Temporal coherence and attention in auditory scene analysis

Shihab A. Shamma[1], Mounya Elhilali[2] and Christophe Micheyl[3]

[1] Department of Electrical and Computer Engineering and Institute for Systems Research, University of Maryland, College Park, MD 20742, USA
[2] Electrical Engineering Department, Johns Hopkins University, Baltimore, MD, USA
[3] Department of Psychology, University of Minnesota, Minneapolis, MN 55455, USA

Humans and other animals can attend to one of multiple sounds and follow it selectively over time. The neural underpinnings of this perceptual feat remain mysterious. Some studies have concluded that sounds are heard as separate streams when they activate well-separated populations of central auditory neurons, and that this process is largely pre-attentive. Here, we argue instead that stream formation depends primarily on temporal coherence between responses that encode various features of a sound source. Furthermore, we postulate that only when attention is directed towards a particular feature (e.g. pitch) do all other temporally coherent features of that source (e.g. timbre and location) become bound together as a stream that is segregated from the incoherent features of other sources.

The auditory scene analysis problem

Humans and other animals routinely detect, identify and track sounds coming from a particular source (e.g. someone's voice, a conspecific call) among sounds emanating from other sources (e.g. other voices, heterospecific calls, ambient music and street traffic) (Figure 1). The apparent ease with which they determine which components and attributes in a sound mixture arise from the same source belies the complexity of the underlying biological processes. By analogy with the scene segmentation problem in vision, this is referred to as the auditory scene analysis problem (Glossary) [1] or more colloquially as the cocktail party problem [2–4]. Understanding how the brain solves this problem is a fundamental challenge facing auditory scientists, because it will shed light on difficulties experienced by the hearing-impaired in multi-source environments [5] and contribute to more effective front-end processors for auditory prostheses and automatic speech recognition [6].

Recent studies have inspired numerous hypotheses and models concerning the neural underpinnings of perceptual organization in the central auditory system, and especially the auditory cortex [7–17]. One prominent hypothesis that underlies most investigations is that sound elements segregate into separate streams whenever they activate well-separated populations of auditory neurons that are selective to frequency or any other sound attributes that have been shown to support stream segregation [18–27]. We refer to this hypothesis as the "population separation"

Corresponding author: Shamma, S.A. (sas@umd.edu).

Glossary

Auditory scene analysis: processes by which sequential and concurrent acoustic events are analyzed and organized into auditory streams.

Auditory stream: series of sounds perceived by the listener as a coherent entity and, as such, can be selectively attended to among other sounds. The word 'stream' emphasizes the fact that sounds usually unfold over time. Although sounds coming from different physical sound sources typically form separate streams, this is not always the case. For example, a choir singing in unison consists of multiple sources heard as a single stream, whereas an audio speaker is a single physical source that usually creates multiple streams. Several objective criteria exist by which it can be determined if a series of sounds is perceived as a stream.

Coherence: temporal coherence between two channels is defined here in the following specific sense: it denotes the average similarity or coincidence of their responses measured over a given time-window. It is computed as the running cross-correlation coefficient at zero lag between the channel responses integrated over relatively long time windows (50–500 ms). Therefore, channels with similar activity over this time interval are highly coherent (a correlation coefficient near 1), such as the synchronous tone pairs in Figure Ib in Box 2. Anti-coherence therefore refers to the relationship between opposite or inverted responses (cross-correlation coefficient near –1) such as the alternating tones in Figure Ia in Box 2.

Complex tone: periodic sound that contains multiple frequencies.

Frequency: number of cycles per unit of time. It is usually expressed in cycles/s, or Hertz (Hz).

Fundamental frequency (F_0): inverse of the period of a harmonic complex tone. It is the highest frequency of which all other frequency components in a harmonic complex tone are integer multiples.

Harmonic: spectral component in a harmonic complex tone.

Noise: strictly speaking, an aperiodic sound. More broadly, it is any undesirable sound.

Pure tone: tone that consists of a single frequency.

Sound token: defined in this article as a burst of sound that rapidly evokes a percept. A token can be as simple as a pure tone, a harmonic complex or a transient acoustic event such as a click, or as complex as a vowel, a syllable or a musical chord. It usually has one or more of the common attributes of sound such as pitch, loudness, location or timbre.

Spectrogram: visual representation of the spectrum of a sound as a function of time. Time is usually shown along the abscissa and frequency along the ordinate, and sound energy (or amplitude) at each time–frequency point is indicated using color or shades of gray.

Spectrum: representation of the frequency content of a signal. It is usually obtained using a Fourier transform and shows the amplitude and/or phase of the different frequency components in a signal.

Streaming: process of forming segregated percepts of auditory sources. In the literature on hearing research, the terms 'streaming' and 'stream formations' are most often reserved to describe sequential grouping or organization of sound segments or tokens over time. In this article, we exclusively use 'streaming' in this sense. There are both subjective and objective criteria to determine whether a stream is perceived or not, although there is no universal agreement on these.

Synchronous stimuli: stimuli that always have a common onset in time when they co-occur.

Tone: periodic sound.

Tone or token sequence: sequence of sound elements that occur at relatively slow rates (<20 Hz). Examples are experimental sequences of pure tones, notes of a musical melody and syllables in running speech.

Tuning curve: usually refers to the selectivity of auditory neurons to acoustic frequencies, often measured using pure tones. It is analogous to the receptive field of a visual neuron.

  doi:10.1016/j.tins.2010.11.002 Trends in Neurosciences, March 2011, Vol. 34, No. 3

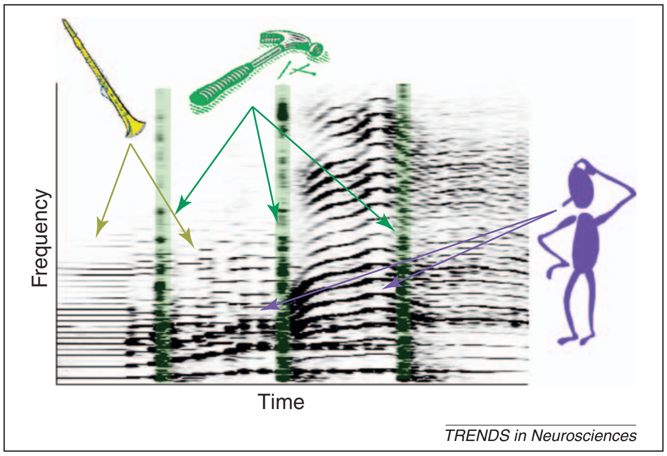

Figure 1. Spectrogram of a complex scene with multiple objects. The figure shows a time–frequency analysis of an acoustic recording of a scene consisting of flute, a human voice and a hammer. The hammer hits are immediately visible as repetitive and transient broadband strips of energy spanning all frequencies. Both the flute and the human voice contain a rich harmonic structure that changes over time. The human voice reveals clear pitch variations and formant transitions, shown as time-course changes in both the pitch and formant locations. Note that the flute and speech give rise to clearly distinct acoustic events that are uncorrelated in time.

hypothesis. Another influential hypothesis is that streams are formed automatically or pre-attentively, in or below the primary auditory cortex [27–29].

In this Opinion piece, we point out shortcomings of these two hypotheses and propose an alternative to each within an overall framework for understanding auditory scene analysis and its neural basis. On the basis of a combination of neurophysiological data, psychophysical observations and computational studies, we argue that the formation of auditory streams depends fundamentally on the temporal coherence of responses of neural populations selective to various sound attributes (e.g. frequency, pitch, timbre, spatial location) in the auditory cortex. In addition, we suggest that attention plays a key role in stream formation, because it biases the auditory system toward a particular grouping or binding of sound-source attributes, depending on the listener's current behavioral or perceptual goals.

Temporal coherence in auditory scene analysis

Problems inherent to auditory scene analysis are similar to those found in visual scene analysis. However, there are a few notable unique aspects. In particular, whereas natural and artificial visual scenes often contain a large proportion of static or slow-moving elements, auditory scenes are essentially dynamic, containing many fast-changing, relatively brief acoustic events (referred to as tokens in Box 1) [30,31]. Therefore, an essential aspect of auditory scene analysis is the linking over time, or streaming, of tokens produced by the same sound source, while simultaneously separating them from others produced by other sources. We explain here why we are of the opinion that the key first step to this process of streaming is the temporal coherence of the tokens within a stream (or equivalently, their incoherence across streams) and not the widely assumed population-separation hypothesis. The Glossary gives a precise definition of coherence.

