The role of primary auditory cortex in sound localisation

Mr. Richard Lanyon, B.A.

March 24, 2003

Abstract

This thesis investigates the involvement of primary auditory cortex (A1) in sound localisation. Experiments were carried out both to assess the effect of A1 inactivation on sound localisation, and to measure the sensitivity of single A1 neurons to sound source location.

Ferrets were trained to localise bursts of broadband noise, of varying intensity and duration, from an array of loudspeakers that spanned 360° in azimuth. Bilateral A1 inactivation caused an impairment on this task, but only for short-duration stimuli. Unilateral A1 inactivation also resulted in an impairment for short-duration stimuli, but this was limited to the side of space contralateral to the inactivation, and was only seen in animals which had been highly trained prior to surgery. A feature of the impairment in all animals was the increased number of "front-back confusions", where the animal's response was on the correct side of the midline but the wrong side of the interaural axis.

Recordings from ferret A1 showed that the firing rate of individual neurons varied little as sound source location was changed. Further, the neurons' location sensitivity was affected by changes in stimulus intensity and duration. However, mathematical techniques were used to measure the information these neurons provided about sound source location, and it was found that this information was not sensitive to intensity or duration changes. The analysis also showed that the amount of information provided by response latency was greater than that carried by firing rate. Similar mathematical treatment tentatively suggested that the information from different neurons was only slightly redundant, so it may be possible to account for whole-animal localisation performance by assuming that the output of large numbers of neurons is considered.

It is concluded that A1 is involved in processing the location of sound sources, but it seems unlikely that sound localisation is A1's primary or only role within the auditory system.

Contents

Part I

Introduction

Chapter 1

Introduction

The physiology of hearing is well understood at the level of the ear and the more peripheral parts of the auditory brainstem, but relatively little is known of the more complex processing which goes on at higher levels of the auditory system. One of the most studied 'high-level' auditory processing areas is primary auditory cortex (A1). Many functions have been attributed to this area of the brain, but none conclusively. Studies which have recorded neuronal activity in A1 have typically concluded that A1 has a role in most types of general auditory processing, including the analysis of pitch, temporal structure, species-specific vocalisations and sound localisation. By contrast, data from animals with A1 lesions points particularly to a role in sound localisation. The goal of this work to investigate the nature of the involvement of A1 in sound localisation. It is hoped that this will result in a better understanding of how the brain as a whole processes high-level auditory information.

1.1 The functional anatomy of the brain

An assumption underlying this work is that different areas of the brain have different functions. This premise is based mainly on the observation that focal brain damage can lead to very specific behavioural problems. However, it is important to understand how this functional segregation might occur, as the explanation will bear on the extent to which A1 has a specialised role in hearing.

1.1.1 Functional specialisation

It seems intuitively obvious that distributing the brain's processing among different areas would be efficient, but it is worth considering exactly why this

is the case.

The brain consists of many neurons, where each neuron can be thought of as a separate information processor and the brain as a whole can be considered as a multiprocessor computer. The connections between different processors can be of two different types: serial connections and parallel connections. A serial connection is one where two processors are connected to one another in series, i.e. one after the other. A parallel connection is one where the processors act alongside one another, in parallel, and are independently connected to the input and output of the system.

Importantly, connecting processors in parallel allows for much faster processing than connecting them in series. In a serial system, each processor has to wait upon the output of another, so the delay between the input and the output of the system will be longer than for a parallel system. Also, a serial chain of processors cannot handle more information simultaneously than a single processor can. With a parallel architecture, by contrast, the input can be split up into smaller sections, each to be dealt with independently. This is why a "division of labour" between different processors (or groups of processors) might be efficient.

However, not all tasks can be processed in parallel. This is because one task will often depend upon the outcome of another task. For example, in order to recognise a word you must first have identified its component syllables. This means that the processor dealing with words will have to be connected to the output of the processor dealing with syllables, i.e. a serial connection is needed. In other words, some types of task processing are intrinsically serial in nature.

Ideally then, processors should be connected in parallel where possible, but connected serially when they are concerned with tasks that are related to one another. This implies a modular brain architecture, where each module of serially connected neurons will process a selection of tasks that are related, but will do so largely independently of other such "assemblies" processing different groups of tasks.

It is this modularity which is believed to underlie functional specialisation within the brain. If this belief is correct, an investigation into how the brain's different assemblies are specialised should provide important information about the way it organises its processing.

1.1.2 Anatomy and development

How does the modular processing suggested in section 1.1.1 relate to the actual anatomy of the brain? If there is a link, it will be for biological rather than computational reasons.

Most obviously, it might be expected that the brain should be wired with all connections as short as possible. This is because neurons, especially neurons with very long neurites, use up a lot of energy. Using the shortest connections possible would reduce the metabolic load of connecting neurons with long neurites. There are also developmental advantages to having short connections. During development, neurons send out processes which make contact with other neurons. Once contact is made, some of these connections will die off and some will survive. Thus the final pattern of connections will depend on the original direction of growth of the neurons, and/or on which connections survive and which do not. Both of these mechanisms are more efficient over short distances: directing processes to a predefined target over long distances is presumably difficult, while sprouting many random processes only to eliminate inappropriate connections later is wasteful.

If all neuron-to-neuron connections were as short as possible, then the distance between neurons would be related to the importance of making a connection between them. We have already seen that the importance of neurons being connected is associated with the inter-relatedness of the tasks they are processing. This means that neurons processing related tasks might be expected to be near to one another.

However, the brain is not just a mass of neuron-to-neuron connections: it has structure on a larger scale and features that can be distinguished without reference to single neurons at all. It is such features which divide the brain into areas like A1[47].

The reasons for the existence of such large-scale structures are probably developmental. Importantly, it is unlikely that any developmental mechanism could specify all of the connections in the mammalian brain in a point-to-point manner. The number of neurons is simply too large for any signalling mechanism to indicate a single intended connection out of all the possible neuron-to-neuron choices.[1] Even if such a signalling system could exist, the amount of information required to specify all the correct connections in the human brain would be huge—more than can be contained in the human genome.[2] This means that, at least to some extent, brain connectivity must be specified at a coarser scale than neuron-to-neuron. The existence of a gross anatomical structure to the brain probably reflects connections being specified between groups of neurons, rather than between individual neurons.

[1] Any noise in such a mechanism would quickly obscure the signal, because the number of degrees of freedom (viz. the number of possible connection patterns) is so large.

[2] However, species with far fewer neurons may have nervous systems whose connections are entirely specified by their genes, an example of this being the nematode *Caenorhabditis elegans*.

1.1.3 Summary

If neural processing is believed to be modular, as argued in section 1.1.1, it is tempting to consider the large-scale brain areas described in section 1.1.2 as functional units. This assumption provides a simple framework allowing scientists both to get a handle on how the brain divides up its processing, and to divide up their research effort in a similar way.

However, while anatomical brain areas are clearly defined and separate from one another, the neural assemblies processing different tasks are poorly defined and may overlap. This makes it difficult to say that any particular group of neurons, such as a brain area, is uniquely responsible for processing a given task or tasks.

Furthermore, the way neural processing is organised is not fixed. This is because it is important for the brain's processing to be tailored to circumstances, e.g. during development, learning, or recovery from damage. It is believed that this plasticity results from a variety of changes in how neurons are connected to one another: these include changes in the strength of connections, and also the growth of new connections and "pruning" of old connections. However, not all aspects of brain structure are equally flexible, and the gross anatomy of the brain appears to remain fairly constant through life. If the brain is plastic mostly on a small scale, this suggests that the particular neurons responsible for a task may change, even though the division of the brain into areas remains the same.

Overall, it would seem that assigning functions to brain areas provides a useful heuristic for understanding how the brain works. However, it is important when doing this to be aware of all the limitations: the fact that various areas may be involved in the same functions, albeit to differing extents, and the fact that the functional organisation of the brain may change even though the large-scale structure of the brain probably will not.

1.2 Neuroscience and brain function

Given that it is valid, at least to an extent, to ascribe a particular function to a specific area of the brain, how do we find out what that function is?

1.2.1 Lesion techniques

Probably the most obvious, and certainly the oldest technique for assigning a role to a brain area is to look at what happens when there is damage to the area in question. This has been investigated in humans with accidental damage (e.g. stroke, trauma, anoxia) or surgical lesions made for medical

reasons (typically to treat epilepsy) as well as in animals with experimental surgical lesions[31]. More modern experimental techniques do not even need to permanently damage the area being investigated. In humans, trans-cranial magnetic stimulation has been used to temporarily disrupt brain activity within a specific area[27]. In animal experiments, cooling probes have been used to inactivate areas of brain acutely[93], whereas the release of drugs from osmotic mini-pumps or polymer implants[129] can be used to inactivate areas chronically.

The idea behind this approach is that lesioning (or otherwise interfering with) an area will produce a behavioural deficit, and the lesioned area is therefore responsible for "doing" whatever behaviour is impaired. Of course, any brain area is likely to be involved in a large number of behavioural tasks, and it is then up to the experimenter to describe the common features of the tasks on which there is an impairment.

A serious problem with the lesion method is that it will only work to the extent that the lesioned area acts independently of the rest of the brain. However, we have already seen (section 1.1) that segregation of function in the brain is far from complete. Lesioning an area will not only disrupt the tasks processed by that area, tasks A and B, say, but will also disrupt all other tasks which depend on tasks A and B.

As well as not being complete, segregation of function in the brain is not static. An example of this is that humans can recover from stroke. This is at least partly because intact areas take over the function of the lesioned area. As recovery can continue over long periods of time (sometimes years, in humans), it is an open question as to when exactly the assessment of any deficit should be made.

1.2.2 Recording techniques

The other major method of assigning a function to a sensory brain area[3] is to relate activity in that area to what the animal is experiencing.[4] There are a number of different ways of assessing the level of activity in an area of the brain. Functional imaging techniques such as PET, fMRI and optical imaging measure how the brain responds to the metabolic demands of active neurons. Functional imaging tends to have quite poor spatial and temporal resolution, but can simultaneously determine activity levels over a large area of the brain. Electrophysiological techniques instead record the actual electrical

[3]The brain is conventionally divided into sensory and motor areas. Different techniques, not covered here, are used to investigate motor areas of the brain.

[4]In fact, in most of these experiments the animal is anaesthetised, but it is still possible to record responses to sensory stimuli.

impulses generated by neurons as they send messages to one another. This can offer greater temporal and spatial resolution than imaging, but recording from multiple sites at the same time is more difficult. Depending on the filtering of the electrical signal, and on the physical size and positioning of the electrodes, such recordings can produce electroencephalograms (EEGs), local field potentials, multi-unit recordings or single-unit recordings (in increasing order of spatial resolution).

The premise of a neurophysiological experiment is that a given neuron (or group of neurons) will show activity patterns which are sensitive to whatever it is that the neuron is processing. For example, an experiment might involve recording the response to several auditory stimuli, where each stimulus differs from the others in level and/or frequency. If the neuronal response varies with changes in stimulus level but does not vary with changes in stimulus frequency, then the claim might be made that the neuron or neurons are processing stimulus level, and that stimulus frequency is being processed elsewhere.

However, as mentioned in section 1.1, different tasks will be processed by different but overlapping neurons. This means that a given group of neurons will be involved in more than one task, and thus may be sensitive to a wide range of stimulus parameters. In the case of the example given above, an auditory neuron may be sensitive to both stimulus level and stimulus frequency, and many other things besides. This makes neurophysiological data quite difficult to interpret.

Furthermore, sensory processing is not merely a matter of being sensitive to changes in stimulus parameters. A useful analogy is with a telephone line: it is true that the signal in a telephone cable depends on the original stimulus that produced it, but most people would hesitate to say that the telephone cable was "processing" the stimulus. By contrast, the telephones at either end of the cable really are processing the stimulus, as the input to a telephone is transformed in some way to get the output. If we regard stimulus processing as being a transformation between input and output, then we cannot draw conclusions about neural processing unless we know how the representation of stimulus parameters changes between the input and output of the neuron(s). This shows how it is important to interpret recordings from one part of the brain in the light of what is already known about recordings made in other areas that are connected to it.

1.3 Sound and sound localisation

In order to understand how primary auditory cortex might be processing sound, we need to consider what information is carried by sound waves, and what strategies the brain might have for analysing sound.

1.3.1 Sound and hearing

Sound consists of changes in pressure which are conducted, usually through air, to the ear. One way of describing such a stimulus would be as a change in air pressure over time. However, most natural sounds consist of air pressure changes which are cyclic or wave-like, i.e. they repeat at regular time intervals. As was pointed out by Laplace (but attributed to his supervisor, Fourier), it is often useful to describe such cyclic phenomena not as a continuous series over time, but as a collection of regularly repeating components. A so-called Fourier transform takes a continuous time series and expresses it instead as a frequency series. This consists of a set of one or more sinusoidal waves of different frequencies, with each one of these frequency components having an associated magnitude and phase. The magnitude of a component is how strong that component is compared to the other components with different frequencies, and the phase of a component indicates where the start of its cycle occurs with respect to the other components.[5]

If naturally-occurring sounds are subjected to a Fourier analysis, most of them are found to contain components at several different frequencies. Often these frequency components are harmonically related, that is, their frequencies are all integer multiples of a 'fundamental' frequency. There also exist 'noisy' sounds whose components are not harmonically related, but are spread at random across the frequency spectrum. It is possible to produce sounds consisting of only a single frequency: these are called 'pure tones', and while they are rarely encountered in nature they are often used in scientific experiments.

The description of sound in terms of waves of different frequencies is almost inescapable in auditory science. This is mainly because there are some similarities between the operation of the auditory system and a Fourier analysis of sound. The initial transduction of sound into neural activity occurs within the cochlea (part of the inner ear), where sound sets up travelling waves along a membrane known as the basilar membrane. This membrane carries hair cells, which respond to movement with changes in the amounts

[5]For a cyclic phenomenon, phase refers to the position between the start and end of a cycle. It is measured in degrees or radians, such that the start of a cycle is $0°/0$ rad, and the end of a cycle is $360°/2\pi$ rad.

of neurotransmitter they release; this in turn affects the firing of auditory nerve fibres[30]. The basilar membrane is narrow and taut at one end (the basal end) and wide and loose at the other end (the apical end). As a result of this, the basilar membrane at the basal end of the cochlea moves more in response to high-frequency sounds, and that at the apical end moves more in response to low-frequency sounds. Thus the pattern of activity within hair cells[6] at different points along the basilar membrane resembles the power at different frequencies in a Fourier analysis of the original sound. This representation of sound—where different neurons are sensitive to different frequency components—is found in many regions of the auditory system, including primary auditory cortex (see section 1.4).

It is important to realise, though, that the cochlea does not perform anything like a true Fourier analysis. There are several reasons for this, not least the fact that a Fourier analysis is linear, while both the movement of the basilar membrane and the firing of auditory nerve fibres are non-linear. Another reason is that the auditory system is operating in "real-time", and so must be performing a continuous series of analyses, each over quite a short period. This is crucial, because it effectively divides auditory phenomena into those that occur on a short time scale and are analysed primarily by their frequency, and those that occur on a longer time scale and are analysed by their timing. These two methods of auditory analysis are known as spectral analysis and temporal analysis, respectively.

1.3.2 Cues to sound source location

The brain has two main sources of information about the position of a sound source. One is the irregular shape of the outer ear, as this means that the effects which the outer ear has on incoming sound will depend upon the location of the sound source. The other is the existence of two ears, as, depending on the location of the sound source, the resulting sounds may be more intense and/or arrive earlier in one ear than the other.

Neither of these cues gives much information about the distance of a sound source, except perhaps for sound sources which are very close to the head. Indeed, little is known about auditory distance perception[10, 14], and most studies of sound localisation (including this one) restrict themselves to the perception of the direction of a sound source.

Sound source direction is usually quantified in terms of azimuth (direction measured along the horizon) and elevation (direction measured in the vertical

[6]It should be noted that hair cells are not neurons, and do not fire "all-or-none" action potentials. Activity in these cells is instead expressed in terms of graded receptor potentials, which influence neurotransmitter release.

plane).

1.3.2a Monaural pinna cues

The shape of the outer ear acts to amplify some frequencies of incoming sounds and to attenuate some others[128]. In particular, the numerous small folds of the pinna and concha can have very dramatic and narrowband effects on the transmission of sound to the eardrum. The pinna and concha are also asymmetric in shape, so these effects vary with the angle of incidence of the sound. In other words, the pattern of which frequencies are attenuated and which amplified will be different for sounds arriving from different directions[16, 113].

Such filtering effects are not always effective as a localisation cue. Low frequency sounds will have wavelengths similar to or greater than the size of the folds in the pinna or concha. As a result, low frequency sounds will simply diffract around the ear, and there will not be any noticeable filtering. Also, detecting that some frequency components are amplified/attenuated more than others requires that the original sound has several different components, i.e. it must be broadband.

There is a further problem which results from needing to compare the amplitude of different frequency components of the same sound. These comparisons are subject to the process of deciding which frequency components are identified with which sound source. This process, known as 'auditory scene analysis'[9, 79] is complex and poorly understood.

Another difficulty is that the brain cannot calculate the true effect of pinna/concha filtering on a sound, as it only has access to the sound after it has passed through the outer ear. Given a particular feature in the sound spectrum detected by the inner ear, there is no way of telling whether it was present in the original sound or whether it was produced by the pinna/concha. It seems likely that the auditory system makes some assumptions either about the spectrum of the original sound or the types of features produced by the outer ear—for example, particular spectral features may be much more likely to be produced by the outer ear than by anything in the environment, or might occur over different frequency ranges when produced by the outer ear than when present in environmental sounds. This remains largely hypothetical, however.

1.3.2b Interaural difference cues

The separation of the two ears on the head means that sound will travel a different path to each ear. This means that not only will a given sound

arrive in one ear before the other, but also the sound at the farther ear will be distorted by the 'acoustic shadow' cast by the head[108].

The most difficult aspect of analysing these interaural differences is that of 'stereo matching', i.e. deciding which components of the signal arriving at one ear should be compared to which components arriving at the other ear. As mentioned previously (section 1.3.2a), calculating which frequency component belongs to which sound source is non-trivial, so it is difficult to match sounds based on their original source. However, it is comparatively simple to calculate interaural differences for each sound frequency, and many areas of the brain which are sensitive to interaural differences are known to be organised on a frequency by frequency basis[52].

One obvious way of using interaural differences is to measure the interaural time difference (ITD) between the arrival of sound at the two ears. If ITDs are analysed frequency by frequency, then in practice this means analysis of interaural phase differences (IPDs). However, there is a problem in using IPDs as a localisation cue for high-frequency sounds. This is that the time taken for sound to travel between the two ears can be greater than the duration of one cycle of a high-frequency sound, and so IPD becomes an ambiguous indicator of ITD. This ambiguity can be removed by making comparisons of IPDs across frequency bands, but cross-frequency comparisons are not straightforward, as mentioned above.

Interaural level differences (ILDs, also known as interaural intensity differences or IIDs) result mainly from the acoustic shadow cast by the head, so that sound intensity is greater for the near ear than the far ear. The different outer ear filtering at the two ears will also contribute to ILDs. As mentioned in section 1.3.2a, sound attenuation is only significant at high frequencies: at low frequencies the wavelength of the sound is greater than the diameter of the head or the structures of the outer ear. Thus low-frequency sound will simply diffract around the head, with little loss of intensity between the near and far ears[108].

Because the head is approximately spherically symmetric with respect to ITDs[143, 69] and, to a lesser extent, ILDs, these cues are mainly useful in localising a sound source with respect to a line between the two ears (the interaural axis). If the head is modelled as a perfect sphere—a gross but useful simplification—then there is a conical locus of positions consistent with a given interaural difference. This locus is known as the "cone of confusion"[143]. For tasks where only azimuth is being judged, and not elevation or distance, the cone of confusion reduces to a series of "front-back confusions".

1.3.2c Using multiple cues

It is important to realise that monaural and binaural cues seldom occur in isolation. Interaural difference cues occur alone only in the case of pure-tone or low-frequency stimuli, while monaural pinna/concha filtering cues exist alone only in cases where a sound is detected in just one ear (e.g. if it is so quiet as to be detectable only in the near ear).

Furthermore, where the two types of cue do exist for the same sound they are not independent of one another: only certain combinations of cue values can be produced by real-world stimuli. This means it is potentially advantageous to process the two types of cue together. In particular, pinna cues are often cited as a means of disambiguating the front-back confusions which may occur in analysing interaural differences[19]. Conversely, comparing the frequency spectra in the two ears may make it possible to determine which features of those spectra are environmental features and which are the results of outer ear filtering.

1.4 Primary auditory cortex (A1)

Cerebral cortex is believed to be the location of the highest-level aspects of neural processing in mammals, including the processes which neuroscientists think of as "consciousness"[121].[7] This belief is based on phenomena such as that of "blindsight", which occurs in human patients whose cortical visual pathways are damaged,[8] but who have intact subcortical visual pathways. These subjects can correctly localise visual targets on a forced-choice test, but are not aware of any visual experience[7].

Auditory cortex, of which primary auditory cortex (A1) is a part, is a loosely-defined group of adjoining cerebral cortical areas all of which show increased neural activity in response to auditory stimuli. In most mammalian species, auditory cortex is located in the left and right temporal lobes of the cerebrum.

1.4.1 A1 and the auditory system

That A1 might be particularly important is indicated by its position in the auditory system, where it forms the main gateway to cerebral cortex. The auditory pathway starts in the cochlea and proceeds to the cochlear nucleus

[7] These are not necessarily the processes that either philosophers or normal people think of as consciousness.

[8] Specifically, these patients have damage to primary visual cortex (V1).

(CN), from where it splits into a large number of parallel pathways, most of which eventually terminate in the inferior colliculus (IC). Many of these pathways involve intermediate structures such as the superior olivary complex and the nuclei of the lateral lemniscus, but there is also a smaller direct projection from CN to IC. The IC projects mainly to the medial geniculate nucleus of the thalamus (MGN), and it is MGN which provides most of the auditory input to cerebral cortex[140]. It is because A1 is the principal cortical projection of the MGN that it is *primary* auditory cortex. However, the MGN does project to other auditory cortical areas as well as A1[3, 50], so it would be wrong to regard A1 as being the only way for auditory information to reach cerebral cortex.

While the IC, MGN and A1 seem to form something like a single auditory pathway into cortex, beyond that point the auditory areas do not seem to be arranged so straightforwardly. Several non-primary areas have been identified in the major experimental species. For example, in the cat, A1 can be found at the dorsal tip of the ectosylvian gyrus, and around it are areas such as secondary auditory cortex (A2), which is ventral to A1, and two areas anterior and posterior to A1 which are usually referred to simply as the anterior and posterior auditory fields (AAF and PAF respectively)[109]. In the rhesus macaque, A1 lies in the middle of the supratemporal plane, between a more rostral and a more caudomedial area (R and CM, respectively)[106]; in addition to these "core" areas, there are also a further three "belt" fields in the lateral belt of the macaque, as well as many "para-belt" areas elsewhere. Unfortunately, it is unclear how the areas in these different species correspond to one another. It is also unclear how these areas are arranged into a processing hierarchy, as they are heavily interconnected, both directly and also indirectly via feedback connections to subcortical areas. It has been argued, largely by analogy to the visual system, that there are separate streams for processing the location and the identity of a stimulus[106]—the "what" and "where" pathways. Others have proposed that the processing of species-specific vocalisations requires a specialised stream to itself, because of the importance of what is seen as a simple form of spoken language[135]. In both of these schemes, A1 is one of the final processing stages before the split into separate streams.

1.4.2 Recording studies

A1 has become a favoured target for auditory electrophysiologists, mainly because it lies immediately underneath the skull in many species and is thus easily accessible once a craniotomy is performed. The vast majority of recording in A1, as elsewhere, is carried out on anaesthetised animals. Traditionally

the anaesthetic has been barbiturate as it gives a stable recording environment, but the results produced are not necessarily representative of the functioning of the awake animal. The experimental species most commonly used are rhesus macaque, cat and ferret. Where comparable data are available from several species we concentrate mainly on ferret, as it is the species we ourselves use.

1.4.2a Sensitivity to sound spectrum and temporal structure

The most common test of a neuron's auditory sensitivity is to measure its response threshold—the minimum stimulus intensity needed to elicit a response—for pure tones of different frequencies. As mentioned earlier (section 1.3.2a), the outer ear can have significant distorting effects on incoming sounds, and it is important to distinguish these from the neuron's own sensitivity to frequency. For this reason, such experiments are typically carried out under "closed-field" conditions where the stimulus is presented directly to the eardrum, via headphones. Sound localisation is usually not of primary interest in these experiments, and so the stimuli used are either identical in the two ears, or are only present in the ear contralateral to the recording site.[9]

Under these conditions the responses of A1 neurons are very similar to those of neurons at lower levels of the auditory system. Like subcortical neurons, A1 neurons in barbiturate-anaesthetised animals respond robustly to pure-tone stimuli. And, like subcortical neurons, most A1 neurons are sharply "tuned" to pure tones of a particular frequency; that is, the response threshold is lowest for tones of a certain frequency and rises sharply as tone frequency is increased or decreased. These findings have been confirmed in a variety of species, including cat[111], ferret[63] and macaque[80]. The frequency at which pure-tone response threshold is lowest is known as the characteristic frequency (CF) or best frequency (BF) of the neuron.

While near-threshold responses in A1 are similar to those in the auditory midbrain and hindbrain, responses to supra-threshold tones can be different. Most neurons in the peripheral auditory system respond to increases in stimulus level with monotonic increases in response magnitude, up to the maximum firing rate of the cell. A much smaller number of peripheral neurons have a "best intensity": increases in stimulus level lead to increases in firing rate up to a certain level, but beyond that level further increases in stimulus level lead to decreases in firing rate. By contrast, neurons like this are much more common in auditory cortical areas and, to a lesser extent,

[9] Contralateral stimuli are used because more A1 neurons respond preferentially to them than to ipsilateral stimuli (section 1.4.2b).

the MGN. Estimates vary as to the proportion of A1 neurons which have non-monotonic rate-level responses, depending on the species, the criteria used and the range of intensites tested: 36% of cat A1 neurons have been reported to be non-monotonic[99], as compared to "the minority" of ferret A1 neurons[101] and 78% of A1 neurons in the awake macaque[96].

With respect to their pure-tone responses, non-primary areas of auditory cortex can be divided into two groups: those that, like A1, receive input from the ventral portion of the MGN, and those that do not. Areas with ventral MGN input show sharp and robust pure-tone frequency tuning, much like A1, e.g. PAF[102] and AAF[100] in the cat, and R in the rhesus macaque[80]. Other areas show higher thresholds and less clear responses to pure tones, e.g. cat A2[117] and macaque CM[80]. That is not to say that all responses derived from ventral MGN are the same, however, as cat PAF shows a considerably higher number of non-monotonic cells than A1[102] (86% *vs.* 36%) and PAF may also contain more cells with complex pure-tone responses than A1 does[70].

A1 neurons are in some ways less sensitive to the temporal structure of simple stimuli than neurons in more peripheral areas. Like many auditory neurons, those in A1 respond mainly to the onset of a stimulus. Also, the timing (first-spike latency) of an A1 neuron's response to tone bursts is dependent on the acceleration of peak pressure during tone onset[42], as it is in the auditory nerve[44]. However, unlike neurons in lower levels of the auditory system, A1 neurons often respond *only* to stimulus onset, rendering them insensitive to features such as the duration of a tone burst. Instead, the magnitude (the number of evoked spikes) of the response depends on the sound pressure at the instant of response initiation[43].

Sensitivity to stimuli with more complicated temporal structures is also less obvious in A1 than it is in areas which are earlier in the processing hierarchy. Responses to click trains[32, 115] and sinusoidal amplitude-modulated (AM) tones[136] show that, while A1 neurons do exhibit responses "locked" to stimulus modulations, the maximum modulation frequency for which this occurs is lower than in the MGN. However, while neuronal activity is not always synchronised to the clicks in a click train, a small number of A1 cells vary their overall firing rate in a way which indicates the click frequency even at high frequencies[72, 71]. It is often inferred that A1 temporal processing, at least in some cells, is taking place over a longer time-scale than in more peripheral areas of the brain.

Less research has been carried out on stimuli with complex spectral structure, either in A1 or elsewhere. Some A1 neurons in the cat show responses to species-specific vocalisations which are more than the sum of the components of the stimulus, and this is true for a higher proportion of A1 neurons

than for MGN neurons[137]. On the other hand, recordings in A1 of the awake monkey show that neurons are "tuned" to the spectral content of harmonic stimuli rather than their pitch[36]—in other words, the neurons were responding to individual components of the sounds rather than to the harmonic structure as a whole. This may not be true at later stages of processing. Furthermore, the proportion of cells that respond preferentially to particular species-specific vocalisations is greater in the macaque lateral belt than it is in A1[131]. Those authors who have proposed the existence of two or more processing streams in auditory cortex usually ascribe the processing of vocalisations to one particular stream[135].

An alternative approach to mapping the sensitivity of A1 neurons involves the use of large stimulus sets varying systematically in both the spectral and temporal dimensions. Such stimuli include "spectral ripples"[126, 125], "random chords"[33, 116] and "temporally orthogonal ripple combinations"[68, 67]. The response of most A1 neurons to these stimuli appears to be a linear combination of the responses to each of the individual spectral and temporal components, with two notable exceptions. One is that neurons have a maximum and minimum firing rate—this obviously introduces non-linearity into a neuron's responses. More interestingly, many A1 neurons show different responses for stimuli containing upwards changes in frequency from those containing downwards changes in frequency[68]. This agrees with earlier claims that neurons which have a preference for the direction of frequency-modulation (FM) are more common in ventral MGN and A1 than elsewhere[136].

The problem with these arbitrary stimuli is that the auditory system may respond to ecologically relevant stimuli in a qualitatively different way. For example, acoustic features of a bird's chirping which are well outside the "spectro-temporal receptive field" of a cat A1 neuron can still have an effect on that neuron's response[8]. Showing even more specificity, some cells in marmoset A1 respond differently to time-reversed recordings of marmoset vocalisations than they do to the unreversed recordings, whereas neurons in cat A1 do not distinguish between the two versions of the marmoset call[135]. This is even though marmosets and cats have a similar range of hearing.

1.4.2b Sensitivity to sound source location

The most obvious way to investigate the processing of sound source location in A1 is simply to play sound stimuli from various possible locations and see how neuronal response changes. Stimuli presented like this, from external speakers rather than headphones, are known as "free-field" stimuli. The sounds in question are usually broadband noise bursts of some description:

the broadband nature of the stimulus being most important, as it provides various localisation cues which are not present when the stimuli are narrow-band (section 1.3.2).

Individual A1 neurons seem surprisingly insensitive to the location of a sound source. Many "omni-directional" cells exist which respond equally well irrespective of location—the precise number depends on the criteria used, but is reported to be ∼20% of tested cells[49, 105]. Even the cells that are not omni-directional often have a very large "receptive field", that is, there are a large number of positions from which a sound will provoke a response. The exact size of this receptive field depends on stimulus intensity, with higher intensities giving a larger receptive field, but it frequently spans 180° in azimuth. In some cases increasing stimulus intensity causes a shift in the receptive field's position, as well as an increase in its size.

All these directional properties are similar to free-field location tuning in the areas immediately upstream of A1 in processing, viz. MGN[51, 23] and inferior colliculus[1, 2]. However, neurons in rhesus macaque area CM, which is believed to be downstream of A1 in processing terms[107], appear to be more selective for spatial location[112]. The same is also true for neurons in the lateral belt area of the macaque, which is even further on in processing[106]. Notably, the neurons most sensitive to species-specific vocalisations (mentioned above) are in a different part of the lateral belt to those which are most sensitive to sound source location.

Apart from its breadth, the other notable feature of A1 tuning to sound source location is its preference for certain positions. Specifically, more cells respond to stimuli in the contralateral hemifield (relative to the recording site) than to stimuli in the ipsilateral hemifield. The exact percentages vary between studies, but up to half of all cells respond better to contralateral than ipsilateral stimuli, with the rest split between cells that respond best to ipsilateral stimuli, cells that respond best to stimuli positioned at the midline and cells that have no directional preference[35, 49, 105]. Similar results hold in the MGN[23], but it is notable that in the inferior colliculus the proportion of cells responding to midline and ipsilateral stimuli is much lower[84, 83].

Of the cues that contribute to location tuning, the easiest to study are interaural cues, and the vast majority (>80%) of A1 neurons show some evidence of input from both ears (in closed-field studies). About 70% of binaural cells are sensitive to ILD cues, and most of these can be shown to have an excitatory input from the contralateral ear and an inhibitory input from the ipsilateral ear[98, 101]. Such cells are thus most responsive to stimuli which are more intense in the contralateral ear. A much smaller number of ILD-sensitive cells receive excitatory input from both ears, and show some sort of multiplicative interaction between the two inputs. These

cells respond most strongly when the stimulus level is nearly the same in both ears. ITD sensitivity is less common than ILD sensitivity, mainly because it is only exhibited in response to low-frequency sounds, and thus only by cells with a low CF. However, among neurons with a CF of less than 2.5 kHz, about 80% are responsive to ITDs[110]. If these cells are presented with ITDs within the range that would occur naturally,[10] they either show a response which increases in magnitude from one end of the ITD range to the other, or there may instead be a single "best" ITD within the natural range, with responses decreasing markedly to either side of this optimum delay. If cells are presented with ITDs outside the physiological range, however, they can be seen to have several optimum delays, each differing from the other by the period of the stimulus[11]; this illustrates how the auditory system cannot distinguish IPDs which differ by an integer number of cycles.

Spectral cues have not received as much attention as interaural difference cues, whether in A1 or elsewhere, perhaps because they are more difficult to produce. However, it has recently become possible to use in-ear microphones to record the effects of the outer ear on free-field stimuli, and then to generate closed-field stimuli which incorporate some of these effects. Studies using these "virtual space" techniques[13] have found that the receptive field of an A1 neuron is predictable on the basis of its binaural inputs, its pure-tone frequency tuning and the measured pinna effects[12, 116].

1.4.2c Distribution of response properties within A1

Although A1 is partly defined by its anatomy, it is also defined by the way its single-cell response properties are distributed: A1 is organised in what is called a 'tonotopic' or 'cochleotopic' fashion[111, 63]. In other words, neuronal CF varies systematically across A1, such that neurons with low CFs are found at one end, and neurons with high CFs at the other. This resembles the organisation of the cochlea, which is more sensitive to low frequencies at the apical end and high frequencies at the basal end (section 1.3.1). However, the variation in CF across A1 is not smooth (unlike the cochlea); instead A1 contains discrete 'isofrequency strips', each with neurons of similar CF.

In fact, almost all subcortical auditory structures are also organised like this, from the auditory nerve all the way up to the MGN. So are those auditory cortical areas which receive input from ventral MGN[109]. Even neurons in the macaque lateral belt can be seen to be organised in a similar manner, if narrowband noise is used instead of pure tone stimuli[106].

[10] The range of naturally-occurring ITDs depends on the size of the animal's head, but is of the order of a few hundred microseconds.

What is particularly obvious with cortical rather than subcortical areas, because they are laid out two-dimensionally as a sheet, is that the tonotopic arrangement only accounts for organisation along one dimension. There is no conclusive evidence of any stimulus parameter being mapped out continuously across the other dimension of A1, though it is sometimes assumed that this is simply because the relevant parameter has not been identified. However, recordings and functional imaging in the gerbil have shown that sensitivity to stimulus periodicity—associated with the fundamental frequency of harmonic sounds—is mapped continuously across A1, in a "horseshoe-like" manner[119]. This periodicity mapping appears to be unrelated to tonotopicity.

Other neuronal properties are arranged non-randomly across A1, though none of them are as systematically distributed as CF. Sensitivity to many parameters is arranged in a "patch-like" manner across A1, whereby neurons with similar response properties tend to occur near to one another. This has been demonstrated for both sharpness of frequency tuning[45, 133] and various features of the intensity-response curve such as threshold, dynamic range and non-monotonicity[46]. It has also been claimed that connections within A1 occur mainly between patches with similar properties[109]. Sensitivity to some other stimulus parameters appears to vary smoothly between the centre of A1 and the border with other areas—this is similar to the arrangement of sensitivity to "best" modulation frequency of AM tones in the cat inferior colliculus[118]. Other, more complex aspects of frequency processing such as FM direction-selectivity do not appear to be distributed in any systematic manner through A1[45].

Cells with similar binaural properties[101, 62] and/or tuning to sound source location[104, 22] also appear together in patches, but there is certainly nothing like a map of auditory space in A1. This is even though maps of auditory space exist elsewhere in the mammalian brain, e.g. in the deep layers of the superior colliculus[66].

1.4.3 Lesion studies

Although lesion techniques might seem the obvious way of investigating the role of a brain area, there are probably fewer lesion studies than electrophysiological studies. For those studying the effects of surgical lesions on animals, the surgery may well be straightforward, but the behavioural training is time-consuming. The limitations on human studies are even more severe, as the number of patients suffering from brain damage is (thankfully) small, and these cases typically show partial damage to several brain areas rather than complete damage to a single area. In particular, it is very rare for human

subjects to suffer accidental damage which is restricted to A1 and which affects both hemispheres of the brain—this is important as it is possible to argue that left and right A1 should be considered together as a whole rather than independently (see below).

1.4.3a Lesion studies in animals

Early studies in animals concentrated on bilateral lesions of much of auditory cortex—both primary and secondary areas—with surprisingly mild results. Even these very large bilateral lesions do not impair cats' performance on a frequency discrimination task[17], although this task is affected by severing the connection between the inferior colliculus and the MGN[39]. Also, while the response pattern of lesioned cats on click rate discrimination (a temporal processing task) was different from normals, their overall accuracy was not impaired[29]. However, while these lesions do not affect discrimination of the frequency of pure tones, they do cause an impairment in discrimination of the pitch of anharmonic sounds[138]. Similar lesions also lead to a deficit on other "general-purpose" auditory tasks such as the ability to discriminate different spectro-temporal patterns[34, 61], and different directions of frequency modulation[65, 61]. Cats with bilateral auditory cortex ablation are also impaired when asked to correctly indicate, by approaching a target, which of two loudspeakers was the source of a sound[85].

Only a few animal studies have looked at large unilateral auditory cortex lesions, and these have been rather inconclusive. The ability to discriminate up and down FM sweeps seems only slightly, if at all, impaired after unilateral ablation[65]. Unilateral lesions also impair the detection of a tone against a background of noise, but only when the tone is contralateral to the lesion and the noise is ipsilateral; this impairment actually disappears if a second operation is carried out to ablate auditory cortex on the intact side[28].

One task that is definitely impaired following unilateral lesions, even if they are limited to A1, is that of sound localisation, although the impairment is more severe if auditory cortex beyond A1 is also affected[59]. Both cats and ferrets with unilateral A1 lesions display a localisation impairment, though it is restricted to sounds which are both brief (<100 ms) and which originate from a source in the contralateral hemifield[55, 59]. Furthermore, ablation of a single isofrequency band in A1 leads to a localisation impairment which is restricted to that frequency[56]. Bilateral A1 lesions also cause impairments on a localisation task, this time in both hemifields, although the ability to discriminate two sound source locations which are in opposite hemifields appears to be intact.

Surprisingly, sound localisation is the only task on which unilateral A1

lesions appear to have an effect, at least in cat and ferret. Bilateral A1 lesions have been shown to cause at least one other deficit, on the task of "gap detection". This involves detecting whether or not long-duration noise bursts are interrupted by a brief "gap", and is designed to test the limits of temporal processing. The minimum detectable gap was of longer duration for ferrets with bilateral A1 lesions than it had been for the same animals prior to surgery[64] (although in many of these animals the lesions were not entirely restricted to A1). That this processing of very short-duration auditory events (gap durations of <40 ms) should be impaired is interesting given that the localisation deficit associated with A1 damage is also restricted to short durations.

However, there is substantial evidence that observed effects of A1 ablation depend as much on the task being used as on the auditory discrimination it involves. In particular, if sound localisation is assessed with an approach-to-target task like those mentioned above, animals are impaired following A1 lesions, but if a withdrawal or "no-go" response is required then animals are not affected[40, 76]. This results, at least in part, from the existence of a subcortical pathway which is involved in processing aversive conditioning tasks[76], but nonetheless shows that A1 is not absolutely necessary for sound localisation.

Furthermore, the time which animals are allowed to recover from A1 damage may be critically important in determining the extent of any impairment. While bilateral ablation of the entire of auditory cortex has no lasting effect on either sound detection or frequency discrimination in the rat[60], temporary bilateral inactivation of A1 alone is sufficient to disrupt both of these tasks for many hours[130]. Thus, while A1 may not be necessary for frequency discrimination, it seems to be involved in a frequency discrimination task under normal circumstances. It is therefore possible that the same may be true for some or all of the other tasks described above.

1.4.3b Lesion studies in humans

In humans, the condition most closely associated with damage to A1 is that of "pure word deafness", which is caused by bilateral lesions[26]. It is probably most well-known because it is highly debilitating, involving an inability to understand speech. However, despite its name, it appears to have very little to do with language processing *per se*, instead being an impairment of temporal processing on the time scale required to distinguish different consonants[6, 97]. This ties in well with the time scale of the gap detection impairment in animals with bilateral A1 lesions (see above)[64].

Discrimination of tone frequencies is usually considered to be unimpaired

in humans with bilateral A1 damage[58, 57], although there is one recorded case which appears to show the opposite[132], and it is possible that earlier studies may have used insufficiently sensitive tests. There is also evidence that in the case of the more complex task of discriminating the pitch of a harmonic stimulus, unilateral damage to right A1 does lead to a deficit[147].

Localisation deficits have been reported in cases of human A1 damage, both bilateral and unilateral. One study found that subjects with unilateral A1 damage were less proficient than normals at detecting either ILDs or ITDs, while those with bilateral A1 damage were entirely unable to detect ITDs[145]. As with tests on animals, it is reported that this impairment is mainly in the side of space opposite to the lesion[24]. However, these effects are not very consistent between patients: there are many patients with cortical lesions apparently including A1 who show no signs of any impairment, while unilateral lesions are sometimes associated with bilateral deficits[148].

Interestingly, humans with callosal agenesis—a congenital absence of the fibre tract connecting the left and right cerebral hemispheres—are also impaired on a free-field localisation task[103]. Although this is not specific to A1 it does suggest that sound localisation processing involves auditory cortex as a whole, rather than the two hemispheres independently.

1.4.4 Summary

What is most suprising about the evidence presented above is that there is an apparent contradiction between the recording data and the lesion data. The most obvious feature of the electrophysiology of A1 is the frequency sensitivity and tonotopic organisation, and yet pure-tone frequency discrimination is apparently not impaired even in animals with large bilateral lesions of the whole of auditory cortex. By contrast, the most well-studied impairment following A1 damage is in sound localisation, and yet A1 neurons have very large and intensity-dependent receptive fields.

One explanation for this inconsistency is that recording studies may have misrepresented the extent of A1 involvement in sound localisation. Neuronal tuning is conventionally measured by simply counting the number of spikes evoked by presentation of a stimulus. However, there is a growing body of evidence that the precise timing, as well as the actual number, of spikes is important in conveying information through the nervous system[114]. It is also important to consider the amount of information carried by A1 as a whole, rather than by individual cells. Both of these approaches require some means of measuring the information carried by a neuron or group of neurons.

An alternative explanation is that the deficit in sound localisation seen in

lesion studies may have been misunderstood. The very task-sensitive nature of the impairment suggests this much. Indeed, sound localisation reflects several different aspects of auditory processing, potentially including such general tasks as spectral analysis and auditory scene analysis, as well as the processing of interaural cues, and any of these could be the "true" cause of A1's involvement. However, behavioural measures used in lesion studies to date have rarely done more than assess the presence or absence of an impairment.

1.5 Plan of this research

Given the mixed implications of electrode recording and lesion studies of A1, this research will attempt to use some new approaches to both types of study.

This work will build on previous studies of animals with A1 lesions by testing localisation ability over a wider range of sound sources. In contrast to previous studies, our apparatus allows animals to be tested over an azimuth range of 360° (chapter 2). This allows them to make errors due to front-back confusions, where previous studies have not, and thus potentially provides information about what exactly is being wrongly processed when errors are made (chapters 3 and 4). It is hoped that the pattern of errors of these animals will shed light on their underlying impairment. Additionally, unlike earlier studies on A1, we use drug-loaded polymer implants as a method of A1 inactivation. This technique is not only reversible, but it should also avoid unintentional surgical damage to non-A1 neurons which terminate in, pass through or lie near A1, e.g. thalamic afferents, "fibres of passage" and underlying white matter tracts.

We also add to earlier work on electrode recording. First, we record simultaneously from more than one neuron in A1 (chapters 5 and 6). Second, we use information theory (chapter 7) to analyse the amount of information contained in the number and timing of spikes in single-cell recordings (chapter 8). This should determine whether the apparent insensitivity of A1 to sound location is merely a result of the traditional methods of analysing electrophysiological data.

Part II

Sound localisation behaviour

Chapter 2

Behavioural methods

2.1 Care and training of animals

All of the animals used in the behavioural experiments were ferrets (*Mustela putorius*) and were pigmented rather than albino. All were housed in rack-mounted cages. They were fed on a high-protein carnivore diet, and were given free access to water except when being trained. They were also weighed regularly, both during training and normally, to ensure their weight stayed within normal limits.

The weighing is necessary because the animals were motivated to perform the various auditory tasks by water-depriving them and then using measured amounts of water as a reward during training. Water deprivation takes at least a day to have an effect on the animals' motivation, but they cannot be deprived for too long for obvious welfare reasons. Thus they were water-deprived and trained in 'runs' of behavioural experiments, each of which would typically last about two weeks. From twenty-four hours before the start of a run, the animal would be water-deprived and switched from normal to dry food. It would then receive all its water during twice-daily training 'sessions', where each session would consist of approximately 100 'trials' (depending on the task, see below). If an animal was not running enough trials to get a healthy amount of water, *ad lib* water would be given in its home cage. Once the run was finished, water deprivation would cease and the animal would be returned to a normal (non-dry) diet.

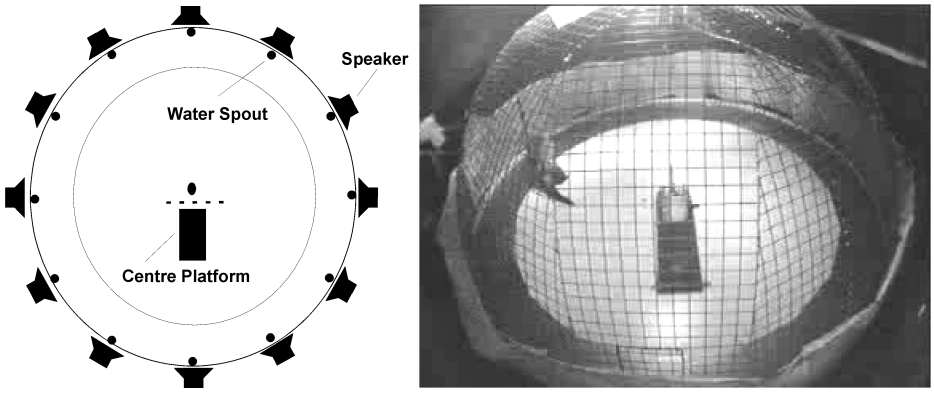

Figure 2.1: Views of the localisation chamber. All tasks described in the text used the layout of speakers and peripheral water spouts shown in the schematic plan (left). A ferret can be seen approaching a peripheral water spout in the top left corner of the photograph (right).

2.2 Localisation chamber

2.2.1 Construction

A diagram of the localisation chamber is shown in Figure 2.1. The apparatus consisted of a circular training area (radius 75 cm) of laminated chipboard, enclosed by a dome of 25 mm^2 wire mesh. This enclosure was in turn contained within a wooden chamber of base dimensions 190 cm by 190 cm. The roof of the chamber was flat over a 100 cm by 190 cm section; this section was 132 cm from the floor, with the rest of the roof sloping at a 33° angle to a minimum height of 72 cm. The chamber was double-walled, each wall being 15 mm thick medium-density fibreboard, with a 4 cm air gap between them. The inside walls and ceiling were lined with 50 mm Melatech sound absorbing foam (The Noise Control Centre, Leicestershire, UK).

2.2.2 Response detection and reward delivery

A water spout was set in the centre of the chamber, just behind which was a raised platform (37 cm by 12 cm by 3 cm) for the ferret to stand on. As well as the central water spout, up to fifteen peripheral response spouts could be mounted around the edge of the wire mesh enclosure, at the height of a ferret head. In the experiments described here the chamber was set up with twelve peripheral water spouts at 30° intervals. A series of metal plates was mounted around the edge of the enclosure, and there was also a metal plate mounted on the centre platform. Licks on a water spout could be detected

automatically as a change in resistance between the spout and the metal plate on which the animal would be standing, and this information about spout licks was then sent to a computer. Water could be delivered from the water spouts by the opening of computer-controlled solenoids (Flo-control, Valaeder Pneumatics, Cambridge UK). The opening times of the solenoids were calibrated so that each spout would deliver the same amount of water per opening (0.05–0.06 ml). The behaviour of the animals during training was monitored using closed-circuit television.

In some experiments, or in training, a wire grid was placed between the platform and the central water spout, with a hole cut into it large enough for the ferret to place its head through. This was intended to encourage the animals to take up a consistent body posture on approaching the centre spout. However, in many experiments this head-grid was deemed undesirable, because it made anterior water-spouts more difficult to approach than lateral or posterior spouts. Consistent body posture in experiments without a head-grid was achieved by removing the front portion of the metal plate on the central platform; this meant that a lick would only be detected if the animal's back feet were on the platform as it licked the spout.

2.2.3 Stimulus presentation

Stimuli were presented from one of up to fifteen speakers (Audax mid-range YN43W) which could be mounted in a variety of positions around the periphery of the enclosure, outside the wire mesh. The precise positioning of the speakers was hidden from the animal by a curtain of thin muslin. For the experiments described here only twelve speakers were used: these were mounted at 30° intervals within a horizontal plane level with the animal's head. The twelve speakers were thus adjacent to the twelve peripheral water spouts mentioned in section 2.2.2.

All stimuli were Gaussian noise bursts with 5 ms cosine-squared onset and offset gating. They were generated digitally on computer and processed using TDT System II hardware (Tucker-Davis Technologies, Gainesville, FL). Stimuli were digitally filtered using an FIR technique to compensate for differences in speaker characteristics. The size of the FIR filter was limited by the time required to compute the filtering, which is longest for long-duration sounds; hence the quality of filtering was highest for the shortest sounds (e.g. a 51-point filter was used for 1000 ms stimuli as compared to a 501-point filter for 40 ms stimuli). In most experiments the desired output was flat white noise (0.5–30 kHz), but the FIR filters could also be modified to create low-pass, high-pass or band-pass noise. The sampling period for DA conversion was 12 μs, and after conversion the resulting signals were fed through an

anti-aliasing filter with nominal cutoff frequency of 30 kHz. They were then attenuated before being fed to the relevant speaker. The amount of attenuation was computer-controlled, and, if desired, the stimulus intensity on each trial could be randomly chosen from one of 5 possible levels: in early experiments these levels were at 2 dB intervals between 66–74 dB SPL, later they were at 7 dB intervals between 56–84 dB SPL. Measurements necessary to calibrate sound levels and produce the required filtering were made using a Brüel and Kjær type 4134 $\frac{1}{2}$-inch condenser microphone attached to a Brüel and Kjær type 1618 measuring amplifier.

This apparatus could also produce 'continuous' noise, starting when the animal triggered a trial and ending when it made a response. Such a stimulus was produced by looping shorter stimuli; the resulting noise is not truly random, but it was only used in localisation training and for 'easy' trials (see section 2.3) so this should not be problematic.

2.3 Localisation procedure

This task was run in the chamber shown in section 2.2.

An animal would initiate a trial by licking the centre spout. In some experiments the animal was required to lick this spout continuously for a length of time known as the 'spout delay'. The purpose of this delay was to temporally separate the onset of licking and the onset of the sound stimulus so that the animal would attend as much as possible to the sound when it was triggered. For experiments that used a spout delay, its duration was determined randomly from within a preset range before every trial; typically this range was 0.5–2 s.

The lick on the centre spout triggered a broadband noise burst from one of the peripheral speakers. To complete the trial, the animal then had to respond by licking one of the peripheral water spouts. The response would be scored as correct if the animal licked the peripheral water spout beneath the speaker from which the sound had been presented. The response would be scored as incorrect if any other peripheral water spout were licked. Licking the centre spout again had no effect until either a response (correct or incorrect) was made or the trial was timed out (see below). Correct responses resulted in a reward being delivered at the spout as it was licked. Incorrect responses elicited no reward. Any response, correct or not, ended the trial: no further licks on the peripheral spouts had any effect, and the animal had to lick the centre spout again to initiate another trial. If the animal did not lick any of the peripheral spouts within 15 s of licking the centre spout, the trial would be timed out and the animal would have to start a new trial.

The stimulus was presented from a randomly chosen speaker on each trial. However, if an animal made an incorrect choice on a trial, the sound would be presented from the same speaker again in each subsequent trial until the animal chose correctly. After 3 such correction trials, the animal would be presented with an 'easy' trial, which consisted of continuous noise rather than a noise burst. Again, the easy trials would be repeated until the animal chose correctly. Correct responses on correction trials were rewarded as on normal trials. However, it is important to note that responses on these correction or 'easy' trials were not considered for the purposes of analysing animals' performance.

Animals were occasionally given rewards at the centre spout in order to encourage them to lick it properly; this would happen at random, occurring once in every 20 trials on average. Again, trials on which this occurred were not used in calculating performance statistics.

Animals were run for approximately 100 trials in a session, not including correction trials or centre spout reward trials.

2.4 Preparation and use of Elvax implants

2.4.1 Manufacture of Elvax implants

All Elvax implants used in this study were supplied to us "ready-made" by Dr. Adam Smith. The manufacturing methods described here are those used in his laboratory. Similarly, the investigation of muscimol release characteristics was carried out by Dr. Adam Smith and Dr. Ian Thompson.

2.4.1a Elvax washing and storage

The polymer Elvax 40-W, also known as Elvax 40P, was provided in the form of beads (a gift from Du Pont UK). These needed to be washed in alcohol in order to remove traces of butylhydroxytoluene. Washing was carried out in a series of alcohol solutions, for 24 hours per solution: the first five changes were all of 95% alcohol, the next five were in increasing concentrations up to 100% alcohol. The beads were then filtered, dried, and stored in the dark at room temperature.

2.4.1b Muscimol Elvax

0.5 g of washed Elvax beads were dissolved in dichloromethane (Sigma-Aldrich, Poole, UK) by vortex mixing at room temperature to give a 10% w/v solution. An aqueous solution of muscimol (Tocris Cookson, Langford,

UK) was produced by dissolving muscimol in one equivalent of 10 M aqueous NaOH, and adding double-distilled water as necessary to give 100 μl of 3.75 M muscimol solution.

The aqueous muscimol was then added to the Elvax/dichloromethane solution to give a 75 mM suspension of aqueous muscimol in Elvax/dichloromethane. 50 μl of 5% aqueous Fast Green (Allied Chemical, Morristown, NJ) was added to the suspension in order to increase the visibility of implants made from this mixture. The suspension of aqueous solutions in dichloromethane was then agitated on a vortex mixer for 10 mins. before pouring without delay into a glass mould embedded in dry ice.

After leaving the Elvax solution to freeze for 20 mins., the resulting sheet was removed from the mould and freeze-dried overnight at -60°C and 10^{-4} atm. The dried block was then sectioned on a microtome into 200 μm slices, and these were stored on filter paper at 4°C.

2.4.1c Blank Elvax

Blank Elvax for control experiments was prepared in the same way as muscimol Elvax, except that the 100 μl of aqueous solution added to the Elvax/dichloromethane solution contained nothing but double-distilled water.

2.4.2 Implantation of Elvax

2.4.2a Elvax preparation

Elvax pieces were rehydrated and trimmed to the desired size well in advance of surgery, as particularly high concentrations of drug are released on initial rehydration, especially at newly-cut edges[129]. The slices were washed in warm (37°C) phosphate buffered saline at least 3 days before use, and trimmed at least 24 hours prior to implantation.

2.4.2b Elvax implantation

All surgical procedures were carried out under aseptic conditions.

Each ferret was anaesthetised with approximately 2 ml kg^{-1} of i.m. Saffan (9 mg ml^{-1} alphaxolone with 3 mg ml^{-1} alphadolone acetate). Once unconscious the animal was maintained in a state of areflexive anaesthesia with further i.m. injections of 0.3–0.4 ml Saffan as needed (roughly every 30 minutes, depending on the animal). A square window was drilled into the skull, and the square piece of bone removed and kept in sterile saline. A horizontal slit was made in the dura above the ectosylvian gyrus. The piece of Elvax was then inserted through the slit and carefully manoeuvred into position

over the estimated position of A1, which is within the ectosylvian gyrus.[1] If possible, an edge of the Elvax piece was inserted under the skull to hold it in place. The previously-removed piece of cranium was replaced, and covered with a small piece of clingfilm—the clingfilm was used to prevent blood clots or tissue from adhering to the loose piece of cranium. Size 4/0 or 5/0 coated vicryl sutures were used to rejoin the temporal muscle to the sagittal crest. Size 3/0 silk sutures were used to close the incision in the scalp. The animal was given 0.1 ml Dopram (20 $\mathrm{mg\,ml^{-1}}$ doxapram hydrochloride), 0.15 ml ampicillin and 0.15 ml vetergesic, all delivered i.m. Marcaine and dermisol cream were also applied to the wound. The animal was then allowed to recover consciousness.

