*Humans use local spectrotemporal correlations to detect rising and falling pitch*

Parisa A. Vaziri[1], Samuel D. McDougle*[2,3], Damon A. Clark*[3,4,5,6,7]

1 – Yale College, Yale University, New Haven, CT 06511
2 – Dept of Psychology, Yale University, New Haven, CT 06511
3 – Wu Tsai Institute, Yale University, New Haven, CT 06511
4 – Dept of Molecular Cellular and Developmental Biology, Yale University, New Haven, CT 06511
5 – Dept of Physics, Yale University, New Haven, CT 06511
6 – Dept of Neuroscience, Yale University, New Haven, CT 06511
7 – Quantitative Biology Institute, Yale University, New Haven, CT 06511
* Equal contributors and lead contacts: samuel.mcdougle@yale.edu, damon.clark@yale.edu

## Abstract

To discern speech or appreciate music, the human auditory system detects how pitch increases or decreases over time. However, the algorithms used to detect changes in pitch, or pitch motion, are incompletely understood. Here, using psychophysics, computational modeling, functional neuroimaging, and analysis of recorded speech, we ask if humans detect pitch motion using computations analogous to those used by the visual system. We adapted stimuli from studies of vision to create novel auditory correlated noise stimuli that elicited robust pitch motion percepts. Crucially, these stimuli possess no persistent features across frequency or time, but do possess positive or negative local spectrotemporal correlations in intensity. In psychophysical experiments, we found clear evidence that humans judge pitch direction based on both positive and negative spectrotemporal correlations. The observed sensitivity to negative correlations is a direct analogue of illusory "reverse-phi" motion in vision, and thus constitutes a new auditory illusion. Our behavioral results and computational modeling led us to hypothesize that human auditory processing employs pitch direction opponency. fMRI measurements in auditory cortex supported this hypothesis. To link our psychophysical findings to real-world pitch perception, we analyzed recordings of English and Mandarin speech and discovered that pitch direction was robustly signaled by the same positive and negative spectrotemporal correlations used in our psychophysical tests, suggesting that sensitivity to both positive and negative correlations confers ecological benefits. Overall, this work reveals that motion detection algorithms sensitive to local correlations are deployed by the central nervous system across disparate modalities (vision and audition) and dimensions (space and frequency).

## Introduction

From discriminating phonemes to being moved by Bach's *Partitas*, detecting changes in pitch over time, or pitch motion, is fundamental to human audition. Indeed, in everyday speech we use both intonation and lexical tones — including complex rising and falling pitches — to signify meaning (1-3). In English, for instance, rising pitch at the end of a sentence signifies a question. In Mandarin Chinese, changes of pitch within words conveys fundamental differences in meaning. But how does the human auditory system detect changes in pitch?

Changes in pitch can, in principle, be detected in at least two ways. First, listeners could identify an auditory "object" corresponding to the pattern of frequencies made by a voice or any other sound source (e.g., a friend's speech, a violin, etc.). If at the next instant in time the object moved to higher frequencies, listeners would infer a rising pitch, or the opposite if the object moved to lower frequencies (**Fig. 1A**). By identifying and tracking auditory objects, listeners can perceive changes in the object's pitch over time. In vision, humans use this "feature tracking" approach as one mechanism for detecting motion (4).

An alternative method for detecting changes in pitch would be to compute local correlations in sound volume over time at nearby frequencies. These local correlations would enable listeners to infer whether pitches are rising or falling without the added burden of first identifying auditory objects. Methods like these are the basis of canonical models for spatial motion detection in vision (5, 6). They can be dramatically revealed by visual illusions involving negative correlations, including "reverse phi" phenomena (5-7). Thus, at least in vision, humans use both object-tracking and intensity correlations to detect motion in the environment (8, 9).

Object tracking is a plausible method for detecting changes in pitch. Humans are clearly adept at identifying and tracking auditory objects: In the well-known "cocktail party" effect, guests at a noisy party can pick out and track a single voice in a sea of other voices (10-12). More generally, listeners can group nearby frequencies into auditory objects, which strongly influences the perception of rising and falling pitch (13). Likewise, the perception of continuity with rising and falling tones is also consistent with tracking auditory objects (14), and psychophysical studies of frequency change detection have tended to use isolated frequencies or persistent sound spectra in which auditory object tracking is possible (15-18). Studies also show that pitch change discrimination can occur over seconds, suggestive of object tracking (19).

What are the neural correlates of detecting rising and falling tones? Neurophysiological studies have shown that both subcortical neurons (20, 21) and cortical neurons (22, 23), including in primates (24), respond selectively to rising or falling tones in a narrow range of frequencies. They achieve this selectivity by nonlinearly combining different frequency inputs at different delays. Moreover, studies of many cortical auditory neurons have characterized complex spectrotemporal receptive fields, which show how responses depend on different frequencies over time (25, 26). Thus, although neural responses to auditory stimuli with local spectrotemporal correlations have not been measured to date, neurons with appropriate spectrotemporal tuning could detect such correlations. Neurons that detect rising or falling tones could in principle support algorithms that detect pitch motion by object tracking but could also, crucially, support those that work by sensing spectrotemporal correlations. It thus remains

unclear whether the human auditory system can use spectrotemporal correlations to perceive directed changes in pitch. In this study, we hypothesize that detecting local spectrotemporal correlations is a fundamental computation of the human auditory system.

## Results

### *Spectral motion without features*

We set out to test whether humans can detect auditory motion based on local spectrotemporal correlations. To do this, we adapted a stimulus used to study visual motion detection (27, 28) to develop new correlated noise auditory stimuli that use increments and decrements in volume to generate local correlations in volume at specific offsets in frequency and time (see **Methods**). We designed four stimuli with positive or negative correlations in volume at an offset of 1/6 second, with the frequency directed either upward or downward by 1/15 octave (**Fig. 1B, S1**). These sounds were inharmonic, so that fundamental frequencies could not be used to judge pitch changes (29, 30). We presented these stimuli to participants for 2 seconds and asked them to report whether they perceived the sound as having a rising or falling pitch profile over time.

Participants reported that upward-directed positive correlations rose in pitch over time, while those with downward directed positive correlations fell in pitch over time (**Fig. 1C, Supp. Movie 1**). This psychophysical result demonstrates that humans can identify rising or falling pitch based on local correlations alone, without persistent auditory objects.

Remarkably, when we presented stimuli with negative correlations in frequency and time, participants reported the opposite percepts (**Fig. 1C, Supp. Movies 1 and 2**). That is, the upward-directed negative correlations sounded like they were falling in pitch, while the downward directed negative correlations sounded like they were rising in pitch. Participants who consistently perceived rising or falling pitch in the stimuli with positive correlations also consistently perceived rising or falling pitch in the stimuli with negative correlations (**Figure S1**). This striking illusion demonstrates that humans are sensitive not just to positive spectrotemporal correlations, but to negative ones as well. This result is a direct analog to illusory reverse-phi visual motion percepts, which have been reported across many species and phyla (5, 7, 31-33).

How does the strength of these spectrotemporal correlations relate to perception? To answer this question, we varied the coherence of the stimulus and again asked participants to judge whether tones were rising or falling in pitch (**Fig. 1D**). We titrated the coherence of the stimuli from 1 to 0 by randomly replacing correlated time-frequency elements with random ones, such that the coherence represented the fraction of original correlations remaining (see **Methods**). With high coherence, participants perceived rising and falling pitches in a pattern similar to the first experiment (**Fig. 1C**). As coherence decreased, however, the probability of judging a sound as rising tended towards chance (0.5). There were no significant differences between the curves for (↑ +) and (↓ -) or (↓ +) and (↑ -) ($p > 0.05$ for each, as measured by a two-way, repeated measures ANOVA), meaning that inverting the stimulus correlation and direction led to indistinguishable percepts. These results reveal a clear monotonic relationship between the

strength of spectrotemporal correlations and the strength of pitch change percepts, both for positive and for negative correlations.

In vision, object tracking can integrate information between the two eyes, while correlation based algorithms rely on correlations within each eye (9). We next asked if spectrotemporal correlations for pitch motion detection are computed monaurally or binaurally. The structure of our correlated noise stimulus is created by summing a random binary mask with itself at a frequency-time offset (**Fig. 1E, Methods**). This allowed us to play one binary mask to the left ear and a shifted one to the right ear, so that neither ear alone would be presented with any correlations. In this context, detecting spectrotemporal correlations can only proceed by integrating information across the two ears. We played all four types of binaural correlations to participants and asked them to judge whether they heard rising or falling sounds. They reported the same pattern of percepts as in the monaural stimuli, though with average reported directions somewhat closer to chance. This demonstrates that the perception of rising or falling pitch can use information from both ears to integrate volume information to compute spectrotemporal correlations. This is consistent with data showing that many cortical auditory neurons integrate signals from both ears (34).

*Tuning of human spectrotemporal correlation detectors*

Our next step was to characterize the spectral and temporal tuning of the correlation sensitivity we had observed. To do this, we designed a different kind of stimulus, one inspired by random dot kinetograms in visual neuroscience (35). In these stimuli, a medium intensity sound that played at all frequencies was interrupted by brief pips at different frequencies, 50 ms in duration (24). These pips either increased the volume of a specific frequency or decreased it to zero (**Fig. 2A**, see **Methods**). After an initial set of pips were placed randomly in frequency and time, we added a second set of pips with a specific delay in time and change in frequency, yielding correlated pip pairs. These pairs had positive correlations when both pips were loud or both were silent, and negative correlations when one was loud and one was silent. This allowed us to create auditory stimuli with upward and downward-directed pairs of pips with positive or negative correlations (**Fig. 2B**). Like the stimuli used in **Fig. 1**, these stimuli had no auditory objects that persisted in time or frequency, but crucially, they allowed us to vary the delay continuously between correlated pips.

We first used these stimuli to map out the sensitivity to different delays between individuated tones. We kept the frequency change at 1/15 octave and swept values of the delay between correlated pips while asking participants to judge whether the pitch was rising or falling over time (**Fig. 2C**). For both negative and positive correlations and upward and downward-directed displacements, we found that peak directional sensitivity occurred at a delay of around 40 ms. This peak did not change appreciably when the pip duration was shortened to 20 ms (**Fig. S2**). According to models for visual motion estimation, this peak sensitivity value reflects the typical relative delays in the circuits detecting local motion signals (27). The delay seen here is on a similar timescale, though is slightly longer than delays measured by similar experiments in human and fly visual systems (27, 36).

We then measured sensitivity to the magnitude of displacements in frequency space. Using a similar method, we set the delay to 40 ms and varied the frequency displacement within a pip pair (**Fig. 2D**). We found that peak sensitivity occurred for tone displacements of 1/15th octave, though there was still significant direction-selectivity at 2/15th octave displacements ($p < 0.05$ for both positive and negative correlations by a paired t-test). This result shows that correlation-based motion detectors in the human auditory system are most sensitive to small shifts in frequency in the vicinity of 1/15th of an octave (4.7% changes in frequency) or less. This result is consistent with peak sensitivity for changes in complex sounds (37) and with the smaller values of frequency discrimination thresholds in humans (15).

*Sensitivity to spectrotemporal volume patterns*

Our positive and negative correlation stimuli each consist of multiple patterns in volume over frequency and time. Upward-directed positive correlation (↑ +) stimuli consist of both loud-loud and soft-soft combinations, whereas the negative versions (↑ –) consist of both loud-soft and soft-loud combinations. Prior work using long-lasting spectrotemporal correlations in auditory stimuli has suggested that humans are selectively sensitive to loud-loud combinations (38). Are humans sensitive to all four pairwise combinations, or to just a subset of them? To address this question, we generated new correlated pip auditory stimuli (**Fig. 3A**) where each stimulus had paired pips of only one of the four types: loud-loud, soft-soft, loud-soft, or soft-loud (see **Methods**). We asked participants to judge whether these different stimuli were rising or falling and recorded their responses (**Fig. 3B**). Participants were sensitive to all four different pairings with both upward and downward displacements.

In visual motion detection, one generalization beyond pairwise correlations involves so-called triplet correlations (39, 40). In vision, triplet correlations are patterns that contain spatiotemporal correlations over three points in space and time, but no pairwise correlations, and can elicit visual motion percepts in humans (39, 41), flies (41, 42), and fish (43). Visual motion detection algorithms are sensitive to this higher-order correlative structure, but is the same true in audition? When participants were presented with auditory analogs of visual triplet correlation stimuli (see **Methods**), they did indeed perceived auditory motion (**Figure S3**) and did so in a pattern much like that found in fly and fish visual perception. This correspondence across both species and modalities points to significant similarities in the neural algorithms used by animals in processing auditory and visual motion.

*Psychophysical and cortical signatures of opponent subtraction of spectral motion signals*

When we presented positively and negatively correlated stimuli, we observed a striking symmetry: Tuning of negative correlation percepts matched the tuning of positive correlation percepts, but in the opposite direction (**Fig. 2**). This clear symmetry is highly suggestive of an opponent architecture. To investigate this, we first built a simple motion energy model unit to describe a hypothetical directionally tuned auditory neuron (**Fig. 4A**). The model unit filtered sound intensity linearly over frequency and time in a pattern that enhanced upward-directed spectral motion, similar to prior suggestions (44), before sending the signal through a quadratic nonlinearity (6). When we presented this model with correlated pip stimuli (**Fig. 2**), it responded at an elevated baseline level but with deviations that depended on the direction and sign of the

stimulus correlation (**Fig. 4B**). As designed, it responded more to upward-directed positive correlations than to downward-directed ones. Since this model relies solely on pairwise correlations, it was also expected that negative correlation stimuli elicited equal and opposite deviations to positive correlation stimuli. Crucially, however, in this model, negatively correlated stimuli exhibit a different tuning from oppositely directed positive stimuli; that is, inverting the correlation is not equivalent to inverting the direction (i.e., the temporal delay).

