One tone, two ears, three dimensions: An investigation of qualitative echolocation strategies in synthetic bats and real robots

V. Ashley Walker

Ph.D.
University of Edinburgh
1997

Abstract

The aim of the work reported in this thesis is to investigate a methodology for studying perception by building and testing robotic models of animal sensory mechanisms. Much of Artificial Intelligence studies agent perception by exploring architectures for linking (often abstract) sensors and motors so as to give rise to particular behaviour. By contrast, this work proposes that perceptual investigations should begin with a characterisation of the underlying *physical laws* which govern the specific interaction of a sensor (or actuator) with its environment throughout the execution of a task. Moreover, it demonstrates that, through an understanding of *task-physics*, problems for which architectural solutions or explanations are often proposed may be solved more simply at the sensory interface — thereby minimising subsequent computation.

This approach is applied to an investigation of the acoustical cues that may be exploited by several species of tone emitting insectivorous bats (species in the families Rhinolophidae and Hipposideridae) which localise prey using systematic pinnae scanning movements. From consideration of aspects of the sound filtering performed by the external and inner ear of these bats, three target localisation mechanisms are hypothesised and tested aboard a 6 degree-of-freedom, binaural, robotic echolocation system.

In the first case, it is supposed that echolocators with narrow-band call structures use pinna *movement* to alter the directional sensitivity of their perceptual systems in the same way that broad-band emitting bats rely on pinnae *morphology* to alter acoustic directionality at different frequencies.

Scanning receivers also create dynamic cues — in the form of frequency and amplitude modulations — which vary systematically with target angle. The second hypothesis investigated involves the extraction of timing cues from amplitude modulated echo envelopes. This mechanism provides an echolocator with a means for creating dramatic temporal cues for target localisation which, unlike inter-aural timing differences, do not degrade with head size.

The final mechanism investigated exploits the cosine-law dependence of the Doppler effect to extract the off-axis angular position of a target *via* inter-aural frequency differences.

In these investigations, targets consisted of reflectors with rotary movement (*e.g.* small computer cooling fans) which, like the fluttering insect targets of rhinolophids and hipposiderids, periodically modulate sound upon reflection. All localisation schemes operate on the spectral sideband energy reflected by targets with periodic motion and, therefore, can be used to disambiguate the position of these targets amongst the many stationary reflectors present in a cluttered laboratory environment. When these same simple localisation strategies are made to operate on energy in the *particular* sidebands reflected by targets with *particular* motions, target selective localisation emerges.

Acknowledgements

I am grateful to a number of people for their contributions to my thesis work and life in Edinburgh over the past few years.

First, I would like to acknowledge John Hallam for both his supervision of this work, and for the friendship and hospitality that he and his family have shown me. Here, I particularly thank him for trusting me to plot my own course, while being always ready with bailing pails and maps when I found myself aground in strange territory. It was my good fortune to have worked with such a diverse scientist who is also an expert in my technical field.

In the middle of my time in Edinburgh, another of the leaders in the small field of airborne SONAR joined our lab. Herbert Peremans not only supervised my project, but we have had an almost daily interaction for more than a year. I have grown immeasurably from working alongside him. He provided me with a model of a scientist who is most careful, competent and generous with his knowledge.

At times over the course of this work, various friends and colleagues have offered invaluable insights and inspirations. In my first year of research, David Willshaw provided me with guidance in a potential PhD project. (Unfortunately, we were unable to finish it (as such) because the work reported in this thesis had not yet been done.) In 1996, I had a very brief conversation with Brock Fenton at the North American Bat Symposium which changed the course of this work. Martin Westhead has taken an enormous amount of time out of his own research to understand and encourage mine. Other individuals who provided me with valuable feedback (particularly in the preparation of this report) include: Gareth Jones, Gillian Hayes, Beau Lotto, and Richard Shaw. In addition to my supervisors, the following people assisted with the design, construction and maintenance of my computing and robotic equipment: Douglas Howe, Andrew Haston, Sandy Coloquon, Neil Wood, Hugh Cameron, David Wyse, Tom Alexander, Tim Colles and Pete Forrester.

I was partially funded by a British Council Overseas Studentship, a stipend from the American Friends of Edinburgh University and a S.M.A.R.T. fellowship from the EC. As these forms of support did not cover all of the costs associated with living and enjoying oneself in Edinburgh, I would like to remember various people who helped me make ends meet over past few years. These include Bob Fisher (for work on HIPR), Isobel Tate and Martin Westhead (for whom I was a sub-warden), and Bob Bagnall and June Cunningham (who employed me as an instructor in the Departments of Continuing Education and Physical Education, respectively).

On a personal note, I am very grateful to my family: Dale, Mae and Whitney Walker — who took several holidays in Edinburgh when I couldn't get away, generously loaned me money when part-time work became infeasible, ... and knew exactly when a long-distance call or the delivery of a jar of peanut butter was most welcome. In that paradoxical way that only families can, mine constantly encouraged me (across many miles) to "stick with it" while also reminding me that "it doesn't matter" if I finish.

Finally, I would not have moved to Edinburgh and taken up post-graduate research had I not known Beau Lotto. I thank him for beginning this journey with me. I would also like to thank Martin Westhead (whose name already appears in so many places on this page) for concluding the journey with me and for sharing so many interesting stops along the way.

Declaration

I hereby declare that I composed this thesis entirely myself and that it describes my own research.

Ashley Walker
Edinburgh
December 15, 1998

Contents

List of Figures

List of Symbols and Terms

Symbols

t	time
r	range
a	transducer aperture radius
c	speed of sound ($c \approx 340$ m/s in air)
v	agent velocity (m/s)
θ	angular position of target relative to transducer normal
θ_n	angle between transducer normal and target surface normal
δ	"Doppler angle" between vectors defining the line-of-sight from observer to target, and the observer (or target) velocity
A	amplitude
f	frequency (Hz)
ω	frequency (rad/s), $\omega = 2\pi f$
β	dispersion index, $\beta = \frac{\Delta f}{f}$
J_n	nth order Bessel function of the first kind
Q_{10dB}	tuning quality factor, $Q_{10dB} = \frac{f}{\Delta f}$ (measured at 10 dB above threshold)
$SONAR$	SOund NAvigation and Ranging
CF	Constant frequency
CF_2	Second CF harmonic
FM	Frequency modulation
AM	Amplitude modulation
IID	Inter-aural intensity disparity
ITD	Inter-aural temporal disparity
IFD	Inter-aural frequency disparity

Terminology

- **perception.** Sensory (and motor) mechanisms underlying behavioural control.
- *Merkwelt.* The perceptual world of a particular agent.
- **active.** Exploiting the principle of sounding — *e.g.* echolocation, controlled illumination in vision, *etc.*
- **dynamic sensing.** Adjustable, mobile, information seeking mechanisms that "accommodate themselves to the possibilities of stimulation" [Gibson 66].
- **matched filtering.** The use of receptor morphology to exploit constraints imposed by the task and the environment in a manner which simplifies the subsequent signal processing required. (The same could be applied to actuator surfaces.)
- **dynamic filtering.** The use of receptor motion to exploit constraints imposed by the task and the environment to simplify the signal processing required.

- **echolocation.** The perception of objects by means of the way in which they reflect a controlled acoustical beam transmitted by the perceiver. (Synonymous with the term "active SONAR".)

- **SONAR horizon.** The region in space wherein inter-aural intensity disparities (at a given frequency) vary most sharply.

- **bionics.** The fusion of mechanisms from **bio**logical sensing and electro**nics** through the construction of "life-like" sensor systems.

- **agent.** An autonomous entity — *e.g.* animal or robot. This term is not used to equate animals and robots at any *particular* level, but to establish a framework to support discussions of similarities which might enrich and broaden our understanding of perception.

- **physical characterisation.** An attempt to conceptualise interactions between agents and their environments in terms of the underlying physical laws governing the operation of a sensor/actuator.

- **world interface mechanism.** The use of any transduction mechanism (*i.e.* sensory or actuatory) to reduce uncertainty in the world. (Often reduced to "sensing mechanism" or "actuation mechanism", though, in many cases, it can be difficult to distinguish the two.)

- **acoustic axes.** The direction (or region in space) in which the acoustic sensitivity of a receiver is greatest.

Overview

Scope

In every field, there are two broad questions that hang over the research: "what?" (subject) and "how?" (methodology). In Artificial Intelligence (AI) there is a third: "why?". This question asks AI researchers to make a claim for their work which is either *engineering* in purpose or *scientific*. Thus, work in AI tends to be viewed on a spectrum where at one pole there is research whose "aim is to construct an artificial system with a high degree of competence" and, at the other, "the goal is to understand the mechanisms underlying behaviour in natural systems by modelling them" [Hallam 87]. In the work reported in this thesis, I have worked very near to the centre of this spectrum. My aim is both to discover and to invent.

Over the past decade, an increasing amount of work has begun moving inward from the poles of the AI research spectrum. Nevertheless, it is still the case that those of us who do not have a single answer to the "why?" question are often called upon to defend what is viewed as an ambiguous position. I believe that this confusion may be alleviated by encouraging a deeper integration (rather than separation) of the scientific and engineering traditions within AI. In the case of robotics, the removal of these distinctions may accelerate the development of a rich vocabulary for discussing levels of equivalence between engineered and evolved problem solving strategies.

In this work, I have tried to do just this by absorbing the "why?" question into the methodological one. In other words, I have adopted an investigative framework in which the perceptual task under investigation is characterised in terms of the underlying, invariant *physical* (or in this case *acoustical*) processes which govern the interaction of an agent (animal or robot) with the physical world. A principled physical characterisation of this sort provides both a level of abstraction at which to understand perceptual mechanisms and an explanatory framework which is interpretable by a large number of communities. Indeed, one of the main contributions of this work is the validation of an investigative methodology whereby the concerns of biology and AI are fused at the level of developing and testing specific hypotheses about animal sensing.

The practical contribution of this work is the development of a 6 degree-of-freedom, binaural, echolocation system capable of localising targets in 3-dimensional space. The system uses a single narrow-band sensory call and employs receiver motion (rather than signal band-width) to create cues for determining target bearing. Furthermore, it is tuned to preferentially localise targets with characteristic motion and, as such, provides a demonstration of the principle that motion can be exploited to achieve more robust detection in acoustically cluttered environments.

On another level, the performance of this artificial echolocation system demonstrated that the underlying dynamic sensory mechanisms — hypothesised as explanations of hunting strategies employed by several species of Old World bats — are sufficient to account for a number of aspects of the bats' behaviour. Thus, these explanations, which are simpler than a number of current explanations, are plausible hypotheses and can act as a basis for study of behaviour and neurophysiology in bats.

The results also support the argument that it is fruitful to study perception in terms of specialised structures that provide agent-environment interfaces capable of registering the information inherent in the environment. This is contrasted with architectural or representational explanations of perception which seek to create information from (often idealised) sensors or actuators. Instead, when perceptual mechanisms are determined by the physical details of an agent-environment interaction, simplifying assumptions can be made which lead to a reduction in the subsequent signal processing required and facilitate robust and rapid performance of the agent with respect to its specific task(s).

Structure

This thesis is divided into four parts: *I. A Perspective on Perception*, *II. Structuring Sound at the Target*, *III. Structuring Sound at the Receiver*, and *IV. Conclusions*.

In Part I, I devote two chapters to addressing the "what?" and "how?" questions, respectively. Specifically, in Chapter 1, I identify what I believe to be the most interesting aspect of intelligence to investigate and, in Chapter 2, I lay out my experimental methodology. While the former sets up the central question of this thesis, the latter defines how I intend to answer it.

Parts II and III explore hypothesised biosonar perceptual mechanisms. Specifically, I review of some of the target *identification* and *localisation* mechanisms (respectively) thought to be used by echolocating bats. The use of two separate Parts to cover the topics of the empirical investigations somewhat over-emphasises the differences between them. Therefore, in this *Overview*, I want to make it clear that I use separate parts for explanatory purposes only. Indeed, an important result of this thesis work was to show that when echo-localisation mechanisms are tuned to operate on attributes of the acoustic energy reflected by a desired target, those targets can be selectively localised without the need for an explicit, intermediate recognition mechanism to trigger localisation. In Part II, I investigate target signatures which can be used in this way and, in Part III, I explore strategies for using observer motion to superimpose intensity, time and frequency cues on echo signature energy so as to introduce 3D localisation cues.

In terms of format, Parts II and III each begin with a *Background* chapter which provides a general overview contrasting the strategies taken by different bat species in exploiting the acoustics underlying detection and localisation of particular targets in particular environments. It is the purpose of these broadly focused chapters to review and raise questions for which the next chapter(s) provides answers. *Background* chapters are followed by one or more *Investigation* chapters. (One target detection and three different localisation hypotheses were investigated). Chapters reporting novel work contain the sections: *Hypothesis*, *Experiment*, *Results* and *Discussion*.

Because of the inter-disciplinary nature of this research, the concluding part (*IV*) of the thesis contains three chapters describing the *Evaluation* of results. The results are first evaluated as hypotheses about perception in bats. I then compare the behaviour of the artificial echolocation system with that of other SONAR systems in AI/robotics. Finally, I discuss this research as an AI modelling investigation of a biological system — *i.e.* I consider the validity of the underlying investigative methodology which prompted the development and use of the system. The thesis concludes with *Summary of Contributions* of this research.

Navigation aids are given throughout the thesis to remind the reader of this structure.

Part I

A Perspective on Perception

"Our objective is to abstract patterns from Nature..., but many proposed patterns do not in fact correspond to the data. Thus all proposed patterns must be subjected to the sieve of critical analysis, [this] and rigid skepticism without a search for patterns, are the antipodes of incomplete science. The effective pursuit of knowledge requires both functions."

—C. Sagan

"... entities should not be multiplied beyond necessity..."

—Occam's razor

1. Introduction

"...there are two different meanings of the verb *to sense*, first, to detect something, and second, to have a sensation" [Gibson 66]. This thesis is about the first kind of sensing which I will refer to as *perception*.

Reid makes this distinction nicely [Reid 65]:

> *"The external senses have a double province; to make us feel and to make us perceive. ... The feeling which goes along with the perception, we call sensation. The perception and its corresponding sensation are produced at the same time. In our experience we never find them disjoined. Hence we are led to consider them as one thing, to give them one name, and to confound their different attributes. It becomes very difficult to separate them in thought, to attend to each by itself, and to attribute nothing to one which belongs to the other."*

This use of the term perception — to describe *sensory mechanisms*[1] underlying behavioural control — is a departure from the traditional one where the senses are investigated

[1] I use the phrase *sensory mechanisms* here with caution. What I really mean is *world-interface mechanisms*, which can involve any transduction device (*i.e.* sensor or actuator) which reduces uncertainty in the world. Like many authors, I will, at times, draw broad distinctions between sensing and actuation to facilitate explanatory clarity. However, although I may refer to them as separate functions, on many levels sensing and actuation are different aspects of the same function. Not only does one often involve the other, but an agent may just as well understand places in the environment, for example, in terms of the motor commands, rather than the sensory signals, they elicit — *e.g.* see [Nehmzow *et al.* 91].

as channels of sensation and perception is defined as the transformation of sense stimuli into conscious experience. It emphasises the point that "the function of perception in controlling action is an essential determinant of the mechanisms of a perceptual system" [Webb 93]. Thus, defining perception in this way allows us to examine it in a wide variety of biological and robotic systems without the need to explain or make assumptions about the sensations or experiences of those agents. Moreover, this use of the term emphasises the *dynamic*[2] role of the perceiver in obtaining the stimulation it receives, as opposed to having it externally imposed.

1.1 The question

A long-term goal of AI is to build an *autonomous system* — *e.g.* a robot. Autonomy has been variously defined, however, most contemporary definitions include the control of behaviour by sensory mechanisms: "Autonomy in this context refers to [agents'] basic and fundamental capacity to *be*, to assert their existence and to bring forth a world that is significant and pertinent without being digested in advance" [Bourgine & Varela 92]. Thus, the quest to build an autonomous robot has directed AI research effort toward understanding perception as defined here — see, for example, [Maes *et al.* 96, Meyer 97, Langton 97]. This, in turn, has had a profound effect on our understanding of the nature of *intelligence*. The capacity of an agent to interact competently with its environment, to extract — from a myriad of fluctuating stimuli — invariant descriptions of the world upon which to behave purposefully, is now regarded by many researchers as a fundamental aspect of intelligent behaviour [Winograd & Flores 86].

At present, the design of a truly autonomous robot capable of intelligent behaviour is still more art than science. There a poverty of support in programming languages for the non-hierarchical program structuring methods required to control autonomous robots [Hoare 85, Lyons 85]. Moreover, we are limited by our lack of understanding of how sensors (and actuators) should be structured and used to provide control signals required for this kind of behaviour. Indeed, may researchers view the

[2] The term "dynamic" is preferred to "active" so that the latter term may be reserved for sensory processes which exploit the principle of sounding — *e.g.* echolocation, controlled illumination in vision, *etc.*

problem of designing robot sensors, actuators and the linkages between them to be nothing less than the problems addressed by animal evolution (and biological learning mechanisms). "Working on autonomously guided vehicles is clearly tantamount to working in a kind of holistic animal micro-world: such work is forced to respect many (but not all) of the constraints that we say would apply to evolved biological systems" [Clark 89]. Indeed, the problem of finding a way of describing purposeful behaviour in an animal appears to be so closely related to the problem of finding a "coherent, structured, and not overly complex way of implementing purposeful behaviour in a robot" [Malcolm 90] that an increasing number of roboticists have begun building biological insights and, in some cases, circuits into the control systems of their robots – *e.g.* [Grey-Walter 53, Albus 81, Moravec 84, Braitenberg 84, Raibert 86, Brooks 86, Maes 89, Anderson & Donath 90, Beer 90, Nehmzow *et al.* 90, Arkin 91, Franceschini *et al.* 91, Webb 93, Kuc 94, Arbib & Lee 94, Kuc 96, Williamson 96].

A structured interaction between AI and the biological sciences can benefit the latter as well. Robotic models provide a platform from which to study hypotheses about the interaction of particular sensory systems with realistically complex environments. Also, as Agre points out, "building and analysing artificial and simulated creatures may help clarify many biological concepts by forcing unarticulated assumptions and unasked questions to the surface" [Agre 95b]. Webb corroborates this view [Webb 93]:

> *"In fact, there are very few biological systems, especially complete perceptual systems, that have been studied in anything like the detail required for translation into electronics and programming code. Looking at biological systems by attempting to build models of them is ... a potential opportunity for AI to contribute to biological understanding of perceptual systems."*

Despite clear and recognised benefits, an effective methodology for investigating animal perception *via* the design, construction and operational analysis of robotic models has not yet been established: "what is needed is a non-naive way of including biological insights" [Pfeifer & Verschure 92]. To state this another way, if an aim of AI is to use "principled characterisations of interactions between agents and their environments to guide explanation and design" [Agre 95b, Agre 95a], then, what we need in biologically-based AI is a clear notion of what constitutes a *principled characterisation*.

We might clarify this idea by looking at examples from related fields and subfields. In control theory, for instance, a principled characterisation is a set of differential equations which allows the *controller* (or *agent*) to relate the changes in control variables to changes in variables recording the state of a *plant* (*environment*). The mathematical characterisation of control theory arose out of the need for safety and guarantees of performance in relatively well-behaved systems. It is only one possibility. A principled characterisation of interaction need not have any of these (formal/mathematical) qualities to provide a useful guide to the design of artificial agents and the explanation of natural ones.

Earlier work in AI offers examples of principled characterisations of action which are often based on various logics. The *Plan based* approach to investigating intelligence [Lashley 51, Miller *et al.* 60, Newell & Simon 63] said that action has the structure it does "because it arises through the execution of things called Plans which have the same structure" [Agre 95b]. In most cases, this characterisation of action was not sufficiently *inter*-active to provide an explanation of perception as I've defined it here. In particular, this theory failed to provide an account of how actions might be adapted to circumstances with the result that *a priori* generated Plans tended to lead to successful behaviour in only very stagnant environments. Nevertheless, it did offer a clear account of why an agent's actions ought to work.

By contrast, the biologically-based approach suggests that a better explanation of intelligent behaviour would rely on *Plan* (or *action*) *selection* in response to environmental cues (*i.e. interaction*). However, researchers often fail to characterise, in a principled way, the underlying perceptual mechanisms through which this is achieved. Biologically-based systems often employ non-linear pattern analysis techniques to partion sensor data or map between different sensory spaces. It is also typical for a biologically-based system to require a skillful engineer to tweak circuits until they "work". Rarely are these solutions translated into a principled account of "why it works". Such an account of interactivity would necessarily include the answer to such questions as: what does the agent do in a particular environment?, in which environments does the agent work?, under what conditions will it achieve its goals?, what forms of interaction require internal architectural features such as memory? If the essential underlying physical principles exploited by a perceptual system are not well characterised, the work is of limited use to the research community and, possibly, to the designers themselves (should they need to improve upon the system).

Therefore, the central question of this thesis is: **How should we characterise perceptual mechanisms so as to generate hypotheses about intelligent behaviour which may be investigated aboard a robot?**

1.2 An answer

The answer to this question is: **in terms of task-physics**. In full, because an agent's interaction with its environment begins at the physical interface between sensors (or actuators) and the environment, we should seek to characterise environmental interactions in terms of the invariant physical laws which govern them. A full account of sensor behaviour over all possible inputs is not called for here. Rather, we should seek to explain how an agent exploits its sensors over the limited range of inputs afforded by the environment in which it performs specific tasks. This involves identifying assumptions and approximations which might be employed to simplify a task (at the expense of some degree of generality and/or precision).

Such a description naturally leads to a blue-print for a robotic model which can be built and tested. Therefore, in essence, my answer to the question of how best to fuse biology and robotics in a study of perception is the establishment of a research methodology which I will refer to as **bionics**[3] linking mechanisms from **bio**logical sensing and electro**nics** through the construction of "life-like" sensor systems.

Modelling sensors, like all modelling endeavours, is an exercise in abstraction. As we are not able to build an exact replica of the system under consideration (which would be of little use in modelling), then we must begin by identifying what are the essential

[3] The word "bionics" comes from the Greek: "bion" — the units of life — and the ending "ics" indicates "life-like". It has been used in scientific literature for the description of artificial SONAR systems [Busnel & Fish 80, Nachtigall & Moore 88] ("biomimetic" is another popular term). In the popular media, the term "bionic" has come to denote powerful sensors and actuators which, from this author's perspective, truly life-like sensors most certainly would be.

mechanisms which facilitate the successful functioning of perception in animal systems. This, of course, is one of the most impressive and far reaching scientific questions of our time, and one whose answer will undoubtedly be arrived at through a combination of research programmes which address this question both *via* empirical studies involving animal system and modelling endeavours. Below I discuss two (of many) perceptual mechanisms — *matched filtering* (Section 1.2.1) and *dynamic filtering* (Section 1.2.2) — which appear to be exploited by the sensory systems of a vast majority of animals (and a few robotic perceptual systems), but which are relatively poorly understood due to the inherent diversity with which they are expressed in different animals. Both mechanisms are highly suitable for investigation aboard a robotic platform and a methodology for doing this is given in the final section of this chapter.

