Analysis of temporal structure in sound by the human brain

Timothy D. Griffiths[1,2,3], Christian Büchel[2], Richard S.J. Frackowiak[2] and Roy D. Patterson[1]

[1] *Centre for the Neural Basis of Hearing, Physiology Department, University of Cambridge, Downing Street, Cambridge, CB2 3EG, UK*

[2] *Wellcome Department of Cognitive Neurology, Institute of Neurology, 12 Queen Square, London, WC1N 3BG, UK*

[3] *Department of Physiological Sciences, Newcastle University Medical School, Newcastle upon Tyne, NE2 4HH,UK*

Correspondence should be addressed to T.D.G. (t.griffiths@fil.ion.ucl.ac.uk)

For over a century, models of pitch perception have been based on the frequency composition of the sound. Pitch phenomena can also be explained, however, in terms of the time structure, or temporal regularity, of sounds. To locate the mechanism for the detection of temporal regularity in humans, we used functional imaging and a 'delay-and-add' noise, which activates all frequency regions uniformly, like noise, but which nevertheless produces strong pitch perceptions and tuneful melodies. This stimulus has temporal regularity that can be systematically altered. We found that the activity of primary auditory cortex increased with the regularity of the sound. Moreover, a melody composed of delay-and-add 'notes' produced a distinct pattern of activation in two areas of the temporal lobe distinct from primary auditory cortex. These results suggest a hierarchical analysis of time structure in the human brain.

SPECTRAL MODELS OF PITCH PERCEPTION

Traditionally, pitch perception is explained in terms of spectral analysis, and in particular the spectral analysis performed by the inner ear, or cochlea, where the frequency components of a sound produce activity in different places along the cochlea in accordance with their frequencies[1]. Sinusoidal sounds are the easiest to explain; they contain only one frequency component and produce a pitch corresponding to that frequency. Musical notes and the vowels of speech are composed of sets of harmonically related sinusoids, and they produce a pitch corresponding to the fundamental of the harmonic series. A short sample of the neural activity pattern[2] (NAP) produced by the cochlea in response to a musical note composed of all the harmonics of 62.5 Hz is shown in Fig. 1a; the pitch of the note is 62.5 Hz, which is close to the C two octaves below middle C on the keyboard. The ordinate in the figure is place along the length of the cochlea, and it is essentially a logarithmic frequency scale. The abscissa is time in milliseconds, so the figure shows the temporal microstructure of the neural response. The repetition rate of the neural pattern corresponds to the pitch of the sound. Figure 1b shows the 'auditory spectrum' of the sound[3] at three points in time: 200, 300 and 400 ms after the onset of the note. The auditory spectrum is the running temporal average of activity along the cochlea. In the low-frequency region, there are peaks corresponding to the first eight harmonics of the note. As frequency increases, however, the resolution of the spectral analysis decreases, and the auditory spectrum becomes smooth above the eighth harmonic. These auditory spectra show that the first eight harmonics in the series are resolved (that is, they produce individual peaks in the auditory spectrum). In spectral models of pitch perception[4], the low-frequency peaks are identified by the brain, and their frequency spacing is calculated to extract the pitch of the sound.

TEMPORAL MODELS OF PITCH PERCEPTION

The problem with the traditional model of pitch is that the sound in Fig. 1 still produces the same pitch, (albeit with reduced strength or salience) when the sound is filtered to remove all the low-frequency components that produce individual peaks in the auditory spectrum. In this case, the lower half of the NAP and auditory spectrum are empty, and there are no spectral peaks with which to calculate the pitch. In spectral models, the temporal information in the NAP is averaged out in the construction of the auditory spectrum. In contrast, single-unit recordings show that temporal regularity in a stimulus is followed with extreme (sub-millisecond) accuracy by auditory nerve impulses[5], and this has prompted the development of temporal models that associate pitch with the dominant time interval in the NAP[6,7]. To emphasize the importance of time-interval information, the modelers developed a stimulus in which a noise is delayed and added to itself repeatedly[6,7]. It is perceived as a tone with a pitch corresponding to the reciprocal of the delay, mixed with a background noise, a bit like a cracked bassoon playing in the wind. The strength of the pitch increases with the number of iterations of the delay-and-add process. Figure 1c and d shows the NAP and auditory spectra produced when the delay is 16 ms and there are eight iterations of the delay-and-add process. The sound has the same pitch as the sound represented in the upper panel, and the pitch is strong relative to the background noise; moreover, the pitch remains when the stimulus is high-pass filtered. For this stimulus, neither the auditory spectrum nor the NAP provides a good explanation of the perceived pitch. There are no stable, harmonically related peaks in the auditory spectra, and although the NAP does contain an excess of time intervals at 16 ms, it is not apparent in the NAP because the envelope does not repeat regularly in time.