We next created an opponent signal by subtracting signals from two model units with opposite directional tuning (**Fig. 4C**). This opponent signal responded to positively correlated stimuli with positive and negative values when they were directed upward and downward (**Fig. 4C**, *green*). Critically, this opponent signal has an important symmetry: Responses to negatively correlated stimuli have the same tuning as positively correlated stimuli in the opposite direction. Thus, upward-directed negative correlation stimuli yield the same responses as downward-directed positive correlation stimuli. We also derived this result analytically (see **Methods**): When motion energy signals are opponently subtracted, negative correlation stimuli elicit mean responses that match oppositely directed positive correlation stimuli.

To demonstrate that our data contained this symmetry, we compared percepts of negative correlation stimuli to percepts of positive correlation stimuli in the opposite direction, for both frequency change and delay time tuning (**Fig. 4D, E**, replotting data from **Fig. 2**). The curves appeared to fully superimpose. ANOVA tests confirmed that there was no measurable difference between the positive correlation curves and the flipped negative correlation curves (see figure legends for statistics). This robust symmetry between positive and negative correlation stimuli has also been found in visual motion detection in fruit flies (27) and in humans (36).

In primate vision, opponent subtraction occurs in visual area V5, also called MT (45, 46), which has been shown to be causally involved in visual motion percepts (47). Similarly, flies also subtract visual motion signals with opposing preferred directions (48). Motivated by our psychophysical results, analogies with vision, proposals for opponent subtraction to determine spectral direction (16), and by spectral direction opponent auditory cells found in bats (49), we reasoned that human auditory cortex might possess signatures of opponent processing.

We followed the logic of previous functional magnetic resonance imaging (fMRI) studies that identified opponent signals in human cortical area MT and used visual stimuli that summed motion in opposite directions (50). To start, we assume that cortical voxels involved in detecting spectral motion contain units that respond preferentially to rising tones and units that respond preferentially to falling tones, but none that respond to both (**Fig. 4F**) (51). Such a voxel should thus respond reliably to stimuli containing either rising or falling tones. The key distinction between a system with or without opponency lies in its response to a summed stimulus that contains superimposed rising and falling tones: If units are opponent, then the summed stimulus should cause a decrease in voxel activity due to a net suppression of signals in units with opponent responses (50). We therefore designed simple stimuli consisting of rising tones, falling tones, or their sum (**Fig. 4G, S4**) and presented them to subjects while measuring blood-oxygen-level-dependent (BOLD) signals via fMRI.

We searched within a broad auditory cortex mask for voxels that responded more to the non-summed (rising or falling) stimuli than to the summed (opponent) stimulus (see **Methods**). Strikingly, at both group and individual levels, a bilateral region within superior temporal cortex was significantly more activated by the non-summed stimuli than by the summed stimulus (**Fig. 4H, I**), consistent with opponency. The group map extended over multiple bilateral functional subregions of the human auditory cortex (52), including core regions A1 and RI, Area 52, and lateral and medial belt regions (**Fig. S4**). According to the opponency hypothesis, activity in opponent voxels should be similar in magnitude for rising and falling stimuli and suppressed for the summed stimulus. Thus, we wanted to ensure that our result followed this symmetry and was not biased by either the rising or falling stimulus alone (see **Methods**). Activity in putative opponent regions was indeed comparable for rising and falling tones (**Fig. 4J**). Overall, our fMRI findings demonstrate that a key result from our behavioral studies — the clear symmetry between positive and negative correlation percepts — lead to a specific neural hypothesis that was borne out in neuroimaging data. To our knowledge, this general region of human auditory cortex has not previously been identified as a potential locus for opponent spectral motion signals.

*Positive and negative correlation spectrotemporal cues signal tone modulation in speech*

Is there an ecological advantage in detecting both positive and negative spectrotemporal correlations? To address this question, we chose to look at human speech, where tone modulation contains critical semantic information in both tonal and non-tonal languages (1-3). Since humans are sensitive to both positive and negative pairwise correlations in frequency and time, we hypothesized that these correlations could convey information about the direction and speed of tone modulation in human speech. Following in the tradition of relating auditory processing to natural sounds (53), we analyzed corpora of spoken English and Mandarin and examined how tone modulation is related to underlying positive and negative pairwise spectrotemporal correlations in volume (**Fig. 5**, **Methods**).

Our analysis took several steps. First, we computed spectrograms for each of the speech recordings (**Fig. 5A**, *top*). We then used an optical flow algorithm to estimate the change in tone at each point in time – that is, the degree to which the sound was rising or falling in frequency at each time (**Fig. 5A**, *bottom*, see **Methods**). Next, we binarized the spectrogram and looked for specific patterns of volume in frequency and time, examining all four combinations of loud and soft: loud-loud, soft-soft, loud-soft, and soft-loud (**Fig. 5B**). We next computed the local net signal for each pattern at each frequency and time by subtracting the downward directed patterns from the upward directed ones (**Figs. 5C, D**). Finally, we averaged these local net signals over all frequencies to obtain a net pattern signal (**Figs. 5C, D**). Computing net pattern signals is consistent with the opponency we observed psychophysically and in fMRI (**Fig. 4**). For the loud-loud patterns, there was a positive correlation between the time trace of the net pattern signal and the tone change. For the loud-soft patterns, the correlation was negative. The clear suggestion is that negative correlations contain information about tone changes that could be useful to listeners in detecting rising and falling tones in speech.

To see whether this result generalized, we analyzed hundreds of speech snippets that totaled over 90 minutes in English and 40 minutes in Mandarin Chinese (**Fig. 5E, F**, see **Methods**). In

English, the tone changes should be dominated by intonation, while in Mandarin Chinese, the tone changes should reflect both intonational and within-syllable changes in tone (1-3). We reproduced the analysis of the different volume patterns, and then correlated the net signal for each pattern with the computed change in tone. In both English and Mandarin Chinese, the two positive correlation patterns (loud-loud and soft-soft) produced a strong positive correlation ($r > 0.5$) with the intonation velocity, whereas the two negative patterns (loud-soft and soft-loud) produced a strong negative correlation ($r < -0.5$) (**Fig. 5E, F**). These results show that all four patterns could be useful in estimating tone changes in speech. The negative stimulus correlation produced an *anti*-correlation with tone changes, which explains why they elicit percepts in the opposite direction: upward directed negative correlations indicate downward directed tone changes. We obtained similar results when we processed with the spectrograms with continuous rather than digital operations to obtain positive and negative spectrotemporal correlations in the speech data (see **Methods**, **Fig. S5**). Thus, this analysis provides an ecological explanation of the observed inverted percepts to negative auditory correlations.

## Discussion

In the studies reported here, we have demonstrated that humans are sensitive to local spectrotemporal correlations in volume over frequency and time as they discern whether a sound is rising or falling in pitch (**Figs. 1-3**). Participants were equally sensitive to both negative and positive spectrotemporal correlations, a pattern that mirrors a powerful visual phenomenon, the reverse-phi illusion, in a different modality (audition) and over a different dimension of motion (frequency). Inspired by our behavioral results showing symmetry between inverting correlation and inverting direction, we hypothesized that the human auditory system might implement opponent subtraction, echoing a similar operation in visual motion detection. Using fMRI, we discovered that, like visual cortex, regions within human auditory cortex show signatures of opponency (**Fig. 4**). Finally, we demonstrated that negative spectrotemporal correlations likely act as reliable cues to assess tone changes in speech (**Fig. 5**).

The stimuli we developed here (**Figs. 1-3**) in some ways resemble Shepard tones (54), which were designed to sound like they are unceasingly rising or falling. However, Shepard tones consist of periodic auditory features that persist over frequency and time (similar to **Fig. 4F**). Thus, the rising or falling of a Shepard tone could be assessed by simply tracking auditory features over time. The auditory stimuli we developed and investigated here, however, have no such persistent features — a rising or falling percept must instead depend on the detection of positive and negative pairwise spectrotemporal correlations within the stimulus. Thus, the strong percepts of rising and falling tones, which depended on the sign of the correlation, reflect an authentic auditory illusion in which there is no true rising or falling tone but only the imposition of specific spectrotemporal correlations in volume.

Sensitivity to spectrotemporal correlations in judging pitch direction likely acts in coordination with other algorithms for judging changes in pitch. In particular, changes in frequency can be judged over gaps of seconds (37), which points to a different system for such judgements. Similarly, judgements about relative pitch can be made using fundamental frequencies in harmonic sounds (29, 55). These examples suggest that auditory spectral motion processing is

similar to visual motion processing, where positional changes can be detected by both local correlational algorithms and by slower, longer range object-tracking algorithms (8).

We found regions in both primary and non-primary auditory cortex across both Heschl's gyrus and the superior temporal gyrus (STG) that may perform opponent computations to resolve net pitch direction (**Fig. 4H, I**). How might this relate to the neural underpinnings of speech perception? Human auditory cortex displays regional specialization, with areas that selectively encode different aspects of speech, primarily in the STG (56-59). Our results are broadly consistent with findings that regions within STG encode variability in speaker intonation and lexical tone (59, 60). Moreover, our observation that significant portions of Heschel's gyrus also showed spectral motion sensitivity is broadly consistent with other work (61), though we saw the effects bilaterally (**Fig. S4**) and, critically, with an opponent signature. Our results thus suggest that opponency may be a signature of pitch direction processing in circuits involved in simple pitch computations (in primary areas) and in more complex perceptual tasks like speech processing (in non-primary areas).

Canonical algorithmic models for motion detection are sensitive to negative correlations (5, 6), and more neurophysiologically-inspired models for motion detection are similarly sensitive to negative correlations (62, 63). At the single neuron level, units in rodent (22), bat (23), and primate (24) auditory cortex display spectrotemporally oriented receptive fields, which should confer sensitivity to both positive and negative spectrotemporal correlations (**Fig. 4**) (6). Our results suggest that neurons with this type of sensitivity could underlie spectrotemporal correlation detection in humans. Meanwhile, our psychophysical and fMRI results also suggest that units in multiple regions of auditory cortex exhibit directional opponency, a property observed in bat auditory neurons (49). This direction opponency could arise in primary motion detectors (64), or be the result of subtracting opposing cortical or subcortical motion signals (21).

There are well-established similarities in the processing of visual motion between invertebrates and vertebrates (65-67), phyla that diverged hundreds of millions of years ago. Our study shows that local correlational algorithms for motion detection also span modalities, since human audition and vision appear to employ similar computational motifs. Audition thus joins olfaction (68) as a non-visual sense where pairwise, local correlations can generate rich motion percepts. In these experiments, sensitivity to pairwise stimulus correlations also includes sensitivity to negative correlations. This sensitivity to negative correlations is due in part to the mathematics of computing correlations (see **Methods**) (6, 40), providing a conceptual framework for understanding the neural detection of motion that spans modality and species.

Lastly, negative correlations sensed in audition likely act as useful cues to infer real-world changes in the frequency domain (**Fig. 5**), just as they may help in visual motion detection (69, 70). Thus, the illusory pitch motion described here is not just an interesting laboratory epiphenomenon. Rather, it reflects neural sensitivity to the statistics of the auditory world, with direct implications for everyday speech and music perception.

## Contributions

PAV and DAC designed auditory stimuli. PAV and SDM acquired data. PAV, SDM, and DAC analyzed and interpreted data. PAV, SDM, and DAC wrote the paper.

**Acknowledgements**

This work was funded by a grant from the Wu Tsai Institute at Yale University. DAC and this work were funded by NIH R01 EY026555. We thank R. Aslin, E. V. Clark, H. H. Clark, I. Yildirim, and J. Zavatone-Veth for helpful discussions and comments on this project.