1.2.1 Matched filtering, or "the right physics"

The morphology of animal receptors exploits constraints imposed by the task and the environment to simplify the neural processing required. Wehner borrows the engineering term "matched filter" and (metaphorically) applies it to arrays of receptors in which "smart" processing is built into the physics of the interface [Wehner 87]:

> *"... a great deal of spatial information used by an animal to guide its activities within its world is handled and processed at rather peripheral levels within the nervous system. In particular, I shall argue that in many cases it is already the spatial design of the receptor layer at the outermost periphery of the nervous system, that 'solves' a particular problem. ... Of course, perceiving the world through such a 'matched filter' severely limits the amount of information the brain can pick up from the outside world, but it frees the brain from the need to perform more intricate computations to extract the information finally needed for fulfilling a particular task. Algebra gives way to geometry."*

Wehner provides a number of examples of animal sensors which exploit this principle. These include two arthropods, the desert ant *Cataglyphis* and the semi-terrestrial crab *Ocypode*, whose eyes are composed of band-like foveal regions extending across the eye which provide their users with high acuity, size-constant perception of objects appearing near the horizon in their exceedingly flat environments. (These "visual streaks" also occur in the single-lens eyes of some vertebrates living in wide open spaces [Walls 42, Hughes 77].) This mechanism is contrasted with the eyes of flying insects, for whom detection of the horizon itself (as opposed to objects lying along it) is important (see review in [Wehner 81]). In this case, relevant information is obtained by tiny photo-receptors (*ocelli*) with highly under-focussed optics and high rates of receptor convergence which provide the flying animal with the heavily blurred images of the skyline necessary for maintaining flight posture.

As I will discuss in detail in Part III of this thesis, lateral separation of acoustic receivers across the head or body of an animal gives rise to directional cues (in 2D) in the form of inter-aural disparities of intensity, time and/or phase. When receivers are decorated with pinnae, broad-band sound arriving at the ear is filtered in an angle-dependent fashion so as to encode the third dimension [Roffler & Butler 68, Shaw 74]. For example, in humans, a single prominent notch has been identified in the spectrum of broad-band sound passing through the external ear — the frequency of which changes systematically with source elevation [Bloom 77].

In robotics, there are isolated examples of engineered sensing systems which illustrate the fact that exploitation of sensor physics can increase the "information rate" (*i.e.* amount of sense-information per time [Peremans 94]) of a sensor. Several of the robots built by Rodney Brooks and his colleagues at M.I.T. demonstrate these properties. These simple robots — built from distributed networks of augmented finite state machines according to the *subsumption architecture* [Brooks 86] — perform a single task (*e.g.* finding soda cans (*Herbert*), or another robot to follow (*Tom and Jerry*) [Brooks 91b]) and this task specialisation allowed the designers to identify a set of particular sensor or motor conditions upon which to execute equally well specified behaviours.

Brooks attributes the success of this approach to his robots' lack of *internal representation* [Brooks 91a]: if "there are no models built, the problem of uncertainty is inherently reduced." Such a claim made a dramatic break with the failing Plan-based approach to designing robots. However, it obscured the important contribution of

Brooks' work: namely, that these robots were successful in their tasks because their designers built a realistic model of the (albeit simple) sensors and actuators into their robotic systems. To such well designed robots, internal representations (*e.g.* a symbolic variable encoding attributes of sensor or actuator states) might be added and, in turn, these symbolic tokens might be used in reasoning processes without running into the *symbol grounding* problem [Malcolm *et al.* 89, Harnad 90]. In other words, when internal models of the world depend upon perception (and not the converse), they can avoid falling out of phase with their environmental referents. A nice example of this is the SOMASS system [Malcolm 97].

As a consequence of his mis-advertisement, Brooks' pioneering work has spawned more interest in de-centralised control architectures than in the physics of sensing and actuation. However, without a principled methodology for investigating the physical coupling of an agent to its environment, the subsumption architecture may prove to be too limited (or too slow) to facilitate the design of robots for tasks which require an ability to utilise more complex sense-information to guide more complex behaviour. More importantly, his research efforts have failed to impress upon most of the community that programming and control must essentially rely on the information provided by sensors. It is not possible to compensate for the lack of this information by computing. Halme states [Halme 97]:

> *"No doubt sensors make up the most difficult problem area in this field. ... What is important to note is that sensors are based on physics (in some cases on chemistry ...) and many of the problems encountered in this field are basically problems of physics."*

Probert reinforces this perspective in a description of the design of a reactive planning system for the Oxford AGV [Probert 94]:

> *"The realisation of reactive planning depends upon architecture design ... However its success is fundamentally limited by the reliability and quality of sensor data."*

Horswill provides a nice example of the integration of these ideas into the architectural discussion [Horswill 95]. He follows a robot design process of incremental refinement in which assumptions about structures of the environment are "folded in" to the robot "architecture" to yield a simpler version of the architecture — *i.e.* one which requires simpler forms of computation or, in the extreme, no explicit computation at all. He uses this framework to characterise the interaction of an indoor robot "tour guide" with its office environment in which the frequency of "markings" in the structure of the tiled floor *vs.* those of other objects simplifies the visual processing necessary to distinguish between them. Of this work, Agre writes [Agre 95b]:

> *"When the unit of analysis for the design and analysis are defined in terms of interactions, the mutual fit between an agent and its environment becomes the most important source of guidance for the design process. A highly general architecture may be able to function well in a wide variety of circumstances, assuming that its computations are not impossibly cumbersome, but this very generality will produce a great deal of 'slack' in the architecture's relationship to the environment. By aiming for a simple machinery, and by shifting the primary explanatory burden to interactions and not to the architecture, designers such as Horswill are forced to pay ever more detailed attention to the environment and the agent's place within it."*

Horswill admits that this method does not provide a mechanical design formula; rather, it takes considerable thought and post-hoc nationalisation to discover which environmental constraints might actually be exploited. In this sense, borrowing the constraints already exploited by the evolutionary design process (once we begin to characterise their physics adequately) may speed up emerging robot design methodologies such as this.

Another important aspect of perception is that particular kinds of motion can be used to alter sensing requirements.

1.2.2 Dynamic filtering

There is a long standing tradition in biological and psychological research to look at afferent processes in isolation from efferent ones. In the same way that the term "sensing" has often been used to describe input processes on their way to an inner mental life ("sensation"), "action" has been treated as an appendix to more interesting aspects of mental life. In reality, sensing and motion are inextricably bound up in an animal's understanding of its world. The dynamic orientating of receptor surfaces directly affects the content and quality of the sensory information passing through to the nervous system.

Controlling the position of sensors during perception constitutes an "important feature of an animal's strategy for making sensory discriminations" [Rasnow *et al.* 88]. When stimuli come about as a result of an agent's actions, the causal link is from response to stimulus as much as from stimulus to response. By creating and using re-afferent (or feed-back) sense *stimuli*[4], agents can exclude certain kinds of stimuli and enhance those which they do receive. "Instead of supposing that the brain constructs or computes the objective information from a kaleidoscopic inflow of sensations, we may suppose that the orienting of the organs of perception is governed by the brain so that the whole system of input and output resonates to the external information" [Gibson 66].

Many authors criticise Gibson's idea of *neural resonance* for being poorly specified and an over-simplification of perceptual problems (*e.g.* in higher visual processes [Marr 82]). While the former is certainly true, the latter is not necessarily so. An example of what many of us might be prepared to call "resonance" can be found in the desert ant *Cataglyphis* whose eyes are not only structured to provide distance-invariant descriptions of the horizontal skyline (as we saw in the previous section), but also contain a filter (*i.e.* photo-receptors sensitive to the plane of polarized light) through which the animal reads the celestial hemisphere as a compass. Rather than computing its direction from the polarisation pattern directly (a hard computation), the insect moves to align itself with the symmetry plane of the sky (sometimes using only a small

[4] The ordinary use of the word "stimuli" is based upon its Latin roots: "goad" or "prick". This usage carries with it the idea that stimuli are imposed on (passive) observers and come from a source external to the agent. Here the definition is extended to include stimulation produced by events inside the agent which may be brought about by the agent's own actions.

patch of blue sky) and then turns to set its course [Wehner 87].

Work in the field of animate vision has shown how simple optic flow and vision-servoed motions radically alter the nature of visual perception — making tractable some problems which are ill-posed in the context of traditional image processing [Gibson 79, Lee 80, Ballard 87, Ballard 89, Brown 89, Cliff 90]). Similarly, vocal motions (*e.g.* altering call characteristics) can simplify echolocation tasks. For example, several species of tone emitting bats continuously adjust the frequency of their calls in flight so as to compensate for some of the the Doppler shift imposed by their own movement. In this way, a bat can alter the frequency of echoes returning from a target so that they arrive within the narrow range of frequencies where cochlear and neural sensitivity is greatest (*i.e.* the "acoustic fovea" [Schuller & Pollak 79, Neuweiler *et al.* 80]). In Part II of this thesis, I discuss how echoes from fluttering insect prey modulate a CF call signal during reflection such that when a flying bat uses Doppler compensation to bring insect echoes into the foveal window, narrowly tuned foveal neurons discharge in synchrony (or "resonate") with the wing-beat of an insect.

That a biological perceptual system can distinguish stimuli obtained from its own action *vs.* those imposed upon it by external events is now well accepted. For some time this was seen to violate the once prominent doctrine of *specific nerve energies*, formulated by Johanne Müller (see [Boring 42]), which asserted that no information whatsoever about the cause of a receptor's arousal can get past the receptor into the nervous system. It may be true that animals have to "deduce the causes of [their] sensations" [Helmholtz 25] when considering the senses as channels of sensation. However, Gibson argues that, when considering sensing within a perceptual system, this is not the case. Action sensitivity in animals does not depend on specialised receptors, rather there are many concurrent loops available for proprioceptive control of action. They operate at different levels, some stay inside the body, some loop out [Gibson 66].

By exploiting those motor-sensory paths which loop through the environment, an agent can use of the "world as its own model" [Brooks 91b]. This simple and robust perceptual principle exploits the fact that stimuli caused by self-produced action can contain a particular informational structure which an observer can use to identify and maintain a spatial relationship with environmental structures. For example, in the case of optic flow, Lee explains [Lee 74]:

"... the statement that the optic flow pattern affords information about the environment and about the observer's movement means that there are certain 'exterospecific' properties of the flow pattern that are specific to properties of the environment, and certain 'propriospecific' properties that are specific to the observer's movement."

The control of behaviour *via* the use of information available in "acoustic flow" is only beginning to receive research attention — *e.g.* controlled braking [Lee *et al.* 92, Lee 94] or obstacle avoidance [Müller & Schnitzler 97]. How a bat might register the sense information available in the acoustic flow field is a question that has not been addressed experimentally. However, this ability "must be grounded on registering the direction of sounds" [Lee *et al.* 92]. In this work, I examine target localisation cues which an observer can generate *via* dynamic, exploratory self-motion.

1.2.3 The physics of actuation

The principle of "let the physics do the walking" [Flynn & Brooks 89] could be applied to any aspect of an agent's interface with the world. Though I investigate mechanisms which I claim are predominantly "sensory" in nature, these ideas have been taken up to an even greater extent in robotic investigations where the main analytical tools are built upon the theories of kinematics and dynamics in physics [Agre 95b]. We see these principles in operation on a number of different levels.

For example, particular morphologies make agents more or less resistant to getting stuck (*e.g.* the mechanical remote centre compliance [Nevins & Whitney 78]) and afford particular kinds of mobility — *e.g.* wings for flight, webbed appendages for aqueous propulsion.

From analysis of the various types of symmetry found in animal gaits, Raibert built walking and running robots which exploit symmetry to free the motion controller from regulating details of the trajectory: "the details are determined passively by the mechanical system" [Raibert 86]. His work demonstrates nicely how a principled characterisation of the physically possible cycles of motion can simplify explanation and design of locomotion control architectures [Raibert 86, Raibert 89].

A characterisation of a mobile agent might also take into account the interactions among objects in the world which the agent's actions set in motion [Mason 86, Malcolm 90]. Malcolm employs this principle in assembly robotics; he notes: it is "possible in some circumstances to take advantage of naturally occurring constraints on motion to provide the feedback to control errors in part motion without requiring sensing at all, such as by pushing or dropping actions" [Malcolm 90].

1.3 Toward bionics

In AI and other computational sciences, several mechanisms underlying animal perception have been investigated. By far the most popular of these is *learning* — both structured skill refinement and less structured evolution of behaviours (see most work in *e.g.* [Maes *et al.* 96, Meyer 97, Langton 97]).

By contrast, a number of researchers, while conceding that adaptability is an important component of autonomy, argue that "adaptability in this sense does not necessarily require adaptability in the sense of an ability to change internal mechanisms in response to environmental changes: an agent may be able to survive because it has fixed internal mechanisms that are not adversely affected by the changes in its normal environment..." [Webb 93]. Hoyle claims that for "the animal kingdom as a whole, the majority of even the most complex behaviours ... fall into the category of instinctive acts. They require no experience of the behaviour in its context, nor learning, for their perfect execution" [Hoyle 84]. Even if this case is somewhat overstated, *computational neuroethology* [Camhi 84, Cliff 90] — which studies perception by implementing electro-mechanical models of the neural control loops underlying behaviours in simple animals (*e.g.* insects) — is also a growing application area for biologically-based robotics.

In both the computational neuroethological and learning/evolutionary approaches to understanding perception (and in Brook's "synthetic neuroethological" or "reactive" approach described in Section 1.2.1), the explicit aim is to investigate architectures for controlling behaviour. Some of this work has resulted in the construction and evaluation of sensory systems structured to the task and the environment of their users. However, it is all too common for researchers (particularly when investigating perception in computer simulation as opposed to aboard a robot) to focus on schemes for processing the information which they would like to extract from sensors rather than exploiting a possible coupling between sensor and environment.

It is my purpose in this thesis to argue that additional research effort in AI should be devoted to developing principled characterisations of sensory systems in terms of the underlying physics involved in sensing — *e.g.* the matched and dynamic filtering principles described above. In doing so, I think we will see that many problems of perception can be solved most elegantly at the sensory interface. The community has resisted this because "there is a widespread misunderstanding between building a system based on a principled model and building a system which itself uses a model to reason" [personal communication, Martin Westhead]. The latter has been seen as one of the failings of robots built according to the Plan-based approach wherein robots wasted most of their compute cycles reasoning about sensor uncertainty so that they could match signals from passive sensors to the predigested world models underlying Plans. By contrast, exploitation of task-physics (through, for example, matched and dynamic filtering) minimizes subsequent processing because a sensor model is effectively built right into the sensory system itself by tuning it for a specific task.

Characterisation of interaction in this way is not easy because our intuitive and analytical tools for understanding physical phenomena are focused more on general, powerful explanations that work in simple environments rather than on particular, simplifying, task-specific solutions which must operate in the natural world. Isolating the essential features of an environment and task which an agent might exploit requires extensive investigation and, even then, an understanding of the whole mechanism may only come retrospectively [Horswill 95]. However, because we have existence proofs of animal systems employing task-physics to solve a number of perceptual problems, I advocate a reverse engineering approach, wherein the *particular* sensing strategies of *particular* animals are characterised and then physically modelled aboard robots. This approach will not be relevant to all engineering problems, nor can we study all animal sensors in this way (many animal-environment systems are too complex or poorly understood for realistic models to be built). Nevertheless, there remain a great many systems which might be usefully studied *via* the construction of a bionic model. The relationship of the bionic approach to several familiar lines of thinking and experimental work is discussed in the following sections.

1.3.1 Philosophical underpinnings

Every research community inherits an extensive network of ideas from its predecessors. Below I briefly discuss three schools of thought which I understand to be particularly relevant to the bionics approach.

Gibson's ecological approach or 'Direct perception' [Gibson 66, Gibson 79].

Though Gibson is criticised for not developing his answers into causal explanations of perceptual mechanisms, he was certainly asking the right questions. He begins his theory of ecological optics with the basic, essential question: "If vision was gradually evolved, ... why? What benefits are conferred by seeing?" [Gibson 66]. Throughout his work he stresses the importance of understanding perception in terms of the tasks of perceiving systems.

Gibson viewed the environment as containing "affordances" (*i.e.* possibilities for action) and he believed that animals were structured — both genetically and through development — to register this "meaning" inherent in the environment. He was critical of both neurophysiology and more formal approaches to optics which he saw as complicating explanations of sensory systems by seeking the wrong kinds of explanations. In the language of ecological optics, an understanding of perception in terms of behavioural control requires the investigator to seek explanations, not in terms of *sense-stimuli* (*e.g.* "the seeing of light" or "the hearing of sound"), but *sense-information* ("the hearing of things by means of sound"). While the former can explain the response of receptors to stimuli, an explanation of perception would demonstrate how stimulus energy can carry or contain information (*i.e.* invariances in stimuli flux).

Enactive perception [Maturana & Varela 80, Varela *et al.* 91].

Like Gibson, Maturana, Varela and colleagues regard behaviour as the *structural coupling* of the agent to its environment. The idea of structured coupling is a biological notion based on evolutionary theory and these philosophers have developed an approach to perception that is rooted in a corresponding self-producing, or "autopoetic", view of behaving systems. The argument says that organisms interact in complex ways with their environments and these interactions affect both the internal structure of the animal and the structure of the environment. Both these structures change over time through mutual adaptations of species and ecosystems. Because these changes have accumulated in such complex ways (*i.e.* animal and environment are so intricately coupled), it is not possible to understand them except in the context of interaction. "This structure — the manner in which the perceiver is embodied — ... demonstrates how the perceiver can act and be modulated by environmental events" [Varela *et al.* 91].

The enactive approach claims that "the reference point for understanding perception is ... the way in which the nervous system links sensory and motor surfaces" [Varela *et al.* 91] and, not surprisingly, claims Brooks as its representative roboticist. In this respect, the bionic approach must precede the enactive approach because it puts an emphasis on understanding the surfaces themselves. (However, admittedly, distinctions between the two can often be difficult to draw.)

Phenomenology [Merleau-Ponty 62, Heidegger 61].

A brief mention is made here of phenomenology. Unlike many other roboticists, I draw attention to this branch of philosophy to round out the discussion, rather than support a particular tenet of AI. Phenomenology is most often called upon in discussions such as these because it has developed a rich vocabulary for discussing embodiment and, therefore, perception (as it is defined here). However, within Heidegger's notion of structure is the idea that environmental objects take on interrelated meanings. *A pencil*, for example, can offer affordances which evoke a writing behaviour. However, Heidegger also makes allowances for us to withdraw this ordinary relationship to our everyday behaviours and to interact with *this particular* pencil by, for example, utilising it for a different purpose (*e.g.* as a measuring device). Agre claims that, in this way, "Heidegger's work is not directly addressed to AI as it exists today, but rather a larger tradition of which AI is one part" [Agre 95b].

Toward the aim of a better integration of research within AI, it is appropriate to mention that several researchers have worked to extend the notion of structured coupling to other spheres. A nice example is Simon's early work on the emergence and organisation of social structure [Simon 57]. More recently, Kirsh describes and categorises a number of ways in which people manage their physical environments to assist in perception [Kirsh 95]. Hammond, Converse and Grass [Hammond *et al.* 95] take this idea further in an examination of longer-term relationships between agents and environments. They discuss how great complexity of behaviour can be achieved by piecing together situations which have been encountered before, and they suggest that people facilitate this through "stabilisation" of their environments — *e.g.* putting tools away when they are finished with a job so that the next job can begin from the same configuration. Agre and Horswill make a case study of the formalisation of cultural practices which proceeds by exploring operations on artifacts such as the movements of kitchen items through various states during preparation of customary Western meals [Agre & Horswill 92].

Although the research presented in this thesis takes a very low level approach to interactivity, I see this work as part of a wider tradition in AI which seeks to explain behaviour *via* various forms of characterisation. I believe physical characterisation to be a necessary first-step which has been largely overlooked.

1.3.2 Examples and direction

There are numerous examples of sensor models outside the AI literature (see, for example, articles in [Koch & Segev 89, Hawkins *et al.* 96]). These tend to be more rigorous in terms of their inclusion of biological detail, and, as such, are extremely valuable vehicles for formalising and communicating ideas. However, they are generally a different class of model than I am concerned with here in that they tend not to be constructed and tested in the world. Instead, their designers experience them more as an extension of their own understanding rather than autonomous systems generating observable behaviour.

While biological detail often needs to be sacrificed for computability, another aspect of realism is possible in robotic implementations because robotic models have an existence which is separate from that of the investigator. Agre sums this up in his statement that when "designing agents that operate in abstract search territories such as search spaces, and that do not have bodies (simulated or not) in any real sense, it can be easy to lose the distinction between what the agent knows about a situation and what the designer knows" [Agre 95b]. In this regard, physically embodying a model aboard a robot forces a rigorous conceptual separation between agent and designer. When this is possible, "the particularities of the agent's relationship to its environment can come out in full relief ... so that new theories about knowledge [*i.e.* perception] can arise that are rooted in the agent's having a body, being located in a physical environment, interacting with artifacts, and so forth" [Agre 95b].

In AI, research which best illustrates the bionic approach of embodying an animal sensory mechanism within a robot includes that of Franceschini and colleagues [Franceschini *et al.* 91], Kuc [Kuc 94, Kuc 96] and Webb [Webb 93, Webb 95]. In the work of Franceschini *et al.*, the authors built a compound eye consisting of an array of miniature hardware analogue circuits performing elementary motion detection (EMD). The receptor layout and spatial processing routines are based on those of the housefly studied by Franceschini [Franceschini *et al.* 85] and one of the stated aims of the research "was to investigate how a compound eye, inspired by insect vision and essentially consisting of motion sensors, might provide a mobile robot with a visual guidance in a realistic closed loop situation" [Pichon *et al.* 89].

Kuc developed a simulation model of a bat sensori-motor system using a pair of non-linear, time-variable, sampled-data controllers to alter the bat's heading by applying yaw and pitch corrections [Kuc 94]. The yaw correction attempts to lateralise prey by nulling the inter-aural intensity difference (received in two laterally displaced circular receiving aperatures) while the pitch correction acts to null the intensity difference between the overtone and fundamental components of the call. This work "illustrates how non-predictive tracking of an ideal prey can be accomplished with a very simple system" [Kuc 94]. Kuc has gone on to develop a robotic implementation of this system wherein the transmission rate and the configuration of the transducers (*i.e.* the pan angle of the transmitter and receivers) is adapted so as to centre a target in the frontal sound field where echo amplitude and bandwidth measurements are optimal [Kuc 96].

Webb positions her work more within the field of computational neuroethology; however, she also succeeds in designing and using a sensor of the type considered here.

Webb studied cricket phonotaxis by building an auditory sensor based upon the phase comparison operations performed by the female cricket *Gryllus bimaculatus*. This sensor was mounted aboard a mobile robot and used to derive steering signals to drive that robot toward a loud speaker playing recorded songs of the male cricket. Webb extensively tested the localisation abilities of the system under different acoustical conditions. The results provide some evidence in support of Schildberger's hypothesis [Schildberger 85] that the transduction and neural processing mechanisms underlying phonotaxis in crickets are tuned to the frequency and temporal (syllable rate) characteristics of the species-specific chirp such that the cricket inherently orientates toward a chirping conspecific.

The other contribution of Webb's work has been the specification of a methodological framework [Webb 93] for investigating perception. I have adopted and modified her methodology as follows.

- **Target system.** Choose a particular biological sensory system as a target.
- **Principled characterisation.** Characterise the relevant physics of the sensor and sensor motion in sufficient detail to facilitate the generation of hypotheses about the perceptual mechanisms underlying the animal's behaviour.
- **Implementation.** Build a model of the sensory mechanism in which to test behavioural hypotheses. Such a model must capture, realistically, the problems of interacting with the real world.
- **Evaluation.** Carry out experimental tests to evaluate the performance of the model system and examine the hypotheses embedded therein.

In the next chapter, I fill in the details of the four key methodological components of this investigation.