In temporal models of hearing, it is assumed that the pitch is extracted with autocorrelation, in which the neural pattern is cor-

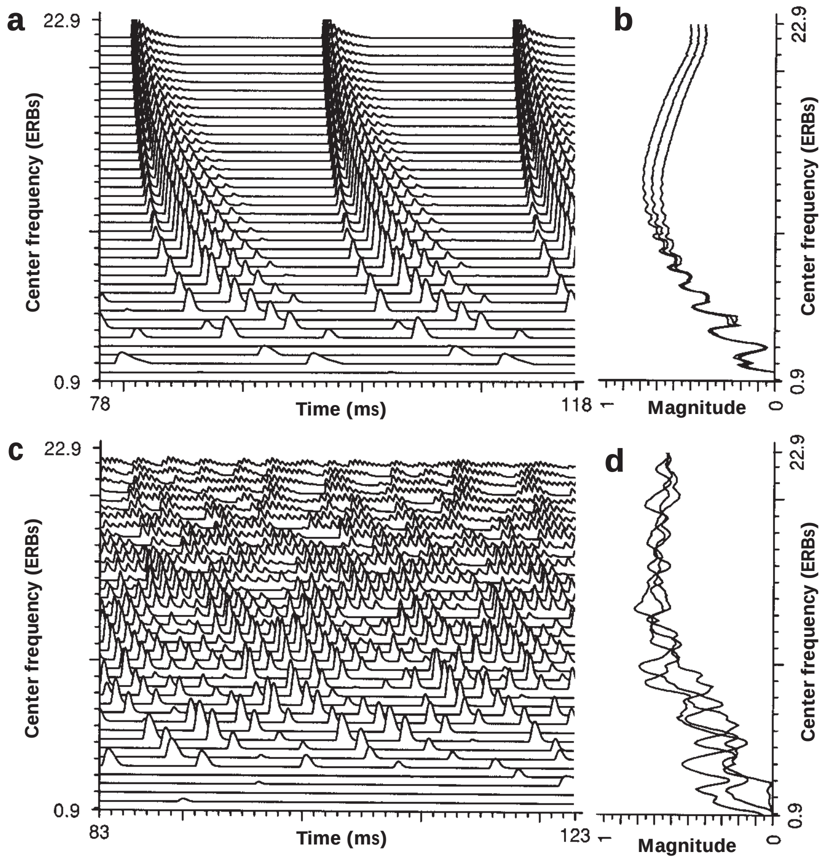

Fig. 1. Spectral representations of complex sounds that produce a pitch perception of 62.5 Hz. The right panels show the activity produced by the sounds across the auditory spectrum, averaged over time. **(a)** Activity pattern in the auditory nerve in response to a musical note. **(b)** The musical note produces a clear harmonic structure at low frequencies. **(c)** Activity pattern in response to the synthetic stimulus used in this experiment. A stimulus with eight delay-and-add cycles, or iterations, is represented. **(d)** The experimental stimulus produces no consistent spectral pattern, especially at the relatively high frequencies corresponding to the passband used in the experiment. The simulation is based on a conventional auditory filter bank implemented in the AIM software[2]. ERB, equivalent rectangular bandwidth (absolute frequency scale based on the ear's frequency analysis.

related with a delayed version of itself [6,7]. When autocorrelation is applied to each channel of the NAPs in Fig. 1, the result is the two 'autocorrelograms' in Fig. 2a and b; they shows that there is a concentration of time intervals at 16 ms in both stimuli in most channels of the NAP. Figure 2c and d show the average autocorrelation across frequency channels; the heights of the peaks at 16 ms show that the pitch is relatively stronger in the harmonic sound, but it is still prominent in the delay-and-add noise. The height of the peak in the 'summary autocorrelogram' rises and falls with the temporal regularity of the NAP, and it predicts the pitch of complex sounds and the relative salience of the pitch[6,7].