# Figures

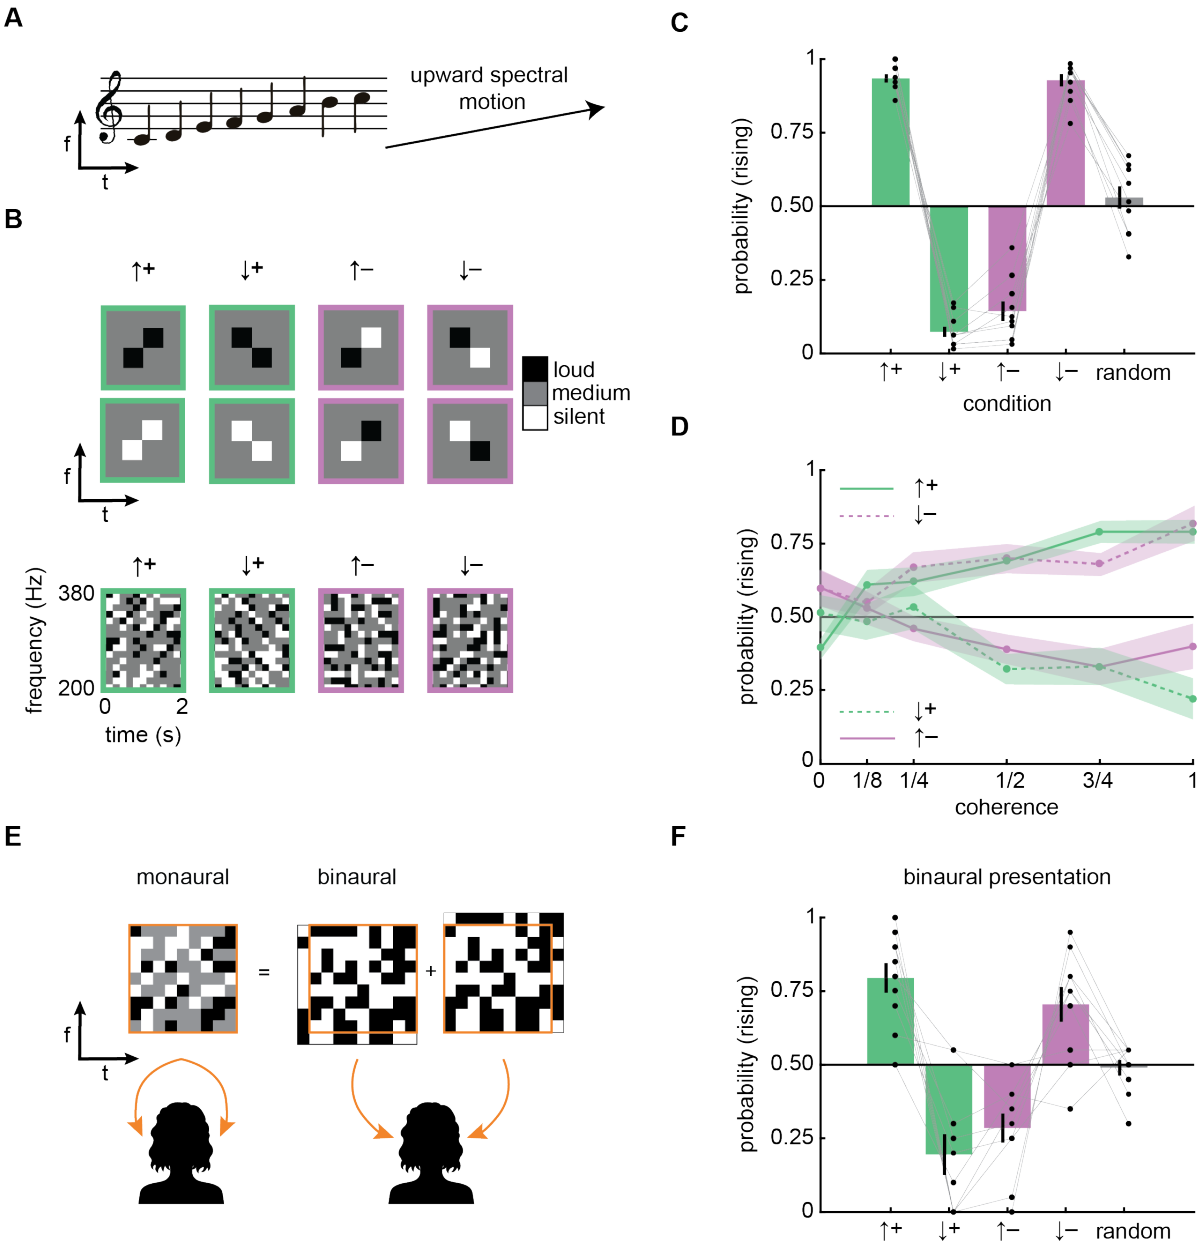


**Figure 1.** Humans detect auditory motion in pairwise frequency-time correlations.

A) Simple schematic of a rising sound written on a music staff and in frequency-time.

B) Diagrams showing sample (*top*) and actual (*bottom*) stimuli. Frequency-time correlations can be directed either upward or downward and be either positively or negatively correlated.

C) Perceived direction of stimuli with varying direction and correlation. Mean ± SEM over N=10 subjects. One-sample t-tests revealed significant deviations from chance (0.50) in pitch direction judgements in all four stimulus conditions (all *p*s < $10^{-5}$). Pitch direction judgements in the random stimulus condition were not significantly different from chance (p = 0.45). Error bars represent mean ± SEM (N = 10).

D) Perceived direction of stimuli with varying degrees of correlation (coherence) in the stimulus. The upward directed positive and downward directed negative curves were not significantly different ($p > 0.05$ by a two-way, repeated measures ANOVA); similarly, the downward directed positive and upward directed negative curves were also not significantly different ($p > 0.05$, same test). Both ANOVAs revealed significant main effects of coherence on pitch direction judgements (all $p$s $< 10^{-5}$). Error shading represents ± SEM (N = 10).

E) Diagram showing how binaural stimuli were presented to each ear.

F) Perceived direction of stimuli with varying directions and correlations using binaural presentation. One-sample t-tests revealed significant deviations from chance (0.50) in pitch direction judgements in all four stimulus conditions (all $p$s $< 10^{-3}$). Pitch direction judgements in the random stimulus condition were not significantly different from chance ($p = 0.72$). Error bars represent mean ± SEM (N = 10).

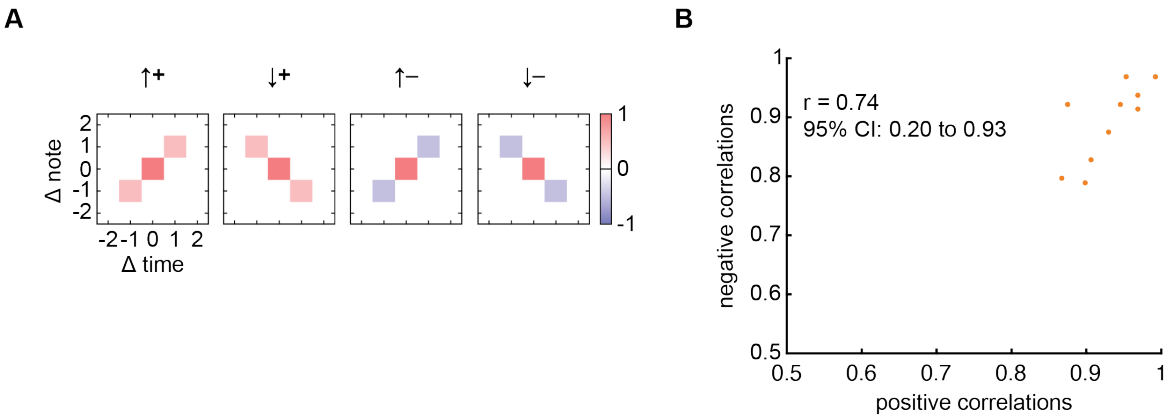


**Supplemental Figure S1.**

A) Stimulus autocorrelation plots at different note and time offsets for the stimuli in **Figure 1B**. The stimuli have positive or negative correlations at a single spectrotemporal offset, directed either upward or downward in frequency over time. These plots are normalized so that the origin has correlation of 1.

B) Correlation between perception of positively correlated and negatively correlated stimuli. To obtain the positive correlation values, we averaged P(rising) for the upward directed, positive correlation stimuli with 1-P(rising) for the downward directed, positive correlation stimuli. To obtain the negative correlation values, we averaged P(rising) for the downward directed, negative correlation stimuli with 1-P(rising) for the upward directed, negative correlation stimuli. Correlation coefficient is the Pearson correlation, and a 95% confidence interval is noted.

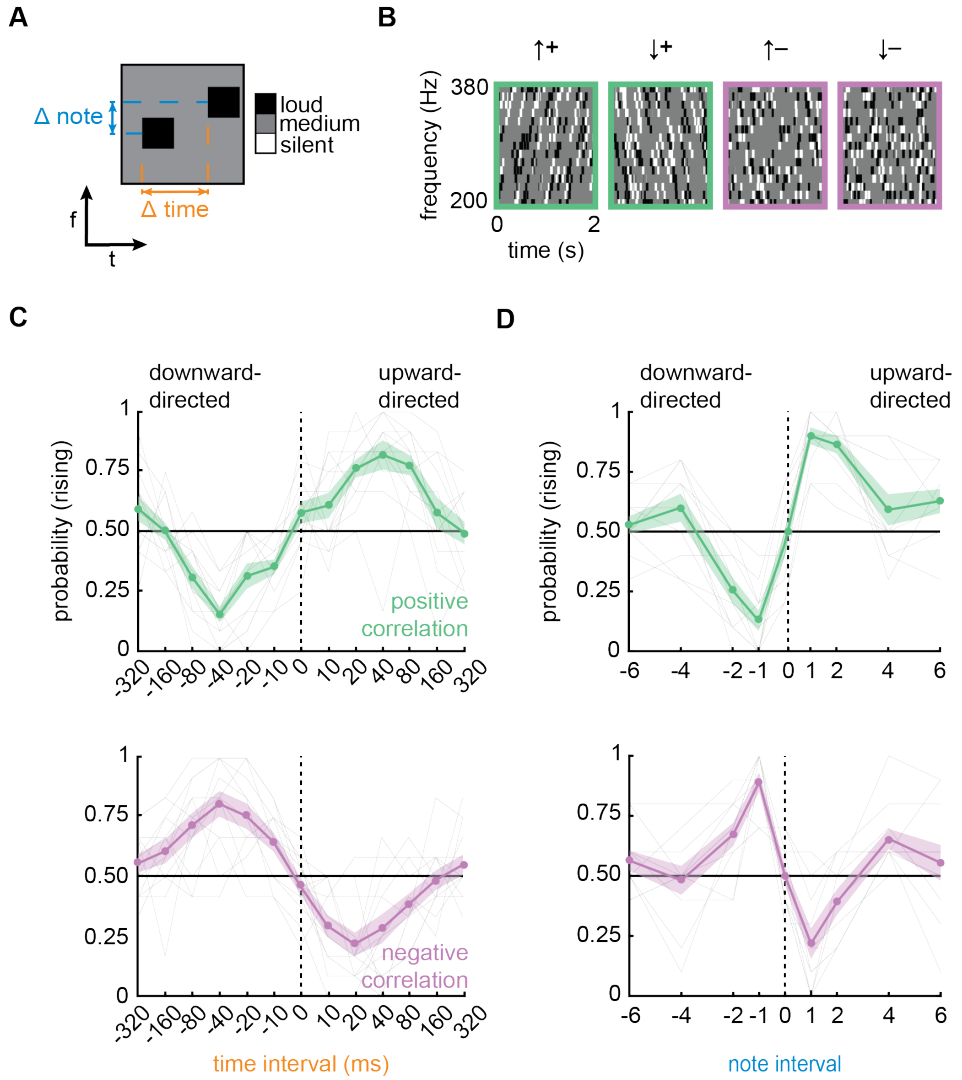


**Figure 2.** Correlation detection is tuned to small frequency changes and short delays in time.

A) Diagram showing a correlated pip pair with a frequency displacement (Δ note) and a delay between pips.

B) Spectrotemporal diagrams of 4 different correlated pip stimuli directed upward and downward with positive and negative correlations. Pip duration in these experiments was 50 ms.

C) Perceived direction of stimuli with Δ note = +1 and varying pip delays; positive pip correlations (*top*) and negative pip correlations (*bottom*). One-way, repeated measures ANOVAs for the positive and negative correlation curves revealed significantly different responses across pip delays (all $ps < 10^{-21}$). Gray lines are individual participant curves. Error shading represents ± SEM (N = 13).

D) Perceived direction of stimuli using varying note intervals and 40 ms pip delays; positive pip correlations (*top*) and negative pip correlations (*bottom*). One-way, repeated measures ANOVAs for the positive and negative correlation curves revealed significantly different responses across note intervals (all $ps < 10^{-12}$). Gray lines are individual participant curves. Error shading represents ± SEM (N = 13).

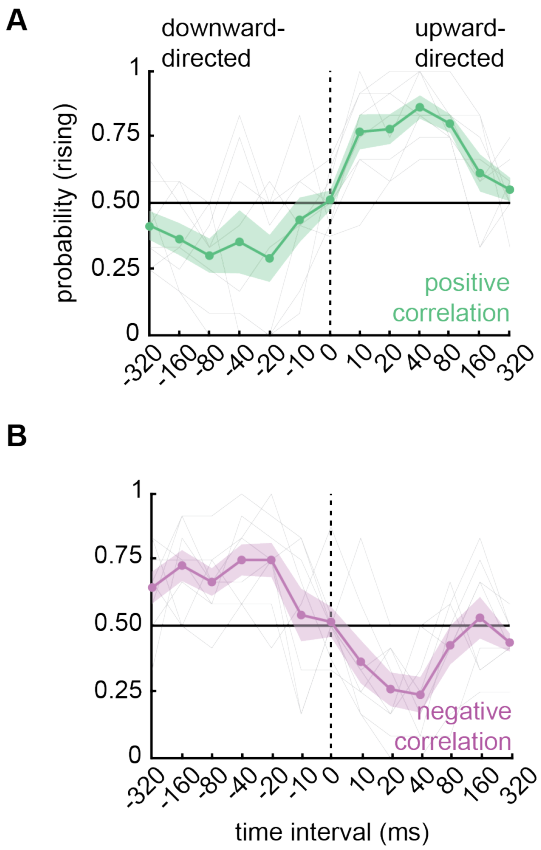


**Supp. Fig. 2.** Interval sweep with a different pip duration.

A) Perceived direction of positively correlated stimuli with varying pip delays and 20 ms pips. Sensitivity tends to peak around 40 ms delays, similar to the data in **Figure 2C**. A one-way, repeated measures ANOVA revealed significantly different responses across pip delays ($p < 10^{-10}$). Error shading represents ± SEM (N = 9).

B) Perceived direction of negatively correlated stimuli with varying pip delays and 20 ms pips. Sensitivity tends to peak around 40 ms delays, similar to the data in **Figure 2C**. A one-way, repeated measures ANOVA revealed significantly different responses across pip delays ($p < 10^{-8}$). Error shading represents ± SEM (N = 9).