2. *Methodology*

In the previous chapter, I argued that, while much of biologically-based AI has embraced agent-environment interactivity as an important research problem, the focus of many current research programmes is on building *control architectures* which couple abstract sensors to abstract motors without incorporating a sufficient understanding of either. I claimed that this research may yield solutions which are overly complicated and, therefore, advocated that more research effort be devoted to investigating how sensory (and motor) systems *themselves* are coupled to the environment. Moreover, I claimed that a fruitful means for doing this involves modelling specific animal sensory mechanisms which have evolved to extract essential sense-information directly — thereby minimising subsequent computation. I refer to such a "life-like" artificial perceptual system as a *bionic* model.

In this chapter, I fill in the details of an investigative methodology (proposed originally by [Webb 93]) which I have adapted to study sensory perception *via* echolocation. In the following sections, I describe the particular animal perceptual system investigated in this work (Section 2.1), a physical characterisation of that system's interaction with its environment (Section 2.2), an implementation (*i.e.* physical model) of the target system (Section 2.3), and tools to evaluate the performance of the model and adequacy of the hypotheses (Section 2.4).

2.1 Target system

In this thesis, I investigate a prey localisation behaviour utilised by several species of Old World bats. Echolocating bats provide excellent systems for robotic investigation for two reasons. First, hearing in bats is one of the better studied biological sensory systems and a corpus of literature exists. Grinnell states [Grinnell 96]:

> *"From the beginning, the study of bats has been characterised by a synergistic partnership between behavioural and physiological approaches. The results have been nothing short of spectacular ... Echolocation, for all of its unanswered questions, is one of the triumphs of neuroethology — better understood perhaps, than any other complex mammalian behaviour, certainly auditory behaviour."*

There is also a great amount of engineering expertise and interest in echolocation systems (*e.g.* artificial SONAR and radar [Skolnik 62, Berkowitz 65]) and, yet, as stated by Peremans: "the large discrepancies between the robustness and performance characteristics of biological echo-perception strategies [Griffin 58] and those employed aboard artificial mobile agents [Everett 89] suggests that the latter are under-exploiting a valuable sensing resource" [Peremans 94].

Analysis

Because bats are dependent upon robust and accurate directional cues for survival in their three-dimensional, low visibility worlds, the mechanisms of sound localisation in echolocation are more highly developed and correspondingly more clearly observed in bats than in other kinds of animals where spatial perception is a less significant aspect of the use of hearing [Lawrence & Simmons 82]. In this regard, Grinnell states [Grinnell 96]:

> *It is not that bats are intrinsically easier to work with or that their nervous systems are more simple. On the contrary, the size of bats, the difficulty of obtaining and maintaining them in captivity, the inaccessibility of their native habitats, and their nocturnal activity in a territory that covers square*

> *miles all combine to make the study of bats daunting. What has made bats such superb subjects for the study of acoustic behaviour and neural processing is that they rely almost exclusively on hearing for the information they need about their environments ...*"

All auditory nuclei in bats have been studied in the single cell recording paradigm and some excellent neurophysiological work has led to the discovery of, for example, topologically ordered echo delay and velocity "maps" in the cortex (*e.g.* [Suga 90a]). Arrays of topologically ordered binaural neurons have been found in the central nucleus of the inferior colliculus which encode the inter-aural intensity disparities occurring across a bat's frontal sound field (see review in [Pollak & Park 96]). In addition to cataloguing the tuning characteristics of individual neurons, neural response to behaviourally relevant stimuli (*i.e.* echoes from prey) has been characterised, *e.g.* [Schuller 84, Ostwald 88].

Bat research has also hosted some ingenious psychoacoustical work. Some of the most famous discrimination experiments in biology have been carried out according to behavioural conditioning techniques, introduced by James Simmons, wherein bats are trained to discriminate between two targets presented simultaneously or sequentially in a 2-alternative forced choice (2-AFC) or yes/no paradigm. (In these studies, the bat is supposed to select a target in every 2-choice trial — which they indicate by crawling to one side or the other of a Y-shaped platform — and their performance level should reflect the relative difficulty of the task.) Real targets may be used, but it is more common for the experimenter to more carefully control the test by using "phantom" targets produced by recording emitted pulses with a microphone near the bat and playing back a modified phantom "echo" at a particular delay and angle from the loud speaker. (See review in [Moss & Schnitzler 96].)

These experiments are inherently sensitive to individual laboratory conditions and some experiments are criticised for not controlling sufficiently against signal level variations or taking into account the masking effects of echoes off the playback equipment itself. However, it appears to me that rather than seeing these phenomena as flaws of the experimental paradigm, they might be viewed as opportunities for investigation of bats' ingenuity (*i.e.* the animal's ability to find easier cues to use in the discrimination task). Indeed, researchers are often able to go back and explain a bat's performance based upon its exploitation of interesting and unexpected echo sources.

I conjecture that the real problem with this experimental paradigm (which is also its great strength) is the ability to control as much of the environment as researchers do. For example, one of the most exciting (and controversial) results to emerge from playback experiments (where phantom echoes were delayed replicas of the original), was that *Eptesicus fuscus* can discriminate jitter in target range of less than < 0.0017 mm (or 10 ns) [Simmons *et al.* 90a]. Few researchers met this amazing observation with the question "why?" — *i.e.* "for what purpose could *Eptesicus fuscus* employ this discriminatory power given that it, like other insectivorous bats, needs only about $1-3$ cm range accuracy [Trappe & Schnitzler 82] to successfully intercept prey? Is this resolution required for robust identification of targets? Is there sufficient evidence of prey selection by *E. fuscus* to support this possibility? Instead of asking questions rooted in *ecology*, reports of this remarkable temporal acuity in bats have provoked researchers to explore *optimality* — *i.e.* to experimentally probe the temporal boundaries of bat audition and analytically explore the possible similarities between bat and artificial echolocation (*e.g.* radar) systems.

Indeed, the role of the bioacoustician has been likened to that of the signal intelligence engineer: both seek to infer, from intercepted signals, the mode of operation and information collection preferences of an intelligence who is unable/unwilling to provide this information more directly [Pye 80]. This approach too often leads one to an analysis of bat calls along the yardstick of optimality used for *ideal receivers*[1]. For example, from comparison of behavioural response data from select playback experiments and autocorrelation representations of bat calls, some authors go so far as to say that evolution of echolocation in bats has driven the analysing capacities of mammalian audition (in both the frequency and time domains) "close to its theoretical limits". While it is possible that bats process echoes with a matched filter — the operations of the peripheral auditory system (*i.e.* bandpass filtering and envelope detecting [Siebert 68, Evans 77]) do not constitute irreversible processes which would

[1] Radar theory is based on the assumptions of an optimum receiver (*i.e.* one employing a matched filter, cross-correlation, the likelihood ratio or the inverse probability criterion) and a high signal to noise ratio ($\frac{2E}{N_0} = 1$).

make matched filtering impossible [Altes 80] — it seems highly unlikely that our notion of "optimal" is the same as that which nature stumbled upon. In other words, unlike the systems for which signal theoretic analyses were developed, echolocators designed according to an evolutionary process are likely to be extracting a complex constellation of acoustic information about which we have an insufficient understanding.

In this regard, an experimental paradigm in which it is possible to observe the interaction of a minimally assumptive receiver model within a realistic environment and task framework provides a valuable alternative experimental platform.

Synthesis

Some developers of artificial echolocation systems have followed in this tradition of over-generalising. We see an extreme in robotics where several researchers have confused themselves and the rest of the community by making claims that their systems possesses "bat-like" properties based on such tenuous links as a loose analogy with the use of time-of-flight to measure range. This is a shame because bats are amongst the most numerous and diverse mammals on the planet (see Chapter 3) and incorporation of knowledge of the particular perceptual strategies (*e.g.* matched and dynamic filtering characteristics) of *particular* species reveals insights that these general approaches cannot identify.

The work reported in this thesis is concerned with target detection and localisation mechanisms which may be used by narrow-band emitting echolocators with mobile receivers. My motivation to address this topic sprang from an interest in several species of Old World bats which are characterised by long constant frequency (CF) calls and highly directional and mobile pinnae. These bats (commonly referred to as "high duty-cycle bats" due to the length of their calls relative to those of other bats) includes species in the Old World families Rhinolophidae and Hipposideridae. (The relatively fixed pinnae *Pteronotus parnellii* — a bat in the New World family Mormoopidae — is also included in the non-taxonomic "high duty-cycle" grouping. See Appendix D for a brief description of these bats.) Most high duty-cycle bats hunt in acoustically cluttered environments and attack only moving prey (*i.e.* fluttering insects). They are uninterested in or unable to localise targets with non-periodic motion. The systematic movement of pinnae in rhinolophids and hipposiderids appears to be essential for good localisation of targets.

Neuroethological descriptions of the "circuits" underlying this behaviour are not presently available. Nevertheless, the behavioural consequences of perception are clear in the case of prey hunting. Moreover, because the acoustical filtering properties of the external and inner ear of these animals are relatively well known, it is possible to investigate this echolocatory behaviour at the informational level. In other words, one can measure the sense-information generated by receiver and target motion, and investigate how this sense-information might be used in target selective localisation. Such a study allows one to address two important questions.

1. What are the invariant cues that an echolocator might use to identify (in a robust, orientation independent fashion) an echo reflected from a fluttering target in an acoustically cluttered environment?
2. How, using only a single frequency echolocation call, can an echolocator localise targets in 3D? (What acoustical cues does receiver motion generate and how might those cues be transformed into spatial percepts?)

In the empirical work reported in this thesis, a 6 degree-of-freedom, binaural echolocation system is used to investigate the acoustic cues generated by receiver motion of the type employed by rhinolophid and hipposiderid bats localising fluttering targets. Moreover, the sufficiency of those cues to engender an echolocator with robust and accurate spatial percepts is tested aboard this model system. The physical characterisations which gave rise to this model are described in the next section.

2.2 Physical characterisation

From analytic consideration of receptor physics (*i.e.* the far field behaviour of ultrasonic transducers), I determined how receivers — moving through the same vertical arc scanning motions employed by rhinolophid and hipposiderids — impose amplitude and frequency modulations on target echoes.

As target motion also modulates echoes in the sensing task considered here, it was necessary to consider the combined effects of target and observer motion. For the study of target motion, reflector models were constructed in simulation so that the reflection process could be explored in more detail under simplified conditions. Specifically, by constructing a series of mathematically well-behaved "stick" insects, I investigated the way in which the spectrum of an echo reflected from a target with periodic motion varies across different target morphologies, motions, and observer viewing angles. Reasoning with these simple targets proved useful in placing a qualitative bound on the detection mechanism needed to identify the acoustic signature of real targets (which, despite their more complicated motions and morphologies, introduce similar patterns of AM and FM).

Based upon these physical characterisations, signal processing mechanisms were hypothesised to (i) extract target echoes in acoustically cluttered environments and (ii) transform these target echoes into relative 3D target angles. The key hypotheses to which these characterisations lead are summarised below.

- **Recognition is lunch (Chapter 4).** Targets with periodic motion amplitude and frequency modulate a long CF carrier so as to impose a pattern of spectral sidebands which encodes the characteristic target motion. Thus, because pursuit of targets with periodic motion involves the evocation of a specific response (turning toward) by a specific stimulus (periodically modulated echo), echolocators hunting exclusively for such targets may be so coupled to their environments that target selection occurs as a result of a single localisation mechanism tuned to compare echo energy in only those frequency channels into which such a favoured target reflects energy. In this way, target selection need not involve an intermediate "recognition" stage wherein echo energy is translated into a 3D target model or matched to some "insect category" *via* a cognitive comparison process.

- **Partial IID maps for complete target localisation (Chapter 6).** In the central nucleus of the bat's inferior colliculus there are binaural neurons which accept inhibitory inputs from lower auditory centres on one side of the head and excitatory inputs from the other so as to compute inter-aural intensity disparities (IIDs). These neurons are laid out topologically according to best IID value such that a sound source moving laterally across the sound field excites neurons in the dorsal region first and this locus of activity moves progressively toward the ventral region. It is hypothesised that upon this substrate, a map of the frontal sound field is assembled. I hypothesise that by moving its ears through opposing arcs, an echolocator equipped with such a "map" can alter the spatial region which the map represents and, in doing so, can capture both target azimuth and elevation.

- **Pinna mobility affords the creation of timing cues which do not degrade with head size (Chapter 7).** Continuously scanning pinnae also set up dynamic amplitude cues [Schnitzler 73, Pye 67] — *e.g.* rates of change of inter-aural disparities and delays in the onset of echo amplitude peak — both of which vary systematically with target elevation. I investigate Schnitzler's hypothesis that IID rates of change can encode target elevation and a hypothesis of my own: namely, that pinna mobility provides an echolocator with a mechanism for creating dramatic temporal cues (*i.e.* delays in the echo amplitude peak) for elevation localisation which, unlike inter-aural timing differences, do not degrade with head size.

- **Doppler shifts for elevation (Chapter 8).** Continuously scanning pinnae also induce frequency modulations which are related to a target's off-axis angular position through the cosine-law. I investigate Pye's hypothesis that, using Doppler cues, an echolocator with fine frequency sensitivity can partition echoes based on Doppler shift and thereby extract the angular position of a target [Pye & Roberts 70].

The model used to generate the data for testing these hypotheses mimics the external auditory apparatus of a bat using functionally analogous transducers, electromechanical motion control circuitry and digital signal processing/filtering. It is not an accurate, to-scale replica of the external auditory system of a bat (it is an order of magnitude larger than a typical high duty-cycle bat's head), nor does it model particular neural circuits in realistic detail (*e.g.* it does not incorporate the time-intensity latency tradeoffs of real neurons). Those details had to be sacrificed in favour of ease of construction and computability. The essential details which it captures include the following: (i) the use of long, ultrasonic CF pulses, (ii) reception of echoes in a pair of mobile receivers, (iii) processing of echoes in an idealised cochlear filter-bank (with appropriate Q_{10dB} values) and (iv) simple comparisons of binaural (and, in one case, monaural) localisation cues based on processing which is, or could be, performed by various populations of neurons in the peripheral auditory nervous system. Furthermore, the model operates in a realistically cluttered environment and localises targets with periodic motion like those used in laboratory experiments with bats.

The bionic system is described briefly below and in more detail in [Peremans *et al.* 97] (available as Appendix B).

2.3 Implementation

The bionic echolocation system used in these studies is a 6 degree-of-freedom ultrasonic stereo head which allows for panning and tilting of the neck, and independent panning and tilting of each receiver. (See Figure 2.1.) The key components include the following:

- **Motors.** The motors driving the different axes are standard radio-control model servos. Control signals, i.e. pulse-width modulated signals, are generated by a transputer.

- **Transmitter.** The ultrasonic transmitter currently mounted aboard the head is a Polaroid series 7000 transducers whose characteristics are summarised in Appendix A (see also [Biber *et al.* 80]). The transmitter driver has inputs for frequency and amplitude modulations to be imposed on a carrier wave. Alternatively, the FM signal generation part of the transmitter module can be bypassed and a digitally generated signal fed directly into the power amplifier. In the work reported here, this latter strategy is employed to create a 50 kHz echolocation signal with well defined phase characteristics.

- **Receivers.** Echo detection and amplification is performed by the receiver modules mounted behind their associated transducers (also Polaroid series 7000). The output signals from the receivers are sampled at 200 kHz. All further pro-

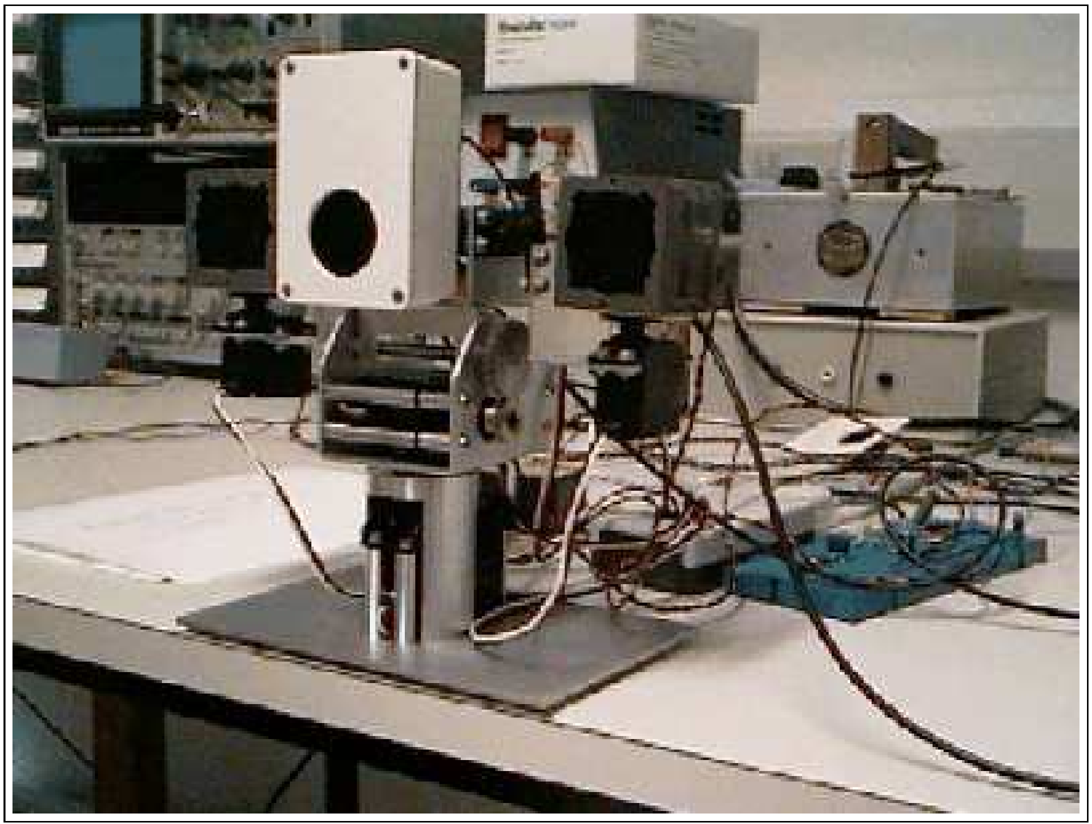

Figure 2.1: **Bionic SONAR system.** The robotic sensor head consists of a central transmitter (fixed to the head) and two independently orientatable receivers. All transducers are mounted in a common elevation plane and laterally separated by 8 cm.

cessing of the right and left received signals is performed on a transputer-based multiprocessor.

- **Signal processing.** The echo signal processing operations are modelled roughly on the processing performed by the mammalian cochlea [Siebert 68, Evans 77, Shamma *et al.* 86]. Received signals are first filtered by narrow bandpass filters centred at behaviourally relevant frequencies and with Q_{10dB} values based upon those of 'filter neurons' in the auditory system of rhinolophids. Bandpass filtered signals are full-wave rectified and then smoothed using a lowpass filter.

- **Simulator.** A simulation tool was developed to test algorithms before implementation on the robotic sensor head. The *3D Echolocation Simulator* (described in full in Appendix C) mimics the functionality of the robotic model, except that the user can control noise and clutter, and can define sensor head attributes (*e.g.* transducer size and separation, call frequency, *etc.*) and target types and positions. Because the simulator allows for the variation of the morphology of the sensor head which is not possible in hardware, results generated in simulation are juxtaposed with measured results where their use provides a more meaningful comparison with the small heads of bats.

Numerous assumptions underlie the operation and use of the bionic system as a model of an animal system. They are justified as they arise in Parts II and III and are reviewed again in Section 9.1. The next section describes how the strength of these assumptions and the behaviour of the system (in which they are embedded) are evaluated.

2.4 Evaluation

In A. Conan Doyle's *A Scandal in Bohemia*, Sherlock Holmes tells Dr. Watson: "It is a capital mistake to theorise before one has data. Insensibly one begins to twist facts to suit theories, instead of theories to suit facts." Indeed, although one may agree (at least in principle) that the greatest utility of a model "is not to fit the data but to sharpen the questions" [Samuel 83], models based upon insufficient data can be misleading.

The extent to which we can infer, from a model, the behaviour of its animal referent (or sub-systems thereof) "is constrained entirely by the mapping between the system and the model" [Webb 93]. Likewise, the extent to which perceptual mechanisms — plucked out of one system (*e.g.* an animal) and implemented in another (*e.g.* a robot) — will facilitate similar behaviour in the second system is dependent upon how well the essential aspects (such as matched and dynamic filtering) have been characterised and incorporated into the artificial model. As argued above, investigating perception *via* a principled physical characterisation and model can provide a powerful platform from which to test the sufficiency of hypotheses about behaviour generating mechanisms; however, determining the strength of the evidence provided by the result of modelling is difficult and dependent on one's ability to answer the following questions [Webb 93].

1. How accurately does the physical model represent the hypothesised system?
2. What results does the model produce and how can those results be interpreted? ("Why does it work?")
3. Can that interpretation be properly compared to the behaviour of the target system?

These questions are addressed in Chapter 9 — where the results of the empirical investigations reported in Parts II and III are evaluated as hypotheses about perception in high duty-cycle bats.

The results of a biological modelling exercise can flow in both directions between the biological sciences and engineering. However, results must be made relevant to a particular field by evaluating the behaviour of the model system so as to extract answers and explanations which are appropriate to each individual field. Therefore, in Chapter 10, the bionic system is reconsidered in terms of its potential contribution to the growing literature on artificial echolocation systems for autonomous robots. Finally, in Chapter 11, the work is discussed as an example of an AI model of a biological system. Benefits arising from this approach to the study of perception are reviewed in that chapter.

Part II

Structuring sound at the target

"I am convinced that there are universal currents of Divine Thought vibrating the ether everywhere and that any who can feel these vibrations is inspired ..."

—Richard Wagner

"... the universe is probably full of music which we cannot perceive ..."

—Sir John Lubbock

"Good, good, good, good vibrations..."

—The Beach Boys

3. Background

The term "echolocation" was coined in the first half of the present century by Donald Griffin to describe bats' ability to localise targets based on the acoustical information contained in reflections of their own emitted sound pulses [Griffin 58]. It is now widely used to describe the active use of sound in sensing by other animals and by artificial systems. Although many echolocation (or *active SONAR*[1] systems, as they are sometimes called) emit ultrasonic frequencies for their directional properties, this is not a necessary characteristic. A more accurate restriction on the use of this mechanism would draw distinctions based upon wave *amplitude* rather than *frequency*. High energy acoustic vibrations can be harnessed to bring about a change in a medium by using the agitation of the waves to carry out this work directly (*e.g.* cleaning and drilling) or indirectly through the phenomenon of cavitation. Echolocation/active SONAR is a low amplitude technique concerned with measuring the effect of an environment on the acoustic wave. Both the terms "echolocation" and "active SONAR" have become heavily associated with range sensing, to the extent that some authors have attempted to impose alternative terms to more fully describe the capabilities of this sensory modality. For example, Neuweiler proposes that because "the term echolocation does not encompass the full capacity of this acoustical information system, ... a term analogous to visualisation, such as audification, would be more appropriate" [Neuweiler 90]. However, to maintain consistency with the established literature, I use the traditional terms in this thesis.