The population-separation theory of auditory streaming

Over the last decade, numerous psychophysical and neurophysiological studies of auditory streaming have concluded that the perceptual organization of sounds into streams is determined by the spatial overlap between responsive neural populations in the peripheral and/or central auditory system. Simply stated, under this hypothesis, sounds that activate distinct (or weakly overlapping) neural populations are heard as separate streams. The tonotopic axis is a major organizational principle throughout the auditory system, so most models based on this population-separation theory of auditory streaming have focused on the frequency dimension [18–27] and have successfully accounted for many important aspects of the perceptual organization of simple tone sequences (Box 2). This hypothesis has been extended to account for stream formation based on other features, such as spectral shape (timbre), periodicity (pitch) and spatial location [32–35]. This requires a multi-feature analysis, which presumably arises from neural responses in the central auditory system that are selective to attributes other than frequency, such as, to various spectral and temporal characteristics of sounds [36–44], sound-source location [45,46] and pitch [47].

However, the spatial separation of neural responses cannot account for the observed influence of the relative timing of sounds on streaming percepts. For example, the population-separation hypothesis predicts that both alternating and synchronous tones (see Figure Ia,b in Box 2) that differ widely in frequency should be heard as separate streams. This prediction is contradicted by psychophysical and neurophysiological data [48] demonstrating that sequences of tones that are separated by an octave or more are still heard as a single stream if the tones are synchronous or, more precisely, fully coherent in time (Box 2 and Glossary). Numerous other psychoacoustical findings indicate that coherence strongly promotes perceptual grouping [49]. To account for these findings, it is necessary to consider the relative timing of the neural responses and, more specifically, their temporal coherence.

Temporal coherence and auditory streaming

Combining multi-feature representations and temporal-coherence analysis leads to a general and flexible framework, which can explain the formation of auditory streams for a wide range of stimuli. This framework is illustrated in Figure 2. It begins with frequency analysis in the cochlea, followed by extraction of a wide variety of spectral and temporal features, including a multi-resolution representation of spectral shapes, harmonicity, temporal periodicity, and inter-aural time and level differences. Some of these features (e.g. harmonicity and inter-aural differences) are related directly to perceptual attributes (e.g. pitch and location).

We next postulate the existence of a temporal-coherence analysis stage that computes correlations among the outputs of the different feature-selective neurons. The correlations are computed over relatively long time windows, ranging in duration between 50 and 500 ms. This range is consistent with the slow dynamics of stimulus-induced fluctuations in spike rate in the auditory cortex (<20 Hz) [39,50]. It is also consistent with the sound-presentation rates over which the formation of streams usually occurs, as well as with the rates of temporal-envelope fluctuations

Box 1. Principles of stream formation and perception

Percepts and processes underlying auditory perceptual organization can be conceptually divided into two categories: instantaneous (sometimes referred to as simultaneous) percepts and sequential processes (or stream formation) [1].

An instantaneous percept refers to that of a sound epoch or token that arises rapidly after its onset and continues throughout its duration. Natural sounds are dynamic and can be conceptualized as sequences of tokens. Each token has associated perceptual attributes (pitch, loudness, timbre and location) that reflect its frequency components and their relationships, for example, whether harmonically related or what their relative amplitudes are. Sound tokens encountered in our environment are endowed with richly varied and complex percepts (some are illustrated in Figure Ia). For instance, a sound token can consist of one or two tones, a perceptually fused harmonic complex, or an inharmonic complex with a 'fractured' multi-tone percept. Tokens can also have attributes other than frequency, such as the pitch of musical notes or a whole chord (Figure Ib), and the perceived location of a point source (Figure Ic). Finally, tokens can have complex attributes such as the timbre of one or more simultaneous vowels (Figure Id) or a highly diffuse sound in a large reverberant hall or that of a large choir singing in unison. All these percepts are extracted relatively early and rapidly in the auditory system by basic neural structures (within a few tens of milliseconds; hence the term instantaneous percepts) and there is a large body of psychoacoustic and neurophysiological results that relate the acoustic parameters of a complex sound to these attributes [92,93].

Sequential organization specifically refers to the sorting of interleaved sound tokens arriving from a mixture of sources into streams that can be selectively attended to and tracked over time. Examples of auditory streams are: two independent interleaved melodies played by a violin and a piano; the melody of a piano within an orchestra; and someone's voice in a crowd. Each stream can be thought of as a sequence of tokens that the listener can attend to and perceive as the target stream or melody. To do so, the listener must distinguish the attributes of the different tokens (instantaneous percepts) and organize them into separate streams (sequential process).

This process has a few basic properties that are addressed in this article. One is that tokens in different streams must be sufficiently incoherent in time. They must also be perceptually distinct enough to reflect the different acoustic characteristics of their sources. Finally, the tokens should remain relatively stable perceptually over time within a stream. For instance, the timbre and pitch of sounds within a stream should not change drastically and quickly, or these sounds will fail to form a coherent auditory stream. This is essentially identical to the continuity principle, which is often invoked as a key ingredient in the learning of object invariance along various dimensions [94,95]. Sequences of tokens unfold relatively slowly over time (>50 ms), so sequential organization (or formation of a stream) is a slow process that can take several seconds to complete, especially when the tokens to be segregated are perceptually close. Finally, it is argued that unlike instantaneous processes that have been demonstrated even in anesthetized animals [96], stream formation engages cognitive processes, such as attention and expectations [19].

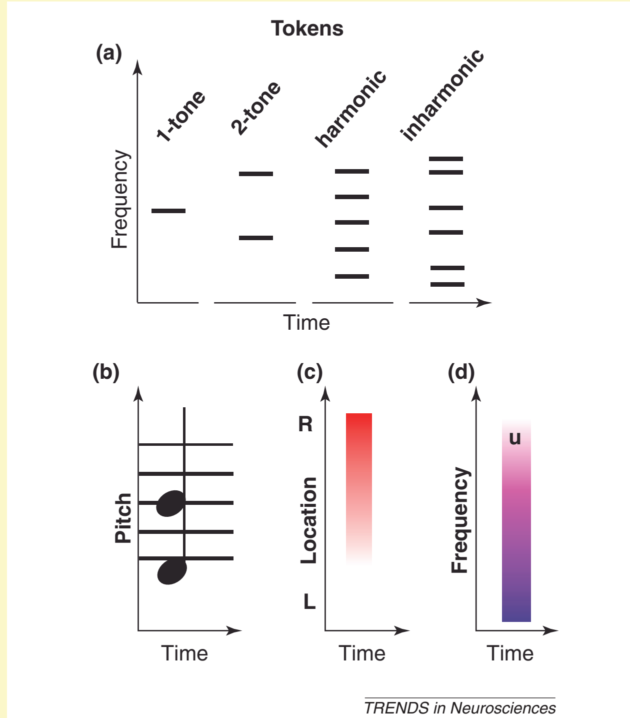

Figure I. Principles and examples of auditory streaming: instantaneous percepts (tokens). Examples of acoustic tokens with different attributes. **(a)** Spectra of tonal tokens: single tone, two tones, harmonic complex, and an inharmonic complex. Tokens are relatively brief and their constituents have a common onset. **(b–d)** Complex tokens. Sound tokens can have various attributes such as (b) the pitch of musical notes or chords, (c) location along the azimuth, and (d) the timbre of a vowel with a specific spectral shape. In each of these panels, the feature value is represented by the pattern of activation along the ordinate. For example, each note in (b) represents the place of activation along the low-to-high pitch values; the activation pattern in (c) has a peak on the right (R) along the left-to-right (L to R) ordinate; the vowel in (d) is represented by its spectral shape along the frequency axis. All these features occur over a brief time interval.

typically encountered in speech (syllabic rate) and music (tempo). Although these coherence computations can take place automatically (pre-attentively), we postulate that active listening or attention is necessary to exploit the results and bind coherent channels into a perceptual stream, while segregating them from the remaining incoherent channels. Clearly, more complex patterns of streaming would arise if channels were partially coherent, and hence might belong to none or to more than one stream simultaneously [51].