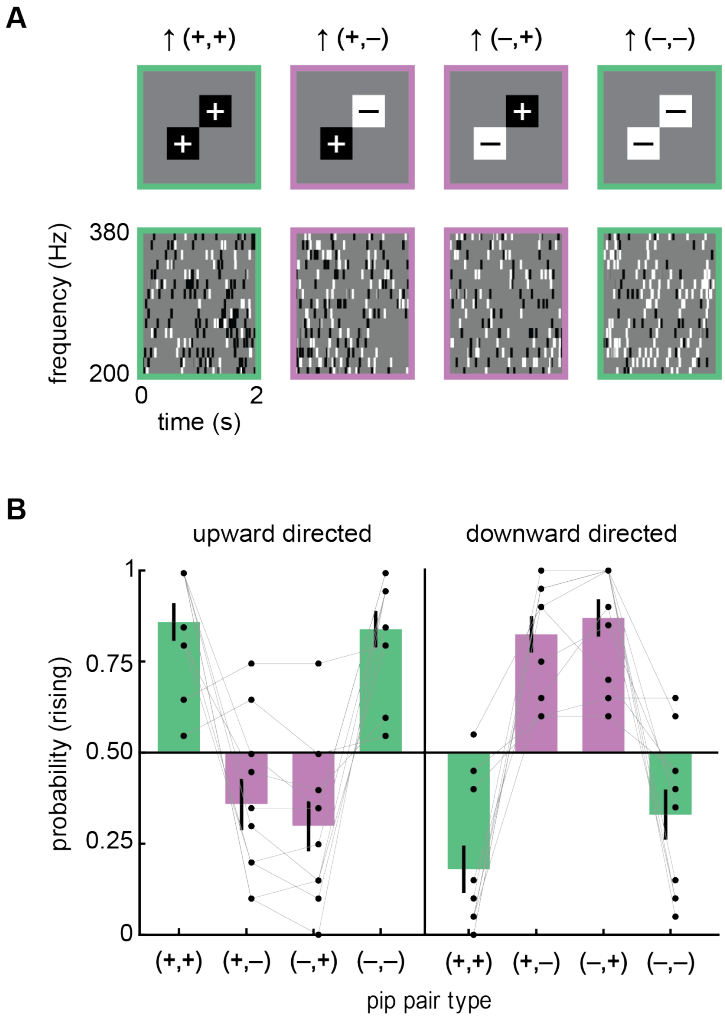


**Figure 3.** Sensitivity to all four pairwise loudness combinations contribute to rising and falling pitch perception.

A) Frequency-time diagram of 4 different pip combinations, presented with 40 ms delays.

B) Probability of perceiving rising pitch for each of the four loudness combinations directed upward (*left*) and downward (*right*). Paired t-tests comparing upward- versus downward-directed stimuli for each matched pair revealed significant direction selectivity across all pitch direction judgements (all $ps < 10^{-4}$). Error bars represent mean ± SEM (N = 10).

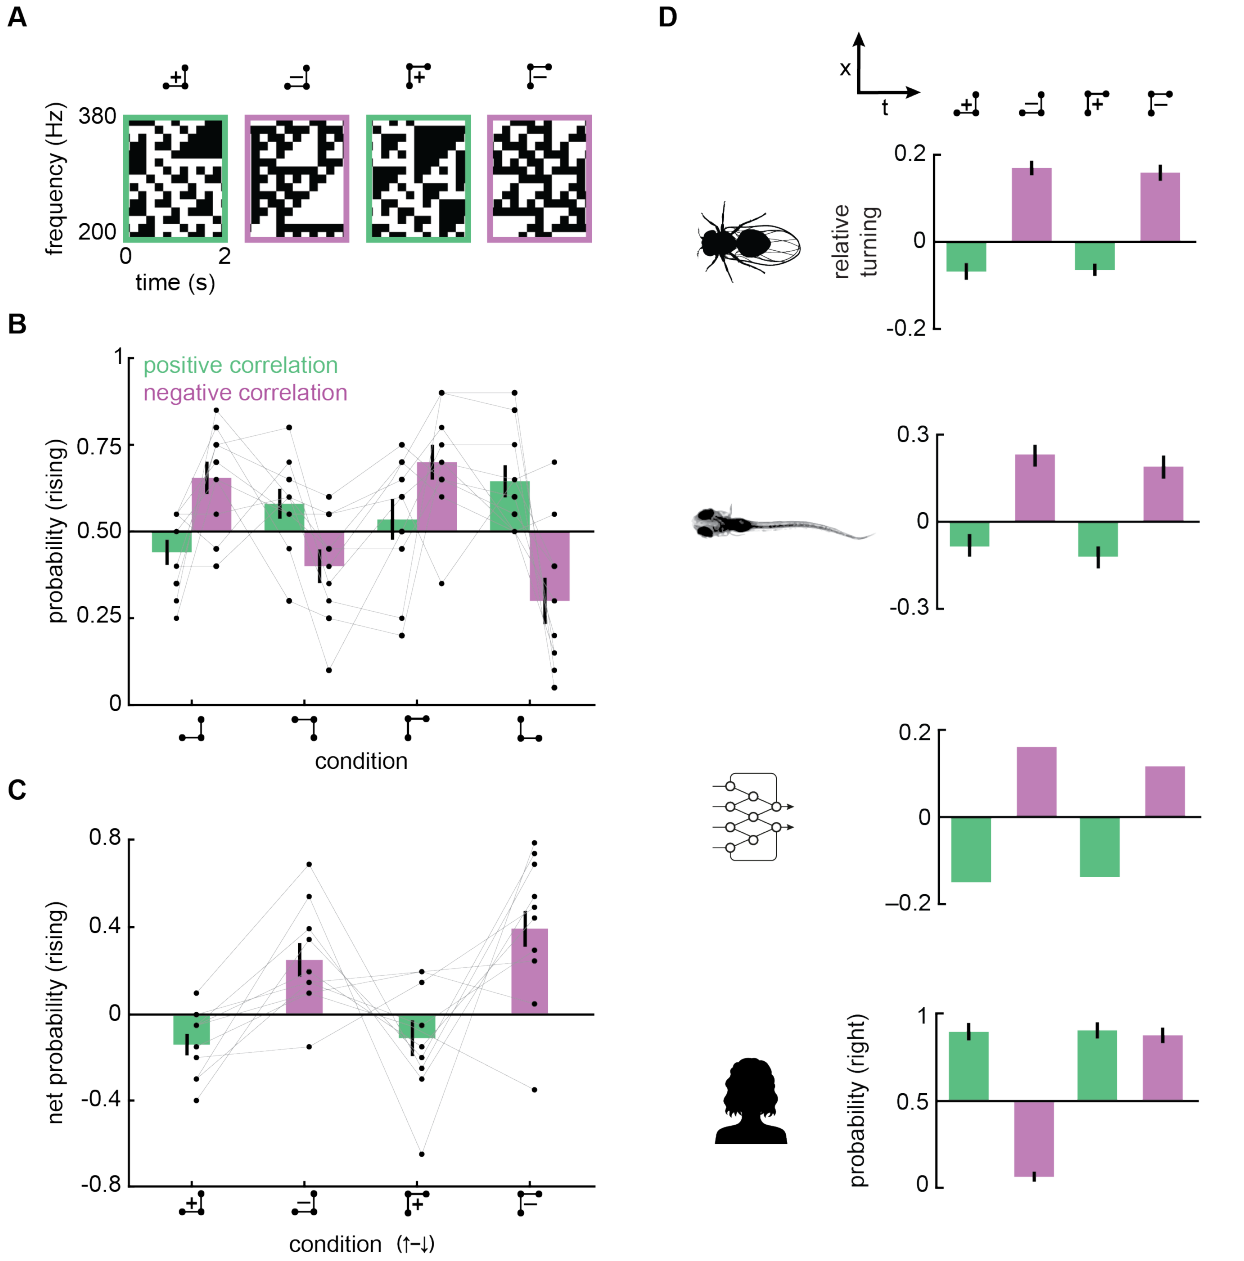


**Supplemental Figure S3.** Human auditory sensitivity to 3-point glider stimuli resembles visual sensitivity in different species.

A) Diagram of 3-point glider stimuli in frequency and time (39). 3-point glider stimuli contain correlations between triplets of points as denoted by the barbell diagrams, and contain no pairwise correlations. Thus, motion percepts with these stimuli would have to rely on correlations beyond pairwise ones.

B) Perceived direction of 3-point glider stimuli. Participants heard rising and falling tones in these triplet correlation stimuli. Error bars represent mean ± SEM (N = 10).

C) Net perceived direction of 3-point glider stimuli with positive and negative correlations. The net probability rising is computed by subtracting the downward directed P(rising) from the upward directed P(rising) in panel (B). Positively correlated stimuli were perceived as falling, while negatively correlated stimuli were perceived as rising. Paired

t-tests revealed significantly different responses to positively and negatively correlated diverging gliders, and to positively and negatively correlated converging gliders (all *p*s < $10^{-3}$). Error bars represent mean ± SEM (N = 10).

D) Net perceived direction of 3-point glider stimuli across various visual systems. Data is replotted from prior publications for fruit flies (41), larval zebrafish (43), a machine learning algorithm (69), and human visual psychophysics (41). Human auditory percepts resemble fruit fly and zebrafish visual percepts and machine learning responses, but not the human visual percepts.

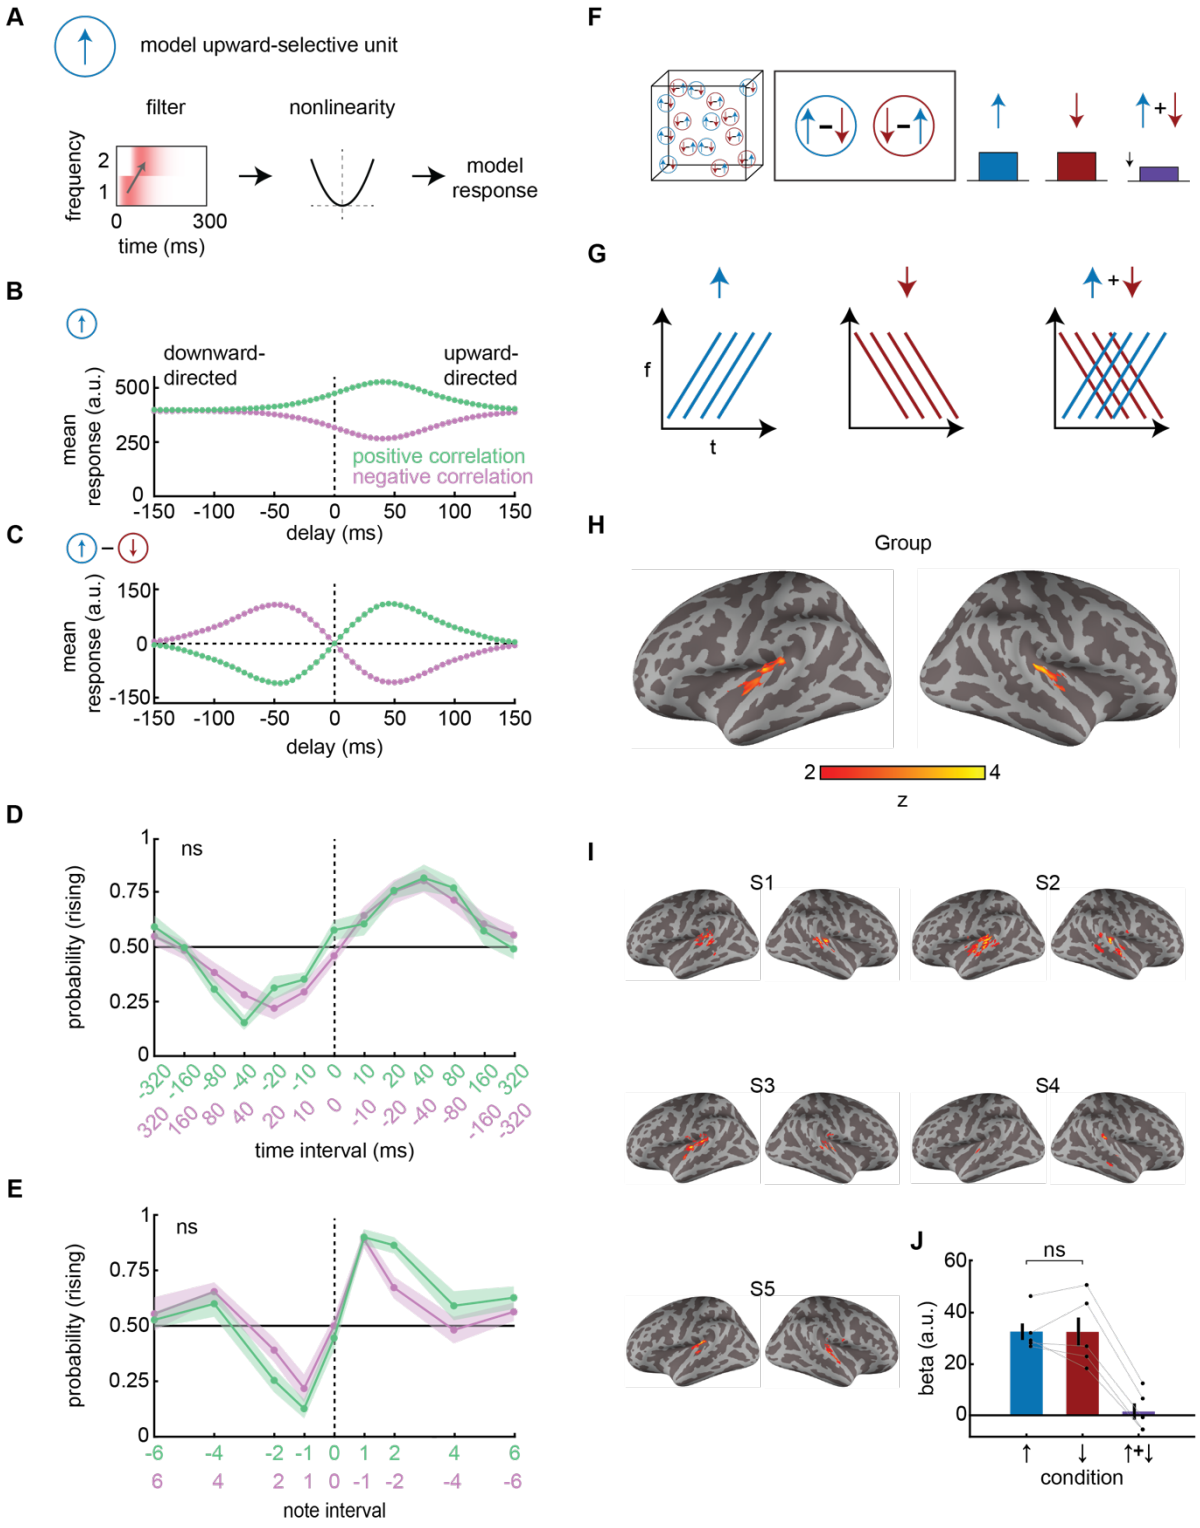


**Figure 4.** Bilateral regions of human auditory cortex show signatures of opponency.

A) A simple model auditory unit that responds more to upward direction spectral motion than downward directed spectral motion. The stimulus spectrogram is convolved with an upward-oriented spectrotemporal filter before the result is squared, as in a motion energy model (6).