Animals were not trained on any behavioural tasks until at least 6 days after implantation, to allow for healing. Sutures were taken out whenever necessary, typically 2–3 weeks after the operation.

2.4.2c Elvax removal

Procedures for Elvax removal were similar to those for implantation. Where the original craniotomy was obvious, the new craniotomy was made in the same place. Clingfilm was not used after replacing the bone window, however. The same sutures and post-operative drugs were used as for Elvax implantation.

2.4.3 Measurement of Elvax properties

2.4.3a Dynamics of muscimol release

To allow the measurement of muscimol release from sheets of Elvax, trace amounts of 19.6 $\mathrm{Ci\,mmol^{-1}}$ tritiated muscimol (NEN Life Science Products, Boston, MA) were incorporated into some muscimol preparations. An aliquot of ^{3}H-muscimol was added to a weighed sample of muscimol in a microcone vial, before freezing and then freeze-drying the vial. The dilution factor of the incorporated radiolabel varied between 1:50 000–1:10 000. After freeze-drying, the contents of the vial were treated as normal (section 2.4.1b).

To measure the muscimol release, sheets containing the radiolabel were incubated in 0.01 M phosphate-buffered saline (pH 7.4) at 37°C, and the bathing solution was changed every 24 hours. The amount of radioactivity

[1] This estimate of the position of A1 is based on previous electrophysiological and imaging studies carried out both in this laboratory[116, 134] and elsewhere[101, 124]. We did not determine the position of A1 in the actual animals under investigation as it is not possible to do so without the risk of causing damage.

released into the bathing solution during each 24 period was determined by liquid scintillation counting. Release measurements were made both on fresh Elvax sheets, to determine the *in vitro* release profile, and also on Elvax sheets that had been removed from *in vivo* placements.

2.4.3b Effects on cortical activity

The actual effects of Elvax implants were assessed electrophysiologically. A total of 14 pigmented adult ferrets were used in a series of experiments on visual cortex, of which 8 received muscimol Elvax implants, 2 received blank Elvax implants, and 4 were controls. The implants were made over primary visual cortex (V1), which has several advantages: V1 neurons are very responsive to stimuli, and they also have notably different response properties to the thalamic afferents which terminate in the same area. Like the Elvax manufacture and measurement of muscimol release properties, this work was carried out in the laboratories of Dr. Adam Smith and Dr. Ian Thompson, and is only summarised here. In addition to the experiments on V1, two other animals were used in similar experiments on A1, in which both were implanted with muscimol Elvax. This later study took place some time after the author finished experimental work on this thesis, and so it is also not described in detail.

In all these experiments, Elvax pieces were first implanted in a manner similar to that described in section 2.4.2b. Then, after a certain amount of time, electrophysiological recordings were made both around the edges and underneath the implants—the latter being possible because the implants had been specially cut so that a window could be removed from the centre. The methods used for A1 electrophysiology were substantially similar to those given in chapter 5, with both pure tones and broadband noise being used as stimuli, while the protocol for V1 recording was as follows:

Animals were anaesthetised with i.m. Saffan, and then given 60 μg of atropine sulphate i.p. Following intubation of the trachea and the insertion of an i.v. line into the radial vein, anaesthesia was maintained with further i.v. injections of Saffan during the initial surgery to expose visual cortex. The animal was then paralysed with i.v. Flaxedil (gallamine triethiodide), injected at 10 $\mathrm{mg\,kg^{-1}\,hr^{-1}}$, and artificially ventilated with 0.5–1.0% Halothane in 3:1 N_2O:O_2 for the duration of recording. Tidal volume and ventilation rate were adjusted to give an end-tidal CO_2 concentration of $\sim$4%, but were typically 10–12 ml and 35–40 $\mathrm{strokes\,min^{-1}}$, respectively. Anaesthesia depth was monitored using an EEG, and core body temperature was maintained between 37–38°C.

Recordings were made using tungsten-in-glass micro-electrodes with a

5 or 10μm tip. Visual stimuli were presented through the eye contralateral to recording, after first dilating the pupil with topical application of 1% atropine sulphate, and then applying a clear zero-power contact lens. Visual receptive fields were mapped using hand-held stimuli, before using computer-generated monochrome sine-wave gratings to quantify units' orientation preferences.

Chapter 3

Performance of normal ferrets on an auditory absolute localisation task

3.1 Introduction

The challenge of psychophysics is to find metrics which are easy to measure, easy to analyse quantitatively, and yet which are also psychologically relevant.

The psychophysical metric most commonly used in studies of free-field sound localisation in animals is the minimum audible angle (MAA). This is a measure of how far apart two sound sources have to be in order for their positions to be reliably discriminated. The main advantage of this task is that the statistics associated with discriminating two stimuli are well-understood. The disadvantage is that discriminating two known stimulus locations is not the same thing as judging the absolute location of a stimulus. In fact, prior knowledge of possible speaker positions is known to affect the perception of stimulus location in humans[94]. Another disadvantage is the large amount of testing needed to provide an accurate MAA measurement. In most studies, the absolute value of the MAA is less important than how localisation ability varies under different conditions, and measuring MAAs for each of these conditions is an inefficient way of determining this.

There are other commonly-used tests of animal sound localisation, however. In one such task, an animal has to indicate the perceived location of a sound source by physically approaching it[55, 56]. This task is usually referred to as "absolute localisation", to contrast it with the "relative localisation" in an MAA task. It has the advantage that the subject's response

is directly related to its perception of sound source location. A drawback, though, is that in order for animals to approach a target, it must be mounted on the floor—this restricts approach-to-target experiments to studies of localisation along the horizon. Another disadvantage is that in order to train animals on a task it is usually necessary to reward them for making "correct" judgements. Measurements of localisation ability are then limited to the resolution of this reward procedure. Because of this, "absolute localisation" is a slightly inaccurate term, as the animal is faced with a limited number of response choices rather than being free to respond anywhere. A way of solving this problem is to quantify head and/or eye movements made towards a sound source[78, 48]. This has the advantage of being a truly absolute measurement of localisation. However, while it is possible to instruct humans to make head or eye movements voluntarily, in animal studies these movements are more likely to be "reflexive" orienting responses, in which case they probably reflect something other than the conscious perception of target location. It is also possible to use aversive conditioning schedules, where animals must avoid an electric shock that is signalled by the location of a sound source in some way[75, 41]. As with orienting head and eye movements, though, it is possible that aversive conditioning measures a less cognitive mechanism of sound localisation than an approach-to-target test.

In this study, we opt for an approach-to-target absolute localisation procedure. This is mainly because the evidence collected here is intended as a control for the effects of A1 inactivation, and it is known that A1 lesions can affect animals' performance on approach-to-target tasks[76].

We also address a fundamental weakness in most animal sound localisation experiments to date, which is that they have only used sound sources positioned in the front hemifield and along the horizon. As discussed elsewhere (section 1.3.2b), several interaural cues are ambiguous, and may only serve to locate a target somewhere in a locus known as the "cone of confusion". However, only one position within this cone of confusion is both in the front hemifield and along the horizon, so a cue which would normally be ambiguous is abnormally informative when sound source location is restricted in this way. Although for practical reasons we follow earlier studies in using sound sources placed along the horizon, we do use positions in both the front and rear hemifields. This means that animals which are using ambiguous interaural cues may exhibit "front-back errors" on our task when in earlier studies they would have performed perfectly.

3.2 Methods

3.2.1 Numerical conventions

Throughout this chapter, azimuth is expressed in degrees, from 0° to 360°: 0° is straight ahead, 90° is to the right, 180° is directly behind, and 270° is to the left. Experimental measurements are expressed either to the precision of measurement, or to 3 significant figures, whichever is lower. Other numeric quantities are treated similarly, although values for statistical probability, p, are sometimes given to 3 decimal places only, for reasons of space and clarity.

In analysis of variance (ANOVA) tables, SS is an abbreviation for Sum of Squares, MS is an abbreviation for Mean Sum of Squares, and significant results are shown in boldface. The values used in ANOVAs are often proportional data, e.g. the proportion of trials to which an animal responds correctly. Such data typically show variance which is maximum near 0.5 and minimum near 0 or 1, and this heterogeneity of variance must be corrected for. We used the "arcsine correction", where the inverse sine of the square root of the proportion is used instead of the proportion itself[139].

3.2.2 Animals

Six ferrets were involved in training and testing, and these can be divided into two groups. Group PILOT consisted of animals F9731, F9732 and F9733, all of which were female. They were trained on a schedule which was designed to assess their performance at its peak. They were also among the first animals to be trained in our absolute localisation chamber, and so were used to evaluate many of the methods used in training. Group NO-OP consisted of animals F9934 (male), F9935 (male) and F9936 (female). These animals were put through a much more well-defined regime designed to assess their performance both during training and at the end of training when they had achieved a high level of performance. None of the animals in either of these groups had been involved in any experiments before the work described here.

3.2.3 Apparatus

The apparatus used was the localisation chamber described in section 2.2 and shown in figure 2.1. Briefly, this was a circular chamber containing one water spout mounted just in front of a central raised platform, as well as twelve more water spouts mounted at 30° intervals around the edge of the chamber. Above each peripheral water spout was a loudspeaker. A computer was used

to trigger presentation of stimuli from the loudspeakers, and to monitor when an animal licked any of the water spouts.

3.2.4 Behavioural training and testing

The task on which all animals were tested was the localisation task described in detail in section 2.3. In this, an animal would trigger a trial by licking the central water spout. This would cause a burst of noise to be presented from one of the twelve peripheral speakers; in the experiments described here the noise used was flat white noise in the range 0.5–30 kHz. The animal's task was then to lick the peripheral water spout mounted underneath whichever speaker had just been activated. A correct response was rewarded with a measured amount of water, presented from the peripheral spout. An incorrect response was not rewarded, and on subsequent trials ("correction trials") the stimulus was presented from the same speaker until a correct response was given. Responses to correction trials were excluded when calculating animals' performance. Also excluded were the occasional trials on which animals were rewarded at the centre spout—such trials were used to encourage the animals to make proper contact with the centre spout.

As explained in section 2.1, both training and testing of animals was carried out in a series of 'runs', each lasting about 2 weeks and during which the animal would be water-deprived. For a particular animal, at least two days (and typically more) would separate each run from the next. During a run, each animal would be put through two training or testing sessions a day, and each session would consist of roughly 100 trials (depending on the precise schedule).

3.2.4a Group PILOT

Spout training Initial training involved getting the animals accustomed to the system of response detection and delivery. Animals would have to lick each water spout in turn, and would receive a set amount of water from each spout for doing so. Particular emphasis was placed on getting them to take up a consistent posture when licking the centre spout, as in the full task this is what triggers sound presentation.

Sound training Once the animals had learnt how to use the water spouts, sound was introduced to the regime. The animals were trained using a simplified version of the full task, so that they would get used to the general procedure of running trials. For the first run only four speakers—at 0°, 90°,

noise duration	number of sessions or trials
2 s	1 session
1 s	300 trials
750 ms	1 session
500 ms	300 trials
400 ms	1 session
300 ms	1 session
200 ms	300 trials
100 ms	300 trials
40 ms	300 trials

Table 3.1: Typical testing schedule for ferrets in Group PILOT. Sessions were run in the order listed—top to bottom. Animals were run at each stimulus duration either for the specified number of sessions, or until they had run approximately the specified number of trials.

180° and 270° azimuth—were used, with a stimulus that consisted of looped noise rather than a brief noise burst.

Subsequent runs introduced more aspects of the full task. The second run introduced more speakers. Shorter stimulus durations, of a few hundred milliseconds, were also introduced; this change was made gradually over a number of sessions. A further run was used to train the animals with stimuli as short as 40 ms. The next run added 'easy trials'. The run following that introduced a 'spout delay', which was also introduced gradually as the animals initially found it confusing.

In all, the animals in this group went through 5 runs of training, including both spout training and sound training, over the course of ∼80 days.

Testing Each of the animals was then tested on three runs of the full task. Stimulus duration was varied during the course of each run. Because performance at the shortest durations was quite poor and thus discouraging for the animals, the stimulus duration was reduced in small steps through the course of a run. Behavioural data were collected from sessions with stimuli lasting 1 s, 500 ms, 200 ms, 100 ms and 40 ms; approximately 300 trials (usually 3 sessions) were run at each of these key stimulus durations. Sessions at intermediate durations were also used to avoid demotivating the animals. The schedule for a typical testing run is shown in table 3.1, although more sessions were run if an animal was performing less well than expected. One animal (F9732) was put through an additional run after the three normal testing runs, in which the lights in the chamber were switched off.

noise duration	number of sessions or trials	
	first run	second and third runs
looped	2 sessions	0 sessions
2 s	2 sessions	1 session
1 s	2 sessions	2 sessions
750 ms	1 session	1 session
500 ms	300 trials	300 trials
400 ms	1 session	0 sessions
300 ms	1 session	0 sessions
200 ms	300 trials	300 trials
100 ms	300 trials	300 trials
40 ms	400 trials	400 trials

Table 3.2: Testing schedule for ferrets in Group NO-OP. Sessions were run in the order listed—top to bottom. Animals were run at each stimulus duration either for the specified number of sessions, or until they had run approximately the specified number of trials.

Stimulus intensity was varied from trial to trial (except for looped noise trials), with the intensity on each trial being randomly selected from the following values: 66 dB SPL, 68 dB SPL, 70 dB SPL, 72 dB SPL, 74 dB SPL. The value of the spout delay was randomly generated for each trial, in the range 0.5–2 s. The headgrid was not used during testing.

3.2.4b Group NO-OP

Spout training Like those in Group PILOT, animals in Group NO-OP were first familiarised with the chamber and its system of water spouts by means of several (6–8) sessions in which sound stimuli were not used. In these sessions the animals were given rewards for visiting each of the chamber's water spouts in turn. Extra rewards were given for standing properly on the platform and licking the centre spout, as the animals need most encouragement to do this. For animals whose stance on the platform was particularly bad, the headgrid was used during this training.

Testing The animals were then put through three runs of the full sound localisation task. There was no intervening "sound training" as the performance of these animals during learning was as important as their peak performance. The number of sessions or trials which were run at each stimulus duration is given in table 3.2.

No spout delay was used during this testing. This is because spout delay

has to be introduced over several sessions, which would make it difficult to use differences in performance between sessions as an indicator of improvements in sound localisation. Stimulus intensity in all these sessions (except those with looped noise) varied randomly from trial to trial, having one of the following values in each trial: 56 dB SPL, 63 dB SPL, 70 dB SPL, 77 dB SPL, 84 dB SPL. The headgrid was not used during testing.

3.3 Results

3.3.1 Overall performance

Ferrets' performance on the localisation task can be seen, in a fairly "raw" form, in figure 3.1. This figure shows the proportion of trials on which each ferret responded correctly, broken down by session. The division of ferret testing into 'runs' of about 10 days worth of twice-daily sessions can clearly be seen. It is also obvious that the animal's performance deteriorates as the stimulus duration is decreased during each run.

Because the main factor determining animals' performance seems to be stimulus duration, a better way of presenting the data in figure 3.1 is to group trials by duration instead of by session. In doing this, we focus on key durations of 1 s, 500 ms, 200 ms, 100 ms and 40 ms. The data from figure 3.1 are re-plotted in figure 3.2, except this time trials in the same run are grouped together by stimulus duration, and trials from the non-key durations are excluded. The effect of stimulus duration can clearly be seen. By contrast, there seems to be very little difference between the different testing runs, at least for the animals in the PILOT group. In particular, F9732's performance on her fourth run, which was in the dark, certainly shows no impairment caused by the lack of visual cues. Two of the animals in the NO-OP group (F9935 and F9936) do appear to show a slight improvement—presumably a learning effect—from the first to the second run, but the same is not true of the third animal (F9934).

We assessed the significance of these results using an analysis of variance (ANOVA). As in figure 3.2, trials were grouped together by run, and by stimulus duration within each run, and the proportion of those trials on which the animal gave a correct response was used as the measure of performance. This performance measure was arcsine-corrected to avoid problems with variance heterogeneity, as described in section 3.2. Stimulus duration and run were then used as within-subjects factors in the ANOVA. The PILOT and NO-OP data were analysed separately, and F9732's last run was excluded as there was no comparable run for the other two animals in the PILOT

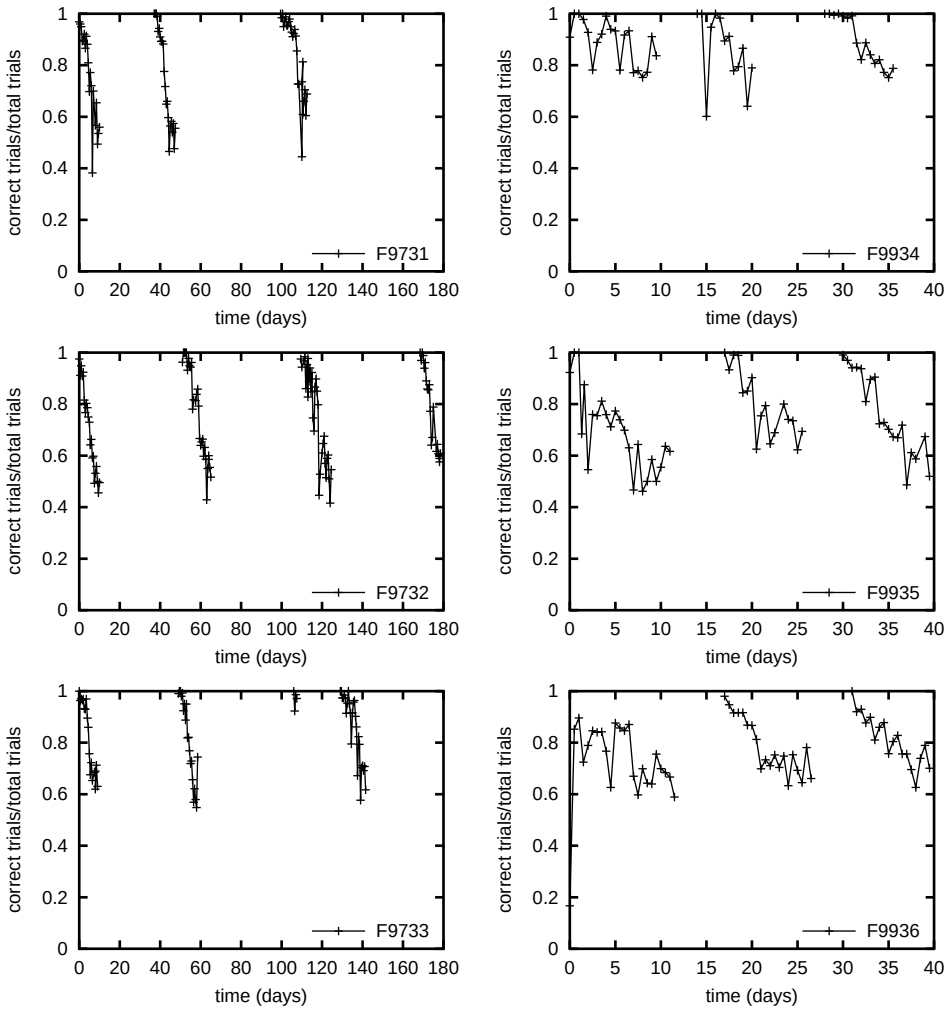

Figure 3.1: The performance of ferrets in group PILOT (left column) and NO-OP (right column), measured as the proportion of trials to which the correct response was given. Each point represents the performance on a particular session, and sessions within the same run are connected by lines. Sessions containing fewer than 10 trials are not shown.

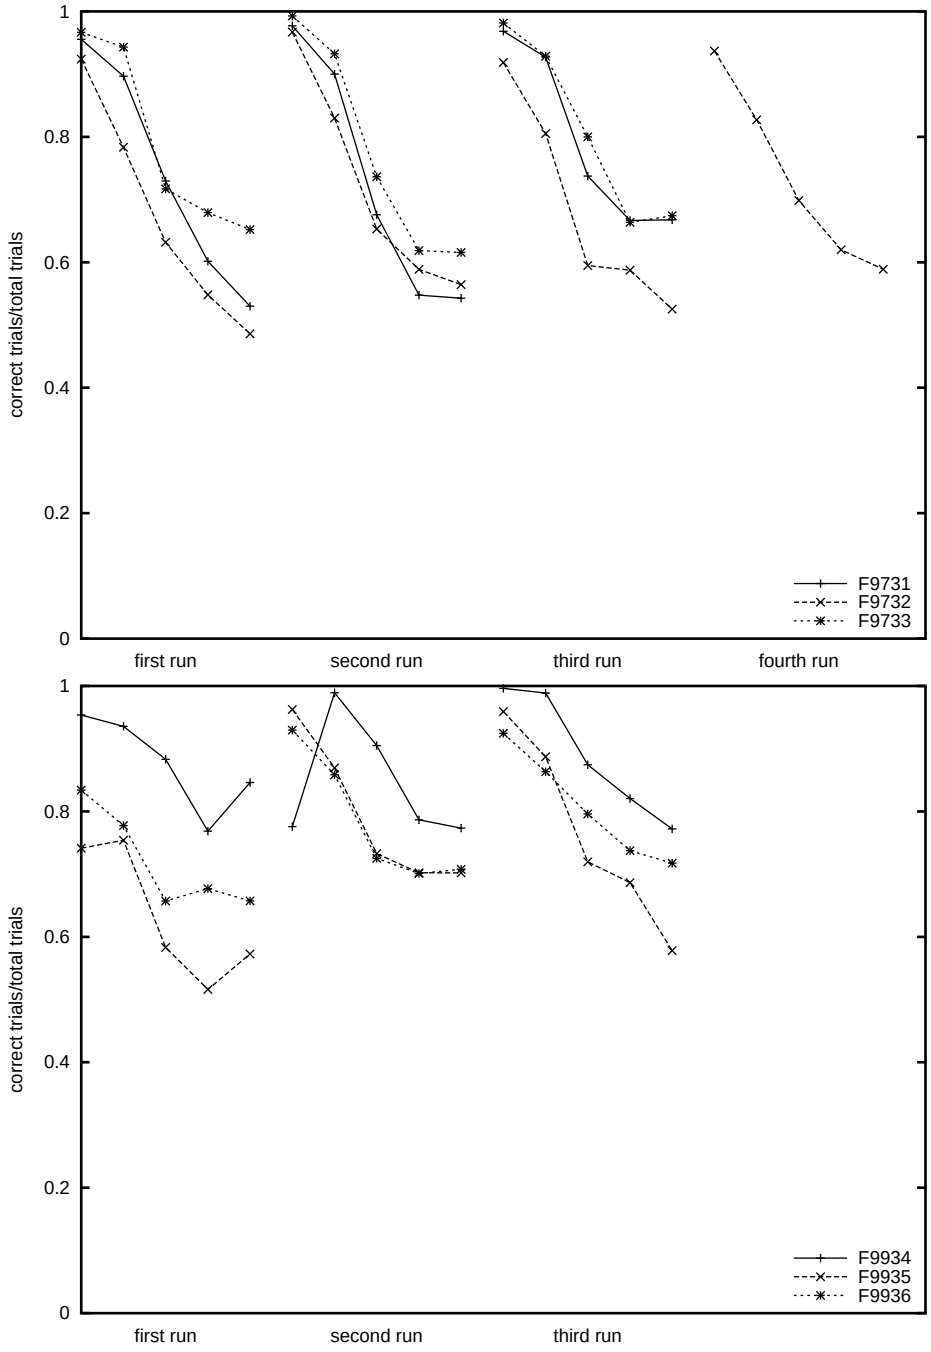

Figure 3.2: The performance of ferrets in group PILOT (top) and NO-OP (bottom) at key durations in each run. Performance is measured as the proportion of trials to which the correct response was given. Within each run, trials are grouped by stimulus duration, so the first data-point in each run represents trials at 1 s, the second represents trials at 500 ms, and so on for 200 ms, 100 ms and 40 ms.

Group PILOT

source	df	SS	MS	F	p
duration	4	1.79	0.447	232	**<0.001**
error(duration)	8	0.0154	0.00192		
run	2	0.00790	0.00395	0.869	0.486
error(run)	4	0.0182	0.00454		
duration*run	8	0.0177	0.00221	3.05	**0.028**
error(duration*run)	16	0.0116	0.000726		

Group NO-OP

source	df	SS	MS	F	p
duration	4	0.618	0.154	19.8	**<0.001**
error(duration)	8	0.0624	0.00780		
run	2	0.0965	0.0483	3.00	0.160
error(run)	4	0.0642	0.0161		
duration*run	8	0.0484	0.00605	1.04	0.446
error(duration*run)	16	0.0928	0.00580		

Table 3.3: ANOVA for animals in Groups PILOT and NO-OP, showing the differences in performance associated with stimulus duration and testing run. Performance is measured as the arcsine-corrected proportion of trials to which the correct response was given.

group. The results, detailed in table 3.3, show that the effect of duration was significant for both the PILOT and NO-OP groups, while the effect of run was not significant in either case, and the interaction between duration and run was only significant for the PILOT animals.

3.3.2 Performance at different positions

As well as looking at overall performance, we are able to break down animals' scores according to the position of the stimulus. Figure 3.3 shows the performance of animal F9731—an animal in the PILOT group— separately for each speaker position. As can be seen, performance tends to be highest at the front midline, and falls off gradually for stimuli which are at either side of the animal. Performance to the animal's rear is lower than to the animal's front, but shows a similar pattern of being slightly higher at the midline than at more lateral positions. Figure 3.4, shows data from the same animal, but this time plotted to show the distribution of responses to different stimuli. It is apparent from this that most of the responses are either to the correct position, or to positions adjacent to the correct position—in other words,

errors are typically small.

Data from another animal, this time F9934 from the NO-OP group, are shown for comparison in figures 3.5 and 3.6. The generally higher performance of this animal, as shown in the overall performance figures (figure 3.2), is reflected in the plots of performance *versus* speaker position (figure 3.5). However, the manner in which performance varies with stimulus position is quite similar to that of F9731; in particular, best performance is directly in front of the animal, and worst performance is at rear lateral positions. However, compared to F9731, more of the responses made by F9934 are to positions which are not adjacent to the correct position—this can be seen in figure 3.6 in the responses to stimuli at 150° and 210°. Such errors fit the profile of "front-back confusions" discussed in greater detail in section 3.3.3.

As this study is intended to act as a control for the effect of unilateral A1 inactivation (among other things), we are particularly interested in whether there are any differences in localisation ability between the right and left hemifields. This is because unilateral A1 inactivation is expected to have an effect primarily in the hemifield contralateral to the inactivation. It is difficult to compare the left and right sides in figures 3.3–3.5, but figure 3.7 shows left and right hemifield scores for all of the normal animals. As can be seen, there is no obvious superiority of performance in one hemifield rather than the other.

This difference, or lack of it, between the performance in the two hemifields was assessed statistically with an ANOVA. Trials were grouped by stimulus duration within each run, and also by stimulus hemifield (trials where the stimulus was directly in front of or behind the animal were excluded). Thus the within-subjects factors for the ANOVA were stimulus hemifield, stimulus duration and testing run. As with the analysis of overall performance, we used the proportion of responses which were correct as an indicator of performance, and then applied an arcsine-correction. The results of the ANOVA are given in table 3.4. As expected, they do not show any significant effect of stimulus hemifield, either as a main effect or in interaction with any other factors. Indeed, the only significant effect in this analysis, for either group, was that of stimulus duration.

Although the ANOVA reveals no effect of stimulus hemifield on performance, it could be the case that the animals do not all have the same "handedness", in which case the between-hemifield differences might have cancelled one another out in the ANOVA. To address this, the animals' scores were also examined individually for signs of between-hemifield differences in performance. As before, the measure of performance used was the proportion of trials where a correct response was given. The raw data were pooled across the different testing runs, although trials using different stimulus durations

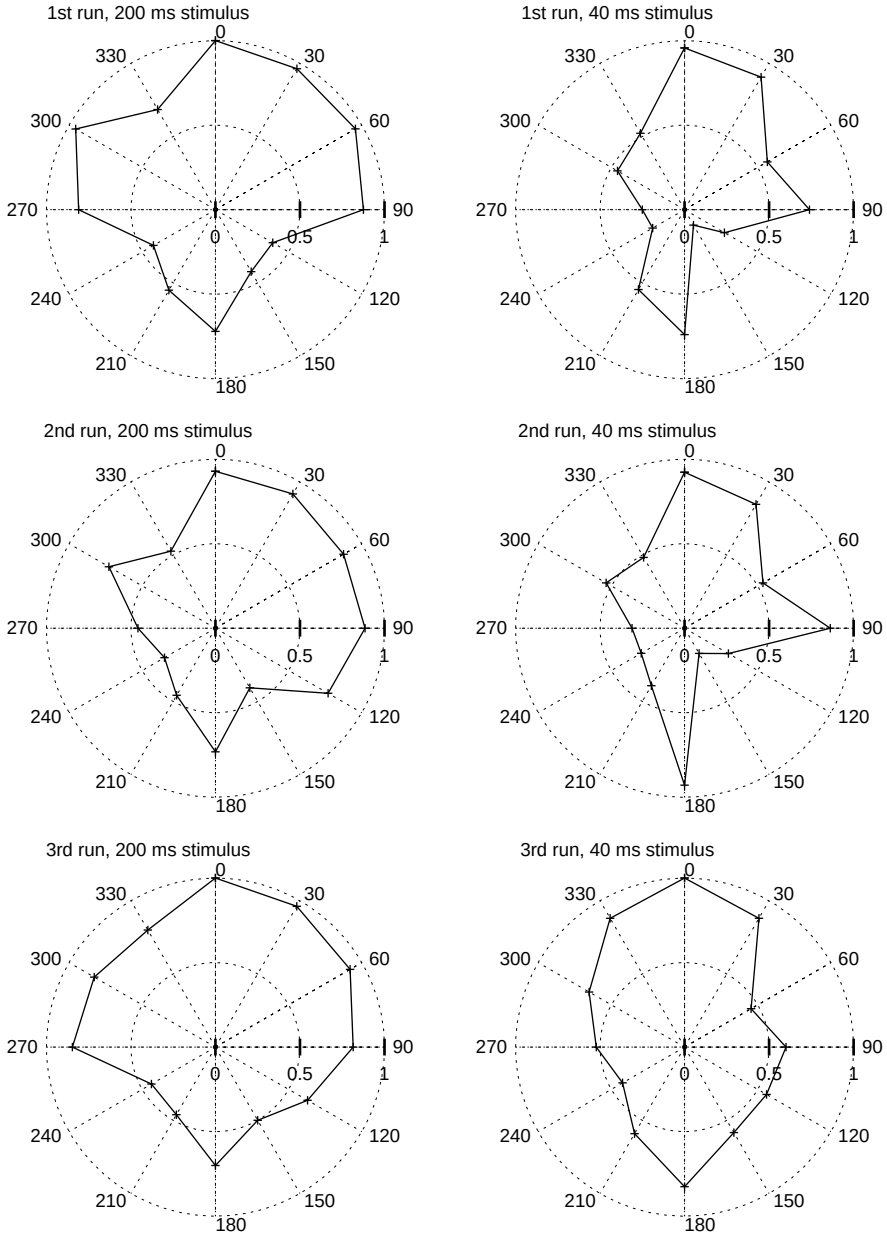

Figure 3.3: The performance of an animal from the PILOT group (F9731) in response to stimuli presented from different positions. Performance is shown for 200 ms (left column) and 40 ms (right column) stimuli, and over all three testing runs (in chronological order, top to bottom). The angle of each point indicates the stimulus position, and the radius of each point represents the proportion of trials on which the correct response was given.

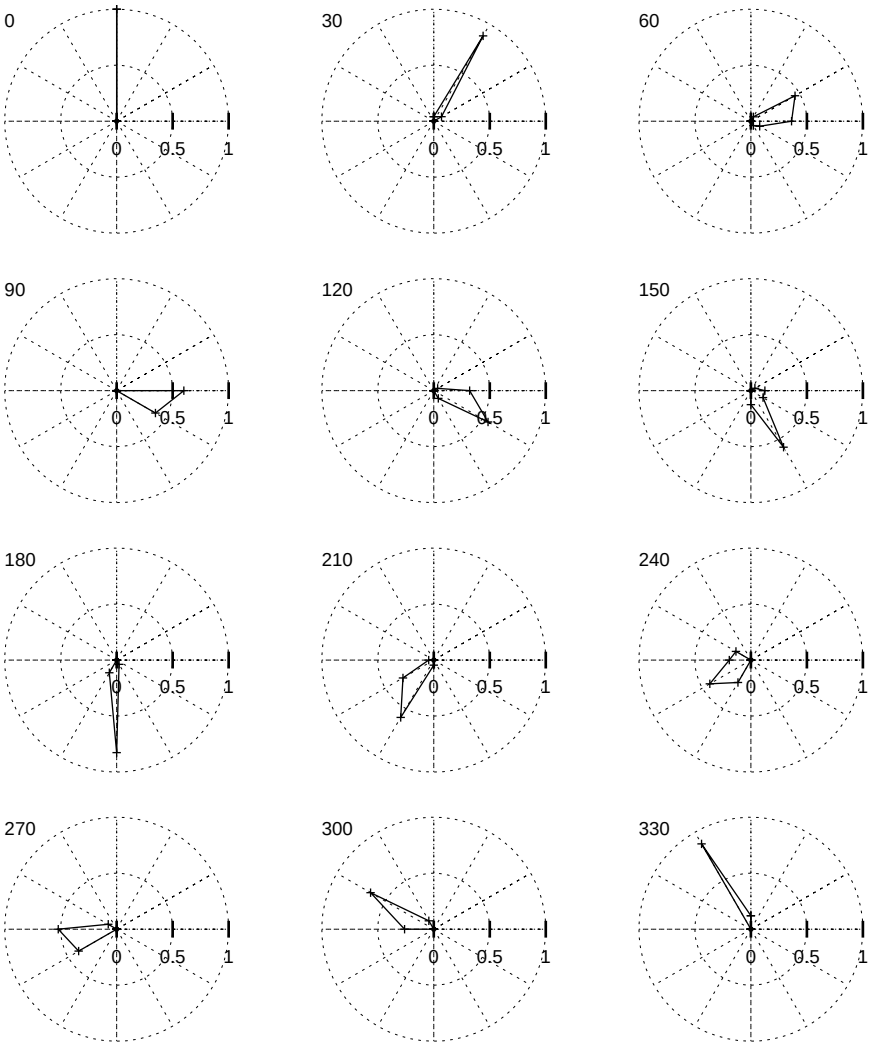

Figure 3.4: The distribution of responses to stimuli emanating from different positions, for an animal from the PILOT group (F9731). Different plots show the responses to different stimulus positions, with the label on each plot giving the stimulus azimuth. The angle of each point indicates the response position, and the radius of each point represents the proportion of trials on which the animal responded at that position. Data shown are from responses to 40 ms stimuli on the third testing run (as in the bottom right panel of figure 3.3).

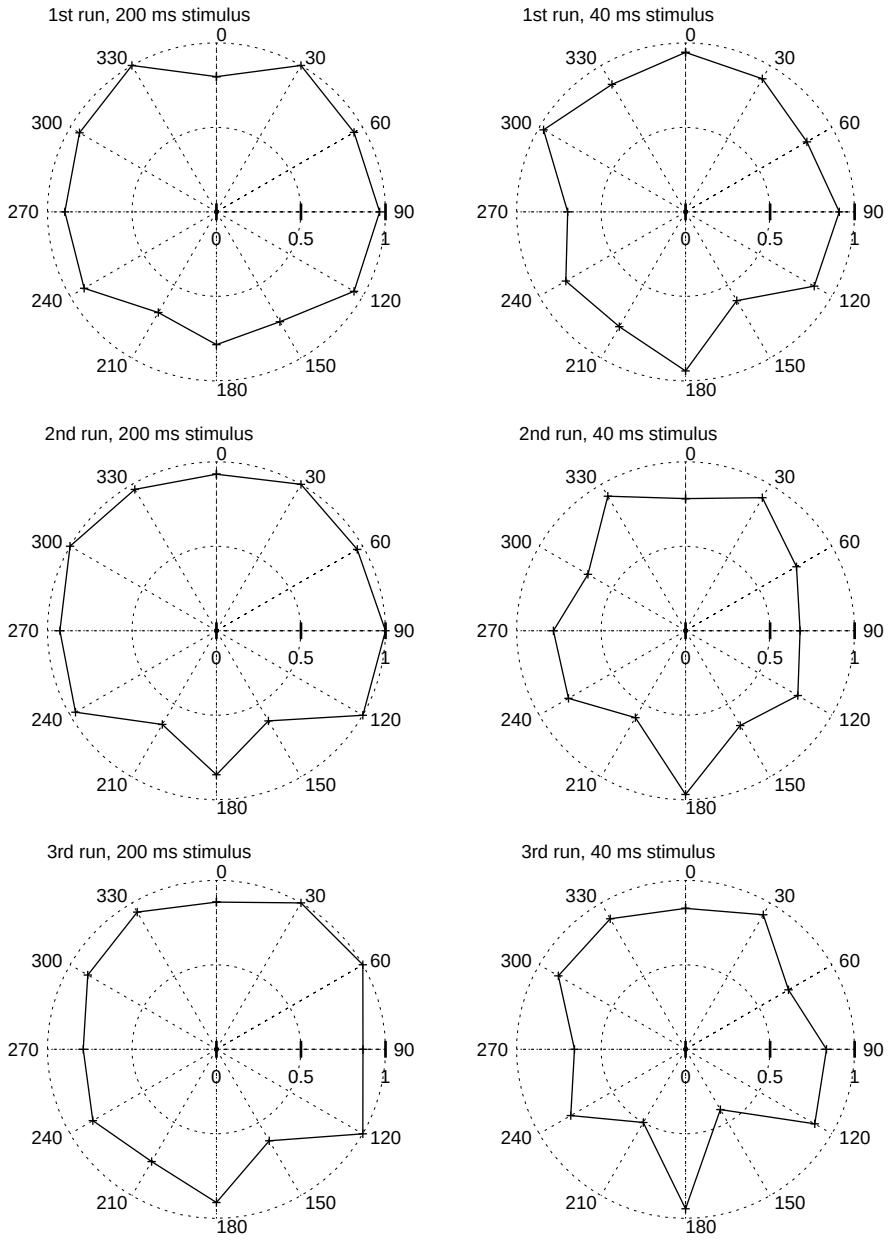

Figure 3.5: The performance of an animal from the NO-OP group (F9934) in response to stimuli presented from different positions. Performance is shown for 200 ms (left column) and 40 ms (right column) stimuli, and over all three testing runs (in chronological order, top to bottom). Data are plotted as in figure 3.3.

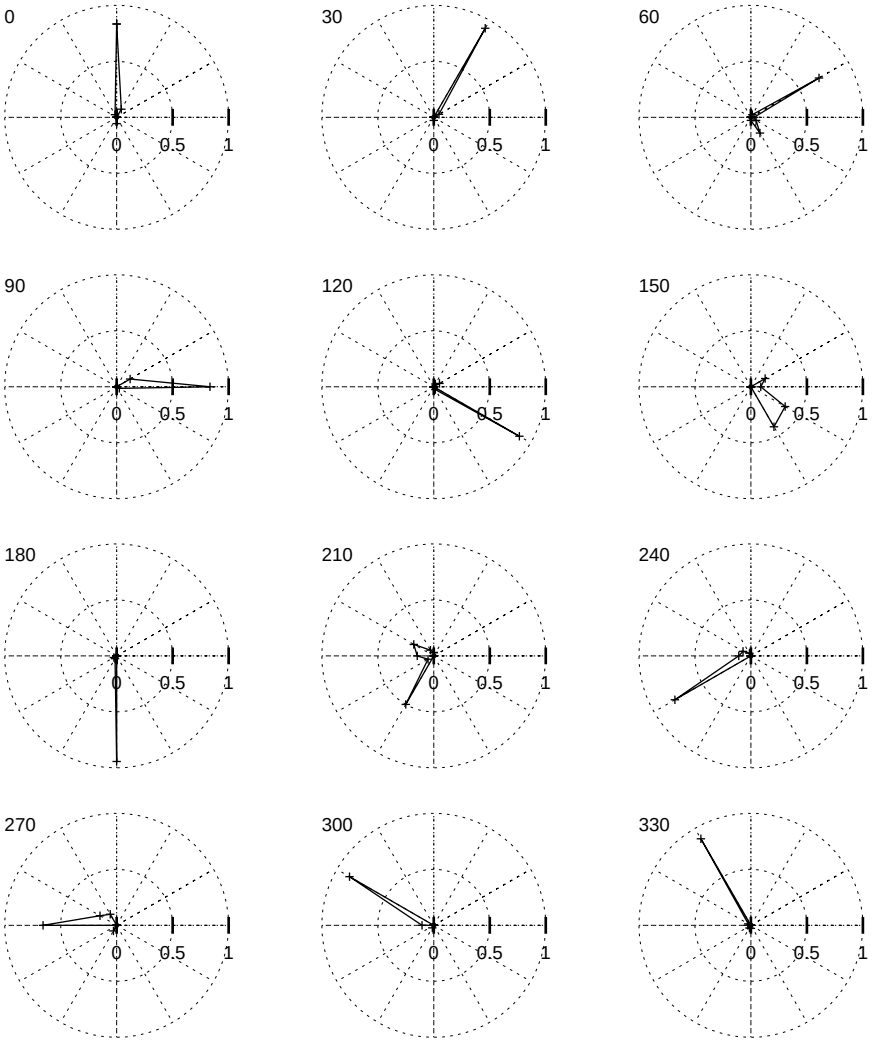

Figure 3.6: The distribution of responses to stimuli emanating from different positions, for an animal from the NO-OP group (F9934). Data shown are from responses to 40 ms stimuli on the third testing run (as in the bottom right panel of figure 3.5). Plotting conventions are as in figure 3.4.

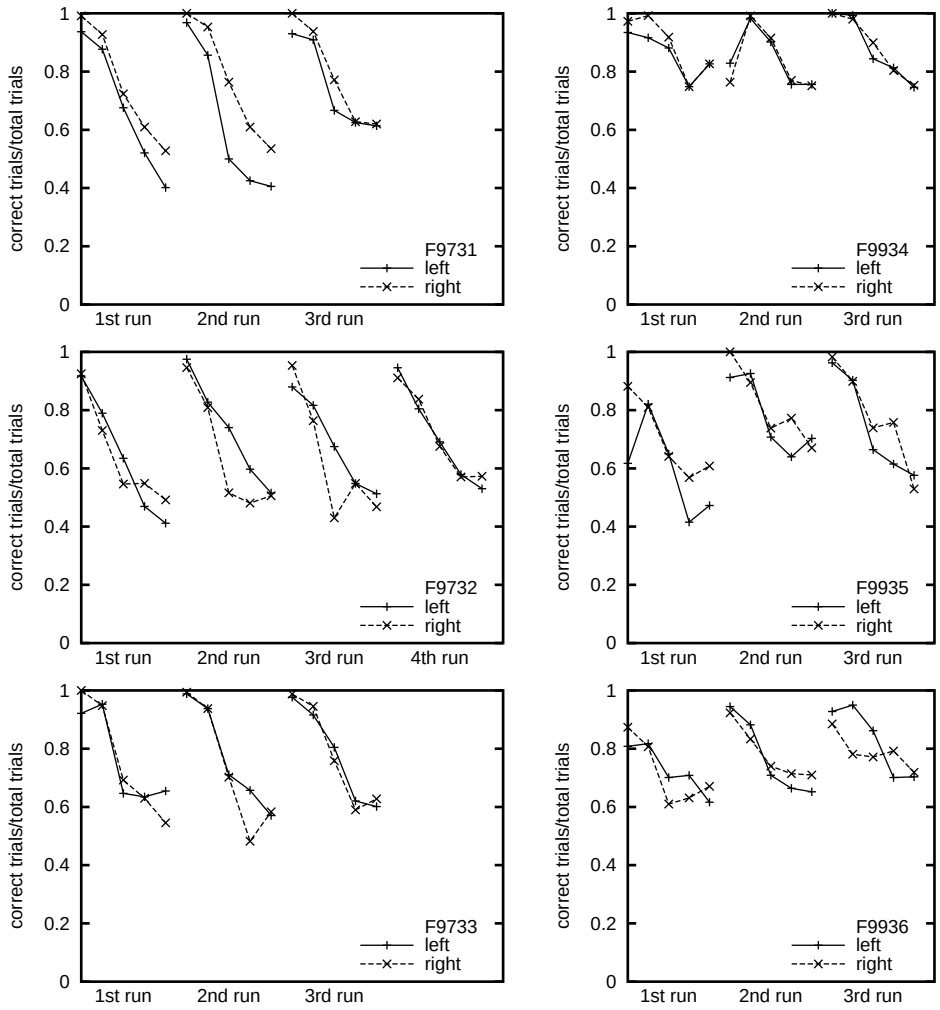

Figure 3.7: The performance of ferrets in groups PILOT (left column) and NO-OP (right column), with the results for stimuli presented in the left and right hemifields shown separately. Performance is measured as the proportion of trials to which the correct response was given. Data from trials where the stimulus was directly ahead of or directly behind the animal are not included in these plots. Other than the separation into two hemifields, trials are grouped by stimulus duration within each run, as in figure 3.2.

Group PILOT

source	df	SS	MS	F	p
duration	4	4.61	1.15	153	**<0.001**
error(duration)	8	0.0604	0.00755		
run	2	0.0150	0.00749	0.900	0.476
error(run)	4	0.0333	0.00832		
hemifield	1	0.0201	0.0201	0.344	0.617
error(hemifield)	2	0.116	0.0582		
duration*run	8	0.0329	0.00411	2.05	0.105
error(duration*run)	16	0.0320	0.00200		
duration*hemifield	4	0.0463	0.0116	2.14	0.168
error(duration*hemifield)	8	0.0433	0.00542		
run*hemifield	2	0.00417	0.00208	0.302	0.755
error(run*hemifield)	4	0.0276	0.00690		
duration*run*hemifield	8	0.0221	0.00276	0.841	0.581
error(duration*run*hemifield)	16	0.0526	0.00329		

Group NO-OP

source	df	SS	MS	F	p
duration	4	1.65	0.412	24.4	**<0.001**
error(duration)	8	0.135	0.0169		
run	2	0.160	0.0802	2.24	0.223
error(run)	4	0.143	0.0357		
hemifield	1	0.0176	0.0176	0.869	0.450
error(hemifield)	2	0.0406	0.0262		
duration*run	8	0.0911	0.0114	0.831	0.589
error(duration*run)	16	0.219	0.0137		
duration*hemifield	4	0.0287	0.00717	0.847	0.533
error(duration*hemifield)	8	0.0677	0.00847		
run*hemifield	2	0.0154	0.00770	5.04	0.081
error(run*hemifield)	4	0.00611	0.00153		
duration*run*hemifield	8	0.0368	0.00460	1.72	0.170
error(duration*run*hemifield)	16	0.0429	0.00268		

Table 3.4: ANOVA for animals in Groups PILOT and NO-OP, showing the differences in performance associated with stimulus duration, stimulus position (right hemifield *vs.* left hemifield) and testing run. Performance is measured as the arcsine-corrected proportion of trials to which the correct response was given.

were examined separately from each other. On the assumption that the underlying distribution is Binomial, the proportions for the right and left hemifields were then compared using a z-test. Of the animals in group PILOT, F9731 shows significantly better performance in the right hemifield than the left, at all durations examined (500 ms: $z = 3.01$ giving $p = 0.00127$; 200 ms: $z = 3.97$ giving $p = 3.53 \times 10^{-5}$; 100 ms: $z = 2.63$ giving $p = 0.00429$; 40 ms: $z = 2.45$ giving $p = 0.00696$). None of the other PILOT animals show any significant hemifield-related differences. Of the NO-OP animals, both F9934 and F9935 show some evidence of better performance in the right than left hemifields, but only at specific stimulus durations. In the case of F9934 significant differences are seen at 500 ms ($z = 1.91$ giving $p = 0.0283$) and 200 ms ($z = 1.75$ giving $p = 0.0403$), while for F9935 the difference is only significant at 100 ms ($z = 4.31$ giving $p = 8.22 \times 10^{-6}$).

3.3.3 Front-back errors

As mentioned in section 3.1, our experimental set-up allows animals to mislocalise stimuli from the front hemifield into the rear hemifield, and *vice versa*, where previous experiments have not. This is important, because some sound localisation cues do not distinguish a sound source in the front hemifield from one in the rear hemifield, and therefore an animal which relies on these cues might make a significant number of "front-back errors". We therefore wished to examine the number of front-back errors made by our animals. We defined front-back errors as being responses which were on the correct side of the animal's midline, or directly on the midline itself, but which were also on the wrong side of the interaural axis—this is a standard definition used in human psychophysical studies[73].

It is important to note that the likelihood of an error being classified as a front-back error is related to the magnitude of that error—for example, an error of only 30° can never be a front-back error according to our definition. We therefore need to make sure that any changes in the incidence of front-back errors are not simply an artefact of changes in error magnitude. For this reason, we compare front-back errors with other errors of the same magnitude. Figure 3.8 shows the full list of front-back errors, as well as the list of errors which are of equal magnitude to the front-back errors, and which can act as controls. Unfortunately, it is not possible to derive any equal-magnitude controls for front-back errors which are made when the target is on the midline—as a result we are forced to exclude such errors when calculating the incidence of front-back errors.

Figure 3.9 shows the incidence of both front-back errors and control errors, as defined in figure 3.8 and expressed as a proportion of the total num-

stim.	front-back errors
0°	120°, 150°, 180°, 210°, 240°

stim.	f-b	ctrl.
330°	180°	120°
	210°	90°
	240°	60°
300°	180°	60°
	210°	30°
	240°	0°
270°	none	
240°	300°	180°
	330°	150°
	0°	120°
210°	300°	120°
	330°	90°
	0°	60°

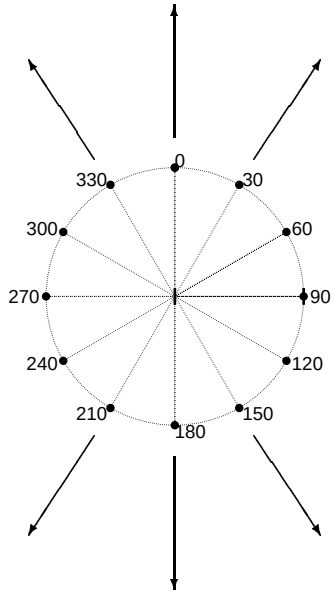

stim.	f-b	ctrl.
30°	120°	300°
	150°	270°
	180°	240°
60°	120°	0°
	150°	330°
	180°	300°
90°	none	
120°	0°	240°
	30°	210°
	60°	180°
150°	0°	300°
	30°	270°
	60°	240°

stim.	front-back errors
180°	0°, 30°, 60°, 300°, 330°

Figure 3.8: All possible front-back errors and control errors, broken down by stimulus position. Loudspeaker and water spout positions are represented by black dots in the central diagram, and labelled in degrees. The list of erroneous response positions refers to these labels. Abbreviations: "stim." = stimulus position, "f-b" = front-back errors, "ctrl." = control errors.

ber of responses. As with the other analyses, these data were subject to an ANOVA, grouping trials together by stimulus duration within each run. The measurement used for the ANOVA was the proportion of responses which were errors of a particular type: either front-back errors, or control errors of equal magnitude. These proportions were then arcsine-corrected, and stimulus duration, error type and testing run were all used as within-subjects factors for the analysis. Table 3.5 gives the results. For both groups of animals, the incidence of front-back errors seems in figure 3.9 to be greater than that of control errors, and the ANOVA shows that this is indeed significant at short stimulus durations, as indicated by an interaction between duration and error type. In fact, for animals in group NO-OP the difference between front-back errors and control errors is significant over all durations, as shown by a main effect of error type. It is, however, important to note that for both groups there is also a main effect of duration, which reflects the fact that the incidence of all types of errors goes up as the task gets more difficult.

As the overall number of errors changes at different durations, it is worth examining the incidence of front-back errors as a proportion of all errors, instead of as a proportion of trials, and this is shown in figure 3.10. This figure reveals two problems in using this metric as a measure of incidence. The first is that F9934, one of the NO-OP ferrets, has both fewer errors than the rest of its group (figure 3.2, bottom), and more front-back errors than the rest of its group (figure 3.9, top right), which means that the frequency of front-back errors as a proportion of the total number of errors is atypically high (figure 3.10, top right). Consequently, if we are to use the NO-OP animals as a control group for other studies (chapter 4), the proportion of trials is probably a more useful measure than the proportion of errors. The second problem with using the proportion of errors is that incidence at the longest stimulus durations can be quite variable, and often confusingly large, because the total number of errors at these durations is so small.

Although we have already shown that there are no significant differences in overall frequency of errors between the two hemifields, we also investigated whether there was any difference in frequency of front-back errors between the hemifields. The resulting ANOVAs are given in tables 3.6 and 3.7, but the corresponding graphs are not shown because, as expected, no significant effect of hemifield was found in either group.

3.4 Discussion

In these experiments we constructed a test of absolute localisation in animals which is measured by an approach-to-target method, and which covers the

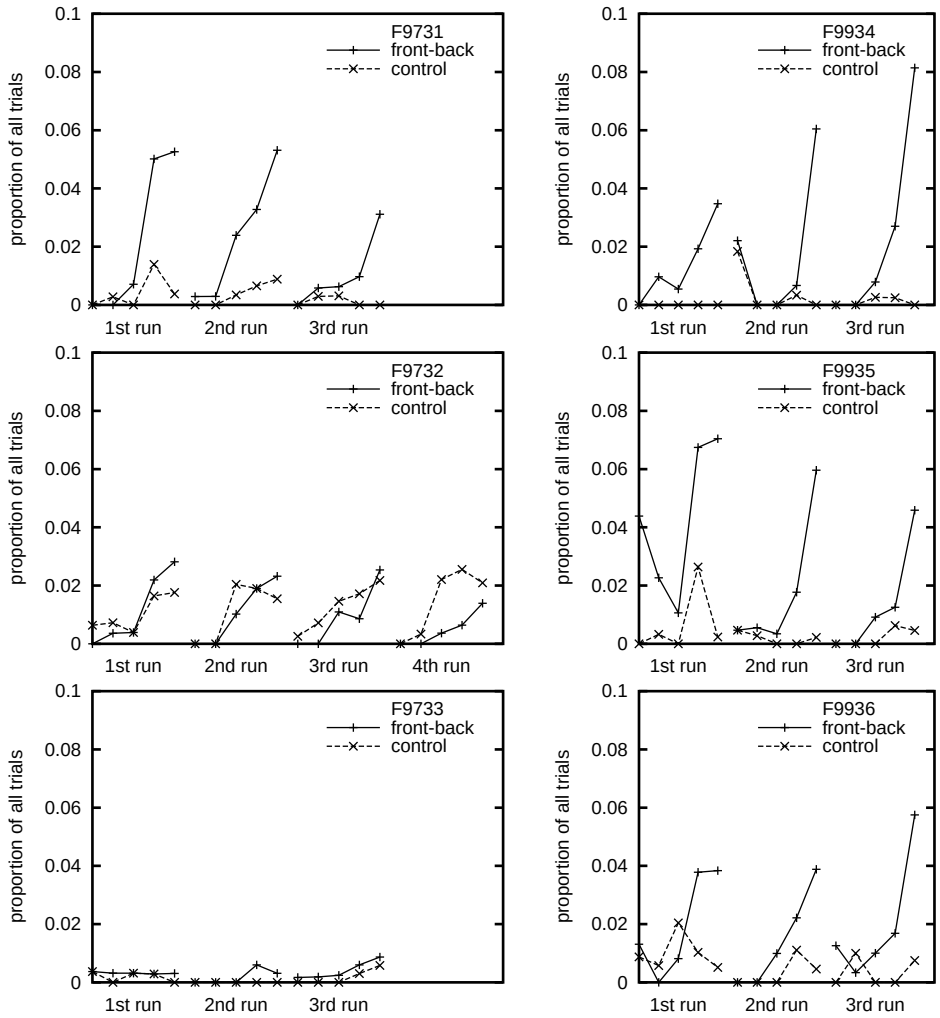

Figure 3.9: The incidence of front-back errors, as a proportion of the total number of trials, in the responses of ferrets in group PILOT (left column) and NO-OP (right column). Errors which are of equal magnitude to front-back errors are shown as a control. Errors are grouped by stimulus duration within each run, in the same way as trials are grouped in figure 3.2.