THE ANATOMICAL BASIS FOR TIME-INTERVAL PROCESSING

Temporal models of pitch do not address the question of where this time-interval processing might occur. Physiological studies show that the accuracy of timing information decreases at higher levels of the auditory system[8], which means that the temporal microstructure must be converted into a more stable code (probably spatial[9]) in the auditory pathway up to the cortex. The process involves short-term temporal integration in which different neurons detect and accumulate activity about specific time intervals in the NAP. Animal studies indicate that the temporal integration may occur in the inferior colliculus[9]; the current study was carried out to delimit the location of the temporal integration mechanism in humans. In models of temporal integration, regular and irregular sounds with the same total energy and the same auditory spectra produce the same firing rate in units located prior to temporal integration but distinctly different rates of firing after the regularity is detected. Specifically, regular sounds will produce greater activity than irregular sounds in units of the temporal array that detect the regularity. Delay-and-add noise enables us to create a set of stimuli, all of which have the same total energy and the same auditory spectra, but which nevertheless, have varying degrees of temporal regularity. These stimuli should produce uniform levels of activity in nuclei prior to temporal integration and very different levels in the temporal integration area and subsequent areas.

In this study, we used positron emission tomography (PET) to address the question of where in the auditory system temporal integration occurs. We have identified brain areas that show increased activation as a function of the degree of temporal structure of the stimulus, which must, therefore, lie at a higher level in the auditory pathway than the temporal-integration mechanism. We also investigated the effect of presenting these stimuli as melodies. This allowed us to define areas that are involved in the analysis of temporal structure at the level of sound sequences (seconds) rather than the level of pitch (milliseconds). We conclude that temporal integration for the representation of time structure at the millisecond level occurs before primary auditory cortex, whereas analysis of longer-term time structure in sound occurs beyond primary auditory cortex.

Results

FINE TIME-STRUCTURE ANALYSIS

First, we defined brain areas activated by the stimulus after the temporal integration mechanism had stabilized the representation of the fine time structure within the individual. Ten normal subjects underwent 12 scans using positron emission tomography to measure regional cerebral blood flow as an index of neural activity while attending to the stimuli. Subjects were required to listen for any changes in the sound but did not perform any output task. Over the 80-second acquisition period of any scan, the subject was presented with a sequence of one-second sounds. The pitch of the individual sounds in each scan, determined by the delay in the delay-and-add process, was varied for all scans over the same pitch range. All sounds in any scan contained the same degree of temporal structure, determined by the number of iterations. Between scans, the number