B) Mean response of the unit to correlated pip stimuli with different delays and correlation signs, corresponding to upward and downward directed positive and negative correlations.

C) As in (B), but for an opponent signal, consisting of an upwardly tuned unit response minus an identical unit tuned to downward motion.

D) Comparison of P(rising) for positive and negative correlation stimuli sweeping time interval, aligning upward directed positive correlation stimuli with downward directed negative correlation stimuli. Data replotted from Figure 2. The curves were not significantly different ($p > 0.05$ by a two-way, repeated measures ANOVA).

E) As in (D) but for sweeping the tone difference. The curves were not significantly different ($p > 0.05$ by a two-way, repeated measures ANOVA).

F) Conceptual schematic of opponency in brain regions. An opponent voxel/region would respond strongly to rising and falling tones but be suppressed by the sum of the two stimuli.

G) Stimulus design. Stimuli were rising, falling, or summed rising and falling.

H) Group level analysis. A bilateral region within auditory cortex responded less to summed stimuli than non-summed stimuli. Cluster-corrected with false positive rate at $p < 0.05$ with a cluster-forming threshold of 20 voxels.

I) Individual level analysis. Regions in auditory cortex across subjects responded less to summed stimuli than non-summed stimuli. Cluster-corrected with false positive rate at $p < 0.05$ with a cluster-forming threshold of 20 voxels.

J) Control analysis showing symmetric beta values in response to rising and falling stimuli in individually defined opponent ROIs ($p > 0.05$ via one-sample t-test). (Note that all beta values are relative to an implicit baseline that includes responses to ambient scanner noise.) Error bars represent mean ± SEM ($N = 5$).

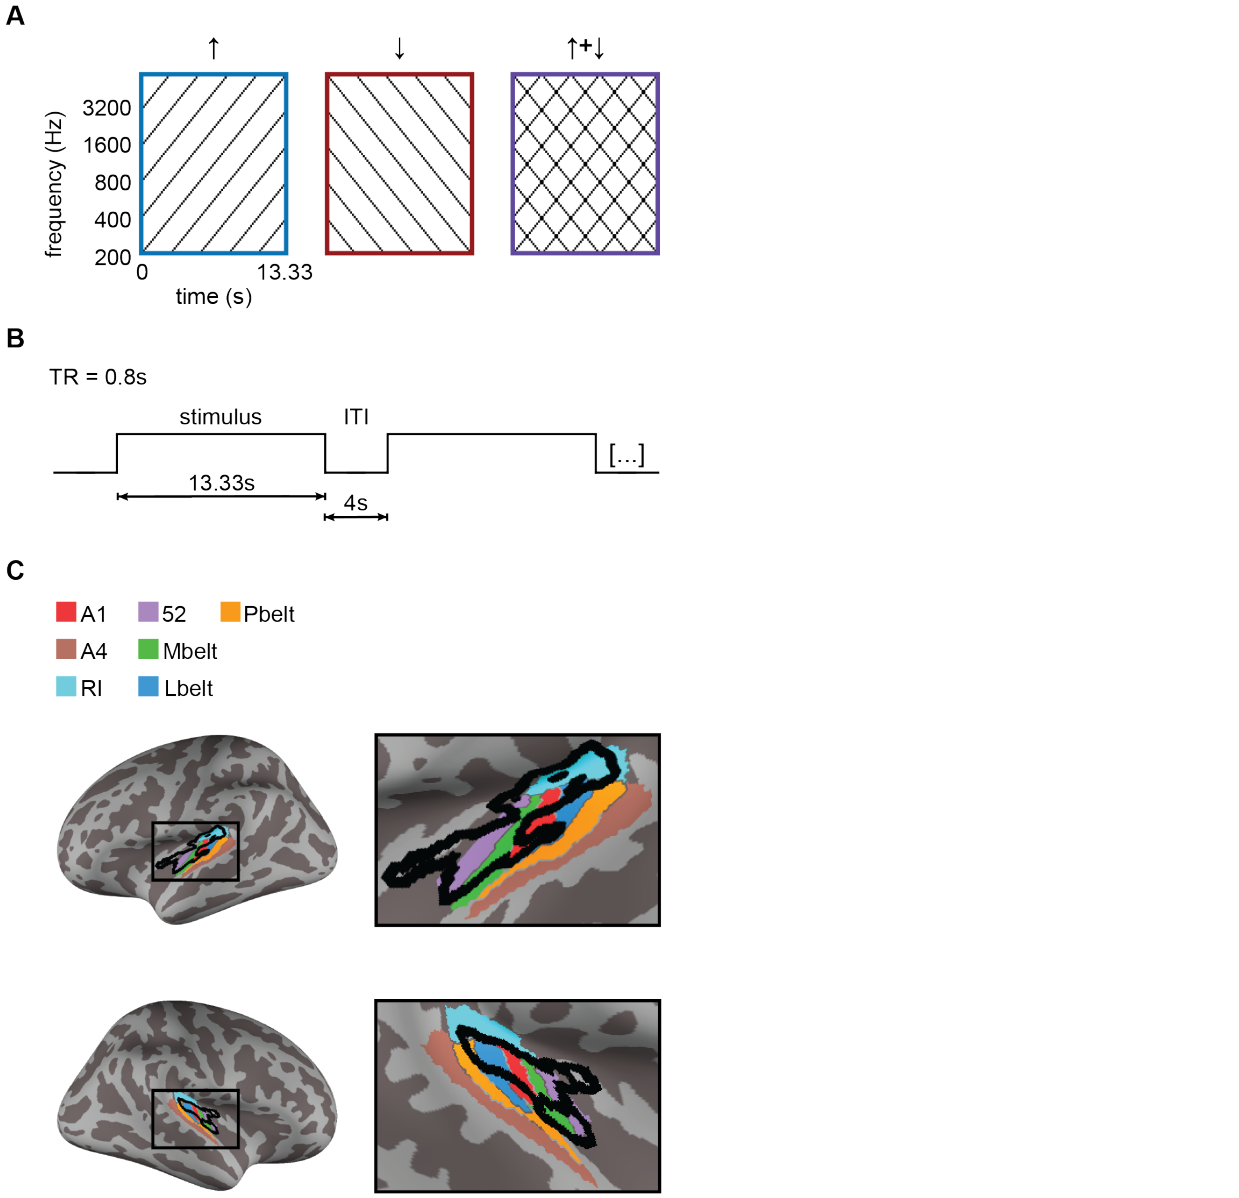


**Supplemental Figure S4.** Further method and result details from fMRI opponency experiment.

A) Depiction of actual stimuli used for the opponency experiment.

B) Time course of fMRI trial structure.

C) Group level analysis showing bilateral regions within auditory cortex that demonstrate significant opponent properties. Black outline reflects significant clusters from **Figure 4H**. Colored patches show cortical regions in accordance with (52). *RI* = retroinsular cortex; *Mbelt* = medial belt of auditory cortex; *Lbelt* = lateral belt of auditory cortex; *Pbelt* = parabelt region.

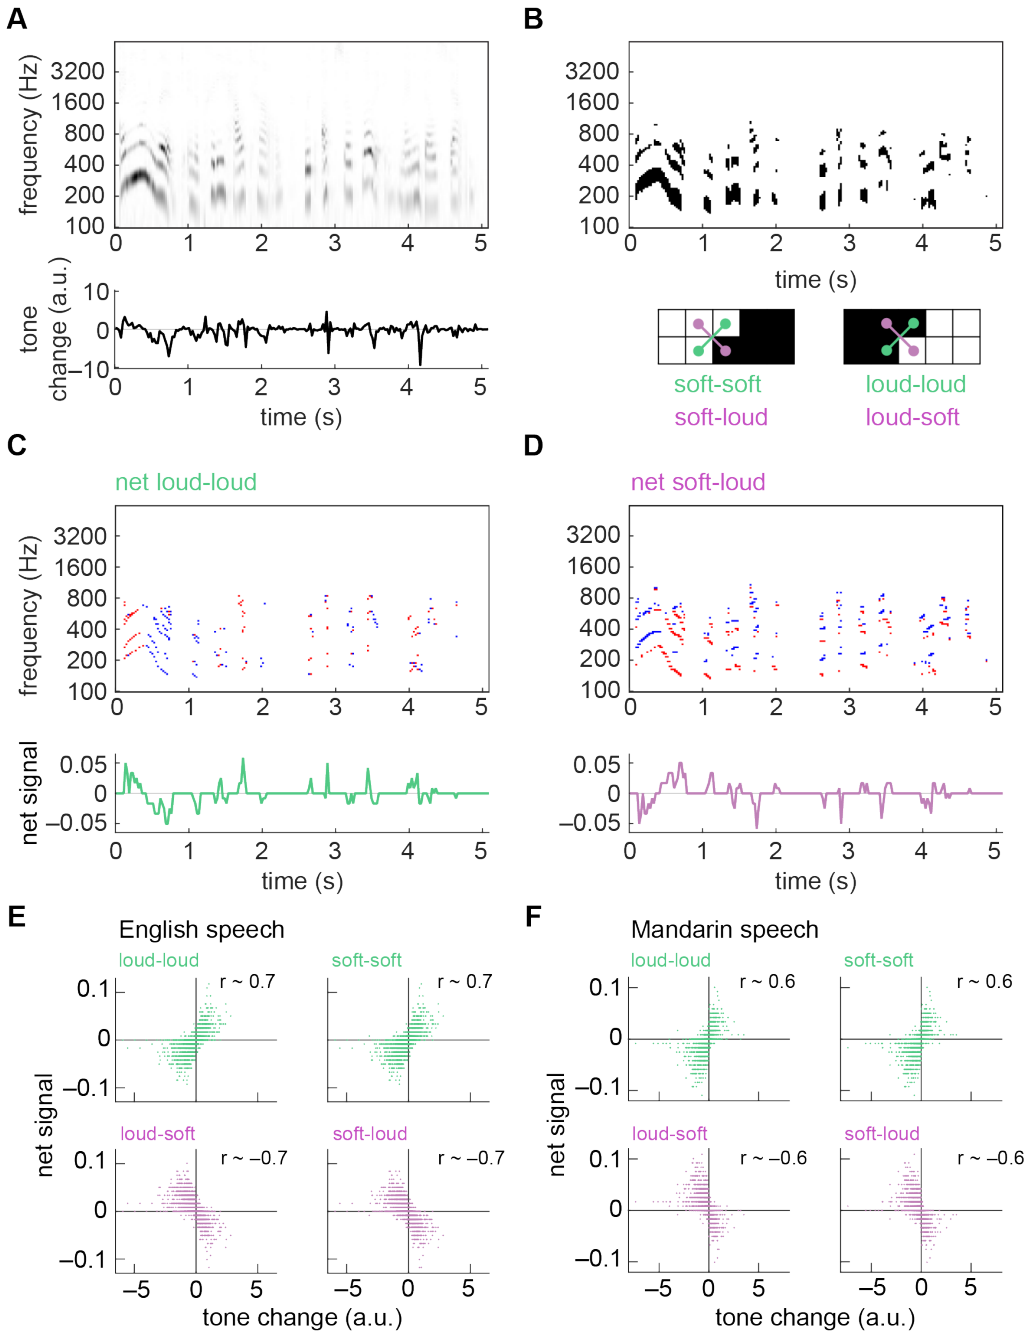


**Figure 5.** Rising and falling tone in spoken language can be detected through both positive and negative pairwise correlations.

A) Spectrogram of voice saying, "Anyone lived in a pretty how town (with up so falling many bells down)" (*top*). Intonation velocity estimate from spectrogram (*bottom*, see **Methods**). Positive tone changes correspond to rising frequencies in the sound.

B) Binarized spectrogram from (A) (*top*). Four distinct loud and soft frequency-time combinations in the binarized spectrogram (*bottom*).

C) Net loud-loud instances at each frequency and time in the binarized spectrogram in (B) (*top*). Red is +1, blue is –1, white is 0. Frequency-averaged net loud-loud signal (*bottom*).

D) Net soft-loud instances at each frequency and time in the binarized spectrogram in (B) (*top*). Red is +1, blue is –1, white is 0. Frequency-averaged net soft-loud signal (*bottom*).
E) Correlations between the tone change estimate at each time and the frequency-averaged net signals for loud-loud, soft-soft, loud-soft, and soft-loud patterns. Data from English speech corpus (71).
F) As in (E) but for Mandarin speech corpus (72).