Temporal coherence solves the auditory binding problem

The principle of grouping by temporal coherence provides an elegant solution to the auditory binding problem, namely the problem of associating different sound features (loudness, pitch, timbre and spatial location) with the correct (i.e. corresponding) sound source, and of linking these features together to produce a unified percept, while keeping them separate from the features of other sources. This is because features of a particular source will, in general, be present whenever the source is active, and absent when it is silent. Furthermore, different sound sources (with all of their associated features) will rarely fluctuate in strength at exactly the same times.

Whether and how temporal-coherence computations are implementedneurally remains unclear. These computations may involve neurons whose responses depend strongly on the temporal coherence of their input spike trains, or combination-sensitive neurons that respond selectively to particular combinations of inputs, such as neurons that respond strongly to two simultaneously presented tones, even though they respond weakly to either tone alone [52–57]. Such neurons would provide evidence for a substrate capable of integrating temporally coherent responses across spatially distributed neural populations.

The hypothesis that temporal coherence across neural populations solves the binding problem is not unique to the

Box 2. Coherence and attention in streaming: examples with tone sequences

The simplest stimuli that illustrate the role of temporal coherence and attention in stream formation are the much-studied sequences of pure tones [1]. To start, tones that alternate repeatedly between two far-apart frequencies are usually heard as two streams (Figure Ia). This, we claim, is not because the responses are widely spaced on the tonotopic axis, but rather because they induce incoherent responses (i.e. as illustrated in the separate auditory channels of A and B). The evidence for this statement is that when channel responses in A and B are made temporally coherent, for example, when the tones are synchronized (Figure Ib), the tones are heard as one stream despite their large separation in frequency [48].

If the alternating tones are brought closer together as in Figure 1c, the A and B channels become highly overlapped and hence are driven by both tones and carry coherent responses. They are therefore heard as a single perceptual stream that oscillates in frequency regardless of tone presentation rates [1,79]. In Figure 1d, two synchronous tone sequences of fixed and variable frequencies are heard as two streams because the coherence between the A and B channels is weak.

This example is generalized by the stimulus in Figure 1e where a target tone sequence is embedded in masker tones. The target sequence evokes responses in channel A that are incoherent with channel B (and other channels too), and hence can be heard streamed from the complex. Finally, Figure 1f illustrates the ''capture and streaming'' of a simultaneous tone pair [1]. Here the tone pair is normally heard as a single complex sound when presented in isolation. However, a preceding sequence of low tones in channel B decorrelates the responses in channels A and B causing them to perceptually segregate. Consequently the tone sequence 'captures' the low tone, separating it from the high tone, which is now heard clearly against the background of the low tones.

Although temporal coherence computations could occur without significant cognitive control (e.g. similar to cochlear frequency analysis), we propose that attentive listening is necessary for subsequent exploitation of the results to bind coherent attributes or group channels into different streams. An experimental finding that is consistent with this claim is that when one attends to the incoherent responses of the alternating tones illustrated in Figure Ia, one initially hears a unified percept that only gradually gives way to two streams (known as the build-up) [1,79–81], which suggests that the incoherence is ignored prior to the onset of attention.

To explain further the relationship between coherence, binding and streaming within the context of the model, consider the percepts evoked by the alternating and synchronous tones (Figure Ia,b) when presented in separate ears (e.g. A, right; B, left ear). Each tone now has two coherent attributes, pitch and location, and so by attending to one (e.g. pitch) it binds perceptually with the other (location) to form one stream. The alternating tones (e.g. as illustrated in Figure Ia) are incoherent, and hence their attributes are also incoherent and will stream apart, making it easy to distinguish and associate each tone with its pitch and ear-of-entry. By contrast, the synchronous tones (e.g. as illustrated in Figure Ib) and all their attributes are coherent, and hence all will bind together into one stream. In this case, we predict that listeners would find it difficult to determine which tone is in which ear even if the frequencies are well separated.

We should emphasize that synchronicity and coherence are different notions. The first is an instantaneous property, whereas the latter is an average measure (a windowed cross-correlation). We propose that only coherence is key to streaming. To illustrate this distinction, consider the closely spaced alternating tones of Figure Ic. These tones are asynchronous, but their channels A and B are sufficiently overlapped that they carry coherent responses and hence are heard as one stream. By contrast, we predict that the synchronous tone sequences illustrated in Figure Id stream apart because the A and B channels have incoherent responses. Exactly the same argument applies to the so-called informational masking (IM) stimulus in Figure Ie [102,104]. Finally, the two synchronous tones in Figure 1f are nevertheless perceived in separate streams [97] because the A and B channels carry incoherent responses.

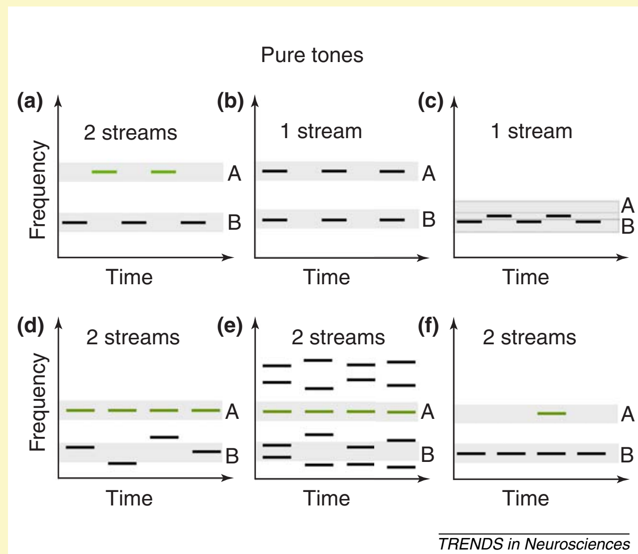

Figure I. Streaming with pure tones. Examples of sequential organization of pure-tone sequences. **(a)** Two alternating tones of widely separated frequencies are usually perceived as two separate streams. The green color indicates a separate stream. The shaded regions denote two hypothetical neural auditory channels activated by the tones. The A and B channels are incoherent. **(b)** Two synchronous sequences are perceived as a single stream because the A and B channels are coherent. **(c)** Alternating (asynchronous) tones of close frequencies are usually heard as a single perceptual stream. **(d)** Two synchronous tone sequences of fixed and variable frequencies. **(e)** 'Release from informational masking' stimulus. **(f)** Capture and streaming of a simultaneous tone pair.

auditory modality [58,59]. Temporal coherence across different sensory modalities might support cross-modal binding (as in lip-reading, for which both visual and auditory inputs are used). Unfortunately, relatively little is known about interactions between auditory and visual or somatosensory inputs in auditory streaming (see [60] for an exception).

It is important to point out that variants of the principle of grouping by temporal coherence have previously been applied to sensory perception problems [61,62], including models of auditory scene analysis [63,64]. However, our model differs from previous ones in that temporal coherence is based on the relatively slow-varying stimulus features (<20 Hz) that induce phase-locked cortical responses. Importantly, in this model, temporal coherence does not rely on intrinsic (i.e. not stimulus-driven) oscillatory activity in the nervous system (e.g. local field-potential oscillations in the gamma frequency range [65]).

Two recent computational studies have implemented some of these ideas to successfully simulate the formation of auditory streams for a wide variety of stimuli [48,66], including simple sequences of regularly repeating tones, stochastic tone sequences and concurrent speech sounds.

The role of attention in auditory stream formation

Is streaming a pre-attentive process?