[1] SOund Navigation and Ranging [Blitz 63]

The perceptual principle underlying echolocation is that echoes contain an altered version of the known frequency, amplitude and temporal structure of an emitted signal (*i.e.* a call) and, therein, encode a description of reflecting target(s). In realistically complex sensing scenarios, matching measured effects unambiguously to causes is not straightforward — *e.g.* target size and texture, as well as target range and angular bearing all conspire to determine echo attenuation. Moreover, there are a variety of interference effects which further increase the uncertainty in attributing causes to effects: in addition to sensor noise, an echolocator has to contend with the effects of acoustical clutter (*i.e.* reflections from superfluous structures), "jamming" signals from conspecifics (and prey), atmospheric attenuation of signals (which itself depends upon a complex constellation of factors including frequency, temperature, humidity and the molecular composition of the atmosphere). The echolocation systems of animals have evolved as niche specific adaptations which let them specialise in extracting a limited set of probable causes from measured effects.

That echoes can deliver rich sense-information at a high rate is demonstrated by echolocating animals (*e.g.* bats and dolphins) — who perform some of the most exacting 3-dimensional, rapid target localisation tasks of which we have knowledge. All species known to employ echolocation are warm blooded and regions of the auditory nervous systems devoted to processing echoes are well developed. In addition to their robust performance, we might also marvel at the variety of evolved echolocation systems. Echolocation is polyphyletic in vertebrates, and it has appeared more than once in both the Aves (*i.e.* orders Caprimulgiformes and Apodiformes) and mammals (*i.e.* orders Insectivora, Cetacea and Chiroptera, and possibly others) [Fenton 84].

Echolocating bats (sub-order Microchiroptera) provide examples of the most sophisticated airborne echolocation systems in existence. (Good general reviews and collections of seminal papers on echolocating bats include [Griffin 58, Busnel & Fish 80, Nachtigall & Moore 88, Fay & Popper 96].) Moreover, the approximately 800 different species of Microchiroptera are among the most diverse of the living mammals: bats range in size from 2 g to more than a kilogram, with wing-spans varying from 20 cm to more than 2 m, and have a range of diets including insects, fish, amphibians, fruit, pollen, nectar, and blood. They are also remarkably well distributed geographically —

with habitats on every continent except polar regions.

The diversity of Microchiroptera is clearly reflected in their echolocation signals. In contrast to the short, broad-band clicks of some birds, shrews, many cetaceans, and the one echolocating megachiropteran (*Rousettus aegyptiacus*), the signals of Microchiroptera consist of a wide variety of different call structures: from broad-band to tonal, short (< 1 ms) to long (> 100 ms), vary in intensity from less than 60 dB SPL to more than 100 dB sound pressure level (SPL) (as measured at 10 cm), and span a range of frequencies from 10 to more than 200 kHz.

There has been some effort to characterise echolocating bats by the predominant time-frequency profile of their calls (*i.e.* as tone emitting "CF bats", or frequency modulating "FM bats"). The use of echolocation calls as the basis for preparing phylogenies of bats has been proposed [Simmons & Stein 80]. However, inter-species overlap in these parameters and intra-species geographic/habitat variability has made this difficult. Moreover some species actively alter the intensity, duration and/or frequency components of their calls in response to their changing information-gathering needs and interference problems [Griffin 58, Novick 77, Simmons & Stein 80]. Nevertheless, for the sake of this discussion, it is useful to generalise. What appears to be the least contentious way of dividing Microchiroptera is according to call length or duty-cycle [Fenton 96].

1. *Low duty-cycle bats.* These bats emit low duty-cycle ($< 20\%$) calls consisting of frequency modulated (usually linear frequency or linear period) pulses. Most typically, pulses are swept through about an octave over < 5 ms. (One or more harmonics may increase this bandwidth.) Pulses may be shorter in duration (< 2 ms), in which case the signal is often described as a low amplitude, multi-harmonic burst. Some species switch between the pulses and the bursts — using the former during the search phase of a prey capture cycle and the latter during attack. (An example of such a pulse train is shown in Figure 3.1 (a).) A third, less common low duty-cycle call structure consists of a shallowly modulated pulse lasting up to 10 ms. Species using this call structure are relatively poorly studied, but it appears that the longer, more narrow-band sound is used primarily in open areas, and that pulses are shortened and become increasingly broad-band upon

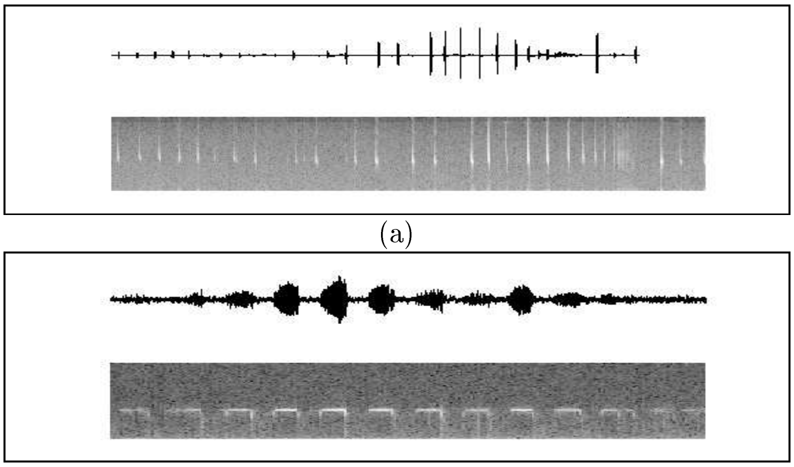

(a)

(b)

Figure 3.1:
Time waveform (above) and spectrogram (below) representations of biosonar pulses. (a) Several calls of *Pipistrellus pipistrellus* (unconfirmed) recorded in northern Italy, summer 1992. (b) Long CF-FM calls of the greater horseshoe bat, *Rhinolophus ferrumequinum*, recorded in northern Italy, fall 1992. (Sounds from the collection of *Centre Interdisciplinare di Bioacustica e Ricerche Ambientali, Universita' degli Studi di Pavia.*)

detection (or in confined areas).

2. *High duty-cycle bats.* A smaller number of species belonging to 3 families (Rhinolophidae, Hipposideridae and Mormoopidae) emit high duty-cycle ($> 45\%$) calls consisting of a tone (or several CF harmonics) of duration $10 - 100$ ms. (Upon detection of a target, duty-cycle may increase to $70 - 80\%$ *via* an increase in the duration of individual pulses or a decrease in inter-pulse intervals.) In many rhinolophid and hipposiderid species, most energy is emitted at second harmonic (CF_2). This is also true of *P. parnellii*; however, this New World bat also emits significant energy at the fundamental, third and fourth harmonics of the call. In most high duty-cycle bats, the CF call component is often followed by a short, downward FM sweep and may be preceeded by a swift upward FM sweep (of linear or curved time course). A typical example of a series of *Rhinolophus ferrumequinum* calls is shown in Figure 3.1 (b).

The remainder of this chapter reviews the way in which different species solve the potential problem of call/echo interference and considers the kinds of target descriptions (or signatures) available within broad- and narrow-band echoes. It is divided into two parts to reflect the two different sensing strategies employed by bats with different duty-cycle calls. Sections 3.1 looks at "frequency subtractive" approaches wherein emitted broad-band calls return filtered such that information about reflecting targets is contained in the frequencies which are "missing" (*i.e.* selectively attenuated) upon reflection. Sections 3.2 looks at how frequencies are added to a tonal call (through amplitude and frequency modulation processes) which arise when there is relative motion between the bat and reflecting surfaces.

3.1 Frequency subtractive approach

3.1.1 Separating call and echo energy

One of the first questions which intrigued researchers in the study of echolocation was how bats prevented themselves from becoming deafened by their own pulse emissions. Like other vocalising mammals, bats reduce the sensitivity of their hearing during emission by contracting middle ear muscles. (Although not considered here, neural inhibition plays a role as well [Suga & Schlegel 72, Suga & Shimozawa 74].) For example, *Tadarida* begins contracting the middle ear muscles a few milliseconds before pulse emission such that reduction in sensitivity is maximum (20−30 dB as measured by cochlear microphonics) at approximately the onset of emission [Henson 65, Suga & Jen 75]. Relaxation begins soon thereafter and sensitivity is gradually restored to maximum levels 5 − 8 ms after the pulse. For an echolocator, this mechanism also provides a form of *automatic gain control* [Kick & Simmons 84] — *i.e.* echo intensity decreases associated with increasing range are proportionately compensated for by this gradual increase in auditory sensitivity. (The full importance of this stabilisation of intensities is not well understood; however, localisation of the angular position of a target based on inter-aural intensity comparisons is one of the features that may be more effectively analysed within an environment of stabilised absolute intensities (See Section 5.1.2.)

To ensure that pulse and echo do not overlap, low duty-cycle species shorten pulse duration and increase pulse repetition rate when approaching a target [Schnitzler & Henson 80]. For example, a bat may emit one echolocation pulse per wing-beat in the early phases

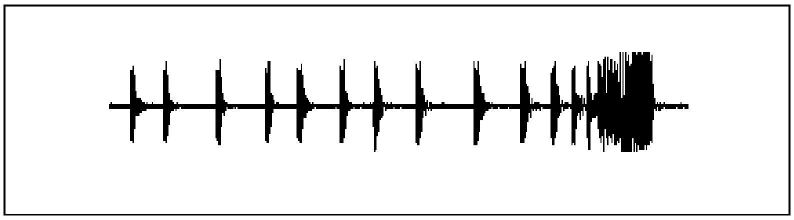

Figure 3.2:
Time waveform of the feeding buzz pulses emitted by ***Lasionycteris noctivagans*** **as it converges on an insect.** (Recording by Burr Betts.)

of a prey capture cycle (*e.g.* one pulse every approximately 50 – 300 ms). Pulse emission rates may increase dramatically — to an approximately 5 ms "feeding buzz" [Griffin *et al.* 60] — as a bat converges on a target. Figure 3.2 shows a heterodyned recording of a *Lasionycteris noctivagans* feeding buzz.

3.1.2 Echoes from targets

In the late 1950's, experiments investigating prey capture by low duty cycle bats were moved into the laboratory where broad-band emitting *Myotis lucifigus* were simultaneously strobe photographed and recorded as they captured small insects and meal worms thrown into the air by a "meal worm gun" [Griffin 58, Webster & Brazier 65, Webster 67]. The results astonished researchers: "Not only could these bats localise airborne targets accurately enough to catch them with the tail membrane or the tip of a purposely outstretched wing, but also identified and caught meal worms in the presence of a number of other targets of similar size but different shape (*e.g.* disks or spheres) that contributed interfering signals ..." [Grinnell 96].

Although much has been discovered over the subsequent 40 years of investigation, it still remains an active research question as to how insects (and other complex 3D range extended targets) are encoded by the nervous system of bats. Targets of the size caught by most insectivorous bats reflect echoes which consist of a series of *glints*[2] reflected from different scattering surfaces across the body of the target. Differences in

[2] The term glint is used here loosely analogously with vision. However, an *acoustic* glint is a delayed and weighted impulse response of a point scattering target [Kober 82]. The term "frequency glint" is used to describe a peak spectral broadening in a Doppler shifted echo.

depth and angular position of the scattering surfaces are so small that glints overlap for most of their duration. (For simplicity, one should imagine that a bat is unable to resolve the angular positions of individual scattering surfaces, so that the multiple-wavefront echo encodes all surface topology as the time delays between wavefronts originating at different scattering points.) The result of this superposition is that constructive and destructive interference between these overlapping echoes produce patterns of amplitude modulation (*i.e.* troughs and peaks in echo intensity) across the bandwidth of the call.

The structure of broad-band echoes reflected from insects have not been well characterised (although insonification studies carried out within the past decade provide some insight [Kober 88, Kober & Schnitzler 90, Moss & Zagaeski 94]). However, it has been shown that two-wavefront generating targets (such as a plate with a fixed depth hole drilled into it) can be characterised by a pattern of periodic minima ("notches") in the spectrum — whose periodicity is related to the time delay (*i.e.* range difference) between wavefronts [Simmons *et al.* 74, Beuter 80]. Figure 3.3 shows an example of the spectral profile resulting from interference of 2 echoes reflected from targets separated by a delay of less than 350 μs (*i.e.* < filter integration time). This type of signal has been used as a stimulus in two sets of 2-AFC experiments aimed at characterising the processing performed by a broad-band emitting bat's "SONAR receiver".

In the first set of investigations, bats were asked to discriminate between two-wavefront echoes that differ in the range/delay of their individual echo components. *Eptesicus fuscus* [Simmons *et al.* 74, Mogdans *et al.* 93], *Myotis myotis* [Habersetzer & Vogler 83] and *Megaderma lyra* [Schmidt 88, Schmidt 92] performed the task with better than 1 mm resolution. However, discrimination performance was not a monotonic function of the difference in wavefront delay separation: discrimination between particular pairs of delay separations, *e.g.* 12 and 36 μs, fell to chance. In this case, the result was explained by the presence of a prominent spectral notch at approximately 42 kHz which results from the superposition of echoes with either of these internal delays. Schmidt proposed that a "spectral correlation" receiver model best accounted for this performance [Schmidt 92]. Similarly, in discrimination experiments involving a one-wavefront echo *vs.* a two wavefront echo (which varies in the range/delay separation of its indi-

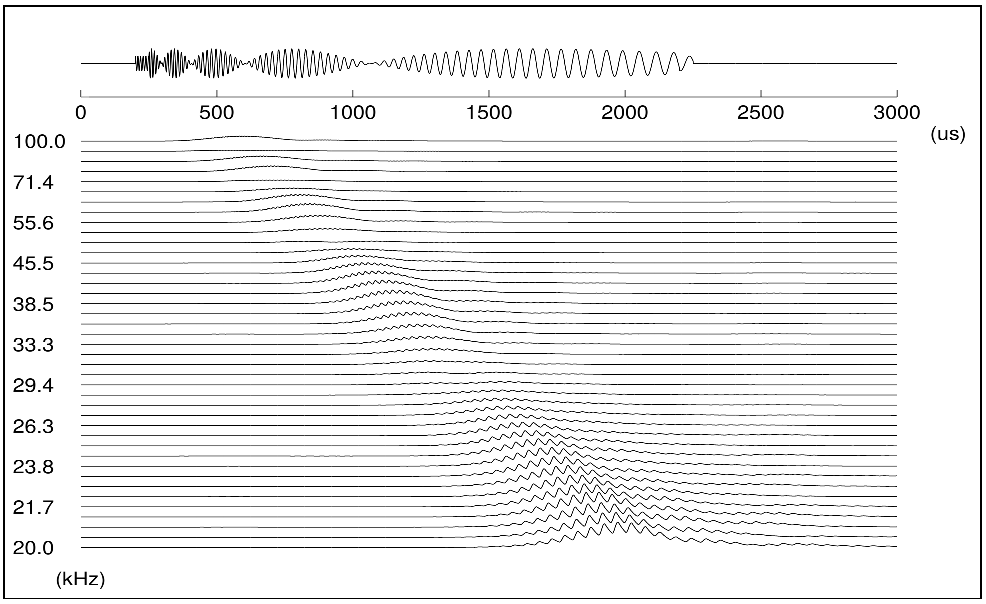

Figure 3.3:
Lowpass outputs of a filter-bank for an input echo containing glints from 2 **simulated point-like reflectors.** Reflectors are separated by a distance corresponding to less than 350 μs (*i.e.* < filter integration time) [Peremans & Hallam 98].

vidual components), Mogdans *et al.* observed discrimination performance which was related to to the magnitude of spectral differences. However, these authors are cautious in their interpretation of an underlying receiver model: "the time-domain and frequency-domain models are mathematically related through the Fourier transform, and ... it is not possible to differentiate between the two models using the bat's performance in the behavioural task" (Mogdans *et al.* reported in [Moss & Schnitzler 96]).

Based on the bat's near chance performance for particular delay offsets, Mogdans and colleagues did conclude that the bat does not perceive the two-wavefront echo as containing distinct glints. By contrast, a similar investigation by Simmons *et al.* showed that discrimination errors increased when a one-wavefront echo is presented at the same delay or amplitude as one of the components of the complex echo [Simmons *et al.* 90b]. From this result, the investigators proposed that the bat's SONAR receiver works to convert the spectral cues into a series of delays. More recently, Simmons and other coworkers presented a model that combines a time domain, coarse time information, and frequency domain, fine time information, analysis of the received signal to arrive at the temporal decomposition [Saillant *et al.* 93]. The SCAT receiver consists of three blocks. A *cochlear block* derives a spectrogram representation from the original signal by analysing it in a number of frequency channels. A *temporal block* uses the timing information about corresponding events in the different frequency channels to derive an estimate for the arrival time of a group of overlapping echoes. This arrival time of the group is then further refined by a *spectral block* which uses the relative amplitudes of the outputs of the different frequency channels to derive estimates for the delays between the individual echoes that make up the group. It is suggested that, in principle, the SCAT model could resolve all delays in a multiwaveform echo, but the present model can only support two wavefronts. (Further work by Peremans and Hallam highlights some other consequences of their model [Peremans & Hallam 98].)

A receiver model determined in this way — were one able to arrive at it — would constitute only a partial explanation of behaviour. It is far from clear how a series of delays (or a spectral profile) might be translated into the neural signals which steer a bat towards a target. However, the predominant view found in the literature is that the astonishing prey capture abilities of bats must be mediated by formal target recognition — *e.g.* the assemblage of a 3D neural target model. Toward this view, Grinnell states [Grinnell 96]:

> *"It is difficult to escape the conclusion that bats are able to form an acoustic image of the world in front of them. It appears that the echoes from each pulse, returning from different directions and delays corresponding to the distance of different targets, are processed by the brain to provide a three-dimensional map. This map is comparable to the visual image produced by the flash of a strobe light and tells the bats simultaneously about the location of potential prey and all other objects within hearing range, such as tree branches, wires, buildings, and the ground."*

If this is true, with increased bandwidth more information is available and, consequently, a target becomes more discriminable. However, some caution must be exercised in considering the proposition that high bandwidth transmitting and receiving capabilities are necessary for robust detection of complex range extended targets. First, this description must explain what effect the superposition of clutter has on the recognisability of these targets. When foraging close to vegetation or ground/water surfaces, bats must detect their prey within a clutter of time-smeared echoes reflected from these structures [Neuweiler *et al.* 80, Schnitzler & Henson 80]. Also, we have assumed a stationary target and stationary bat when in reality both may be in continuous motion. Does this characterisation encode an insect description which is motion invariant?

Though the characterisation discussed in this section is the best developed and most general, other explanations of the interactions of *particular* bat species with *particular* targets suggest the use of different cues. In this regard, echolocating bats provide examples of a variety of alternative strategies for dealing with motion and clutter. For example, just prior to attacking targets against structured background, some low duty-cycle bats do not employ a feeding buzz but, instead, rely on prey generated sounds (*e.g.* insect scrambling) to capture targets [Fenton 96]. Another type of motion detection is used by *Noctilio leporinus* (a species which pursues only jumping fish): interference between echoes from the body of the fish and the water disturbances it creates by jumping produce a spectral pattern that the bat appears to use to identify its prey [Schnitzler *et al.* 94]. Moreover, several high duty-cycle bats hunt exclusively for insects and exploit the fact that when a long, predominantly CF, probing beam falls on an insect's wings, the periodic changes in position (and velocity) of those reflecting surfaces rhythmically modulate the amplitude and frequency structure of the call in synchrony with the insect's wing-beat [Roeder 63, Schnitzler 78, Schnitzler & Henson 80, Schnitzler *et al.* 83, Schuller 84]. The latter is an example of the frequency additive approach which I will discuss in the next section.

3.2 Frequency additive approach

3.2.1 Separating call and echo energy

Analysis of periodic signal modulation improves with increased signal length. In the case of high duty-cycle bats, calls which are long enough to capture at least one complete insect wing-beat cycle ($20 - 100$ Hz for many species of nocturnal European insects) will overlap with echoes under most hunting conditions. (This explanatory order need not imply that evolution worked in this order, *i.e.* long calls may have originally evolved for another purpose.) Whereas echolocators using short signals may achieve call/echo separation in the time domain by only "turning on" their receivers after the pulse has been transmitted, echolocators which employ long pulses ($t > 2r/c$, where r is target range and c is the speed of sound) must find alternative means for separating calls and echoes. Call/echo isolation in a stationary echolocation system (*e.g.* ground based radar) may be achieved by physically separating transmitting and receiving transducers, inserting absorbing materials and/or mechanical baffles between transducers, and/or manipulating transducer directionality. Mobile systems have too many other constraints on their morphology to find complete solutions *via* these mechanisms. However, a solution is available in the signals themselves (or rather, in the physics of the sensory paradigm) in that, although long calls and echoes may overlap in the time domain, they may not overlap in the frequency domain. In other words, an elegant means for separating emission and echo can be constructed by exploiting the spectral broadenings caused by relative motion between target and sensor. This appears to be the approach of the *high duty-cycle* bats, whose duty-cycles can be as large as 80% with pulse-echo overlap tolerated (even exploited). (Note, pulse/echo overlap is only tolerated for the CF portion of the call. High duty-cycle bats decrease the FM component of their signals as they converge on a target such that the FM components of call and echo do not overlap [Schnitzler 68].)

The cochleae of these bats are larger than those of less specialised low duty-cycle bats and possess traits which appear to have come about during relatively recent evolutionary developments [Kössl & Vater 96]. (Relevant descriptions of the major auditory nuclei discussed in this section are given in Section D.2.) Specifically, there are very narrowly tuned bandpass "filter neurons" [Neuweiler & Vater 77] with *best frequencies* (BF)[3] just above the frequency of the dominant second CF harmonic of the call [Suga *et al.* 75, Suga *et al.* 76, Suga & Jen 77, Vater *et al.* 85, Feng & Vater 85, Kössl & Vater 90]. For example, filter neurons in *R. ferrumequinum* can reach Q_{10dB} *values*[4] up to $400 - 500$ [Neuweiler 80], as compared to $Q_{10dB} = 5 - 30$ in neurons of

[3] Best frequency is defined as the frequency corresponding to the neuron's minimum threshold.

[4] Q_{10dB} is defined as the quotient of BF and bandwidth of the frequency threshold curve 10 dB above

other frequency ranges [Suga *et al.* 76, Möller *et al.* 78]. In addition to sharp tuning, the cochleae of Rhinolophidae, Hipposideridae and *Pteronotus parnellii* (though this has not been observed in other Mormoopidae) deviate from the uniform pattern of logarithmic frequency spacing characteristic of many mammalian cochleae in that the narrow frequency band around CF_2 is represented on the basilar membrane (BM) in a vastly expanded fashion covering approximately 30% of its length (a length otherwise used for a complete octave). Moreover, the cochlea of high duty-cycle bats is densely innervated by the auditory nerve in the basal region representing CF_2 and this region is separated from the ordinary hearing range by a sharp peak of insensitivity [Grinnell 70, Neuweiler 70]. Due to the tonotopic organisation of the primary auditory pathway, this over-representation and sharp tuning (and, therefore, fine frequency analysis) is conserved throughout the bat's auditory system, as the cochlear surface is re-mapped upon each succeeding auditory region [Suga *et al.* 76, Schuller & Pollak 79, Schweizer 81, Pollak & Bodenhammer 81, Ostwald 84, Pollak *et al.* 86].

The presence of these unique features has caused authors to remark that the auditory systems of high duty-cycle bats contain a frequency magnifier or "acoustic fovea" [Schuller & Pollak 79]. The analogy with vision is even clearer when the integrated actions of both sensory and motor systems are considered. In vision, changing the gaze keeps images of interest fixated on the fovea. In echolocation, high duty cycle bats use a behaviour known as *Doppler-shift compensation* [Schnitzler 68, Schuller *et al.* 74, Neuweiler *et al.* 80, Schnitzler & Henson 80] to reduce some of the Doppler shift induced by their own flight speed and thereby ensure that the reflected CF_2 component of the echo arrives at a frequency which can be heard by the foveal neurons. The Doppler shift compensation mechanism corrects for positive shifts up to $4-8$ kHz, which would allow the bat to correct for approximately twice its average flight speed, or its own speed and that of an approaching target. (Negative shifts are not compensated for.) The sophisticated feedback system for locking the echo carrier frequency to the reference frequency only functions when emitted sounds and echoes overlap [Schuller 77]. There is some evidence to show that, in *R. ferrumequinum*, emissions in the range $78-82$ kHz facilitate foveal neurons ($82-86$ kHz) by increasing firing rates and

minimum threshold.

lowering thresholds by as much as 20 dB [Neuweiler 80].