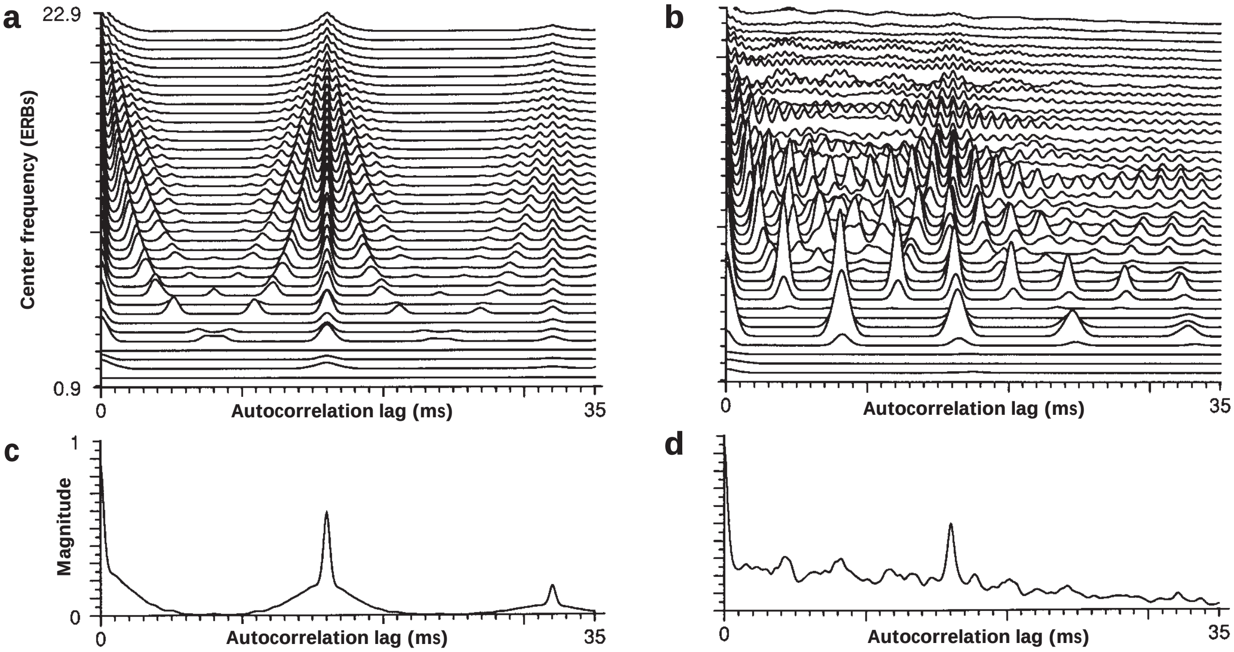

Fig. 2. Time-interval representations (autocorrelograms) of complex sounds that produce a pitch perception of 62.5 Hz. **(a)** The autocorrelogram of activity in the auditory nerve corresponding to the musical note in Fig. 1. **(b)** The autocorrelogram of activity in the auditory nerve corresponding to the synthetic stimulus used in this experiment. In both cases temporal regularity is apparent at 16 ms, which is clearly shown by the peaks in the summary autocorrelation for the musical note **(c)** and the experimental stimulus **(d)**.

of iterations was varied from 0–16 in a parametric design. Thus, although the scans were 80 seconds long, the only differences between the scans were in the fine temporal structure at the millisecond level. This allowed us to seek brain areas that occur after temporal integration; models of temporal integration[2] predict increasing activation as a function of temporal regularity after integration has occurred. This experiment was therefore specifically designed to delimit the location of brain regions occurring after temporal integration, rather than to demonstrate pitch mapping as in previous experiments[10].

The first analysis sought the predicted increase in activity with increasing temporal regularity in three areas for which there is reasonable anatomical specification: the inferior colliculus, medial geniculate body and primary auditory cortex. Figure 3 is a statistical parametric map showing areas where there was a main effect of temporal structure (number of iterations) on activity. That is, the map shows areas where regional brain activity increases with increased temporal regularity in the sound stimulus. Taking the accepted position of the human primary auditory cortices on each side[11] as our prior anatomical hypothesis, the relationship in both areas was significant using the spatial extent method of Friston[12]. The areas of activation show anterior posterior asymmetries, with the right area of activation being placed more anteriorly, which is consistent with previous reports of the location of the primary auditory cortex[11]. The area of activation on the left is situated more laterally and may also correspond to secondary auditory cortex rather than primary auditory cortex . No corresponding effects were demonstrated in the inferior colliculus or medial geniculate body. It was noteworthy that no significant activation as a fuction of temporal regularity was found in any other cortical area.

Table 1. Areas showing significant iteration–melody interaction

Area	Talairach coordinates			*Z* score
Right posterior temporal	72	-40	6	5.68
Left anterior temporal	-54	10	-18	4.99
Left posterior temporal	-58	-42	-2	4.73
Right anterior temporal	58	12	-26	4.71

LONGER-TERM TIME-STRUCTURE ANALYSIS

Second, we defined areas where the pattern of activation depends on the temporal structure of the sound at the level of sound sequences (seconds), rather than at the level of milliseconds, by presenting the individual sounds as different patterns of pitch (as melodies and as simpler pitch patterns). Specifically, we sought areas where the pattern of activation varied with the degree of sound-sequence complexity. Each subject underwent two scans for each level of fine-structure regularity (number of iterations). The pitch range in each scan was the same. However, the pitch was varied in two distinct ways at each level (Fig. 4). In the no-melody condition, the pitch was increased in a systematic 'staircase' pattern, whereas in the melody condition, the same variation in pitch over the scan was achieved by presenting novel diatonic melodies composed specifically for the experiment (Methods).

This second analysis sought brain areas involved in processing the longer-term temporal structure of the sound, determined by the type of sound sequence. We hypothesized that areas involved in the analysis of sound sequences would show greater activation as a function of iteration for the melody condition than for the no-melody condition. Sound-sequence analysis depends on detecting the pitches of individual sounds, and areas involved in sequence analysis will show a more marked dependence on pitch strength for more complex pitch sequences like the melody condition. Such an interaction was shown in four areas (Table 1; Fig.