A widely held view by which has emerged from electrophysiological studies in humans [28,29,67–70], is that auditory streams are formed pre-attentively in the auditory system, much like the extraction of low-level features in early precortical stages. Depending on the listener's intentions and guided by representations of previously encountered audi-

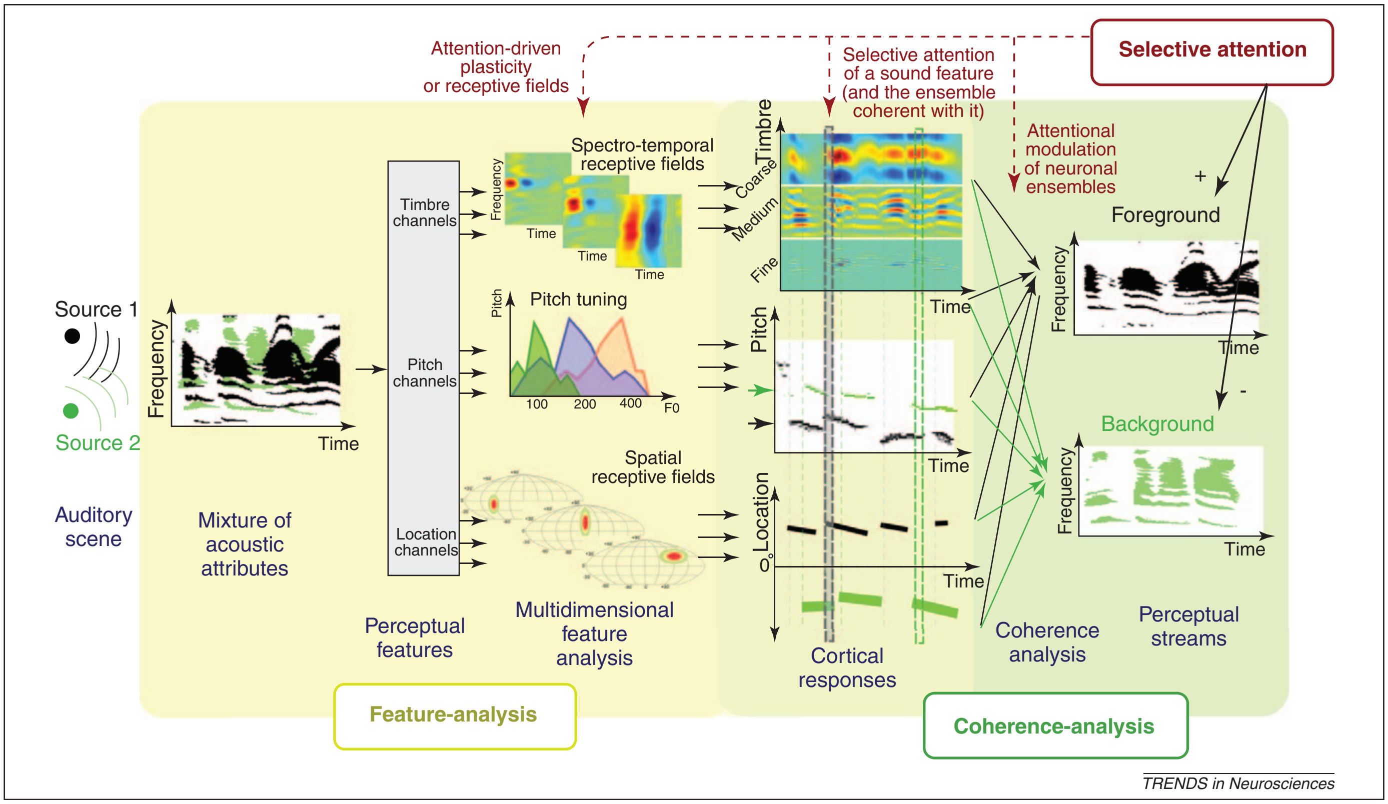

Figure 2. Schematic of the proposed model of auditory stream formation. From left to right: Multiple sound sources constitute an auditory scene, which is initially analyzed through a feature analysis stage. This stage consists of a cochlear frequency analysis followed by arrays of feature-selective neurons that create a multidimensional representation along different feature axes. The figure depicts timbre, pitch and spatial location channels. Note that for computational convenience and illustration purposes, these feature maps are shown with ordered axes when in fact such orderly representations are neither known nor are essential for the model. The outcome of this analysis is a rich set of cortical responses that explicitly represent the different sound features, as well as their timing relationships. The second stage of the model performs coherence analysis by correlating the temporal outputs of the different feature-selective neurons and arranging them based on their degree of coherence, hence giving rise to distinct perceptual streams. Complementing this feed-forward bottom-up view are top-down processes of selective attention that operate by modulating the selectivity of cortical neurons. This feature-based selective attention translates onto object-based attentional mechanisms by virtue of the fact that selected features are coherent with other features that are part of the same stream.

tory objects (or streams) that are now stored in memory, attention would simply serve to enhance the perception of a particular stream in the auditory scene, while suppressing others [71–75]. Thus, in this view, attention is involved in stream selection rather than in stream formation [76], which remains essentially a pre-attentive process. This is reminiscent of similar views proposed earlier for the visual modality [77,78].

We refer to this as the object-based attention theory of auditory scene analysis. An important challenge for this hypothesis is that complex auditory scenes can usually be organized perceptually in many different ways. For instance, when listening to an orchestra, one can listen to the ensemble, to a particular instrument (e.g. the trumpet or flute), or to a group of instruments (e.g. the strings or the woodwinds). In the first case, the orchestra will be heard as a single stream; in the other cases, different streams will be heard, corresponding to individual instruments or to groups of instruments. It seems unlikely that the brain would waste resources representing large numbers of potential decompositions of auditory scenes into streams before (and independently of) attentional selection.

Attention influences stream formation

The hypothesis that attention can only operate on neural representations of already formed auditory objects is contradicted by psychophysical findings. First, when listening to sound sequences such as those illustrated in Figure Ib in Box 2, the frequency separation required to induce a percept of two separate streams is usually much smaller if the listener is actively trying to hear out the high-pitch tones than if she is listening less selectively [79]. This finding indicates that active engagement in the task, as well as the implicit attention brought to bear during the task's performance, does not merely serve to select one among several already formed streams; instead, attention can influence the stream-formation process itself [80,81].

At the neural level, attention can influence auditory stream formation in at least two important ways. First, it can enhance responses to different features, and thus modify the neural representation and ultimately the perceptual saliency of these features. During the last decade, several studies have demonstrated rapid task- and attention-dependent changes in the spectrotemporal receptive fields of the auditory cortex [82]. Preliminary results of a study that sought to test this hypothesis in awake behaving animals performing streaming tasks indicate that during behavior, responses to the attended stream become better segregated compared to those in response to background sounds [83].

In addition, attention can influence streaming by modulating the temporal coherence of neural populations [84]. Recent findings indicating that temporal coherence between distinct populations of neurons tuned to a target