These adaptations have a number of interesting consequences. First, as stated previously, they facilitate frequency-domain separation of call and echo. Using Doppler compensation, a high duty-cycle bat can position echoes within the fovea while the call is heard by less sensitive neurons tuned to lower frequencies [Long & Schnitzler 75, Schuller 80]. For example, *R. ferrumequinum* emits pulses having a CF_2 component of between $81 - 83$ kHz, regulated to an accuracy of ± 100 Hz in an individual bat (Schuller *et al.* refer to this as the "resting frequency" [Schuller *et al.* 74]). In flight, echoes will be Doppler shifted and the bat will lower the emission frequency to hold the frequency of the returning echoes at a "reference frequency" within the range of BFs of the acoustic foveal neurons ($82 - 86$ kHz). This is done with some precision. In *R. ferrumequinum*, the reference CF_2 echo frequency is maintained with an accuracy of ± 50 Hz (or a reference deviation of 0.06%) [Schuller *et al.* 74].

Secondly, analysis of echoes within the fine frequency window of the acoustic fovea confers on the bat an ability to partition echoes from different targets on the basis of subtle frequency differences. For example, targets located at different relative positions will experience different Doppler shifts and it is possible that a bat extracts directional cues for obstacle avoidance [Müller & Schnitzler 97] and target localisation (see [Pye & Roberts 70] and Chapter 8) by exploiting the cosine-law dependence underlying the Doppler shift phenomena.

The remainder of this and the next chapter are devoted to describing a kind of echo partitioning which arises from target, rather than observer, motion. (In this regard, Doppler compensation serves to decouple frequency addition due to target motion from that due to the bat's motion.) Specifically, I look at how the detection of periodically modulated target echoes within the high resolution window provided by the acoustic fovea may facilitate the identification of fluttering prey *via* mechanisms which are robust to the potential problems associated with clutter and target motion.

3.2.2 Echoes from targets

Clutter resistance

The characteristic amplitude and frequency modulation pattern that a fluttering insect reflects is distinct from all other echo modulations in nature — *e.g.* movements of non-insect reflectors or non-fluttering insects. Amplitude peaks occur when the reflective portion(s) of the wings are in a position perpendicular to the impinging sound waves [Schnitzler *et al.* 83]. At the moment of glint, reflections may be 20 dB (up to 40 dB in the extreme) greater than those off stationary body parts [Kober & Schnitzler 90]. Furthermore the relative velocity of reflective portions of the wing introduces periodically changing shifts in frequency due to the Doppler effect. As in amplitude modulation, there are peaks in the Doppler shift profile ("frequency glints") which occur when the component of the wing velocity relative to the bat is highest. The waveform describing the Doppler modulating signal takes on negative (*i.e.* drops below the carrier frequency) as well as positive values, depending upon the relative direction of wing movement (*i.e.* away from or toward the bat) during the wing-beat cycle. Degrees of spectral broadening may be as much as $\pm 1 - 3$ kHz in typical modulation processes involving insect prey such as moths, beetles and flies with wing lengths ranging between $3 - 30$ mm and wing-beat rates between $20 - 100$ Hz [Kober 88, von der Emde & Menne 89, Kober & Schnitzler 90, von der Emde & Schnitzler 90].

If a bat could discriminate these different kinds of echo modulations, it would have a reliable cue for prey detection even among strong background clutter. Foveal frequency resolution suggests this kind of analysis is possible — given the Doppler compensatory mechanism to ensure that interference from reflections off stationary and slowing moving background objects do not enter the fovea. Moreover, the Doppler compensatory mechanism acts sufficiently sluggishly that the bat only compensates for Doppler shifts of complete echoes and never reacts to short shifts, occurring within the pure tone, which actually provide the information that the bat needs to analyse [Schuller *et al.* 75]. (When the modulation frequency is cyclical and the rate is above approximately 20 Hz, the emitted frequency is kept at the lower end of the compensation range so that the highest frequency components in the echoes stay near the reference frequency.)

Thus, by using their calls as carriers for the frequency and amplitude modulations imposed by wing flutter, high duty-cycle bats can partition echoes into periodically and non-periodically modulated echoes so as to detect the faint echoes reflected from insects in environments where they might otherwise be obscured by prominent background clutter [Pye 67, Schnitzler 70, Schuller 72, Johnson *et al.* 74, Neuweiler 83, Neuweiler *et al.* 88, Pollak & Casseday 89]. Various observations on the feeding behaviour of high duty-cycle bats show that, indeed, they do hunt in most acoustically cluttered environments. Rhinolophids usually forage near the ground, close to vegetation, and along walls [Eisentraut 50, Brosset 66, Jones & Rayner 89]. Hipposiderids, *H. speoris* and *H. bicolor*, for example, enter into the foliage of trees [Neuweiler 83] and *P. parnellii* has been observed hunting mainly within 3.5 m of the ground under a canopy of foliage [Bateman & Vaughan 74]. These environments are rich in insect populations, but detection is made exceedingly difficult by the presence of numerous other strong reflecting targets.

Neuweiler summarises how the use of motion cues provides the bat with an inherently clutter resistant mechanism [Neuweiler *et al.* 80]:

> *"... a bat searching for flying insects close to foliage or any other dense background will receive a sequence of time smeared echoes destroying signal structure (e.g., FM-sweeps) and deteriorating the detectability of the prey. However, the time smeared echoes of a pure tone signal retain the signal structure, namely a pure tone, as long as the background is stationary whereas any moving object in front of the dense target will 'pop up' in the pure tone as a complex small modulation. Even when the background is randomly moving (e.g., foliage) wing-beating prey may be easily discriminated against the random modulations of the echo (caused by background movements) by the periodicity of modulations coming from beating wings. The latter argument may be also the reason why these bats use very long pure tones in search flight."*

Motion invariance

High duty-cycle bats appear to have become so specialised for the task of extracting periodic target motion that some species only pursue periodically oscillating targets. In the laboratory, *R. ferrumequinum* catches flying moths and moths fluttering their wings while attached to a surface; moths which do not flutter are not pursued [Schnitzler & Henson 80, Trappe & Schnitzler 82]. *H. ruber* [Bell & Fenton 84] and *P. parnellii* [Goldmann & Henson 77] pursue only fluttering prey (in flight or on surfaces) and abort pursuit when insects stop fluttering. *R. rouxi* and *H. speoris*, when presented with stationary insects or insects scrambling on the ground, do not react; however, when presented with wing-less tethered coachroach nymphs vibrating vertically up and down at an appropriate rate, they catch the prey (although they prefer wing-beating cockroaches to the yo-yoed nymphs [Link *et al.* 86]). Periodically vibrating loud speaker membranes attract rhinolophids [Schnitzler & Flieger 83, von der Emde & Schnitzler 86] and a propeller rotating so as to produce insect-like fluttered echoes has been shown to attract *R. ferrumequinum*, *P. parnellii* and *H. ruber* (see also reviews in [Schnitzler & Henson 80, Ostwald *et al.* 88, Moss & Schnitzler 96]).

The specialisations which increase flutter sensitivity may also facilitate *prey selectivity.* Because insects of the same species share similar flutter characteristics (*e.g.* wing morphology, flap rates and flap mechanics), fine frequency discrimination within the acoustic fovea may facilitate both the detection and identification of an insect [Schuller 72, Johnson *et al.* 74]. There are a variety of observations suggesting that high duty-cycle bats do engage in an active prey selection [Goldmann & Henson 77, Trappe & Schnitzler 82, Schnitzler & Ostwald 83, von der Emde 88, Nitsche 87, von der Emde & Menne 89, von der Emde & Schnitzler 90, Kober & Schnitzler 90, Jones 90]. Most of these studies were performed on rhinolophids and *P. parnellii.* In bats with lower duration calls, such as some hipposiderids, flutter discrimination is not as well developed and there are fewer reports of active prey selection among these species [Bell & Fenton 84]. (Note, reports on prey selectivity for Microchiroptera in the field are scarce. Although such behaviour has been observed, conclusions are contradictory and more studies are needed to understand under what circumstances these bats exhibit prey selection.)

The cues which a high duty-cycle bat may use to differentiate between species are far from clear, nor have available cues from fluttering targets been sufficiently characterised. From aural observations (*i.e.* listening to slowed-down versions of insect echoes) and visual inspection of echo oscillograms and spectrograms reflected from insonified fluttering insects, it is readily apparent to human classifiers that the "structure of the glints as well as the [portion of the] echo between glints is rarely identical among insect species" [von der Emde & Schnitzler 90]. However, if the bat is to identify an insect of a particular species, it has to overcome the problem that echoes from the same insect "change dramatically when the insect turns in space ... a species specific parameter, present in all oscillograms and spectrograms of an insect is not obvious to a human observer" [von der Emde & Schnitzler 90].

Although neurophysiological investigations with natural echoes have not revealed the existence of *grand-mother* cells (*i.e.* units specialised to discharge for particular insect species), filter neurons — even at the periphery of the auditory system — discharge preferentially to a specific range of information-bearing parameters (*e.g.* echo intensity, modulation rate, extent, percent *etc.*) [Ostwald *et al.* 88]. Like neurons elsewhere in the mammalian auditory system, filter neurons respond to periodically modulated signals with discharges that are tightly locked to the modulating waveform. (In most studies, sinusoidally amplitude or frequency modulated (SAM or SFM, respectively) signals are used as stimuli.) Filter neurons typically respond either with a synchronised "tonic" discharge (*i.e.* a sustained discharge which roughly follows the sinusoidal stimuli) or a synchronised "phasic"/"transient" discharge (*i.e.* firing of a few spikes to a distinct part of the modulating waveform) [Suga & Jen 77, Schuller 79, Pollak & Schuller 81, Bodenhamer & Pollak 83, Ostwald *et al.* 88, Lesser *et al.* 90] when the frequency of an echo sweeps into (or simply oscillates within) the narrow confines of their tuning curves. The characteristic which particularly distinguishes filter neurons in the acoustic fovea of a high duty-cycle bat's auditory system is that they are far more sensitive to periodic modulations (particularly at low echo intensities) than are more broadly tuned units [Pollak & Schuller 81, Vater 82, Bodenhamer & Pollak 83]. Single unit studies in the cochlea, inferior colliculus and cortex of *R. ferrumequinum* and *P. parnellii* demonstrate that the filter neurons can follow SFM down to modulation depths below 20 Hz [Schuller 79, Pollak & Schuller 81, Bodenhammer & Pollak 83, Pollak *et al.* 86] and SAM down to modulation indices below 25% [Schuller 79, Vater 82]. Of this phenomenon, Pollak writes [Pollak *et al.* 86]:

> *"The bat's perception of such phase locked discharge trains must be dramatic. A visual analogy would be facilitating the detection and identification of a firefly by actively eliminating as much light as possible from the environment while maintaining images in the ... field of the acoustic fovea. The luminescent pattern of a firefly would stand in sharp relief against a jet black background, rendering the insect readily detectable as well as identifiable from its flashing pattern."*

Sharp tuning and these dramatic changes in spike count response will also allow filter neurons to encode subtle changes in frequency which may prove to be equally important in target discrimination of this sort. The spectrum of a modulated echo will be neurally encoded in the sites of activity along a neural substrate containing subpopulations of units varying continuously in their tuning and modulation preferences.

In the next chapter, I demonstrate that an orientation invariant, target-specific cue exists in the positions of the spectral sidebands reflected from a target with periodic motion (*i.e.* model insects and a variety of rotating fans). Given the presence of such a cue, that chapter goes on to demonstrate how spectral sideband energy can be used directly by a localisation mechanism which derives a steering signal from time and intensity cues available in the low frequency modulated echo envelopes. Moreover, *target specific* localisation can be achieved if this localisation mechanism is made to operate on energy in only those frequency bands which contain the spectral signature reflected by a desired target. In this way, the selection of targets with periodic motion need not be mediated by formal recognition (or matching to an internal model). Rather, in the case of an echolocator highly coupled to its environment and task, it might arise naturally out of the physics of the sound generation, reflection and reception processes.

4. Investigation: Target Selection

Generally two separate mechanisms are thought to underlie target selective behaviour in animals: (i) a mechanism for recognising desired targets on the basis of which (ii) a mechanism for localising targets is invoked. Separation of the two can be difficult, as many psychophysical experiments ask an animal to indicate its preference for a target by moving towards it. This chapter explores the possibility that only one mechanism is required to achieve selection of targets with periodic motion *via* long CF echolocation signals.

4.1 Hypothesis

Motivation

To appreciate this idea, I recall the recommendation, put forward in the introductory part of this thesis, that we should seek to explain (and design) sensory mechanisms by characterising the physical interaction of an agent with its environment throughout the execution of its particular task. In the case of target acquisition by long CF echolocators, this characterisation involves consideration of the signal transduction and filtering processes underlying echo reception and generation, and an exploitation of the particular constraints imposed by the environment (*e.g.* clutter and flutter) and task (clutter rejection, flutter localisation). Viewed within this framework, the prey selective behaviour of high duty-cycle bats might be accomplished by a localisation mechanism which is tuned to specific patterns of energy in an echo, and, therefore, can most successfully localise those targets which reflect that pattern. This is reasonable to suppose because high duty-cycle bats prey *only* on targets with periodic motion and, therefore, do not require an independent ability to localise other echo sources in this way. Similarly, although echolocation may be used in *avoidance* behaviours, *approach* (or *target selective*) behaviours do not require the identification of echoes that need not be located. Such a localisation mechanism provides a means for inherently rejecting echoes from stationary targets and targets with non-periodic motion, and can facilitate selection of fluttering targets on the basis of the extent to which they reflect the pattern of echo energy which the echolocator can localise.

Mechanisms sub-serving the localisation of sounds by their modulation waveforms are not well understood (in this respect, human psychophysical studies provide some of the richest descriptions [Kay 82]). In order to consider the directional cues available in a modulating waveform, it is necessary to understand the more general mechanisms underlying sound source localisation. (Here I will consider lateralisation, or sound source localisation in 2D; 3D localisation is the topic of Part III of this thesis.)

Localisation of tones

Animals are able to localise sound sources because of the way in which sound is propagated in the environment. A point source pulsating at a fixed frequency cyclically compresses molecules of surrounding medium, which, in turn, results in sinusoidal changes in pressure. Thus, an acoustical wave is a compression (or *longitudinal*) wave which propagates in the surrounding medium at a rate which is determined by that medium ($c \approx 344$ m/s in air). The location of a sound source can be perceived by measuring properties of the passing wave. For example, near the source, particle movements contain sufficient energy to be detected [Ewing 89]. In the far field, timing and intensity characteristics of the sound may be used.

Sound converging on a receiver from different positions in the far field will be characteristically delayed and attenuated as a result of its unique travel paths. The amplitude of the wave decreases as the inverse of the squared distance due to spherical spreading (frequency dependent absorption is also significant in many media) and the phase of the wave varies with listener position (cyclically rather than monotonically). In the case of sound amplitude, a spreading wave can be further modified by the *shadowing* effects of nearby objects or parts of the listeners' anatomy itself (*e.g.* head or external ears). Moreover, the shape of pressure sensitive receivers causes otherwise identical waves arriving from different directions to affect them differently — *e.g.* a finite extent diaphragm, unlike a point receiver, will respond most strongly to sound arriving normal to its surface while sound waves arriving from oblique angles interfere across the surface, in ways determined by the ratio of sound wavelength to diaphragm radius.

Just as sound arriving from different source positions is characteristically delayed and attenuated, sound *originating* from a single location — but *received* at different positions in the field — can be characterised by particular delay and intensity differences. By employing two laterally separated receivers, animals lateralise sound based on inter-aural disparities of intensity (IID), arrival time (ITD) and/or phase. The relative magnitudes and, therefore, the relative importance of temporal *vs.* amplitude cues varies with sound frequency, and the acoustic properties and dimensions of the head and pinnae of an animal.

Although an ITD can provide a relatively large head with precise angle-to-target measurements, it is an inherently less flexible mechanism than an IID comparison because the calibration constant — *i.e.* the speed of sound — is fixed by the environment. By contrast, intensity is an adjustable calibration constant: it can be manipulated by characteristics of the receiver such as pinna size and orientation aboard the head. Using IIDs, even small binaural receiving systems can enjoy the stereo effects of large heads.

It is widely assumed that this is the strategy of echolocating bats. The heads of bats are small and maximum ITDs (*i.e.* for targets located at $90°$ off-axis) are less than 50μs. Likewise, phase differences are unlikely to afford bats with sufficient directional cues, since the half wavelength of ultrasonic sounds is shorter than inter-aural distances. Although the heads of bats are small, they can generate large IIDs due to the effects of the well developed pinnae and head shadowing: *i.e.* measured IIDs are typically $25 - 30$ dB across the frontal sound field (see review in [Pollak & Park 96]). Moreover, binaural neurons in the bat's auditory system are most sensitive to IIDs [Schlegel 77, Fuzessery & Pollak 85, Pollak 88]. Reported neuronal time-intensity trading ratios (*i.e.* the ratios of those inter-aural time and intensity differences which result in the same neuronal effect) are, on average, 47 μs per dB [Harnischfeger *et al.* 85, Pollak 88]. Thus, in the lateralisation of targets, intensity effects can easily dominate.

Localisation of modulated tones

The classical theory whereby localisation using inter-aural time (or phase) differences is restricted to the low frequency range and localisation of high frequencies requires inter-aural intensity differences holds specifically for pure tones [Mills 60]. If high frequency tones are modulated by lower frequency waveforms (introduced, say, during the reflection or reception process), sources can be more easily localised on the basis of the inter-aural timing/phase differences between the modulating waveforms. For example, in headphone experiments with humans, changes in inter-aural time differences between the common amplitude envelopes of two modulated tones can be detected as movement of an apparently fused sound source, even when the carrier frequencies delivered to the ears differ [Henning 74]. Minimum detectable time/phase differences are comparable to those corresponding to the lateralisation of pure tones of the same low frequency as the modulating tones. It is important to note that, as the phase of an envelope is inextricably linked with the instantaneous mean amplitude of a sound, it is not clear whether inter-aural amplitude differences, rather than phase, were the salient cues.

Mechanism

A similar ability has been hypothesised for bats. Specifically, it has been suggested that certain ON/OFF units in the peripheral nervous system [Lesser *et al.* 90, Covey *et al.* 91, Grothe 94] — which respond phasically to rapid transients in echo envelopes (*e.g.* glints) — may mark the signals entering the right and left ears such that binaural comparison of the phase of the (relatively low frequency) echo envelopes may be performed. It is intriguing to wonder whether the peculiar medial superior olive (MSO) might have evolved in microchiroptera to play a role in such a marking mechanism. If the MSO is connected to pre-motor areas (as has been speculated for the lateral superior olive (LSO) [Covey *et al.* 91]), it is possible that it could provide steering signals to control yaw movements toward a fluttering target. (Section D.2

reviews this possibility further.)

In this chapter, I investigate the directional cues contained in modulated echo envelopes *via* analysis and computer simulation (Section 4.2) and in a robotic "prey" capture experiment (Section 4.3). The former section considers, for a simple fluttering reflector, how periodic reflector motion re-distributes the spectral energy in a long CF call into particular spectral sidebands — the position of which is determined by the characteristic motion of the target. A computer simulation is then used to demonstrate that (for targets whose reflector surfaces are continuously visible) this cue is invariant to target aspect angle and, moreover, that modulated echo envelopes available in corresponding sidebands arriving at the right and left ears provide timing cues which can be used to steer the echolocator toward the fluttering target.

In Section 4.3, a target selective localisation mechanism (one that localises targets with particular spectral signatures) is tested in the real world using rotating fan targets. Here the bionic sensor was mounted aboard a robot and the inter-aural differences in echo envelopes used to determine motor commands to drive the robot toward a designated Target Fan (*i.e.* target reflecting a designated spectrum). These experiments were performed in a cluttered laboratory environment containing both stationary reflectors and reflectors with other types of periodic and non-periodic motions.

4.2 Fluttering targets as directional beacons

In order to motivate this discussion, I introduce the notion that a target's motion creates a *signature signal*. On their own, the physical events out of which these signature signals arise do not necessarily generate sufficient acoustic energy to be detected by a distant listener. Moreover, signature signals may be relatively low frequency signals and, therefore, may not have the desired directional propagation properties to facilitate localisation. However, signature signals can be communicated over a distance when they modulate the amplitude and frequency characteristics of a higher frequency (and intensity) echolocation call signal. The effectiveness of this communication depends upon the echolocator's ability to extract the essential features of the modulation from a received echo(es). (In this regard, the more an echolocator exploits the modulations which come about from the essential locomotive motions of a target, the more effective it can be at this game.)

Numerous factors contribute to a target's signature signal — including target shape, size, and number of reflecting surface(s); the rate and mechanics of motion; and the echolocator's viewing perspective. Here I consider targets with simple, well defined motions and morphologies in order to examine how the spectral composition of an insonifying carrier tone is altered upon reflection. Two broad assumptions are employed in the target characterisations used here: (i) oscillations are truly periodic and (ii) reflecting surfaces are continuously visible — *i.e.* reflectors are composed of one or more individual point-like scattering surfaces which each generate echoes that are visible throughout the hemisphere normal to the surface. Upon reflection, such a target will simultaneously amplitude and frequency modulate an echolocation signal. The resulting echo can be described as the real part of a complex signal with time varying amplitude and phase components [Vakman 72, Stremler 90, Panter 65, Loughlin & Tacer 95]. Considering a target with several moving reflective points, we have:

$$\phi_{AMFM}(t) = Real\ \{[A_c(1 + \sum_{m=1}^{M} f_{AM,m}(t))]\ exp[j(\int_0^t \omega_c(1 + \sum_{m=1}^{M} f_{FM,m}(t))dt)]\} \quad (4.1)$$

where A_c is the carrier signal amplitude, ω_c is the carrier frequency, and f_{AM} and f_{FM} are the amplitude and frequency modulation waveforms and M is the number of reflecting surfaces.

For simplicity, now consider a "stick" insect consisting of one diffuse reflecting point at the tip of its wing which is fluttered so as to sinusoidally amplitude and frequency modulate an echolocation tone. (Simplifications of the complex motions and morphologies of real insects which are implied by the characterisations used in this chapter are discussed in Section 9.1.9.) Strictly speaking, the FM waveform can be sinusoidal but the AM waveform should be rectified, as a negative amplitude glint is physically impossible. This can be expressed as follows:

$$\phi_{AMFM}(t) = Real\ \{[A_c(1 + |sin\ \omega_{AM} t|)]\ exp\ [j(\omega_c t + \beta\ sin\ \omega_{FM})t]\} \quad (4.2)$$

where β is the *dispersion index* — which, for periodic modulations, is defined as the depth of modulation divided by the rate $\frac{\Delta\omega}{\omega_{FM}}$).