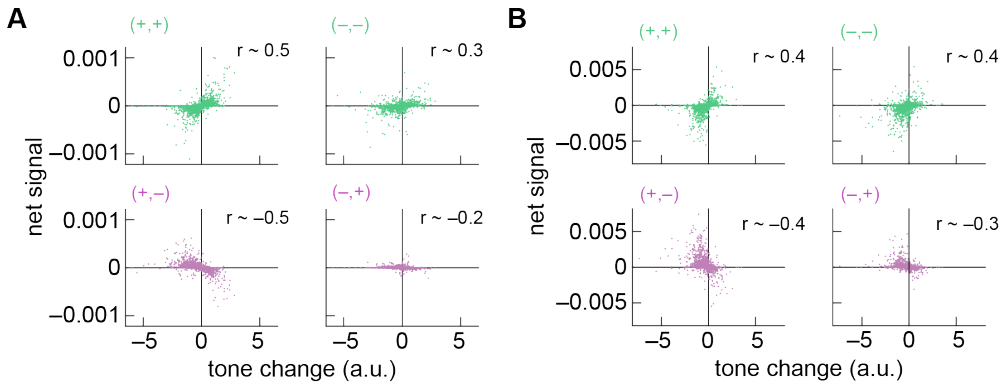


**Supplemental Figure S5.** Multiplicative interactions of amplitude derivatives are informative about intonation direction.

A) Correlations between the tone change estimate at each time and a continuous correlator model using only positive signals (+,+), only negative signals (–,–), and mixtures of the two (+,– and –,+) (see **Methods**). The correlations comprising the net signal were obtained by taking the derivative of the spectrogram amplitude in time, then multiplying derivatives of neighboring frequencies with a time-step delay and subtracting a mirror image product. Signals were rectified before multiplication to obtain the four pairs of multiplied signals, which together add up to a full correlator model. The net signals computed from (+,+) and (–,–) pairs correlated positively with tone change, while the net signals from (+,–) and (–,+) pairs correlated negatively with tone change. Data from English speech corpus (71).

B) As in (A) but with data from Mandarin speech corpus (72).

**Supp. Movie 1.** Demonstration of positive and negative pairwise correlations using ternary correlated noise stimuli, as in Figure 1.

**Supp. Movie 2.** Demonstration of positive and negative pairwise correlations using ternary correlated noise stimuli, analogous to the stimuli in **Supp. Movie 1** but in visual motion detection (27).

## Methods

### *Psychophysical measurements*

All participants (N = 33; 12 female; mean age: 23.3 years, range of 18 years to 32 years) provided informed, written consent in accordance with procedures approved by the Yale University Institutional Review Board. To measure human psychophysical curves (**Figures 1-3**), we recruited participants with self-reported normal hearing from within the university population. Participants were seated in a quiet room, wearing headphones (Model DT 770 PRO, Beyerdynamic, Heilbronn, Germany) to listen to various sound stimuli and make perceptual judgments. The sounds were created in Matlab and presented using Psychtoolbox (73-75) on a Macbook Pro, using its native soundcard. Participants adjusted the volume to a comfortable level, which we estimated to typically be around 60 dB. Each sound was played for 2 seconds, after which participants were cued to judge, to the best of their ability, whether it sounded like a rising or falling tone. To ensure they understood the task, participants went through several example sounds with the researcher before beginning the experiment. Participants usually completed two experiments lasting approximately 15 minutes each. The data was analyzed using custom code written in Matlab. The code to produce the sounds, all anonymized data, and the code used to analyze the data and produce Figure 1-3 are all publicly available at: [GitHub repository here, to be made available on publication].

### *Creating correlated sounds*

We created complex sounds containing multiple frequencies, following the design of visual stimuli that have been informative in that field. To do this, we created a comb of constant carrier frequencies, with frequencies ranging over 6 octaves from 200 Hz to 6400 Hz, with 15 frequencies per octave, equally spaced in log-space. The sampling frequency was chosen to be 20 kHz for all experiments. Each carrier frequency was then multiplied by a slower, time varying envelope, before the frequencies were summed to make the overall waveform for that sound. Mathematically, the sound waveform, $w(t)$, looks like:

$$w(t) = \sum_{i=1}^{N} \theta_i m_i(t) \sin(2\pi f_i t)$$

Where the $f_i$ is the indexed carrier frequencies, $t$ is sampled at 20 kHz, and the value $\theta_i$ was chosen to roughly equalize the perceptual salience of the different frequencies, using the ISO standard 226 at 60 dB. (We note that in various tests in lab, this perceptual salience scaling was not critical for the percepts we measured; since we included it in initial experiments, we included it for all stimuli in this study.) It remains to compute the suite of $m_i(t)$ envelope functions to create each sound. The envelope functions were computed as outlined below. All envelope functions are computed to have non-negative binary or ternary values, and were filtered with a 25 ms low-pass filter in the ternary stimuli (Fig. 1) and at 0.5 ms low-pass filter in the pip stimuli (Figs. 2 and 3) to eliminate sharp transitions. After all waveforms $w(t)$ were created, they were scaled to have a minimum value of –1 and maximum value of +1.

Ternary pairwise correlations (**Figure 1B-D**).

To create sounds with only local, pairwise correlations between specific frequency and time offsets, we followed a protocol used in prior visual experiments (27, 28, 76). Based on informal experiments attempting to optimize our own percepts, we discretized frequencies into 15 notes per octave and time into 1/6 second frames. This change in frequency is similar to the most salient change in frequency in a prior study (37). We then created an initial binary mask in this coarse-time representation, $B_{ij}$, where $i$ indexed the frequency and $j$ the time step in 1/6 second intervals. In each trial, each element of $B$ was chosen from a Bernoulli distribution with probability 0.5, then centered to have values of $\pm 1/2$ instead of 0 and 1. A ternary mask, $M$, was created by the following formula:

$$M_{i,j} = B_{i,j} + PB_{i+d,j+1}$$

The mask is thus the binary matrix added back to itself with a displacement in frequency of $d = \pm 1$ for upward and downward directed correlations. The mask is ternary, with values of 0, 1, and –1. The correlation parity is chosen by $P = \pm 1$, so that the offset matrices are added to create positive correlations and subtracted to create negative correlations. The discrete autocorrelation function of this mask $M$ is equal to:

$$C_{m,n} = \frac{1}{2}\delta_{m,0}\delta_{n,0} + \frac{1}{4}P\left(\delta_{m,d}\delta_{n,1} + \delta_{m,-d}\delta_{n,-1}\right)$$

Where the $\delta_{i,j}$ terms are Kronicker delta functions (see **Fig. S1**). Importantly, the elements in the mask are not deterministically the same or different at the spectrotemporal offset of the correlated displacement, so that spectral patterns vary substantially at each temporal update of the stimulus.

A continuous time expression for the autocorrelation function is available in a prior work describing similar stimuli in vision (28).

The coarse-time matrix $M$ was recentered to have values of 0, 0.5, and 1, then up-sampled to the sampling frequency $F_s$ to create $m_i(t)$ at each frequency. The masks were filtered with a 25 ms low-pass filter to eliminate sharp transitions.

To create the stimuli with varying coherence, we replaced a fraction of mask elements with random ternary stimuli, drawn from the values (0, 0.5, 1) with probabilities (0.25, 0.5, 0.25). The fraction replaced was equal to $(1 - C)$ where $C$ is the coherence value.

Binaural pairwise correlations (**Figure 1D, E**).

To play sounds such that correlations only existed by integrating across the ears, we simply played $B_{i,j}$ in one ear and $PB_{i+d,j+1}$ in the other ear, for the correlations as described above to create the ternary pairwise correlations. To play these binary masks, we created two masks $M_{i,j} = B_{i,j}$ and $M_{i,j} = PB_{i\pm d,j+1}$ to play to the two ears. The matrices were recentered to have values of 0 and 1, then up-sampled to the sampling frequency. The masks were filtered with a 0.5 ms low-pass filter to eliminate sharp transitions.

Correlated pips with time and frequency offsets (**Figure 2**).
To create the correlated pip stimulus, we discretized frequency space into 15 tones per octave. We first initialized our masks $m_i(t)$ to be 0 for all times, sampled at the sampling frequency $F_s$. We then placed initial delta-function pips in a Poisson distribution across all frequencies and times in our sound, at a rate of 4 pips per frequency per second. Positive and negative pips were equally probable, represented by mask values of ±1. We then created a second set of pips offset by the selected change in frequency and delay time, according to the two different correlation types. After imposing the correlations, the overall pip rate became 8 pips per frequency per second. We then convolved this event-trace with a boxcar function with the length of the pip duration to create the mask at $F_s$. Pips had a duration of 40 ms in Figure 2 and 20 ms in Figure S2. Last, the masks were linearly transformed to be between 0 and 1 and filtered with a 0.5 ms low-pass filter to eliminate sharp transitions. The loud values corresponded to values of 1 in the mask, the soft to values of 0, and the background to values of 0.5.

Correlations between loud and soft pips (**Figure 3**).
These stimuli were generated similarly to the correlated pips stimulus above. However, only two thirds of all pips were in correlated pairs of loud-loud, soft-soft, loud-soft, or soft-loud. In the case of the loud-loud correlated pips, the remaining third of pips consisted of randomly placed soft pips. In the case of soft-soft correlated pips, the remaining third of pips consisted of randomly placed loud pips. And in the cases of soft-loud and loud-soft, the remaining third were equally distributed between soft and loud pips. Thus, the four types had equal numbers correlated pairs in each stimulus. The overall rate of pips for all stimuli was 6 pips per frequency per second.

Triplet correlations (**Figure S3**).
We made triplet correlation binary masks, discretized in frequency at time, following prior procedures (39, 41). The frequency was discretized in 15 tones per octave and time was discretized into 1/6 second frames. The frequencies began at 200 Hz and ranged over 5 octaves. The masks $m_i(t)$ were linearly transformed to have values of 0 and 1 and were filtered in time with a 0.5 ms low-pass filter to eliminate sharp transitions.

Rising, falling, and opponent tones (**Figure 5**).
To create the rising, falling, and opponent tones used in our fMRI experiment, we used frequencies discretized into 1/16 octave steps and time discretized into 1/6 second steps. Ascending tones were created from a binary mask equal to an ascending line of time-frequency elements in this discretized space (**Fig. S5**) and descending tones consisted of a descending line of time-frequency elements. The summed ascending plus descending was the sum of the two masks. All masks were filtered in time with a 0.5 ms low-pass filter to eliminate sharp transitions. We switched to 16 steps per octave for this experiment so that the ascending and descending stimuli never played the same frequency simultaneously, making the addition of the stimuli more straightforward.

Code to generate the sounds used in these experiments is available at [GitHub repository on publication].

*Model motion energy unit (**Figure 4**)*

We created a model motion energy unit by convolving a linear filter with a sound spectrogram, then squaring the result. That is:

$$r(t) = \left((f_1 * S_1)(t) + (f_2 * S_2)(t)\right)^2$$

The filters were chosen to be:

$$f_1(t) = \frac{t}{\tau^2} e^{-t/\tau}$$
$$f_2(t) = f_1(t - T)\Theta(t - T)$$

Where $f_2$ is just a time-shifted version of $f_1$ with a time shift of $T$ = 40 ms. The function $\Theta$ is a Heaviside step function. The two filters are applied to adjacent frequencies in the spectrogram, $S_1(t)$ and $S_2(t)$, so that the filter enhances signals directed upward over time.

We computed the mean of $r(t)$ over time to get the mean response for a given stimuli. Stimuli were created to match the correlated pip-style stimuli in Figure 2. The opponent response was computed as

$$r_{\text{opp}}(t) = \left((f_1 * S_1)(t) + (f_2 * S_2)(t)\right)^2 - \left((f_2 * S_1)(t) + (f_1 * S_2)(t)\right)^2$$

The second, negative term is the same as the first term but with the filter flipped in frequency space, so that it corresponds to a downward selective unit. This response was likewise averaged over time to produce the plots in Figure 4.

Matlab code to create **Figures 4B, C** is available at [Github repository on publication].

*Speech analysis*

Spoken language databases were analyzed to ask how spectrotemporal correlations could act as indicators for rising and falling tones in speech. Using Matlab, we first loaded short snippets of speech from two databases: 438 snippets constituting a total of 91 minutes of data from Librispeech, a corpus of read English (71); and 749 snippets constituting a total of 52 minutes of data from Magicdata Mandarin Chinese Read Speech Corpus (72), a corpus of read Mandarin. We computed a spectrogram for each snippet of speech using the Matlab command `spectrogram`; we extracted the spectral amplitude at a resolution of 40 samples per second with no overlap between samples, at 20 evenly spaced frequencies per octave from 100 Hz to 6400 Hz (**Figure 5A**). We estimated the rising/falling intonation change of the sound at each point using the Matlab command `opticalFlowHS`, which uses the Horn-Schunck method (77) to estimate directional local flow (typically optic flow) between frames. We averaged the calculated flow over frequencies to compute an estimate of the frequency "flow" with arbitrary units, which we termed tone change (**Figure 5A**). This method does not make strong assumptions about how changes in speech tone or frequency should be computed. It should work

to extract tone changes from most complex sounds. We then examined estimators of this tone change as follows:

1) To compute binary correlations in frequency and time, we first binarized the spectrogram using Otsu's method (Matlab command `imbinarize`) (78), which maximizes the variance between the binarized time-frequency element amplitudes while minimizing variance within each of the two categories (**Figure 5B**). We made 8 new binary frequency-time data arrays, containing Boolean values at each point in time and frequency, $V_{t,f,\uparrow,\pm,\pm} \coloneqq \left(\{A_{t,f}, A_{t+1,f+1}\} = \{\pm 1, \pm 1\}\right)$ and an equivalent one for downward directed volume patterns. These matrices are records of the existence of each pattern of sound intensity at each time and frequency. From these, we computed the net signal of each pattern at each frequency by subtracting the downward directed matrix from the upward directed one. We last found the mean net signal over all frequencies for each pattern (**Figure 5C, D**). We computed the correlation between these mean net signals at each time point with the calculated upward or downward flow velocity (**Figure 5E, F**). Note that the sum of these net pattern signals sum to 0 over the four different patterns (±,±), so that the 4 signals are not independent.
2) To generate non-binarized correlation plots, we first linearly filtered the spectrogram amplitudes, $A_{t,f}$, to take temporal derivatives: $F_{t,f} = A_{t,f} - A_{t-1,f}$. We then used these derivatives, $F_{t,f}$, which have positive and negative values, as inputs to a Hassenstein-Reichardt correlator model (Hassenstein and Reichardt 1956, Fitzgerald and Clark 2015). We then computed the net (+,+) correlations, for instance, as $N_{t,f,+,+} = \left[F_{t,f}\right]_+ \left[F_{t+1,f+1}\right]_+ - \left[F_{t+1,f}\right]_+ \left[F_{t,f+1}\right]_+$, where $[x]_+ = x$ when $x > 0$ and $[x]_+ = 0$ otherwise. A similar process computed the net (–,–), (+,–) and (–,+) correlations. We averaged these signals over frequency to obtain a single indicator of velocity at each point in time. These indicators were then correlated with the estimated tone change of the sound snippet at that point in time (**Figure S5**).