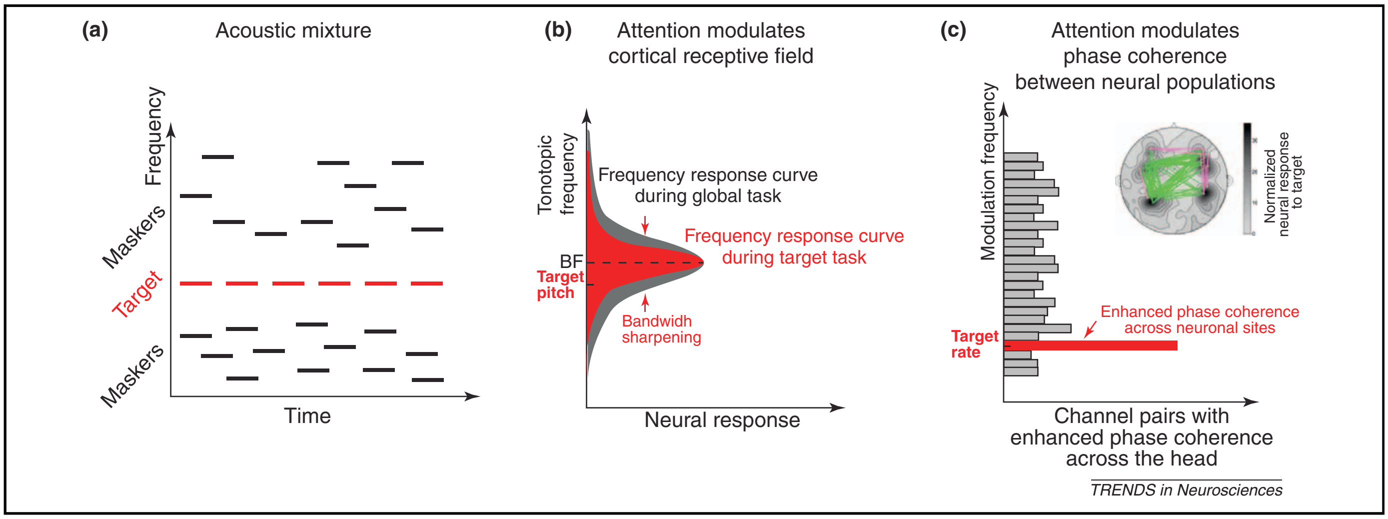

Figure 3. Schematic of the influence of attention on the cortical selectivity of sound features and the representation of coherent features of an attended stream. **(a)** Schematic of the time–frequency distribution of an acoustic mixture with a regularly repeating tone sequence (target) among a background of random tones (maskers). Perception of the target depends critically on a number of parameters, including the frequency separation between the target and closest masker components, the repetition rate of the target and the overall sequence duration. **(b)** Illustration of the frequency–response curve for a single unit recorded in the primary auditory cortex of a behaving ferret and the changes observed under two different behavioral tasks. When the animal attends to the repeating target tone (target task, red curve), the receptive field tuned to the target frequency sharpens in a direction that enhances segregation of the target from the background of the maskers. When the animal performs a listening task that involves attending to the entire sound mixture (global task, grey curve), the tuning curve shows a much broader shape relative to the selective attention state (adapted from [83]). **(c)** Phase coherence between distinct neural populations as measured by distributed magnetoencephalography (MEG) channels recording neural activity in human subjects. The phase coherence contrasts a selective attention task (in which subjects attended to the repeating target tone) to a global attention task (in which subjects paid attention to the background maskers). Such recordings reveal that enhancement of phase coherence occurs exclusively at the attended target repetition rate (in this case 4 Hz) (adapted from [75]). The inset represents an example of the MEG magnetic field distribution for a single listener, illustrating that MEG channel pairs have robust phase coherence in response to the rate of the target tone sequence. Channel pairs with enhanced phase coherence are shown in green and channel pairs with reduced coherence are shown in pink.

is augmented during attention (Figure 3) are consistent with this hypothesis. Enhanced phase coherence between distributed neuronal clusters helps to resolve the competition between different acoustic features in a sound mixture by facilitating temporal coherence analysis, thereby heightening the perceptual boundary between the currently attended stream and the background. Evidence of such a general mechanism by which attention influences the timing of neural responses has been found in the auditory system [74] and in the visual [85] and somatosensory modalities [86].

Temporal coherence reconciles feature- and object-based attention

Temporal coherence can help bind the diverse features of a stream in a manner that highlights an elegant synergy

Box 3. Hearing out sounds within a stream

The perceptual segregation of streams – a process of sequential organization – should not be confused with the hearing out of a component out of many simultaneous components in a sound complex such as a musical chord. When listening to a complex sound token, one can listen 'analytically', 'hear out' individual sound components, and even attend selectively to one of these components. For example, normal listeners can readily hear out a mistuned component in a harmonic complex [Figure Ia(iii)], perceptually attend to the different notes in a chord or to one of a pair of synchronous pure tones that are far apart in frequency (see Figure Ib in Box 2). These percepts, however, are not examples of streams because they do not arise through any sequential processes or organization and, moreover, they fail objective psychoacoustic criteria for streaming percepts (discussed below). A simple example is the case of the simultaneous tone sequences discussed earlier (Box 2), which, despite being readily heard as distinct tones, are nevertheless perceived as a single stream [48]. We claim that the same arguments apply to the sequences of mistuned harmonics illustrated in Figure Ia(iii), the double vowels illustrated in Figure Ib, and the two directional sounds shown in Figure Ic. In each of these cases, the distinct sound heard out of the complex mixture of sounds is nevertheless part of the same one stream because it produces coherent responses, the fundamental criterion for streaming.

A more complex example is the musical fragment illustrated in Figure Id, where we predict that the two opening bars are heard as a single, rich stream with all instruments playing in a temporally coherent fashion, just like an orchestra playing in unison. In the subsequent bars, two streams diverge as the oboe and the violins play incoherently. Another example involves multiple talkers (Figure Ie), as might occur during a cocktail party. It is generally agreed that the segregation of simultaneous voices in this case is largely facilitated by the temporal incoherence of their syllabic segments, which enables the listener to 'glimpse' (or gather snapshots) of the target voice during 'dips' in the other voice [98]. Viewed abstractly, the alternating bursts of the perceptually distinct green and pink speech patterns illustrated in Figure Ie are analogous to the alternating tones illustrated in Figure Ia of Box 2.

The proposed distinction between hearing-out components in a complex versus streams raises the important question of how to objectively measure listeners' perception of streams. A commonly used approach involves measuring listeners' ability to detect (or discriminate) differences in the relative timing of sounds. It has been found that when listeners hear different sounds as belonging to separate streams (subjectively), they lose the ability to detect (or discriminate) small differences in the relative timing of those sounds [97,99–101]. For instance, they can no longer tell if one sound, which is perceived as part of one stream, starts before or after another sound, which is perceived as part of another stream. Other approaches involve the detection or discrimination of changes in some attribute (e.g. pitch) of sounds in one stream in the presence of irrelevant (e.g. random) changes in the same or a different attribute in another stream [102–104].