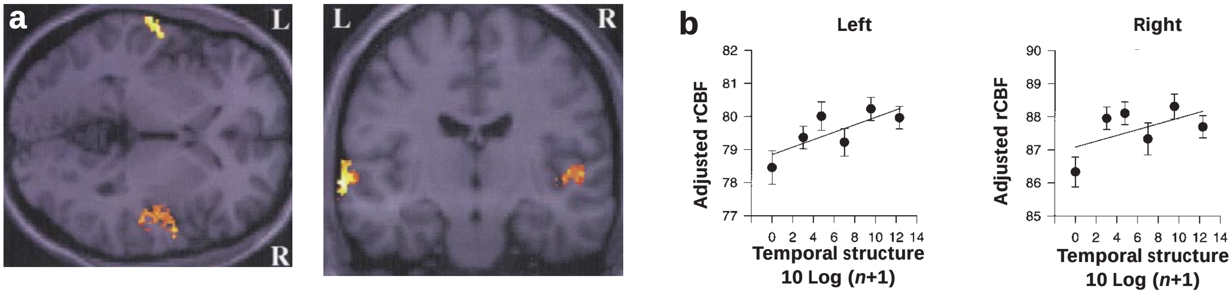

Fig. 3. Main effect of temporal structure within the individual sounds (iteration). **(a)** Statistical parametric map showing areas where brain activity increased with increasing temporal structure of the sound, as determined by the number of delay-and-add cycles (iterations) for the noise stimulus. The areas of significant effect are shown in color, rendered onto axial and coronal MRI brain sections. Activation is demonstrated in the left superior temporal planum (Talairach coordinates -70 -16 -2, $Z = 3.45$) and in the right superior temporal planum (58 -2 -4, $Z = 2.8$) at the level of the primary auditory cortices. **(b)** Relationship between activity and temporal structure in the two areas showing a significant effect in (a). Temporal structure is measured by 10 log ($n + 1$), where n is the number of iterations.

5a and b): the right and left posterior superior temporal gyri and the right and left anterior temporal lobes. These areas are not contiguous with primary auditory cortex, as are most of the primate auditory association cortices[13–15]. Figure 5c shows the relationship between activation and iteration for the no-melody and melody conditions considered separately, in the four areas of interaction. In all areas, the blood flow decreased with iteration in the no-melody condition and increased with iteration in the melody condition.

Discussion

Fine time-structure analysis

This is the first demonstration of activity in the human auditory cortex that varies with the degree of temporal regularity independently of the auditory spectrum, indicating that temporal integration occurs at or before the auditory cortex. In contrast, most previous functional imaging studies of human auditory cortex have demonstrated a mapping of spectral features[16–18]. Our data show that the human primary auditory cortices subserve both timing and spectral analyses. This is in accord with animal work suggesting orthogonal representation of both timing and spectral information in the ascending pathway in the inferior colliculus[9]. Although the data are consistent with temporal integration at or before the auditory cortex, we suspect that this process has occurred before the cortex; lesions of the human auditory cortex produce deficits at longer temporal intervals (tens of milliseconds) than those involved in pitch perception (milliseconds)[19]. Two previous human studies using magnetoencephalography are also consistent with auditory cortex subserving perception based on timing information[10,20], but they used stimuli with peaks in the auditory spectrum as well as temporal regularity. The latter study, additionally, suggested orthogonal coding of timing and spectral information in the auditory cortices, given the constraint on the temporal interpretation above. The current study indicates that iterated rippled noise in combination with parametric functional imaging provides a powerful method for studying auditory temporal processing in the absence of spectral cues.

Longer-term time-structure analysis

The analysis of longer-term time structure demonstrated strikingly symmetrical areas in the anterior and posterior temporal lobes where the pattern of activity varies according to the complexity of the sound sequence. These areas are hypothesized to be involved in sound-sequence analysis. The differences in sequence complexity occur at the level of seconds, which is relevant both to melodic perception and to the suprasegmental (prosodic) analysis of speech. We suggest that the areas shown are analyzing longer-term time structure in a manner that is not specific to music or speech. Human lesion studies[21] also support bilateral involvement of the superior temporal cortices in both melodic and prosodic analysis, although melodic analysis deficits without dysprosody can occur with right unilateral lesions[22,23]. The areas shown are likely to be activated via cortico-cortical projections from primary auditory cortex. We suggest that they are acting on 'pitch tokens' (symbolic representation of pitch) and are not involved in the primary processing of the pitch itself.

Stimuli such as these where the spectrum is carefully controlled may be useful in future studies of sound sequencing and melodic perception, as they allow manipulation of the long-term temporal structure without altering the auditory spectrum over time. In previous imaging studies of pitch processing in speech and music[24–26], the output tasks have been varied with identical stimuli to demonstrate different neural substrates for different aspects of speech and music processing. This represents a more 'cognitive' approach to musical

Fig. 4. Pitch sequences used as stimuli. Excerpts from the beginning and end of **(a)** the no-melody pitch sequence and **(b)** one of the melody pitch sequences. The greater sequence complexity in the melody condition is apparent.