Group PILOT

source	df	SS	MS	F	p
duration	4	0.115	0.0288	7.96	**0.007**
error(duration)	8	0.0289	0.00362		
run	2	0.00200	0.00100	0.323	0.741
error(run)	4	0.0124	0.00309		
error-type	1	0.0174	0.0174	1.206	0.387
error(error-type)	2	0.0289	0.0144		
duration*run	8	0.0124	0.00156	0.908	0.534
error(duration*run)	16	0.0274	0.00171		
duration*error-type	4	0.0196	0.00491	3.99	**0.046**
error(duration*error-type)	8	0.00986	0.00123		
run*error-type	2	0.00112	0.000560	1.33	0.361
error(run*error-type)	4	0.00169	0.000422		
duration*run*error-type	8	0.00211	0.000264	0.405	0.902
error(duration*run*error-type)	16	0.0104	0.000651		

Group NO-OP

source	df	SS	MS	F	p
duration	4	0.135	0.0338	46.2	**<0.001**
error(duration)	8	0.00584	0.000730		
run	2	0.0144	0.00719	1.19	0.392
error(run)	4	0.0241	0.00603		
error-type	1	0.135	0.135	47.1	**0.021**
error(error-type)	2	0.00574	0.00287		
duration*run	8	0.0219	0.00273	0.747	0.651
error(duration*run)	16	0.0586	0.00366		
duration*error-type	4	0.0854	0.0214	12.7	**0.002**
error(duration*error-type)	8	0.0134	0.00168		
run*error-type	2	0.00223	0.00112	0.265	0.780
error(run*error-type)	4	0.0169	0.00421		
duration*run*error-type	8	0.00674	0.000842	0.614	0.754
error(duration*run*error-type)	16	0.0220	0.00137		

Table 3.5: ANOVA for animals in Groups PILOT and NO-OP, showing the differences in error likelihood associated with stimulus duration, error type (front-back *vs.* control) and testing run. Error likelihood is measured as the arcsine-corrected proportion of responses which are of the given type.

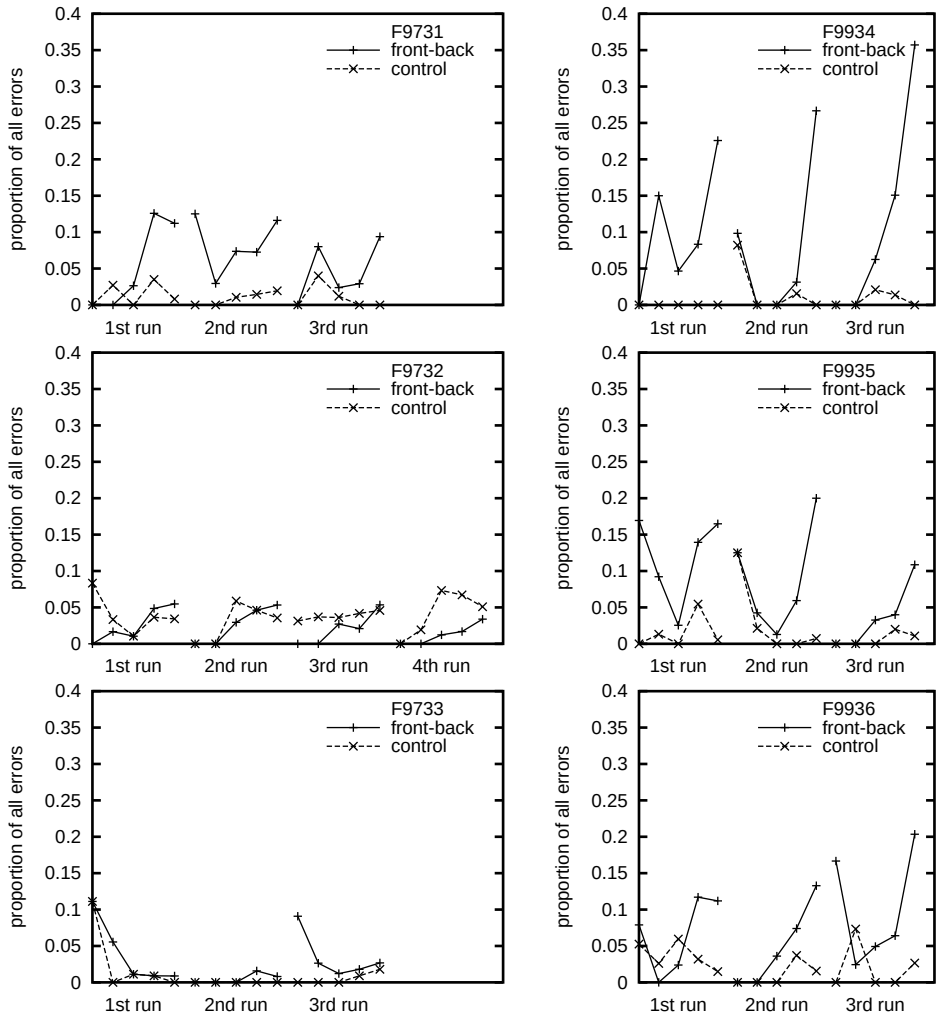

Figure 3.10: The incidence of front-back errors, as a proportion of the total number of errors, in the responses of ferrets in group PILOT (left column) and NO-OP (right column). Errors which are of equal magnitude to front-back errors are shown as a control. Errors are grouped by stimulus duration within each run, in the same way as trials are grouped in figure 3.2.

Group PILOT

source	df	SS	MS	F	p
hemifield	1	0.00874	0.00874	0.653	0.504
error(hemifield)	2	0.0268	0.0134		
error-type	1	0.0393	0.0393	1.15	0.396
error(error-type)	2	0.0685	0.0343		
run	2	0.00303	0.00151	0.319	0.744
error(run)	4	0.0190	0.00474		
duration	4	0.270	0.0676	7.44	**0.008**
error(duration)	8	0.0726	0.00908		
hemifield*error-type	1	0.00749	0.00749	0.384	0.599
error(hemifield*error-type)	2	0.0390	0.0195		
hemifield*run	2	0.000441	0.000220	6.04	0.959
error(hemifield*run)	4	0.0208	0.00521		
error-type*run	2	0.00250	0.00125	1.27	0.375
error(error-type*run)	4	0.00395	0.000988		
hemifield*error-type*run	2	0.00522	0.00261	3.98	0.112
error(hemifield*error-type*run)	4	0.00262	0.000655		
hemifield*duration	4	0.0154	0.00385	0.651	0.642
error(hemifield*duration)	8	0.04735	0.00592		
error-type*duration	4	0.0493	0.0123	3.37	0.068
error(error-type*duration)	8	0.0292	0.00366		
hemifield*error-type*duration	4	0.00402	0.00100	0.602	0.672
error(hemifield*error-type*duration)	8	0.0133	0.00167		
run*duration	8	0.0263	0.00329	0.971	0.491
error(run*duration)	16	0.0542	0.00339		
hemifield*run*duration	8	0.0131	0.00164	0.804	0.609
error(hemifield*run*duration)	16	0.0326	0.00204		
error-type*run*duration	8	0.00457	0.000571	0.659	0.719
error(error-type*run*duration)	16	0.0139	0.000867		
hemifield*error-type*run*duration	8	0.0156	0.00195	1.44	0.255
error(hemifield*err-type*run*duration)	16	0.0217	0.00136		

Table 3.6: ANOVA for animals in Group PILOT, showing the differences in error likelihood associated with stimulus duration, stimulus position (right *vs.* left hemifield), error type (front-back *vs.* control) and testing run. Error likelihood is measured as the arcsine-corrected proportion of responses which are of the given type.

Group NO-OP

source	df	SS	MS	F	p
hemifield	1	0.000365	0.000365	0.021	0.899
error(hemifield)	2	0.0354	0.0177		
error-type	1	0.303	0.303	32.1	**0.030**
error(error-type)	2	0.0189	0.00944		
run	2	0.0309	0.0155	1.20	0.391
error(run)	4	0.0516	0.0129		
duration	4	0.356	0.0891	54.2	**<0.001**
error(duration)	8	0.0131	0.00164		
hemifield*error-type	1	0.0137	0.0137	6.53	0.125
error(hemifield*error-type)	2	0.00420	0.00210		
hemifield*run	2	0.00487	0.00244	1.48	0.329
error(hemifield*run)	4	0.00656	0.00164		
error-type*run	2	0.00439	0.00219	0.292	0.761
error(error-type*run)	4	0.0300	0.00751		
hemifield*error-type*run	2	0.00689	0.00344	0.593	0.595
error(hemifield*error-type*run)	4	0.0232	0.00580		
hemifield*duration	4	0.0101	0.00252	0.513	0.729
error(hemifield*duration)	8	0.0393	0.00492		
error-type*duration	4	0.235	0.0587	18.6	**<0.001**
error(error-type*duration)	8	0.0253	0.00316		
hemifield*error-type*duration	4	0.00368	0.000920	0.180	0.942
error(hemifield*error-type*duration)	8	0.0408	0.00510		
run*duration	8	0.0391	0.00489	0.758	0.643
error(run*duration)	16	0.103	0.00645		
hemifield*run*duration	8	0.0183	0.00229	0.878	0.555
error(hemifield*run*duration)	16	0.0418	0.00261		
error-type*run*duration	8	0.0197	0.00247	0.977	0.487
error(error-type*run*duration)	16	0.0404	0.00253		
hemifield*error-type*run*dur	8	0.0242	0.00302	1.14	0.388
error(hemifield*err-type*run*duration)	16	0.0423	0.00264		

Table 3.7: ANOVA for animals in Group NO-OP, showing the differences in error likelihood associated with stimulus duration, stimulus position (right *vs.* left hemifield), error type (front-back *vs.* control) and testing run. Error likelihood is measured as the arcsine-corrected proportion of responses which are of the given type.

full 360° range of azimuths. We found that normal ferrets are capable of performing well on this task, even with no prior training. The main factor affecting their performance was the duration of the stimulus, while stimulus position also had some effect.

What is encouraging about this method is how little training it appears to require, by comparison to a similar MAA task[92]. Although the animals in the PILOT group were extensively trained before systematic testing began, the same was not true of the animals in the NO-OP group, who had only minimal acclimatisation to the experimental chamber prior to testing. Despite this, neither the PILOT nor the NO-OP animals showed significant improvement over the course of testing, suggesting that they were already near the peak of their performance. This is despite the fact that the testing regime provides plenty of feedback that could be used to guide an animal's learning: rewards were contingent upon correct responses throughout testing, and correction trials were always presented following incorrect responses. The speed with which the NO-OP animals pick up the task suggests that this is an intuitively obvious task for them. Not only that, but all the animals appeared to find the task quite easy—even during sessions when performance was at its worst, about 50% of responses were correct, whereas responding at random would only achieve an $8.\dot{3}$% correct rate.

Given the high level of performance on this task, it is important to stress that we excluded possible non-spatial cues from our setup, in order to be certain that our animals were responding only on the basis of location cues. To prevent animals learning the characteristics of the different loudspeakers in the task, we used individual filters for each one so that the sound spectrum produced from each was the same. We also varied stimulus intensity pseudo-randomly, so that the ferrets could not reliably use the gross sound-attenuating effect of their bodies in differentiating posterior and anterior sound sources. We also masked the loudspeakers from view, and in one case ran an animal in the dark, to reduce the effect of any visual cues on performance.

There is a drawback to our experimental set-up, however, and that is its relatively coarse sampling of the range of azimuths—we use speakers and targets spaced at 30°, while the acuity of a ferret on an MAA task can be as low as 10° for stimuli directly ahead of the animal, depending on stimulus duration[92]. This means that our measures of performance are potentially insensitive to small differences in localisation ability, something which may be important in assessing the effect of A1 inactivation on this task, as we do in chapter 4.

The duration of the stimulus that the animals had to locate was the most important factor in determining performance on this task. Performance for

1 s stimuli was frequently close to 100% correct, while the same figure for 40 ms stimuli was around 50%. Presumably the longer the stimulus, the longer the brain has to analyse the location cues which it contains; in effect, the brain has the opportunity to "sample" the incoming sound several times. However, there is another mechanism that may give rise to an effect of stimulus duration. It has frequently been postulated that during long-duration sounds an animal has the opportunity to make "scanning" head movements; among other things, this could help disentangle the sound filtering effects of the outer ear from the spectrum of the stimulus itself. There is, though, little evidence that head movements take place in this way: observations on other animals in a similar task[92] reveal an initial orienting response, about 100–250 ms after stimulus onset, after which the head was kept steady until either stimulus offset or until the animal left the central platform to make a response. Given that we find localisation ability to be maximum for stimuli which are located along the front midline, the orienting response could have more to do with keeping the target straight ahead than with creating relative motion cues.

As already mentioned, the location of the stimulus did appear to have some effect on localisation ability. We find, as have previous studies, that animals' performance was best at the frontal midline, declining at more lateral positions[78, 41]. The difference between midline and lateral positions has been attributed to several factors, in particular the directional specificity of various filtering effects of the outer ear[78, 48]. Other authors have stressed, not the directionality of the cues themselves, but the rate of change of these cues with changing stimulus position, both in the case of outer ear-related spectral cues[18] and interaural difference cues.

Our population of subjects showed no overall group difference in localisation ability between the right and left hemifields, but animals' performance did show a certain amount of individual variation in this. Of our six animals, one showed a consistent localisation advantage for stimuli in the right hemifield. A further two showed a less consistent, and smaller, advantage, again in the right hemifield. However, the differences between animals show no signs of being qualitative rather than quantitative, and there do not appear to be distinct groups of "right-handed" and "left-handed" ferrets with respect to sound localisation.

Unlike many previous animal studies, we are in a position to examine performance in the posterior as well as the anterior hemifield. We found that performance tends to be lower in the posterior hemifield than the anterior hemifield, but otherwise follows a similar pattern of being higher near the midline than at more lateral positions. This is similar to the results of studies in humans[73, 86].

Given that one of the advantages of our circular chamber is the ability to observe so-called "front-back errors", it is worth examining how often they occur in these animals. Although the results differ slightly between our groups, it does seem that, particularly when the animals' performance is low, front-back errors are more common than other errors of equal magnitude. This suggests that such errors are worth considering as a class in themselves, as has been done in many previous studies[15, 73, 86]. However, there are notable differences between the frequency of front-back errors in different animals, and if some animals are indeed more susceptible to front-back confusions than others, then this suggests that different animals may have different strategies for sound localisation.

Chapter 4

Effects of A1 inactivation on auditory absolute localisation in ferrets

4.1 Introduction

Primary auditory cortex (A1) is the main gateway through which auditory information reaches the cerebral cortex—an anatomical position which suggests it has both an important and general-purpose role in auditory processing. It is thus surprising that animals' performance on tasks which might be regarded as being at the core of auditory perception—frequency discrimination, for example—is unaffected by A1 lesions[17]. In fact, the most convincingly-demonstrated effect of A1 damage in animal experiments is an impairment in sound localisation[55, 59]. As a result, much research has concentrated on investigating exactly what it is about sound localisation and/or sound localisation tasks that is so special and so important.

A1-lesioned animals have been shown to be impaired on at least two sound localisation tasks: the approach-to-target minimum audible angle (MAA) task and the approach-to-target absolute localisation task. In both of these the animal—usually a cat or a ferret—must approach a target in order to be rewarded for making a correct choice. In the MAA task the animal chooses between two targets, indicating which of two loudspeakers was the source of a noise, while in the absolute localisation task the animal must actually approach the sound source directly. In both cases the targets must be attached to the floor, and hence the potential loudspeaker positions are usually restricted to the horizontal plane—these are thus essentially *azimuth* locali-

sation tasks.[1]

On an MAA task, unilateral ablation of A1 led to a reduction in ferrets' ability to distinguish two nearby sound sources, but only when both sources were positioned contralateral to the lesion, and when the sounds used were brief (<100 ms)[59]. Acuity for sources near the front midline, or in the front quadrant ipsilateral to the lesion, remained intact. Bilateral A1 ablation appears to have an effect on MAA scores in both lateral front quadrants, but, as with unilateral A1 ablation, midline acuity remains intact.

Similarly, on a seven-choice absolute localisation task, cats given unilateral A1 lesions are impaired at locating the source of clicks when that source is in the front quadrant contralateral to the lesion[55]. The same animals are not impaired with clicks presented from the midline or from loudspeakers in the ipsilateral front quadrant. Crucially, unilateral A1 lesions which are restricted to a particular isofrequency band lead to an impairment on this task which is not only limited to the contralateral hemifield, but is also limited to short-duration tone bursts of certain frequencies[56]. The frequency range over which the animals are impaired matches the frequency sensitivity of the destroyed cells.

The most obvious interpretation of these results is that A1 damage destroys the cortical representation of the side of auditory space contralateral to the lesion. One problem with this view is that not all sound localisation tasks are affected by A1 lesions[76]. However, this may simply be because the tasks which are not affected can be processed adequately at the subcortical level. A more serious problem is that notions of a representation of auditory space do not accord well with the electrophysiological data. There is no "map" of auditory space laid out across A1, for example, although it is true that the majority of A1 neurons respond more vigorously to stimuli in the contralateral than the ipsilateral hemifield.

If A1 does not contain a unified representation of auditory space, then it may still contain representations of various auditory cues which are important for localisation. If this is the case, then the obvious question is: what are these cues, and how are they important to locating sound sources?

As discussed in section 1.3.2, there are two main auditory features which provide cues to sound localisation: interaural differences and pinna/concha-related spectral filtering. Although cues of both types are produced by free-field stimuli, selective disruption of one or the other might be expected to produce different patterns of errors. In particular, interaural difference cues are approximately symmetric about the interaural axis. Thus a selective

[1] In an MAA task the targets and loudspeakers need not be next to one another, so it is possible to have loudspeakers that are not in the horizontal plane[92].

disruption of spectral cues might lead to an increase in mistakes which are also symmetric about the interaural axis. Previous experiments will not have revealed this because they have restricted themselves to locations along the horizon in the front hemifield.

We intend to test the effects of A1 inactivation on an approach-to-target absolute localisation task, using sound sources in both the front and rear hemifields. This should allow us to measure the frequency of "front-back errors", and hence draw conclusions about the sound localisation cues that A1 is processing. Our methodology also uses several other novel features. Instead of physically damaging A1, we use implants that slowly release an inactivation agent (muscimol) over the course of weeks. This is a valuable technique for many reasons: it is reversible, causing little or no permanent damage to the area under study, it should avoid any damage to afferent neurons, fibres of passage or underlying white matter tracts, and it may minimise knock-on physiological effects associated with physical lesions[141]. Also, as well as inactivating A1 in trained animals, as has been done in previous studies, we inactivate A1 in some animals *before* training them on the task. This should begin to address whether A1 is more involved in learning a new task or in performing an old task.

4.2 Methods

4.2.1 Numerical conventions

Throughout this chapter, azimuth is expressed in degrees, from 0° to 360°: 0° is straight ahead, 90° is to the right, 180° is directly behind, and 270° is to the left. Experimental measurements are expressed either to the precision of measurement, or to 3 significant figures, whichever is lower. Other numeric quantities are treated similarly, although values for statistical probability (p) are sometimes given to 3 decimal places only, for reasons of space and clarity.

In analysis of variance (ANOVA) tables, SS is an abbreviation for Sum of Squares, MS is an abbreviation for Mean Sum of Squares, and significant results are shown in boldface. The values used in ANOVAs are often proportional data, e.g. the proportion of trials to which an animal responds correctly. Such data typically show variance which is maximum near 0.5 and minimum near 0 or 1. As in chapter 3, this heterogeneity of variance is corrected using an "arcsine correction". This is where the inverse sine of the square root of the proportion is used instead of the proportion itself[139].

4.2.2 Animals

The ferrets were divided into several experimental groups. Group PILOT consisted of animals F9731, F9732 and F9733—these animals had been previously trained and tested on the localisation task as described in chapter 3. The NO-OP group animals from chapter 3 were not used in the experiments described here, although they do form a control group for this study. Group UNI consisted of animals F9907, F9908 and F9909 (all female) and Group BI consisted of animals F9922, F9923 and F9924 (all male). The animals in both these groups were totally untrained before the work described here. Group UNI animals were used to test the effect of implanting muscimol Elvax over left A1, while Group BI animals were given muscimol Elvax implants over both left and right A1. Two other animals, F9448 (female) and F9535 (male), made up Group BLANK: these would be implanted with blank Elvax as a control for the effects of surgery. At the start of this study the animals in Group BLANK had been trained and tested on several auditory tasks, including auditory masking task and an MAA task, but they had no experience in the chamber used for absolute localisation.

4.2.3 Apparatus

All training and testing was carried out in the localisation chamber described in section 2.2 and shown in figure 2.1. This was a circular chamber containing a water spout in the centre, mounted just in front of a raised platform. Around the edge of the chamber were mounted twelve peripheral water spouts at 30° intervals. Above each peripheral water spout was a loudspeaker. A computer was used to trigger presentation of stimuli from the loudspeakers, and to monitor when an animal licked any of the water spouts.

4.2.4 Behavioural training and testing

The task used to test the animals was the localisation task described in detail in section 2.3 and also used in chapter 3. Briefly, an animal would lick the central water spout, which would cause a burst of flat white noise to be presented from one of the twelve peripheral speakers. In most cases the frequency range of the noise was 0.5–30 kHz, but more restricted frequency ranges were used in some cases, as indicated in the testing schedules below. The animal's task was to lick the peripheral water spout mounted underneath whichever speaker had delivered the noise burst. A correct response was rewarded with water. An incorrect response was not rewarded, and on subsequent trials ("correction trials") the stimulus was presented from the

same speaker until a correct response was given. Responses to these correction trials were excluded when calculating animals' performance. Also excluded were trials on which animals were rewarded at the centre spout.

As explained in section 2.1, the training and testing of animals was carried out in a series of 'runs', each lasting about 2 weeks and during which the animal was water-deprived. For a particular animal, at least two days (and typically more) separated each run from the next. During a run, each animal was put through two training or testing sessions a day, and each session consisted of roughly 100 trials (depending on the schedule).

4.2.4a Group PILOT (practised animals, left A1 inactivation)

Training These animals were already experienced on the auditory localisation task, having previously been trained and tested as described in chapter 3, so no further training was needed.

A1 surgery (implantation) All the animals in this group had pieces of muscimol Elvax implanted over left A1. The procedures for manufacture and implantation of muscimol Elvax are described in section 2.4. The animals were not tested until at least 6 days following surgery, to allow the surgical wound to heal.

Testing The animals were put through several testing runs in the absolute localisation chamber. Most of these runs followed one of the three schedules listed in table 4.1, although it was common for extra sessions to be added *ad hoc*. The 'full' and 'short' schedules differ only in the use of stimulus durations which would not normally be analysed (2 s, 750 ms, 400 ms and 300 ms), but the 'partial' schedule concentrates on 40 ms stimuli and omits several key stimulus durations. In addition to the trials detailed in table 4.1, sessions were also run using 40 ms stimuli with restricted bandwidth—these took place at the end of a run, and the results were analysed separately from the rest of the data.

The actual schedules used for each animal on each run are shown in table 4.2. Of the two testing runs marked as 'other' in table 4.2, F9732's second run was a schedule intermediate between the 'full' and 'short' schedules described in table 4.1, while F9732's third run was similar to the 'partial' schedule, but contained no sessions at 40 ms other than those using restricted-bandwidth stimuli.

As with the testing of these animals in chapter 3, stimulus intensity was varied from trial to trial. On most sessions, the intensity on each trial was randomly selected from the following values: 56 dB SPL, 63 dB SPL,

noise duration	number of sessions or trials		
	'full' run	'short' run	'partial' run
2 s	1 session	1 session	0 sessions
1 s	3 sessions	3 sessions	1 session
750 ms	1 session	0 sessions	0 sessions
500 ms	300 trials	300 trials	1 session
400 ms	1 session	0 sessions	0 sessions
300 ms	1 session	0 sessions	0 sessions
200 ms	300 trials	300 trials	2 sessions
100 ms	300 trials	300 trials	0 sessions
40 ms	300 trials	300 trials	300 trials

Table 4.1: Typical testing schedules for ferrets in Group PILOT while implanted with muscimol Elvax. Sessions were run in the order listed—top to bottom. Animals were run at each stimulus duration either for the specified number of sessions, or until they had run approximately the specified number of trials.

ferret	1st run	2nd run	3rd run	4th run	5th run	6th run
F9731	full	full+	short	partial	full	
F9732	partial	other	other+	short	full	full
F9733	full	full+	short	partial	full	

Table 4.2: Testing schedule for ferrets in Group PILOT while implanted with muscimol Elvax. The table indicates which schedule was used for which animal during which run. The names of schedules are those given in table 4.1. Runs marked as "other" are described in the text. The '+' sign indicates runs which include sessions with restricted-bandwidth stimuli.

70 dB SPL, 77 dB SPL, 84 dB SPL. However, for F9731 and F9733, the first two runs used values over a smaller range: 66 dB SPL, 68 dB SPL, 70 dB SPL, 72 dB SPL, 74 dB SPL. A spout delay was used throughout, being determined at random before the start of each trial from within the range 0.5–2 s. The headgrid was not used during testing.

As mentioned above, restricted-bandwidth noise was used in some testing sessions. Four different stimuli, all 40 ms in duration, were randomly interleaved in these sessions: low-pass noise with a cutoff of 10 kHz, high-pass noise with a cutoff of 20 kHz, band-pass noise with a pass-band of 10–20 kHz, and broadband noise as normal. By contrast to normal testing, stimulus intensity was not varied within these sessions.

A1 surgery (removal) The pieces of Elvax were then removed from all of these animals, before a final round of tests. Procedures for implant removal can be found in section 2.4.2c. As with implantation, the animals were not tested until at least 6 days following surgery.

Testing Following removal of muscimol Elvax, the animals were re-tested using the schedules given in table 4.3. F9731 and F9733 were only given one post-removal testing run, whereas F9732 was given two. Stimulus intensity was randomly chosen on each trial from among the following values: 56 dB SPL, 63 dB SPL, 70 dB SPL, 77 dB SPL, 84 dB SPL. Other aspects of the post-removal testing were as for the testing during implantation.

4.2.4b Group UNI (unpractised animals, left A1 inactivation)

Training Animals in this group were familiarised with the chamber and its system of water spouts by means of several (6–8) sessions in which sound stimuli were not used. In these sessions the animals were given rewards for visiting each of the chamber's water spouts in turn. Extra rewards were given for standing properly on the platform and licking the centre spout, as the animals needed most encouragement to do this. For animals whose stance on the platform was particularly bad, the headgrid was used during this training.

A1 surgery All the animals in this group had pieces of muscimol Elvax implanted over left A1. The procedures for manufacture and implantation of muscimol Elvax are described in section 2.4. The animals were not tested until at least 6 days following surgery, to allow the surgical wound to heal.

noise duration	number of sessions or trials		
	F9731/F9733	F9732	
2 s	1 sessions	2 session	1 session
1 s	3 sessions	3 sessions	1 session
750 ms	1 session	1 session	0 sessions
500 ms	300 trials	300 trials	2 sessions
400 ms	1 session	0 sessions	0 sessions
300 ms	1 session	0 sessions	0 sessions
200 ms	200 trials	300 trials	0 sessions
100 ms	200 trials	300 trials	0 sessions
40 ms	500 trials	500 trials	900 trials

Table 4.3: Testing schedules for ferrets in Group PILOT after removal of muscimol Elvax. Sessions were run in the order listed—top to bottom. Animals were run at each stimulus duration either for the specified number of sessions, or until they had run approximately the specified number of trials. Of the 2 schedules listed for F9732, the left column was used for the first post-removal run, and the right column for the second post-removal run.

Testing The animals were then put through three full runs of the absolute localisation task. The number of sessions or trials that were run at each stimulus duration is given in table 4.4; this schedule is intentionally similar to that used for the Group NO-OP and Group BI animals.[2]

The rest of the testing procedure was also the same as for animals in Groups NO-OP and BI. No spout delay was used in these sessions. Stimulus intensity was varied randomly from trial to trial, being chosen from among the following values on each trial: 56 dB SPL, 63 dB SPL, 70 dB SPL, 77 dB SPL, 84 dB SPL. The headgrid was not used during testing.

4.2.4c Group BI (unpractised animals, bilateral A1 inactivation)

Training Animals in this group were familiarised with the apparatus in the same way as those in Group UNI.

A1 surgery The animals in this group had pieces of muscimol Elvax implanted over both left and right A1. Surgery on both hemispheres was carried out in the same operation. The procedures for manufacture and implantation

[2]There are small differences between the schedule used for Group UNI and that used for Groups NO-OP and BI. This is because Group UNI was the first group of animals to be tested without having previously been trained with a sound stimulus—as a result, the schedule was more cautious than for Groups NO-OP and BI which were trained later.

noise duration	number of sessions or trials	
	first run	all other runs
looped	2 sessions	0 sessions
2 s	2 sessions	1 session
1 s	2 sessions	2 sessions
750 ms	1 session	1 session
500 ms	300 trials	300 trials
400 ms	1 session	1 session
300 ms	1 session	1 session
200 ms	300 trials	300 trials
100 ms	300 trials	300 trials
40 ms	500 trials	400 trials

Table 4.4: Testing schedule for ferrets in Group UNI. Sessions were run in the order listed—top to bottom. Animals were run at each stimulus duration either for the specified number of sessions, or until they had run approximately the specified number of trials.

of muscimol Elvax are described in section 2.4. The animals were not tested until at least 6 days following surgery, to allow the surgical wound to heal.

Testing The animals were then put through four full runs of the absolute localisation task. The number of sessions or trials that were run at each stimulus duration is given in table 4.5; this schedule is the same as that used for the training of Group NO-OP animals.

The rest of the testing procedure was also the same as for animals in Groups NO-OP and UNI. No spout delay was used in these sessions. Stimulus intensity was varied randomly from trial to trial, being chosen from among the following values on each trial: 56 dB SPL, 63 dB SPL, 70 dB SPL, 77 dB SPL, 84 dB SPL. The headgrid was not used during testing.

4.2.4d Group BLANK (blank Elvax)

A1 surgery Both the animals in this group had pieces of blank Elvax (i.e. Elvax containing no drug) implanted over right A1 prior to any training or testing on this task. The procedures for manufacture and implantation of blank Elvax are described in section 2.4. The animals were allowed to rest for at least a week following surgery, to allow the surgical wound to heal.

Training and testing The animals in Group BLANK were highly familiar with the use of water spouts due to their training as part of earlier studies, so

noise duration	number of sessions or trials	
	first run	all other runs
looped	2 sessions	0 sessions
2 s	2 sessions	1 session
1 s	2 sessions	2 sessions
750 ms	1 session	1 session
500 ms	300 trials	300 trials
400 ms	1 session	0 sessions
300 ms	1 session	0 sessions
200 ms	300 trials	300 trials
100 ms	300 trials	300 trials
40 ms	400 trials	400 trials

Table 4.5: Testing schedule for ferrets in Group BI. Sessions were run in the order listed—top to bottom. Animals were run at each stimulus duration either for the specified number of sessions, or until they had run approximately the specified number of trials.

they only needed one session of "spout training" in the absolute localisation chamber. They were also trained and tested during the same run: initial training involved several (4–10) sessions with looped noise, and then one or two sessions with either a 2 s or 1 s noise burst. Testing then commenced with a similar schedule to that used for the PILOT animals in chapter 3. The exact schedule is not given here as it was determined on an *ad hoc* basis and differed between the two animals in the group. However, 300 trials' worth of data was collected at each of the key stimulus durations of 500 ms, 200 ms, 100 ms and 40 ms. After a break, these animals were put through another testing run, giving two runs worth of data in total.

During testing, stimulus intensity varied from trial to trial, being randomly chosen on each trial to be one of the following values: 66 dB SPL, 68 dB SPL, 70 dB SPL, 72 dB SPL, 74 dB SPL. The headgrid was not used for either training or testing, and no spout delay was used.

4.2.5 Measuring properties of muscimol Elvax

4.2.5a Dynamics of muscimol release

The release of muscimol from Elvax pieces was measured using the methods described in section 2.4.3. This consisted of incorporating radiolabelled muscimol into Elvax sheets, and tracking the release of the radiolabel into a bathing solution of phosphate-buffered saline by using scintillation-counting.

In some cases this was done for pieces of Elvax which had been implanted into ferrets for several days or weeks, in order to gauge how much of the muscimol had already been released.

4.2.5b Effects on cortical activity

In addition to tracking the release of muscimol from Elvax, measurements were also made to confirm the effectiveness of the implants as cortical blocking agents. This is described in greater detail in section 2.4.3. In summary, though, the experiments involved implanting muscimol Elvax into ferrets, then, after a certain period of time, using electrophysiological techniques to record the cortical activity in the region of the implant. In fact, the implants used for this were specially cut so that a circular window could be removed from the centre of the implant—this allowed recordings to be made in the cortex underneath the Elvax sheet.

4.3 Results

4.3.1 Properties of muscimol Elvax

The dynamics of muscimol release from Elvax sheets are shown in figure 4.1, as calculated using radiolabelled muscimol. While muscimol release *in vitro* appears to be considerably higher in the first 2–3 weeks than it is thereafter, considerable amounts of drug are still being released from Elvax even after 6 weeks in phosphate-buffered saline. Furthermore, pieces of Elvax which have been implanted *in vivo* still contain measurable amounts of muscimol more than 160 days after implantation. This suggests that release into cortex is slower and more prolonged than it is into saline, probably because of the growth of new tissue around the implant.

For the animals which were implanted with muscimol Elvax, figure 4.2 shows how long after surgery the various testing runs took place. Comparison with figure 4.1 confirms that the implants are capable of releasing measurable amounts of muscimol for all the testing runs, even in the case of group PILOT.

That the muscimol actually has a physiological effect is shown in figure 4.3, from which it can be seen that muscimol-containing implants in visual cortex (V1) cause massive blockade of cortical activity for at least 6 weeks. During this time, the only activity remaining below the implants is insensitive to stimulus orientation, indicating that it is due to thalamic afferents rather than cortical neurons. Even 16 weeks after implantation the blockade below the implant was considerable. The lateral spread of the muscimol-induced inactivation appears to be of the order of 1 mm for a

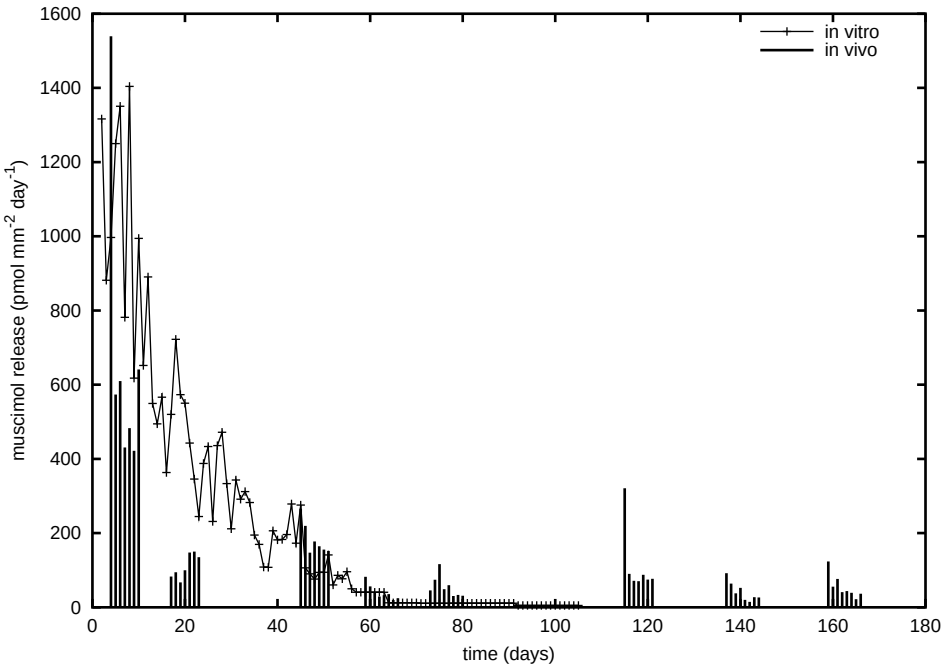

Figure 4.1: The release of muscimol from Elvax into phosphate-buffered saline at 37°C. One group of slices were kept *in vitro* throughout, and for these slices the abscissa represents the time elapsed since immersion in saline. Another group were first implanted *in vivo* before being removed and placed in saline; for these slices the abscissa represents the time since implantation. The thin line represents the mean release from *in vitro* slices ($n = 4$), while each group of 7–8 adjacent bars represents release from a single *in vivo* slice ($n = 8$).

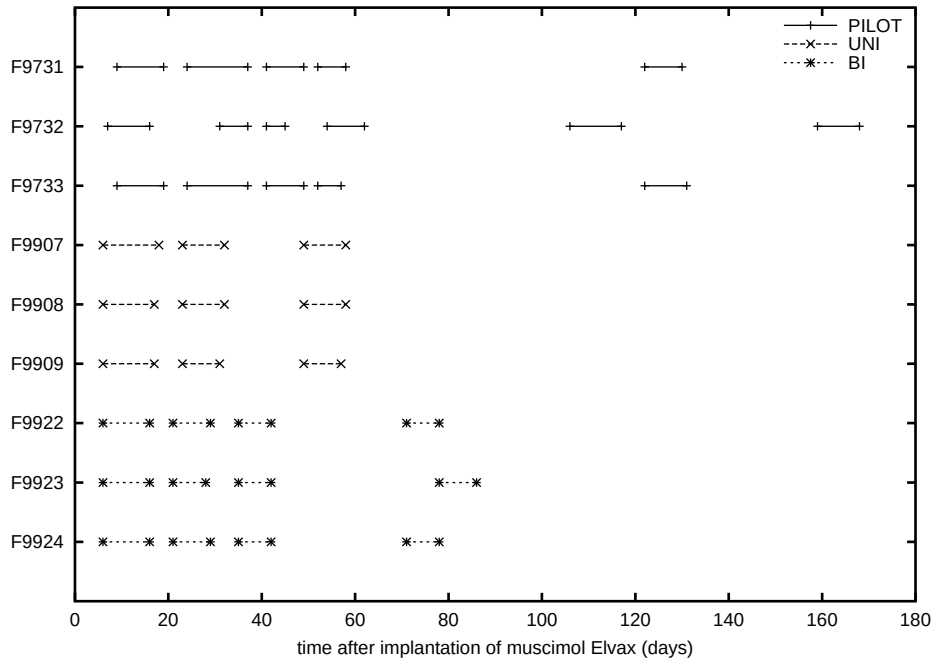

Figure 4.2: The chronology of behavioural testing with respect to muscimol Elvax implantation. The abscissa represents time after implantation surgery, on the same scale as in figure 4.1. Labels on the ordinate give the number of each animal. Each horizontal line represents a single testing run. Pre-implantation and post-removal testing runs are not shown.

2 week-old implant, decreasing thereafter, but it is restricted to the upper cortical layers. The recordings made with blank Elvax show that the effects on activity were down to the muscimol, and not the Elvax itself—while there may be an effect of blank Elvax in the very top 0.1–0.2 mm of cortex, it is small if it is present at all.

As well as these experiments in visual cortex, similar recordings were made in A1. The effect of the muscimol is more difficult to discern in this case, as the overall level of neural activity is lower in A1 than in V1. For this reason recordings were made every 50 μm in depth instead of every 0.1 mm, to make sure that no activity was missed. The results are shown in figure 4.4, and, as with the V1 recordings, there is hardly any activity beneath the implant. The small patch of stimulus-related activity in the 3-week recording is nearly 2.5 mm beneath the cortical surface, and so is almost certainly a subcortical structure (possibly the insula).

Given the localised nature of the effect of muscimol, the position of the Elvax implants is of some importance, and this is shown for one ferret in figure 4.5. As can be seen, the implant covers most of the middle ectosylvian gyrus, as well as some of the posterior ectosylvian gyrus, but leaves exposed a small area along the caudal edge of the suprasylvian sulcus. A1 in the ferret is associated with the middle ectosylvian gyrus[59, 124, 134] (also see figure 6.2).

4.3.2 Group PILOT

4.3.2a Performance after Elvax implantation

Figure 4.6 shows the overall performance of the animals in the PILOT group before, during and after implantation of muscimol Elvax over left A1. It suggests that there may be a small effect of muscimol Elvax at short stimulus durations, but it is expected on the basis of previous studies that these animals will show impaired localisation primarily in the side of space opposite to the implant. The scores for the right and left hemifields are shown separately in figure 4.7, and it is apparent that, at least on the first post-implantation run, responses to 40 ms stimuli are less accurate if the stimulus is on the right than if it is on the left. The timescale of these figures is given elsewhere (figure 4.2).

In interpreting figures 4.6 and 4.7 it should be noted that the data available differs from run to run, in that some key stimulus durations were not used in some of the post-implantation runs. This reflects the testing schedule used, as indicated in section 4.2.4a. Partly because of this, and partly because of the differences between the post-implantation testing of the different ani-

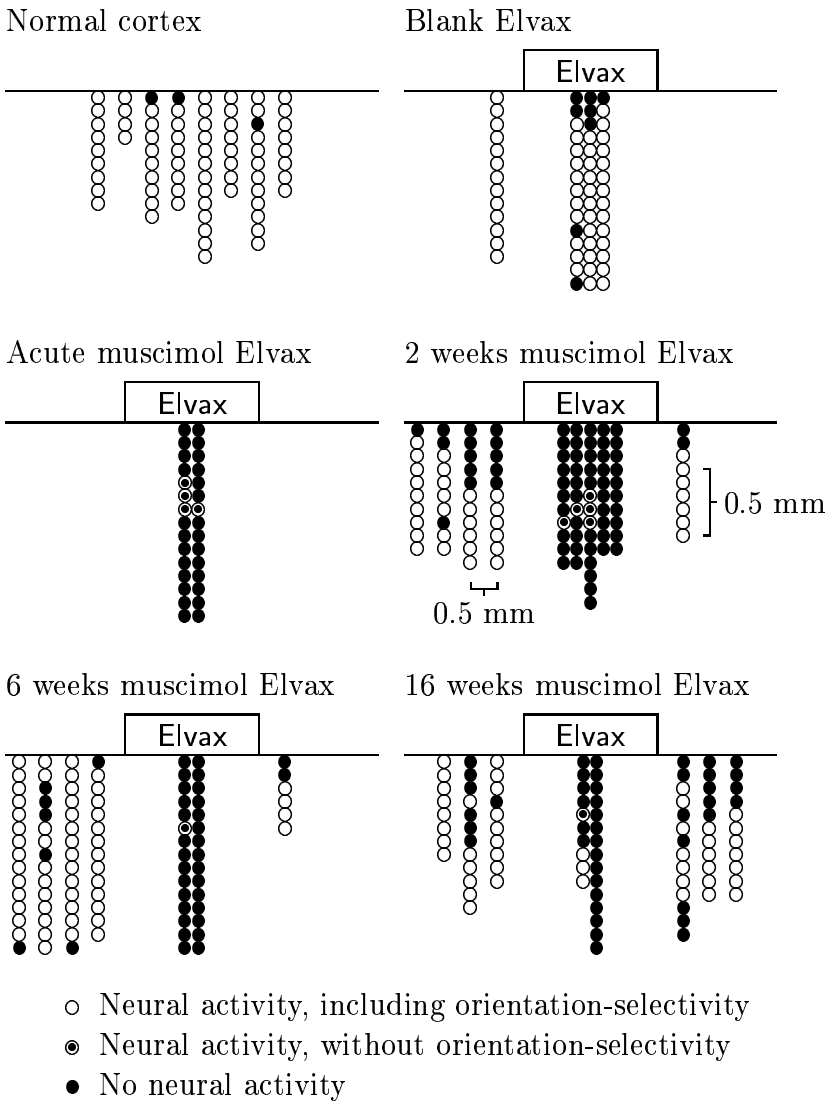

Figure 4.3: The spatial extent of the effect of muscimol Elvax, measured in V1. Each vertical column shows results from a single electrode penetration, with recordings taken every 0.1 mm in depth. Electrode penetrations through Elvax implants were made by first removing a small piece of the implant. Other electrode penetrations were made at 0.5 mm intervals from the edge of the implant. Each diagram represents recordings from 2 animals, except for the "normal cortex" condition which shows data from 4 animals.

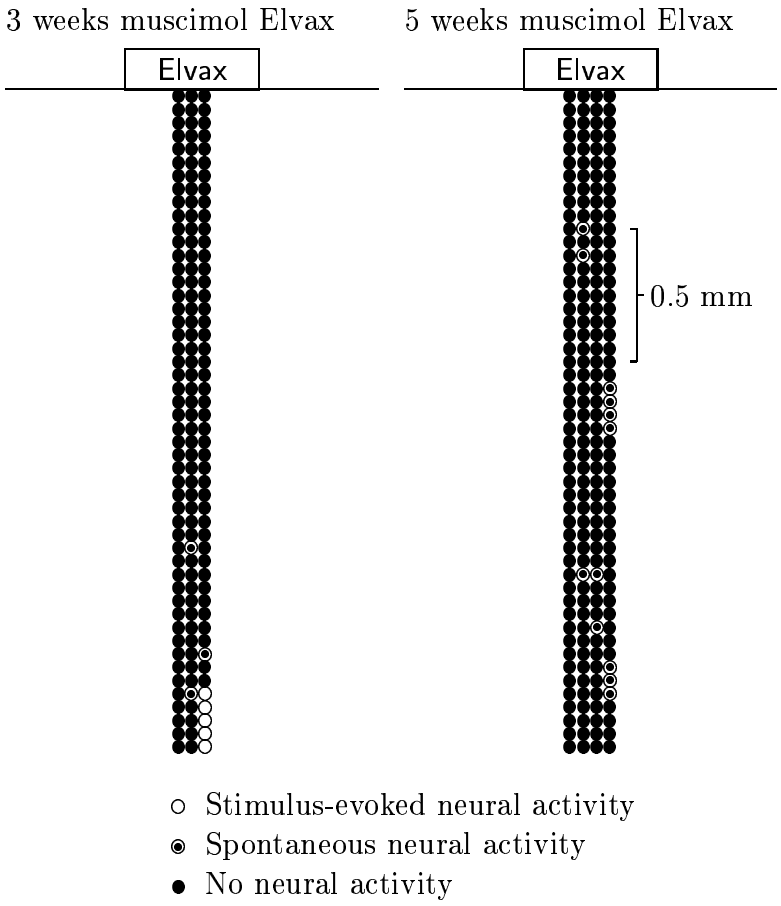

Figure 4.4: The spatial extent of the effect of muscimol Elvax, measured in A1. Each vertical column shows results from a single electrode penetration, with recordings taken every 50 μm in depth. Electrode penetrations through Elvax implants were made by first removing a small piece of the implant. Each diagram represents recordings from a single animal.

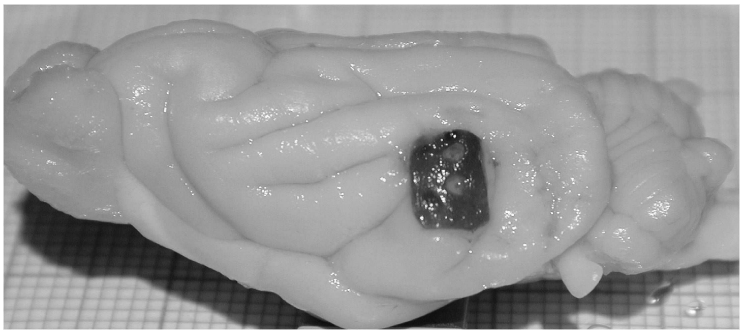

Figure 4.5: The brain of a ferret which had received a muscimol Elvax implant intended to cover A1. The piece of Elvax can clearly be seen. Two holes are visible in the Elvax—these were made to allow electrode recording underneath the implant, indicating that this animal is part of the methodological study described in the text.

mals in the group, it is not sensible to analyse these data in the same systematic manner as for the pre-implantation data. A slightly more basic measure of the effect of muscimol Elvax is the difference between pre-implantation and post-implantation performance for each animal, where performance is the proportion of trials on which the correct response is given. The mean of this difference, across the three animals, can be used as a test statistic. If the values of overall performance are used, the mean difference between pre-implantation and post-implantation scores is not significantly different from zero at any of the key durations. However, if the scores are calculated separately for the left and right hemifields (ignoring stimuli along the midline), then there is a significant difference in pre and post-implantation scores for 40 ms stimuli, but only in the right hemifield ($t = 6.00$, $p = 0.0267$ with 2 d.f. for a 2-tailed test).

4.3.2b Distribution of errors

Because of our 360° testing chamber, we are in a position to investigate whether the extra errors caused by the muscimol Elvax implants are "front-back errors" or not. Front-back errors are those which lie on the correct side of the animal's midline but the wrong side of its interaural axis, and figure 4.8 shows the mean frequency of these errors in response to 40 ms stimuli, both for the last pre-implantation run and the first post-implantation run. Because the classification of an error as being a front-back error is not independent of the error magnitude, the frequency of control errors of equal size is also shown; figure 3.8 gives a comprehensive list of error classifications.

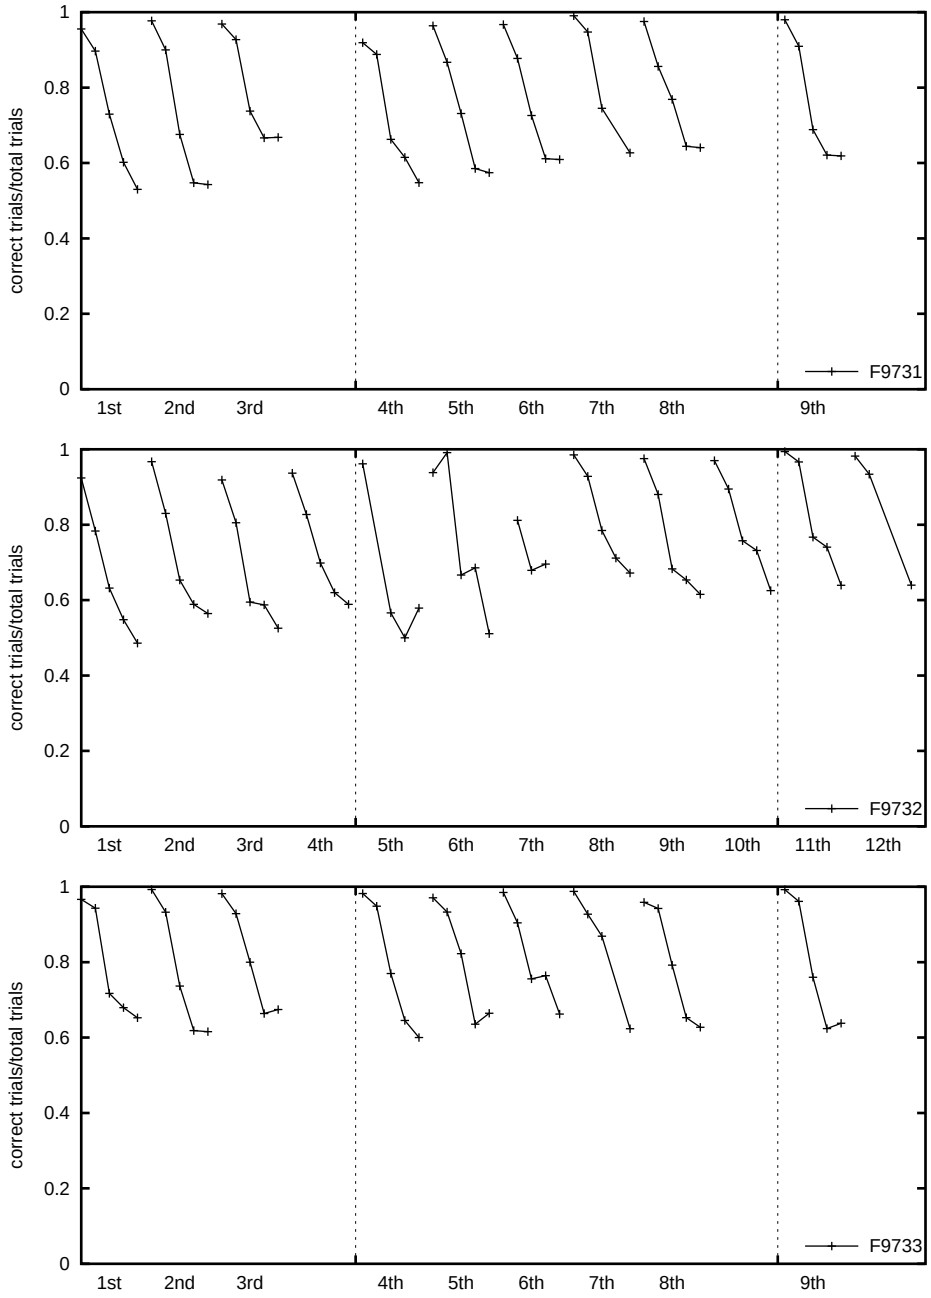

Figure 4.6: The performance of ferrets in group PILOT at key stimulus durations in each run. Performance is measured as the proportion of trials on which the correct response was given. Durations examined are 1 s, 500 ms, 200 ms, 100 ms and 40 ms. Within each run, trials are grouped by stimulus duration and plotted in descending order of duration. Vertical dotted lines represent the implantation (left) and removal (right) of muscimol Elvax.

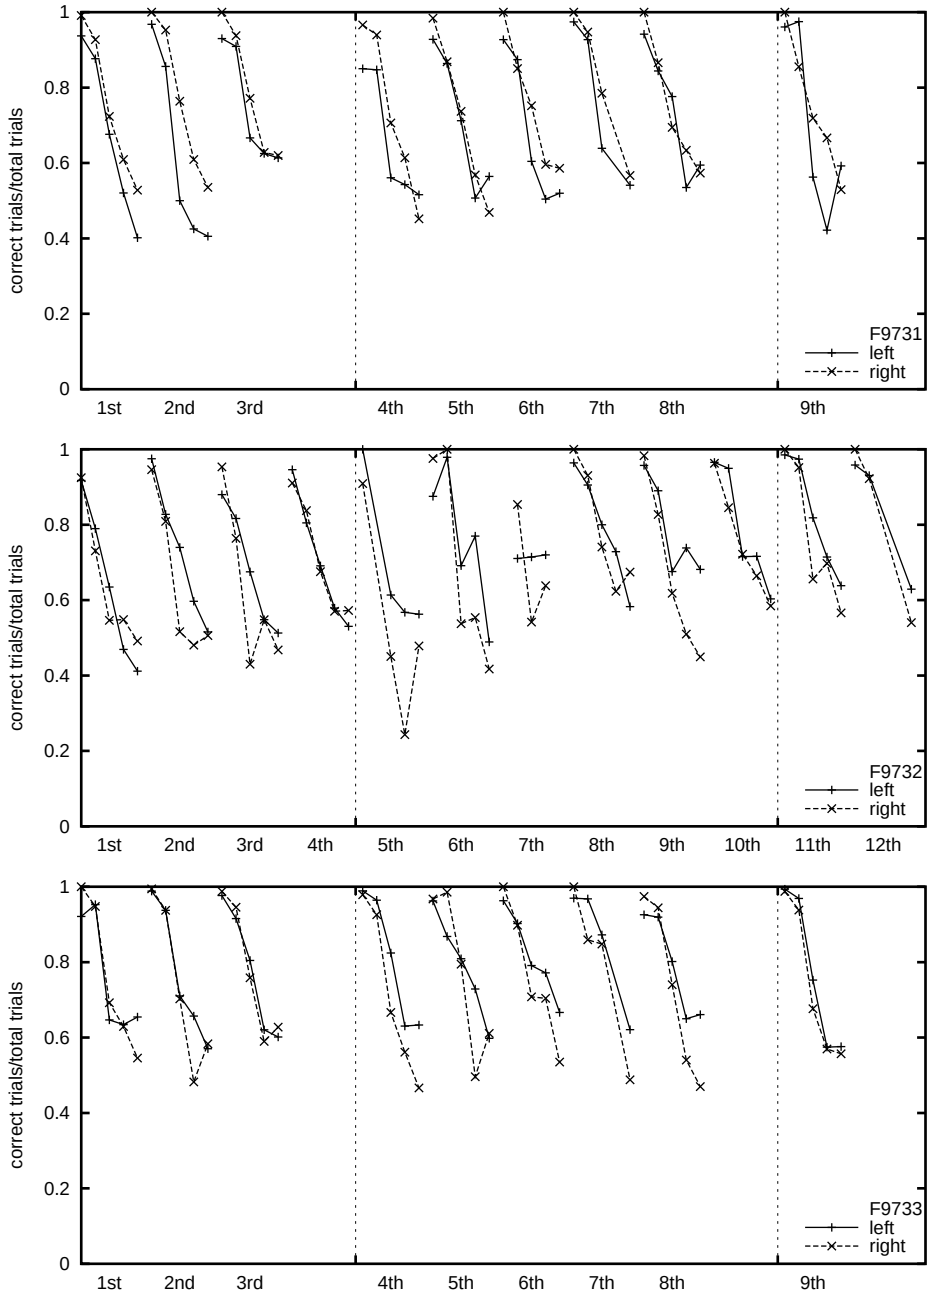

Figure 4.7: The performance of ferrets in group PILOT, with the results for stimuli presented in the left and right hemifields shown separately. Performance is measured as the proportion of trials to which the correct response was given. Trials where the stimulus was directly ahead of or behind the animal are ignored. Trials are grouped by stimulus duration as in figure 4.6. Vertical dotted lines represent the implantation (left) and removal (right) of muscimol Elvax.

