase-locked neural activity binds features belonging to the same object by regulating the excitability of neurons, such that a given neural network will be more responsive when processing features from the corresponding auditory object (22). Furthermore, such a rhythmic regulation of neuronal excitability may also contribute to the segmentation of continuous speech into perceptual units (e.g., syllables) (16).

Hierarchical Processing of Auditory Objects in Auditory Cortex. Of the two major neural response components that track the speech envelope, the M100_{STRF} is much more strongly modulated by attention than the M50_{STRF}. These two neural response components track the speech envelope with different latencies and are generated from distinct neural sources. Based on their positions relative to the neural source of the M100 (29), the centers of neuronal current generating the M50_{STRF} and M100_{STRF} are dominantly from Heschl's gyrus and the planum temporale, respectively (31). The latency and source location of the two components demonstrate a hierarchy of auditory processing (32, 33), and the representation of the attended object becomes dominant from shorter to longer latency activity and from core to posterior auditory cortex. Therefore, although auditory object representations may start to emerge as early as primary auditory cortex (7), the top-down attentional modulation of the large-scale, object-based neural representation may emerge most strongly with later processing.

The routing of the neural processing of the attended auditory object into posterior auditory cortex may generally underlie the selection of auditory information when there are competing spectrotemporally complex auditory objects. MEG studies have shown that selectively listening to sound embedded in a complex auditory scene modulates longer latency (~100–250 ms) responses in association auditory cortex but not the shorter latency (~50 ms) steady-state response in core auditory cortex (21, 34, 35), as is also the case for MEG/EEG responses to transient sounds (13, 36–38). PET studies also indicate that the areas posterior to core auditory cortex are more activated when speech is interfered by temporally modulated noise than stationary noise (39, 40), because modulated noise contains speech-like features and requires additional processes of information selection. Furthermore, a recent fMRI study has shown that in a multitalker environment, the planum temporale is increasingly activated when the number of information sources (i.e., speakers) increases (41). Taken together, these results support the idea that posterior auditory cortex plays a major role in the generation of auditory objects (42, 43) and the selection of information based on the listener's interest (25, 44).

Neural Adaptation to the Intensity of Individual Auditory Object. The recognition of speech relies on the shape of its spectrotemporal modulations and not its mean intensity. This study demonstrates that cortical activity is precisely phase locked to the temporal

modulations but insensitive to the mean intensity of the speech streams, and therefore is effectively encoding only the shape of the modulations. Intensity gain control has been demonstrated in multiple stages of the auditory system (15) and constitutes an auditory example of neural normalization, which has been suggested as a canonical neural computation (45). For example, in the cochlear nucleus, neurons encode the shape of the spectral modulation of speech (e.g., a vowel) invariantly to its mean intensity (46).

Critically different from these previous studies, however, the encoding of temporal modulations seen here is invariant to the intensity of each speech stream rather than the overall intensity of the mixture. In the Varying-Loudness experiment, the intensity of one speaker changes, whereas the other is kept constant. Maintaining a stable representation despite the altered speech requires the observed neural adaptation to the sound intensity of the specific speech stream. The stable representation of the constant speaker, in contrast, requires the observed lack of adaptation to the overall intensity of the sound mixture, which covaries with the intensity of the altered speech. These both contrast with the simpler mechanism of global intensity gain control, which would require the neural representation of both speech streams to be modulated in the same way based on the overall intensity of the acoustic stimulus. Therefore, the data strongly suggest the existence of an object-specific intensity gain control, which normalizes the strength of neural activity based on the intensity of individual auditory objects.

In sum, this study demonstrates the key signatures of an object-specific neural representation arising from the analysis of a complex auditory scene. Such an object-specific neural representation is phase locked to the slow rhythms (<10 Hz) of the encoded auditory object, and it adapts to the intensity of that object alone. Under the modulation of top-down attention, the auditory response in posterior auditory cortex (latency near 100 ms) dominantly represents the attended speech, even if the competing speech stream is physically more intense. This object-specific auditory representation provides a bridge between feature-based, precisely phase-locked sensory responses and interference-resilient cognitive processing and recognition of auditory objects.

Methods

Subjects, Stimuli, and Procedure. Twenty normal-hearing, right-handed, young adult native speakers of American English (between 18 and 26 y old) participated in the experiment in total. Eleven (5 female) participated in the Equal-Loudness experiment, 6 (3 female) participated in the Varying-Loudness experiment, and 3 (2 female) participated in the Same-Sex experiment.

The stimuli contain three segments from the book *A Child's History of England* by Charles Dickens, narrated by three different readers (2 female).

Each speech mixture was constructed by mixing the two speakers digitally in a single channel and was then divided into sections with a duration of 1 min. All stimuli were delivered identically to both ears using tube phones plugged into the ear canals. In the Equal-Loudness and Varying-Loudness experiments, the two simultaneous speakers were of opposite sex (mean pitch separation of 5.5 semitones). In the Equal-Loudness experiment, the two speakers were mixed, with equal rms values of sound amplitude. The subjects were instructed to focus on only one speaker and to answer questions related to the comprehension of the passage they focused on. In the Varying-Loudness experiment, the stimuli were mixtures of the same two speakers. The intensity of one speaker was fixed, whereas the intensity of the other speaker was either the same or was -5 or -8 dB weaker. The Same-Sex experiment used mixtures of speakers of the same sex. Details of the three experiments are included in *SI Methods*.