The AM term in Equation 4.2 can be expanded in a series straight-forwardly. Likewise, the FM term[1] can be expanded in a Fourier series with coefficients given by the nth order Bessel functions of the first kind, J_n, which can be evaluated numerically in terms of the parameters n and β. Doing so yields:

$$\phi_{AMFM}(t) = Real\ \{[A_c(1 + \sum_{l=-\infty}^{+\infty} \frac{exp\,[j(2l\omega_{AM}t)]}{\pi(1-4l^2)})] \sum_{n=-\infty}^{+\infty} J_n(\beta)\, exp\,[j(\omega_c + n\,\omega_{FM}t)]\} \tag{4.3}$$

From this expression, it is apparent that amplitude modulation, a linear modulation process, effectively multiplies the time domain functions describing the modulating and call signal or, equivalently, convolves their spectra such that the amplitude modulating component of the signature signal (producing sidebands at $2l\,\omega_{AM}$, where $l \in \mathcal{Z}$ and $\mathcal{Z}$ is the set of integers) is symmetrically translated (from $\omega = 0$) through a distance ω_c. Frequency modulation has the effect of distributing carrier energy and each of the AM sidebands to spectral positions $\omega_c + n\,\omega_{FM}$, where $n \in \mathcal{Z}$. The $1/\pi\,(1-4l^2)$ and $J_n(\beta)$ terms control the contribution of the AM and FM sidebands. In the case of the latter, $J_n(\beta)$ reaches its peak at approximately $\beta = n + 1$.

This characterisation allows us to see that a particular modulation waveform produces sidebands at well defined positions and a spectral profile (*i.e.* distribution of energy among the sidebands) which is stretched or compressed for different β. However, the characterisation thus far has blurred the distinction between the waveform characterising the *actual* movement patterns of the wing — which is determined by

[1] The frequency modulation term given in Equation 4.2 is borne out of the following simple expression:

$$\phi_{FM}(t) = Real\ \{A\, exp\,(j\psi(t))\}$$

where $\psi(t)$ is the phase angle (note, I have ignored an initial phase offset). Thus, $\psi(t)$ is related to the instantaneous frequency as follows:

$$\psi(t) = \int_0^t (\omega_c + \Delta\omega\, cos\,\omega_{FM})dt = \omega_c t + \beta\, sin\,\omega_{FM} t$$

The function $cos\ sin(\omega t)$ occurs in many physical problems and has been extensively tabulated — see, for example, [Stremler 90].

species specific morphological and flight specialisations — and the perceived, or *relative*, movement profiles — which are determined by both target-specific factors and viewing orientation. Even with the assumption of a fixed reflecting point(s) visible from all orientations, the waveforms describing the reflected frequency and amplitude variations will vary somewhat with target aspect angle. In the next subsection, I look at how these waveforms are affected by the relative orientation of echolocator and target.

Signature signals illustrated

This subsection aims to provide an intuitive explanation of the reflection process considered thus far. The echoes presented here were generated in a 3D Echolocation Simulator environment containing the simple generic insect already described and a simple model bat (see Figure 4.1). (Simulated bat and insect are described in full in Appendix C.) The wing moves within a vertical stroke plane through an angular excursion of 180° at a rate of 100 Hz — see Figure 4.2. (No twisting or forward-back motion is considered.) At each time step in the simulation, a new wing angle is calculated as $f(\omega_{FM}t) \times 90°$, where f is a modulating function which varies periodically between ± 1 (in this example, $f = sin$). Thus, the velocity varies sinusoidally between zero at the top and bottom of the stroke ($\pm 90°$) and its maximum value — determined by the wing length and flap rate — through the middle of the stroke (0°).

The reflecting point at the tip of the insect's wing is diffuse (*i.e.* audible from any position in its frontal hemisphere); however, the glint process is modelled by a scatter function wherein an angle θ_n between the transducer axis and the normal to the wing is used to scale the magnitude of the reflection (as $|cos(\theta_n)|$) so as to vary the strength of the reflection between a maximum and a minimum (DC) reflection value. (In this regard, one might argue that the reflector becomes somewhat specular, but this is a moot point.) In Figure 4.2, the wing positions producing maximum and minimum amplitude glints (for a particular observer location/orientation) are numbered consecutively according to the order in which the wing moves through them during a single wing-beat cycle.

Figures 4.4-4.8 illustrate how the amplitude and frequency modulating waveforms re-

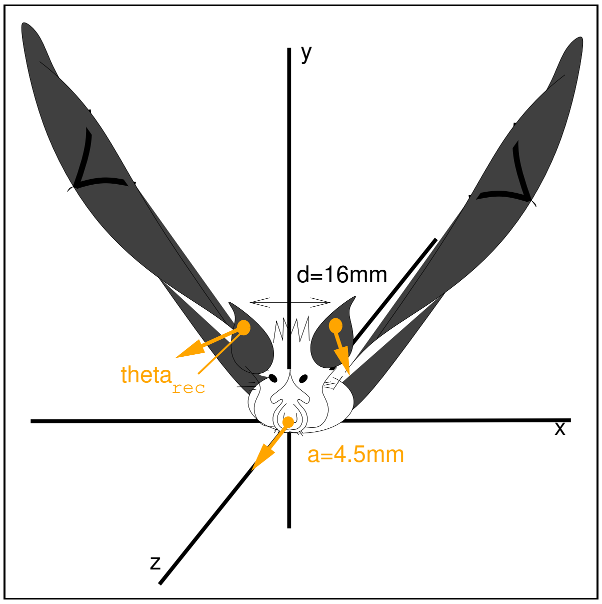

Figure 4.1:
Simulated bat head consisting of an emitting and two receiving piston-like transducers. Inter-aural distances, transducer size and acoustic axes are shown. (The vertical separation between receivers and emitter is set to zero in all simulations.)

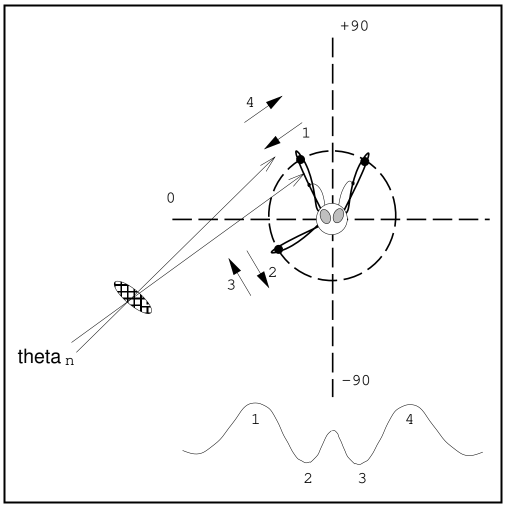

Figure 4.2:
Motion of the model insect. Numbered are the positions producing the maximum (1,4) and minimum (2,3) amplitude reflections for an observer located at the position and orientation indicated by the iconised transducer. Below, the relative amplitudes at each of these points are sketched.

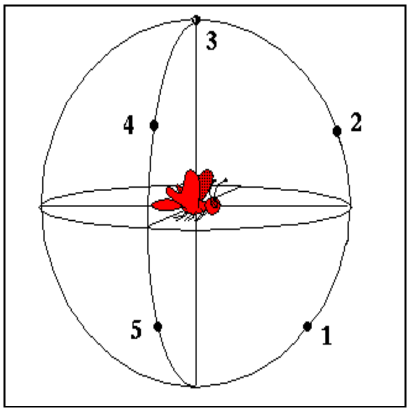

Figure 4.3:
Target visibility angles. (Numbered angles are used in interpreting subsequent figures.) Separation between target and viewer is 0.5 m. Viewing angles along front- and side-facing arcs are: (1) $-60°$, (2) $30°$, (3) $90°$, (4) $30°$, and (5) $-60°$. Emitter is orientated directly at the target. Echoes arriving at the right simulated ear (which is slightly offset from the emitter, but oriented in the same direction) are shown in Figures 4.4-4.8. (The insect model and other aspects of the *3D Echolocation Simulator* are described in more detail in Appendix C.)

flected by this model insect vary with viewing angle. There, the oscillograms (top), filtered time waveforms (middle) and summed spectral energy (bottom) are plotted for a 50 ms echo received in the right simulated ear of a bat echolocating directly at the target from each of the five viewing perspectives of Figure 4.3. Echoes are filtered in overlapping frequency channels (centre frequency spacings of 50 Hz) which are realistically tuned (average filter widths are 200 Hz, $Q_{10dB} = 415$). The signals shown in each channel of the filter-bank have been rectified and lowpass filtered ($f_c = 1000$ Hz).

Considering first the three echoes received from observation points along the front-facing arc of Figure 4.3 (positions $1 - 3$ in Figure 4.3, echoes in Figures 4.4-4.6), the waveforms are fairly simple when seen from these viewing perspectives which are perpendicular to the flap plane. Within each of the 5 wing-beat cycles captured by the 50 ms pulse, the amplitude peaks twice — once on the up-stroke and again on the down-stroke (see oscillograms). The amplitude and frequency modulating waveforms maintain their shape but grow in magnitude as the viewer moves toward the top of the visibility sphere along this arc.

As seen from this front-facing arc, the insect wings are moving in a plane which is

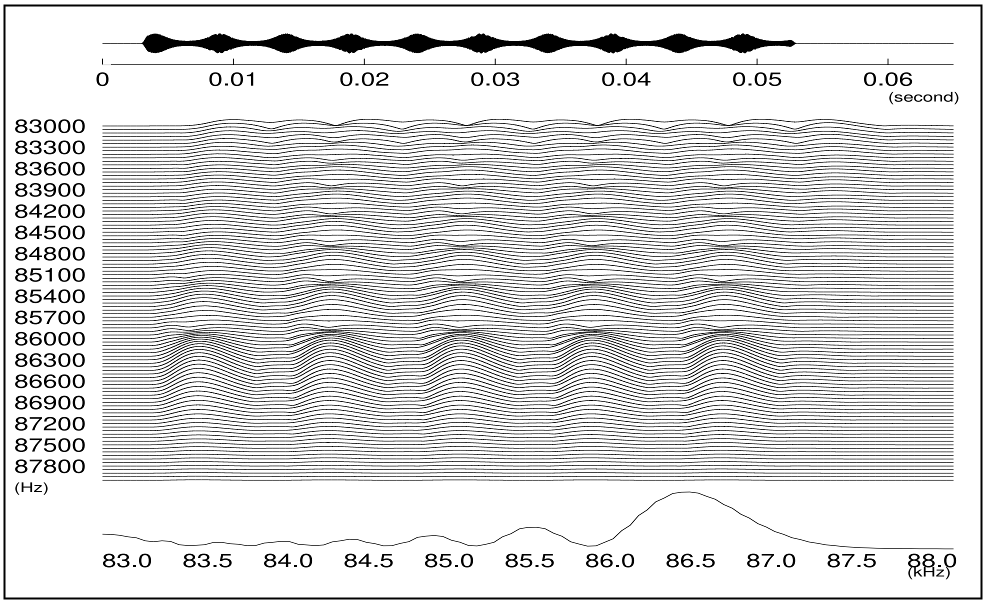

Figure 4.4:
Oscillogram, filter-bank and spectral representations of target as seen from Position 1 in Figure 4.3. In Figures 4.4-4.8, simulated echoes are show as oscillograms (top), the response across a filter-bank (middle) and as summed spectral energy (bottom). The filter-bank performs the following operations: bandpass filtering around the carrier, rectification, followed by lowpass filtering. Individual bandpass (7th order Butterworth) filters have an average filter width of 200 Hz ($Q_{10dB} = 415$) and are separated by 50 Hz. Lowpass filters are 2nd order Butterworth filters with $f_c = 1000$ Hz. Spectral information is normalised to highlight the detail. Only half of the symmetric spectrum is shown in the interest of space. The target is a single reflecting point at the tip of one 9 mm long wing, flapping with a sinusoidally varying velocity through an 180° vertical flap plane, at a wing-beat rate of 100 Hz.

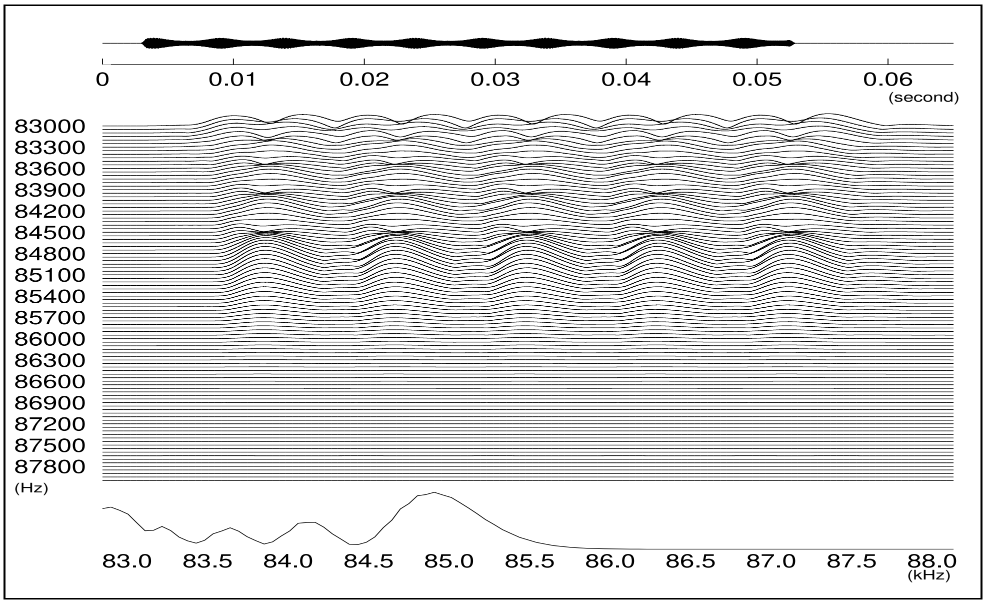

Figure 4.5:
Oscillogram, filter-bank and spectral representations of target seen from Position 2 in Figure 4.3.

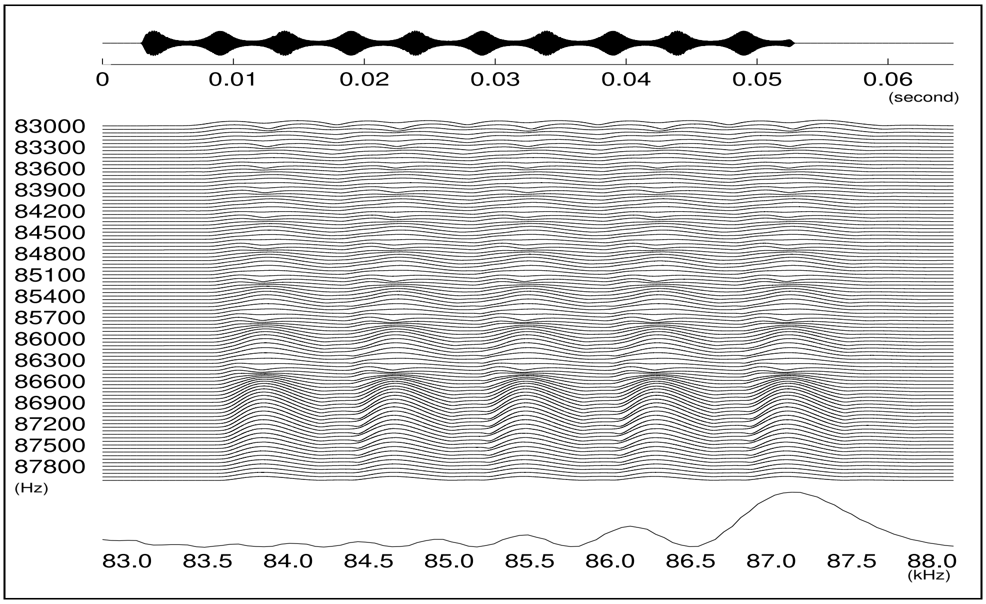

Figure 4.6:
Oscillogram, filter-bank and spectral representations of target seen from Position 3 in Figure 4.3.

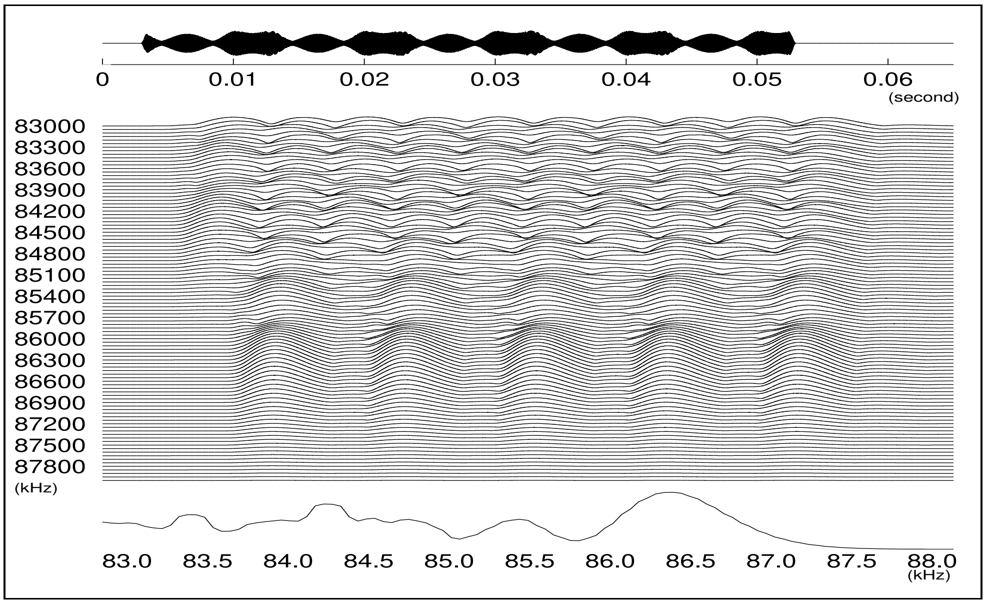

Figure 4.7:
Oscillogram, filter-bank and spectral representations of target seen from Position 4 in Figure 4.3.

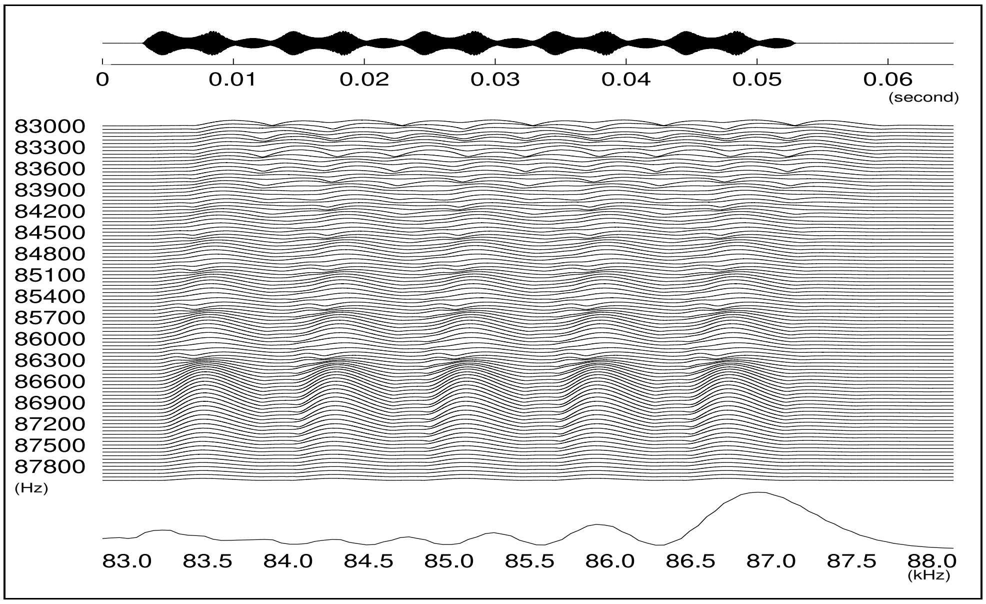

Figure 4.8:
Oscillogram, filter-bank and spectral representations of target seen from Position 5 in Figure 4.3.

largely perpendicular to the viewer direction such that changes in Doppler-shift are brought about primarily through changes in wing velocity throughout the wing-beat cycle (rather than through changes in the (*relative*) velocity projected onto the line-of-sight between target and observer). At the top of the arc, the frequency modulation is greatest because the viewer has moved into the flap plane — consequentially, β reaches its peak and spectral energy is pushed furthest from the carrier. Also, at this viewing angle the two velocity components are in phase: the wings are moving with their greatest velocity *relative* to the observer at the same time that they are moving with their greatest *absolute* velocity.

The frequency modulating waveforms seen from within the flap plane (positions $4-5$ in Figure 4.3) become distorted as the viewer moves toward the horizontal plane where the two Doppler contributing factors (*i.e.* projected *vs.* absolute wing velocity) are most out of synchronisation. Nevertheless, the position of the sidebands so not change. One can see from Equation 4.1 that this will be true as long as the FM *rate* does not change.

Within the flap plane, aspect angle determines when the maximum and minimum amplitude reflections occur within the wing-beat cycle — thereby giving rise to additional amplitude glints. (Figure 4.2 depicts the situation as seen from *Position 4*.) These distortions are visible in the time waveforms atop Figures 4.7-4.8. Unlike FM, distortions in the AM waveform effect sideband position as well (see Equation 4.1). However, the filter-bank response (and Equation 4.1) shows that the AM contribution to the spectrum dies out quickly, while the target specific harmonic structure (*i.e.* the relative position of the sidebands) is preserved at higher frequencies.

To an echolocator with a complex signal processing system capable of registering echo intensity, modulation depth, percent *etc.*, target features other than modulation rate are, no doubt, salient and useful cues. In this work, only the invariant relationship between sideband harmonics is used to identify reflectors on the basis of their motion. Clearly, this limits the number of targets which can be distinguished; however, discrimination performance is not simply related to the visible or reported rate (*e.g.* the species specific wing-beat rate) of the target. The reflected rate of an oscillating/rotating/fluttering target can be a complicated function of target morphology and motion. For example, the rate of a fan with similar blades is the fan's reported rotation rate multiplied by the number of blades. Reflectors with morphological symmetries may give rise to situations wherein a visible rate is effectively doubled in the acoustic domain during reflection. In the case of insects, it is very likely that factors such as wing stiffness contribute to the effective, reflected rate such that different species with putatively similar wing-beat rates reflect echoes containing quite different sideband positions.

Although not investigated here, an echolocator clearly might detect the orientation of a target by monitoring the distribution of energies across the spectrum. As shown in the spectral plots at the bottom of Figures 4.4-4.8, the spectral profile is stretched and compressed as the viewer moves into different Doppler receiving angles.

Using signatures to localise their source

The echo envelopes oscillating at low (*i.e.* effective the wing-beat) frequencies in the filter-bank channels of Figure 4.4-4.8 contain time/phase and intensity cues which may be used to localise targets *via* binaural comparison. In the case of temporal cues, a correlation-like process can be used to determine the temporal offset between waveforms in corresponding frequency channels of the right and left filter-banks. However, as in the case of comparisons of tones, echolocators with small heads will run into a resolution problem because a simple temporal correlation comparison yields angular cues whose resolution is inversely proportional to inter-aural separation. In performing correlation, echolocators with large heads have a different problem to contend with; namely, as receivers are separated, a correspondence problem arises. The correspondence problem is a fundamental physical constraint which may only be overcome by reorientating the receivers (which, one might argue, requires some degree of knowledge about how to solve the correspondence problem). By contrast, echolocators with small inter-aural dimensions may overcome their resolution problem at the signal processing level by using a filtering mechanism which trades intensity for time. As discussed in the previous section, in bats, neuronal latencies (*i.e.* time-intensity trading ratios) work such that a 1 dB decrease in echo intensity yields firing delays which approach the maximum possible ITDs.