As explained in section 3.3, control errors do not exist for stimulus locations of 0° or 180°, so front-back errors made at these positions are excluded from our analysis. It is obvious from figure 4.8 that the frequency of front-back errors in the right hemifield—the side contralateral to the implant—is higher after implantation than before, while the same is not the case for the control errors. This is confirmed by a within-subjects ANOVA carried out on the 40 ms trials of the first post-implantation run. The incidence of front-back errors, and equal-magnitude control errors, is measured as a proportion of the number of responses, subject to an arcsine-correction. Error type (front-back *vs.* control) and stimulus hemifield (right *vs.* left) are the within-subjects factors for this analysis, and the interaction between these two factors is the only significant effect ($F(1,2) = 40.9$, $p = 0.024$).

Although the data reveal a significant and hemifield-specific increase in the frequency of front-back errors following A1 inactivation, this is not a large effect. Figure 4.9 gives the distribution of responses for F9731, one of the animals in the PILOT group. It is clear from this that the errors made after after muscimol Elvax implantation are similarly distributed to those before surgery—this is a quantitative rather than qualitative difference.

4.3.2c Restricted-bandwidth stimuli

All of the animals in the PILOT group were tested with high-pass, low-pass and band-pass stimuli while muscimol Elvax was in place. This is important because if for some reason we have only partially inactivated A1 then the resulting impairment may be specific to a particular range of frequencies[56]. As noted in section 4.2.4a, the restricted-bandwidth stimuli were 40 ms in duration, and were presented in additional sessions at the end of an otherwise normal run: F9731 and F9733 were tested at the end of their second post-implantation run, while F9732 was tested at the end of her third post-implantation run.

The different restricted-bandwidth stimuli were all presented along with full-bandwidth stimuli (as a control) in a randomly interleaved manner, so a straightforward comparison can be made between the different conditions. Results for each stimulus are shown in the top left panel of figure 4.10, which suggests that performance with the 10–20 kHz band-pass stimuli is generally worse than for any of the other conditions. For each of the animals, performance on the various restricted-bandwidth conditions was compared statistically to the performance with the full-bandwidth stimulus. A z-test, based on a Binomial distribution, was used for this. All three animals showed a significant decrement in performance on the band-pass condition compared to the all-pass condition (F9731: $z = 3.86$, giving $p = 5.70 \times 10^{-5}$; F9732:

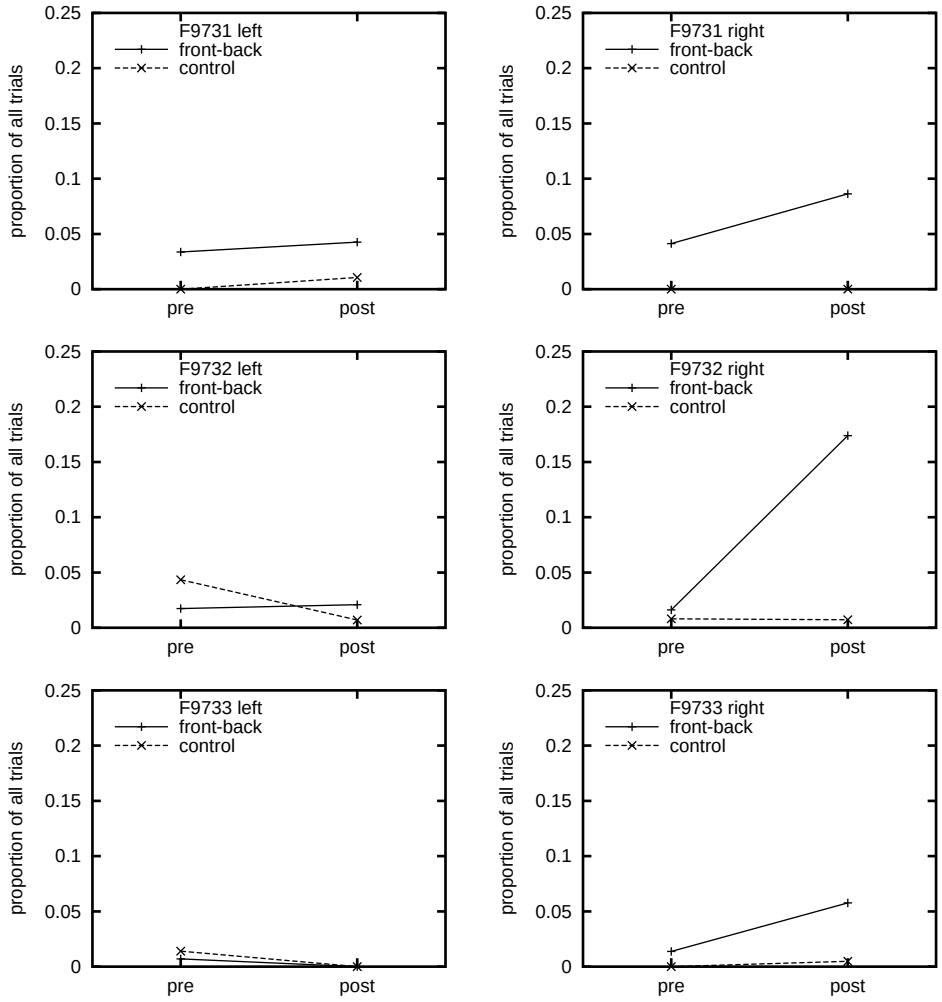

Figure 4.8: The mean incidence of front-back errors in the responses of the PILOT group to 40 ms stimuli, on the runs immediately before and immediately after implantation of muscimol Elvax over left A1. Error frequency is shown separately for the left and right hemifields. Errors of equal magnitude to front-back errors are shown as a control.

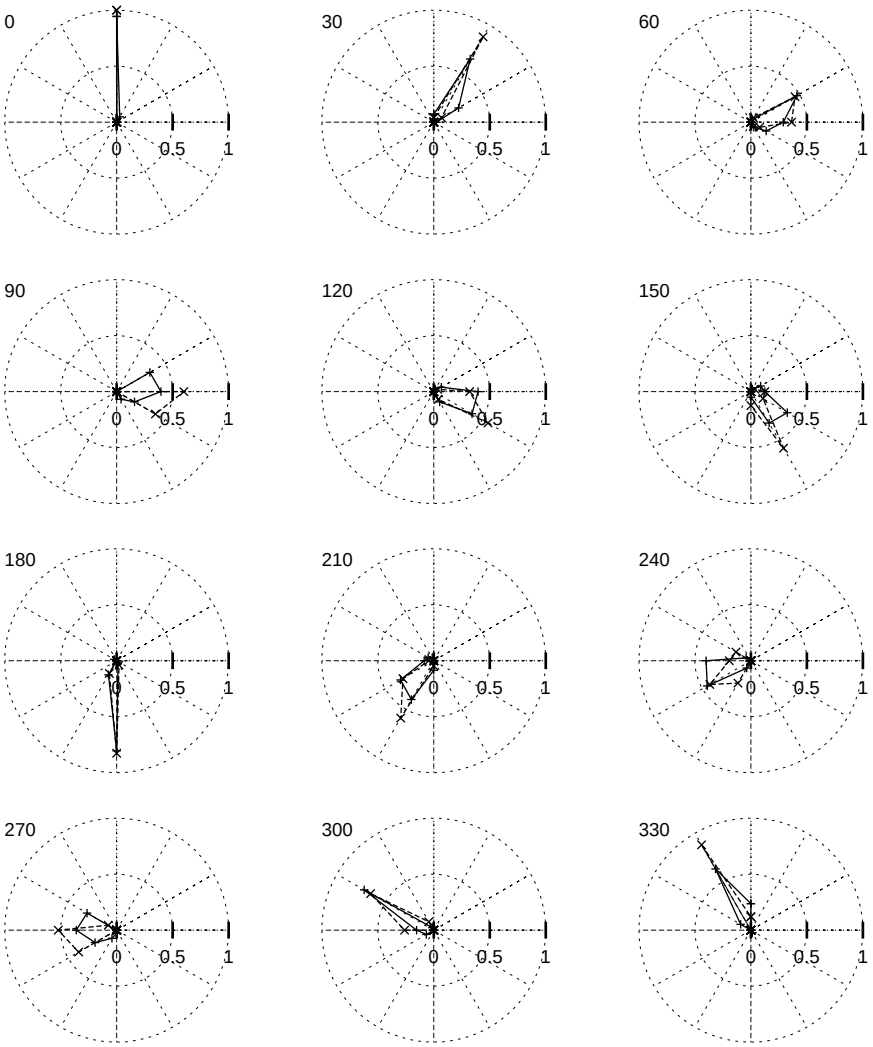

Figure 4.9: The distribution of responses to stimuli emanating from different positions, for an animal from the PILOT group (F9731). Solid lines show responses to 40 ms stimuli on the first post-implantation testing run, while dotted lines represent data from the last pre-implantation run (from figure 3.4). Different plots show the responses to different stimulus positions, with the label on each plot giving the stimulus azimuth. The angle of each point indicates the response position, and the radius of each point represents the proportion of trials on which the animal responded at that position.

$z = 1.90$, giving $p = 0.0285$; F9733: $z = 3.89$, giving $p = 5.02 \times 10^{-5}$). None of them showed significant effects of either the high-pass or low-pass conditions.

However, narrowing the bandwidth of a stimulus would be expected to remove valuable localisation cues, so there is nothing about the above results which implies a link to the inactivation of A1. Of greater importance is whether changes in stimulus bandwidth have a specific effect in the hemifield opposite the inactivated A1. The left and right hemifields are displayed separately in the remaining panels of figure 4.10. Neither F9731 nor F9733 shows a large difference between the two hemifields, but F9732 appears to be impaired on band-pass stimuli in the right hemifield but unaffected in the left. Again, z-tests were used to evaluate the significance of the results, this time comparing performance in the left hemifield with performance in the right hemifield. None of the animals showed any significant hemifield-related differences in any of the stimulus conditions.

4.3.2d Performance after Elvax removal

Although it does not bear on the behavioural effects of A1 inactivation, it is interesting to consider what happens after the removal of the muscimol Elvax. Comparing the difference in performance between the last pre-implantation run and the first post-removal run shows no significant difference at any of the key durations, either in terms of overall performance, or when looking specifically at the right hemifield.

4.3.3 Group BLANK

The function of this group was primarily to check that the process of surgery and the chronic presence of a drug-free implant did not itself cause any substantial behavioural impairment. On the basis of figure 4.3 we do not expect any cortical inactivation, but it is possible to imagine some other medical complication which might have an effect on animals' performance.

It is not possible to directly compare the scores for these animals with those from any other groups, as their training routine was so different. In particular, the BLANK animals did not have the extensive pre-implantation training of the PILOT animals, and yet they were not completely naïve like the animals in the other groups, as they had been trained on a different task which used similar apparatus. However, given that the animals in the PILOT group show a difference in the incidence of front-back errors between the two hemifields, it is reasonable to look for this as evidence of an effect on A1. Figure 4.11 shows the incidence of errors in the right and left hemifields.

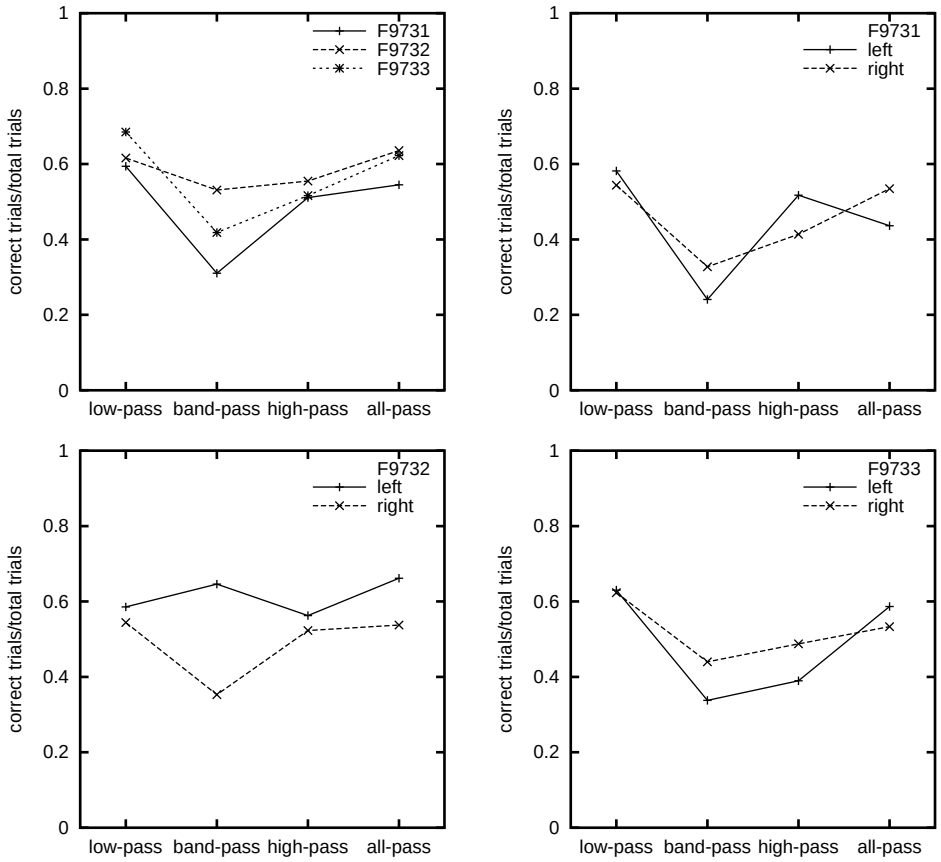

Figure 4.10: The performance of ferrets in group PILOT in response to differently filtered 40 ms stimuli. Performance is measured as the proportion of trials on which the correct response was given. The top left panel gives the overall performance score for each animal, while the other panels show the performance calculated separately for the left and right hemifields in each animal.

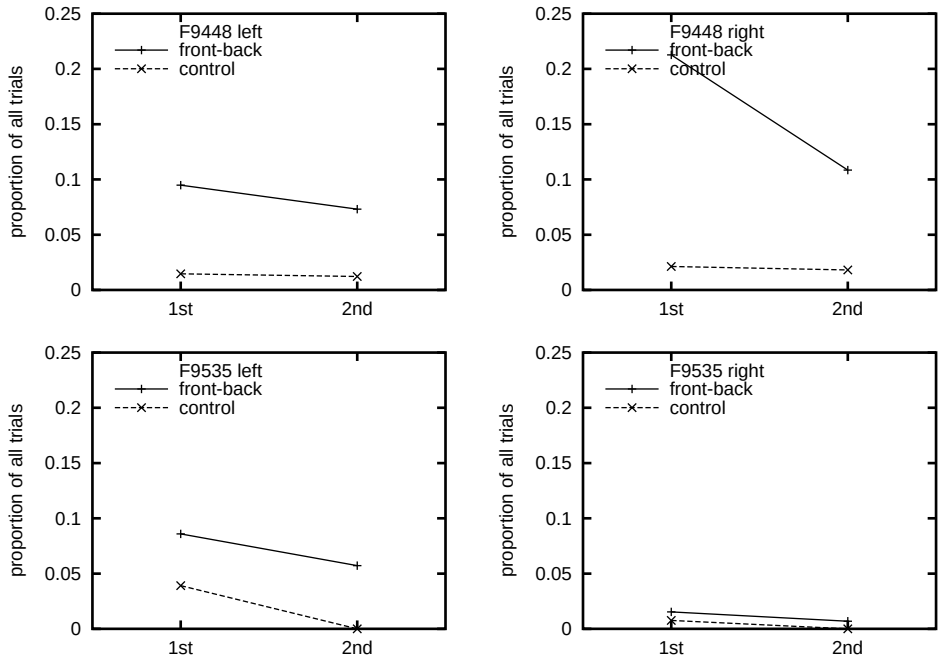

Figure 4.11: The mean incidence of front-back errors in the responses of the BLANK group to 40 ms stimuli, on the two runs following implantation of blank Elvax over right A1. Error frequency is shown separately for the left and right hemifields. Errors of equal magnitude to front-back errors are shown as a control.

While the number of errors of all types is greater than that for the PILOT groups (cf. figure 4.8), the differences between hemifields are in the opposite direction in the two animals, implying that these are individual differences rather than any effect of the implant. An ANOVA of the front-back errors for 40 ms stimuli, with error type and stimulus hemifield as within-subjects factors, shows no significant effects.

4.3.4 Groups UNI and BI

As the testing programme for the animals in these groups was very similar, they are considered here together. In fact, the analysis of the data from both of these groups is the same as that for the animals in Group NO-OP, and results for those animals (detailed in section 3.3) are quoted here throughout.

4.3.4a Overall performance

The first test carried out was to examine whether there were any changes in overall performance over the course of the 3 or 4 runs the animals were put through. Performance was measured as the proportion of trials to which the animal responded correctly. Trials from the same run were grouped together according to stimulus duration, ignoring those trials which were not from the key durations of 1 s, 500 ms, 200 ms, 100 ms or 40 ms. These figures were arcsine-corrected, and then subject to an analysis of variance (ANOVA), using both stimulus duration and run as within-subjects factors. The results, given in table 4.6, include a significant effect of run, reflecting an improvement in the animals' performance from one run to the next. In the case of the UNI animals the improvement appears to be different at different durations, as not only was the effect of run significant, but also the interaction between stimulus duration and run. The animals in the BI group did not show this interaction, although they did show a significant effect of run alone. The apparent changes in performance with testing run are by contrast to the data for animals in the NO-OP group (section 3.3), who do not show any such changes. However, as might be expected, the effect of stimulus duration is significant for all three groups of animals.

These data are shown graphically in figure 4.12, which covers a period of ~60 days of testing for the UNI animals and ~80 days for the BI animals (see figure 4.2). In both UNI and BI groups the run-related improvement in performance appears to take place mainly between the first and second runs. It is also interesting that there appears to be a rather sharp improvement between the 1 s trials and the 500 ms trials on the first run of the animals in the BI group—this is not something which is seen in the equivalent figure for the NO-OP animals (figure 3.2, bottom). With respect to this, it is worth remembering that stimuli are always used in decreasing order of stimulus duration during the course of a run, and so the 1 s trials on the first run represent the animals' first exposure to the task. Another apparent difference between groups is that performance by the BI group seems worse than that of the animals in either NO-OP or UNI groups, at least at short durations (see section 4.3.5).

4.3.4b Right-left differences

Given that the animals in the PILOT group show impaired performance which is specific to one hemifield, we also consider the effect of stimulus hemifield on these animals. Of particular interest is the UNI group, which might be expected to show a similar asymmetry to the PILOT group as both

Group UNI

source	df	SS	MS	F	p
duration	4	0.689	0.172	98.8	**<0.001**
error(duration)	8	0.0140	0.00174		
run	2	0.247	0.123	39.2	**0.002**
error(run)	4	0.0126	0.00315		
duration*run	8	0.0896	0.0112	5.39	**0.002**
error(duration*run)	16	0.0332	0.00208		

Group BI

source	df	SS	MS	F	p
duration	4	2.65	0.663	132	**<0.001**
error(duration)	8	0.0403	0.00504		
run	3	0.282	0.0939	72.3	**<0.001**
error(run)	6	0.00779	0.00130		
duration*run	12	0.264	0.0220	1.57	0.169
error(duration*run)	24	0.337	0.0140		

Table 4.6: ANOVA for animals in Groups UNI and BI, showing the differences in performance associated with stimulus duration and testing run. Performance is measured as the arcsine-corrected proportion of trials to which the correct response was given.

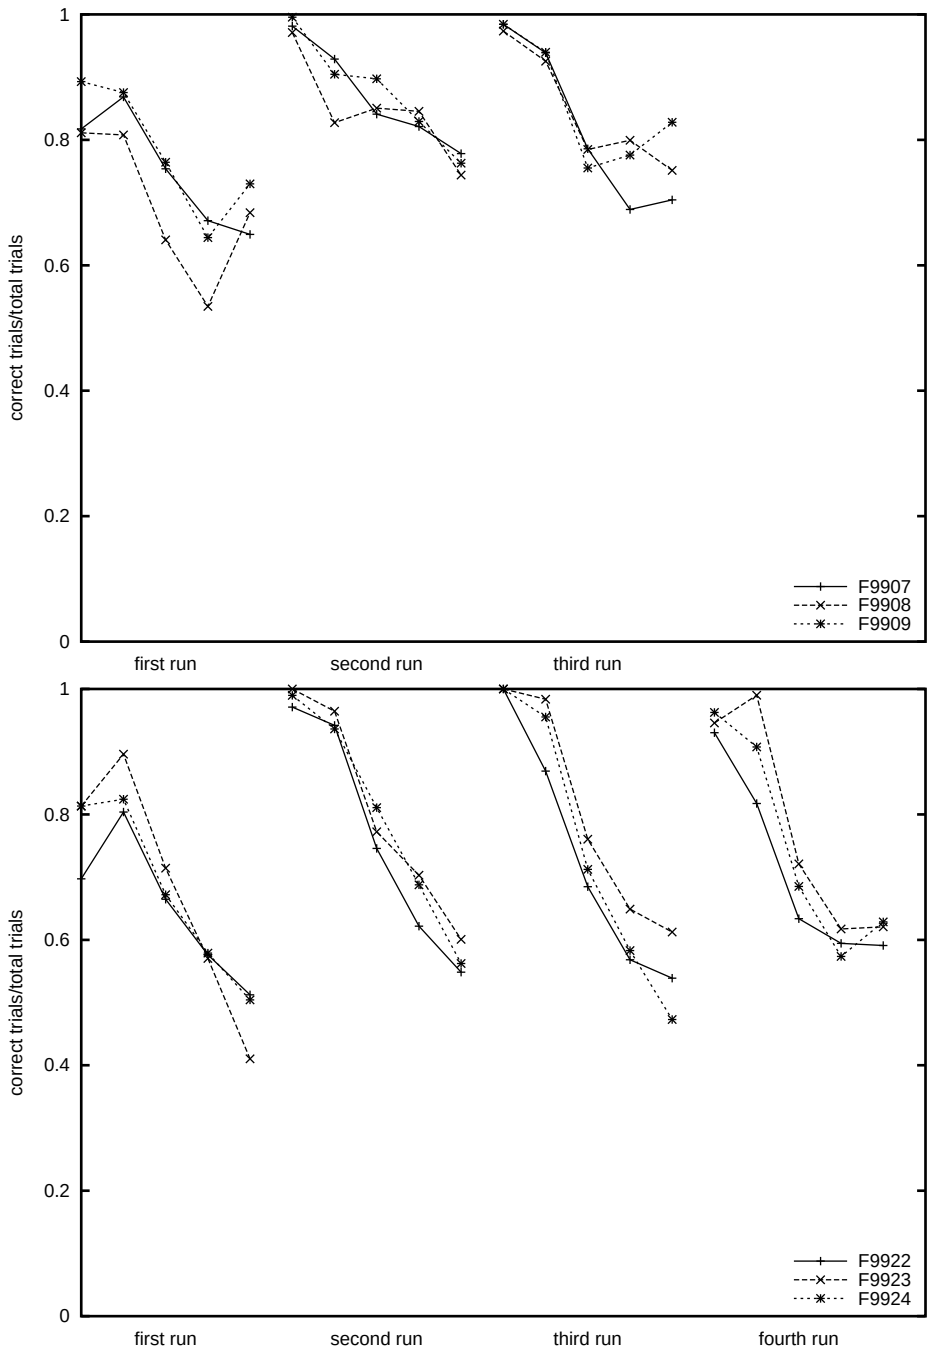

Figure 4.12: The performance of ferrets in group UNI (top) and BI (bottom) at key durations in each run. Performance is measured as the proportion of trials to which the correct response was given. Within each run, trials are grouped by stimulus duration, so the first data-point in each run represents trials at 1 s, the second represents trials at 500 ms, and so on for 200 ms, 100 ms and 40 ms.

have unilateral muscimol Elvax implants. The measure of performance used was the same as that used elsewhere, viz. the proportion of trials to which the correct response is given, calculated separately for trials on which stimuli were in the left and right hemifields. Trials where the stimulus was either directly ahead of or behind the animal were ignored. This measure was then arcsine-corrected and used as the basis for a within-subjects ANOVA, with stimulus duration, stimulus hemifield and testing run as the within-subjects factors—the results are given in table 4.7.

For both groups of animals, table 4.7 shows that the effects of stimulus duration, testing run and the interaction between duration and run are significant—these findings are similar to those given in table 4.6 for the animals' overall performance. However, the results for group UNI also show a three-way interaction between duration, run and hemifield, something which is not the case for the BI group (or the NO-OP group). This suggests that for these animals, there may be differences between the left and right hemifields for a specific stimulus duration and on a particular run.

Graphs of these right-left comparisons are shown in figure 4.13. It is interesting to note that in the case of the UNI group, performance in the right hemifield for the very first trials (1 s stimuli in the first run) was consistently worse than that in the left hemifield under the same conditions. Surprisingly, however, on the second and third runs it is performance for the left hemifield which is worse, at least for F9908 and F9909 (figure 4.13, middle left and bottom left).

4.3.4c Front-back errors

Front-back errors are mislocalisations which lie on the correct side of the animal's midline but the wrong side of its interaural axis. As with the animals in the NO-OP group, we examine the incidence of these errors as a proportion of the total number of responses made by the animals. Because the classification of an error as being a front-back error is not independent of the error magnitude, we compare the incidence of front-back errors with the incidence of control errors of equal size. A list of front-back errors and control errors is given in figure 3.8. Errors made to stimuli along the midline are excluded from this analysis as no control errors exist for these positions.

Table 4.8 shows a within-subjects ANOVA for the (arcsine-corrected) frequencies of these errors, using the type of error, the stimulus duration and the testing run as within-subjects factors. Animals in both groups showed a significant difference between the incidence of front-back errors and control errors, as well as an interaction between error-type and stimulus duration; both these effects are also shown by animals in the NO-OP group

Group UNI

source	df	SS	MS	F	p
duration	4	1.76	0.440	77.2	**<0.001**
error(duration)	8	0.0456	0.00570		
run	2	0.474	0.237	21.0	**0.008**
error(run)	4	0.0450	0.0113		
hemifield	1	0.0368	0.0368	1.48	0.348
error(hemifield)	2	0.0497	0.0249		
duration*run	8	0.194	0.0243	3.00	**0.029**
error(duration*run)	16	0.130	0.00811		
duration*hemifield	4	0.0571	0.0143	1.69	0.244
error(duration*hemifield)	8	0.0676	0.00844		
run*hemifield	2	0.0339	0.0170	2.16	0.231
error(run*hemifield)	4	0.0313	0.00783		
duration*run*hemifield	8	0.133	0.0167	3.97	**0.009**
error(duration*run*hemifield)	16	0.0671	0.00419		

Group BI

source	df	SS	MS	F	p
duration	4	6.46	1.61	148	**<0.001**
error(duration)	8	0.0873	0.0109		
run	3	0.607	0.202	64.6	**<0.001**
error(run)	6	0.0188	0.00313		
hemifield	1	0.172	0.172	13.5	0.067
error(hemifield)	2	0.0255	0.0128		
duration*run	12	0.476	0.0397	7.34	**<0.001**
error(duration*run)	24	0.130	0.00541		
duration*hemifield	4	0.0431	0.0108	2.35	0.141
error(duration*hemifield)	8	0.0367	0.00458		
run*hemifield	3	0.0663	0.0221	3.50	0.090
error(run*hemifield)	6	0.0379	0.00631		
duration*run*hemifield	12	0.105	0.00878	2.06	0.064
error(duration*run*hemifield)	24	0.103	0.00427		

Table 4.7: ANOVA for animals in Groups UNI and BI, showing the differences in performance associated with stimulus duration, stimulus position (right hemifield *vs.* left hemifield) and testing run. Performance is measured as the arcsine-corrected proportion of trials to which the correct response was given.

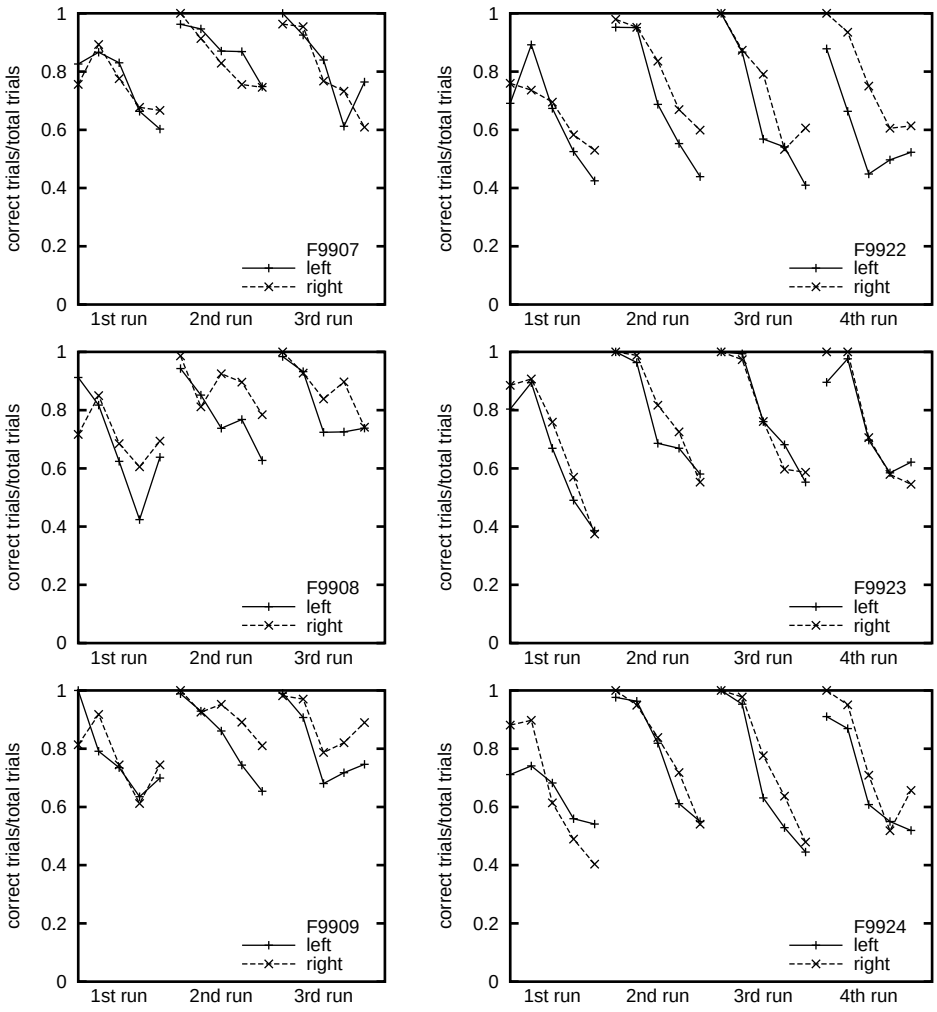

Figure 4.13: The performance of ferrets in group UNI (left column) and BI (right column), with the results for stimuli presented in the left and right hemifields shown separately. Performance is measured as the proportion of trials to which the correct response was given. Data from trials where the stimulus was either directly ahead of or directly behind the animal are not included in these plots. Other than the separation into two hemifields, trials are grouped by stimulus duration within each run, as in figure 4.12.

(section 3.3). However, the animals in group UNI also showed a three-way interaction between error-type, duration and testing run, while this was not the case for either the animals in the BI group or those in the NO-OP group.

The data underlying this analysis are shown in figure 4.14. In both groups it appears that front-back errors are more common than control errors of equal magnitude, particularly at short stimulus durations—this reflects the effect of error-type and its interaction with duration shown by the ANOVA (table 4.8). It is also obvious that front-back errors at short durations are considerably more common in the BI group animals than in the UNI group animals or the NO-OP animals (figure 3.9, right), while the proportion of control errors is similar in each of the groups (see also section 4.3.5).

To show the incidence of front-back errors in context, figures 4.15 and 4.16 show the distribution of all responses made by two example animals. Both F9907 (figure 4.15) and F9922 (figure 4.16) make some front-back errors—this is shown by, for example, the responses at 150° made by F9907 or F9922's responses at 240°. A similarly-trained normal animal also makes similar errors, however, as is shown in figure 3.6. The main conclusion borne out by these figures is that the animal from the BI group simply makes more mistakes than the animals from the other groups.

As well as examining the overall frequency of front-back errors, it is of particular interest to group the front-back errors by stimulus hemifield. This analysis was carried out for both UNI and BI groups, although it is particularly pertinent to the UNI group. As before, the arcsine-corrected error frequencies were used as the basis for an ANOVA, but this time there were four within-subjects factors: stimulus duration, testing run, error-type and stimulus hemifield. The resulting ANOVA tables are given in tables 4.9 and 4.10, for groups UNI BI, respectively. The tables are quite complicated due to the large number of factors, but most of the effects simply replicate the analyses already shown. For the BI group animals no hemifield-specific effects are expected, and indeed there are none in table 4.10. For the UNI animals, though, table 4.9 reveals both an interaction between stimulus hemifield and error-type, and another between stimulus hemifield, testing run and stimulus duration. There is no significant four-way interaction for these animals, however, which implies that the interaction between stimulus hemifield and error-type is not specific to short-duration stimuli. The underlying data for the UNI animals are given in figure 4.17, which shows that the number of errors of all types appears to be quite variable between conditions. It is difficult to pick out any general trends, but *post hoc* analysis shows that front-back error incidence was slightly higher in the right hemifield than the left, while the reverse is the case for control error incidence. However, because this was not specific to short durations, it is probably not related to A1 inactivation.

Group UNI

source	df	SS	MS	F	p
duration	4	0.137	0.0343	36.0	**<0.001**
error(duration)	8	0.0762	0.000952		
run	2	0.00701	0.00351	0.761	0.525
error(run)	4	0.0184	0.00461		
error-type	1	0.0642	0.0642	134	**0.007**
error(error-type)	2	0.000961	0.000481		
duration*run	8	0.0339	0.00423	3.236	**0.022**
error(duration*run)	16	0.0209	0.00131		
duration*error-type	4	0.0727	0.0182	19.4	**<0.001**
error(duration*error-type)	8	0.00749	0.000937		
run*error-type	2	0.000838	0.000419	0.242	0.795
error(run*error-type)	4	0.00691	0.00173		
duration*run*error-type	8	0.00987	0.00123	3.23	**0.022**
error(duration*run*error-type)	16	0.00611	0.000382		

Group BI

source	df	SS	MS	F	p
duration	4	0.596	0.149	81.3	**<0.001**
error(duration)	8	0.0147	0.00183		
run	3	0.0432	0.0144	7.04	**0.022**
error(run)	6	0.0123	0.00205		
error-type	1	0.218	0.218	255	**0.004**
error(error-type)	2	0.00171	0.000853		
duration*run	12	0.0972	0.00810	6.06	**<0.001**
error(duration*run)	24	0.0321	0.00134		
duration*error-type	4	0.160	0.0400	18.3	**<0.001**
error(duration*error-type)	8	0.0175	0.00219		
run*error-type	3	0.00474	0.00158	2.95	0.121
error(run*error-type)	6	0.00322	0.000536		
duration*run*error-type	12	0.0105	0.000875	0.756	0.686
error(duration*run*error-type)	24	0.0278	0.00116		

Table 4.8: ANOVA for animals in Groups UNI and BI, showing the differences in error likelihood associated with stimulus duration, error type (front-back *versus* control) and testing run. Error likelihood is measured as the arcsine-corrected proportion of responses which are of the given type.

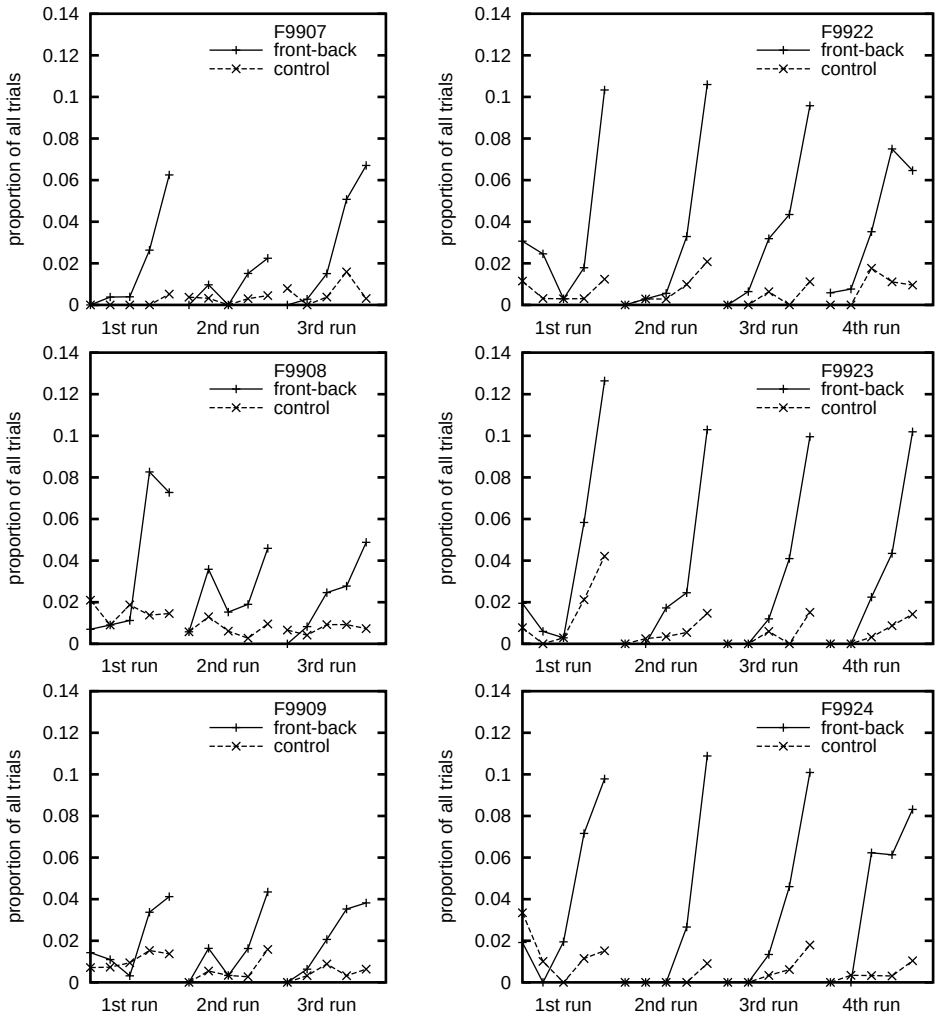

Figure 4.14: The incidence of front-back errors, as a proportion of the total number of responses, in the responses of the ferrets in group UNI (left column) and BI (right column). Errors which are of equal magnitude to front-back errors are shown as a control. Errors are grouped by stimulus duration within each run, in the same way as trials are grouped in figure 4.12.

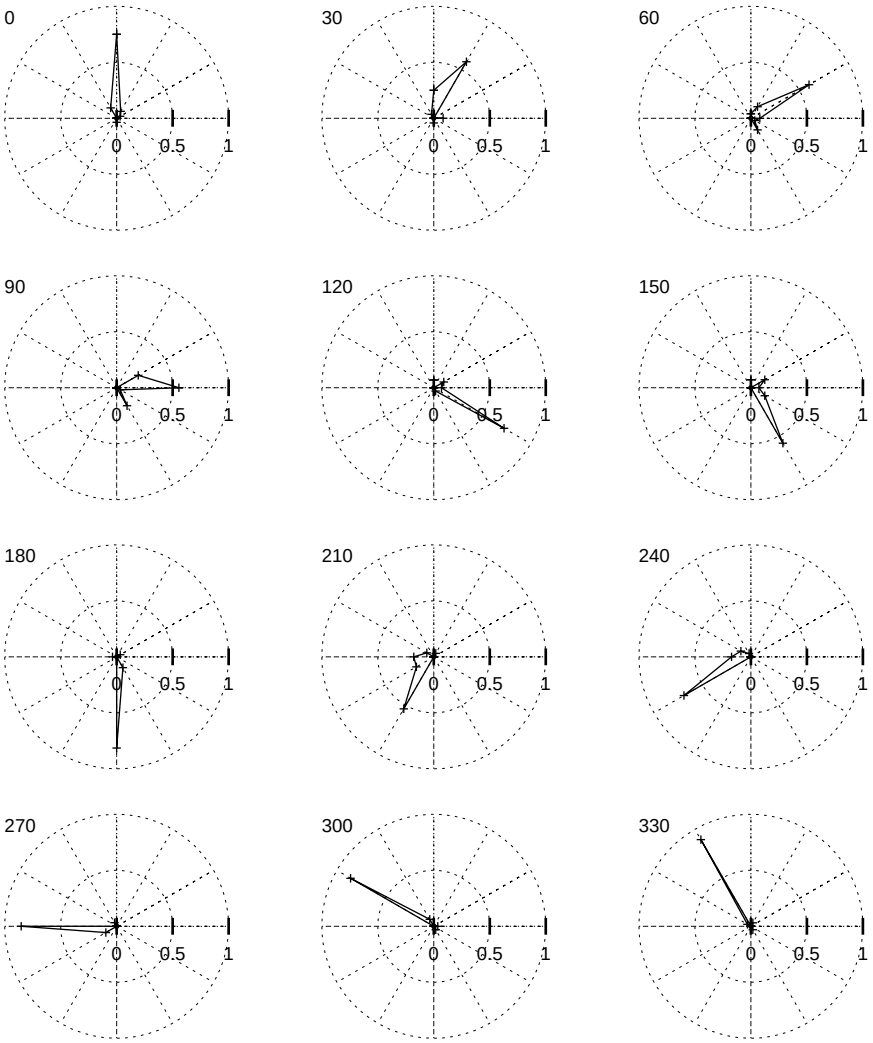

Figure 4.15: The distribution of responses to stimuli emanating from different positions, for an animal from the UNI group (F9907). Data shown are from responses to 40 ms stimuli on the third testing run. Plotting conventions are as in figure 4.9.

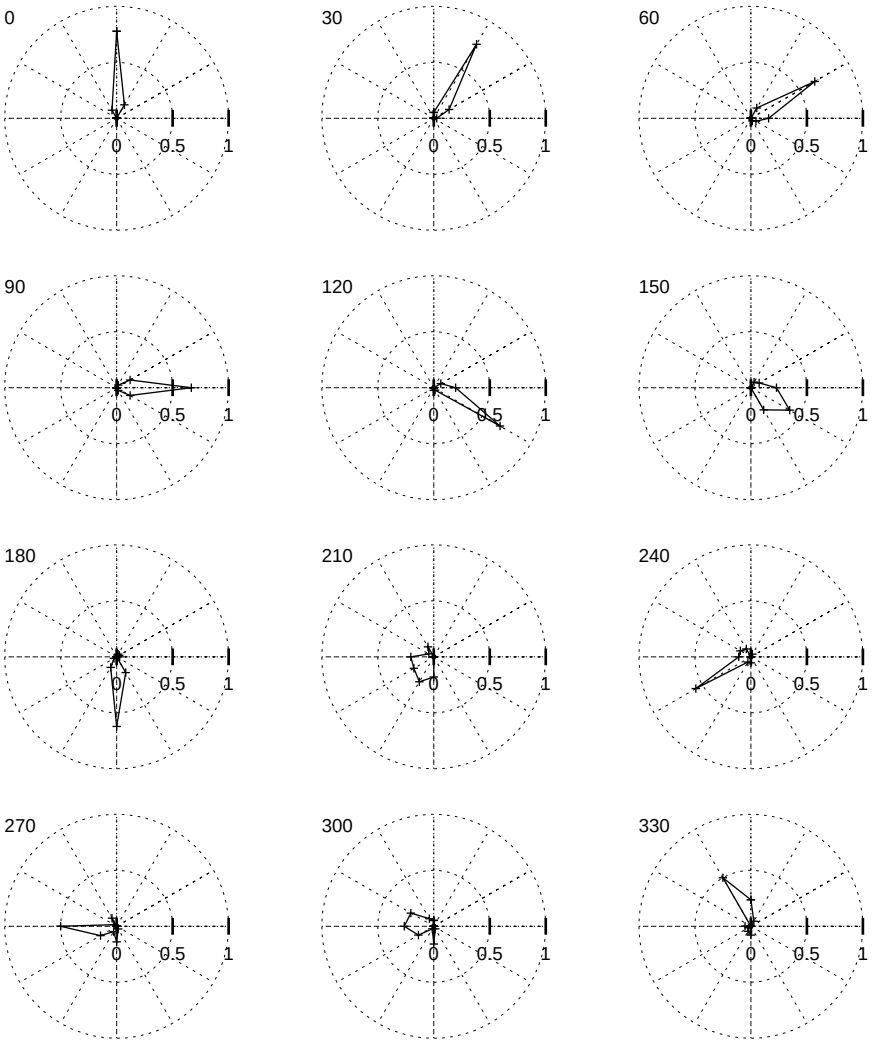

Figure 4.16: The distribution of responses to stimuli emanating from different positions, for an animal from the BI group (F9922). Data shown are from responses to 40 ms stimuli on the third testing run. Plotting conventions are as in figure 4.9.

Group UNI

source	df	SS	MS	F	p
hemifield	1	9.96×10^{-6}	9.96×10^{-6}	0.002	0.971
error(hemifield)	2	0.0122	0.00609		
error-type	1	0.186	0.186	151	**0.007**
error(error-type)	2	0.00247	0.00123		
run	2	0.0225	0.0113	1.54	0.319
error(run)	4	0.0293	0.00732		
duration	4	0.388	0.0970	37.4	**<0.001**
error(duration)	8	0.0208	0.00259		
hemifield*error-type	1	0.0459	0.0459	20.1	**0.046**
error(hemifield*error-type)	2	0.00456	0.00228		
hemifield*run	2	0.0112	0.00560	0.879	0.483
error(hemifield*run)	4	0.0255	0.00637		
error-type*run	2	0.00250	0.00125	0.282	0.768
error(error-type*run)	4	0.0178	0.00444		
hemifield*error-type*run	2	0.0114	0.00570	6.65	0.054
error(hemifield*error-type*run)	4	0.00343	0.000858		
hemifield*duration	4	0.00713	0.00178	0.572	0.691
error(hemifield*duration)	8	0.0249	0.00312		
error-type*duration	4	0.182	0.0454	26.6	**<0.001**
error(error-type*duration)	8	0.0137	0.00171		
hemifield*error-type*duration	4	0.0127	0.00317	1.62	0.260
error(hemifield*error-type*duration)	8	0.0157	0.00196		
run*duration	8	0.0569	0.00711	3.40	**0.018**
error(run*duration)	16	0.0334	0.00209		
hemifield*run*duration	8	0.0669	0.00836	3.32	**0.020**
error(hemifield*run*duration)	16	0.0403	0.00252		
error-type*run*duration	8	0.0223	0.00279	2.50	0.056
error(error-type*run*duration)	16	0.0178	0.00111		
hemifield*error-type*run*duration	8	0.0346	0.00432	1.84	0.143
error(hemifield*error-type *run*duration)	16	0.0376	0.00235		

Table 4.9: ANOVA for animals in Group UNI, showing the differences in error likelihood associated with stimulus duration, stimulus position (right *vs.* left hemifield), error type (front-back *vs.* control) and testing run. Error likelihood is measured as the arcsine-corrected proportion of responses which are of the given type.

Group BI

source	df	SS	MS	F	p
hemifield	1	0.103	0.103	4.91	0.157
error(hemifield)	2	0.0418	0.0209		
error-type	1	0.554	0.554	412	**0.002**
error(error-type)	2	0.00269	0.00135		
run	3	0.0790	0.0263	6.21	**0.029**
error(run)	6	0.0254	0.00424		
duration	4	1.47	0.368	94.1	**<0.001**
error(duration)	8	0.0313	0.00391		
hemifield*error-type	1	0.00415	0.00415	2.32	0.268
error(hemifield*error-type)	2	0.00359	0.00179		
hemifield*run	3	0.00391	0.00130	0.234	0.870
error(hemifield*run)	6	0.0335	0.00558		
error-type*run	3	0.0125	0.00416	3.99	0.070
error(error-type*run)	6	0.00625	0.00104		
hemifield*error-type*run	3	0.00388	0.00129	0.389	0.765
error(hemifield*error-type*run)	6	0.0199	0.00332		
hemifield*duration	4	0.0463	0.0116	2.82	0.099
error(hemifield*duration)	8	0.0329	0.00411		
error-type*duration	4	0.438	0.110	24.7	**<0.001**
error(error-type*duration)	8	0.0355	0.00444		
hemifield*error-type*duration	4	0.00724	0.00181	1.66	0.252
error(hemifield*error-type*duration)	8	0.00875	0.00109		
run*duration	12	0.179	0.0149	5.54	**<0.001**
error(run*duration)	24	0.0645	0.00269		
hemifield*run*duration	12	0.0770	0.00642	1.74	0.119
error(hemifield*run*duration)	24	0.0884	0.00368		
error-type*run*duration	12	0.0327	0.00273	1.11	0.396
error(error-type*run*duration)	24	0.0589	0.00246		
hemifield*error-type*run*duration	12	0.0254	0.00212	2.09	0.061
error(hemifield*error-type *run*duration)	24	0.0244	0.00102		

Table 4.10: ANOVA for animals in Group BI, showing the differences in error likelihood associated with stimulus duration, stimulus position (right *vs.* left hemifield), error type (front-back *vs.* control) and testing run. Error likelihood is measured as the arcsine-corrected proportion of responses which are of the given type.

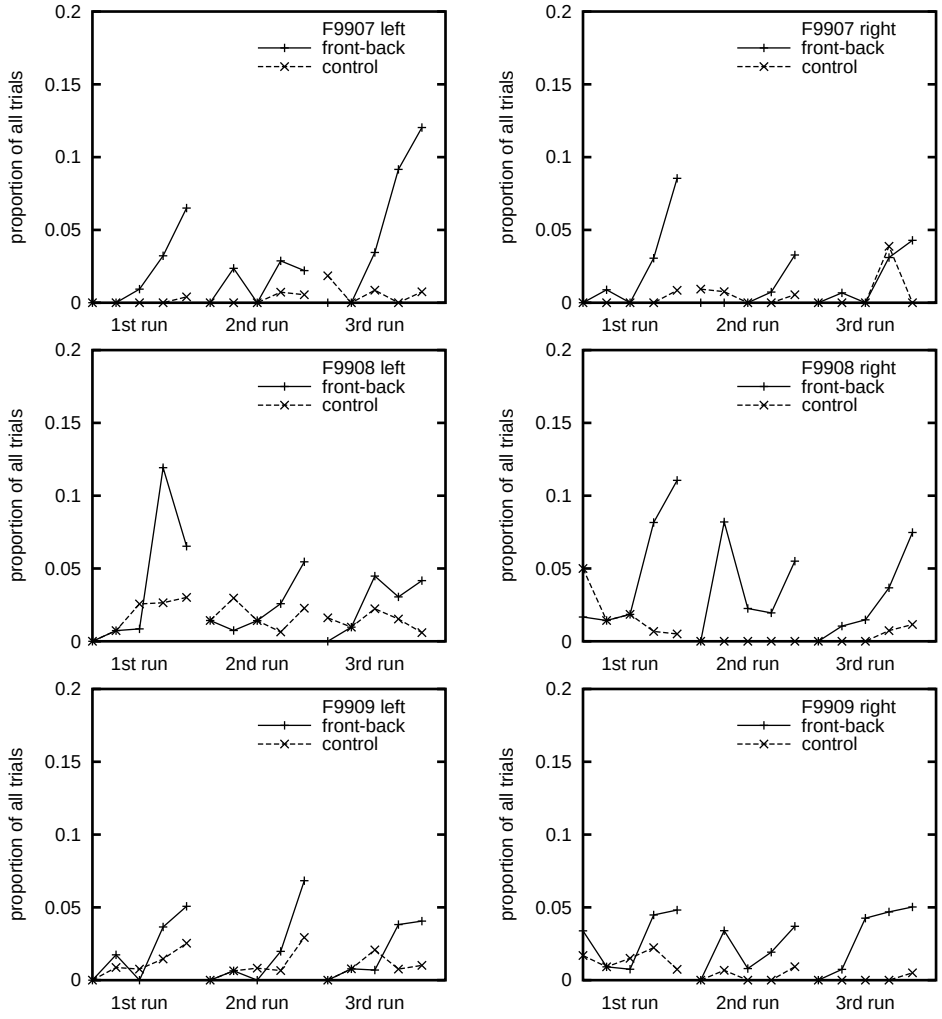

Figure 4.17: The incidence of front-back errors in the responses of ferrets in group UNI, with the left and right hemifields shown separately (left and right columns, respectively). Incidence is expressed as a proportion of the total number of responses, and errors of equal magnitude to the front-back errors are shown as a control. Errors are grouped by stimulus duration within each run, in the same way as trials are grouped in figure 4.12.

source	df	SS	MS	F	p
group	2	0.115	0.0574	0.944	0.440
error(group)	6	0.365	0.0608		
duration	4	3.10	0.775	197	**<0.001**
duration*group	8	0.306	0.0383	9.74	**<0.001**
error(duration)	24	0.0944	0.00393		
run	2	0.582	0.291	42.2	**<0.001**
run*group	4	0.0427	0.0107	1.55	0.251
error(run)	12	0.0827	0.00689		
duration*run	8	0.240	0.0300	8.66	**<0.001**
duration*run*group	16	0.135	0.00846	2.44	**0.009**
error(duration*run)	48	0.166	0.00346		

Table 4.11: ANOVA table for comparison of the overall performance of ferrets between Groups NO-OP, UNI and BI. Performance is measured as the arcsine-corrected proportion of trials to which the correct response was given.

4.3.5 NO-OP *vs.* UNI *vs.* BI

As the animals in groups NO-OP, UNI and BI have all been tested on the same schedule, it is possible to make a direct statistical comparison between them using a partial within-subjects ANOVA. In such an analysis, group forms a between-subjects factor, while the other factors under consideration (stimulus duration, testing run, etc.) are within-subjects factors. Table 4.11 shows the results of such an ANOVA for the ferrets' overall performance. Performance was measured in the same way as in tables 3.3 and 4.6, although the fourth run of the animals from the BI group was ignored as neither the NO-OP animals nor the UNI animals had been put through a fourth run.

In interpreting this analysis, it is important to take into account that most of the factors shown have already been analysed as part of the group by group ANOVAs presented earlier. What is important here is whether there is any effect of the group the animal was in (NO-OP, UNI or BI), either alone or in interaction with other effects. In fact, the main effect of group is not significant, but the interaction between group and stimulus duration is. *Post hoc* tests show this is attributable to the difference in performance for 40 ms stimuli between the animals in the BI group and the animals in the other groups. There is also a significant three-way interaction between group, stimulus duration and testing run—this ties in with the observation that the UNI and BI animals improve from one run to the next, while the NO-OP animals do not.

Table 4.12 shows a similar analysis, but this time including stimulus

source	df	SS	MS	F	p
group	2	0.296	0.148	1.33	0.333
error(group)	6	0.668	0.111		
duration	4	7.64	1.91	208	**<0.001**
duration*group	8	0.633	0.0791	8.61	**<0.001**
error(duration)	24	0.221	0.00919		
run	2	1.13	0.565	33.5	**<0.001**
run*group	4	0.110	0.0274	1.63	0.231
error(run)	12	0.202	0.0169		
hemifield	1	0.105	0.105	6.52	**0.043**
hemifield*group	2	0.00535	0.00268	0.167	0.850
error(hemifield)	6	0.0963	0.0161		
duration*run	8	0.484	0.0605	6.68	**<0.001**
duration*run*group	16	0.211	0.0132	1.45	0.158
error(duration*run)	48	0.435	0.00906		
duration*hemifield	4	0.0283	0.00708	0.879	0.491
duration*hemifield*group	8	0.0754	0.00942	1.17	0.356
error(duration*hemifield)	24	0.193	0.00805		
run*hemifield	2	0.0149	0.00745	1.50	0.261
run*hemifield*group	4	0.0379	0.00994	2.01	0.158
error(run*hemifield)	12	0.0594	0.00495		
duration*run*hemifield	8	0.0540	0.00675	1.93	0.078
duration*run*hemifield*group	16	0.152	0.00952	2.72	**0.004**
error(duration*run*hemifield)	48	0.168	0.00351		

Table 4.12: ANOVA table for comparison of the performance in the right and left hemifields between ferrets in Groups NO-OP, UNI and BI. Performance is measured as the arcsine-corrected proportion of trials to which the correct response was given.

position—left *versus* right hemifield—as a within-subjects factor. The measure of performance is the same as that in tables 3.4 and 4.7. Although there are several significant results here, most of them are the same as for the analysis of overall performance (table 4.11). The notable difference is that, instead of showing a significant three-way interaction between group, stimulus duration and testing run, there is instead a significant four-way interaction between group, hemifield, stimulus duration and testing run. *Post hoc* analysis shows that this reflects differences in performance between the animals in the BI group and the others—it does not reflect any significant hemifield-specific impairment in the UNI animals.