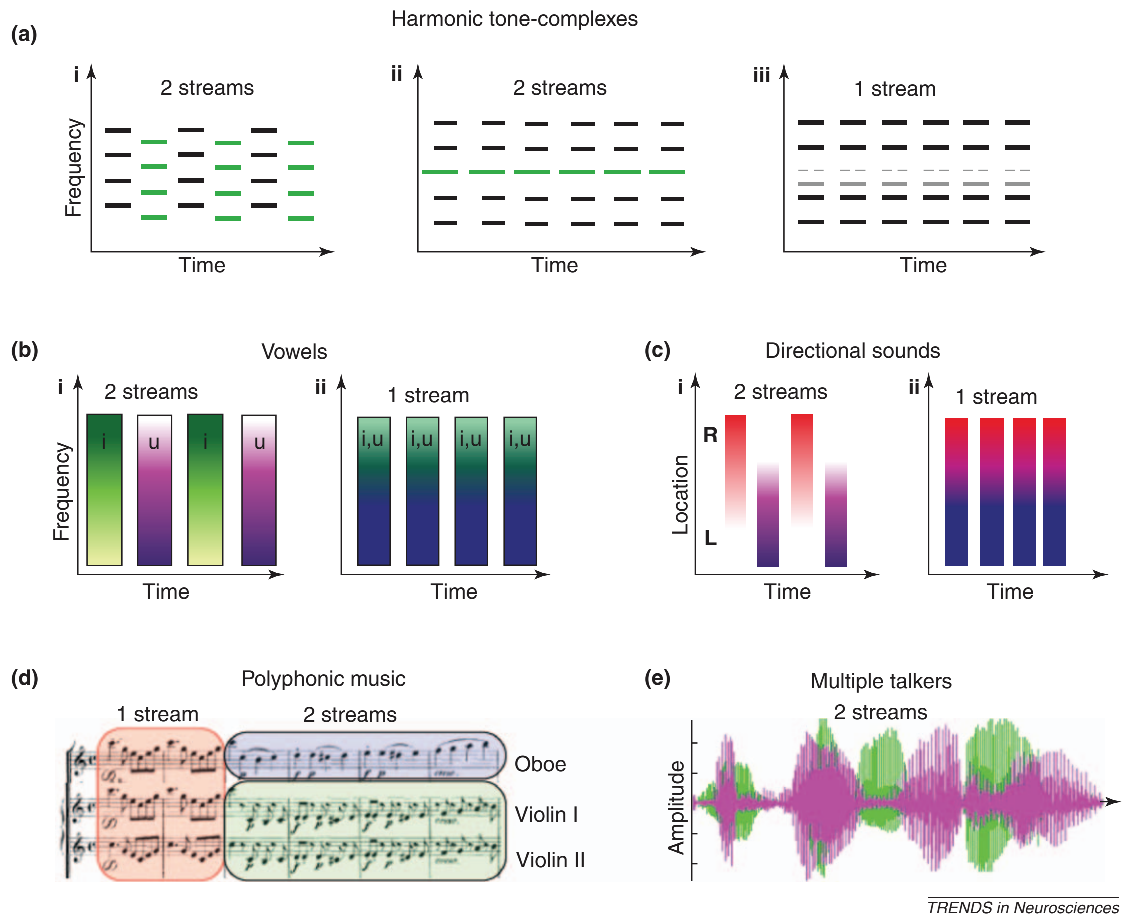

Figure I. Streaming with complex sounds. Principles of sequential organization apply equally well to complex stimuli that evoke responses in feature-selective channels (analogous to the frequency-tuned channels for tones). Several examples are illustrated. **(a)** Streaming with harmonic complexes. Harmonic complexes are usually perceived as a fused sound with a pitch at the frequency of the fundamental (bottom) component of each complex. **(i)** Two alternating complexes (green and black) stream apart, just like alternating pure tones [32]. **(ii)** A harmonic complex is perceptually fractured when one component begins earlier (e.g. the green harmonic). Because of its temporal incoherence, this component forms a separate stream from the rest of the complex (the black tones). **(iii)** A harmonic complex also becomes perceptually fractured when one component (the grey tone) is mistuned from a harmonic relationship and pops out from the complex. However, in this case the two percepts within the token continue to belong to a single stream because they remain temporally coherent. **(b)** Streaming of vowels. A sequence of vowel pairs is perceived either as two streams or as one, depending on the temporal coherence of the vowels [1]. **(i)** The alternating pair of vowels /i/ and /u/ is represented schematically by different spectra. These vowels (just like the alternating tones) segregate into two streams [1]. **(ii)** As for synchronous tones, when the vowels are played simultaneously they can still be individually recognized but are nevertheless heard as a single stream. **(c)** Streaming of sounds from different locations. Two sounds from the left (L) and right (R) stream apart when **(i)** they are played alternately [1,35], but form a single stream when **(ii)** played coherently. In the latter case, we predict that the sound is heard as a single stream from (indeterminate) multiple locations. **(d)** Streaming of musical instruments. The beginning of Mozart's Concerto K299 is illustrated here. The first two bars are heard as a single rich stream because all the instruments are playing coherently, despite the distinct timbres of the oboe and the violin, and the different notes (pitches) played by the two violins. In the subsequent bars, two streams diverge because the oboe and the violins play incoherently. **(e)** Streaming of two simultaneous talkers. When the waveforms for two different spoken sentences (represented by pink and green) are overlaid, they often appear as alternating sound tokens. This incoherence between the two waveforms (each with its own distinct timbre, pitch or even location) facilitates their streaming apart. In a choir singing in unison, the waveforms for the all singers would completely overlap and hence are heard as one rich stream (analogous to a piano playing a sequence of chords).

between object-based and feature-based attention in stream formation. When a feature is selectively attended to, it effectively serves as the anchor that points to and can be used to bind other features that are temporally coherent with it. For example, when attempting to hear out a female talker in the presence of a concurrent male talker, a listener might choose to attend to the high pitch (i.e. the female voice), and then through this particular feature perceptually access all other voice attributes that are coherent with it (e.g. location and timbre). If, instead, the listener has access to the approximate location of the female talker (e.g. based on visual information), he could attend selectively to the channels encoding this region of auditory space, and subsequently access other coherent attributes (e.g. pitch and timbre) of the female voice. Thus, as long as one distinctive feature of a target stream is sufficiently salient to be attended to by the listener, he could have access to and the ability to distinguish all other features of the target stream, owing to their temporal coherence. This process is, in some ways, similar to that invoked to explain the finding that observers who attend selectively to task-relevant visual features, learn to better discriminate not just these features, but also all other task-irrelevant features that occur concomitantly, even when they are too weak to be consciously perceived [87].

Memory in auditory scene analysis

Up to this point, our focus has been on the postulate that sequential auditory grouping processes utilize dynamic cues to stream sounds and render them perceptually as auditory objects. One might ask why the emphasis has been placed on dynamic cues, given that static scenes such as images have been the primary vehicle for the study of segmentation and identification of visual objects. We propose that in the absence of dynamic cues, recognition of objects in static scenes must combine memory (i.e. priors or heuristics) with low-level perceptual primitives such as edges, edge continuity, texture analysis and color. Such perceptual primitives are analogous to the percepts of harmonicity and binaural disparities in audition (referred to as tokens or instantaneous percepts in Box 1). For example, identification of a complex assemblage of oval shapes, straight and curved edges, and multiple colors and textures on a canvas as a face or a tree must invoke preexisting (either learned or hardwired) templates of these objects. The same logic applies to static auditory scenes: determination of whether a sustained (or steady) sound from a throat singer or from two simultaneous choir singers is either one source or two is essentially arbitrary and depends on the listener's expectations and contextual cues (memory) and not on sensory evidence alone. However, once dynamic cues are introduced, as when the two voices become dynamically modulated (coherently or incoherently) in pitch, loudness or timbre, sensory evidence becomes the key to perceptual streaming of the sound either into one complex source (i.e. a source being composed of two elements) or into two separate sources (Box 3).

To summarize, listening for sources in natural environments often engages hardwired preferences for conspecific vocalizations and memories of familiar sounds that are important to the animal for survival or reproduction [88–91]. However, in many common situations when sources are novel (such as speech produced by an unfamiliar speaker or musical notes of a novel melody) or when the acoustic environment is complex and cluttered, dynamic cues (temporal coherence) play the primary role in enabling attention to bind coherent attributes and organize them into streams.

Summary

Here, we proposed two ideas within an overall framework to explain the perception of auditory scenes. The first is that auditory stream formation is critically dependent on the temporal coherence between neural responses to sounds in the auditory cortex. Specifically, when stimulus-induced cortical responses are temporally coherent, the features they represent can potentially become perceptually unified (or bound) as one stream, distinct from other temporally incoherent responses. This principle explains stream formation and perception of a wide range of stimuli, including spectrally and temporally complex natural sounds such as voices and music. The second hypothesis is that attention influences stream formation by initiating the binding process and modulating the neural representations of the acoustic features and/or of temporal coherence patterns among these features. Both of these hypotheses remain under intensive scrutiny and experimentation. Nevertheless, they are already proving useful as a theoretical framework to broaden and guide future investigations (Box 4) of the neural basis of auditory scene analysis.

Box 4. Future directions

A number of questions regarding the perceptual organization of complex auditory scenes remain unresolved, ranging from neuronal mechanisms to behavior. Here, we highlight several of the key topics that are the subject of current and future investigations.

(1) Neural circuitry of auditory scene analysis
- What are the neural underpinnings of streaming in non-primary auditory and non-sensory cortical areas?
- Is there explicit evidence of temporal-coherence computations carried out at some level of auditory cortical processing?
- What is the neural signature of the emergence of auditory streams?

(2) Role of attention and behavior
- What are the neural correlates of streaming in behaving animals?
- Do attention-induced neuronal changes at the level of the auditory cortex show a causal effect with improved behavioral performance during streaming tasks?
- How does attention modulate the binding of acoustic features into perceptual streams?

(3) Scene analysis across modalities
- If confirmed by empirical evidence, does the principle of temporal coherence reveal a fundamental principle underlying scene analysis across sensory modalities?

Acknowledgments

This work was supported by the following grants to the authors: NIH R0107657, MURI N000141010278, AFOSR FA9550-09-1-0234 and NSF CAREER award IIS-0846112.