Code to analyze the spoken language databases and produce the panels in Figure 5 is available at [GitHub repository].

*fMRI recordings and analysis*

Whole-brain imaging was performed at the Brain Imaging Center at Yale University, on a Siemens 3 T Prisma MRI scanner using a 32-channel head coil. Functional data were acquired with a gradient-echo echoplanar pulse sequence (TR = 0.80 s, TE = 30 ms, flip angle = 52˚, voxel size = 2.4 mm × 2.4 mm × 2.4 mm, MB acc. factor = 6). T1-weighted MP-RAGE anatomical images were collected as well (TR = 2.5 s, TE = 2.0 ms, flip angle = 8˚, 208 slices, voxel size = 1.0 mm isotropic). Functional imaging in our sample (N=5; 1 female; mean age: 26.2 years; authors PAV and SDM were participants in the fMRI study) was performed in ~5-minute runs, with the total number of functional runs per participant ranging from 3-5. Fifteen auditory stimuli were presented per run in an event-related design (5 each of three stimulus types: rising, falling, and summed). Each stimulus lasted for 13.33 s, separated by an inter-trial interval (ITI) of 4 s. The order of the three stimulus types was randomized in each run. Participants passively listened to the tones and were not required to render any responses. MRI-

optimized noise-canceling headphones (Optoacoustics OptoACTIVE III) were used to limit effects of background scanner noise and the noise-cancelling software was trained on the EPI sequence sound features before each session using a brief calibration run.

The fMRI-Prep toolbox was used for preprocessing (79). The anatomical image was corrected for intensity non-uniformity (INU) with N4BiasFieldCorrection (80) and used as T1w-reference. The T1w-reference was then skull-stripped with a Nipype implementation of the antsBrainExtraction.sh workflow in ANTs, and tissue segmentation of cerebrospinal fluid (CSF), white-matter (WM), and gray-matter (GM) was performed on the brain-extracted T1w using FFAST (FSL 6.0.5) (81). Volume-based spatial normalization to standard (MNI) space was performed through nonlinear registration with antsRegistration (ANTs 2.3.3). For each of the BOLD runs, a reference volume and its skull-stripped version were generated using a custom methodology of fMRIPrep. Head-motion parameters were estimated using MCFLIRT (FSL 6.0.5) (82) and BOLD time-series were resampled into native space by applying the transforms to correct for head-motion, and the BOLD reference was co-registered to the anatomical reference using mri_coreg (FreeSurfer) followed by FLIRT. Co-registration was configured with 6 DOF. Several confounding time-series were calculated based on the preprocessed BOLD: framewise displacement (FD), DVARS and three region-wise global signals. The BOLD time-series were resampled into standard space, and volumetric resamplings were performed using ANTs.

Our main analyses involved constructing general linear models (GLMs) to quantify the effects of the three stimulus types within auditory cortex. GLM analyses were performed using Nilearn (83). Confound regressors of no interest (generated using fMRIPrep, see above) were entered into each GLM. These included six standard motion regressors, the framewise displacement time course, and white matter and global signal time courses. Each stimulus type (rising, falling, and summed) was modeled using boxcar regressors over the entire stimulus presentation phase (13.33 s) of the relevant trials, and was convolved with the canonical double-gamma hemodynamic response function. The main contrast of interest at the group and individual levels compared BOLD responses to the non-summed directional stimuli (i.e., rising and falling) to the summed stimuli (i.e., superimposed rising + falling). The contrast was designed to highlight deviations from a null hypothesis of equivalent responses between directional and opponent stimuli. Individual subject runs were combined in a fixed effects analysis and then brought to the group level for mixed-effect analyses, where we controlled the false positive rate at $p < 0.05$ with a cluster-forming threshold of 20 voxels. Critically, individual-level results for all subjects were also analyzed and displayed, using the same thresholding parameters. All contrasts were performed within an *a priori* anatomical mask that consisted of any voxels crossing the 50% probability threshold within a combined bilateral probabilistic atlas (Harvard-Oxford) that included both the STG and Heschel's gyrus. Individual and group results were projected onto the standard (MNI) cortical surface (FreeSurfer) for visualization.

A simple control analysis was also performed to ensure that the non-summed > summed results were not driven by a single non-summed stimulus (e.g., rising or falling) having a proportionally larger response, but rather by symmetric responses to the rising and falling stimuli. To perform this control analysis, we first extracted individualized regions of interest (ROIs) from the non-summed > summed contrast (using the threshold described above), and then extracted average

beta values within that ROI for each stimulus type. We note that while this was of course not an unbiased ROI relative to the hypothesis that non-summed stimuli would on average show stronger activity than summed, it was unbiased relative to the hypothesis of symmetric responses to rising versus falling tones.

*Opponency implies a symmetry in responses with opposite correlations in opposite directions*

The motion energy model uses pairwise correlations to extract motion information from input stimuli and seems to accurately represent important aspects of cellular physiology (6). In the motion energy model, stimuli over space and time, $S(x,t)$, are convolved with a space-time oriented linear filter, $H(x,t)$. (In this section, we will derive results in space, but a frequency variable $f$ could substitute for $x$ and this approach would apply sound intensity over frequency rather than light intensity over space.) The result of the convolution is squared to obtain a response:

$$r(x,t) = \left( \iint dx'dt'H(x',t')S(x-x',t-t') \right)^2$$

This response is stronger, on average, to stimuli with motion in the preferred direction than in the null direction. The preferred direction corresponds to the orientation of the filter $H$ in space time, which amplifies signals when the motion direction aligns with the filter orientation. When the response is averaged over time and space, it yields a pleasing form in Fourier space, such that the mean response is the dot product of the stimulus power with a weighting function (6):

$$\langle r \rangle = \iint dkd\omega \left|\tilde{H}(k,\omega)\right|^2 \left|\tilde{S}(k,\omega)\right|^2$$

Where $\tilde{H}$ and $\tilde{S}$ are the Fourier transforms of $H$ and $S$. Therefore, to understand responses of this model, it is useful to compute the power spectrum of the stimulus.

For a random dot kinetogram in which the dots are displaced by $\Delta x$ in space and $\Delta t$ in time, the covariance density, $C$, of the stimulus is a function of the offsets in time and space, $x$ and $t$:

$$C(x,t) = \beta\delta(x,t) + \alpha\delta(x-\Delta x, t-\Delta t) + \alpha\delta(x+\Delta x, t+\Delta t)$$

Where the first term is the stimulus autocovariance and the remaining two terms correspond to correlations in the stimulus at offsets of $(\Delta x, \Delta t)$ and $(-\Delta x, -\Delta t)$. For random dot kinetograms, $\beta < 1$ and $\alpha$ can take on positive or negative values for positively and negative correlated random dot kinetograms. This derivation is in continuous space, using Dirac delta function correlations; a similar result with discrete time and frequencies was found earlier in the methods for the ternary stimuli. The power spectrum of the stimulus is the Fourier transform of this covariance function:

$$\left|\tilde{S}(k,\omega)\right|^2 = \iint dxdte^{ikx}e^{i\omega t}C(x,t) = \beta + \alpha\cos(\omega\Delta t + k\Delta x)$$

The power is highest/lowest along lines of constant phase in cosine, or when $\omega\Delta t + k\Delta x = n\pi$. When the $\alpha$ is negative, for negative correlation stimuli, this effectively changes the phase of the cosine by 180 degrees. The motion energy model says the mean response to such a stimulus, for a unit with filter $H$, is:

$$\langle r \rangle = \iint dk d\omega \left|\widetilde{H}(k,\omega)\right|^2 (\beta + \alpha \cos(\omega\Delta t + k\Delta x))$$

This is the type of curve shown in **Figure 4B**, in which there is a baseline response determined by $\beta$ and the integral of $\left|\widetilde{H}(k,\omega)\right|^2$. There is a modulatory term that depends on $\alpha$ and the dot product of $\left|\widetilde{H}(k,\omega)\right|^2$ with $\cos(\omega\Delta t + k\Delta x)$, which gives the modulation a directional tuning. This form means that the modulation inverts when the sign of the correlation (sign of $\alpha$) inverts. If there is a peak response to a stimulus with correlation $\alpha$ at a specific $\Delta t$ and $\Delta x$, then the peak will be equal and opposite when $\alpha$ is inverted. Importantly, however, the peak is not the same when the direction of the stimulus is inverted, that is when $\Delta x \rightarrow -\Delta x$.

However, if we compute an opponent response, in which we subtract the response with one filter orientation from the response with the opposite filter orientation (inverting the $k$ in the Fourier domain), then we find:

$$\langle r_{opp} \rangle = \iint dk d\omega \left(\left|\widetilde{H}(k,\omega)\right|^2 - \left|\widetilde{H}(-k,\omega)\right|^2\right)\left(\beta + \alpha \cos(\omega\Delta t + k\Delta x)\right)$$

$$\langle r_{opp} \rangle = \alpha \iint dk d\omega \left(\left|\widetilde{H}(k,\omega)\right|^2 - \left|\widetilde{H}(-k,\omega)\right|^2\right)\left(\cos(\omega\Delta t + k\Delta x)\right)$$

Here, we see that the opponent subtraction causes the $\beta$ term to drop out entirely so that the remaining term is just proportional to $\alpha$, the correlation in the stimulus. The mean opponent response can be computed for correlation stimuli with parameters $\alpha$, $\Delta t$, and $\Delta x$: $\langle r_{opp}(\alpha, \Delta t, \Delta x)\rangle$. Because of the directional opponency, the response inverts when the stimulus is reversed in space:

$$\langle r_{opp}(\alpha, \Delta t, \Delta x)\rangle = -\langle r_{opp}(\alpha, \Delta t, -\Delta x)\rangle$$

And because of the proportionality with the correlation, the response inverts when the stimulus correlation is inverted:

$$\langle r_{opp}(\alpha, \Delta t, \Delta x)\rangle = -\langle r_{opp}(-\alpha, \Delta t, \Delta x)\rangle$$

Therefore, for an opponent signal, inverting the correlation is equivalent to inverting the direction of the signal:

$$\langle r_{opp}(-\alpha, \Delta t, \Delta x)\rangle = \langle r_{opp}(\alpha, \Delta t, -\Delta x)\rangle$$

For any set of filters, as long as they are opponently subtracted, inverting the sign of the correlation is identical to inverting the direction of the stimulus, when computing the

spatiotemporal average response. So when stimuli can be generated that have autocovariance structures like those in the ternary scintillator (**Fig. 1**) or in a random dot kinetogram (**Fig. 2**), if the computation is based on pairwise correlations and is opponent, the equations above show that the response will always be inverted when the stimulus correlation is inverted, and always be equivalent to inverting the direction of the stimulus. Therefore, opponency implies the sort of inversion symmetries we observed in our data, where inverting the correlation sign generates percepts with the same tuning as inverting the direction of the stimulus (**Fig. 4D, E**, but also visible in **Figs. 1-3**). Opponency also implies the sort of consistent symmetries between positive and negative correlation stimuli observed in human motion perception (36). We note that it is also possible to achieve this kind of symmetry using precisely defined filters that lead to opponent properties in single units, without a subtractive step (64).

## References


1. D. Hirst, A. Di Cristo, Intonation systems. *A survey of Twenty Languages* (1998).

2. J. Gandour, Tone perception in Far Eastern languages. *Journal of phonetics* **11**, 149-175 (1983).

3. M. J. W. Yip, *Tone* (Cambridge University Press, 2002).

4. Z.-L. Lu, L. A. Lesmes, G. Sperling, Perceptual motion standstill in rapidly moving chromatic displays. *Proc. Natl. Acad. Sci. USA* **96**, 15374-15379 (1999).

5. B. Hassenstein, W. Reichardt, Systemtheoretische Analyse der Zeit-, Reihenfolgen-und Vorzeichenauswertung bei der Bewegungsperzeption des Rüsselkäfers Chlorophanus. *Zeits. Naturforsch.* **11**, 513–524 (1956).

6. E. Adelson, J. Bergen, Spatiotemporal energy models for the perception of motion. *JOSA A* **2**, 284-299 (1985).

7. S. M. Anstis, B. J. Rogers, Illusory reversal of visual depth and movement during changes of contrast. *Vision Res.* **15**, 957-IN956 (1975).