The final analysis, given in table 4.13 is that of front-back errors (cf. tables 3.5 and 4.8). This shows an interaction between group, stimulus du-

source	df	SS	MS	F	p
group	2	0.0326	0.0166	2.00	0.217
error(group)	6	0.0500	0.00834		
duration	4	0.641	0.160	130	**<0.001**
duration*group	8	0.0901	0.0113	9.10	**<0.001**
error(duration)	24	0.0297	0.00124		
run	2	0.0538	0.0269	6.16	**0.014**
run*group	4	0.0101	0.00253	0.581	0.683
error(run)	12	0.0524	0.00436		
error-type	1	0.335	0.335	258	**<0.001**
error-type*group	2	0.00994	0.00497	3.83	0.085
error(error-type)	6	0.00780	0.00130		
duration*run	8	0.0634	0.00792	3.74	**0.002**
duration*run*group	16	0.0543	0.00339	1.60	0.105
error(duration*run)	48	0.102	0.00212		
duration*error-type	4	0.262	0.0654	52.7	**<0.001**
duration*error-type*group	8	0.0236	0.00295	2.37	**0.048**
error(duration*error-type)	24	0.0298	0.00124		
run*error-type	2	0.00452	0.00226	1.02	0.390
run*error-type*group	4	0.00135	0.000337	0.152	0.958
error(run*error-type)	12	0.0266	0.00222		
duration*run*error-type	8	0.00929	0.00116	1.10	0.381
duration*run*error-type*group	16	0.0133	0.000833	0.789	0.690
error(duration*run*error-type)	48	0.0508	0.00106		

Table 4.13: ANOVA table for comparison of the incidence of front-back errors, and errors of equal magnitude, between ferrets in Groups NO-OP, UNI and BI. Incidence is measured as the arcsine-corrected proportion of responses which are errors of the given type.

ration and error-type (front-back *vs.* control) which presumably reflects the large increase in front-back errors at short durations in the BI group animals. Indeed, *post hoc* comparisons show significant differences, on all three runs, between the number of front-back errors made at 40 ms by the animals in the BI group and those made by the animals from the other two groups. Importantly, there is no interaction here between group and testing run, in contrast to the analysis of performance given above. This shows, as does the *post hoc* analysis, that the incidence of front-back errors is elevated for the BI animals throughout all the testing runs, while the reduction in performance is only obvious in the first few runs.

4.4 Discussion

We have investigated the effect on an absolute sound localisation task of chronic A1 inactivation, both unilaterally and bilaterally, and in both highly-trained and untrained animals. Overall, our results suggest that bilateral implantation causes an impairment on our task when short-duration stimuli are used. In the case of animals with unilateral implants, those who have had previous training on this task show a mild impairment which is limited to the contralateral hemifield at short stimulus durations, while those with no previous training show little evidence of any impairment.

That our implants cause their effects in the manner intended—by inactivating cortex—is indicated by the evidence given in figures 4.1–4.4. The figures show that the implants release an effective amount of muscimol over a prolonged time period, and that the release is sufficient to cause a shutdown of intrinsic neural activity in cerebral cortex underlying the implant. They also show that drug-free Elvax had either very minor or no effects on cortical activity. We also implanted drug-free Elvax unilaterally in two ferrets (Group BLANK), with no obvious adverse effects on behaviour. However, subsequent testing of the effects of unilateral muscimol Elvax on inexperienced animals showed no impairment either, so it is possible, albeit unlikely, that this absence of a behavioural effect is a result of our testing regime.

Importantly, figure 4.3 also shows that the cortical blockade caused by muscimol Elvax does not extend far beyond the edges of the Elvax sheet. This means that, provided our implants were correctly placed over A1, inactivation should not have spread much into other cortical areas. Having said that, we erred on the side of larger rather than smaller Elvax implants, as we were mindful of results from lesion studies where only a small amount of A1 needed to be intact in order for the performance to be unimpaired[56]. It is thus possible that we partially affected secondary auditory cortical areas. It is also possible that, because the Elvax implants lay flat against the cortical surface, we did not affect parts of A1 which may have been within the suprasylvian sulcus.

The training and testing procedure for the pilot study (Group PILOT) was designed to resemble that used in previous studies of the effects of A1 lesions. That is, animals' localisation ability was tested both before and after inactivation of A1, the differences being that in earlier work the inactivation was by means of a lesion, and the testing took place only in the front hemifield. For the most part, our findings replicate those of earlier studies—unilateral A1 inactivation appears to cause a deficit in the localisation of short-duration stimuli in the hemifield contralateral to the affected side of cortex. There are novel aspects to our results, though: first, the impairment

in this study seems to be much more mild than in lesion studies, and second, there appears to be an increase in the number of front-back errors, errors which are not possible on the tasks in earlier studies.

There are several reasons which could explain the apparent mildness of the lesion effects in this case. One is the problem noted in chapter 3—that with speakers and targets placed at 30° intervals, we can only very coarsely sample the localisation abilities of our animals. An animal could be seriously impaired on an MAA task and yet remain perfectly able to distinguish speakers which are 30° apart. However, not all previous studies have used MAA tasks, and animals with unilateral A1 lesions have shown apparently severe impairments on a similar task to this one[56].

Another notable point, though not a problem as such, is that the use of muscimol Elvax is potentially much less damaging and much more specific than the use of aspiration lesions. Muscimol is a $GABA_A$ agonist, so its effects are simply to inhibit neuronal firing—this means that there should be no effect on underlying white matter tracts or fibres of passage, no damage to neurons whose axons terminate in A1, and reduced knock-on effects in areas which receive input from A1. In other words, the effect of muscimol Elvax should be targetted towards A1 itself, while some of the results of lesion studies may reflect side-effects in other areas. In particular, A1 lesions are known to cause retrograde degeneration in the thalamus, notably in the ventral division of the MGN[59]. As discussed in chapter 1, ventral MGN provides direct input to several auditory cortical areas besides A1, including areas such as AAF and PAF in the cat[109], so all of these areas are potentially affected by lesions.

On the other hand, as noted above and shown in figure 4.5, the disadvantage of this method of inactivation is that the implant cannot be placed precisely over A1, and if we have failed to inactivate some of A1 then it is possible that this could account for the limited impairment. To investigate this possibility we tested the PILOT animals with restricted-bandwidth stimuli: if the edges of A1 had not been covered by the Elvax, then this would leave the representation of the highest and lowest frequencies intact, and so filtering out these frequencies from our stimulus should reveal the underlying impairment. The results are somewhat equivocal, given both the absence of similar testing on normal animals, and the fact that narrowband stimuli are intrinsically more difficult to locate than broadband stimuli. The PILOT animals did perform worse than normal when stimuli were band-pass filtered, but there was no significant difference between performance in the left and right hemifields, and therefore no evidence that this effect is related to the muscimol Elvax implant.

The increased number of front-back errors is in some respects more sur-

prising than the mildness of the impairment. Front-back errors are believed to be caused when subjects are unable to disambiguate cues which are symmetric about the interaural axis, e.g. ITD cues. This disambiguation must involve cues which are not symmetric in this way, the most obvious being outer ear-related spectral cues. But if the effect of A1 inactivation was solely to damage the processing of a single cue type, then the effect on a task which only involved the front hemifield—where there is a large degree of redundancy between different cues—might be expected to cause only a small impairment, which is not the case. An alternative explanation arises from the fact that animals with bilateral A1 ablation were were unable to locate sound sources within a hemifield, but still able to distinguish sound sources on opposite sides of the midline[59]. The errors that we use as controls for front-back errors are either on the midline or on opposite side of the midline to the original sound source. Thus the disproportionate increase in front-back errors could reflect a problem in distinguishing sound sources within the same hemifield, rather than an inability to disambiguate symmetric cues. Whatever the cause, it is interesting to note that the elevated incidence of front-back errors seems to be longer-lasting—in terms of the number of testing runs—than the the impairment in overall performance, which at the very least suggests that the number of front-back errors provides a sensitive test of impaired localisation.

While the effect of unilateral A1 inactivation on highly trained animals is largely what might be expected, the same is not true for naïve animals (Group UNI). These animals show no evidence of an effect of Elvax implantation on overall performance, differences between hemifields which, if anything, suggest an improvement in the contralateral hemifield, and no increase in the overall number of front-back errors. The only hint of an effect is that performance on the first few trials, when the animals are still learning the task, seems to be consistently worse in the hemifield opposite the implant. This suggests that training before A1 inactivation causes A1 to acquire a role on this task, which is subsequently revealed when inactivation causes an impairment, whereas training after A1 inactivation may result in slower learning, but leads to some other brain area taking over A1's role. In this context it is important to point out that what the animals are learning can hardly be expected to be sound localisation *per se*—these are adult animals who are presumably capable of locating sound sources—but is likely to be the sensorimotor connections relevant to this task. Therefore it is this sensorimotor role which is being ascribed to A1. This is consistent with evidence discussed in chapter 1, which indicates that some tasks involve sensorimotor circuits which appear to be entirely subcortical, and which are not impaired by A1 lesions[76].

If some other brain area takes over from A1 in the case of the untrained,

unilaterally-inactivated animals, the obvious candidate for this other brain area is the A1 in the hemisphere opposite the implant. This possibility is supported by the findings that bilateral inactivation of A1 (Group BI), even in untrained animals, causes a significant impairment in localisation ability, and an accompanying increase in the number of front-back errors. That the two sides of A1 might form a functional unit is further suggested by the observation that humans with no corpus callosum are deficient at sound localisation[103], and the fact that callosal connections between the two A1 areas are much more common than the equivalent connections between primary visual cortices[25, 123, 122].

One aspect of the behaviour of both groups which show an impairment (Groups PILOT and BI) is that they appear to show some sort of improvement while the implants are in place. This could be caused by several things: it could be a decrease in the amount of muscimol release, a decrease in the effectiveness of muscimol on cortex, a case of some other brain area taking over the role of A1, or it could simply reflect that the animals are getting better at the task. The first two possibilities seem to be disproved by the data on muscimol release already mentioned—it seems to be the case that cortical blockade continues for at least 6 weeks, if not longer, while most of the improvement for the bilaterally-inactivated animals seems to take place within this time period. The effect of Elvax removal in the pilot study suggests that the fourth explanation is not the case—performance after muscimol Elvax removal was the same as performance before implantation, suggesting that no learning had taken place in the meantime. This leaves the suggestion that some other area is taking over the function of A1 while the muscimol Elvax is in place. Such recovery is in contrast to the effects of lesions, which have been reported to be permanent[56], possibly because the knock-on effects of A1 lesions damage the areas which would be capable of taking over from the lesioned area.

Part III

Sound localisation electrophysiology

Chapter 5

Electrophysiological methods

5.1 Use of animals

As with the behavioural experiments, all animals used were pigmented ferrets (*Mustela putorius*). Prior to recording, all had been housed in in rack-mounted cages and fed on a high-protein carnivore diet, with free access to water.

5.2 Surgery

The ferret was first anaesthetised with 2 $\mathrm{ml\,kg^{-1}}$ Saffan (9 $\mathrm{mg\,ml^{-1}}$ alphaxolone with 3 $\mathrm{mg\,ml^{-1}}$ alphadolone acetate), delivered i.m. The animal was then given 0.1 $\mathrm{ml\,kg^{-1}}$ of atropine sulphate (600 $\mu\mathrm{g\,ml^{-1}}$) s.c. to prevent pulmonary congestion, and 0.15 ml of Dopram (20 $\mathrm{mg\,ml^{-1}}$ doxapram hydrochloride) i.m. to counter the respiratory suppressant effect of the anaesthetic. An i.v. line was inserted in the radial vein, and 0.1 ml top-ups of Saffan were delivered as needed (approx. every 15–30 minutes, depending on the animal). Some animals were also given 1 $\mathrm{ml\,kg^{-1}}$ of Dexadreson (2 $\mathrm{mg\,ml^{-1}}$ dexamethasone, as sodium phosphate) i.v., in small doses over 15 minutes—this is to reduce brain oedema later in the experiment. A tracheotomy was performed, and the animal then placed in a stereotaxic frame using blunt ear bars. A rectal probe was inserted and used to monitor the animal's body temperature, which was maintained near 39°C using a heating blanket. Heart rate was also monitored, by measuring the ECG.

The skull was then exposed on both sides, attaching a head-holder to the side opposite recording. A window was drilled into the skull on the recording side, and the dura reflected. The exposed cortex was covered with either 2% agar or mineral oil, to reduce pulsations and prevent it from drying out. The

ferret was then transferred to gaseous anaesthesia, using 0.5–1% isoflurane in a 1:1 mix of N_2O:O_2 at a rate of approx. 0.7 litres min^{-1}.

The stereotaxic frame was then removed, and the animal transferred to the recording chamber. Further doses of Dopram and atropine were administered every 12 hours. The animal was also given i.v. fluids continually, using an infusion pump, to avoid dehyrdation.

Experiments were terminated by humanely killing the animal with an i.v. overdose of sodium pentobarbitone (Euthatal).

5.3 Apparatus

All experiments were carried out in a sound-proofed, anechoic chamber. During an experiment, the anesthetised animal would be placed on a small metal table in the centre of the chamber. The presentation of stimuli and recording of responses were both controlled from outside the chamber using a PC-compatible computer, the "Brainware" software suite and TDT System II digital signal processing hardware (Tucker-Davis Technologies, Gainesville, FL). The chamber also contained pipes through the walls to allow anaesthetic gases and i.v. fluids to be delivered from outside the chamber during experiments (see section 5.2).

Auditory stimuli were generated on computer as 80 kHz, 16-bit digital samples. These were sent to a D/A converter and the resulting analog signal was put through an anti-aliasing filter with nominal cutoff frequency of 30 kHz. The filtered stimuli were then amplified by a QUAD 240 power amplifier and sent to a KEF T27 loudspeaker within the chamber.

Inside the chamber, the loudspeaker was attached to a semicircular metal rail, 0.65 m in radius, which was in turn attached to the floor at one end and to the ceiling at the other end. The position of this semicircular rail was such that the centre of the semicircle was where the animal's head would lie during an experiment. The azimuth of the loudspeaker could be varied by a motor which rotated the rail about the points where it was attached to the floor and ceiling. The elevation of the loudspeaker could be moved by means of a second motor which allowed it to "crawl" up and down the rail, although in all the experiments described here the loudspeaker was positioned along the horizon (0° elevation). Both motors were controlled by computer from outside the chamber[4].

Neural responses were recorded using tungsten-in-glass micro-electrodes, typically of 10–20 μm tip length. The electrode or electrodes were held in an assembly which was fixed to the metal table on which the animal was placed. This assembly consisted of a manual micro-manipulator in which was held a

motorised and remotely-controlled microdrive. The micro-manipulator was used to set the gross position of the whole assembly, while the microdrive was used to move the electrode(s) back and forth in small increments. In experiments where only a single electrode was used at a time, the electrode was simply inserted into a small headstage, and this was attached to the microdrive. In experiments involving simultaneous recordings from several electrodes, up to four electrodes were held in the microdrive directly, but these were also connected by flexible control wires to a separate 4-way microdrive/headstage assembly (Alpha Omega Engineering, Israel). All electrodes could be moved at once by moving the main microdrive or the micro-manipulator, and electrodes could also be moved independently of one another by means of the flexible wires and the 4-way microdrive.

The signals picked up by each micro-electrode were amplified (approx. 10 000-fold) and band-pass filtered (typically 0.5–5 kHz) before being digitised (8 bits, 25 kHz sampling rate) using an A/D converter and analysed on computer. This process was synchronised to stimulus presentation, such that the signals from the micro-electrodes were digitised and recorded for the 1 s immediately following stimulus onset. The next stimulus was not presented until after the completion of this 1 s of continuous recording; typically, inter-stimulus intervals were between 1 s and 1.5 s.

5.4 Stimulus presentation

During recording, various stimulus sets were used to investigate the sensitivities of A1 neurons. The term 'unit' is used here to refer to the source of a series of electrophysiological recordings, where that series is believed to originate from a single neuron. The word 'neuron' is not used as it is impossible to be sure that these are not recordings from multiple neurons.

5.4.1 Noise detection threshold

An approximate noise detection threshold was determined for each A1 unit that was investigated. This was needed in order to set an appropriate intensity when measuring location tuning, because location tuning for near-threshold sounds is known to be very different from location tuning for more intense sounds. This is at least in part because sounds which are near-threshold in the ear nearest to the sound source are likely to be sub-threshold in the ear farthest from the sound source—in effect they are monaural.

Noise detection thresholds were determined by presenting a battery of 100 ms bursts (with 5 ms cosine-squared gating) of broadband noise at a

variety of intensities. This noise was generated on computer as white noise, with a 30 Hz–30 kHz spectrum, but no compensation was made for speaker characteristics, so the final stimulus spectrum would not have been flat. At least 20 stimuli were presented at each of the intensities used. When presenting these stimuli, the loudspeaker was positioned approximately along the "acoustic axis" on the side contralateral to the recording site: typically at an azimuth 80° from straight ahead. The choice of intensities and the order in which stimuli were presented were both under the experimenter's control and were decided *ad hoc*, but intensities varying in 10 dB steps were typically used, and by default the stimuli were presented in a pseudo-random order.

5.4.2 Frequency sensitivity

A very rough measure of frequency sensitivity was also made for each unit. The main purpose of this was to confirm the position of the cell within A1, and to determine whether ITDs or ILDs would be more likely to be the major binaural cue to sound source location, so precision was not of great importance.

Frequency sensitivity was simply measured by presenting several 100 ms tone bursts (with 5 ms cosine-squared onset and offset gating) at a variety of different frequencies and intensities. However, "intensity" in this case refers to the strength of the signal being sent to the loudspeaker—neither the characteristics of the loudspeaker nor the shape of the outer ear were taken into account in generating pure-tone stimuli, so the absolute intensity of the sound reaching the animal's eardrum was unknown and would be different for different frequencies. As with the determination of noise detection thresholds, the loudspeaker was positioned at an azimuth $\sim$80° from straight ahead, on the side contralateral to the recording site. 20 stimuli were presented at each of the frequencies used, and the stimuli were usually presented in a pseudo-random order, but this could be overridden by the experimenter.

5.4.3 Spatial sensitivity to noise

The main focus of these experiments was on each unit's sensitivity to the azimuth of a sound source. Recordings were made of the responses to noise presented from different positions along the horizon (0° elevation). In most recordings these positions were at 20° intervals in azimuth, from 160° on one side of the midline to 160° on the other (inclusive). However, in some cases technical problems meant it was not possible to go beyond $\pm$140° from the midline. Either 40 ms or 100 ms broadband noise bursts were used, with

5 ms cosine-squared onset and offset gating. As with the stimuli used in determining noise threshold (section 5.4.1), the noise was generated as flat in the range 30 Hz to 30 kHz, but no compensation was made for speaker characteristics. The sound intensity of the stimulus was always at least 25 dB greater than the noise detection threshold determined earlier, but was set at the experimenter's discretion.

It was not possible to randomise the order of stimulus presentation, as moving the speaker under remote control from one location to the next takes time, and doing this between each stimulus would have been impractical. Instead it was necessary to present stimuli in an order which was grouped by speaker position: several stimuli would be presented at one position, then the speaker would be moved by 20° and several more stimuli presented, etc. For experiments where multiple stimulus durations or intensities were used, the duration and intensity were kept constant throughout each "azimuth sweep", and multiple sweeps were made: then if the neuron became unresponsive or was otherwise "lost" during a later sweep it would at least be possible to analyse all the data from the earlier sweep(s).[1]

The non-random order of stimulus presentation was a problem, as stimulus position could then become confounded with other factors which change during the course of an azimuth sweep, e.g. habituation of the cell's response, changes in the relative position of electrode and neuron, or small changes in the depth of anaesthesia. Thus the neuron's response to a few appropriately-chosen stimuli was checked after the end of each azimuth sweep—the resulting data were not necessarily analysed, but were simply used to see whether the neuron's responsivity had changed during a sweep.

Where it was intended that the data from these experiments would be analysed using information theory, it was important to ensure that the number of stimulus presentations was sufficiently large to allow an accurate calculation of transmitted information. This is covered in detail in chapter 7, but is mentioned here as it is the reason for using 40 presentations of each stimulus—a larger number than is typically the case in electrophysiological studies.[2]

[1] Neurons were frequently lost during recording. Small brain pulsations cause continual movement of the neuron with respect to the electrode, and this often led either to the neuron moving out of recording range or to the neuron becoming impaled on the electrode and dying.

[2] Although in the first pilot study only 20 repeats of each stimulus were used.

5.5 Response analysis

In most cases it is not possible to simply read out the activity of a single neuron from the signal picked up by an electrode. First, not everything recorded by microelectrodes is neuronal activity: there is also a large amount of noise. Second, a given electrode will often pick up signals from more than one neuron at a time. These two problems are addressed in separate stages of analysis.

In order to separate out neural spikes from background noise, a simple test was used. Where the "raw" voltage signal showed a biphasic response—peak followed by trough or *vice versa*—crossing experimenter-determined thresholds in both phases, and where the two phases were within 1 ms of one another, the digitised record was stored as a potential neural spike.

All potential neural spikes were then analysed further, a process which could take place either during or after the actual recording. The "Brainware" software suite measures various aspects of each spike waveform, e.g. peak voltage, trough voltage, peak-to-peak depth, time between peak and trough. If these measurements are plotted against one another for each spike, forming a series of scatterplots, then clusters of spikes with similar waveforms can be picked out. This process is known as "cluster cutting" (the term "cluster" in this context referring to a cluster of spikes, not a cluster of neurons). Depending on how similar the spikes in each cluster appear to be, these can be classified as either "single-unit" or "multi-unit" clusters, where a single-unit cluster is taken to represent the activity of a single neuron.

Much of the detailed cluster cutting was carried out well after the actual experiment took place. This meant that the estimate of noise detection threshold which was made during the experiment was sometimes revealed to be inaccurate by later analysis.

Chapter 6

Simultaneous single-cell recordings from ferret A1

6.1 Introduction

Despite the popularity of electrode recording in A1, and the quantity of data which it has produced, it has proved difficult to draw any conclusions about the function of A1 from electrode recording studies alone.

First, A1 neurons are sensitive to many auditory stimulus parameters, including the frequency and intensity of pure tones, the modulation frequency of AM and FM tones, and the location of noise sources[21]. It does not seem to be the case that any one parameter stands out in terms of A1 sensitivity, and, in any case, it is difficult to know how to compare sensitivity to different parameters. It is true that the frequency tuning of a neuron[1] is related to its position within A1, with neurons sensitive to the lowest frequencies along one edge and those tuned to the highest frequencies along the opposite edge[111, 63]. However, this organisation offers little insight into the function of A1, as most aspects of hearing involve the processing of stimulus frequency in some way. Similarly, while other aspects of neuronal sensitivity may be distributed in a patch-like manner across A1[109], these include the bandwidth of frequency tuning[45, 133] and various features of the intensity-response curve[46], none of which suggest a particular processing role.

There is also little that qualitatively distinguishes the electrophysiology of A1 from that of areas either before it or after it in the auditory processing chain. It is only possible to distinguish the properties of A1 neu-

[1] A neuron's "characteristic frequency" (CF) is the most oft-used measure of frequency tuning. It refers to the frequency of a pure-tone stimulus for which response threshold is lowest.

rons from other neurons in quantitative terms: A1 contains *more* cells with non-monotonic rate-level functions than subcortical areas, cells in A1 are *on average* less sensitive to very simple temporal modulations than those in subcortical areas[32, 115, 136], etc. Similarly, the "tonotopic" organisation of A1 into strips of similar CF is found in almost all subcortical auditory structures, as well as some secondary auditory cortical areas.

The biggest hint as to the role of A1 in auditory processing comes not from electrode recording, but from lesion studies, where the most notable effect of A1 damage is a deficit in sound localisation[76]. However, there is little electrophysiological evidence to support a claim that A1 is primarily involved in sound localisation *per se*. In fact, A1 neurons are tuned in a rather broad and level-dependent manner to sound source location, which might be taken to indicate the opposite[49, 105]. Furthermore, there is evidence that the sensitivity of A1 neurons to sound location is mainly a result of their frequency tuning[116], suggesting that these neurons do not have any special properties which predispose them to the processing of auditory space.

There are, though, good reasons why electrophysiological studies might have underestimated the sensitivity of A1 neurons to sound source location (or, indeed, to any other stimulus parameter). One is that it may be a mistake to simply count the number of spikes evoked by the presentation of a particular stimulus, as most studies do. There is some evidence to suggest that neurons in general may encode aspects of a stimulus in the timing as well as the number of spikes[114]. Another is that sound source location may be better represented by the firing of groups of neurons than by a single neuron.

While these possibilities are enticing, they involve measuring how much information a neuron (or group of neurons) transmits about the location of a sound source. Such measurements are not entirely straightforward, as they require either a mathematical analysis using information theory[88, 87] or a related approach such as the use of artificial neural networks[81, 82]. Thus the work described here is mainly aimed at providing data for the analysis of information described elsewhere (chapter 8). In order to acquire the necessary data for this, we intend to make recordings from more than one neuron at a time, by simultaneously using up to three electrodes. Multi-electrode recording of this kind is interesting in itself, as it is not yet a commonplace technique.

6.2 Methods

6.2.1 Numerical conventions

Throughout this chapter, azimuth is expressed in degrees, from $-180°$ to $+180°$, with $0°$ being straight ahead. The plus and minus signs relate the azimuth to the side of the brain from which recordings are being made—positive azimuths are on the same side of the animal's midline as the side of recording, while negative azimuths are on the other side of space. Experimental measurements are expressed either to the precision of measurement, or to 3 significant figures, whichever is lower. Other numeric quantities are treated similarly.

6.2.2 Animals

A total of six ferrets were used in this study, of which F9805 formed a pilot experiment and F9929, F9930, F9903, F9940 and F9941 formed the main experimental group. With the exception of F9903, which had been briefly involved in a non-invasive behavioural study of binaural hearing, none of the animals had been used for any experiments prior to those detailed below. All animals were female except for F9929.

6.2.3 Surgery

The surgery followed the procedure described in section 5.2. Briefly, the ferret was initially anaesthetised with i.m. Saffan in order to insert an i.v. line and a tracheal cannula. Thereafter it was kept anaesthetised with a gaseous mixture of isoflurane, nitrous oxide and oxygen for the remainder of the experiment, although it was not artificially ventilated. Fluids were provided via the i.v. line. A craniotomy was performed on the left side of the animal's head, making an attempt to maintain the pre-surgical positions of the pinnae. The dura was reflected and the exposed cortex was covered with agar or mineral oil before transferring the animal to the recording chamber.

6.2.4 Apparatus

The apparatus is detailed in section 5.3, and consisted of a sound-proofed, anechoic chamber containing a small table and a single loudspeaker attached to a moveable semi-circular hoop. All stimuli were presented via the loudspeaker, and were generated by computer. The position of the loudspeaker was adjustable using a series of motors, and was also under computer control.

The micro-electrodes used were tungsten-in-glass. In the pilot experiment (F9805), only a single electrode was used. This was placed in a combined electrode-holder/headstage pre-amplifier, which was held by a microdrive. The microdrive was in turn held by a micro-manipulator, which was fastened to the table in the centre of the chamber. In the other experiments, up to three electrodes were used at once. These were held in an electrode holder, which was attached to a microdrive as before. The electrodes could also be moved independently by using a series of flexible push-rods (one per electrode), each driven by a further microdrive. A separate headstage pre-amplifier was then connected to each push-rod. The microdrives were all controlled from outside the chamber, but other positioning was carried out manually.

6.2.5 Recording

Having positioned the electrode holder over the craniotomy by hand, the door of the anechoic chamber was then closed and the electrode advanced into the cortex using the remotely-controlled microdrive and/or push-rods. While this was being done, auditory stimuli were presented at regular intervals; this way, detection of a response was used to determine when the electrode was within recording distance of a neuron. The stimulus used in this search for a recording site was typically a broadband noise burst of 60–70 dB SPL intensity,[2] positioned along the horizon (0° elevation) between −70° and −80° azimuth.[3]

Once a neural response was detected to our search stimulus, an approximate response threshold to noise was determined (as detailed in section 5.4.1). This was necessary in order to determine what stimulus intensity was appropriate for use in investigating location tuning. An approximate measure of frequency tuning was also determined (section 5.4.2), although frequency tuning in many cases was surprisingly neither sharp nor consistent, and so this was often difficult.

The main aim of recording was to determine the response to noise bursts emanating from different positions along the horizon (section 5.4.3). The positions used were at 20° intervals between 160° and −160° azimuth, although for technical reasons some experiments only used positions between 140° and −140°. Because remote-controlled movement of the loudspeaker was slow it would have been impractical to randomise the order of stimulus presenta-

[2]Pure tones were initially used as a search stimulus, but the surprisingly poor response to pure tones led us to switch to noise.

[3]This position should be at or near the acoustic axis of the ear contralateral to the recording site, although the actual position of the acoustic axis is frequency-dependent[66].

tion. As a result, stimuli were presented in azimuth order, starting at one end of the range of azimuths and proceeding in 20° steps to the other end. A number of stimuli were presented at each position: 20 stimuli per position in the pilot experiment (F9805) and 40 stimuli per position thereafter.

Where possible, location tuning was recorded in response to stimuli of differing intensities and/or durations, although the entire range of azimuths was probed at each intensity or duration before moving onto the next. The stimulus intensity used for a particular "azimuth sweep" was chosen during the experiment, but was at least 25 dB above the noise threshold determined earlier, while stimulus duration was either 40 ms or 100 ms.[4]

Having recorded as much useful data as possible at one electrode position, an attempt was made to find other recording positions in the same track (i.e. with the micro-manipulator settings the same), although these had to be at least 100 μm from the original recording position, and from each other, to ensure that multiple recordings were not being made from the same neurons. Between 1–3 such positions were usually found in each track. Once a particular track had been exhausted, the micro-manipulator was repositioned (manually) and the procedure was repeated.

6.2.6 Analysis

6.2.6a Spike discrimination

In order to separate out neuronal signals from noise, and to distinguish signals derived from different neurons, a certain amount of analysis of the "raw" voltage waveforms was needed before it was possible to look at the data as trains of spikes. This involved two stages, described in greater detail in section 5.5: first, the rejection of signals that did not fit the profile of a spike, and, second, the grouping together of spikes that had a similar waveform. The former process was carried out during recording, while the latter process, known as "cluster cutting" was carried out both during and after recording. The clusters of spikes that resulted were described as either "single-unit" or "multi-unit" clusters: single-unit clusters are those which are very well-defined, while multi-unit clusters have the appearance of containing several smaller, badly-defined clusters. The spikes in a single-unit cluster are generally considered to have originated from a single neuron.

[4]The values of 40 ms and 100 ms were chosen to match the durations used in the behavioural experiments detailed in chapters 2–4.

6.2.6b Measures of response

The main measures of cluster response used throughout this study were first-spike latency and evoked spike count. First-spike latency is, very simply, the time elapsed between stimulus onset and the occurrence of the first spike in the cluster's response. Evoked spike count is more complicated because it is possible to make a correction for each cluster's spontaneous activity and background noise—something which is not possible for first-spike latency.

In deriving a value for evoked spike count, two different time "windows" were defined for each cluster of spikes: one to measure the response to the stimulus and one to measure the spontaneous/background activity. Both of these windows were defined in relation to the time of stimulus onset. The boundaries of the response window were set for each cluster separately, based on a visual examination of the cluster's PSTH[5], while the spontaneous/background window was in all cases the period between 500 ms and 1000 ms after stimulus onset.

The number of spikes occurring in each of these windows was then counted for every stimulus presentation, and for every cluster of spikes. The raw spike count was the number of spikes occurring within the response window. The spontaneous/background spike count was the number of spikes from the same cluster, in response to the same stimulus presentation, occurring in the spontaneous/background window. Because the two windows were not the same size, the spontaneous/background count could not sensibly be subtracted directly from the raw spike count. Instead, the factor used to correct the raw spike count was the spontaneous/background *rate*, multiplied by the duration of the response window. It is important to note that this correction factor was not an integer, and so the final value of evoked spike count was also not an integer. Another notable point is that while the contents of each window, i.e. the spike counts, were calculated separately for each stimulus presentation, the position of the response window was defined only once for each cluster of spikes.

6.3 Results

6.3.1 Unit selection

In this study we were only interested in well-defined, single-unit clusters, so clusters were cut to achieve the best possible degree of isolation, rather than

[5]PSTH is an abbreviation for post-stimulus time histogram, a graph showing how spikes are distributed over the time following stimulus presentation.

to maximize the spike count. Several other criteria were used to select the clusters which were considered for further analysis.

- It had to be possible to determine each cluster's response threshold to noise, at least to within 5–10 dB. This was necessary as we intended to ignore near-threshold stimuli (<25 dB suprathreshold) in our measurements of sensitivity to stimulus position.

- It had to be possible to determine the typical latency of the cluster's response, in order to set the "response windows" which are used to calculate evoked spike count. This required that the cluster's PSTH, from which the response windows were set, showed an obvious response.

- The measurements of azimuth sensitivity had to be free of any evidence of adaptation or habituation. This was checked by comparing the cluster's response to a given stimulus with its response to a repeat of that stimulus made at a later point during recording.

Having done this, each cluster which remained was considered to be a single unit for further analysis. The total number of such units was 84, of which 8 were recorded in ferret F9805, 8 in F9929, 17 in F9930, 14 in F9903, 13 in F9940 and 24 in F9941. The number of simultaneously-recorded "pairs" of units was 25: 1 in F9805, 1 in F9929, 7 in F9930, 6 in F9903, 2 in F9940 and 7 in F9941. A "pair" in this context is a unique pairing of two different units where the responses of both were recorded at the same time. Thus where three different units are recorded at the same time this counts as three different pairs—the first unit paired with the second, the second with the third, and the first with the third. Of the 25 pairs, in 4 cases the two units were actually recorded on the same electrode, and distinguished from one another by cluster cutting, while the remaining 21 pairs were composed of units recorded on different electrodes. In no case were more than three units recorded simultaneously, and no more than two units were ever recorded simultaneously from the same electrode.

6.3.2 Unit location

Figure 6.1 shows the positions of recording sites in animals where digital photographs were taken during recording. Positions were determined by using an operating microscope to visualise the electrode tip with respect to pattern of blood vessels on the cortical surface. In multiple-electrode experiments, the distance between the tips of different electrodes was approximately 0.5–1.0 mm—some effort was made to ensure that distance was no less than this,

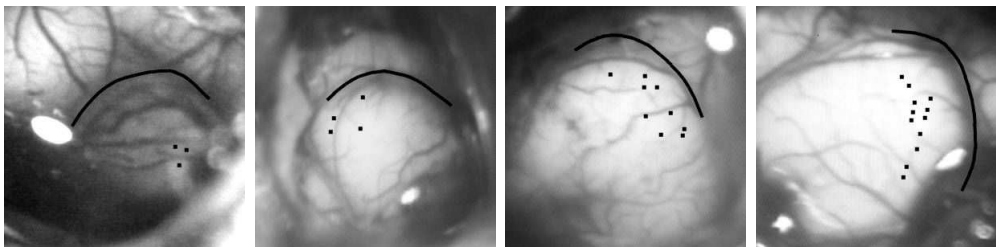

Figure 6.1: Recording sites in, from left to right, F9805, F9903, F9940 and F9941. Square markers show where the electrodes were seen to touch the cortical surface; only sites which yielded recordings are shown. The curved line in each photograph indicates the suprasylvian sulcus. Medial positions are to the top of each photo, and rostral positions to the left, except in the case of F9941 where medial positions are to the top right corner and rostral positions to the top left.

because of the damage that would occur to the electrode tips if they were to accidentally contact one another.

Although it is difficult to confirm this, due to the problems with poor pure-tone responses, all electrode positions appear to be well within A1, and in the area normally characterised by mid to high CFs (8 kHz and above). This is partly based on previously-published single-cell recording studies[101, 124], and partly based on optical imaging studies from our own laboratory[134]. An example of the latter is shown in figure 6.2, which can be directly compared to figure 6.1.

The depth of these various recording sites was not determined anatomically. However, it was often possible to identify the point at which the electrode broke the cortical surface, due to the characteristic change in the signal being recorded. Thus for most sites a rough estimation of recording depth could be made using readings taken from the microdrive controller. In the majority of these cases these readings suggest a depth of less than 0.6 mm, and for at least two-thirds of sites it was less than 0.8 mm. This suggests that most of our recordings were in cortical layers II/III and IV. However, the accuracy of these depth measurements is quite poor, e.g. because electrode penetrations may have been at an oblique angle to the cortical surface.

6.3.3 Unit characteristics

6.3.3a Noise threshold

Unit thresholds to noise, which are a measure of sensitivity, varied widely from $\sim$5 dB SPL to $\sim$55 dB SPL (threshold estimates are only accurate to

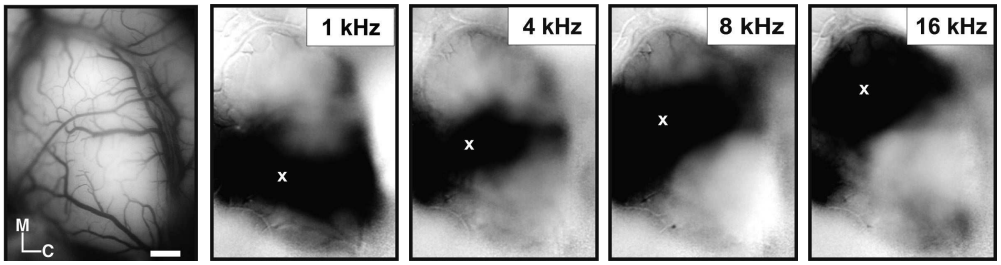

Figure 6.2: The typical extent of ferret A1, as defined by tonotopicity and assessed using intrinsic optical imaging[134]. The leftmost image is a photograph of A1, and the other images reflect cortical activity, in the same animal, in response to narrowband noise with the indicated centre frequency. The white scale bar indicates a length of 1 mm. White 'X's represent centroids of activity. This diagram is taken from unpublished work by Drs. Mrsic-Flogel, Versnel and King.

$\pm$10 dB). The distribution of unit thresholds is shown in figure 6.3. The mean threshold is $\sim$30 dB SPL.

Informally, we found that the responsiveness of A1 seemed to fluctuate a great deal with changes in anaesthetic depth. In deep anaesthesia, very little auditory response could be discerned. At slightly less anaesthesia, responses were easy to elicit, whereas at even less anaesthesia the background level of neuronal activity was such that it could be difficult to isolate a single neuron. All of these stages took place while the animal was sufficiently unconscious as to be areflexive.

6.3.3b Frequency sensitivity

We found that the majority of units did not respond well to pure-tone stimuli, with the responses being weak and labile. Only 30/84 (36%) units showed a clear preference for tones within a particular frequency range. This is quite unusual—one of A1's main features in barbiturate-anaesthesia preparations is the reliable response to pure tones, to the extent that tones are usually used as a search stimulus. Figure 6.4 shows the frequency tuning for two example units, with response measured in terms of evoked spike count (see section 6.2). One of the units has an obvious preference for tones with frequencies near 20 kHz, while the other shows no clear pattern of frequency-sensitivity. Importantly, these two units were both recorded in the same electrode penetration, so it is not simply the case that some positions were responsive to tones and others were not.

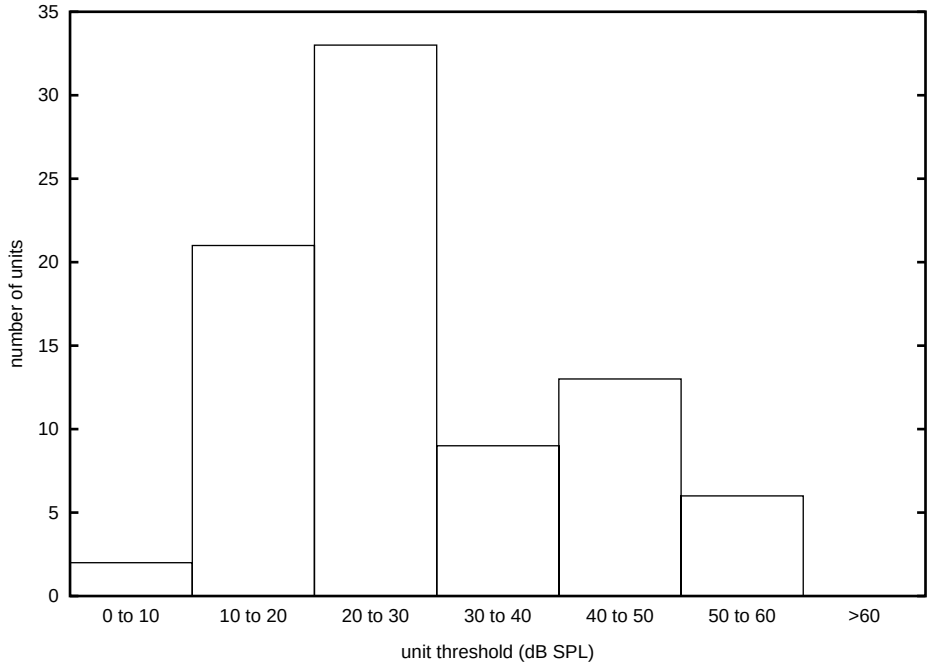

Figure 6.3: The distribution of single-unit noise thresholds. Threshold estimates for each unit are only accurate to within 10 dB.

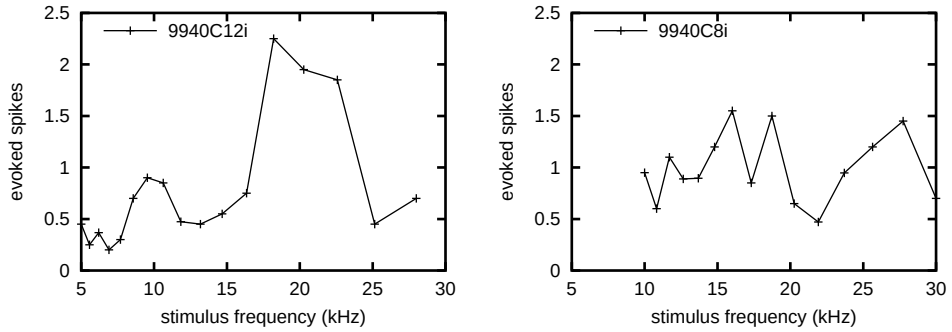

Figure 6.4: The responses of two example units to pure tones of differing frequency. Response is measured as the number of evoked spikes per stimulus presentation, and each point represents the average of 20 stimulus presentations.

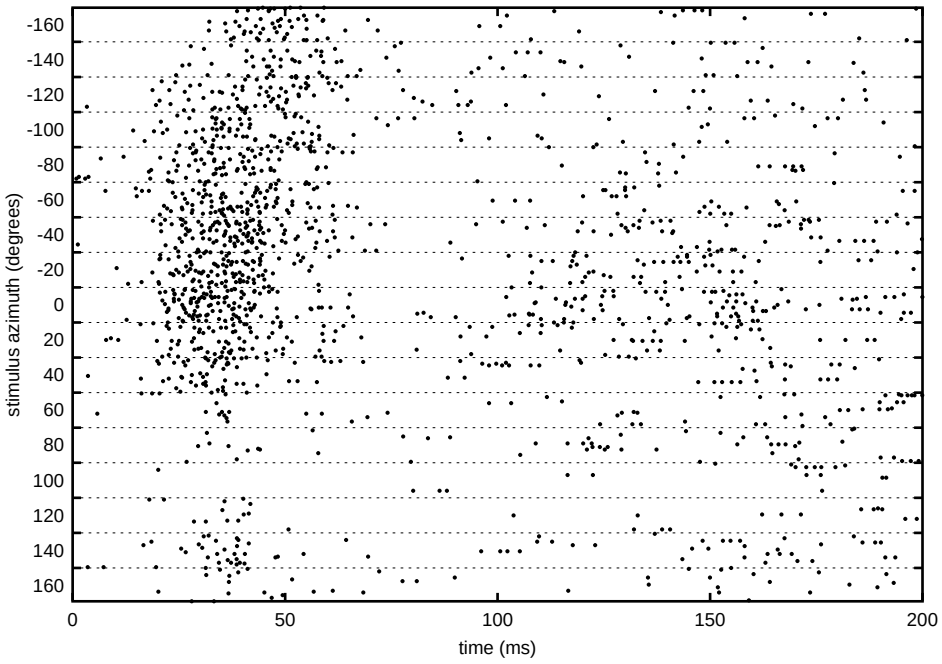

Figure 6.5: Responses of a typical unit (9940C10) to noise bursts from various locations. Each dot represents a single spike. Position along the abscissa represents the latency of the spike, and dots on the same horizontal line are spikes occurring in response to the same stimulus presentation. The azimuth associated with each presentation is given on the ordinate, with negative values indicating the hemifield contralateral to the side of recording.

6.3.4 Sensitivity to sound source location

Figure 6.5 shows the response of a typical unit (Unit C10 in animal F9940) to stimuli of different azimuths. As can be seen, the unit responds strongly to stimuli in the contralateral front quadrant, responds more weakly, and with longer latency, to stimuli in the contralateral rear quadrant, and responds very little to stimuli in the ipsilateral hemifield. The same data are shown as a series of post-stimulus time histograms in figure 6.6 (although note that this shows the full 1 s of recording while figure 6.5 only shows the first 200 ms).

PSTHs derived from two other recordings are shown as examples in figures 6.7 and 6.8. The quality of the recording in figure 6.7 is worse than the others shown here, but it is probably more representative of the typical level of background noise. The unit shown (Unit C61 in F9941) is also quite typical in that it responds to stimuli coming from almost all directions, but responds particularly strongly to stimulus azimuths in the front contralateral

quadrant. The response profile shown in figure 6.8 is less common in that it responds primarily to ipsilateral stimuli. This unit (Unit C23 in F9940) also shows a very strong secondary response peak in the PSTH—this is quite common, and such a peak exists for many of the units we recorded, although usually to a lesser extent. In fact, a small secondary PSTH peak can be discerned in some panels of figure 6.6.

A traditional analysis of the azimuth sweep shown in figures 6.5 and 6.6 involves calculating the evoked spike count for each stimulus presentation, and plotting the mean spike count against stimulus azimuth to give a "receptive field". As described in section 6.2, we determined the evoked spike count by manually defining a time window for the response, counting the number of spikes occurring within that window, and then subtracting an estimate of the spontaneous firing rate. The response window was determined separately for each unit, on the basis of a PSTH. It should be noted that the PSTH used for this was compiled from all the results in an azimuth sweep, rather than being a stimulus-specific PSTH like those shown in figure 6.6. As an example of the full process, the response window for the unit shown in figures 6.5 and 6.6 was 10–62 ms after stimulus onset, and the resulting receptive field is shown in figure 6.9 (left). For comparison, figure 6.9 (right) shows the reciprocal of first-spike latency for the same unit, and it is clear that high spike counts are associated with short latencies.

It is important to show that the magnitude and timing of responses seen in figures 6.5–6.9 are typical, as both magnitude and timing are analysed further in chapter 8. Figure 6.10 displays the response windows for all of the units we recorded from, and, as can be seen, most windows were set with their start 10–30 ms after stimulus onset, and their end <100 ms later. A histogram of first-spike latency, given in figure 6.11, confirms that initial responses most frequently commence 10–20 ms after stimulus onset, while the number of spikes evoked by a stimulus varies mainly between 1–5 spikes, as shown in the same figure. Given the differences in response window size shown in figure 6.10, it could be argued that a measure of spike *rate* would be more appropriate than a measure of spike *count*. However, for our purposes the absolute magnitude of the response was less important than the way in which response differed for different stimuli. As a result, using spike rate instead of spike count would only have had an effect if the response window had been calculated separately for each stimulus. This was not done, because it is difficult to determine a response window in cases where there is little neural response.

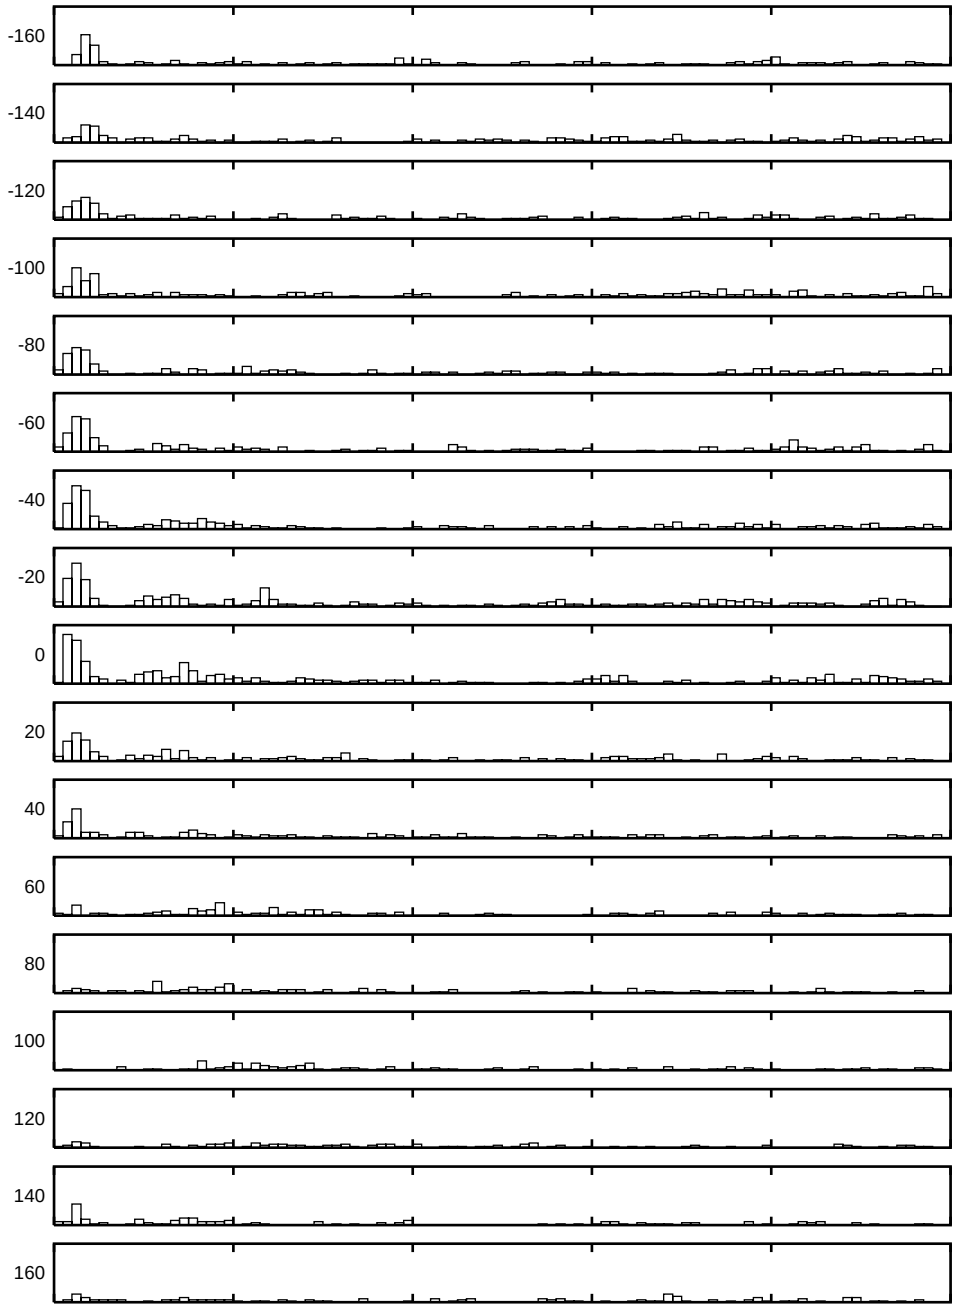

Figure 6.6: PSTHs compiled from the same recordings as in figure 6.5. Each PSTH shows data from forty presentations of the same stimulus. Time after stimulus onset is given on the abscissa (0–1000 ms), and total number of spikes is given on the ordinate (0–50 spikes). Histogram bins are 10 ms in width, and labels to the left of each plot give the stimulus azimuth in degrees.

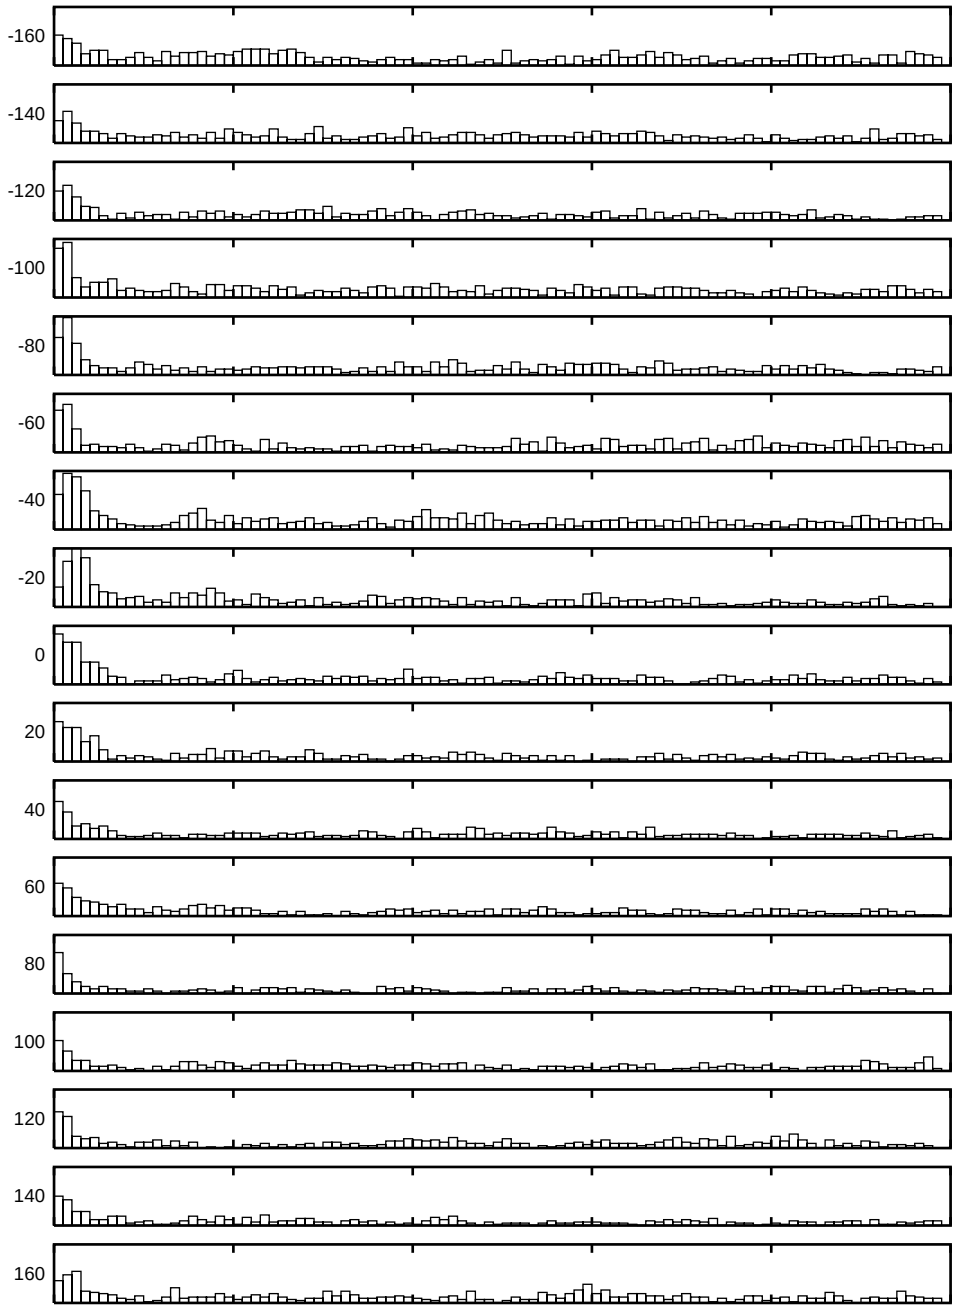

Figure 6.7: PSTHs representing the response of Unit C61, in animal F9941, to stimuli at different positions. The axes and other graphical conventions are the same as in figure 6.6.