References

1 Bregman, A.S. (1990) *Auditory Scene Analysis: The Perceptual Organization of Sound*, MIT Press
2 Cherry, E.C. (1953) Some experiments on the recognition of speech, with one and two ears. *J. Acoust. Soc. Am.* 25, 975–979
3 Bee, M.A. and Micheyl, C. (2008) The cocktail party problem: what is it? How can it be solved? And why should animal behaviorists study it? *J. Comp. Psychol.* 122, 235–251
4 McDermott, J.H. (2009) The cocktail party problem. *Curr. Biol.* 19, R1024–1027
5 Marrone, N. *et al.* (2008) Evaluating the benefit of hearing aids in solving the cocktail party problem. *Trends Amplif.* 12, 300–315
6 Wang, D. and Brown, G.J. (2006) *Computational Auditory Scene Analysis: Principles, Algorithms, and Applications*, Wiley-IEEE Press
7 Carlyon, R.P. (2004) How the brain separates sounds. *Trends Cogn. Sci.* 8, 465–471
8 Nelken, I. (2004) Processing of complex stimuli and natural scenes in the auditory cortex. *Curr. Opin. Neurobiol.* 14, 474–480
9 Sinex, D.G. (2005) Spectral processing and sound source determination. *Int. Rev. Neurobiol.* 70, 371–398
10 Alain, C. (2007) Breaking the wave: effects of attention and learning on concurrent sound perception. *Hear. Res.* 229, 225–236
11 Micheyl, C. *et al.* (2007) The role of auditory cortex in the formation of auditory streams. *Hear. Res.* 229, 116–131
12 Fay, R.R. (2008) Sound source perception and stream segregation in nonhuman vertebrate animals. In *Auditory Perception of Sound Sources (Springer Handbook of Auditory Research, Vol. 29)* (Yost, W.A. *et al.*, eds), pp. 307–323, Springer
13 Bidet-Caulet, A. and Bertrand, O. (2009) Neurophysiological mechanisms involved in auditory perceptual organization. *Front Neurosci.* 3, 182–191
14 Snyder, J.S. and Alain, C. (2009) Toward a neurophysiological theory of auditory stream segregation. *Psychol. Bull.* 133, 780–799
15 Winkler, I. *et al.* (2009) Modeling the auditory scene: predictive regularity representations and perceptual objects. *Trends Cogn. Sci.* 13, 532–540
16 Shamma, S. and Micheyl, C. (2010) Behind the scenes of auditory perception. *Curr. Opin. Neurobiol.* 20, 361–366

17 Micheyl, C. and Oxenham, A.J. (2010) Pitch, harmonicity and concurrent sound segregation: Psychoacoustical and neurophysiological findings. *Hearing Res.* 266, 36–51
18 Hartmann, W. and Johnson, D. (1991) Stream segregation and peripheral channeling. *Mus. Percep.* 9, 155–184
19 Beauvois, M.W. and Meddis, R. (1996) Computer simulation of auditory stream segregation in alternating-tone sequences. *J. Acoust. Soc. Am.* 99, 2080–2270
20 McCabe, S. and Denham, M.J. (1997) A model of auditory streaming. *J. Acoust. Soc. Am.* 101, 1611–1621
21 Fishman, Y.I. *et al.* (2001) Neural correlates of auditory stream segregation in primary auditory cortex of the awake monkey. *Hear. Res.* 151, 167–187
22 Kanwal, J.S. *et al.* (2003) Neurodynamics for auditory stream segregation: tracking sounds in the mustached bat's natural environment. *Network* 14, 413–435
23 Bee, M.A. and Klump, G.M. (2004) Primitive auditory stream segregation: a neurophysiological study in the songbird forebrain. *J. Neurophysiol.* 92, 1088–1104
24 Fishman, Y.I. *et al.* (2004) Auditory stream segregation in monkey auditory cortex: effects of frequency separation, presentation rate, and tone duration. *J. Acoust. Soc. Am.* 116, 1656–1670
25 Bee, M.A. and Klump, G.M. (2005) Auditory stream segregation in the songbird forebrain: effects of time intervals on responses to interleaved tone sequences. *Brain Behav. Evol.* 66, 197–214
26 Micheyl, C. *et al.* (2005) Perceptual organization of tone sequences in the auditory cortex of awake macaques. *Neuron* 48, 139–148
27 Pressnitzer, D. *et al.* (2008) Perceptual organization of sound begins in the auditory periphery. *Curr. Biol.* 18, 1124–1128
28 Sussman, E. *et al.* (1999) An investigation of the auditory streaming effect using event-related brain potentials. *Psychophysiol.* 36, 22–34
29 Sussman, E.S. *et al.* (2007) The role of attention in the formation of auditory streams. *Percept. Psychophys.* 69, 136–152
30 Attias, H. and Schreiner, C.E. (1997) Temporal low-order statistics of natural sounds. *Adv. Neural Inf. Process. Syst.* 9, 27–33
31 Singh, N.C. and Theunissen, F.E. (2003) Modulation spectra of natural sounds and ethological theories of auditory processing. *J. Acoust. Soc. Am.* 114, 3394–3411
32 Vliegen, J. *et al.* (1999) The role of spectral and periodicity cues in auditory stream segregation, measured using a temporal discrimination task. *J. Acoust. Soc. Am.* 106, 938–945
33 Grimault, N. *et al.* (2002) Auditory stream segregation on the basis of amplitude-modulation rate. *J. Acoust. Soc. Am.* 111, 1340–1348
34 Moore, B.C.J. and Gockel, H. (2002) Factors influencing sequential stream segregation. *Acta Acustica* 88, 320–333
35 Boehnke, S. and Phillips, D. (2005) The relation between auditory temporal interval processing and sequential stream segregation examined with stimulus laterality differences. *J. Percept. Psychophys.* 67, 1088–1101
36 Schreiner, C.E. and Sutter, M.L. (1992) Topography of excitatory bandwidth in cat primary auditory cortex: single-neuron versus multiple-neuron recordings. *J. Neurophysiol.* 68, 1487–1502
37 Schreiner, C.E. (1994) Order and disorder in auditory cortical maps. *Curr. Opin. Neurobiol.* 5, 489–496
38 Versnel, H. *et al.* (1995) Ripple analysis in the ferret primary auditory cortex. III. Topographic and columnar distribution of ripple response parameters. *Aud. Neurosci.* 1, 271–286
39 Kowalski, N. *et al.* (1996) Analysis of dynamic spectra in ferret primary auditory cortex. I. Characteristics of single-unit responses to moving ripple spectra. *J. Neurophysiol.* 76, 3503–3523
40 Kowalski, N. *et al.* (1996) Analysis of dynamic spectra in ferret primary auditory cortex. II. Prediction of unit responses to arbitrary dynamic spectra. *J. Neurophysiol.* 76, 3524–3534
41 Schreiner, C.E. *et al.* (1997) Temporal processing in cat primary auditory cortex. *Acta Otolaryngol. Suppl.* 532, 54–60
42 Schreiner, C.E. (1998) Spatial distribution of responses to simple and complex sounds in the primary auditory cortex. *Audiol. Neurootol.* 3, 104–122
43 Miller, L.M. *et al.* (2002) Spectrotemporal receptive fields in the lemniscal auditory thalamus and cortex. *J. Neurophysiol.* 87, 516–527
44 Sutter, M.L. (2005) Spectral processing in the auditory cortex. *Int. Rev. Neurobiol.* 70, 253–298
45 Middlebrooks, J.C. *et al.* (1980) Binaural response-specific bands in primary auditory cortex (AI) of the cat: topographical organization orthogonal to isofrequency contours. *Brain Res.* 181, 31–48
46 Mrsic-Flogel, T.D. *et al.* (2005) Encoding of virtual acoustic space stimuli by neurons in ferret primary auditory cortex. *J. Neurophysiol.* 93, 3489–3503
47 Bendor, D. and Wang, X. (2005) The neuronal representation of pitch in primate auditory cortex. *Nature* 436, 1161–1165
48 Elhilali, M. *et al.* (2009) Temporal coherence in the perceptual organization and cortical representation of auditory scenes. *Neuron* 61, 317–329
49 Sheft, S. (2007) Envelope processing and sound-source perception. In *Auditory Perception of Sound Sources* (Yost, W.A. *et al.*, eds), pp. 233–280, Springer
50 Lu, T. *et al.* (2001) Temporal and rate representations of time-varying signals in the auditory cortex of awake primates. *Nat. Neurosci.* 4, 1131–1138
51 Shinn-Cunningham, B.G. *et al.* (2007) A sound element gets lost in perceptual competition. *Proceedings of the National Academy of Sciences of the United States of America* 104, 12223–12227
52 Jiang, D. *et al.* (1996) Frequency extent of two-tone facilitation in onset units in the ventral cochlear nucleus. *J. Neurophysiol.* 75, 380–395
53 Atencio, C.A. *et al.* (2009) Hierarchical computation in the canonical auditory cortical circuit. *Proc. Natl. Acad. Sci. U. S. A.* 106, 21894–21899
54 Luczak, A. *et al.* (2004) Multivariate receptive field mapping in marmoset auditory cortex. *J. Neurosci. Methods* 136, 77–85
55 Bizley, J.K. *et al.* (2009) Interdependent encoding of pitch, timbre, and spatial location in auditory cortex. *J. Neurosci.* 29, 2064–2075
56 Sadagopan, S. and Wang, X. (2009) Nonlinear spectrotemporal interactions underlying selectivity for complex sounds in auditory cortex. *J. Neurosci.* 29, 11192–11202
57 Pienkowski, M. and Eggermont, J.J. (2010) Nonlinear cross-frequency interactions in primary auditory cortex spectrotemporal receptive fields: a Wiener–Volterra analysis. *J. Comput. Neurosci.* 28, 285–303
58 Blake, R. and Lee, S.H. (2005) The role of temporal structure in human vision. *Behav. Cogn. Neurosci. Rev.* 4, 21–42
59 Alais, D. *et al.* (1998) Visual features that vary together over time group together over space. *Nat. Neurosci.* 1, 160–164
60 Rahne, T. *et al.* (2008) A multilevel and cross-modal approach towards neuronal mechanisms of auditory streaming. *Brain Res.* 1220, 118–131
61 Eggermont, J.J. (1990) *The Correlative Brain: Theory and Experiment in Neural Interaction*, Springer
62 von der Malsburg, C. (1994) The correlation theory of brain function. In *Models of Neural Networks II* (Domany, E. *et al.*, eds), pp. 95–119, Springer
63 von der Malsburg, C. and Schneider, W. (1986) A neural cocktail-party processor. *Biol. Cybern.* 54, 29–40
64 Wang, D.L. and Brown, G.J. (1999) Separation of speech from interfering sounds based on oscillatory correlation. *IEEE Trans. Neural Netw.* 10, 684–697
65 Lakatos, P. *et al.* (2005) An oscillatory hierarchy controlling neuronal excitability and stimulus processing in the auditory cortex. *J. Neurophysiol.* 94, 1904–1911
66 Elhilali, M. and Shamma, S.A. (2008) A cocktail party with a cortical twist: how cortical mechanisms contribute to sound segregation. *J. Acoust. Soc. Am.* 124, 3751–3771
67 Alain, C. and Arnott, S.R. (2000) Selectively attending to auditory objects. *Front. Biosci.* 5, D202–212
68 Alain, C. *et al.* (2002) Neural activity associated with distinguishing concurrent auditory objects. *J. Acoust. Soc. Am.* 111, 990–995
69 Dyson, B.J. and Alain, C. (2004) Representation of concurrent acoustic objects in primary auditory cortex. *J. Acoust. Soc. Am.* 115, 280–288
70 Snyder, J.S. and Alain, C. (2005) Age-related changes in neural activity associated with concurrent vowel segregation. *Brain Res. Cogn. Brain Res.* 24, 429–492
71 Hillyard, S.A. *et al.* (1973) Electrical signs of selective attention in the human brain. *Science* 182, 177–180
72 Tiitinen, H. *et al.* (1993) Selective attention enhances the auditory 40-Hz transient-response in humans. *Nature* 364, 59–60
73 Woldorff, M.G. *et al.* (1993) Modulation of early sensory processing in human auditory cortex during auditory selective attention. *Proc. Natl. Acad. Sci. U. S. A.* 90, 8722–8726