Data Recording and Processing. The neuromagnetic signal was recorded using a 157-channel, whole-head MEG system (KIT) in a magnetically shielded room, with a sampling rate of 1 kHz. The ongoing neural response (excluding the first second) during each 1-min stimulus was filtered between 1 and 8 Hz (21), downsampled to 40 Hz, and was then used for the decoding and STRF analysis.

The temporal envelope of each speaker in the stimulus was decoded by linearly integrating the spatial-temporal brain activity. The decoder was optimized so that the decoded envelope was maximally correlated with the speaker to decode and minimally correlated with the other speaker (*SI Methods*). All correlations in this study were measured by the absolute value of the Pearson's correlation coefficient. The decoder optimized in this way was a discriminative model that reconstructed an envelope similar to one speaker but distinct from the other, and it was therefore used to explore the neural code unique to each speaker. The decoder was applied to individual trials, and the percentage of trials where decoding was successful (decoded envelope being more correlated with the intended speaker) is always reported as the grand average. This decoding approach effectively explores both the spatial and temporal information in MEG and avoids the sometimes ill-posed problem of estimating the neural source locations.

The STRF was used to model how speech features are encoded in the MEG response (21), in contrast to how the decoders transform MEG activity (backward) to speech features. A single STRF transforms the spectrotemporal features of speech to a single-response waveform. Therefore, because of the multichannel nature of MEG, a complete forward model is described as a 3D spatial-STRF model (MEG sensor position \times frequency \times time). The MEG response was modeled as the sum of the responses to individual speakers, with each having unit intensity (more details about STRF estimation are provided in *SI Methods*). The MEG data were averaged over trials in the STRF analysis for each stimulus and attentional condition.

ACKNOWLEDGMENTS. We thank David Poeppel, Mary Howard, Monita Chatterjee, Stephen David, and Shihab Shamma for comments and Max Ehrmann for technical assistance. This research was supported by the National Institute of Deafness and Other Communication Disorders Grant R01-DC-008342.

- Bregman AS (1990) *Auditory Scene Analysis: The Perceptual Organization of Sound* (MIT Press, Cambridge, MA).
- Shinn-Cunningham BG (2008) Object-based auditory and visual attention. *Trends Cogn Sci* 12(5):182–186.
- Shamma SA, Elhilali M, Micheyl C (2011) Temporal coherence and attention in auditory scene analysis. *Trends Neurosci* 34(3):114–123.
- Griffiths TD, Warren JD (2004) What is an auditory object? *Nat Rev Neurosci* 5: 887–892.
- Fishman YI, Steinschneider M (2010) Formation of auditory streams. *The Oxford Handbook of Auditory Science: The Auditory Brain*, eds Rees A, Palmer A (Oxford Univ Press, New York), Vol 2, pp 215–245.
- Snyder JS, Gregg MK, Weintraub DM, Alain C (2012) Attention, awareness, and the perception of auditory scenes. *Front Psychol* 3:15.
- Nelken I, Bar-Yosef O (2008) Neurons and objects: The case of auditory cortex. *Front Neurosci* 2(1):107–113.
- Wang D, Brown GJ (2006) *Computational Auditory Scene Analysis: Principles, Algorithms, and Applications* (Wiley-IEEE, New York).
- Cherry EC (1953) Some experiments on the recognition of speech, with one and with two ears. *J Acoust Soc Am* 25:975–979.
- Brungart DS (2001) Informational and energetic masking effects in the perception of two simultaneous talkers. *J Acoust Soc Am* 109:1101–1109.
- Elhilali M, Ma L, Micheyl C, Oxenham AJ, Shamma SA (2009) Temporal coherence in the perceptual organization and cortical representation of auditory scenes. *Neuron* 61:317–329.
- Fritz J, Shamma S, Elhilali M, Klein D (2003) Rapid task-related plasticity of spectrotemporal receptive fields in primary auditory cortex. *Nat Neurosci* 6:1216–1223.
- Hillyard SA, Hink RF, Schwent VL, Picton TW (1973) Electrical signs of selective attention in the human brain. *Science* 182:177–180.
- Xiang J, Simon J, Elhilali M (2010) Competing streams at the cocktail party: Exploring the mechanisms of attention and temporal integration. *J Neurosci* 30:12084–12093.
- Robinson BL, McAlpine D (2009) Gain control mechanisms in the auditory pathway. *Curr Opin Neurobiol* 19:402–407.
- Luo H, Poeppel D (2007) Phase patterns of neuronal responses reliably discriminate speech in human auditory cortex. *Neuron* 54:1001–1010.
- Abrams DA, Nicol T, Zecker S, Kraus N (2008) Right-hemisphere auditory cortex is dominant for coding syllable patterns in speech. *J Neurosci* 28:3958–3965.
- Ahissar E, et al. (2001) Speech comprehension is correlated with temporal response patterns recorded from auditory cortex. *Proc Natl Acad Sci USA* 98:13367–13372.
- Koskinen M, et al. (2012) Identifying fragments of natural speech from the listener's MEG signals. *Hum Brain Mapp*, 10.1002/hbm.22004.
- Kerlin JR, Shahin AJ, Miller LM (2010) Attentional gain control of ongoing cortical speech representations in a "cocktail party". *J Neurosci* 30:620–628.
- Ding N, Simon JZ (2012) Neural coding of continuous speech in auditory cortex during monaural and dichotic listening. *J Neurophysiol* 107:78–89.
- Koschroder CE, Lakatos P (2009) Low-frequency neuronal oscillations as instruments of sensory selection. *Trends Neurosci* 32:9–18.
- Kay KN, Naselaris T, Prenger RJ, Gallant JL (2008) Identifying natural images from human brain activity. *Nature* 452:352–355.