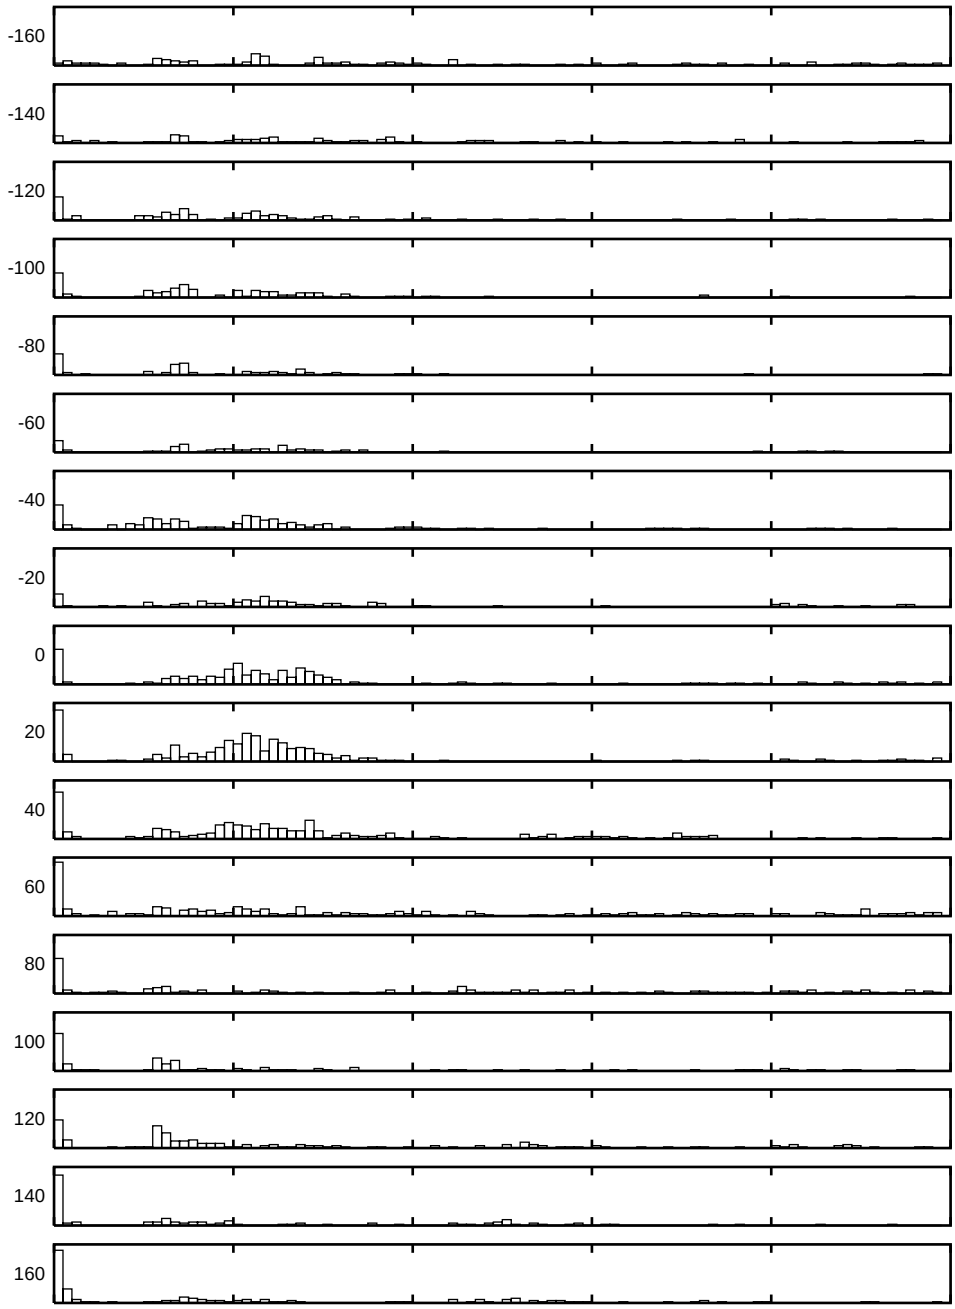

Figure 6.8: PSTHs representing the response of Unit C23, in animal F9940, to stimuli at different positions. The axes and other graphical conventions are the same as in figure 6.6.

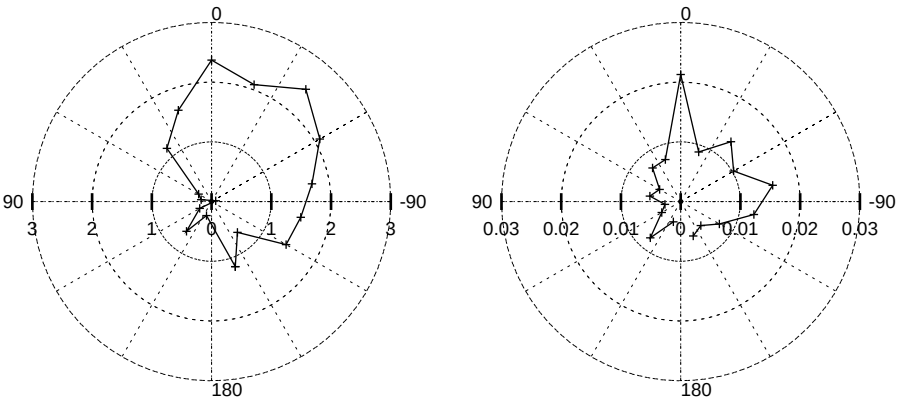

Figure 6.9: The spatial receptive field for the unit in figure 6.5. The angle of each point represents stimulus azimuth, and the radius of each point represents response magnitude—either evoked spike count (left) or the reciprocal of first-spike latency (right). Each point is an average of forty presentations.

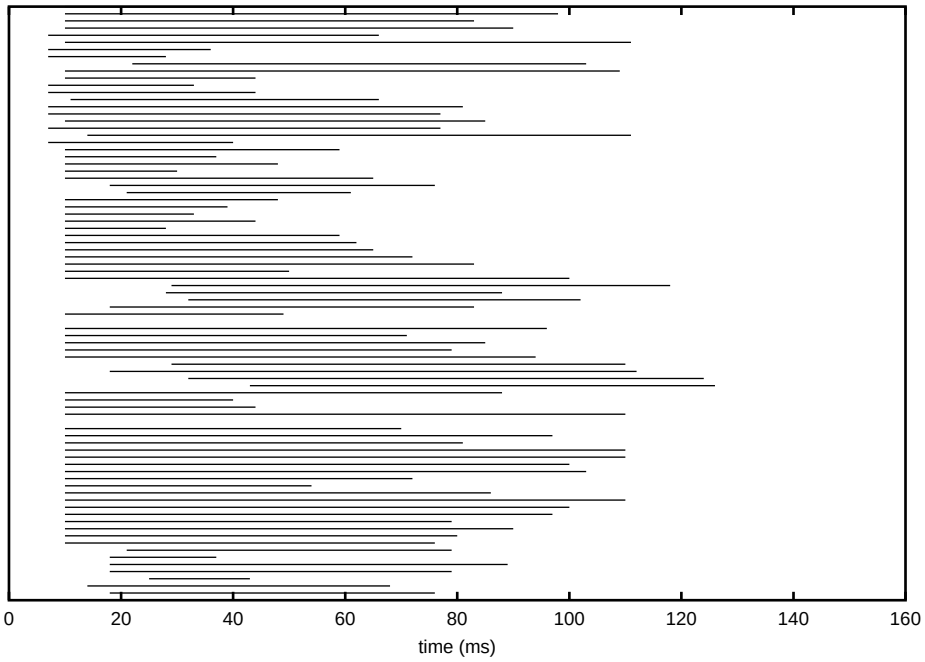

Figure 6.10: The response windows used in calculating evoked spike count, for every unit recorded. Each horizontal line represents a different unit, with its position and extent along the abscissa indicating the timing and duration of the window with respect to stimulus onset.

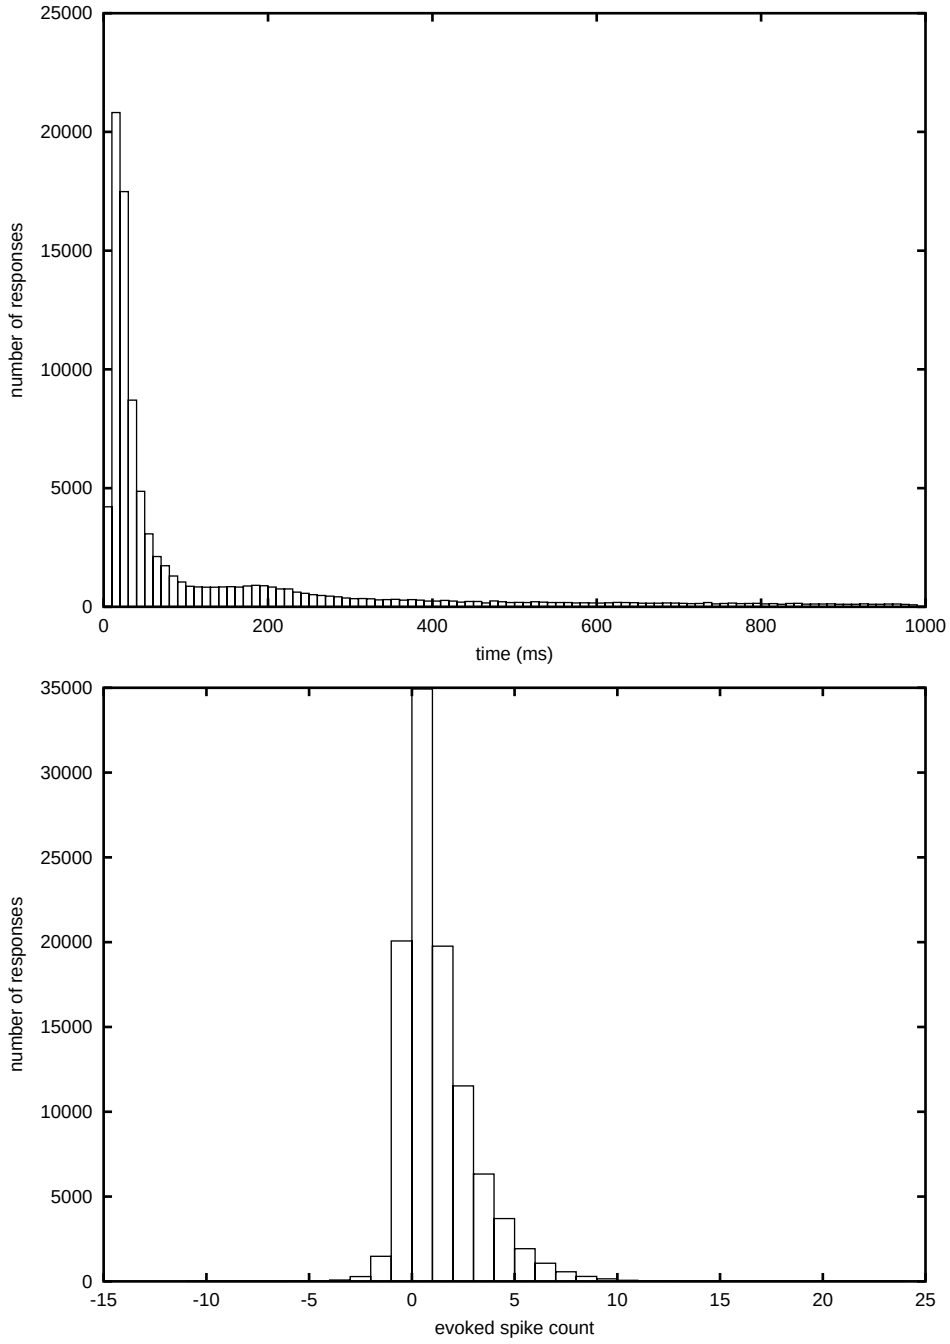

Figure 6.11: Histograms showing the first-spike latency (top) and number of evoked spikes (bottom) in responses to noise stimuli. Bins are 10 ms or 1 spike in width. All responses from all units are counted, with the exception of responses collected during threshold determination or responses to pure-tone stimuli.

6.3.4a Quantifying location sensitivity

To show how receptive field properties are distributed through our population of units, a more quantitative assessment of location tuning is needed. So, for each unit, and for each combination of stimulus intensity and stimulus duration, two measures of location tuning were used, both of which were derived from evoked spike counts. The first of these is a measure of the direction to which the unit is mainly tuned, and the second is a measure of the breadth or bandwidth of that tuning.[6]

The direction to which each unit was tuned was quantified by considering the unit's receptive field as a series of vectors. The response at each stimulus position was considered to be a vector, with direction equal to the direction of the stimulus and length equal to the average evoked spike count at that position. These vectors were then summed, and the direction of the resulting sum was used as a measure of the unit's directional sensitivity. A slight complication results from the fact that we were unable to place our loudspeaker directly behind the animal—this means that we have not evenly sampled the full 360° range of azimuths. To avoid this sampling bias, which would affect the vector sum, we used the average of the responses at −160° and 160° as an estimate of the response directly behind the animal. This estimate was then used, along with all the other response vectors, to calculate the vector sum for the unit. Because of similar problems with the unevenness of sampling, we did not attempt to calculate a vector sum for units which had only been tested with locations from −140° to 140°.

The breadth of a unit's tuning was more simply calculated as the total angle over which the unit's response was greater than 50% of its peak response. Where the responses at adjacent positions were either side of the 50% mark, linear interpolation was used to decide how much of the angle between the two positions should count as being greater than 50%. Where stimuli in two non-adjacent areas produce a greater than half-maximum response, both areas are included in the calculation of bandwidth. Because stimuli could not be positioned directly behind the animal, the maximum measurable bandwidth is the reflex angle between −160° and 160° i.e. 320°.

Sound intensity is known to affect neurons' sensitivities to sound location, and so, because we did not use a single "standard" stimulus intensity in our experiments, it is difficult to directly compare values of either vector direction or half-height bandwidth between units. However, the vector direction and

[6]Both of these measures are appropriate when applied to units with a single discrete responsive area within their receptive field. Not all A1 neurons respond like this and, strictly speaking, terms like "bandwidth" should not be applied to neurons with other patterns of responsiveness. However, we follow common usage in doing so.

half-height bandwidth for each unit are presented as a function of stimulus intensity in figure 6.12; the stimulus duration for all the data presented in these figures is 100 ms. As can be seen, the majority of units show values of vector direction which are in the front contralateral quadrant, and almost all are in the front hemifield. With respect to the breadth of spatial tuning, this is quite variable between units, but for many units is of the order of an entire hemifield (180°).

6.3.4b Effect of stimulus level

Some of the effects caused by changes in stimulus level are already visible in figure 6.12: in particular, the majority of units tested at more than one level seem to show an increase in half-height bandwidth at higher sound levels. Because units were typically tested at stimulus levels which were 20 dB apart, it is possible to group together the effects of a 20 dB increase in level on both vector direction and bandwidth. These data are shown in figure 6.13. What can be seen from this figure is that increases in stimulus level are apparently associated with an ipsilateral shift in vector direction, and a broadening of unit bandwidth. The mean change in vector direction is a 21.4° shift towards the ipsilateral side ($n = 31$), and the mean increase in bandwidth is 29.0° ($n = 39$). However, while the increase in bandwidth is statistically significant ($t = 2.15$, giving $p = 0.0379$ with a 2-tailed test and 38 d.f.), the shift in vector direction is not ($t = 1.44$, giving $p = 0.159$ with a 2-tailed test and 30 d.f.).

6.3.4c Effect of stimulus duration

Several units ($n = 30$) were tested with 100 ms stimuli and 40 ms stimuli of the same stimulus intensity. Figure 6.14 shows the relationship between the values of vector direction and bandwidth derived at 40 ms and at 100 ms from each of these units. The difference between the two different values of vector direction is that the value at 40 ms is slightly more contralateral than that at 100 ms (mean difference = 9.33°), but this is not significant ($t = 0.837$, giving $p = 0.411$ with a 2-tailed test and 24 d.f.). However, switching from a 100 ms to a 40 ms test does result in a decrease in bandwidth (mean change = -27.3°) which is significant ($t = 2.79$, giving $p = 0.00929$ with a 2-tailed test and 29 d.f.).

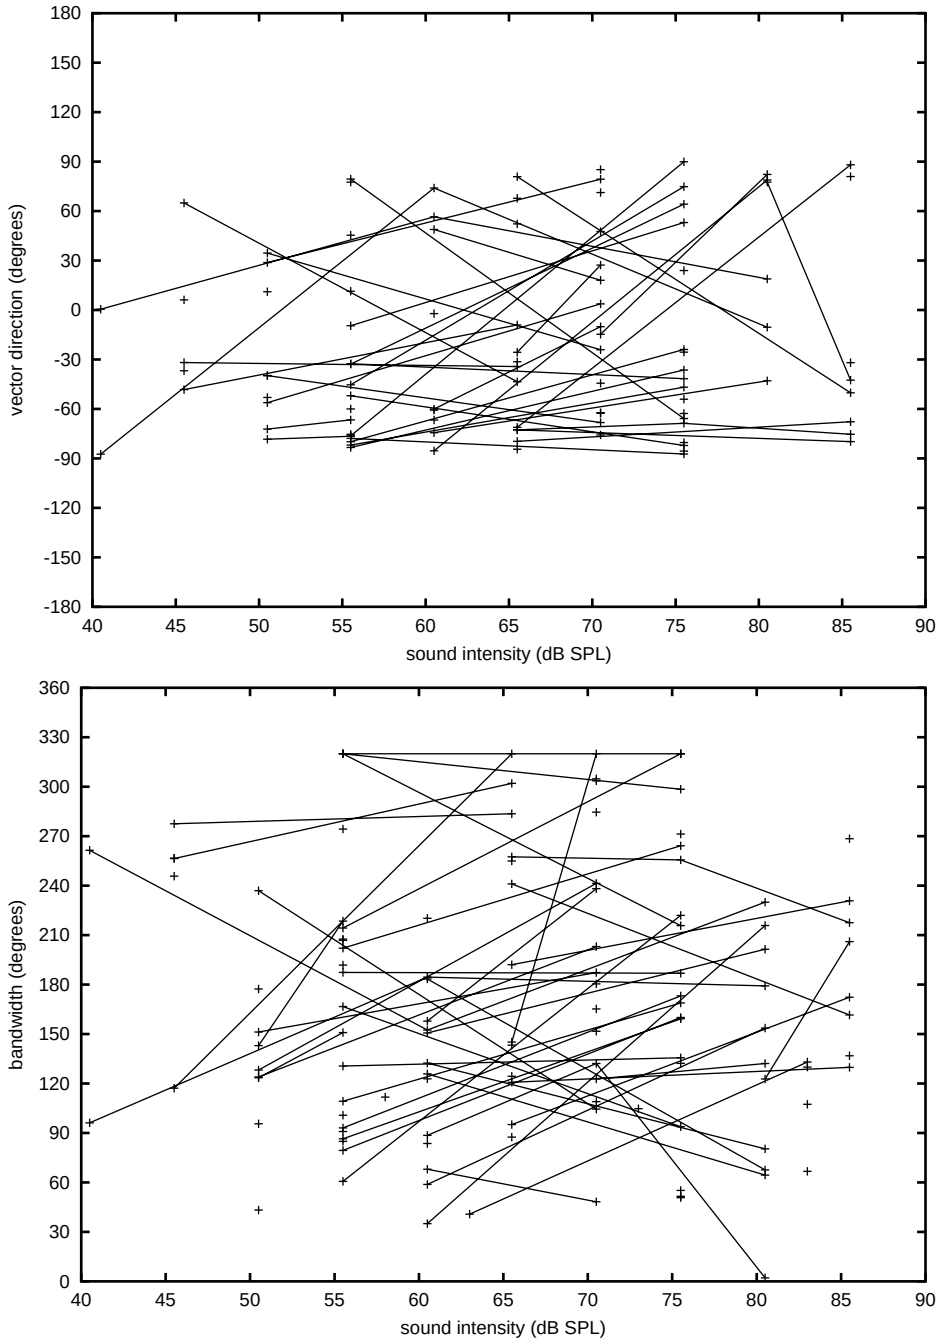

Figure 6.12: A graph showing the vector direction (top) and half-height bandwidth (bottom) of each unit *versus* the stimulus intensity at which the unit's tuning or bandwidth was determined. Points joined by lines show measurements of vector direction taken from the same unit at two different stimulus intensities. Stimulus duration was 100 ms in all cases.

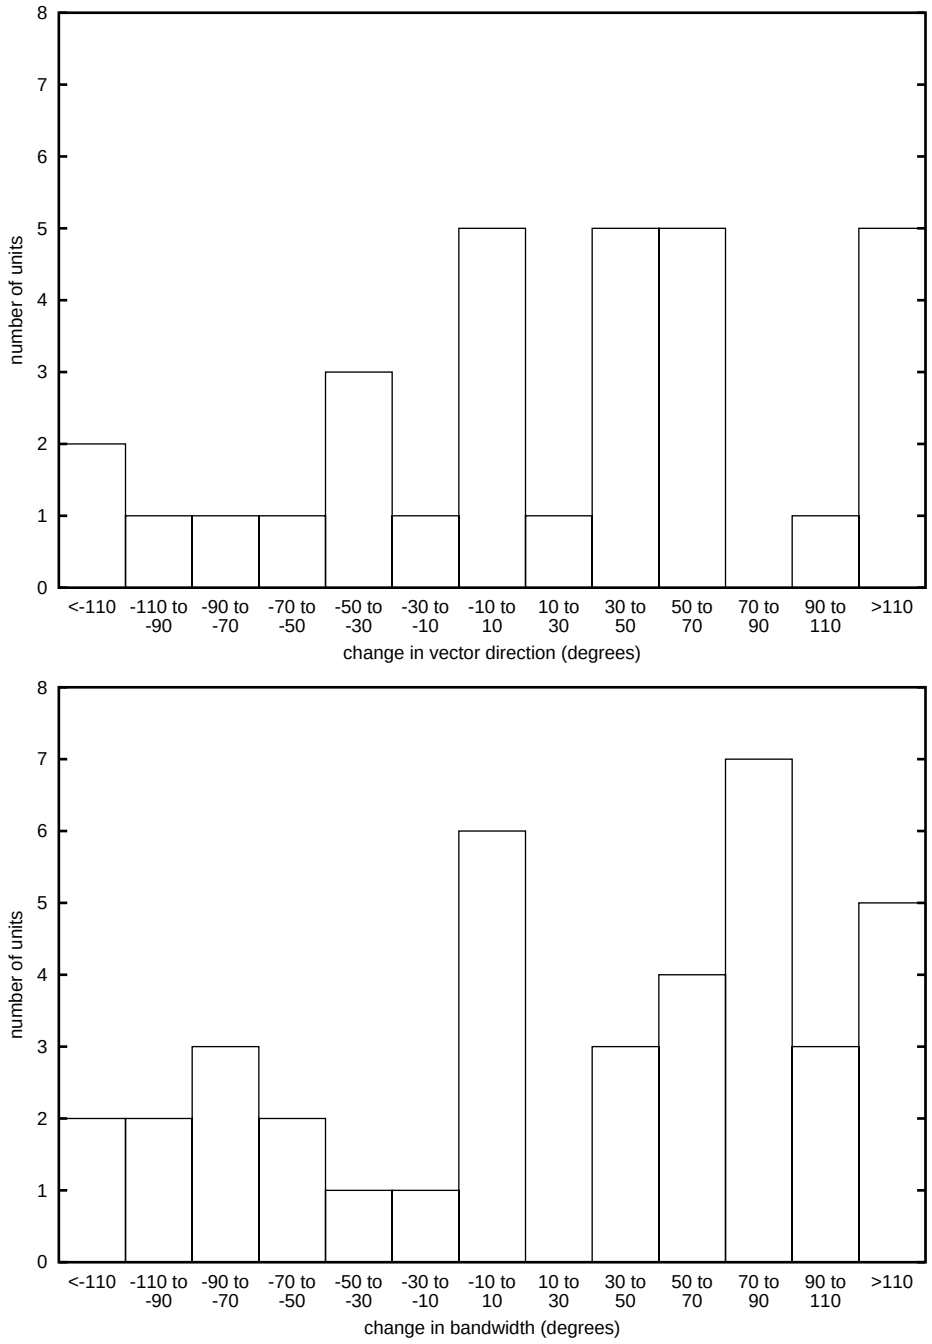

Figure 6.13: Histograms showing the change in each unit's vector direction (top) and half-height bandwidth (bottom) produced by a 20 dB increase in stimulus level. Only responses to 100 ms stimuli are considered.

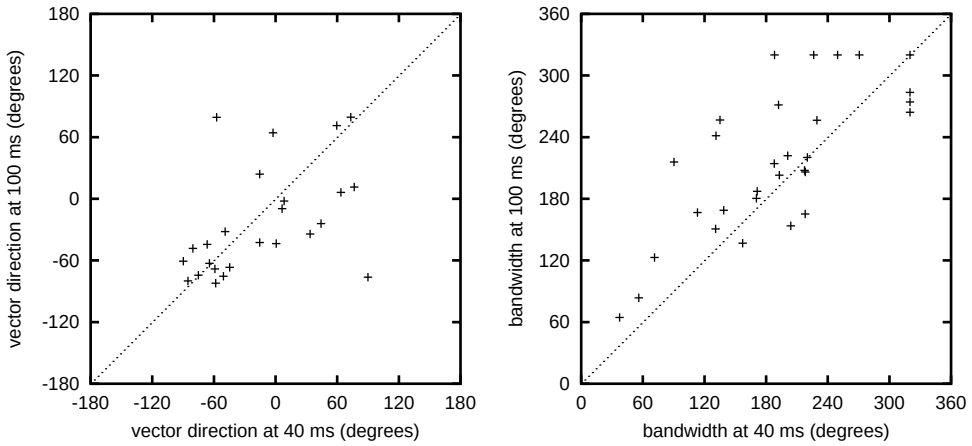

Figure 6.14: The effect of stimulus duration on vector direction (left) and half-height bandwidth (right). Values plotted on the abscissa are those measured with a 40 ms stimulus, and values on the ordinate were measured with a 100 ms stimulus of the same intensity.

6.3.5 Distribution of properties across A1

This study involves a detailed examination of a relatively small number of units in each experimental animal, and so the resulting data are not well-suited to building up a functional "map" of A1. Furthermore, because we do not have a well-defined measure of frequency sensitivity for most of our units, we are not able to investigate the relationship between tonotopy and other neuronal properties.

We can, however, examine the similarity of units which are recorded in the same electrode penetration. If units with similar properties are clustered in the same area of A1, then the variation of those properties within an electrode penetration ought to be less than the variation across the entire sample of units.

Not all properties can be examined in this way. We have already shown that unit bandwidth is significantly affected by changes in both stimulus level and stimulus duration, so, because we do not have a "standard" stimulus level which we have used for all units, it is difficult to make comparisons between units. By contrast, our measure of a unit's "preferred direction" is relatively unaffected by changes in stimulus level and duration, so it is possible to compare this property between different units under different conditions. We can also compare the value for noise detection threshold between different units, as this always measured under the same conditions.

The results are shown in figures 6.15 and 6.16, which compare the vari-

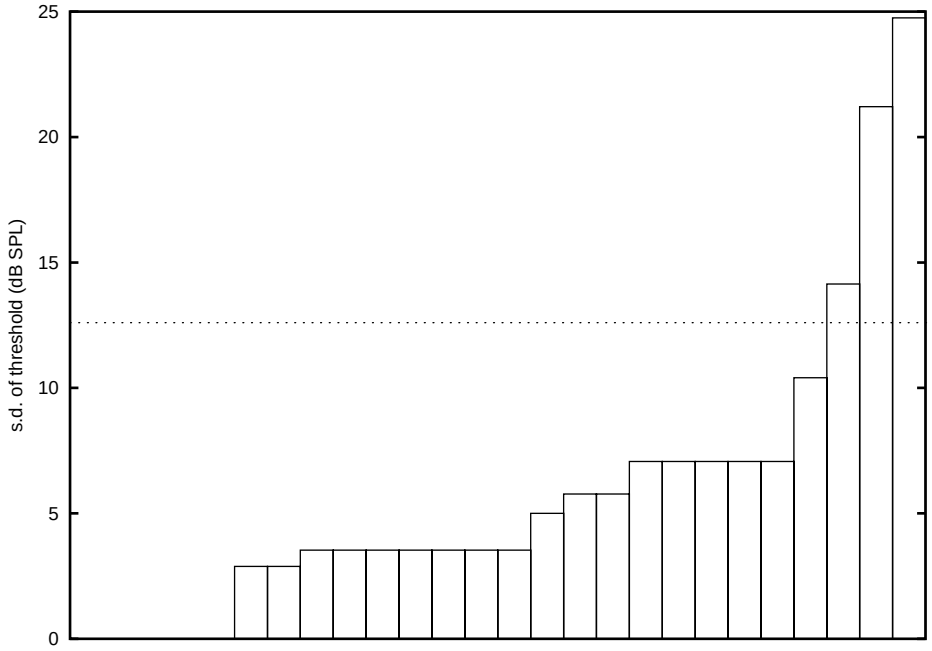

Figure 6.15: The standard deviation of noise detection threshold within each electrode penetration. Values are plotted in ascending order of standard deviation. Only those electrode penetrations which contain more than one unit are shown. The horizontal dotted line represents the standard deviation of noise detection threshold for all units in all penetrations.

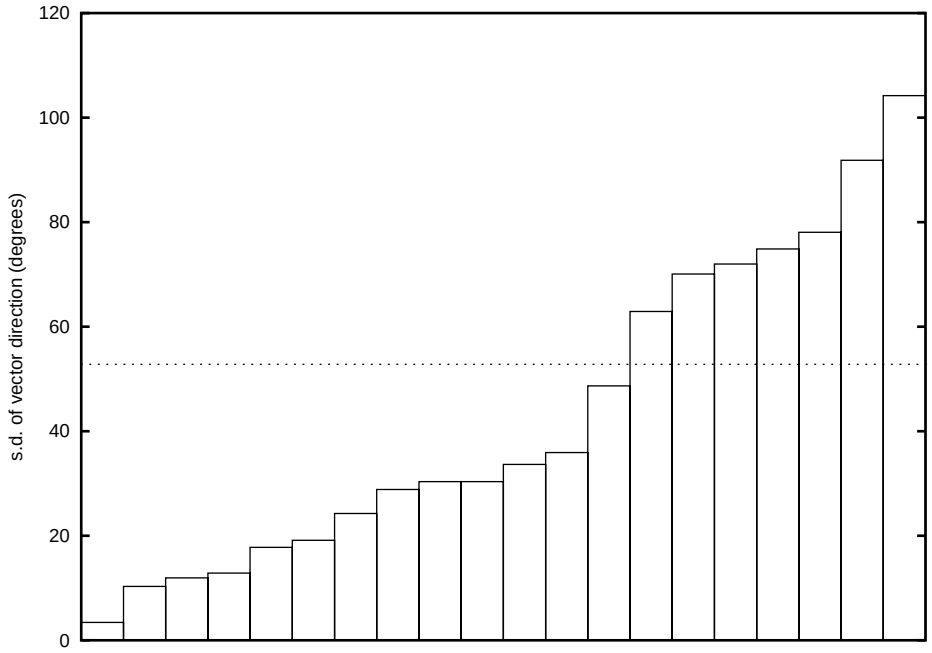

Figure 6.16: The standard deviation of vector direction within each electrode penetration. Values are plotted in ascending order of standard deviation. Only those electrode penetrations which contain more than one unit are shown. The horizontal dotted line represents the standard deviation of vector direction for all units in all penetrations. Where vector direction has been measured more than once for the same unit, only the first measure is used.

ation of unit properties within an electrode penetration to the variation of unit properties across the entire sample of units. As can be seen, for the majority of electrode penetrations, unit properties vary less within the penetration than within the whole population of units. This is particularly the case for unit threshold, where only 3/26 (12%) of penetrations show more variation in threshold than the population standard deviation; in the case of vector direction the equivalent figure is 7/20 (35%). These figures are, unfortunately, not very amenable to statistical analysis, as the number of units within each penetration is both small and different for different penetrations, and many penetrations contain recordings from only a single unit.

6.3.6 Simultaneous responses in multiple units

As mentioned in section 6.3.1, where several units were recorded from simultaneously, the units were grouped into pairs for analysis. Although the main intention was to examine the properties of these pairs using information theory (chapters 7 and 8), it is possible to draw some conclusions from a basic analysis of the correlation between firing in the two units of each pair. For each pair of units, we derived a correlation coefficient from a linear regression between the values of evoked spike count in one unit, and the simultaneously-recorded values of evoked spike count in the other unit. A linear regression is not ideal for this purpose, as it assumes that spike counts are distributed following a normal distribution, which they are not. However, we use it solely as a commonly-understood indicator of correlation, and we do not claim any significance for the absolute values of correlation coefficients.

In examining the similarity of response between units, it is important to consider how much of any correlation is simply down to the fact that the units are being presented with the same stimulus. Therefore, as well as measuring the correlation between units' responses to the same presentation of each stimulus, we also measured the correlation between responses which were made to different presentations of the same stimulus. Specifically, where a pair consisted of units A and B, we matched unit A's response to the nth presentation of a stimulus with unit B's response to presentation $n + 1$ of the same stimulus. This matching between one presentation and the next also "wrapped around", so that unit A's response to the last presention of a stimulus was matched to unit B's response to the first presentation of the same stimulus. We refer to presentations which have been mismatched in this way as "shifted" presentations. Correlation coefficients derived from shifted presentations should provide a control to those derived from simultaneous presentations, as the stimulus parameters are the same in both cases.

Figure 6.17 shows the correlation between the spike counts of the two

units in a each pair, as a function of the average of the two spike counts; the top graph represents the true correlation, and the bottom graph is the correlation between shifted presentations. The correlations between simultaneous presentations are quite variable, ranging from over 0.8 (highly correlated) through 0 (uncorrelated) to -0.4 (slightly negatively correlated), and the highest correlations occur at low firing rates. By contrast, the correlations between shifted presentations are lower and less variable, with hardly any correlations above 0.4; these correlations also seem to be unrelated to the number of evoked spikes.

While figure 6.17 shows a separate correlation coefficient for every different stimulus position, in figure 6.18 a single correlation coefficient is calculated for each pair across all different stimulus positions. As might be expected, the more extreme correlation values are "averaged out" by this, but coefficients for simultaneous presentations (top graph) are still considerably higher than those for shifted presentations (bottom graph), being mainly between 0 and 0.4 instead of 0 and 0.1. As a guide to the statistical relevance of these results, most points in figure 6.18 reflect a linear regression on the basis of 680 data points,[7] in which case any correlation coefficient of greater than $\sim$0.08 would reflect a significant correlation. Of course, as mentioned above, this assumes that the variables being correlated are distributed normally, which they are not.

Figure 6.18 also shows the effect of stimulus intensity on correlations, as a separate correlation coefficient is calculated for each stimulus intensity to which a pair is exposed. It would seem that, for most pairs, increases in stimulus intensity cause either little change or a small decrease in the correlation between spike counts in the two units. This holds true for both the simultaneous and shifted presentations.

Increases in stimulus duration appear to have a similar effect, or lack of it, on the correlation between activity as do increases in stimulus intensity. In figure 6.19 increases in duration either seem to have little effect, or lead to a small decrease in the amount of correlation.

One of the reasons the correlation between unit firing is interesting is that it is possible that the correlation itself is "tuned" to particular stimulus properties. This is investigated for a few example pairs in figure 6.20, which shows how both evoked spike count and correlation between spike counts vary as a function of stimulus azimuth. The first thing to notice is that the

[7] The number of data points used to calculate a correlation coefficient in figure 6.18 is the number of stimulus positions multiplied by the number of stimulus presentations at each position. In most cases this is $17 \times 40 = 680$, although for recordings in F9805 only 20 presentations were used per position, and in some recordings in other animals only 15 positions were used.

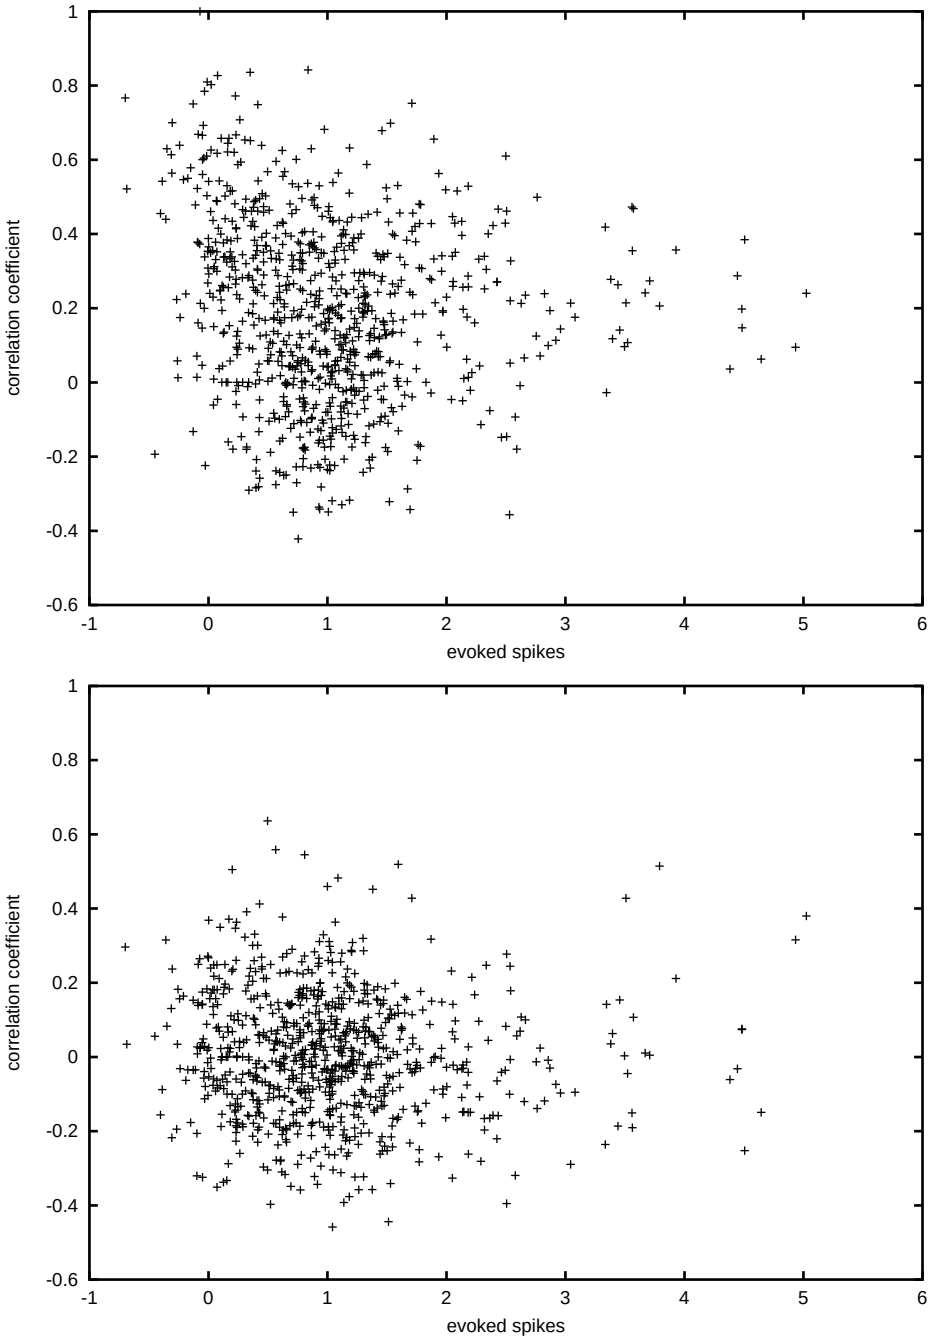

Figure 6.17: The correlation between evoked spike counts for simultaneously-recorded units, as a function of the average count. Correlations are shown for simultaneous (top) and shifted presentations (bottom). Coefficients are derived from a linear regression between all responses of the selected units to a stimulus of particular intensity, duration and azimuth.

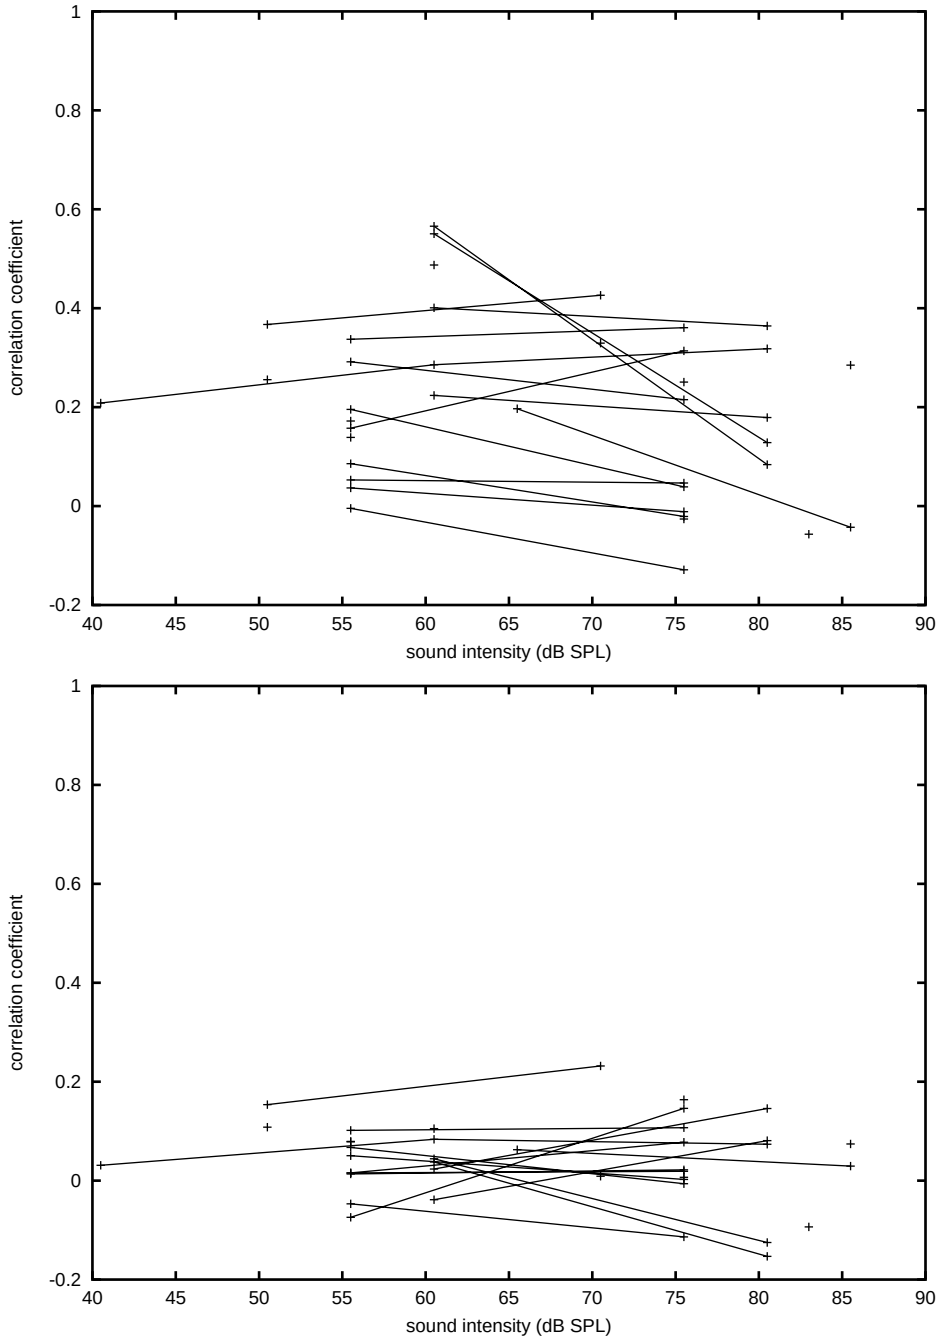

Figure 6.18: The correlation between spike counts of simultaneously-recorded units, as a function of stimulus intensity. Correlations are shown for simultaneous (top) and shifted presentations (bottom). Coefficients are derived from a linear regression between all responses, at all azimuths, to a 100 ms stimulus of the indicated intensity. Points joined by lines are recordings from the same units.

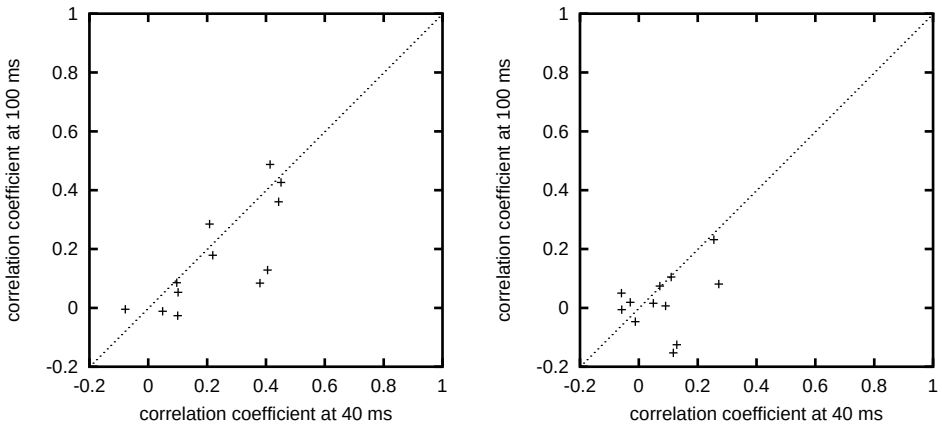

Figure 6.19: The effect of changes in stimulus duration on the correlation between evoked spike counts in simultaneously-recorded units. Correlations are shown for simultaneous (left) and shifted presentations (right). Coefficients are derived from a linear regression between all responses, at all azimuths, to 40 ms and 100 ms stimuli of the same intensity.

correlations for shifted presentations are all close to zero. This indicates that stimulus-driven similarities in units' responses are not in themselves particularly correlated. By contrast, the correlations for simultaneous presentations can be either considerably higher (figure 6.20, top left) or slightly lower (figure 6.20, bottom right) than zero, although both of these occurrences are rare. In fact, the pair presented in the bottom right of figure 6.20 is the only example of a consistently negative correlation between simultaneously-recorded units. More common are pairs where correlations are high, or at least higher than for shifted presentations, but only for a few stimulus locations (figure 6.20, top right and bottom left). These areas of increased correlation tend to be those where firing rates in general are low.

6.4 Discussion

Many studies have recorded responses to free-field stimuli in A1, and aside from the information theoretic analysis presented in chapter 8, the major differences in methodology between this and earlier work are the use of ferret rather than cat, the use of isoflurane as an anaesthetic, and the use of multiple electrodes. We used isoflurane in the hope that it might have less of a central depressant effect than barbiturate anaesthesia, but responses to pure tones were surprisingly poor, a finding which has subsequently been confirmed

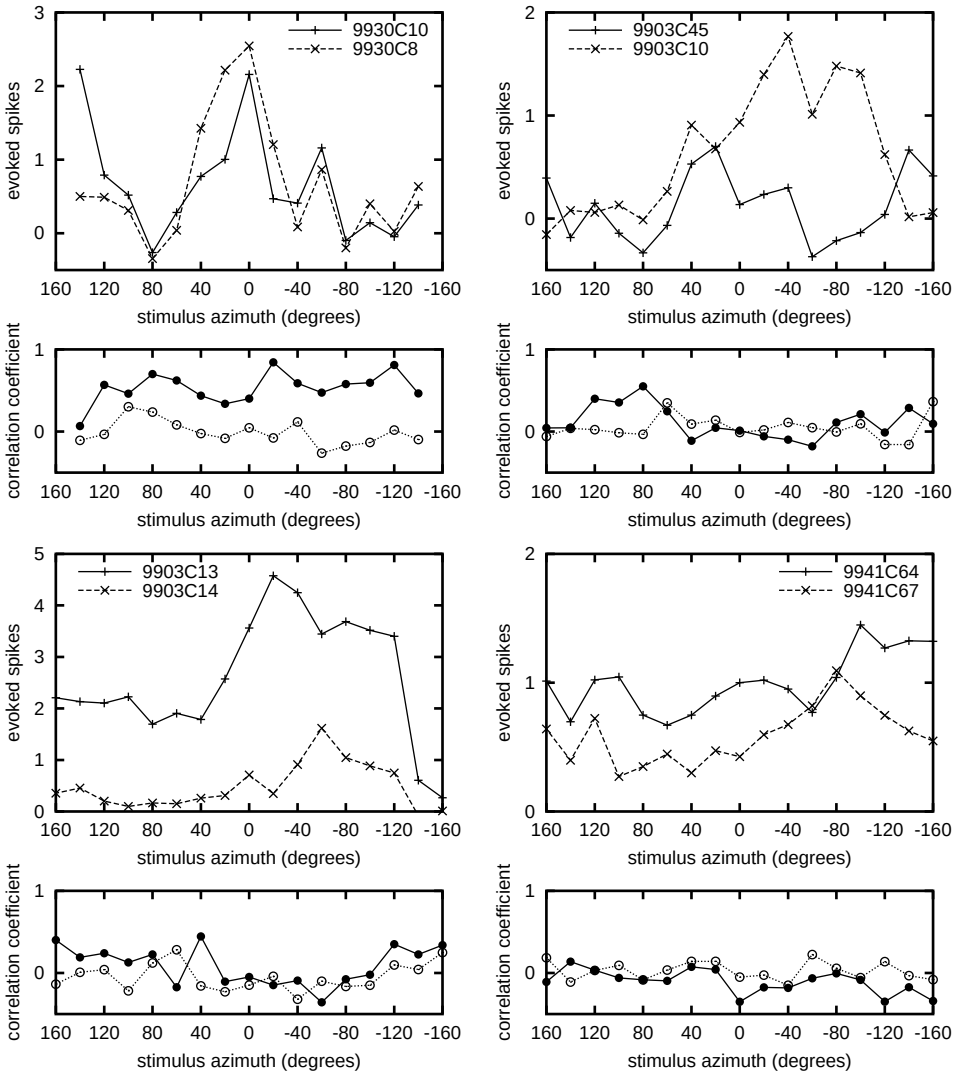

Figure 6.20: Four examples of the relationship between stimulus azimuth, the response of each unit within a pair, and the correlation of the units' responses. Filled circles (•) indicate correlations for simultaneous presentations and open circles (○) indicate correlations for shifted presentations. Each set of results represents only a single stimulus intensity and duration.

elsewhere[20].

Most of our results relating to location tuning follow those of previous research[49, 22, 105]. We find that a large number of units respond maximally to stimuli which are in the front contralateral quadrant. These units are, however, not very sensitive to location, and respond at a reasonably high level to noises emanating from a wide range of positions. It is observations like these which have questioned the involvement of A1 in the process of sound localisation.

We also concur with earlier results suggesting that the receptive field size of A1 cells is not fixed, but depends on the intensity of the stimulus being used to measure it[22, 12]. We find that increases in stimulus intensity lead to an increase in the number of positions to which a unit will respond. Because units have a maximum firing rate, increases in intensity have a greater effect at stimulus locations where the response is small than they do at locations where the response is large. The result is that the number of positions which respond at greater than 50% of maximum will increase.

Previous studies have gone beyond this in finding that increases in level can cause an actual shift in the direction to which a neuron is tuned, towards more ipsilateral positions[12]. While we do see a similar effect in our experiments, it is not statistically significant. This may be because we made only moderate changes in stimulus intensity, and we might have found receptive field shifts if larger variations had been used. Alternatively, it is possible that such shifts only occur in a different part of a unit's dynamic range, e.g. in the change from near-threshold to supra-threshold intensities, whereas we have always used stimuli which were at least 25 dB supra-threshold.

Whatever its extent, the intensity-dependence of receptive field properties could be taken to imply that A1 units are simply encoding stimulus level, and that any apparent direction-sensitivity is simply a result of the units' frequency tuning and binaural properties. Of course, this argument can be turned round—if these neurons were really representing the intensity of a stimulus, presumably they would do so in a manner which was independent of stimulus position. Also, the fact that these units show binaural interactions at all means that they are encoding localisation cues to some extent. We suggest that the main implication of these receptive field properties is that the stage in the auditory system where intensity and location have been disentangled from one another is later on in processing than A1.

We also make a novel finding—that increases in stimulus *duration*, as well as intensity, can cause an increase in the breadth of a unit's receptive field. This further supports the hypothesis that A1 cells are simply responding to the total amount of sound energy in the the stimulus.

Our methods are not well-suited to investigating how unit properties are

distributed across A1. Partly this is because, within each animal, we have concentrated on examining a few units in considerable detail, leaving us without enough of a spread of recordings to come to any conclusions about A1 topography. Also, we are not in a position to use the frequency tuning of a unit as a "shorthand" for its position with A1, due to our problems with weak tonal responses. We have examined the variation of neural properties within an electrode track, and both unit sensitivity and location tuning seem to vary less within a track than they do across the population as a whole. This finding is consistent with earlier work[22]. However, noise detection threshold varies with anaesthetic depth, and so the clustering of unit sensitivity may be an artefact of slow changes in anaesthesia levels.

Probably the most interesting aspect of this experiment is our investigation of the correlation between spike counts in simultaneously-recorded units. What we are actually interested in is whether, in principle, the information transmitted by evoked spike count in different units can simply be added up to give the total amount of information transmitted by A1 as a whole. This is a question which can best be answered by information theory (chapters 7 and 8), but it is necessary to detail the correlation between evoked spike counts if we are to understand the information theory analysis. What we find is that correlations between different units seem to be highest when firing rates are low. Furthermore, the correlations between non-simultaneous stimulus presentations are generally lower than those between simultaneous presentations, showing that the correlations that do exist are not directly due to the fact that the same stimulus is being used. All this suggests that it is the non-stimulus-driven activity, or "noise", which is correlated between units. This is interesting in light of evidence from modelling studies that small amounts of correlated noise can increase the information-carrying properties of a network[54].

However, it is important to distinguish between "noise" in modelling studies and in actual recordings. This is because real-world recordings are not a perfect record of the firing of neurons—often multiple neurons will be classified as a single unit, and some spikes will actually be non-neural events such as bloodflow artefacts or electrical noise. This is important because "recording noise" may well be correlated between electrodes in a way that the underlying firing of neurons will not, and the correlation will show up when the actual firing of neurons is low.

Chapter 7

Mathematical methods

7.1 Introduction

Experiments in neurophysiology typically aim to reveal simple relations between the value of some stimulus parameter and the number of spikes evoked by a stimulus; yet in many cases a parameter which is expected to be important appears to cause either a small or non-straightforward change in evoked spike count. This could be because the area of the brain under investigation is not involved in processing the parameter being varied, but there is often good evidence from non-physiological studies to suggest otherwise. An alternative explanation is that simple examinations of neural tuning curves do not tell the whole story about the information processing capabilities of a neuron.

One problem with the conventional method of interpreting tuning curves is that it relies on many assumptions about what a tuning curve 'ought' to look like. These assumptions may be at least partly contradictory: often a very sharply tuned, typically non-linear neuron is perceived to be particularly informative,[1] whereas other studies cite linearity of response as an indicator of neural processing.

Further problems rest with the response measure used to actually calculate conventional tuning curves, viz. evoked spike count. First, there is no *a priori* reason that high firing rates should be intrinsically more important than low firing rates, as is typically assumed; from a processing point of view what is important is that the responses to different stimuli should be distinguishable from one another. Second, there is evidence to suggest that the

[1] These are often referred to as 'grandmother cells', the idea being that there exist neurons whose responses are so sharply tuned that they would respond only to one's own grandmother.

time-scale over which spikes are typically counted (hundreds of milliseconds) is too long, and other response measures such as inter-spike interval may be more appropriate[114].

What is needed is some quantitative, 'assumption-free' measure of how well the output of a neuron can be used (by some decoding process) to determine the input stimulus.[2] This measure would at least tell us how much information there is in the firing of a neuron, even if it cannot tell us whether the nervous system uses this information. The measure should also be applicable to a variety of different response 'codes', not just evoked spike count.

It is in producing such a measure that information theory may be useful. Information theory is a branch of mathematics originally developed in the telecommunications industry, and it allows for the quantitative measurement of information transfer. The potential use of information theory in psychology was discovered quite early on, where researchers made attempts to discover how much information the human brain (as a whole) could transmit about various different sets of stimuli[5]. The first major application of information theory to neurophysiology was probably that of Richmond and Optican[88], who used it to investigate whether a neuron's 'spike rate' or 'spike timing' contained more information about a standardised stimulus set. These two types of studies show the use of measures of transmitted information in comparing both stimulus sets and response codes.

Sadly, most research has ignored information theory, probably in part because of the mathematical peculiarities associated with it. In particular, the fairly limited sample sizes which are used in most neurophysiological experiments lead to systematic errors in estimates of transmitted information[90]. Even some papers which have attempted to correct for this 'small-sample bias' have used correction techniques which are themselves biased. Another problem is that many comparisons of stimulus or response sets have not controlled for mathematical differences in the constructions of the sets being compared.

Here the mathematics behind information theory is described, and used to show how information theory can (and cannot) be of use to neurophysiology.

[2] Although the goal of the nervous system is not necessarily to determine what the input stimulus was. The important thing is that the organism in question makes a sensible response to the stimulus.

7.2 Quantifying information

7.2.1 Information content

7.2.1a Information and entropy

A given probability distribution, whatever it represents, can be considered to contain a certain amount of information. Or, to put it another way, it requires a certain amount of information to describe a particular probability distribution. The goal of information theory is to quantify this information, and the most commonly-used method of doing this is that adopted by Shannon[127].[3]

An easy way to convey the concepts behind Shannon information theory is to consider something like a series of coin-flips. How much information does it take to describe such a series? There are two factors that affect the answer to this question. First, the longer the series, the larger the amount of information. Second, if the coin is loaded to fall on one side more often than the other, the outcome of a particular coin-flip will be less informative than if the coin is entirely fair. This is because in the extreme case the coin always lands on the same side and, if it is known which side it is loaded towards, it takes no information at all to describe the series of coin-flips.

Viewed in this way, information can be considered as analogous to a quantity in use in physics and chemistry called 'entropy' or 'disorder'. This makes some intuitive sense, as a very ordered system takes little information to describe.[4] It also means that a certain amount of quantitative mathematics, designed for dealing with entropy, can be appropriated for the study of information.

The basic equation for the entropy of a discrete variable can thus be rewritten as a description of information. The result is shown in equation 7.1, where H is the information of a discrete variable x, x_i is the ith value of x, $p(x_i)$ is the probability of the ith value of x occurring, and there are X different such values.