74 Bidet-Caulet, A. *et al.* (2007) Effects of selective attention on the electrophysiological representation of concurrent sounds in the human auditory cortex. *J. Neurosci.* 27, 9252–9261
75 Elhilali, M. *et al.* (2009) Interaction between attention and bottom-up saliency mediates the representation of foreground and background in an auditory scene. *PLoS Biol.* 7, e1000129
76 Shinn-Cunningham, B.G. (2008) Object-based auditory and visual attention. *Trends Cogn. Sci.* 12, 182–186
77 Duncan, J. (1984) Selective attention and the organization of visual information. *J. Exp. Psychol. Gen.* 113, 501–517
78 Desimone, R. and Duncan, J. (1995) Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* 18, 193–222
79 van Noorden, L.P. (1975) *Temporal Coherence in the Perception of Tone Sequences*, Eindhoven University of Technology
80 Cusack, R. *et al.* (2004) Effects of location, frequency region, and time course of selective attention on auditory scene analysis. *J. Exp. Psychol. Hum. Percept. Perform.* 30, 643–656
81 Carlyon, R.P. *et al.* (2001) Effects of attention and unilateral neglect on auditory stream segregation. *J. Exp. Psychol. Hum. Percept. Perform.* 27, 115–127
82 Fritz, J.B. *et al.* (2007) Auditory attention – focusing the searchlight on sound. *Curr. Opin. Neurobiol.* 17, 437–455
83 Yin, P. *et al.* (2007) Primary auditory cortical responses while attending to different streams. In *Hearing: From Sensory Processing to Perception* (Kollmeier, B. *et al.*, eds), pp. 257–266, Springer
84 Niebur, E. *et al.* (2002) Synchrony: a neuronal mechanism for attentional selection? *Curr. Opin. Neurobiol.* 12, 190–194
85 Kim, Y.J. *et al.* (2007) Attention induces synchronization-based response gain in steady-state visual evoked potentials. *Nat. Neurosci.* 10, 117–125
86 Steinmetz, P.N. *et al.* (2000) Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature* 404, 187–190
87 Watanabe, T. *et al.* (2001) Perceptual learning without perception. *Nature* 413, 844–848
88 Suied, C. *et al.* (2010) Why are natural sounds detected faster than pips? *J. Acoust. Soc. Am. Exp. Lett.* 127, EL105–110
89 VanRullen, R. *et al.* (2005) Spike times make sense. *Trends Neurosci.* 28, 1–4
90 Roye, A. *et al.* (2010) Is my mobile ringing? Evidence for rapid processing of a personally significant sound in humans. *J. Neurosci.* 30, 7310–7313
91 Kriegstein, K. and Giraud, A. (2010) Implicit multisensory associations influence voice recognition. *PLoS Biol* 4, e326
92 Moore, B.C.J. (2003) *An Introduction to the Psychology of Hearing*, Academic Press
93 Winer, J.A. and Schreiner, C.E., eds (2009) *The Auditory Cortex*, Springer
94 DiCarlo, J.J. and Cox, D.D. (2007) Untangling invariant object recognition. *Trends Cogn. Sci.* 11, 333–341
95 Best, V. *et al.* (2008) Object continuity enhances selective auditory attention. *Proc Natl Acad Sci U S A.* 105, 13174–13178
96 Lopez-Poveda, Enrique A. *et al.*, eds (2010) *The Neurophysiological Bases of Auditory Perception*, Springer
97 Micheyl, C. *et al.* (2010) Auditory stream segregation and the perception of across-frequency synchrony. *J. Exp. Psychol. Hum. Percep. Perf.* 36, 1029–1039
98 Qin, M.K. and Oxenham, A.J. (2003) Effects of simulated cochlear-implant processing on speech reception in fluctuating maskers. *J. Acoust. Soc. Am.* 114, 446–454
99 Bregman, A.S. and Rudnicky, A.I. (1975) Auditory segregation: stream or streams? *J. Exp. Psychol. Hum. Percept. Perform.* 1, 263–267
100 Roberts, B. *et al.* (2002) Primitive stream segregation of tone sequences without differences in fundamental frequency or passband. *J. Acoust. Soc. Am.* 112, 2074–2085
101 Roberts, B. *et al.* (2008) Effects of the build-up and resetting of auditory stream segregation on temporal discrimination. *J. Exp. Psychol. Hum. Percept. Perform.* 34, 992–1006
102 Kidd, G., Jr *et al.* (1994) Reducing informational masking by sound segregation. *J. Acoust. Soc. Am.* 95, 3475–3480
103 Micheyl, C. *et al.* (2005) Performance measures of auditory organization. In *Auditory Signal Processing: Physiology, Psychoacoustics, and Models* (Pressnitzer, D. *et al.*, eds), pp. 203–211, Springer
104 Micheyl, C. *et al.* (2007) Hearing out repeating elements in randomly varying multitone sequences: a case of streaming? In *Hearing – From Sensory Processing to Perception* (Kollmeier, B. *et al.*, eds), pp. 267–274, Springer