8. Z. L. Lu, G. Sperling, Three-systems theory of human visual motion perception: review and update. *JOSA A* **18**, 2331-2370 (2001).

9. Z.-L. Lu, G. Sperling, The functional architecture of human visual motion perception. *Vision Res.* **35**, 2697-2722 (1995).

10. J. H. McDermott, The cocktail party problem. *Curr. Biol.* **19**, R1024-R1027 (2009).

11. E. C. Cherry, Some experiments on the recognition of speech, with one and with two ears. *The Journal of the acoustical society of America* **25**, 975-979 (1953).

12. J. K. Bizley, Y. E. Cohen, The what, where and how of auditory-object perception. *Nat. Rev. Neurosci.* **14**, 693-707 (2013).

13. A. S. Bregman, J. Campbell, Primary auditory stream segregation and perception of order in rapid sequences of tones. *Journal of experimental psychology* **89**, 244 (1971).

14. V. Ciocca, A. S. Bregman, Perceived continuity of gliding and steady-state tones through interrupting noise. *Perception & Psychophysics* **42**, 476-484 (1987).

15. J. M. Sinnott, R. N. Aslin, Frequency and intensity discrimination in human infants and adults. *The Journal of the Acoustical Society of America* **78**, 1986-1992 (1985).

16. L. Demany, C. Ramos, On the binding of successive sounds: Perceiving shifts in nonperceived pitches. *The Journal of the Acoustical Society of America* **117**, 833-841 (2005).

17. R. N. Aslin, Discrimination of frequency transitions by human infants. *The Journal of the Acoustical Society of America* **86**, 582-590 (1989).

18. K. Siedenburg, J. Graves, D. Pressnitzer, A unitary model of auditory frequency change perception. *PLoS Comp. Biol.* **19**, e1010307 (2023).

19. L. Demany, C. Semal, Automatic frequency-shift detection in the auditory system: A review of psychophysical findings. *Neuroscience* **389**, 30-40 (2018).

20. H.-W. Lu, P. H. Smith, P. X. Joris, Mammalian octopus cells are direction selective to frequency sweeps by excitatory synaptic sequence detection. *Proc. Natl. Acad. Sci. USA* **119**, e2203748119 (2022).

21. R. I. Kuo, G. K. Wu, The generation of direction selectivity in the auditory system. *Neuron* **73**, 1016-1027 (2012).

22. C.-q. Ye, M.-m. Poo, Y. Dan, X.-h. Zhang, Synaptic mechanisms of direction selectivity in primary auditory cortex. *J. Neurosci.* **30**, 1861-1868 (2010).

23. S. Andoni, N. Li, G. D. Pollak, Spectrotemporal receptive fields in the inferior colliculus revealing selectivity for spectral motion in conspecific vocalizations. *J. Neurosci.* **27**, 4882-4893 (2007).

24. R. C. DeCharms, D. T. Blake, M. M. Merzenich, Optimizing sound features for cortical neurons. *science* **280**, 1439-1444 (1998).

25. F. E. Theunissen, K. Sen, A. J. Doupe, Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds. *J. Neurosci.* **20**, 2315-2331 (2000).

26. L. M. Miller, M. A. Escabí, H. L. Read, C. E. Schreiner, Spectrotemporal receptive fields in the lemniscal auditory thalamus and cortex. *J. Neurophysiol.* **87**, 516-527 (2002).

27. E. Salazar-Gatzimas *et al.*, Direct measurement of correlation responses in Drosophila elementary motion detectors reveals fast timescale tuning. *Neuron* **92**, 227-239 (2016).

28. S. Roy, R. d. R. van Steveninck, Bilocal visual noise as a probe of wide field motion computation. *J. Vis.* **16**, 8-8 (2016).

29. M. J. McPherson, J. H. McDermott, Relative pitch representations and invariance to timbre. *Cognition* **232**, 105327 (2023).

30. M. J. McPherson, J. H. McDermott, Diversity in pitch perception revealed by task dependence. *Nature human behaviour* **2**, 52-66 (2018).

31. D. A. Clark, L. Bursztyn, M. A. Horowitz, M. J. Schnitzer, T. R. Clandinin, Defining the computational structure of the motion detector in Drosophila. *Neuron* **70**, 1165-1177 (2011).

983 32. M. B. Orger, M. C. Smear, S. M. Anstis, H. Baier, Perception of Fourier and non-Fourier
984 motion by larval zebrafish. *Nat. Neurosci.* **3**, 1128-1133 (2000).

985 33. B. Krekelberg, T. Albright, Motion mechanisms in macaque MT. *J. Neurophysiol.* **93**,
986 2908 (2005).

987 34. J. B. Kelly, P. W. Judge, Binaural organization of primary auditory cortex in the ferret
988 (Mustela putorius). *J. Neurophysiol.* **71**, 904-913 (1994).

989 35. K. H. Britten, M. N. Shadlen, W. T. Newsome, J. A. Movshon, The analysis of visual
990 motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* **12**,
991 4745-4765 (1992).

992 36. R. Bours, M. Kroes, M. Lankheet, Sensitivity for reverse-phi motion. *Vision Res.* **49**, 1-9
993 (2009).

994 37. L. Demany, D. Pressnitzer, C. Semal, Tuning properties of the auditory frequency-shift
995 detectors. *The Journal of the Acoustical Society of America* **126**, 1342-1348 (2009).

996 38. J. Allik, E. Dzhafarov, A. Houtsma, J. Ross, N. Versfeld, Pitch motion with random
997 chord sequences. *Perception & Psychophysics* **46**, 513-527 (1989).

998 39. Q. Hu, J. D. Victor, A set of high-order spatiotemporal stimuli that elicit motion and
999 reverse-phi percepts. *J. Vis.* **10** (2010).

1000 40. J. E. Fitzgerald, A. Y. Katsov, T. R. Clandinin, M. J. Schnitzer, Symmetries in stimulus
1001 statistics shape the form of visual motion estimators. *Proc. Natl. Acad. Sci. USA* **108**,
1002 12909-12914 (2011).

1003 41. D. A. Clark *et al.*, Flies and humans share a motion estimation strategy that exploits
1004 natural scene statistics. *Nat. Neurosci.* **17**, 296-303 (2014).

1005 42. J. Chen, H. B. Mandel, J. E. Fitzgerald, D. A. Clark, Motion estimates in flies include
1006 higher-order correlations that cancel noise induced by the structure of natural scenes. *In*
1007 *preparation* (2018).

1008 43. T. Yildizoglu, C. Riegler, J. E. Fitzgerald, R. Portugues, A Neural Representation of
1009 Naturalistic Motion-Guided Behavior in the Zebrafish Brain. *Curr. Biol.* (2020).

1010 44. G. J. Brown, M. Cooke, Computational auditory scene analysis. *Computer Speech &*
1011 *Language* **8**, 297-336 (1994).

1012 45. R. J. Snowden, S. Treue, R. G. Erickson, R. A. Andersen, The response of area MT and
1013 V1 neurons to transparent motion. *J. Neurosci.* **11**, 2768-2785 (1991).

1014 46. N. Qian, R. A. Andersen, Transparent motion perception as detection of unbalanced
1015 motion signals. II. Physiology. *J. Neurosci.* **14**, 7367-7380 (1994).

47. C. D. Salzman, C. M. Murasugi, K. H. Britten, W. T. Newsome, Microstimulation in visual area MT: effects on direction discrimination performance. *J. Neurosci.* **12**, 2331-2355 (1992).

48. A. S. Mauss *et al.*, Neural circuit to integrate opposing motions in the visual field. *Cell* **162**, 351-362 (2015).

49. S. Andoni, G. D. Pollak, Selectivity for spectral motion as a neural computation for encoding natural communication signals in bat inferior colliculus. *J. Neurosci.* **31**, 16529-16540 (2011).

50. D. J. Heeger, G. M. Boynton, J. B. Demb, E. Seidemann, W. T. Newsome, Motion opponency in visual cortex. *J. Neurosci.* **19**, 7162-7174 (1999).

51. B. Tian, J. P. Rauschecker, Processing of frequency-modulated sounds in the lateral auditory belt cortex of the rhesus monkey. *J. Neurophysiol.* **92**, 2993-3013 (2004).

52. M. F. Glasser *et al.*, A multi-modal parcellation of human cerebral cortex. *Nature* **536**, 171-178 (2016).

53. N. C. Singh, F. E. Theunissen, Modulation spectra of natural sounds and ethological theories of auditory processing. *The Journal of the Acoustical Society of America* **114**, 3394-3411 (2003).

54. R. N. Shepard, Circularity in judgments of relative pitch. *The journal of the acoustical society of America* **36**, 2346-2353 (1964).

55. M. J. McPherson, J. H. McDermott, Time-dependent discrimination advantages for harmonic sounds suggest efficient coding for memory. *Proc. Natl. Acad. Sci. USA* **117**, 32169-32180 (2020).

56. H. G. Yi, M. K. Leonard, E. F. Chang, The encoding of speech sounds in the superior temporal gyrus. *Neuron* **102**, 1096-1110 (2019).

57. G. Hickok, D. Poeppel, The cortical organization of speech processing. *Nat. Rev. Neurosci.* **8**, 393-402 (2007).

58. W. A. de Heer, A. G. Huth, T. L. Griffiths, J. L. Gallant, F. E. Theunissen, The hierarchical cortical organization of human speech processing. *J. Neurosci.* **37**, 6539-6557 (2017).

59. Y. Li, C. Tang, J. Lu, J. Wu, E. F. Chang, Human cortical encoding of pitch in tonal and non-tonal languages. *Nature communications* **12**, 1161 (2021).

60. C. Tang, L. Hamilton, E. Chang, Intonational speech prosody encoding in the human auditory cortex. *Science* **357**, 797-801 (2017).

1049 61. I. S. Johnsrude, V. B. Penhune, R. J. Zatorre, Functional specificity in the right human
1050 auditory cortex for perceiving pitch direction. *Brain* **123**, 155-163 (2000).

1051 62. J. A. Zavatone-Veth, B. A. Badwan, D. A. Clark, A minimal synaptic model for direction
1052 selective neurons in Drosophila. *J. Vis.* **20**, 2-2 (2020).

1053 63. C.-H. Mo, C. Koch, Modeling reverse-phi motion-selective neurons in cortex: double
1054 synaptic-veto mechanism. *Neural Comput.* **15**, 735-759 (2003).

1055 64. B. A. Badwan, M. S. Creamer, J. A. Zavatone-Veth, D. A. Clark, Dynamic nonlinearities
1056 enable direction opponency in Drosophila elementary motion detectors. *Nat. Neurosci.*
1057 **22**, 1318-1326 (2019).

1058 65. D. A. Clark, J. B. Demb, Parallel computations in insect and mammalian visual motion
1059 processing. *Curr. Biol.* **26**, R1062-R1072 (2016).

1060 66. J. R. Sanes, S. L. Zipursky, Design principles of insect and vertebrate visual systems.
1061 *Neuron* **66**, 15 (2010).

1062 67. A. Borst, M. Helmstaedter, Common circuit design in fly and mammalian motion vision.
1063 *Nat. Neurosci.* **18**, 1067-1076 (2015).

1064 68. N. Kadakia *et al.*, Odour motion sensing enhances navigation of complex plumes. *Nature*
1065 **611**, 754-761 (2022).

1066 69. J. E. Fitzgerald, D. A. Clark, Nonlinear circuits for naturalistic visual motion estimation.
1067 *eLife*, e09123 (2015).

1068 70. E. Salazar-Gatzimas, M. Agrochao, J. E. Fitzgerald, D. A. Clark, The Neuronal Basis of
1069 an Illusory Motion Percept Is Explained by Decorrelation of Parallel Motion Pathways.
1070 *Curr. Biol.* **28**, 3748-3762 (2018).

1071 71. V. Panayotov, G. Chen, D. Povey, S. Khudanpur (2015) Librispeech: an asr corpus based
1072 on public domain audio books. in *2015 IEEE international conference on acoustics,*
1073 *speech and signal processing (ICASSP)* (IEEE), pp 5206-5210.

1074 72. MagicData, Retrieved from https://www.openslr.org/123/ on March 28, 2023. (2019).

1075 73. M. Kleiner *et al.*, What's new in Psychtoolbox-3. *Perception* **36**, 1 (2007).

1076 74. D. H. Brainard, The psychophysics toolbox. *Spatial vision* **10**, 433-436 (1997).

1077 75. D. G. Pelli, The VideoToolbox software for visual psychophysics: Transforming numbers
1078 into movies. *Spatial vision* **10**, 437-442 (1997).

1079 76. R. van Steveninck, W. Bialek, M. Potters, R. Carlson, G. Lewen (1996) Adaptive
1080 movement computation by the blowfly visual system. in *Natural & Artificial Parallel*
1081 *Computation: Proceedings of the Fifth NEC Research Symposium* (SIAM), p 21.

77. B. K. Horn, B. G. Schunck, Determining optical flow. *Artif. Intell.* **17**, 185-203 (1981).

78. N. Otsu, A threshold selection method from gray-level histograms. *Automatica* **11**, 23-27 (1975).

79. O. Esteban *et al.*, fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nat. Methods* **16**, 111-116 (2019).

80. N. J. Tustison *et al.*, N4ITK: improved N3 bias correction. *IEEE Trans. Med. Imaging* **29**, 1310-1320 (2010).

81. Y. Zhang, M. Brady, S. Smith, Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans. Med. Imaging* **20**, 45-57 (2001).

82. M. Jenkinson, P. Bannister, M. Brady, S. Smith, Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage* **17**, 825-841 (2002).

83. A. Abraham *et al.*, Machine learning for neuroimaging with scikit-learn. *Frontiers in neuroinformatics* **8**, 14 (2014).