[3]Shannon's definition of information is not the only one. Fisher information is also commonly used, especially in modelling circles as it defines limits for codes which are 'efficient' in informational terms. However, the discussion here is restricted to Shannon information, which is more often used for the analysis of experimental data.

[4]It is perhaps less intuitive that a very disordered system (e.g. white noise) takes a large amount of information to describe, because most people are unable to distinguish one very disordered system from another. Not all information (in the formal sense) is actually informative (in the colloquial sense).

$$H(x) = \sum_{i=1}^{X} p(x_i) \log_2 \frac{1}{p(x_i)} \tag{7.1}$$

It is worth showing that equation 7.1 really does fit the criteria for a measure of information as outlined above. First, equation 7.1 implies that information is large where the probabilities associated with each value of x are equal to one another. This is consistent with the coin-flip example given above: with a fair coin all series of heads and tails are as likely as one another, and the information needed to describe a particular length of series is maximal. Second, equation 7.1 predicts that, where the probabilities associated with each value of x are equal, larger values of X will result in larger values of H. This also fits with the coin-flip example, where the longer the series of coin-flips, the larger the amount of information. The value of X for a single coin-flip is just two (heads or tails), whereas a long series will have a much larger X (2^n, where n is the length of the series) and hence contain more information. Substituting the relevant numbers into equation 7.1 for a single (fair) coin-flip shows that $H = \log(\frac{1}{0.5})$. By the conventions of the 'silicon age' of binary computing, information is measured in 'bits' (i.e. binary digits) and hence the logarithm in equation 7.1 is taken to the base 2. This means that a single coin flip conveys exactly $\log_2(\frac{1}{0.5}) = 1$ bit of information.

7.2.1b Multiple dimensions

Information is calculated by considering the probabilities of all the different values of a particular variable (equation 7.1). The underlying dimensionality of this variable is irrelevant: the information will be the same for a single coin flipped three times (a 3-dimensional variable) as for an eight-sided die[5] rolled once (a 1-dimensional variable), as in both cases there are 8 possible outcomes and they are all equally likely. More formally, a two-dimensional distribution with X possible values along the x dimension and Y possible values along the y dimension can be considered as a one-dimensional distribution with XY possible values. Thus the information for a two-dimensional discrete variable is given by equation 7.2.

$$H(x,y) = \sum_{i=1}^{X} \sum_{j=1}^{Y} p(x_i, y_j) \log_2 \frac{1}{p(x_i, y_j)} \tag{7.2}$$

[5] Anyone who does not believe such dice exist obviously did not mis-spend their youth playing role-playing games.

7.2.2 Information transmission

7.2.2a Transmitted information

As we can describe the information contained by two variables both independently (equation 7.1) and together (equation 7.2), it is possible to say something about the relation between the two. If there is no relation between two variables, then the information of the two variables considered together will be no more than the sum of the information of the two variables independently. In other words, for two variables x and y, $H(x)$ plus $H(y)$ will be equal to $H(x,y)$. On the other hand, if there is any relation between the two, there will be some information "in common" between the two variables; the information of the two variables together will then be less than the sum of them both separately. This common information is called *mutual* or *transmitted* information, and is denoted by $T(x;y)$.[6] It is calculated as in equation 7.3. $T(x;y)$ can be thought of as being the amount of information about variable y which is conveyed by knowing the value of variable x.

$$T(x;y) = H(x) + H(y) - H(x,y) \tag{7.3}$$

The value $T(x;y)$ is the *average* information transmitted between x and y. It is also possible to calculate the transmitted information that is *conditional* on a particular value of x (x_i in equation 7.4). This is the amount of information about variable y which is conveyed if that specific value of x is known to be the case.

$$T(x_i;y) = \sum_{j=1}^{Y} \frac{p(x_i,y_j)}{p(x_i)} \log_2 \frac{p(x_i,y_j)}{p(x_i)p(y_j)} \tag{7.4}$$

If the conditional transmitted information is being calculated as in equation 7.4, it is easy to express the average transmitted information as the weighted sum of each of the conditional transmitted information values (equation 7.5).

$$T(x;y) = \sum_{i=1}^{X} p(x_i) T(x_i;y) \tag{7.5}$$

7.2.2b Multiple dimensions

Equation 7.5 deals with a 2-dimensional case, but it is relatively straightforward to carry a similar analysis on into more dimensions. Consider the

[6] The term "transmitted information" should not be considered to imply anything about the direction of any causal link between two variables, as $T(x;y) = T(y;x)$.

case where two dimensions (x and y) represent stimulus attributes, and a third dimension (z) is a response dimension. We are interested in how much information the two stimulus dimensions transmit separately ($T(x;z)$ and $T(y;z)$), or together ($T(x,y;z)$), to the response dimension. $T(x;z)$ and $T(y;z)$ can be calculated as for the two-dimensional case, simply ignoring the third dimension. $T(x,y;z)$ can be calculated by considering the variables x and y as a single variable with XY possible values (equation 7.6, see also section 7.2.1a).

$$T(x,y;z) = H(x,y) + H(z) - H(x,y,z) \tag{7.6}$$

It is also possible to calculate the information transmitted between one variable and another when the value of the third variable is known. This is notated as $T_x(y;z)$ and is derived as shown in equation 7.7.

$$T_x(y;z) = T(x,y;z) - T(x;z) \tag{7.7}$$

7.3 Information in neurophysiology

7.3.1 Why use transmitted information?

In neurophysiology, what is of interest is how the brain might use the output of a neuron to determine the stimulus that had produced it. Transmitted information is a measure of how much information about one variable is provided if a value for another variable is known (section 7.2.2). Thus, if the two variables are taken to be a stimulus parameter and a neural response parameter, the resulting transmitted information would seem to be an appropriate measure for neurophysiological use. Transmitted information values would then provide a measure of the discriminability of neural responses to different stimuli.

The use of transmitted information values can lead to conclusions which differ from those resulting from a more traditional analysis. Take the case where evoked spike count is used as the neural response parameter. Conventional wisdom holds that the maximum value of spike count, and the stimulus which produces it, are of key importance in understanding the cell's function. However, it is possible to imagine a case where a cell fires at maximum for nearly half the stimuli in a stimulus set and at minimum for most of the other stimuli, in which case the most informative values of the response are actually those which are between maximum and minimum, because these occur least frequently.[7]

[7] This assumes that all of the stimuli are equally likely.

46	2	0	0	0	48
3	44	5	0	0	52
1	4	38	7	1	51
0	0	6	40	4	50
0	0	1	3	45	49
50	50	50	50	50	250

r (rows) $\quad$ s (columns)

Table 7.1: An experimental frequency table. Each column represents a different stimulus, and each row represents a different response. Each cell thus shows the frequency of a particular combination of stimulus and response. The marginal totals along the bottom edge of the table show the frequency of each stimulus, and those along the right edge of the table show the frequency of each response. The bottom right corner shows there were 250 trials overall. See section 7.3.2a for further details.

Another counter-intuitive effect is that sometimes broadly-tuned neurons will transmit more information than sharply-tuned neurons, and sometimes the opposite. On one hand, a broadly-tuned cell responds differently to a large number of different stimuli, potentially transmitting a large amount of information. On the other hand, the responses of a broadly-tuned cell to different stimuli can be quite similar, so if there is a lot of noise or variability in the system then the differences will be swamped and very little information will be carried. Conversely, a very sharply-tuned cell may have, at the extreme, only two different response levels: "on" and "off". This obviously limits its transmitted information to a maximum of 1 bit, but the information transmission will be very resilient to noise or variability within the system.

7.3.2 Estimating transmitted information

7.3.2a Quantisation of stimuli and responses

Experimental transmitted information values can easily be obtained by constructing a contingency table. Stimuli are binned into S different bins, depending upon the value of some stimulus parameter s (an example of s might be stimulus intensity). Each of the S different stimuli are presented n_s times, and the responses are binned, according to the value of a response parameter r (e.g. spike count), into R different response bins. These results will then fill an $S \times R$ frequency table. Such a table is shown in table 7.1 for an imaginary

experiment where $S = 5$, $R = 5$ and $n_s = 50$. Information values for these data can then be calculated by converting the frequency values into probabilities, and substituting into the relevant equations from section 7.2.2. The derivation of average transmitted information for the dataset in table 7.1 is shown in figure 7.1; conditional transmitted information can also be calculated in a similar manner. Average and conditional transmitted information values are referred to as $\hat{T}(s;r)$ and $\hat{T}(s_i;r)$ respectively; the hats indicate that these are experimental estimates. The discussion below only refers to $\hat{T}(s;r)$, but the same arguments could equally be applied to $\hat{T}(s_i;r)$.

A result of this procedure is that the value for $\hat{T}(s;r)$ is entirely dependent upon the way that s and r are binned. Anything which affects the number or occupancy probabilities of either the stimulus or response bins will affect the value of $\hat{T}(s;r)$. The choice of binning procedure thus becomes crucially important.

One factor which affects the binning of s and r is the resolution with which either can be measured. The lower the resolution, the lower the maximum value of $\hat{T}(s;r)$. This is often described as 'quantisation loss' relative to some 'real' underlying value of $T(s;r)$. Furthermore, $\hat{T}(s;r)$ is limited by the *smaller* of the information in the stimulus set ($H(s)$) and the information in the response set ($H(r)$), and thus by whichever of stimulus resolution and response resolution is worse.

However, a more obvious factor which determines binning procedures is the limited number of trials which can be presented in any experiment. The larger the number of bins, the fewer trials that will fall into each bin. Thus with more bins, the estimate of the occupancy probability of each bin will be less accurate, and so will be the value of $\hat{T}(s;r)$. In fact, information theoretic measures estimated from too few trials (or too many bins) are not simply inaccurate due to this undersampling; they are also systematically biased upwards (section 7.3.2b) due to the increased susceptibility to noise.

7.3.2b Correcting for small-sample bias

On average, a sample estimate of transmitted information ($\hat{T}(s;r)$) will be greater than the corresponding population value ($T(s;r)$). To see why, consider a limiting case where, irrespective of some stimulus s, the value of some response r varies randomly from trial to trial. There is obviously no transmitted information between s and r. However, because there is a certain amount of variation in r, for a small sample size there will be differences in the distributions of r values recorded for different values of s. Hence the experimental estimate of transmitted information will be non-zero. It is important to note that the variation in r in this example could either come

0.184	0.008	0	0	0	0.192
0.012	0.176	0.02	0	0	0.208
0.004	0.016	0.152	0.028	0.004	0.204
0	0	0.024	0.160	0.016	0.2
0	0	0.004	0.006	0.180	0.196
0.2	0.2	0.2	0.2	0.2	1

r (rows) ; s (columns)

$$
\begin{aligned}
H(s) &= \sum_{i=1}^{S} p(s_i) \log_2 \frac{1}{p(s_i)} \\
&= 5 \times \left(0.2 \log_2 \frac{1}{0.2} \right) \\
&= 2.323
\end{aligned}
$$

$$
\begin{aligned}
\hat{H}(r) &= \sum_{i=1}^{R} p(r_i) \log_2 \frac{1}{p(r_i)} \\
&= \left(0.192 \log_2 \frac{1}{0.192} \right) + \left(0.208 \log_2 \frac{1}{0.208} \right) + \left(0.204 \log_2 \frac{1}{0.204} \right) \\
&\quad + \left(0.2 \log_2 \frac{1}{0.2} \right) + \left(0.196 \log_2 \frac{1}{0.196} \right) \\
&= 2.322
\end{aligned}
$$

$$
\begin{aligned}
\hat{H}(s,r) &= \sum_{i=1}^{S} \sum_{j=1}^{R} p(s_i, r_j) \log_2 \frac{1}{p(s_i, r_j)} \\
&= \left(0.184 \log_2 \frac{1}{0.184} \right) + \left(0.008 \log_2 \frac{1}{0.008} \right) \\
&\quad + \cdots + \left(0.180 \log_2 \frac{1}{0.180} \right) \\
&= 3.054
\end{aligned}
$$

$$
\begin{aligned}
\hat{T}(s;r) &= H(s) + \hat{H}(r) - \hat{H}(s,r) \\
&= 2.323 + 2.322 - 3.054 \\
&= 1.591
\end{aligned}
$$

Figure 7.1: Calculation of average transmitted information from the data in table 7.1. Frequency values are converted into probabilies (top) before information values are calculated (bottom). Most of the working for $\hat{H}(s,r)$ has been omitted. Hats indicate experimental estimates. For further details see section 7.3.2a

from the underlying distribution of r itself or come from experimental noise. This suggests that a worse signal-to-noise ratio will produce a larger bias.

There are several different techniques for estimating, and hence correcting for, small-sample bias. Panzeri and Treves[90] review many of the current approaches to this problem, including their own, and it is their analysis which is used here. They point out that many other bias estimation techniques contain systematic inaccuracies, and previous studies which have used them may thus be unreliable.

Panzeri and Treves derive a description of small-sample bias in terms of a series expansion. In their notation the leading, and most important, term in this series is called C_1, and for most practical purposes C_1 can be used as an estimate of the bias. Consider a situation with a number of stimulus bins S and a number of response bins R. The joint probability distribution will fill an $S \times R$ contingency table. Where all of the bins of the contingency table are occupied, the value of C_1 for that table depends on the number of trials, n, and the degrees of freedom of the contingency table, $(S-1)(R-1)$, as shown in equation 7.8.

$$C_1 = \frac{1}{2n \log 2} \{(S-1)(R-1)\} \tag{7.8}$$

However, where the contingency table contains many empty bins, the number of effective degrees of freedom is smaller. In fact, C_1 depends not on the absolute number of bins in the contingency table, but instead depends on the number of 'relevant' bins. A 'relevant' bin is here defined as a bin whose underlying occupancy probability (*not* its experimental occupancy probability) is non-zero. Typically all of the stimulus bins will be relevant, as there is no *a priori* reason for assuming that any one of the stimulus bins should be more likely than any of the others. However, not all of the response bins will be relevant. In particular, it is usually the case that only some response bins are relevant for each stimulus, and equation 7.8 must be modified to take this into account. This results in equation 7.9, where $\tilde{R}$ is the number of relevant response bins overall, and $\tilde{R}_{s_i}$ is the number of response bins which are relevant when the ith stimulus is presented.

$$C_1 = \frac{1}{2n \log 2} \left\{ \sum_{i=1}^{S} \tilde{R}_{s_i} - \tilde{R} - (S-1) \right\} \tag{7.9}$$

The problem then arises of how to estimate the numbers of relevant bins ($\tilde{R}$ and $\tilde{R}_{s_i}$ in equation 7.9) from their experimental values. Here, Panzeri and Treves used a Bayesian statistical approach[142]. In essence, a value for $\tilde{R}$ is chosen and then used to calculate a Bayes estimate of the true

probabilities, given the experimental frequencies and a plausible set of prior assumptions.[8] This is repeated for different values of $\tilde{R}$ and the one chosen is that which gives an expectation value of the number of occupied bins closest to the experimental value. The same procedure is used for $\tilde{R}_{s_i}$, except that the number of trials per stimulus is used in calculations in place of the total number of trials.

As deriving a value for C_1 requires the numbers of relevant bins to be estimated from the data, it might seem that experimental estimates of C_1 might be just as subject to bias as estimates of transmitted information itself. However, closer examination reveals that this is not the case. While both C_1 and $\hat{T}$ must be estimated from the data, C_1 depends only on the *number* of relevant bins, whereas $\hat{T}$ depends both on the number of bins and on their associated probabilities. Panzeri and Treves show by computer simulation that estimates of C_1 are accurate provided the number of trials per stimulus is at least as large as R.

More subtly, it might be argued that the process of estimating a value for C_1 involves making a certain number of assumptions, and this violates the whole idea of transmitted information as an 'assumption-free' statistic. While this is obviously true to some extent, the assumptions being used are fairly weak. There are no assumptions as to *which* bins are occupied, only to *how many* bins are occupied. There are also no assumptions about the occupancy probabilities of the relevant bins.

The formula for C_1 given in equation 7.9 is the correction for average transmitted information; it is also possible to calculate a correction for the conditional transmitted information. However, the formula is more complex, the derivation less clear and it is likely that more trials per stimulus would be needed to provide an accurate experimental estimate of its value. For that reason the equation is not given here; readers are referred to Panzeri and Treves' own work.

7.3.3 Experimental design

As mentioned above (section 7.3.2a), information measures are entirely dependent on the way data are binned. This means that the choice of binning procedure must be made very carefully. Consider an experiment with S different stimuli, each presented n_s times, where responses are binned into R different bins.

[8] The prior probability function is chosen to be constant among the $\tilde{R}$ non-empty bins, and is a lower constant for the $R - \tilde{R}$ empty bins, constrained by requiring that the probability of a bin being empty is $n/\tilde{R}$ times larger than the probability of it being occupied.

- To avoid underestimating the true information transmission capabilities of a given system, the stimulus information ($H(s)$) should be as high as possible. Thus S should be set as large as possible, and each stimulus should be presented an equal number of times so that the occupancy probability of each stimulus bin is equal.

- To avoid underestimating $T(s;r)$ for a given value of $H(s)$, the response information ($H(r)$) should be at least as large as $H(s)$ (section 7.3.2a). Thus R should be at least as large as S, and the response binning procedure should be such that each response bin should be equally likely to be occupied.

- To avoid overestimating $T(s;r)$ due to small-sample bias, R should not be set greater than n_s (section 7.3.2b).

Given the above, the 'optimal choice' for the numbers of bins would appear to be to set $S \sim R \sim n_s$. It would also appear to be 'optimal' to ensure that each stimulus bin is equally occupied, and each response bin is equally occupied. This may not actually be possible in the case of the response bins, because response binning is not entirely under the experimenter's control. To take an extreme case for the sake of simplicity, if the response measure only has two possible values, e.g. the number of spikes in a given 1 ms time window, there can be only two response bins. The experimenter therefore has no choice over binning procedure, and consequently cannot control the occupancy probability of the bins.

In any case, these choices are optimal only in the sense of maximising $\hat{T}(s;r)$ in an attempt to approximate the 'real' underlying value of $T(s;r)$. However, the 'real' value for $T(s;r)$ is that where the binning procedure used to analyse the results is the same as that hypothetically used by the nervous system. This means that some sort of neurally plausible binning procedure may be more useful than one which maximises $H(s)$ or $H(r)$. Also, the low binning resolution required to avoid small-sample bias means that experimental transmitted information values are almost always going to be far from maximal in any case. Although many experimenters would wish otherwise, attempting to measure some *absolute* information transmission value for a neural system may be futile.

In fact, a better use of information theory is probably to *compare* transmitted information values, either comparing different stimulus sets while keeping the response binning procedure constant or *vice versa*. This does not mean completely disregarding the points about binning procedure outlined above, but it does shed a slightly different light on them.

- As $\hat{T}(s;r)$ values are limited by the smaller of $H(s)$ and $H(r)$, it is important to ensure that the smaller value does not obscure the comparison being made. For example, if two different stimulus sets are being compared, it is a bad idea to set $H(s) > H(r)$, as $T(s;r)$ values will then be potentially subject to a ceiling effect imposed by the value of $H(r)$.

- It is important to make sure that the two stimulus or response sets being compared have the same information; if they do not, then the value for $T(s;r)$ could simply be higher for the set with the higher information.

- To avoid differences in small-sample bias from contaminating the comparison being made, R should not be set greater than n_s.

The second of these three points runs up against the problem mentioned above: that response binning is not entirely under the control of the experimenter. It is straightforward enough to ensure that different stimulus sets have the same $H(s)$, but it is more difficult to equalise $H(r)$ for different response measures. While the number of bins can be set to any value by the experimenter (provided $R \leq n_s$), the occupancy probability of the response bins can only be controlled through the choice of binning procedure, and this choice may be quite limited. Thus in many cases it will be difficult to make comparisons of different response measures. It should be noted that, in this context, "different response measures" can refer to the same response code derived from different neurons.

7.4 Conclusions

The use of transmitted information opens up the possibility of measuring the strength of the relation between a stimulus set and a neuron's response, without having to make any *a priori* assumptions about what that relation ought to look like. However, as section 7.3.3 shows, that does not mean that performing an information theory analysis involves *no* assumptions. The choice of stimulus sets, response measures and the binning of both must be made by the experimenter, and these completely determine the final measure of transmitted information. It should be noted that these same caveats apply to techniques like those employed by Middlebrooks[81, 82, 144] which measure information implicitly, but where readers may not be so wary of the mathematical problems outlined here.

The main point of quantifying information in this way is to allow conclusions to be drawn about a neuron's involvement in certain sorts of processing. For example, an experiment which compared different stimulus sets might conclude that the stimuli which a neuron is 'designed' to deal with are those about which the neuron's output transmits most information. Similarly, if some response codes could be validly compared then it might be concluded that the response patterns ("the neural code") a neuron is 'designed' to produce are those which provide most information about the neuron's input. Comparisons could also be made using response measures derived from different neurons, hoping to discover which neuron is more involved in the processing of a particular stimulus set.

However, there are problems with these arguments about a neuron's 'design', as it is possible that large values of transmitted information are not the ultimate goal of a well-designed nervous system. Research on neural networks has shown that changes to network parameters which increase other desirable characteristics, e.g. memory capacity, learning speed and network stability, may decrease information transmission.

Indeed, the 'assumption-free' nature of transmitted information outlined above may in some cases be a weakness rather than a strength, because there may be good design reasons for a neural network to use particular sorts of tuning curve. For instance, with some learning rules (e.g. the original formulation of the Hebb learning rule) high firing rates are more important than low firing rates, irrespective of their discriminability, as it is high firing rates which drive synaptic learning.

While these factors suggest that the information transmitted by a well-designed neuron has to be balanced against other factors, it still seems implausible that a neuron would transmit more information about a task it was not involved with than about a task which it was involved with. This is another argument for comparing transmitted information between stimulus sets, response measures or different neurons, rather than simply drawing conclusions from the absolute $\hat{T}(s;r)$ values.

It thus seems that information theory, correctly applied, can be a useful tool in comparing the involvement of neurons in processing different sets of stimuli or producing different response codes. Similar comparisons can also be made between neurons. All of these comparisons might then offer some indication as to the roles of the neurons being investigated. However, all such comparisons must be made with the mathematical basis of information theory in mind.

Chapter 8

Sound location information in single-cell recordings from ferret A1

8.1 Introduction

While single-cell recordings can produce large quantities of data, it is often not clear how the data should be analysed. Traditionally, the numbers of spikes evoked in a single neuron by each stimulus would simply be plotted on a graph, and possibly fitted to a simple curve. The meaning of such a "tuning curve" in the sensory periphery is not too difficult to work out, as there are well-understood models of how sensory receptor cells work. However, for recording sites like A1, which are many levels of processing away from either the sensory or motor periphery, the meaning of a tuning curve is much less obvious.

One possible solution is to analyse recordings using information theory, as described in chapter 7. This provides a quantitative assessment of how much information is conveyed by a particular measure of neural response about a particular set of stimuli. Not only that, but it does so without requiring that the neuron's response fits any particular model. This approach was pioneered in studies of high-level areas of the primate visual system, where very little is known about neuronal tuning. In those studies, information theory was used to show that more information about a set of visual stimuli was present in the timing of a neuron's response than was present in a simple count of the number of spikes that were evoked[88, 87].

Used in this way, information theory could help to reconcile the electrophysiology of A1 with the results of behavioural studies. The problem

is that behavioural evidence implicates A1 in sound localisation (see chapters 3 and 4), while A1 neurons seem to be quite insensitive to sound source location (chapter 6). However, this apparent insensitivity of A1 neurons could be deceptive—it is possible for a broadly tuned neuron to carry a great deal of information, as long as its firing pattern is reliable and noise-free. Even if the number of spikes fired by an A1 neuron does not carry much information, it is possible that other aspects of the response (e.g. spike timing) carry more, as would appear to be the case in the visual studies cited above.

That such an approach might work in the auditory system is suggested by studies of AES, a secondary auditory area in the cat.[1] These show that an artificial neural network can derive more information about the location of a sound source from the timing of a neuron's response—particularly the latency of the first spike—than from the overall spike count[37]. The same research group has also shown that the errors made by such an artificial network are reduced if the network is provided with the responses of several neurons[38], although this study did not use information theory explicitly. Theory-based studies have also demonstrated that, even if individual neurons transmit rather little information about stimulus position, small groups of A1-like neurons could in principle transmit a great deal more information[53], even if the neurons' firing is partly correlated[54].

We therefore intend to analyse the sensitivity of A1 neurons to sound source location by applying information theory to the recording data presented in chapter 6. In particular, we will investigate whether significant information is carried by the timing (rather than just the number) of spikes evoked by a stimulus, and/or whether it is carried by the combined firing of several different neurons.

8.2 Methods

8.2.1 Numerical conventions

Throughout this chapter, azimuth is expressed in degrees, from $-180°$ to $+180°$, with $0°$ being straight ahead. The plus and minus signs relate the azimuth to the side of the brain from which recordings are being made—positive azimuths are on the same side of the animal's midline as the side of recording, while negative azimuths are on the other side of space. Experimental measurements are expressed either to the precision of measurement, or to 3 significant figures, whichever is lower. Other numeric quantities are treated similarly.

[1] AES is the anterior ectosylvian sulcus, also known as fAES.

8.2.2 Electrophysiology

The methods and protocols used for single-cell recording from ferret A1 are given in detail in chapter 5 and section 6.2. In brief, recordings were made from the left A1 of six isoflurane-anaesthetised ferrets: F9805, F9929, F9930, F9903, F9940 and F9941. Stimuli were presented in the free-field in an anechoic room. The sound source—a loudspeaker—was positioned along the horizon (0° elevation), at a distance of 0.65 m from the ferret's head, and could be placed at any azimuth from +160° to −160°. The stimuli used were 40 ms or 100 ms noise bursts, with an intensity at least 25 dB above the noise threshold of the clusters being recorded from.

Stimuli were presented in "azimuth sweeps"—starting at one extreme of the range of azimuths and proceeding in order to the other extreme—with stimulus intensity and duration being kept constant throughout a sweep. The position of the loudspeaker was moved in 20° steps, with 40 stimuli being presented after each step (or 20 stimuli in the case of F9805). At the end of the sweep, the intensity or duration of the stimulus would be changed before making another azimuth sweep in the opposite direction.

8.2.3 Analysis

8.2.3a General procedures

Our analysis uses information theory in the form described in chapter 7; there are other information theoretical techniques, and this one was selected for its theoretical and computational simplicity. The first step in performing this type of analysis is to create a contingency table, showing the probability that each possible response is associated with a particular stimulus.

Such an approach implies a set of discrete stimuli and responses, so both the stimulus and response values from our experiments needed to be assigned to discrete "bins". The process of stimulus binning was quite straightforward: although stimulus azimuth is potentially a continuous variable, in fact we only ever used 17 different speaker positions (as described above). Chapter 7 suggests that the number of stimulus bins should not be more than the number of presentations per bin, and as we used either 20 or 40 repeats per azimuth, it was possible to treat each of the 17 loudspeaker positions as a different stimulus bin.

By contrast, it is much less obvious what the binning procedure should be for measures of neuronal response. The response measures in question are evoked spike count and first-spike latency, as calculated in section 6.2. The only definite restriction on their binning, as suggested in chapter 7, is that it is desirable to have the same number of response bins as there are

stimulus bins (in this case, 17). The actual binning procedures we used were as follows:

- Evoked spike count was split into bins 1 spike wide, starting at values less than 1 evoked spike, then 1–2 spikes, 2–3 spikes, etc., up to the last bin which covered values greater than or equal to 16 evoked spikes.
- First-spike latency was split into bins 10 ms wide, starting at values less than 10 ms, then 10–20 ms, 20–30 ms, etc., up to the last bin which covered values greater than or equal to 160 ms.
- We also used a second measure of first-spike latency, which was split into 5 ms bins starting at values less than 5 ms, then values from 5–10 ms, 10–15 ms, etc., up to the last bin which included values greater than or equal to 80 ms.
- In all cases, a value is assigned to a particular bin if is greater than or equal to the lower bound of the bin, and less than the upper bound of the bin.

The precise bin boundaries were chosen mainly for reasons of computational simplicity, but the boundaries of the first and last bins in each procedure were chosen to cover the range of typical response values, as shown in figure 6.11. It is important to note that the same binning procedure was used for every unit—there is no requirement from information theory that this be the case, but it makes comparison between units more straightforward.

8.2.3b Single unit analysis

Once it was decided how stimuli and responses were to be binned, each unit's response to a particular stimulus presentation could be assigned to a bin in a contingency table. A separate contingency table was constructed for each unit, and also for each azimuth sweep to which the unit was subjected. In addition, different contingency tables were constructed for first-spike latency and for evoked spike count.

A value of transmitted information was then calculated for each contingency table, using the equations given in section 7.2. Such calculations are subject to a certain amount of systematic error, in that experimental estimates of transmitted information ($\hat{T}$) tend to be higher than the actual underlying values (T). However, it is possible to estimate the size of this bias, and correct for it. The method of bias correction which we use is that of Panzeri and Treves[90], as summarised in section 7.3.2b. Specifically, we

use formula 7.9 to calculate the bias estimate (C_1). As part of this calculation, a procedure involving Bayesian statistics is used to decide which bins in the contingency table are "relevant", and which bins simply contain "noise".[2] We mention this because a slightly different bias correction technique is used for the multiple unit analysis described below.

8.2.3c Multiple unit analysis

In the case of simultaneous recordings, contingency tables were generated for each pair of units, as well for each of the individual units. While the response of a given pair is intrinsically a two-dimensional measure, with the responses of the two units being separate dimensions, the contingency table was constructed by "flattening" the two dimensions into one. Thus the responses of two single units, each of which would separately be represented using R different response bins, were combined to give a single response measure with R^2 different bins. Transmitted information was then derived from this, as described in section 7.2.

However, the contingency tables which result from this analysis are very large, and this prevents information theory from being applied in the same way as it is for single units on their own. This is because Panzeri and Treves' bias correction is only accurate if the number of response bins is less than or equal to the number of presentations of each stimulus. We presented each stimulus 40 times at each position (20 times in the case of F9805), thus it is impossible to compensate accurately for small-sample bias if we used $17^2 = 289$ response bins.

To get round this, we reduced the size of the contingency tables which were produced. The method of table reduction was suggested to us by Dr. Israel Nelken, and it involves removing rows or columns from the table, one at a time. First, the extra-large contingency table is used to calculate an estimate of transmitted information ($\hat{T}$), and, using the Panzeri and Treves method, an estimate of the bias (C_1). Next, the table is examined to find the row or column with the least occupancy probability. This row or column is then merged with one of its two neighbours—whichever has the lesser occupancy probability. Then, values of $\hat{T}$ and C_1 are calculated for this new table, before reducing its size again. This is done until the contingency table consists of a single row or column. The table which is actually used in further analysis is that which gives the maximum value of $\hat{T} - C_1$.

It is important to realise that this table-reduction procedure may itself have introduced a bias into the measurement of transmitted information.

[2] The computer code for this Bayesian procedure was supplied to us by Dr. Stefano Panzeri.

This much is suggested by the fact that the procedure involves maximising $\hat{T} - C_1$, as it is possible that this simply selects the table for which bias correction is least effective. On the other hand, by removing relevant rows and columns from the contingency table, it is possible that this method actually underestimates transmitted information. This issue can only fully be resolved by computer simulation, which lies outside of the scope of this study. However, any results from simultaneous recordings were compared to those from "shifted", non-simultaneous responses to the same stimulus. The construction of these shifted response sets is described in section 6.3.6. Both simultaneous and non-simultaneous data sets should show similar amounts of bias, but if there is any information encoded in the correlation between units then this will only be seen in the simultaneous recordings.

Another problem presented by multiple unit recordings is that the procedure required to calculate C_1 (the bias estimate) from formula 7.9 is not applicable here. This is because one of the assumptions used to estimate the number of "relevant" bins is that neighbouring bins have similar occupancy probabilities—this is not the case when a two-dimensional response matrix is flattened into a single dimension. Also, as the table-reduction effectively removes bins which are not relevant, it is inappropriate to use a bias-estimation procedure which involves an estimate of the number of relevant bins. For this reason, we use formula 7.8 to calculate C_1 in calculations where table-reduction is involved; this formula assumes that all bins in the contingency table are relevant.

8.3 Results

8.3.1 Unit selection

The units analysed here are the same as those in chapter 6; the process of unit selection is described in detail in section 6.3.1.

8.3.2 Information transmitted by single units

As with the measures of location tuning used in chapter 6, because we did not use a single "standard" stimulus intensity in our experiments, it is difficult to directly compare the information values between different units. However, the information transmitted from stimulus position to evoked spike count is shown in relation to stimulus intensity in figure 8.1, while figure 8.2 shows similar graphs for information derived from first-spike latency. As can be seen, information transmitted by evoked spike count is generally lower than

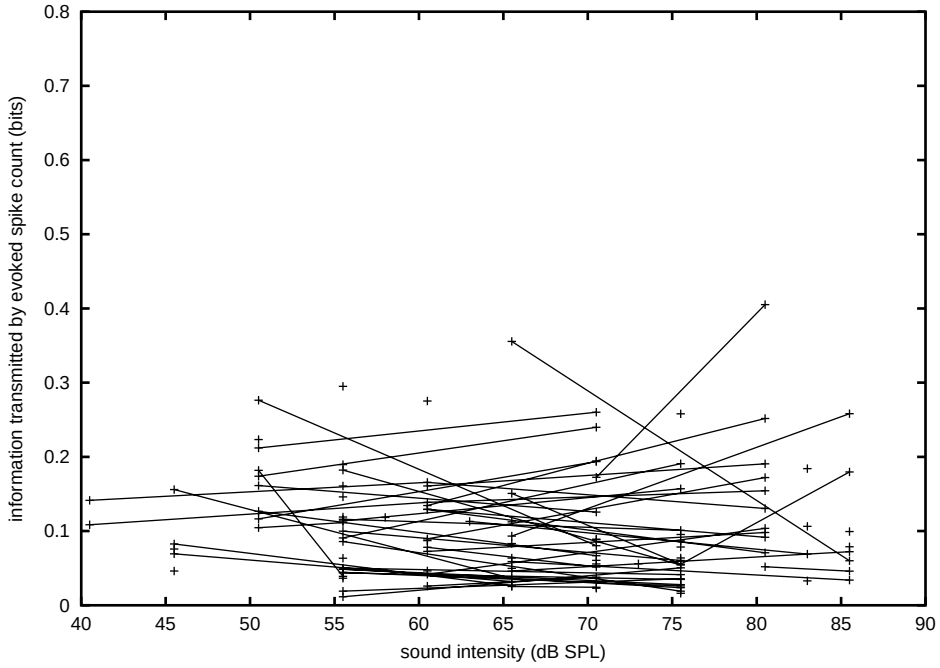

Figure 8.1: The effect of stimulus intensity on the information transmitted by evoked spike count about stimulus position. Each point represents a single azimuth sweep for a single unit. Points joined by lines are measurements made on the same unit at different stimulus intensities. Stimulus duration was 100 ms in all cases.

that transmitted by first-spike latency. By contrast, the resolution with which first-spike latency is measured appears to make little difference to the amount of information transmitted, although the range of information values is perhaps slightly smaller at 5 ms resolution than it is at 10 ms resolution. A direct comparison between the information provided by these different response measures is given in figure 8.3.

8.3.2a Effect of stimulus level

There appears to be no effect of stimulus level in figure 8.1, and no consistent effect in figure 8.2, and this impression is confirmed by quantitative analysis. Although different stimulus levels were used for different units, many units were tested at stimulus levels which were 20 dB apart, and so it is possible to analyse the effects of a 20 dB increase in level. The data from these units ($n = 39$) are shown in figure 8.4; for reasons of clarity, this figure only shows first-spike latency data from measurements at 10 ms resolution.

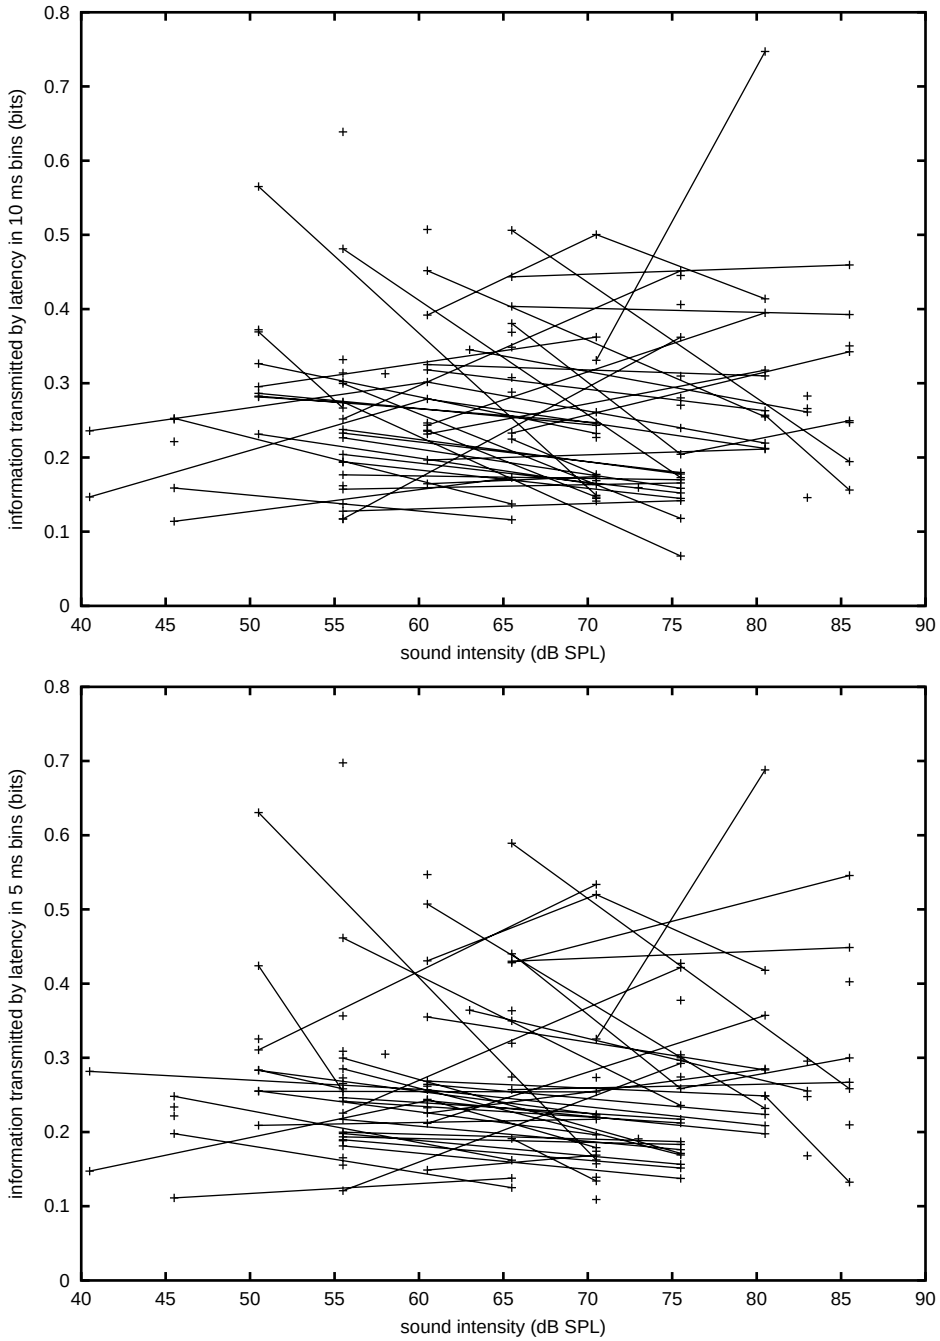

Figure 8.2: The effect of stimulus intensity on the information transmitted by first-spike latency, at resolutions of 10 ms (top) and 5 ms (bottom). Each point represents a single azimuth sweep for a single unit. Points joined by lines are measurements made on the same unit at different stimulus intensities. Stimulus duration was 100 ms in all cases.

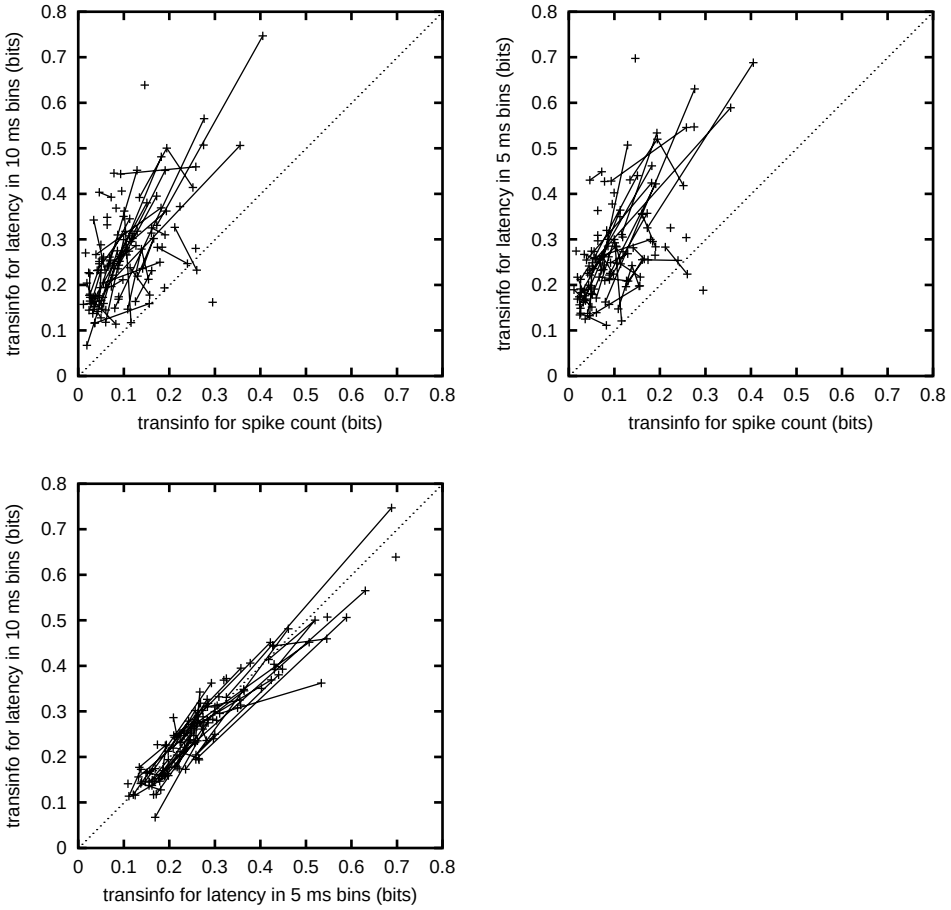

Figure 8.3: Comparisons between the information about stimulus location transmitted by different response measures. Each point represents a single azimuth sweep for a single unit. Points joined by lines show measurements at different stimulus intensities. Stimulus duration was 100 ms in all cases.

Although the figure suggests a small decrease in transmitted information with increasing level, for both evoked spike count and first-spike latency (a mean change of -0.00858 bits and -0.0402 bits, respectively), neither of the differences is significant (paired t-tests: $t = 0.646$ for evoked spike count, giving $p = 0.522$ with a 2-tailed test and 38 d.f.; $t = 1.88$ for first-spike latency, giving $p = 0.0681$ with a 2-tailed test and 38 d.f.).

8.3.2b Effect of stimulus duration

Figure 8.5 shows how values of transmitted information are affected by changes in stimulus duration. As above, only 10 ms resolution first-spike latency data are shown. On average ($n = 30$), transmitted information appears to be slightly higher for 40 ms stimuli than for 100 ms stimuli, but the difference is small (0.0127 bits for evoked spike count; 0.0114 bits for first-spike latency) and not significant (paired t-tests: $t = 0.835$ for evoked spike count, giving $p = 0.411$ with a 2-tailed test and 29 d.f.; $t = 0.606$ for first-spike latency, giving $p = 0.549$ with a 2-tailed test and 29 d.f.).

8.3.3 Information transmitted by pairs of units

We recorded from 25 simultaneously-recorded pairs of units, as described in chapter 6. The information transmitted by the combined evoked spike counts of each pair of units is shown in figure 8.6. What is immediately apparent is that the typical value for transmitted information, which is less than 0.2 bits, is very similar to that for a single unit's evoked spike count. Additionally, the information transmitted by the combined latencies of each pair of units, shown in figure 8.7, is similar to that for the latencies of single units. However, as noted in section 8.2, the methods for deriving information from the responses of pairs of units are different from those used for single units.

In order to make a fair comparison, it is necessary to use the paired-unit analytic methods on the single-unit data. The results are shown in figure 8.8 for evoked spike count, and in figure 8.9 for first-spike latency. As can be seen, the combination of a pair's spike counts yields the same amount of information as the sum of the information transmitted by each unit within the pair. Indeed, the difference between paired-unit information and the sum of single-unit information is not significantly different from zero (paired t-test: $t = 0.515$, giving $p = 0.609$ with a 2-tailed test and 39 d.f.). The same is not true for first-spike latency, however, where the combination of latencies yields slightly less information than the information transmitted by the two units alone. This difference is significant, both for values derived using 10 ms

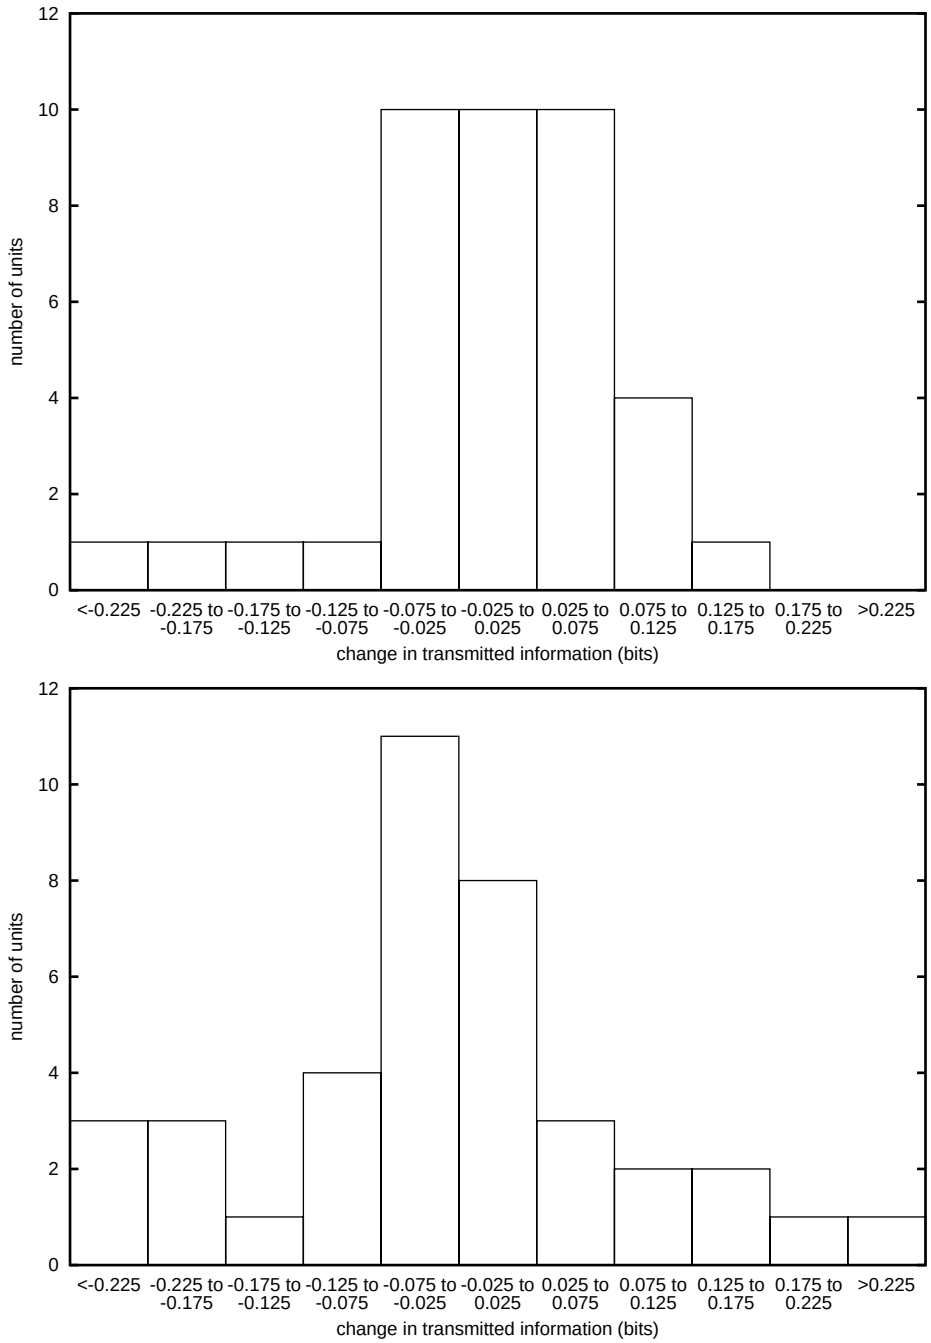

Figure 8.4: The effect of a 20 dB increase in stimulus intensity on the information transmitted between stimulus position and either evoked spike count (top) or first-spike latency (bottom). Only responses to 100 ms stimuli are shown.

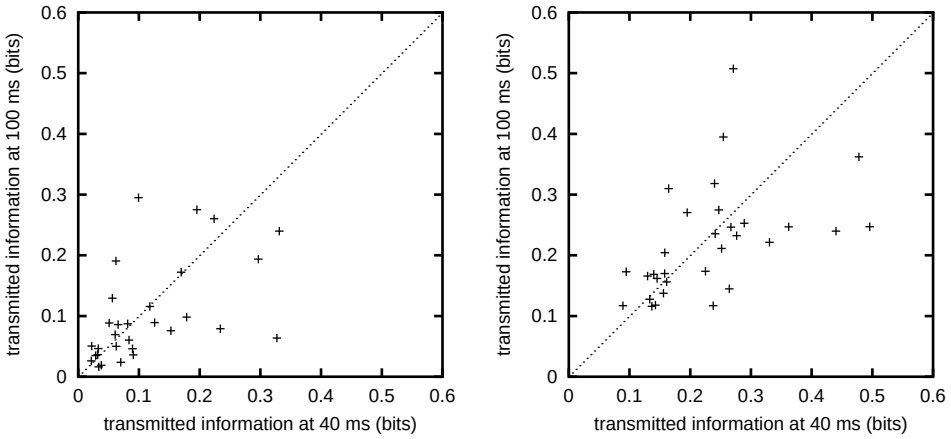

Figure 8.5: The effect of stimulus duration on the information transmitted between stimulus location and evoked spike count (left) and between stimulus location and first-spike latency (right). Each point represents a single unit, with transmitted information for a 40 ms stimulus on the abscissa and transmitted information for a 100 ms stimulus on the ordinate.

resolution (paired t-test: $t = 5.97$, giving $p = 5.67 \times 10^{-7}$ with a 2-tailed test and 39 d.f.) and those using 5 ms resolution (paired t-test: $t = 6.37$, giving $p = 1.57 \times 10^{-7}$ with a 2-tailed test and 39 d.f.).

The results from shifted presentations are also shown, and are very similar to those for simultaneous presentations. For none of the three response measures was there any significant difference between the data from simultaneous and shifted presentations (paired 2-tailed t-tests with 39 d.f.: $t = 0.674$ for evoked spike count, giving $p = 0.504$; $t = 0.236$ for first-spike latency in 10 ms bins, giving $p = 0.815$; $t = 0.257$ for first-spike latency in 5 ms bins, giving $p = 0.799$). This suggests that per-presentation correlations between responses do not carry any useful information.

The effect of stimulus intensity on the information transmitted by a pair of units is difficult to gauge, as the number of pairs with recordings at more than one stimulus intensity is small ($n = 15$). Where multiple intensities had been used, however, they were always 20 dB apart, so the effect of a 20 dB change can be easily assessed. As might be expected from inspection of figure 8.6, the mean effect of such a change is a small increase in transmitted information, of 0.0331 bits, but this is not significant (paired t-test: $t = 1.76$, giving $p = 0.0990$ with 15 d.f.[3] on a 2-tailed test). With first-spike latency

[3]There are 15 degrees of freedom because one of these pairs was recorded at three different stimulus intensities, and thus produced two different 20 dB changes. Using only

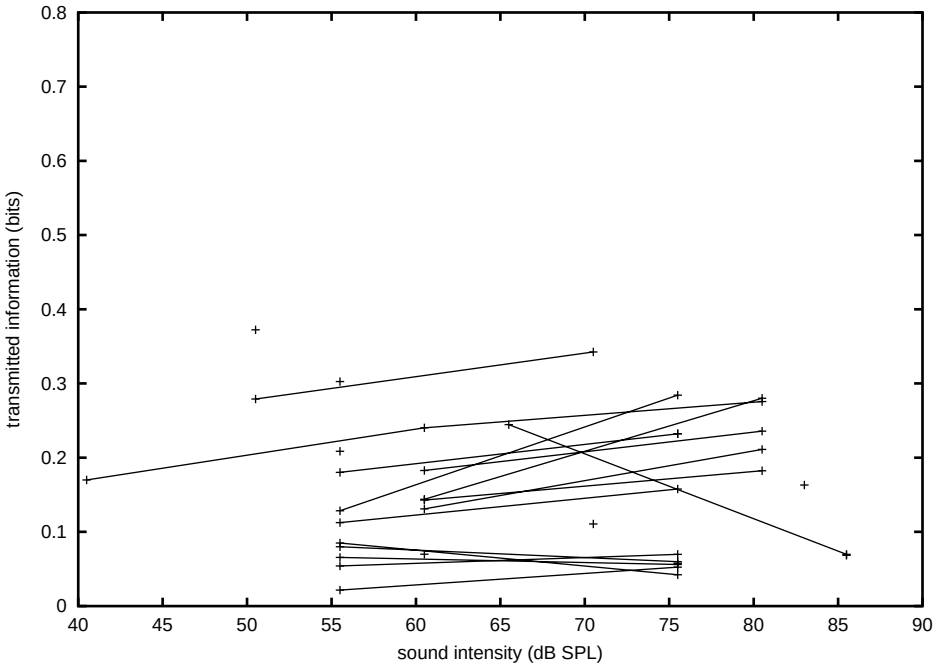

Figure 8.6: The information transmitted from stimulus position to paired-unit evoked spike count, as a function of stimulus intensity. Each point represents a single azimuth sweep for a pair of units. Points joined by lines are measurements made on the same pair at different stimulus intensities. Stimulus duration was 100 ms in all cases.

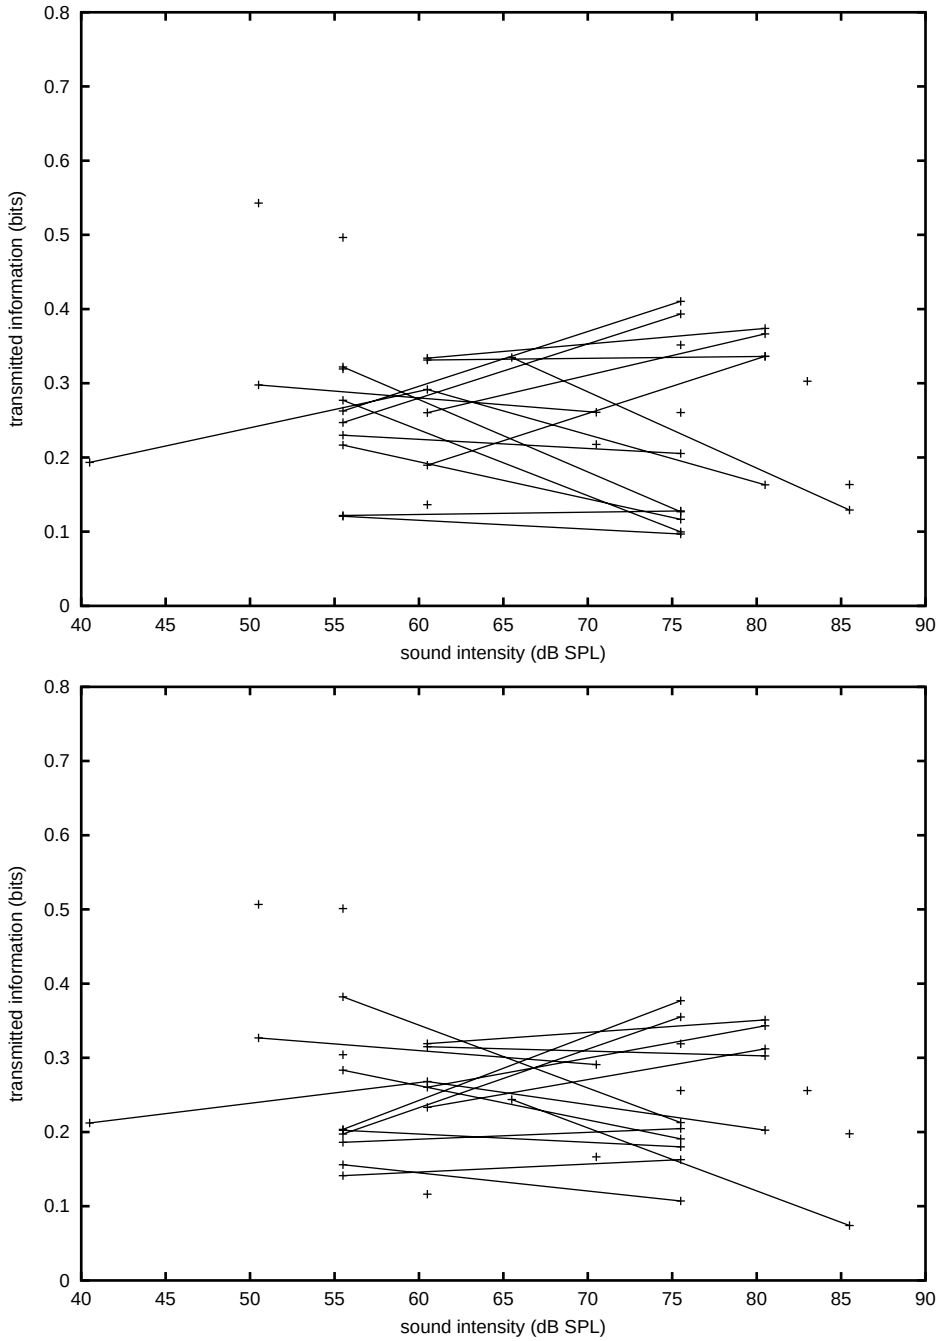

Figure 8.7: The information transmitted from stimulus position to paired-unit first-spike latency, as a function of stimulus intensity. Latency resolutions of 10 ms (top) and 5 ms (bottom) are shown. Each point represents a single azimuth sweep for a pair of units. Points joined by lines are measurements made on the same pair at different stimulus intensities. Stimulus duration was 100 ms in all cases.