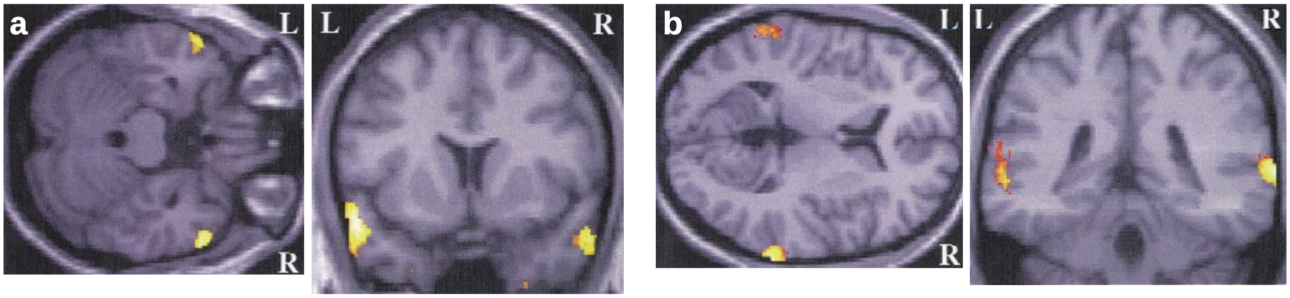

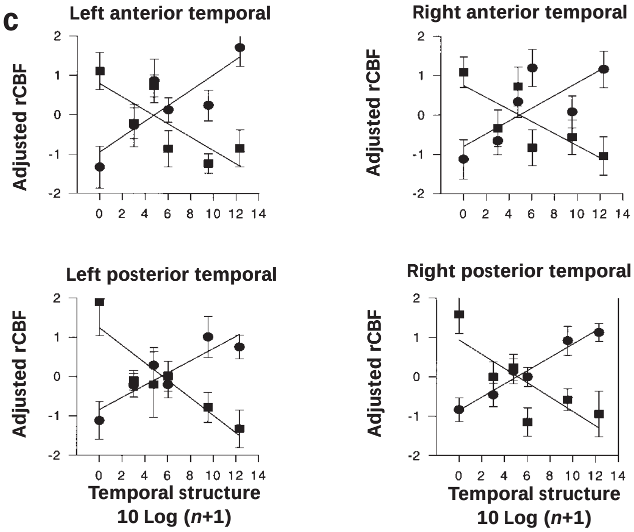

Fig. 5. Iteration–melody interaction. Statistical parametric map showing areas where brain activity increased more with iteration for the melody condition than for the no-melody condition. **(a, b)** Bilateral activation of the anterior (a) and posterior (b) temporal cortices is demonstrated. The rostrocaudal (z) and anteroposterior (y) levels for the sections are $z = -20$, $y = 0$ (a) and $z = -10$, $y = -40$ (b). **(c)** Patterns of brain activity (measured by adjusted regional cerebral blood flow) for the two conditions with different long-term temporal structure, in the four areas showing significant interaction in (a) and (b). For each condition, regional cerebral blood flow is plotted against the temporal structure in the individual sounds. Temporal structure is measured by 10 log (n + 1), where n is number of iterations. Squares show average regional cerebral blood flow for ten subjects for the no-melody condition, and circles show average regional cerebral blood flow for the melody condition. Adjusted regional cerebral blood flow is the value relative to a defined mean of 50 ml per dl per min. Error bars show one standard error. The regression line for each condition is plotted.

analysis than in the present study; we have demonstrated activation beyond primary auditory cortex without varying the output task. This suggests that we are observing a low-level mechanism for sound sequencing, and we would argue that the present study shows activation that is relevant to the temporal analysis of sound in general.

HIERARCHICAL ANALYSIS OF TEMPORAL STRUCTURE

This work suggests that the processing of temporal structure in the human brain occurs at different anatomical levels. Fine temporal structure associated with the perception of pitch is processed in the auditory pathway up to and including primary auditory cortex, whereas the analysis of pitch sequences involves cortical areas distinct from primary auditory cortex. Processing of emergent temporal properties, such as pitch sequences in sound, is allowed by cortico-cortical connections from the primary auditory cortex and suggests analysis in the human brain based on a hierarchy of temporal levels.