In the present implementation, a time-intensity trade-off is achieved by marking the phase of echo envelopes according to when their instantaneous amplitude cycles through a fixed filter threshold. Figure 4.9 (a) and (b) show examples of envelopes output from one of the filter-bank channels for a target located at an azimuth bearing of $+12°$ within the $0°$ elevation plane. This figure illustrates a situation wherein the echo received at the right simulated ear (a) is quieter than that heard at the left (b) such that some portions of the cycle fail to rise above threshold and become marked. Figure 4.9 (c) shows this as unit markers with the right (orange) encoding overlaying left (black) encoding.

In deriving an angle from the pulse encoded signals of Figure 4.9 (c), a correlation process must compute a temporal offset based upon the information available in both the *missing* and *delayed* pulses. Figure 4.10 provides an example of one way in which this might be achieved. In Figure 4.10 (a) a steering signal is built up over time as $\sum_{t=0} pulse_{right}(t) - pulse_{left}(t)$. In this figure, it can be seen how, for echoes arriving from increasing off-axis azimuth angles, the thresholding process causes increasing numbers of pulses to drop out of one envelope encoding with respect to the other, and an uncorrected steering signal accumulates for the duration of the echo. Figure 4.10 (b) shows this same information for pinnae orientated by $\pm 15°$ away from the mid-line (as shown in Figure 4.1). (In this case, the receivers are most sensitive to targets at a $\pm 7.5°$ azimuth offset, rather than straight ahead.) This effectively increases resolution around the mid-line — as targets located around $0°$ azimuth yield echoes with larger inter-aural intensity disparities (causing larger discrepancies in the number of pulses dropping out at the thresholding stage). Figure 4.10 (c) shows the slopes of the steering signals (least squares fit to the steering signal time waveforms) measured across $15°$ of the frontal sound field for both cases of pinnae orientation. This figure also shows how increasing the threshold (dashed line) can increase resolution around the mid-line; however, forcing pulses to drop out in this way yields a signal which is useful over a restricted portion of the sound field only. Better results can be achieved by driving the signal levels apart acoustically (*i.e. via* pinna reorientation).

For each azimuth angle depicted in Figure 4.10, the steering signals are an average of those computed in a number of channels across the filter-bank: channels ranging from

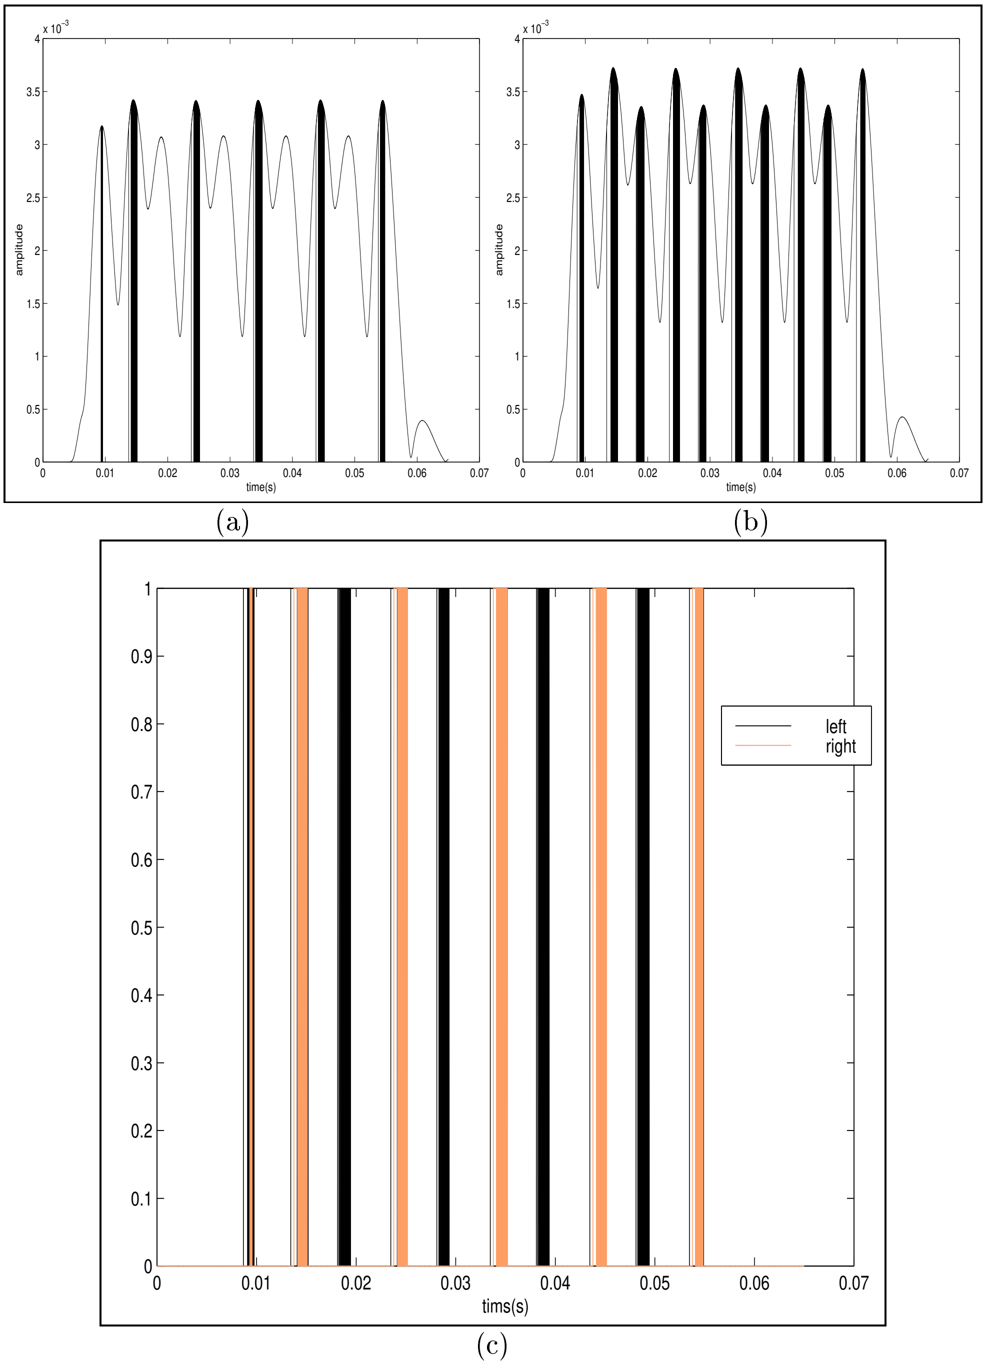

Figure 4.9:
Pulse encoding of echoes. Echoes received at right (a) and left (b) receivers after filtering. The filtered envelopes are marked by vertical lines which, for display, have been scaled to match the amplitudes of the filtered signals. (c) A unit marker is plotted for each upward portion of the envelope cycle during which time the signal is above threshold (set, in this example, to 3.25). Sub-threshold readings are zeroed.

$83.0 - 85.5$ kHz and separated by 100 Hz. (The steering signal computed through binaural comparison of signals varies somewhat between filter-bank channels due to the fact that the Doppler effect varies the amount of energy projected into each frequency channel.) As the number of filter-bank channels contributing to the average increases, the steering signals become more smooth. Deriving a more reliable signal through an integration process such as this works when a target is isolated so that all energy entering the filter-bank corresponds to the target. When other reflectors are present, the echo energy across the channels of the filter-bank is not inherently weighted according to target's reflectivity spectrum, and, therefore, a weighting must be imposed in order for a reliable steering signal to be derived. In the next section, I show how an echolocator which knows the characteristic spectrum (*i.e.* sideband positions) reflected by a favoured target can achieve target specific localisation by only employing energy in particular filter-bank channels during the binaural comparison process.

4.3 Target selective localisation

In this section, I test (aboard a robot) the hypothesis that an echolocator's exploitation of a target's characteristic signature signal takes the form of a coupling between agent and environment in which the localisation mechanism works on signature energy. In this way, targets are localised according to the amount of the desired signal present in an echo such that stationary reflectors and reflectors with different motion signatures are effectively ignored.

Using the scheme discussed above for trading time for intensity so as to facilitate an ITD comparison proved to be difficult to implement on the larger robotic head as differences in projection of target velocity onto the two receivers altered the time waveforms arriving in each frequency channel of right and left filter-banks in ways that did not yield a predictably varying phase offsets. However, since the time-based strategy utilised above essentially exploited differences in echo envelope amplitude, binaural comparison of modulated echo energy is used aboard the robot to serve a similar purpose in localisation.

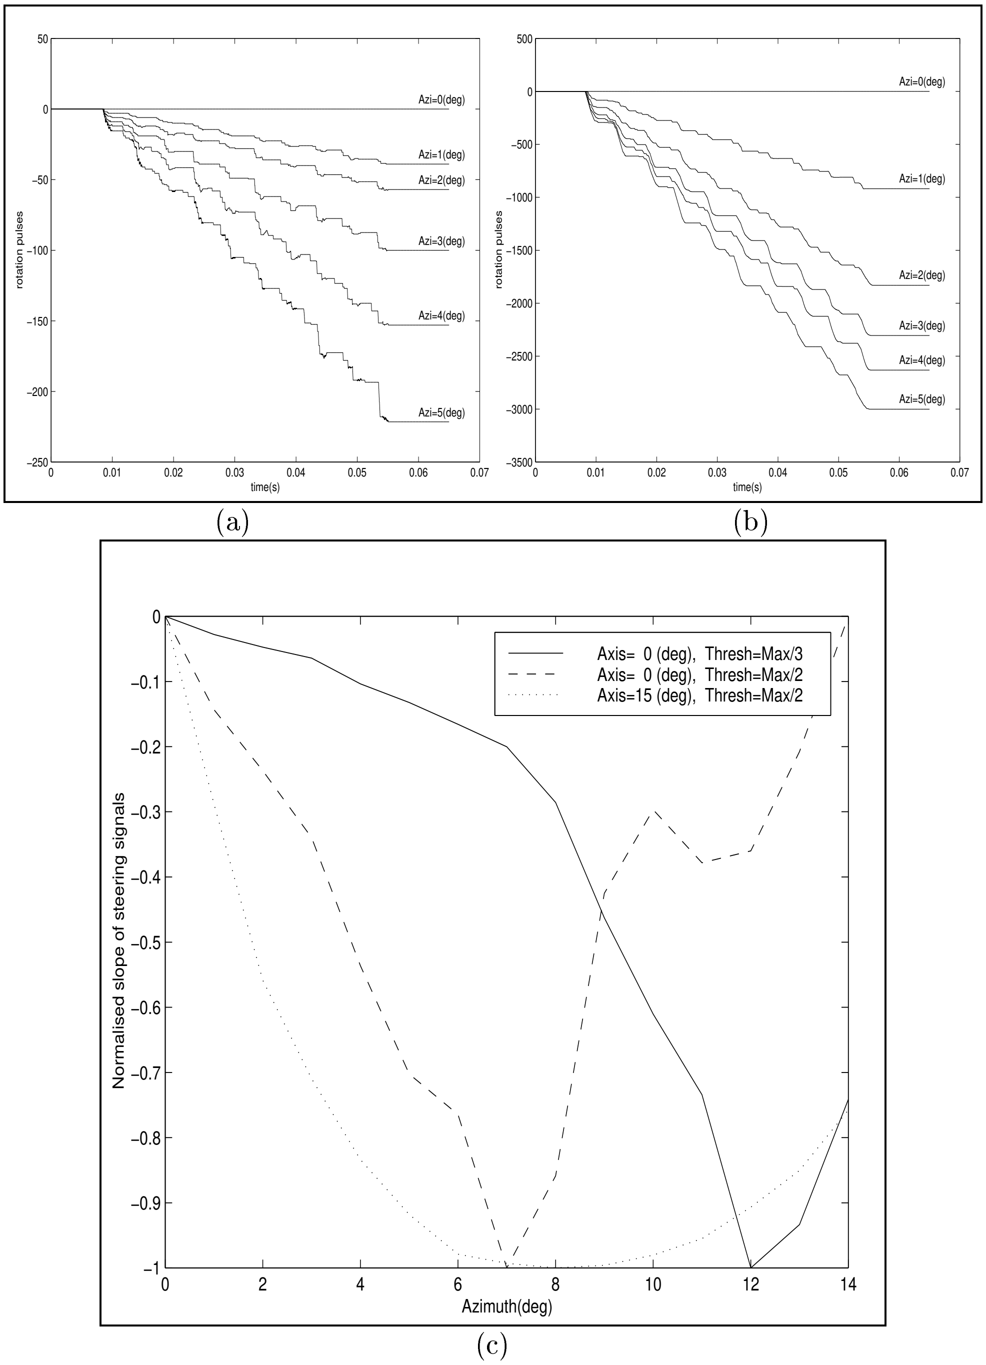

Figure 4.10:
Steering signals. Steering signals computed across time for (a) receivers pointing forward and (b) rotated outward by $\pm 15°$, as in Figure 4.1 (a). (c) Angular resolution as depicted by the slope of the steering signal (least squares fit to steering signals) for different pinnae orientations and encoding thresholds.

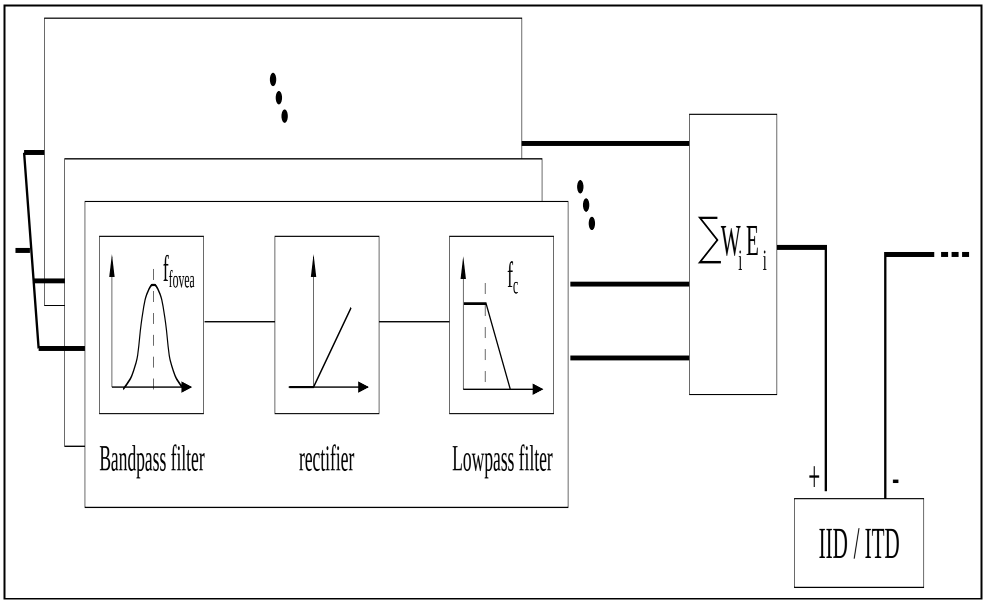

Figure 4.11:
Hypothesised localisation mechanism. The filter-bank depicts an idealised cochlear filtering operations to which all sound entering the mammalian ear is subjected. Outputs of the filter-bank are used in a target selective binaural localisation mechanism. As shown, target selectivity is achieved by weighting energy in the channels emerging from the filter-bank so that inter-aural comparisons are preferentially performed on energy in those spectral sidebands containing the target's spectral signature.

4.3.1 Experiment

Figure 4.11 shows the proposed signal processing scheme which takes echo envelopes output from a filter-bank and performs a binaural comparison. Between filtering and inter-aural comparisons, echoes are passed through a target selective function. Here I depict that as a function which weights spectral energy according to the distribution of expected sideband energies which characterise a desired fluttering target.

In the robotic experiments described below (performed with Herbert Peremans), the bionic sensor was mounted aboard a mobile robot (*i.e.* the in-house built, transputer-controlled mobile robot shown in Figure 4.12 and described in [Forster 91]) and the localisation mechanism was used to convert echo information into appropriate rotation and translation commands which were sent to the motors during the subsequent sensing cycle. As the robot navigated from a given starting position and orientation, a motor log was kept to recreate the path traversed. If the robot failed to reach the target in a reasonably time (approximately 20 minutes or less), the failed run was logged, and the robot rebooted and returned to its starting point.

All runs were performed in a cluttered, busy laboratory consisting of a number of stationary echo generating surfaces (*e.g.* walls, filing cabinets, *etc.*) as well as moving surfaces (students and other robots). The localisation mechanism was tested under the following conditions.

- **No target present.**
- **Target Fan present.** Here the Target Fan was a small Pentium heat sink fan consisting of seven 10 x 6 mm blades, twisted through an angle of approximately 30°. The effective rotation rate of this fan is 540 Hz. (This and the other fan targets reflected echoes which adhere to the assumptions underlying the target model disucssed in Section 4.2 — namely, they have a constant rate of rotation and their moving surfaces are continusouly visible to an insonifying probing beam. In the case of these rotating targets, it is not a single blade which is visible, but a series of identical blades — each of which comes into view, reflects sound over a limited range of angles, and then is replaced by an identical neighbouring blade.)
- **Target Fan plus Clutter Fan 1.** Here, the Cluttering fan was a second small computer fan consisting of five 15 x 6 mm blades twisted through approximately 45°. Effective oscillation rate is 450 Hz.
- **Target Fan plus a Clutter Fan 2.** The slow Clutter Fan was a large office fan consisting of three 10 x 14 cm blades (shallowly twisted). Effective oscillation rate is 60 Hz.

In order to capture many cycles of the oscillating fan, the transmitter emitted long pulses (120 ms) at its optimal frequency of 50 kHz. Receivers were stationary during reception (*i.e.* no pan/tilt). Further details of the filtering process specific to this task and robot are given in Section 4.3.2.

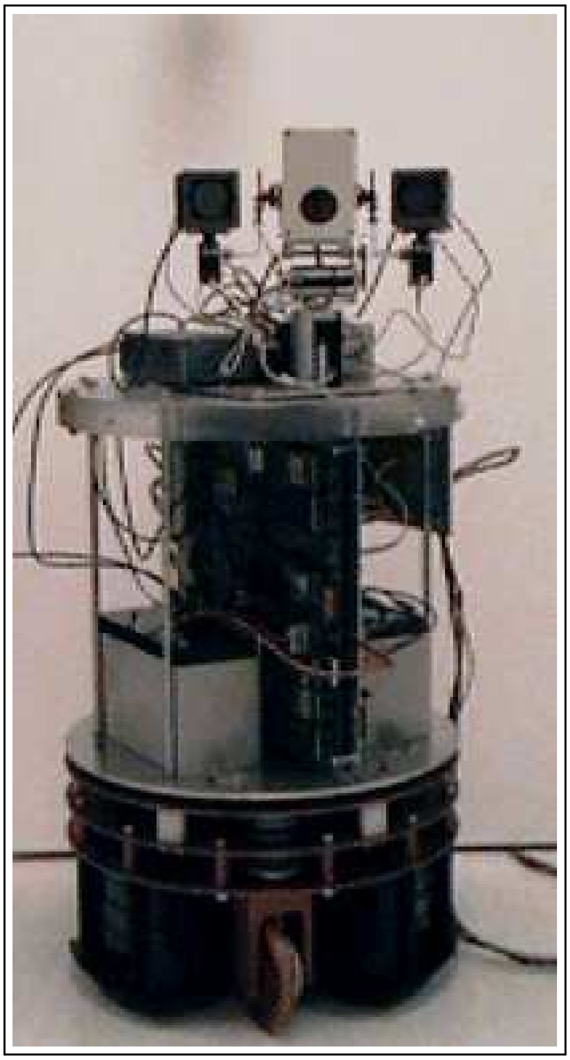

Figure 4.12:
Mobile robot Ben Hope. Ben consists of some in-house built sensors (currently only the bionic echolocation sensor) mounted aboard a B12 Real World interface base. All sensor signal processing and motor control calculations are performed aboard an array of transputers [Forster 91].

4.3.2 Results

Figures 4.13-4.17 show oscillograms, filter-bank responses and spectra of echoes arriving at the left transducer from the designated Target Fan located at several different azimuth and elevation positions within the frontal sound field. (Filter-bank parameters are similar to those used in the previous section.) The effective rotation rate of this fan is 540 Hz and, as shown, it produces sidebands at integer multiples of this rate. These figures show how energy in the sidebands rises and falls as the target moves horizontally (Figures 4.13-4.15) or vertically (Figures 4.16-4.17) away from the transducer. Movement of the target from an angular position in front of the transducer (Figure 4.13) to one at the horizontal periphery (Figure 4.15) yields echoes whose total energy increases (due to an increase in reflections off stationary targets adjacent to the fan), but sideband energy decreases. This is visible as a steady increase in the amplitude of the oscillograms (all oscillograms are scaled with respect to signal in loudest channel) and a movement of the spectral energy towards 50 kHz.

The apparent target aspect angle varies with changes in target bearing and, for a great many targets with moving surfaces, this phenomenon induces variations in β that influence the distribution of energy amongst the sidebands. (This effect was clearly visible in Figures 4.4-4.8 as a movement of the main spectral energy up and down the spectrum.) Aboard a head with large inter-aural transducer separations, right and left receivers may experience different Doppler viewing angles (*i.e.* different βs). In this case, differences in energy in corresponding frequency channels of the right and left filter-banks will be due to target orientation as well as bearing. However, in all cases examined in this work, the magnitude of the effect of orientation on sideband energy was small compared with differences in echo amplitude brought about by transducer directionality. Thus, target bearing could be computed reliably based on comparisons of energy in corresponding sideband frequency channels.

The signals collected for Figures 4.13-4.17 were taken under controlled conditions, where the Target Fan was held at a range of approximately 0.5 m. In the navigation experiment described below, the robot was positioned at $r = 2.0 - 4.0$ m away from the Target Fan and, from most initial orientations, received reflections off nearby cluttering

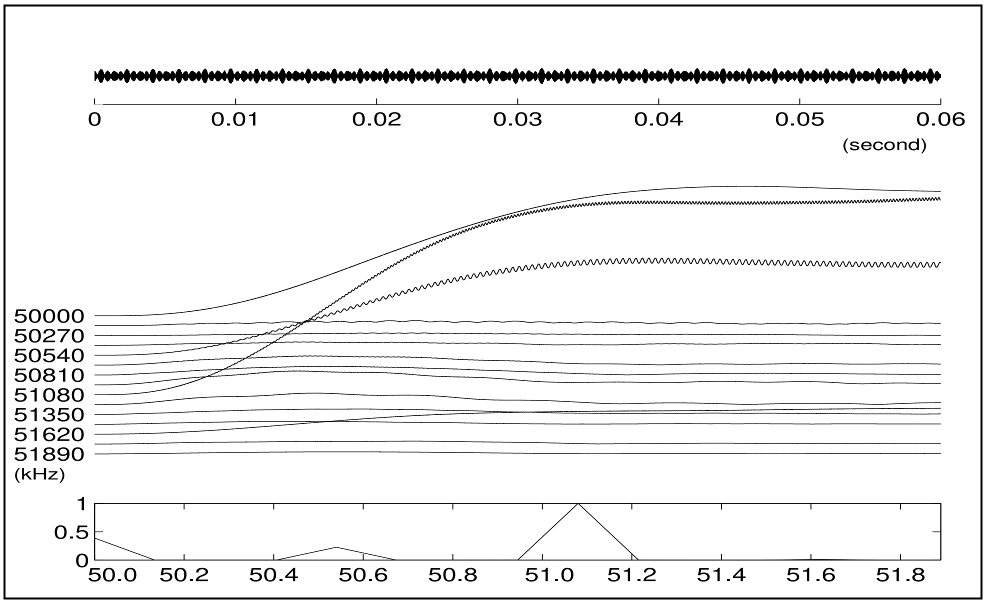

Figure 4.13:
Oscillogram, filter-bank and spectral representations of Target Fan at $(0°, 0°)$ **with respect to transducer normal.** All angular positions are given with respect to the position of the receiver. Oscillograms in Figures 4.13-4.17 are normalised with respect to loudest echo heard by the receiver (shown in Figure 4.15). For all these figures, $Q_{10dB} \approx 500$ for the (7th order Butterworth) bandpass filters. After rectification, signals are lowpass filtered in 2nd order Butterworth filters with $f_c = 1000$ Hz. Filter-bank and spectral plots are scaled according to the energy in the largest spectral channel of each individual echo.