24. Pasley BN, et al. (2012) Reconstructing speech from human auditory cortex. *PLoS Biol* 10:e1001251.
25. Mesgarani N, Chang EF (2012) Selective cortical representation of attended speaker in multi-talker speech perception. *Nature* 485:233–236.
26. Depireux DA, Simon JZ, Klein DJ, Shamma SA (2001) Spectro-temporal response field characterization with dynamic ripples in ferret primary auditory cortex. *J Neurophysiol* 85:1230–1234.
27. deCharms RC, Blake DT, Merzenich MM (1998) Optimizing sound features for cortical neurons. *Science* 280:1439–1443.
28. Yang X, Wang K, Shamma SA (1992) Auditory representations of acoustic signals. *IEEE Trans Inf Theory* 38:824–839.
29. Lütkenhöner B, Steinsträter O (1998) High-precision neuromagnetic study of the functional organization of the human auditory cortex. *Audiol Neurootol* 3(2–3):191–213.
30. Sheft S (2007) Envelope processing and sound-source perception. *Auditory Perception of Sound Sources*, eds Yost WA, Popper AN, Fay RR (Springer, New York), pp 233–279.
31. Steinschneider M, Liégeois-Chauvel C, Brugge JF (2011) Auditory evoked potentials and their utility in the assessment of complex sound processing. *The Auditory Cortex*, eds Winer JA, Schreiner CE (Springer, New York), pp 535–559.
32. Hickok G, Poeppel D (2007) The cortical organization of speech processing. *Nat Rev Neurosci* 8:393–402.
33. Rauschecker JP, Scott SK (2009) Maps and streams in the auditory cortex: Nonhuman primates illuminate human speech processing. *Nat Neurosci* 12:718–724.
34. Gutschalk A, Michéyl C, Oxenham AJ (2008) Neural correlates of auditory perceptual awareness under informational masking. *PLoS Biol* 6:e138.
35. Okamoto H, Stracke H, Bermudez P, Pantev C (2011) Sound processing hierarchy within human auditory cortex. *J Cogn Neurosci* 23:1855–1863.
36. Woods DL, Hillyard SA, Hansen JC (1984) Event-related brain potentials reveal similar attentional mechanisms during selective listening and shadowing. *J Exp Psychol Hum Percept Perform* 10:761–777.
37. Woldorff MG, et al. (1993) Modulation of early sensory processing in human auditory cortex during auditory selective attention. *Proc Natl Acad Sci USA* 90:8722–8726.
38. Ahveninen J, et al. (2011) Attention-driven auditory cortex short-term plasticity helps segregate relevant sounds from noise. *Proc Natl Acad Sci USA* 108:4182–4187.
39. Scott SK, Rosen S, Beaman CP, Davis JP, Wise RJS (2009) The neural processing of masked speech: Evidence for different mechanisms in the left and right temporal lobes. *J Acoust Soc Am* 125:1737–1743.
40. Scott SK, Rosen S, Wickham L, Wise RJS (2004) A positron emission tomography study of the neural basis of informational and energetic masking effects in speech perception. *J Acoust Soc Am* 115:813–821.
41. Smith KR, Hsieh I-H, Saberi K, Hickok G (2010) Auditory spatial and object processing in the human planum temporale: No evidence for selectivity. *J Cogn Neurosci* 22:632–639.
42. Griffiths TD, Warren JD (2002) The planum temporale as a computational hub. *Trends Neurosci* 25:348–353.
43. Zatorre RJ, Bouffard M, Ahad P, Belin P (2002) Where is 'where' in the human auditory cortex? *Nat Neurosci* 5:905–909.
44. Zion Golumbic EM, Poeppel D, Schroeder CE (2012) Temporal context in speech processing and attentional stream selection: A behavioral and neural perspective. *Brain & Language*, 10.1016/j.bandl.2011.12.010.
45. Garandini M, Heeger DJ (2012) Normalization as a canonical neural computation. *Nat Rev Neurosci* 13:51–62.
46. Young ED (2008) Neural representation of spectral and temporal information in speech. *Philos Trans R Soc Lond B Biol Sci* 363:923–945.