Methods

STIMULI AND PET SCANNING. Twelve scans were carried out on each subject using the oxygen-15-labeled water bolus technique. A Siemens scanner was used in three-dimensional mode. Stimuli were add-same iterated rippled noise[7] (IRNS) presented over insert earphones at 75 dB SPL total power. Sounds were 1 s long and presented with a repetition period of 1.1 s. Within any given scan, the number of iterations for all the sounds was fixed at one of six values: 0, 1, 2, 4, 8 or 16. The passband for the sounds was 1–3.5 kHz. High-pass filtering at 1 kHz precludes spectral processing in the auditory system for these low-pitched stimuli. At each of the six values of iteration, subjects underwent two scans, during which the pitch of the sounds was varied between 50 and 250 Hz in one of two different ways. In the ‘melody’ condition, pitch variation was used to produce novel diatonic melodies consisting of six four-bar phrases that presented a theme, development and recapitulation in the traditional way (Fig. 4b). The sounds all had the same duration, and so there was no rhythmic information except in the last bar of each phrase, where there was a double gap before the first note and no gap between the last two notes. This modification suggests an ‘end-of-phrase’ pause, which is important to the sense of music, without actually varying the mean rate of sounds. In the no-melody condition (Fig. 4a), pitch variation was strictly limited; the pitch was fixed for 12 notes and then it was increased by two notes on the diatonic scale. The initial pitch was always A1 (55 Hz), and over the 8 sets of 12 sounds, the pitch rose slowly over the full range used in the melodies. There was no rhythm information whatsoever in the no-melody condition; the sounds were delivered at a constant rate.

ANALYSIS. Data were analyzed with Statistical Parametric Mapping software [SPM96, http//:www.fil.ion.ucl.ac.uk/spm]. Scans were realigned and spatially normalized[27] to the standard stereotactic space of Talairach[28]. The standard template of the Montreal Neurological Institute was used[29]. The data were smoothed with a Gaussian filter (full width at half maximum of 16 mm). All analyses were also carried out using a Gaussian filter with a full width at half maximum of 8 mm, to maximize the chance of demonstrating activation in small subcortical structures. Analysis of covariance was used to correct for differences in global blood flow between the scans and implement a regression analysis to find areas where blood flow increased linearly with temporal structure (represented by log (n + 1), where n is number of iterations). The significance of this regression was assessed with the t statistic at each voxel. These statistics (after transformation to a Z score) constitute an SPM{Z}[30]. Implementation of second- and third-order polynomial fit did not demonstrate a better fit to the data. The activations demonstrated in the primary auditory cortices as a function of iteration were significant at $p < 0.05$ using the spatial extent method of Friston[12] to take the prior anatomical hypotheses into account.

The interaction analysis sought areas where the relationship between regional cerebral blood flow and iteration had a greater slope for the melody than for the no-melody condition. The interactions in the areas listed were significant at the level of $p < 0.05$ after correction for multiple comparisons using Gausssian field theory, in the absence of a specific prior anatomical hypothesis. A main effect of melody was also demonstrated involving the primary auditory cortices and extending along the superior temporal lobes bilaterally. However, the interaction analysis is more informative than the simple main effect: the main effect demonstrates areas involved in sequence analysis and areas with an augmented response to sounds that simply vary more from sound to sound.

The consistent difference in blood flow between the zero-iteration conditions [$10 \log (n + 1) = 0$] in each panel of Fig. 5c is due to the simple rhythm cue used to indicate 'end of phrase' in the melody condition. Although the rhythm was constant across iterations within the melody and no-melody conditions, it could still contribute to the iteration.

Acknowledgements

T.D.G., C.B. and R.S.J.F. are supported by the Wellcome Trust. R.D.P. is supported by the MRC (UK).