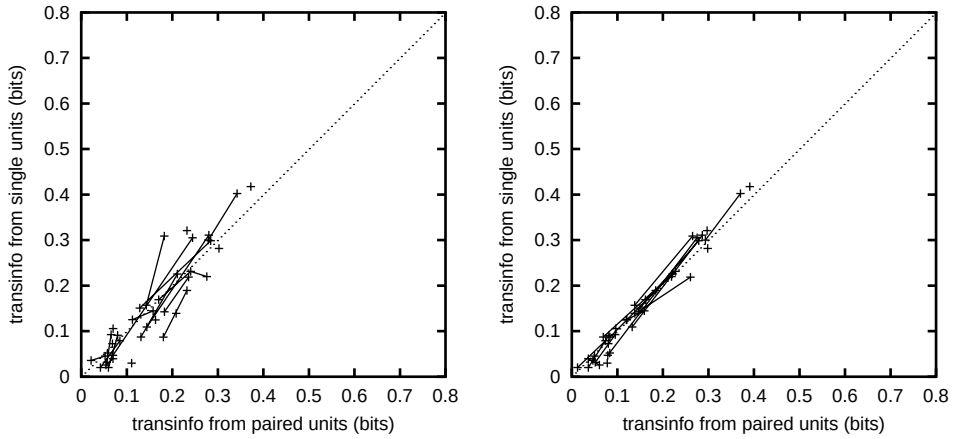

Figure 8.8: The information about stimulus location transmitted by paired-unit spike counts, *versus* the sum of the information transmitted by single-unit counts. Results are shown for both simultaneous (left) and shifted (right) presentations. Each point represents a single azimuth sweep for a pair of units. Points joined by lines show measurements made on the same pair at different stimulus intensities. Stimulus duration was 100 ms in all cases.

the differences are even smaller, as shown in figure 8.7, and are also not significant (paired t-test: $t = 0.398$, giving $p = 0.696$ with 15 d.f. on a 2-tailed test); for simplicity only data from 10 ms binning are shown.

For completeness, the effect of stimulus duration on the information transmitted by pairs of units is also shown, in figure 8.10 (as above, only 10 ms resolution first-spike data are shown). Unfortunately, the number of pairs involved here is even smaller ($n = 12$) and the large variation in information values makes it difficult to draw any conclusions from this figure. There is certainly no significant difference between the 40 ms and 100 ms values (paired t-test: $t = 1.04$ for evoked spike count, giving $p = 0.319$ with 11 d.f. on a 2-tailed test; $t = 0.179$ for first-spike latency, giving $p = 0.861$ with 11 d.f. on a 2-tailed test).

8.4 Discussion

In this study, we have used the techniques of information theory to analyse the results of electrophysiological recordings in A1. While we have made several important observations, the use of information theory in this way

one value per pair makes no difference to the significance of the test.

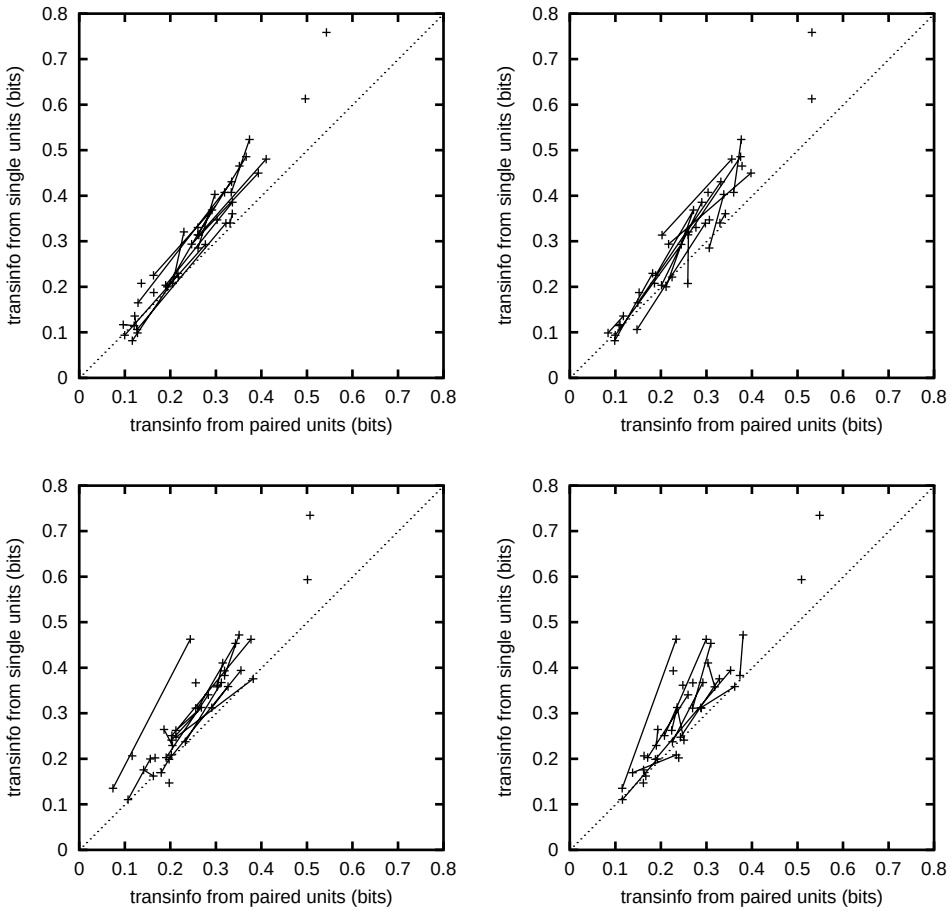

Figure 8.9: The information about stimulus location transmitted by paired-unit first-spike latencies, *versus* the sum of the information transmitted by single-unit latencies. Results are shown for both simultaneous (left column) and shifted (right column) presentations, and for both 10 ms (top row) and 5 ms (bottom row) resolutions. Each point represents a single azimuth sweep for a pair of units. Points joined by lines show measurements made on the same pair at different stimulus intensities. Stimulus duration was 100 ms in all cases.

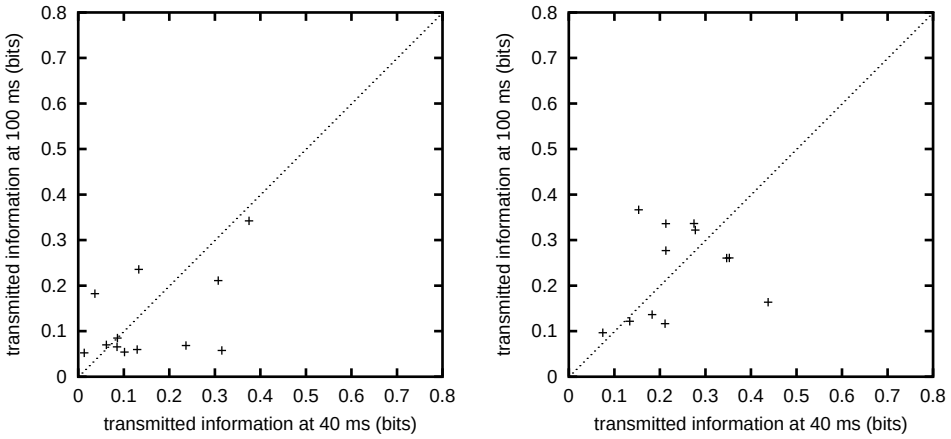

Figure 8.10: The effect of stimulus duration on the information transmitted between stimulus location and either paired-unit evoked spike count (left) or paired-unit first-spike latency (right). Each point represents a pair of units, with transmitted information for a 40 ms stimulus on the abscissa and transmitted information for a 100 ms stimulus on the ordinate.

relies on key assumptions which require justification.

The main problem with information theory as we have applied it is that information values depend to a certain extent on the binning procedure used. Changing the binning of, say, evoked spike count would have an effect on the amount of response information, and this could in turn have a knock-on effect on the amount of information which is transmitted between stimulus and response. In particular, our decision to use the same binning procedure for every unit may not be appropriate: on the one hand, it allows us to make a straightforward comparison between results from different units, but some units may have different ranges of responses than others, and by using a "one size fits all" binning procedure we may be underestimating information. It should be mentioned that "bin-less" methods of deriving measures of information are not free of this problem. Binless methods effectively work by fitting some continuous function to the recording data, and the choice of such a function will have as much effect as the choice of a particular binning method.

Neither are problems of binning restricted to measures of response. Although we have not used a stimulus binning procedure as such, simply choosing a stimulus set has much the same effect on transmitted information: if we had used a different number of stimuli, this would obviously have had an effect on stimulus information, and thus on transmitted information. Choosing

more or less discriminable stimuli would also have had an effect on transmitted information. In this regard it is important to note that we have split up our stimuli into azimuth sweeps—recordings made using particular values of stimulus intensity and duration—and measured a separate value of transmitted information for each one. This was done in order to allow comparisons between different intensities and durations. However, given that one of the perceived problems with location processing in A1 is that it is confounded with intensity processing, it is controversial to measure how A1 distinguishes stimuli which differ in azimuth but not intensity. It would be possible to combine all the different azimuth sweeps that were used for each unit, and derive a single transmitted information value from this. The problem is that the resulting number would not in itself tell us very much—we are more interested in comparing different information values with one another than we are in the absolute values themselves.

A slightly different problem is that single-cell recording may overestimate the unreliability of the processes that we are measuring. Extracellular recording introduces various problems—electrical noise, blood-flow artefacts, cross-talk from other units—that are presumably not relevant to the nervous system itself. While we did take care to exclude very noisy recordings from our analysis, it remains possible that "true" information values for the nervous system may be higher than the results we present here.

The most oft-made comparison made in information theory studies is between the information transmitted by evoked spike count and the information transmitted by some measure of spike "timing"[88, 114]. This is a rather loose distinction—spike count is in some senses a measure of spike timing, as it assesses the number of spikes occurring in a particular time window after stimulus presentation, but the general implication is that a measure on a finer time-scale might be more informative.

In this study, we compare evoked spike count with first-spike latency, finding that first-spike latency carries significantly more information than evoked spike count. Although this conclusion is subject to the dependence on response binning that we have already mentioned, the effect of altering the binning of first-spike latency is surprisingly small. It is also difficult to imagine how evoked spike count could be binned in any other way, as spike count is intrinsically an integer[4] and we have used bins which are 1 spike wide. One possible explanation is that latency has a higher information capacity than spike count. After all, many units have a maximum spike count of only 2 spikes, and so for these units there are only three different

[4]Although, as we have tried to correct for "spontaneous" or background activity, the resulting values of evoked spike count are not integral.

possible responses to a stimulus: no spikes, one spike or two spikes. First-spike latency, on the other hand, is much more variable. However, even a response measure with only three outcomes can potentially transmit more than 1 bit of information, while the actual information transmitted about stimulus position by spike count is typically less than 0.2 bits—it does not seem as though the information capacity of the different response measures is the limiting factor here.

Thus the advantage of spike latency over spike count appears to be quite robust, and indeed our results agree with those from studies in other brain areas[37, 89]. First-spike latency is a rather unusual response measure, in that it is not itself a plausible neural code. This is because the brain does not know when stimulus onset is, and so has no access to absolute values of latency. However, it is often assumed that the brain does have access to relative latencies—either relative to the latencies of other neurons, or relative to previous "spontaneous" spikes that have occurred in the same unit. In the latter case, first-spike latency could be regarded as a reflection of a code such as inter-spike interval. What this result therefore suggests is that the timing of individual spikes on a scale of 5–10 ms (the resolution of latency binning) is more informative than the timing of spikes on a scale of 50–100 ms (the size of the response window in spike counting). The potential advantage of this coding mechanism is that the next stage of processing does not have to wait and see how many spikes will be fired over a long time period, but can respond more quickly.

While changing between response measures has a significant effect on the information carried about stimulus location, changes in the stimulus itself do not. We find that neither stimulus duration nor stimulus level affect the amount of information transmitted about stimulus location, irrespective of whether evoked spike count or first-spike latency is used as the measure of response. Unlike the comparison between spike count and latency, this is less subject to problems in interpreting neural responses.

This finding is interesting as increases in both duration and level appear to increase the breadth of neural tuning to location (chapter 6). This broadening of tuning would typically be used to conclude that less information about sound source information is available at higher stimulus levels/durations, because the shape of the receptive field is more like that of an omni-directional cell. This shows how intuitions based on measures of bandwidth fail to take into account the discriminability of the response at different positions, and reflect the overall shape of the response curve rather than its reliability. Indeed, it is quite plausible that evoked spike count is more reliable at higher stimulus intensities, as the signal-to-noise ratio will improve.

However, our observations on the effects of changing stimulus levels may

not only apply outside the range of levels that we used. We have deliberately restricted ourselves to supra-threshold stimulus intensities, and it is possible that near-threshold intensities would give a different result. This is because such stimuli may be so attenuated by the head that they become effectively monaural, and thus provide fewer location cues to the auditory system. Indeed, studies where near-threshold stimuli are included show seemingly larger changes in receptive field properties than those described here[12].

Our other main interest in this study is the relationship between the information transmitted by single units and that transmitted by multiple units. This is because the information transmitted by individual units is very low, certainly less than the amount of information required to uniquely specify 1 out of 17 stimulus positions, which is $\log_2(17) \approx 4.09$. We have already discussed how it is dangerous to read too much into absolute values of transmitted information. However, it is clear that if A1 does code for sound source location in any useful manner, it must do so using a multi-neuron code.

In the case of evoked spike count, the information about location transmitted by a pair of units is the same as if the two units' firing was independent. In the case of first-spike latency, considering two units together gave slightly less information than considering them separately, i.e. the information provided by the two latencies was slightly redundant. The small redundancy is not surprising—a certain amount of redundancy is to be expected where the tuning of two neurons is similar[95]. As noted in chapter 6, most units respond preferentially to the contralateral side of space, so to that extent their responses are similar and provide slightly redundant information. What is surprising is that there is redundancy for first-spike latency and not for evoked spike count. This suggests that the similarity in stimulus-related responses between units is greater on short time scales than on long time scales. However, a potential problem with this entire analysis is that it is not known for certain that the accuracy of estimating transmitted information—particularly correcting for bias—is the same for the large matrices produced by paired responses as it is for the small matrices of individual neurons.

A less problematic result emerges from mismatching the responses of a pair of units, so that one unit's response to a stimulus is paired with the "wrong" response of the other unit (albeit to the "right" stimulus). This treatment does not affect the amount of information transmitted, despite the fact that there are correlations between units that are broken up by this "shifting" (chapter 6). This shows that such correlations do not carry any useful information about stimulus location, a conclusion which contradicts some theoretical models in which assemblies of cortical units signal stimulus properties by varying the correlation of their responses. Having said that,

most of our pairs of units were some distance apart within A1, so would possibly not be expected to be involved in the same assemblies in any case.

Part IV

Conclusions

Chapter 9

Conclusions

This study has been aimed at elucidating the function of A1 in auditory processing: in particular, assessing whether A1 can be regarded as being specialised for the processing of sound source location.

We found that unilateral inactivation of A1 in animals which are experienced on a localisation task caused a deficit on that task, and that this deficit was limited to the side of space opposite to the inactivated area. Similar unilateral A1 inactivation in untrained animals caused no notable impairment, but bilateral A1 inactivation in untrained animals did cause a very obvious impairment. The number of front-back errors—responses which were on the correct side of the midline but the wrong side of the interaural axis—seemed to be a particularly sensitive indicator of A1 inactivation. In all cases, the impairments that we saw were only for short-duration stimuli (40 ms), and these animals still performed at well above chance level. While this definitely suggests that A1 is involved in our sound localisation task, it is far from implying that all aspects of sound localisation require an intact A1.

One drawback of our behavioural work is that, where we unilaterally inactivated A1 in our behavioural studies, it was always left A1 which was inactivated. While there are no reported differences between the effects of a right A1 lesion as compared to a left A1 lesion in the ferret[59], in human studies it has been suggested that right A1 lesions cause a more profound sound localisation impairment than do left A1 lesions, with right-side lesions affecting both sides of space rather than just the contralateral side[148]. It would thus be valuable to examine the effect of right A1 inactivation in future experiments. Another technical change which could benefit our behavioural apparatus would be the addition of more speakers, as it is possible that using speakers which are closer together could make the task a more sensitive indicator of sound localisation ability, although it might also make it more difficult for normal animals. Having said that, the speaker separation used

here was comparable to that in previous studies of A1 lesions[56].

In fact, our behavioural techniques do offer several advantages over previous studies. Most obvious is the use of stimulus and response positions which range over 360° in azimuth—without this we would not have been able to observe the distribution of front-back errors as we did in this study. Equally important, though, was the use of muscimol Elvax to inactivate A1, instead of the more traditional aspiration lesions. Aspiration lesions are known to have many side-effects, including incidental damage to fibres of passage, white matter tracts and afferent axons, but also more complicated metabolic changes[141]. In the case of A1, lesions are known to cause retrograde degeneration in the ventral division of the MGN[59], an area which provides input to many cortical areas (see chapter 1) and so this may itself have had behavioural consequences. Muscimol Elvax implants have their own set of drawbacks—particularly the problem of precisely targetting an anatomical area—but their effects should be much more selective. In this regard, the apparent mildness of the impairment by contrast to previous studies may reflect previous over-estimation of the involvement of A1 in sound localisation.

In a parallel set of studies on the properties of A1 neurons, we found that individual neurons were quite insensitive to sound location. Not only that, but the location sensitivity of these neurons could alter with changes in either stimulus intensity or stimulus duration. It is possible to measure the amount of location-related information carried by each cell, using techniques from a branch of mathematics known as information theory. An information theoretic analysis showed that the amount of information carried by each cell, while small, was not affected by changes in stimulus intensity or duration. Also, slightly more information was available from the timing of the first post-stimulus spike than was available from the number of evoked spikes. The information contributed by different cells was only redundant to a limited extent, so it is likely that the auditory system could gain large amounts of information about stimulus location by comparing the responses of many neurons.

While the absence of an effect of intensity or duration suggests that neurons' information-carrying is robust, it seems to sit oddly with the behavioural data. The performance of normal ferrets on the localisation task is believed to be similar at moderately different sound levels[91], but we have found that it is very different at different durations. Of course, A1 inactivation only causes an impairment at short durations, so it could be argued that the duration-sensitivity is, like the performance at long durations, down to some other part of the brain. If true, this is quite different from claiming that A1 is an area specialised for sound localisation—it would perhaps be

more accurate to describe it as being specialised for short-duration stimuli, or for temporal processing on a timescale of tens of milliseconds. This is an interpretation which is supported by the association between A1 damage and impaired gap-detection in ferrets[64], as well as the disorder known as "pure word deafness" in humans[6, 97]. An alternative is to suggest that, where stimuli are of long duration, a conscious, behaving ferret has access to cues which A1 in the anesthetised ferret does not—the obvious example being cues that arise from head movements.

As with the behavioural work, though, our electrophysiological methods could use some refinement. The choice of isoflurane as anaesthetic was possibly a poor one, as it may have been the cause of our poor pure-tone responses. An agent such as halothane or α-chloralose might be more appropriate, although possibly not barbiturate as there are suggestions that latency coding is not seen to any large extent in barbiturate anaesthesia. Further, the use of a single loudspeaker on a remotely-controlled hoop proved to be a serious limiting factor on the number of stimuli we could present, due to the slow movement of the hoop. The number of stimulus presentations was in turn the limiting factor in our simple application of information theory. It is possible to get around this limitation, with more complicated mathematical techniques such as the Nelken table-reduction method described in chapter 8, or Panzeri's time series expansion[89], but this is not ideal as these techniques may have their own problems. As well as influencing the number of stimuli that could be presented, use of the hoop also forced us to use our stimuli in a particular order, with stimuli being presented one after another at each position, before moving to an adjacent position. This order has the potential to lead to a significant amount of neuronal adaptation or habituation, both in the responses to a particular position and across the whole range of azimuths. If this is the case, it might undermine our comparison of spike count and latency, as there is no guarantee that adaptation would affect them both in the same way. It would be much better to present stimuli in a pseudo-random order, but this is only possible with either a large array of speakers—which might adversely affect the anechoic nature of the chamber—or else the use of virtual auditory space stimuli.

More important than the practical issues described above, the main problem with this type of study is that it is difficult to distinguish the processing that is intrinsic to A1 from the processing going in areas to which A1 is connected. Areas of the brain do not act in isolation: even ignoring possible accidental damage in the case of aspiration lesions, the inactivation of a richly-connected part of a neural network will have an effect on activity throughout the network[146]. The converse is also true: a richly-connected area will be affected by lesions or inactivation in a large number of different

areas. The most obvious example of this is a disconnection effect—where the inactivation of an area has an effect by robbing a downstream area of some relevant sensory input. To use an analogy described elsewhere (see section 1.2.2), cutting a telephone's cable will certainly prevent the telephone from converting electrical signals into speech, but that does not mean that it is the wire that is responsible for the conversion, and we know that it is in fact the handset.

Thus it is possible that A1 is not actually processing sound source location at all, and the consequences of A1 inactivation could reflect disruption of activity in an area to which A1 projects. This is plausible given A1's apparent position as the gateway into cortex, as well as the observation that, in the macaque, several areas downstream of A1 contain neurons which are more sharply tuned to stimulus position than A1 itself[112, 106]. Similarly, just because we have shown electrophysiologically that A1 neurons are carrying information about sound source location, it does not mean that this information is being processed within A1.

The most convincing evidence that A1 is processing something directly related to a sound's location is the fact that the deficit caused by unilateral A1 inactivation is limited to the contralateral hemifield. Whatever it is that A1 is processing must therefore be associated in some way with that particular area in space. This ties in with the observation that most A1 neurons respond more vigorously to the contralateral side of space. It may also be related to the increased number of front-back errors in A1-inactivated animals, as this could reflect an impaired ability to locate sound sources within a hemifield. Furthermore, sound localisation is the only task which is known to be impaired in animals with unilateral rather than bilateral lesions, although this may not be true in humans[147]. Notwithstanding the possibility that A1 in the two hemispheres act as a single processing unit, this shows that sound source location is an important part of the A1's functional organisation. However, it remains difficult to describe A1 as an essential part of processing sound localisation when the impairments produced by inactivation are small and restricted to short-duration stimuli.

Another possibility, given that sound localisation is a large and complex task, is that A1 may be involved in only a small part of the necessary processing. If we assume that A1 is not processing sound source location *per se*, but is instead processing one or more cues for sound localisation, then it is important to consider that different cues may be processed at different levels of the system. As mentioned in chapter 1, interaural cues can be calculated on a frequency-by-frequency basis, while spectral cues have to be examined in association with a certain amount of auditory scene analysis, and integrating the cue types together also requires sophisticated scene analysis. Interaural

time and difference cues are initially processed at quite an early stage in the auditory system—the superior olivary complex of the brainstem—while there is some evidence that spectral cues may be first processed as early as the cochlear nucleus[77]. This suggests that A1 is involved in the subsequent integration of the different cue types and/or scene analysis.

Some sort of high-level role in the integration of auditory cues is consistent with the observed impairment in pitch discrimination of both animals[138] and humans[147] with A1 lesions. This is because computing pitch requires grouping harmonic frequency components together. The increased number of front-back errors in A1 inactivation could then represent a failure to integrate localisation cues across frequencies, because it is simple frequency-by-frequency ITDs/ILDs which are approximately symmetric about the interaural axis. While it could be argued that A1 neurons are sensitive only to a limited range of frequencies, and must therefore be processing stimuli on a frequency-by-frequency basis themselves, there is an increasing amount of evidence that the activity of these neurons can be influenced by stimulus components which lie outside the "classical" frequency response area[8, 120]. Indeed, if the integration of frequency components is part of a temporal analysis—of "periodicity"[119]—then this would be consistent with the temporal processing impairments seen with bilateral A1 lesions. This does not explain why an A1 lesion results in impaired localisation performance when all the targets are along the horizon in the front hemifield[56], as it might be expected that subcortical ITD/ILD processing would be entirely adequate to carry out this task. It is possible that in this case the impairment is just the result of a disconnection syndrome, because performance is not impaired on a similar task with a different, possibly non-cortical, motor response[76]. This suggests an arrangement where some localisation cues are actively processed by A1, and some others are simply passed on to other auditory areas.

The results of behavioural studies to date cannot convincingly be used to identify a particular aspect of localisation processing in which A1 might be involved. This is because none of the tasks that have been used require much in the way of complicated scene analysis. Although the localisation of broadband noise stimuli potentially requires some comparison between the different components of the noise in order to score optimally, in the absence of any competing sound sources it might easily be possible for an animal to get by without this. Thus I believe it is important for future studies to examine the localisation performance of A1-inactivated animals in the presence of masking or competing sound sources. It might also be possible to examine other tasks, such as the use of harmonic structure to distinguish sound sources, as this also relies on matching between frequency components.

As regards future electrophysiological research, an increase in the un-

derstanding of when and where non-linearity occurs in the responses of A1 neurons seems likely to shed light on its function. This is not because there is anything intrinsically important about non-linearity, but because it seems that it is the increasing amounts of non-linearity which distinguish cortical from subcortical auditory processing. Indeed, in Marr's influential description of visual processing at the "algorithmic" level, simple linear filtering—analogous to the view of an auditory neuron as a spectro-temporal filter—was the very first step in processing, followed by the detection of orientated "edges", and then trying to gradually link together the result into simple features[74]. If such a framework is applied to the auditory system, A1's position in the processing hierarchy would suggest that the stages of simple linear filtering have already been accomplished, and that A1 is trying to link together basic "primitives" into larger features. The fact that there are many linear aspects to the properties of A1 neurons does not undermine this interpretation—the properties of a given neuron do not solely represent its own processing, after all, but also the processing of every prior stage in the chain.

In conclusion, there is some evidence to support the involvement of A1 in processing sound localisation cues, but little to suggest that sound localisation should be regarded as its primary function. It is certainly not the only function which is impaired in the case of bilateral A1 damage. That it is the main casualty of unilateral A1 damage may reflect the fact that each cortex primarily processes information about sounds in the contralateral side of space. In any case, a sound localisation task actually involves several layers of processing, including several stages in which cues pertaining to the same stimulus must be correctly matched to one another. This is a general problem in auditory processing, and few studies to date have attempted to detail how it might be solved. However, we tentatively suggest that some of the relevant processing is both sufficiently general and sufficiently high-level that an area like A1 might be involved.

Bibliography

[1] L. M. Aitkin, J. D. Pettigrew, M. B. Calford, S. C. Phillips, and L. Z. Wise. Representation of stimulus azimuth by low-frequency neurons in inferior colliculus of the cat. *Journal of Neurophysiology*, 53(1):43–59, 1985.

[2] L. M. Aitkin, J. D. Pettigrew, M. B. Calford, S. C. Phillips, and L. Z. Wise. The representation of stimulus azimuth by high best-frequency azimuth-selective neurons in the central nucleus of the inferior colliculus of the cat. *Journal of Neurophysiology*, 57(4):1185–1200, 1987.

[3] R. A. Andersen, P. L. Knight, and M. M. Merzenich. The thalamocortical and corticothalamic connections of AI, AII, and the anterior auditory field (AAF) in the cat: evidence for two largely segregated systems of connections. *Journal of Comparative Neurology*, 194(3):663–701, 1980.

[4] C. S. J. Annetts, B. P. A. Howse, M. E. Hutchings, and A. J. King. A digitally controlled vertical hoop system for the presentation of free-field auditory and visual stimuli. *Journal of Physiology*, 382:21, 1987.

[5] F. Attneave. *Applications of information theory to psychology: a summary of basic concepts, methods, and results.* Henry Holt and company, New York, 1959.

[6] S. H. Auerbach, T. Allard, M. Naeser, M. P. Alexander, and M. L. Albert. Pure word deafness. Analysis of a case with bilateral lesions and a defect at the prephonemic level. *Brain*, 105:271–300, 1982.

[7] P. Azzopardi and A. Cowey. Blindsight and visual awareness. *Consciousness and Cognition*, 7(3):292–311, 1998.

[8] O. Bar-Yosef, Y. Rotman, and I. Nelken. Responses of neurons in cat primary auditory cortex to bird chirps: effects of temporal and spectral context. *Journal of Neuroscience*, 22(19):8619–8632, 2002.

[9] A. S. Bregman. *Auditory scene analysis: the perceptual organization of sound*. MIT Press, Cambridge, MA, 1990.

[10] A. W. Bronkhorst and T. Houtgast. Auditory distance perception in rooms. *Nature*, 397(6719):517–520, 1999.

[11] J. F. Brugge and M. M. Merzenich. Responses in neurons in auditory cortex of the macaque monkey to monaural and binaural stimulation. *Journal of Neurophysiology*, 36:1138–1158, 1973.

[12] J. F. Brugge, R. A. Reale, and J. E. Hind. The structure of spatial receptive fields of neurons in primary auditory cortex of the cat. *Journal of Neuroscience*, 16(14):4420–37, 1996.

[13] J. F. Brugge, R. A. Reale, J. E. Hind, J. C. Chan, A. D. Musicant, and P. W. Poon. Simulation of free-field sound sources and its application to studies of cortical mechanisms of sound localization in the cat. *Hearing Research*, 73(1):67–84, 1994.

[14] D. S. Brungart, N. I. Durlach, and W. M. Rabinowitz. Auditory localization of nearby sources. II. Localization of a broadband source. *Journal of the Acoustical Society of America*, 106(4 Pt 1):1956–1968, 1999.

[15] R. A. Butler. The bandwidth effect on monaural and binaural localization. *Hearing Research*, 21:67–73, 1986.

[16] R. A. Butler. An analysis of the monaural displacement of sound in space. *Perception and Psychophysics*, 41:1–7, 1987.

[17] R. A. Butler, I. T. Diamond, and W. D. Neff. Role of auditory cortex in discrimination of changes in frequency. *Journal of Neurophysiology*, 20:108–120, 1957.

[18] S. Carlile. The auditory periphery of the ferret. II. The spectral transformations of the external ear and their implications for sound localisation. *Journal of the Acoustical Society of America*, 88:2196–2204, 1990.

[19] S. Carlile and A. J. King. Monaural and binaural spectrum level cues in the ferret: acoustics and the neural representation of auditory space. *Journal of Neurophysiology*, 71(2):785–801, 1994.

[20] S. W. Cheung, S. S. Nagarajan, P. H. Bedenbaugh, C. E. Schreiner, X. Wang, and A. Wong. Auditory cortical neuron response differences under isoflurane versus pentobarbital anesthesia. *Hearing Research*, 156(1–2):115–127, 2001.

[21] J. C. Clarey, P. Barone, and T. J. Imig. Physiology of thalamus and cortex. In A. N. Popper and R. R. Fay, editors, *The Mammalian Auditory Pathway: Neurophysiology*, chapter 5, pages 232–334. Springer-Verlag, New York, 1992.

[22] J. C. Clarey, P. Barone, and T. J. Imig. Functional organization of sound direction and sound pressure level in primary auditory cortex of the cat. *Journal of Neurophysiology*, 72:2383–2405, 1994.

[23] J. C. Clarey, P. Barone, W. A. Irons, F. K. Samson, and T. J. Imig. Comparison of noise and tone azimuth tuning of neurons in cat primary auditory cortex and medical geniculate body. *Journal of Neurophysiology*, 74(3):961–80, 1995.

[24] S. Clarke, A. Bellmann, R. A. Meuli, G. Assal, and A. J. Steck. Auditory agnosia and auditory spatial deficits following left hemisphere lesions: evidence for distinct processing pathways. *Neuropsychologia*, 38(6):797–807, 2000.

[25] R. A. Code and J. A. Winer. Commissural neurons in layer III of cat primary auditory cortex (AI): pyramidal and non-pyramidal cell input. *Journal of Comparative Neurology*, 242(4):484–510, 1985.

[26] H. B. Coslett, H. R. Brashear, and K. M. Heilman. Pure word deafness after bilateral primary auditory cortex infarcts. *Neurology*, 34:347–352, 1984.

[27] A. Cowey and V. Walsh. Tickling the brain: studying visual sensation, perception and cognition by transcranial magnetic stimulation. *Progess in Brain Research*, 134:411–425, 2001.

[28] J. L. Cranford. Role of neocortex in binaural hearing in the cat. I. Contralateral masking. *Brain Research*, 100(2):395–406, 1975.

[29] J. L. Cranford, M. Igarashi, and J. H. Stramler. Effect of auditory neocortex ablation on identification of click rates in cats. *Brain Research*, 116:69–81, 1976.

[30] P. Dallos, A. N. Popper, and R. R. Fay, editors. *The Cochlea*. Springer-Verlag, New York, 1996.

[31] H. Damasio and A. R. Damasio. *Lesion analysis in neuropsychology*. Oxford University Press, 1989.

[32] F. de Ribaupierre, M. H. Goldstein, Jr., and G. Yeni-Komshian. Cortical coding of repetitive acoustic pulses. *Brain Research*, 48:185–204, 1972.

[33] R. C. deCharms, D. T. Blake, and M. M. Merzenich. Optimizing sound features for cortical neurons. *Science*, 280:1439–1443, 1998.

[34] I. T. Diamond and W. D. Neff. Ablation of temporal cortex and discrimination of auditory patterns. *Journal of Neurophysiology*, 20:300–315, 1957.

[35] L. M. Eisenman. Neural encoding of sound location: an electrophysiological study in auditory cortex (AI) of the cat using free field stimuli. *Brain Research*, 75:203–214, 1974.

[36] Y. I. Fishman, D. H. Reser, J. C. Arezzo, and M. Steinschneider. Pitch vs. spectral encoding of harmonic complex tones in primary auditory cortex of the awake monkey. *Brain Research*, 786:18–30, 1998.

[37] S. Furukawa and J. C. Middlebrooks. Cortical representation of auditory space: Information-bearing features of spike patterns. *Journal of Neurophysiology*, 87:1749–1762, 2002.

[38] S. Furukawa, L. Xu, and J. C. Middlebrooks. Coding of sound-source location by ensembles of cortical neurons. *Journal of Neuroscience*, 20(3):1216–1228, 2000.

[39] J. M. Goldberg and W. D. Neff. Frequency discrimination after bilateral section of the brachium of the inferior colliculus. *Journal of Comparative Neurology*, 116:265–290, 1961.

[40] H. E. Heffner and R. B. Masterton. Contribution of auditory cortex to sound localization in monkey (*Macaca mulatta*). *Journal of Neurophysiology*, 38:1340–1358, 1975.

[41] R. S. Heffner and H. E. Heffner. Sound localization acuity in the cat: effect of azimuth, signal duration, and test procedure. *Hearing Research*, 36(2–3):221–32, 1988.

[42] P. Heil. Auditory cortical onset responses revisited. I. First-spike timing. *Journal of Neurophysiology*, 77:2616–2641, 1997.

[43] P. Heil. Auditory cortical onset responses revisited. II. Response strength. *Journal of Neurophysiology*, 77:2641–2660, 1997.

[44] P. Heil and D. R. F. Irvine. First-spike timing of auditory-nerve fibers and comparison with auditory cortex. *Journal of Neurophysiology*, 78:2438–2454, 1997.

[45] P. Heil, R. Rajan, and D. R. Irvine. Sensitivity of neurons in cat primary auditory cortex to tones and frequency-modulated stimuli. II: Organization of response properties along the 'isofrequency' dimension. *Hearing Research*, 63(1–2):135–56, 1992.

[46] P. Heil, R. Rajan, and D. R. Irvine. Topographic representation of tone intensity along the isofrequency axis of cat primary auditory cortex. *Hearing Research*, 76(1–2):188–202, 1994.

[47] C. C. Hilgetag, G. A. Burns, M. A. O'Neill, J. W. Scannell, and M. P. Young. Anatomical connectivity defines the organization of clusters of cortical areas in the macaque monkey and the cat. *Philosophical Transactions of the Royal Society of London B*, 355:91–110, 2000.

[48] A. Y. Huang and B. J. May. Sound orientation behavior in cats. II. Mid-frequency spectral cues for sound localization. *Journal of the Acoustical Society of America*, 100(2 Pt 1):1070–80, 1996.

[49] T. J. Imig, W. A. Irons, II, and F. A. Samson. Single-unit selectivity to azimuthal direction and sound pressure level of noise bursts in cat high-frequency primary auditory cortex. *Journal of Neurophysiology*, 63:1448–1466, 1990.

[50] T. J. Imig and A. Morel. Topographic and cytoarchitectonic organization of thalamic neurons related to their targets in low-, middle-, and high-frequency representations in cat auditory cortex. *Journal of Comparative Neurology*, 227(4):511–539, 1984.

[51] W. A. Irons, II. *Directional selectivity of single units in the auditory thalamus of cats.* PhD thesis, University of Kansas, 1989.

[52] D. R. F. Irvine. Physiology of the auditory brainstem. In A. N. Popper and R. R. Fay, editors, *The Mammalian Auditory Pathway: Neurophysiology*, chapter 4, pages 153–231. Springer-Verlag, New York, 1992.

[53] R. L. Jenison. Models of sound location acuity with broad cortical spatial receptive fields. In J. F. Brugge and P. W. Poon, editors, *Central Auditory Processing and Neural Modeling*, pages 161–174. Plenum Press, New York, 1998.

[54] R. L. Jenison. Correlated cortical populations can enhance sound localization performance. *Journal of the Acoustical Society of America*, 107(1):414–421, 2000.

[55] W. M. Jenkins and R. B. Masterton. Sound localization: Effects of unilateral lesions in central auditory system. *Journal of Neurophysiology*, 47:987–1016, 1982.

[56] W. M. Jenkins and M. M. Merzenich. Role of cat primary auditory cortex for sound-localization behavior. *Journal of Neurophysiology*, 52(5):819–47, 1984.

[57] J. Jerger, L. Lovering, and M. Wertz. Auditory disorder following bilateral temporal lobe insult: report of a case. *Journal of Speech and Hearing Disorders*, 37:523–535, 1972.

[58] J. Jerger, N. J. Weikers, F. W. Sharbrough, and S. Jerger. Bilateral lesions of the temporal lobe: a case study. *Acta Oto-laryngologica Supplement*, 258:1–51, 1969.

[59] G. L. Kavanagh and J. B. Kelly. Contribution of auditory cortex to sound localization by the ferret (*Mustela putorius*). *Journal of Neurophysiology*, 57(6):1746–66, 1987.

[60] J. B. Kelly. *The effects of lateral lemniscal and neocortical lesions on auditory absolute thresholds and frequency difference thresholds of the rat.* PhD thesis, Vanderbilt University, 1970.

[61] J. B. Kelly. The effects of insular and temporal lesions in cats on two types of auditory pattern discrimination. *Brain Research*, 62:71–87, 1973.

[62] J. B. Kelly and P. W. Judge. Binaural organization of primary auditory cortex in the ferret (*Mustela putorius*). *Journal of Neurophysiology*, 71(3):904–13, 1994.

[63] J. B. Kelly, P. W. Judge, and D. P. Phillips. Representation of the cochlea in primary auditory cortex of the ferret (*Mustela putorius*). *Hearing Research*, 24(2):111–5, 1986.

[64] J. B. Kelly, B. J. Rooney, and D. P. Phillips. Effects of bilateral auditory cortical lesions on gap-detection thresholds in the ferret (*Mustela putorius*). *Behavioral Neuroscience*, 110(3):542–50, 1996.

[65] J. B. Kelly and I. C. Whitfield. Effects of auditory cortical lesions on discriminations of rising and falling frequency-modulated tones. *Journal of Neurophysiology*, 34(5):802–816, 1971.

[66] A. J. King and S. Carlile. Responses of neurons in the ferret superior colliculus to the spatial location of tonal stimuli. *Hearing Research*, 81:137–149, 1994.

[67] D. J. Klein, D. A. Depireux, J. Z. Simon, and S. A. Shamma. Robust spectrotemporal reverse correlation for the auditory system: optimizing stimulus design. *Journal of Computational Neuroscience*, 9:85–111, 2000.

[68] D. J. Klein, S. A. Shamma, J. Z. Simon, and D. A. Depireux. Linear and non-linear responses to dynamic broad-band spectra in primary auditory cortex, February 1999. Available to download from `http://www.isr.umd.edu/CAAR/posters/ARO99-cortex.pdf`.

[69] G. F. Kuhn. Model for interaural time differences in the azimuthal plane. *Journal of the Acoustical Society of America*, 62:157–167, 1977.

[70] W. C. Loftus and M. L. Sutter. Spectrotemporal organization of excitatory and inhibitory receptive fields of cat posterior auditory field neurons. *Journal of Neurophysiology*, 86(1):475–491, 2001.

[71] T. Lu, L. Liang, and X. Wang. Temporal and rate representations of time-varying signals in the auditory cortex of awake primates. *Nature Neuroscience*, 4(11):1131–1138, 2001.

[72] T. Lu and X. Wang. Temporal discharge patterns evoked by rapid sequences of wide- and narrowband clicks in the primary auditory cortex of cat. *Journal of Neurophysiology*, 84(1):236–246, 2000.

[73] J. C. Makous and J. C. Middlebrooks. Two-dimensional sound localization by human listeners. *Journal of the Acoustical Society of America*, 87(5):2188–200, 1990.

[74] D. Marr. *Vision: a computational investigation into the human representation and processing of visual information.* W. H. Freeman, San Francisco, 1982.

[75] R. L. Martin and W. R. Webster. The auditory spatial acuity of the domestic cat in the interaural horizontal and median vertical planes. *Hearing Research*, 30:239–252, 1987.

[76] R. B. Masterton. Role of the mammalian forebrain in hearing. In J. Syka, editor, *Acoustical Signal Processing in the Central Auditory System*. Plenum Press, New York, 1997.

[77] B. J. May. Role of the dorsal cochlear nucleus in the sound localization behaviour of cats. *Hearing Research*, 148:74–87, 2000.

[78] B. J. May and A. Y. Huang. Sound orientation behavior in cats. I. Localization of broadband noise. *Journal of the Acoustical Society of America*, 100(2 Pt 1):1059–69, 1996.

[79] D. K. Mellinger and B. M. Mont-Reynaud. Scene analysis. In H. L. Hawkins, T. A. McMullen, A. N. Popper, and R. R. Fay, editors, *Auditory Computation*, chapter 7, pages 271–331. Springer-Verlag, New York, 1996.

[80] M. M. Merzenich and J. F. Brugge. Representation of the cochlear partition on the superior temporal plane of the macaque monkey. *Brain Research*, 50:275–296, 1973.

[81] J. C. Middlebrooks, A. E. Clock, L. Xu, and D. M. Green. A panoramic code for sound location by cortical neurons. *Science*, 264(5160):842–4, 1994.

[82] J. C. Middlebrooks, L. Xu, A. C. Eddins, and D. M. Green. Codes for sound-source location in non-tonotopic auditory cortex. *Journal of Neurophysiology*, 80:863–881, 1998.

[83] D. R. Moore, M. E. Hutchings, P. D. Addison, M. N. Semple, and L. M. Aitkin. Properties of spatial receptive fields in the central nucleus of the cat inferior colliculus. II. Stimulus intensity effects. *Hearing Research*, 13:159–174, 1984.

[84] D. R. Moore, M. N. Semple, P. D. Addison, and L. M. Aitkin. Properties of spatial receptive fields in the central nucleus of the cat inferior colliculus. I. Responses to tones of low intensity. *Hearing Research*, 13:159–174, 1984.

[85] W. D. Neff, J. F. Fisher, I. T. Diamond, and M. Yela. Role of auditory cortex in discrimination requiring localization of sound in space. *Journal of Neurophysiology*, 19:500–512, 1956.

[86] S. R. Oldfield and S. P. A. Parker. Acuity of sound localisation: a topography of auditory space. I. Normal hearing conditions. *Perception*, 13:581–600, 1984.

[87] L. M. Optican, T. J. Gawne, B. J. Richmond, and P. J. Joseph. Unbiased measures of transmitted information and channel capacity from multivariate neuronal data. *Biological Cybernetics*, 65(5):305–10, 1991.

[88] L. M. Optican and B. J. Richmond. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *Journal of Neurophysiology*, 57(1):162–78, 1987.

[89] S. Panzeri, R. S. Petersen, S. R. Schultz, M. Lebedev, and M. E. Diamond. The role of spike timing in the coding of stimulus location in rat somatosensory cortex. *Neuron*, 29:769–777, 2001.

[90] S. Panzeri and A. Treves. Analytical estimates of limited sampling biases in different information measures. *Network*, 7:87–107, 1995.

[91] C. H. Parsons, M. Halloran, O. Kacelnik, D. R. Moore, and A. J. King. Effects of sound level on the accuracy of auditory localization by ferrets. *International Journal of Audiology*, 41:246, 2002.

[92] C. H. Parsons, R. G. Lanyon, J. W. H. Schnupp, and A. J. King. Effects of altering spectral cues in infancy on horizontal and vertical sound localization by adult ferrets. *Journal of Neurophysiology*, 82:2294–2309, 1999.

[93] B. R. Payne, S. G. Lomber, A. E. Villa, and J. Bullier. Reversible deactivation of cerebral network components. *Trends in Neurosciences*, 19(12):535–542, 1996.

[94] S. Perrett and W. Noble. Available response choices affect localization of sound. *Perception and Psychophysics*, 57(2):150–158, 1995.

[95] R. S. Petersen, S. Panzeri, and M. E. Diamond. Population coding of stimulus location in rat somatonsensory cortex. *Neuron*, 32:503–514, 2001.

[96] B. E. Pfingst and T. A. O'Connor. Characteristics of neurons in auditory cortex of monkeys performing a simple auditory task. *Journal of Neurophysiology*, 45:16–34, 1981.

[97] D. P. Phillips and M. E. Farmer. Acquired word deafness, and the temporal grain of sound representation in the primary auditory cortex. *Behavioural Brain Research*, 40:85–94, 1990.

[98] D. P. Phillips and D. R. F. Irvine. Responses of single neurons in physiologically defined area A1 of the cat cerebral cortex: sensitivity to interaural intensity differences. *Hearing Research*, 4:299–307, 1981.

[99] D. P. Phillips and D. R. F. Irvine. Responses of single neurons in physiologically defined primary auditory cortex (AI) of the cat: frequency tuning and responses to intensity. *Journal of Neurophysiology*, 45:48–58, 1981.

[100] D. P. Phillips and D. R. F. Irvine. Properties of single neurons in the anterior auditory field (AAF) of cat cerebral cortex. *Brain Research*, 248:237–244, 1982.

[101] D. P. Phillips, J. W. Judge, and J. B. Kelly. Primary auditory cortex in the ferret (*Mustela putorius*): neural response properties and topographic organization. *Brain Research*, 443:281–294, 1988.

[102] D. P. Phillips and S. S. Orman. Responses of single neurons in posterior field of cat auditory cortex to tonal stimulation. *Journal of Neurophysiology*, 51:147–163, 1984.

[103] P. Poirier, S. Miljours, M. Lassonde, and F. Lepore. Sound localization in acallosal human listeners. *Brain*, 116(Pt 1):53–69, 1993.

[104] R. Rajan, L. M. Aitkin, and D. R. F. Irvine. Azimuthal sensitivity of neurons in primary auditory cortex of cats. II. Organization along frequency-band strips. *Journal of Neurophysiology*, 64(3):888–902, 1990.

[105] R. Rajan, L. M. Aitkin, D. R. F. Irvine, and J. McKay. Azimuthal sensitivity of neurons in primary auditory cortex of cats. I. Types of sensitivity and the effects of variations in stimulus parameters. *Journal of Neurophysiology*, 64(3):872–887, 1990.

[106] J. P. Rauschecker and B. Tian. Mechanisms and streams for processing of "what" and "where" in auditory cortex. *Proceedings of the National Academy of Sciences*, 97(22):11800–11806, 2000.

[107] J. P. Rauschecker, B. Tian, T. Pons, and M. Mishkin. Serial and parallel processing in rhesus monkey auditory cortex. *Journal of Comparative Neurology*, 382(1):89–103, 1997.

[108] Lord Rayleigh. On our perception of sound direction. *Philosophical Magazine*, 13:214–232, 1907.

[109] H. L. Read, J. A. Winer, and C. E. Schreiner. Functional architecture of auditory cortex. *Current Opinion in Neurobiology*, 12(4):433–440, 2002.

[110] R. A. Reale and J. F. Brugge. Auditory cortical neurons are sensitive to static and continuously changing interaural phase cues. *Journal of Neurophysiology*, 64(4):1247–1260, 1990.

[111] R. A. Reale and T. J. Imig. Tonotopic organization in auditory cortex of the cat. *Journal of Comparative Neurology*, 192:265–291, 1980.

[112] G. H. Recanzone. Spatial processing in the auditory cortex of the macaque monkey. *Proceedings of the National Academy of Sciences*, 97(22):11829–11835, 2000.

[113] J. J. Rice, B. J. May, G. A. Spirou, and E. D. Young. Pinna-based spectral cues for sound localization in the cat. *Hearing Research*, 58:132–152, 1992.

[114] F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek. *Spikes: exploring the neural code*. MIT Press, Cambridge, MA, 1997.

[115] E. Rouiller, Y. de Ribaupierre, A. Toros-Morel, and F. de Ribaupierre. Neural coding of repetitive clicks in the medial geniculate body of the cat. *Hearing Research*, 5:81–100, 81.

[116] J. W. H. Schnupp, T. D. Mrsic-Flogel, and A. J. King. Linear processing of spatial cues in primary auditory cortex. *Nature*, 414:200–204, 2001.

[117] C. E. Schreiner and M. S. Cynader. Basic functional organization of second auditory cortical field (AII) of the cat. *Journal of Neurophysiology*, 51:1284–1305, 1984.

[118] C. E. Schreiner and G. Langner. Periodicity coding in the inferior colliculus of the cat. II. Topographical organization. *Journal of Neurophysiology*, 60(6):1823–1840, 1988.

[119] H. Schulze, A. Hess, F. W. Ohl, and H. Scheich. Superposition of horseshoe-like periodicity and linear tonotopic maps in auditory cortex of the Mongolian gerbil. *European Journal of Neuroscience*, 15(6):1077–1084, 2002.

[120] H. Schulze and G. Langner. Auditory cortical responses to amplitude modulations with spectra above frequency receptive fields: evidence for wide spectral integration. *Journal of Comparative Physiology*, 185(6):493–508, 1999.

[121] J. R. Searle. Consciousness. *Annual Review of Neuroscience*, 23:557–558, 2000.

[122] M. A. Segraves and A. C. Rosenquist. The afferent and efferent callosal connections of retinotopically defined areas in cat cortex. *Journal of Neuroscience*, 2(8):1090–1107, 1982.

[123] M. A. Segraves and A. C. Rosenquist. The distribution of the cells of origin of callosal projections in cat visual cortex. *Journal of Neuroscience*, 2(8):1079–1089, 1982.

[124] S. A. Shamma, J. W. Fleshman, P. R. Wiser, and H. Versnel. Organization of response areas in ferret primary auditory cortex. *Journal of Neurophysiology*, 69(2):367–83, 1993.

[125] S. A. Shamma and H. Versnel. Ripple analysis in ferret primary auditory cortex. II. Prediction of unit responses to arbitrary spectral profiles. *Auditory Neuroscience*, 1:255–70, 1995.

[126] S. A. Shamma, H. Versnel, and N. Kowalski. Ripple analysis in ferret primary auditory cortex. I. Response characteristics of single units to sinusoidally rippled spectra. *Auditory Neuroscience*, 1:233–54, 1995.

[127] C. E. Shannon. A mathematical theory of communication. *Bell Systems Technical Journal*, 27:379–423, 623–656, 1948.

[128] E. A. G. Shaw. The external ear. In W. D. Keidel and W. D. Neff, editors, *Handbook of Sensory Physiology: Auditory System*, pages 455–489. Springer-Verlag, New York, 1974.

[129] A. L. Smith, R. G. Lanyon, C. J. Akerman, G. E. Baker, A. C. Dempster, C. H. Parsons, I. D. Thompson, and A. J. King. Functional consequences of persistent, localized and reversible inactivation of cerebral cortex by muscimol. In preparation, 2002.

[130] S. K. Talwar, P. G. Musial, and G. L. Gerstein. Role of mammalian auditory cortex in the perception of elementary sound properties. *Journal of Neurophysiology*, 85:2350–2358, 2001.

[131] B. Tian, D. Reser, A. Durham, A. Kustov, and J. P. Rauschecker. Functional specialization in rhesus monkey auditory cortex. *Science*, 292(5515):290–293, 2001.

[132] M. J. Tramo, G. D. Shah, and L. D. Braida. Functional role of auditory cortex in frequency processing and pitch perception. *Journal of Neurophysiology*, 87:122–139, 2002.

[133] H. Versnel, N. Kowalski, and S. A. Shamma. Ripple analysis in ferret primary auditory cortex. III. Topographic distribution of ripple response parameters. *Auditory Neuroscience*, 1:271–85, 1995.

[134] H. Versnel, J. E. Mossop, T. D. Mrsic-Flogel, B. Ahmed, and D. R. Moore. Optical imaging of intrinsic signals in ferret auditory cortex: responses to narrowband sound stimuli. *Journal of Neurophysiology*, 88:1545–1558, 2002.

[135] X. Wang. On cortical coding of vocal communication sounds in primates. *Proceedings of the National Academy of Sciences*, 97(22):11843–11849, 2000.

[136] T. Watanabe. Fundamental study of the neural mechanism in cats subserving the feature extraction process of complex sounds. *Japanese Journal of Physiology*, 22:569–583, 1972.

[137] T. Watanabe and Y. Katsuki. Response patterns of single auditory neurons of the cat to species-specific vocalization. *Japanese Journal of Physiology*, 24:135–155, 1974.

[138] I. C. Whitfield. Auditory cortex and the pitch of complex tones. *Journal of the Acoustical Society of America*, 67(2):644–647, 1980.

[139] B. J. Winer, D. R. Brown, and K. M. Michels. *Statistical principles in experimental design*. McGraw-Hill, New York, third edition, 1991.

[140] J. A. Winer. Thalamus and cortex. In D. B. Webster, A. N. Popper, and R. R. Fay, editors, *The Mammalian Auditory Pathway: Neuroanatomy*. Springer-Verlag, New York, 1992.

[141] O. W. Witte, I. Buchkremer-Ratzmann, K. Schiene, T. Neumann-Haefelin, G. Hagemann, M. Kraemer, K. Zilles, and H. J. Freund. Lesion-induced network plasticity in remote brain areas. *Trends in Neurosciences*, 20(8):348, 1997.

[142] D. H. Wolpert and D. R. Wolf. Estimating functions of probability distributions from a finite set of samples. *Physical Review series E*, 52:6841–6854, 1995.

[143] R. S. Woodworth. *Experimental Psychology*. Holt, New York, 1938.

[144] L. Xu, S. Furukawa, and J. C. Middlebrooks. Sensitivity to sound-source elevation in nontonotopic auditory cortex. *Journal of Neurophysiology*, 80:882–894, 1998.

[145] K. Yamada, K. Kaga, A. Uno, and M. Shindo. Sound lateralization in patients with lesions including the auditory cortex: Comparison of interaural time difference (ITD) discrimination and interaural intensity difference (IID) discrimination. *Hearing Research*, 101:173–180, 1996.

[146] M. P. Young, C. C. Hilgetag, and J. W. Scannell. On imputing function to structure from the behavioural effects of brain lesions. *Philosophical Transactions of the Royal Society of London B*, 355:147–161, 2000.

[147] R. J. Zatorre. Pitch perception of complex tones and human temporal-lobe function. *Journal of the Acoustical Society of America*, 84:566–572, 1988.

[148] R. J. Zatorre and V. B. Penhune. Spatial localization after excision of human auditory cortex. *Journal of Neuroscience*, 21(16):6321–6328, 2001.