RECEIVED 29 APRIL; ACCEPTED 16 JULY 1998

1. von Helmholtz, H. L. F. *On the Sensations of Tone* (Longmans, London, 1885).
2. Patterson, R. D., Allerhand, M. H. & Giguerre, C. Time-domain modeling of peripheral auditory processing: a modular architecture and a software platform. *J Acoust. Soc. Am*. 98, 1890–1894 (1995).
3. Patterson, R. D. The sound of a sinusoid: Spectral models. *J. Acoust. Soc. Am*. 96, 1409–1418 (1994).
4. Cohen, M. A., Grossberg, S. & Wyse, L. A spectral network model of pitch perception. *J. Acoust. Soc. Am*. 98, 862–879 (1995).
5. Rose, J. E., Brugge, J. F., Anderson, D. J. & Hind, J. E. Phase-locked response to low-frequency tones in single auditory nerve fibres of the squirrel monkey. *J. Neurophysiol.* 30, 769–793 (1967).
6. Patterson, R. D., Handel, S., Yost, W. A. & Datta, A. J. The relative strength of the tone and the noise components in iterated rippled noise. *J. Acoust. Soc. Am*. 100, 3286–3294 (1996).
7. Yost, W. A., Patterson, R. & Sheft, S. A time domain description for the pitch strength of iterated rippled noise. *J. Acoust. Soc. Am*. 99, 1066–1078 (1996).
8. Rouilly, E., deRibaupierre, Y. & deRibaupierre, F. Phase-locked responses to low frequency tones in the medial geniculate body. *Hear. Res*. 213–226 (1979).
9. Langner, G. & Schreiner, C. E. Periodicity coding in the inferior colliculus of the cat. I. neuronal mechanisms. *J. Neurophysiol.* 60, 1799–1822 (1988).
10. Pantev, C., Hoke, M., Lutkenhoner, B. & Lehnertz, K. Tonotopic organization of the auditory cortex: pitch versus frequency representation. *Science* 242, 486–488 (1989).
11. Penhune, V. B., Zatorre, R. J., MacDonald, J. D. & Evans, A. C. Interhemispheric anatomical differences in human primary auditory cortex: probabalistic mapping and volume measurement from magnetic resonance scans. *Cereb. Cortex* 6, 661–672 (1996).
12. Friston, K. J. Testing for anatomically specified regional effects. *Hum. Brain. Mapp.* 5, 133–136 (1997).
13. Merzenich, M. M. & Brugge, J. F. Representation of the cochlear partition on the superior temporal plane of the macaque monkey. *J. Neurophysiol.* 24, 193–202 (1973).
14. Rauschecker, J. F., Tian, B., Pons, T. & Mishkin, M. Serial and parallel processing in rhesus monkey auditory cortex. *J. Comp. Neurol.* 382, 89–103 (1997).
15. Pandya, D. N. Anatomy of the auditory cortex. *Rev. Neurologique* 151, 486–494 (1995).
16. Lauter, J. L., Herscovitch, P., Formby, C. & Raichle, M. E. Tonotopic organisation in the human auditory cortex revealed by positron emission tomography. *Hear. Res.* 20, 199–205 (1985).
17. Talavage, T. M., Ledden, P. J., Sereno, M. I., Rosen, B. R. & Dale, A. M. Multiple phase-encoded tonotopic maps in human auditory cortex. *Neuroimage* 5, S8 (1997).
18. Wessinger, C. M., Buonocore, M. H., Kussmaul, C. L. & Mangun, G. R. Tonotopy in human auditory cortex examined with functional magnetic resonance imaging. *Hum. Brain Mapp.* 5, 18–25 (1997).
19. Phillips, D. P. & Farmer, M. E. Acquired word deafness and the temporal grain of sound representation in the primary auditory cortex. *Behav. Brain Res.* 40, 85–94 (1990).
20. Langner, G., Sams, M., Heil, P. & Schulze, H. Frequency and periodicity are represented in orthogonal maps in the human auditory cortex; evidence from magnetoencephalography. *J. Comp. Physiol.* 181, 665–676 (1997).
21. Peretz, I. *et al.* Functional dissociations following bilateral lesions of auditory cortex. *Brain* 117, 1283–1301 (1994).
22. Zatorre, R. J. & Halpern, A. R. Effect of unilateral temporal lobe excision on perception and imagery of songs. *Neuropsychologia* 31, 221–232 (1993).
23. Griffiths, T. D. *et al.* Spatial and temporal auditory processing deficits following right hemisphere infarction. A psychophysical study. Brain 120, 785–794 (1997).
24. Zatorre, R. J., Evans, A. C., Meyer, E. & Gjedde, A. Lateralization of phonetic and pitch discrimination in speech processing. *Science* 256, 846–849 (1992).
25. Zatorre, R. J., Evans, A. C. & Meyer, E. Neural mechanisms underlying melodic perception and memory for pitch. *J. Neurosci.*14, 1908–1919 (1994).
26. Platel, H. *et al.* The structural components of musical perception. A functional anatomical study. *Brain* 120, 229–243 (1997).
27. Friston, K.J. *et al.* Spatial registration and normalisation of images. *Hum. Brain Map.* 2, 165–169 (1995).
28. Talairach, P. & Tournoux, J. *A Stereotactic Coplanar Atlas of the Human Brain* (Thieme, Stuttgart, 1988).
29. Evans, A. C., Kamber, M., Collins, D. L., & Macdonald, D. in *Magnetic Resonance Scanning and Epilepsy* (eds Shorvon, S., Fish, D. Andermann, F., Bydder, G. M. & Stefan, H.) 263–274 (Plenum, 1994).
30. Friston, K.J. *et al.* Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* 2, 189–210 (1995).