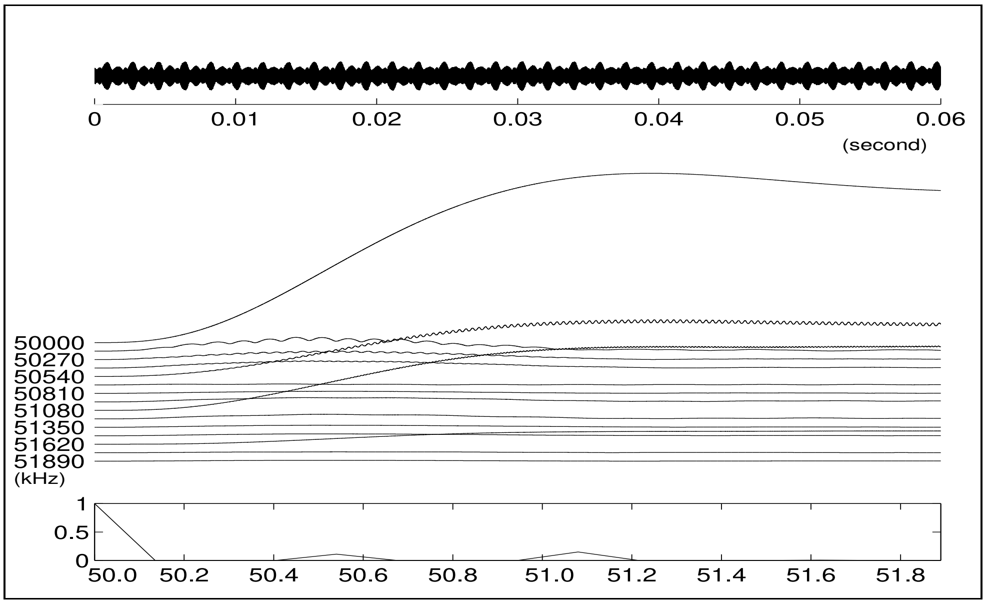

Figure 4.14:
Oscillogram, filter-bank and spectral representations of Target Fan at $(4°, 0°)$ **off-axis.**

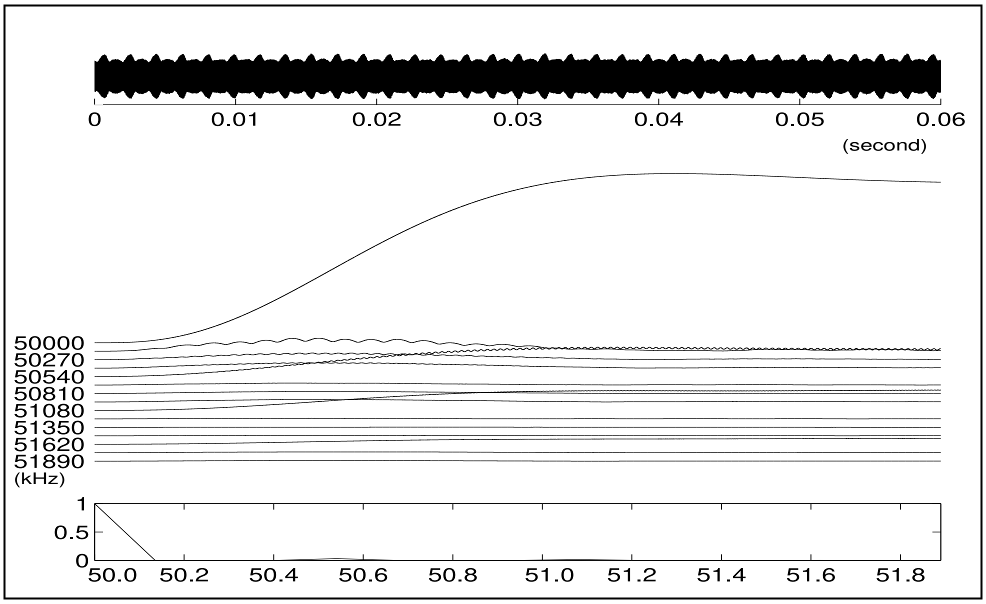

Figure 4.15:
Oscillogram, filter-bank and spectral representations of Target Fan at $(8°, 0°)$ **off-axis.**

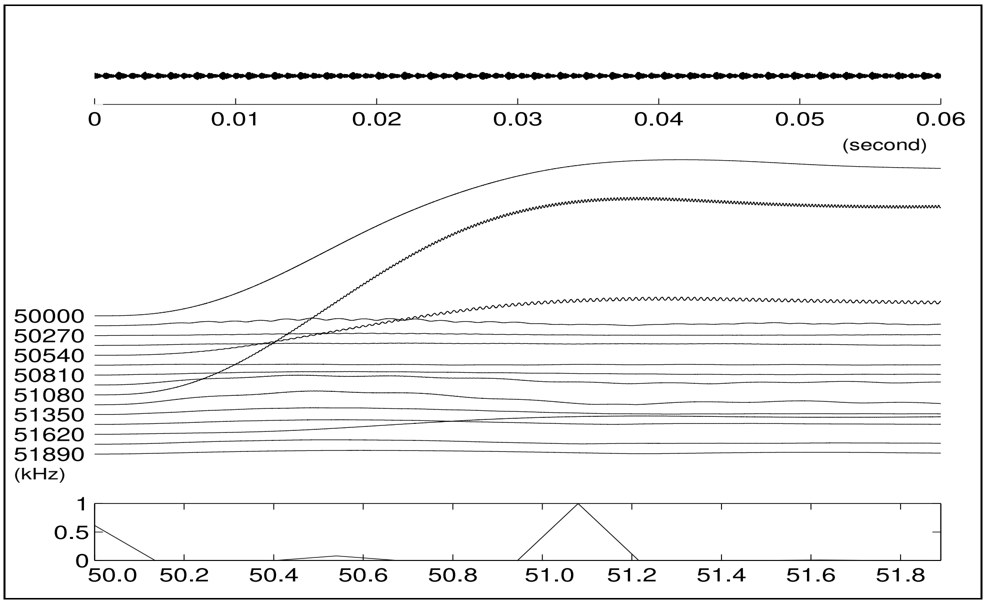

Figure 4.16:
Oscillogram, filter-bank and spectral representations of Target Fan at $(0°, 4°)$ **off-axis.**

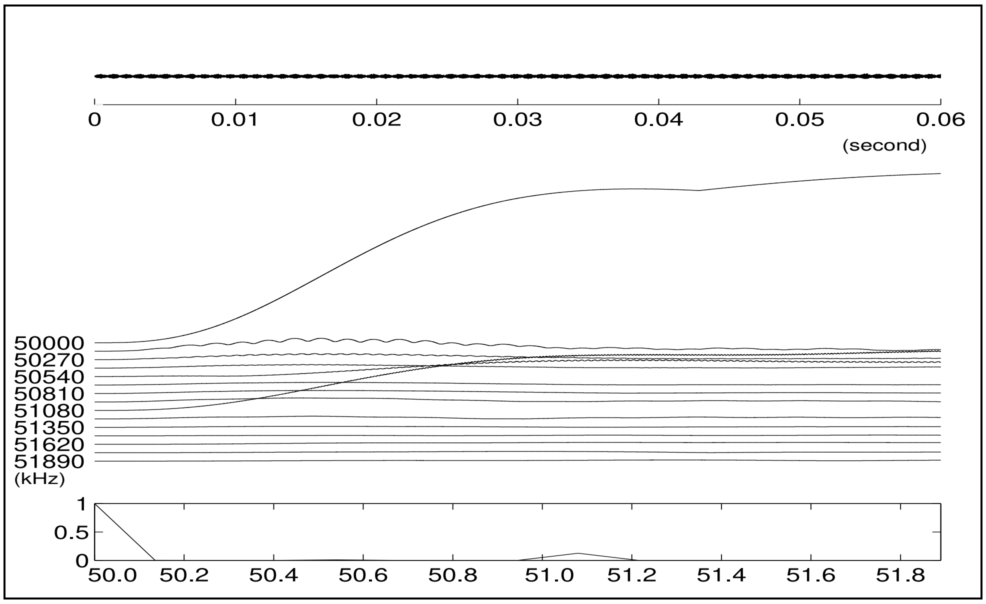

Figure 4.17:
Oscillogram, filter-bank and spectral representations of Target Fan at $(0°, 8°)$ **off-axis.**

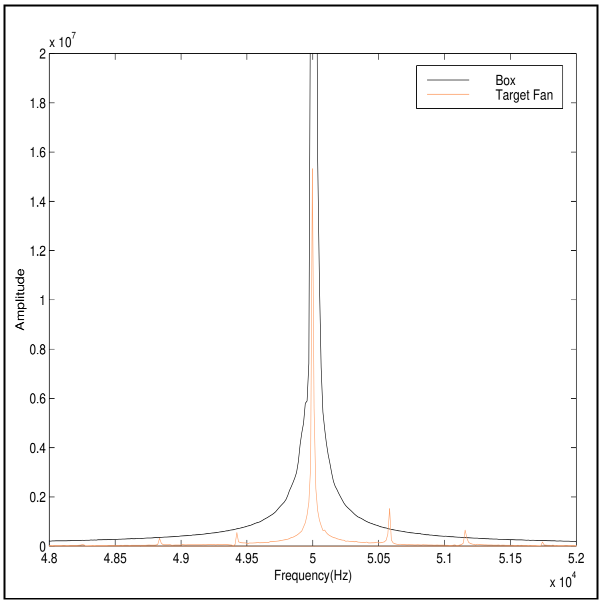

Figure 4.18:
Spectrum of echoes from stationary surface and Target Fan. Range of fan and box surfaces are 0.5 m.

surfaces which yielded echoes which were $1-3$ orders of magnitude larger than those from the small rotating Target Fan. In these cases, it is possible for spectral energy from the 50 kHz channel to spill over into adjacent channels, as shown in Figure 4.18. Unfortunately, the robotic system did not move sufficiently quickly to employ Doppler shift compensation to overcome this phenomena. Moreover, it did not have the signal processing resources to bandpass filter echoes in parallel across a large number of frequency channels. If the latter were a possibility, it might, for example, detect reflections from the fan *via* a weighting function which integrated energy at sideband positions (multiples of 540 Hz) and subtracted off energy in channels between sideband positions.

In our implementation, because the long pulses necessary to capture several cycles of the fan were slow to process, we used the simplest possible version of this scheme — *i.e.* we restricted the number of filtering channels to two (50 and 51.1 kHz). The localisation mechanism computed a ratio of echo energy $r^i = \frac{E(\phi)^i_{51.1}}{E(\phi)^i_{50.0}}$ (where i is used to denote the left or right receivers) which was then compared to a threshold (τ) in order to determine which of the following control laws to execute:

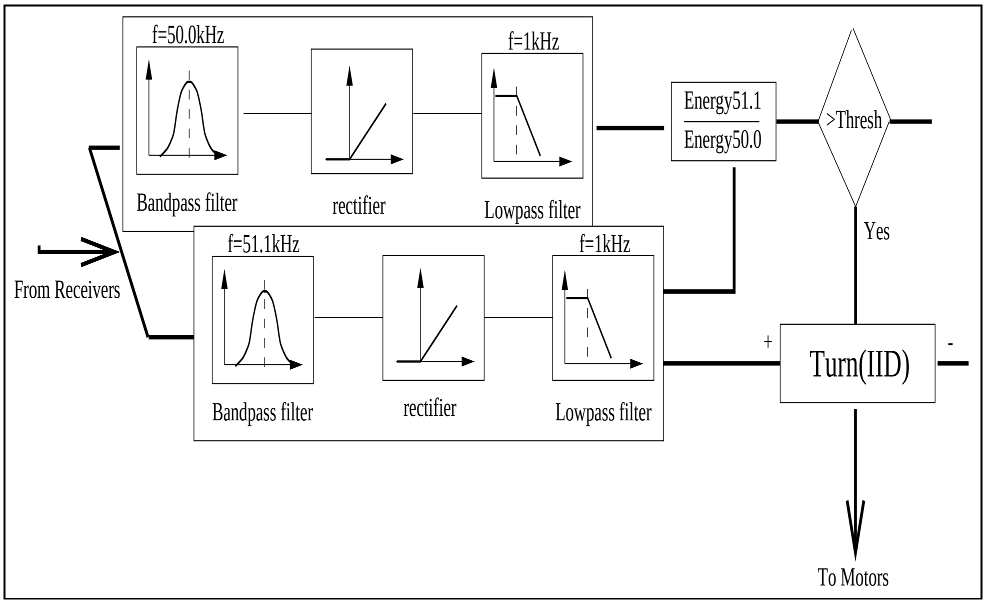

Figure 4.19: **Localisation mechanism used aboard the robot.**

$(r^{right}$ and $r^{left}) < \tau$:	$(r^{left}$ or $r^{right}) > \tau$:
$P_{translate} = 0\ cm$ $\theta_{rotate} = 3°$	$P_{translate} = 10\ cm$ $\theta_{rotate} \propto 20\ \log \frac{E(\phi)^{left}_{51.1}}{E(\phi)^{right}_{51.1}}$

Were we able to compute more filters so as to select targets on the basis of ratios of sidebands, the threshold would not be necessary. However, it proved useful for the experiments discussed below in that the mechanism yielded supra-threshold returns when receiving echoes from the Target Fan (high energy in the 51.1 kHz channel, *i.e.* numerator) or when echolocating into free space (low energy in the 50 kHz channel, denominator). In cases where the robot was in a position or orientation from which it was unable to detect the Target Fan, this artifact of the thresholding scheme facilitated exploration into open regions of the room where it often came into range of the Target Fan.

The series of figures and discussions in the following subsections demonstrate how the selective localisation of signature energy in the 51.1 kHz channel gives rise to target selective behaviour.

Target Fan

Figures 4.20 show 10 trajectories followed by the robot as it hunted for the Target Fan ((a) runs $T1-T4$, (b) $T5-T8$ and (c) $T9-T10$). In each case, the robot was started at a randomly selected orientation at a distance > 2 m from the Target Fan. In all but one case ($T1$), the robot was out of range (or orientated away from) the Target Fan such that it performed search rotations until a supra-threshold echo was received. Plus symbols plotted along the trajectory indicate positions/orientations where the robot did not receiver supra-threshold returns and, therefore, performed one or more of these rotations.

As the Target Fan was reliably detectable at a range of ≤ 2 m, many of the trajectory segments which are both at a greater range from the Target Fan than this and missing plus symbols correspond to situations wherein the robot was oriented into free space and therefore receiving supra-threshold returns due to a lack of energy in the 50 kHz channels (as opposed to the presence of significant energy in the 51.1 kHz channel). This artifact of the localisation mechanism served to drive the robot along a trajectory leading into free-space until (i) it caught echoes off the Target Fan ($T2$, $T3$, $T5$, $T6$, $T7$) or (ii) it came in close range of a stationary obstacle which drove up energy in the 50 kHz channel and invoked search rotations which eventually turned the robot toward the Target Fan ($T4$, $T8$, $T9$, $T10$).

In this regard, runs $T9$ and $T10$ are particularly illustrative. Here, the robot was initially orientated to the left of the Target Fan and made search excursions into the left upper (free-space) quadrant of the room. At a range of approximately 2 m from the upper wall, the sensor started receiving echoes which drove the ratios below threshold and caused the robot to make large, left-turning, circular search manoeuvre(s). As it emerged from these search rotations, the robot was attracted to free space on the right part of the room and moved in that direction until reaching a cluttered area (wherein the ratios fell below threshold again) and the robot swung back to find the Target Fan. These runs illustrate nicely how the target signature is salient from diverse approach angles.

No Target Fan present

Figures 4.21 show the tracks of the robot over approximately 20 minutes as it ran the localisation mechanism in the presence of no Target Fan. As before, the robot was

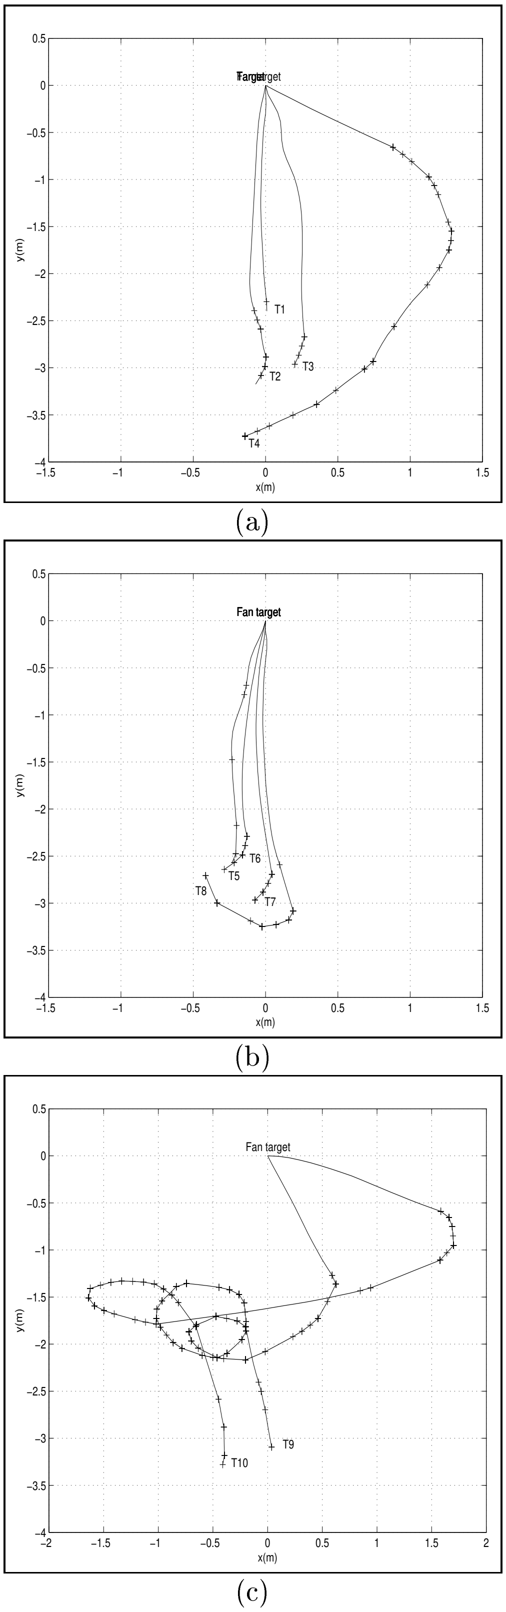

Figure 4.20:
Localisation: Target Fan amongst stationary clutter. Tracks of the robot in an $x - y$ coordinate system as it attempted to follow a sensory gradient toward the Target Fan. The Target Fan was located at the indicated position — dangling from the lens of a camera which simultaneously recorded the motion for demonstration purposes. (Distance between the upper (*i.e.* behind the Target Fan) and lower walls of the room is shown approximately to scale, but the side walls of the room extend much beyond the boundaries of the plot. Plus symbols associated with each track show portions of the trajectory where the system was searching.

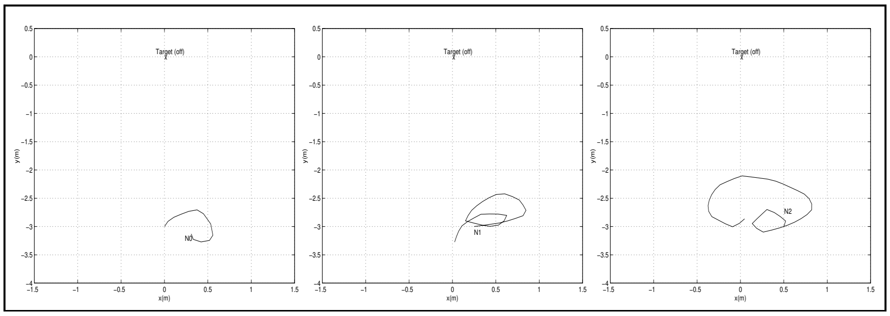

(a) (b) (c)

Figure 4.21:
Localisation(?): No Target Fan. Tracks of the robot in an $x-y$ coordinate system as it searched for a sensory gradient to follow toward the Target Fan.

booted at arbitrary initial orientations. In all three trials, the robot received several supra-threshold returns which enabled translation commands to be sent to the motors; however, because no Target Fan energy was present, movements were mainly search rotation manoeuvres which led to spiralling tracks. Indeed, the segments appear short because the robot spent a great deal of time in fixed-position search spins. (These tracks suggest that the echoes from the Target Fan received in the lower ($r > 2$ m) portion of the room, although they are unreliable, did encourage the robot to explore the upper half of the room in the previous experiment.)

Target Fan with Clutter Fan (450 Hz)

Placing another fan with a different rotation rate (another small computer cooling fan with 5 blades oscillating at an effective rate of 450 Hz) in the sound field has little effect on energy in the filter-bank 51.1 kHz channel. (Figure 4.22 shows that their spectra do not overlap.) Therefore, this Clutter Fan had no effect on the ability of the echolocating system to localise the Target Fan. Figures 4.23 (a) and (b) show the tracks of several trials wherein the robot navigated toward the Target Fan in the presence of this second fan located at distances of (a) 0.5 m and (b) 0.2 m from the Target Fan.

In this experiment, the high Q_{10dB} tuning of the filters was necessary in order for the system to distinguish between the fans. However, were it possible to run other filters in parallel, targets could be selected on the basis of more complex sideband patterns.

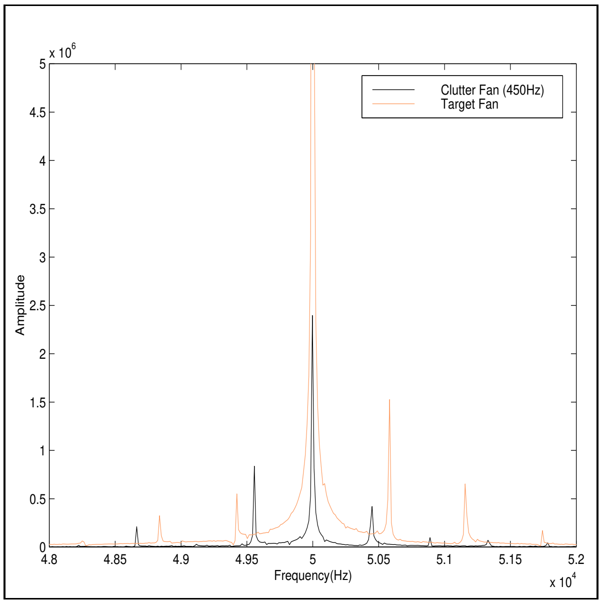

Figure 4.22:
Spectrum of echoes from the Target Fan and a second computer cooling fan. Both fans at $r = 0.5$ m.

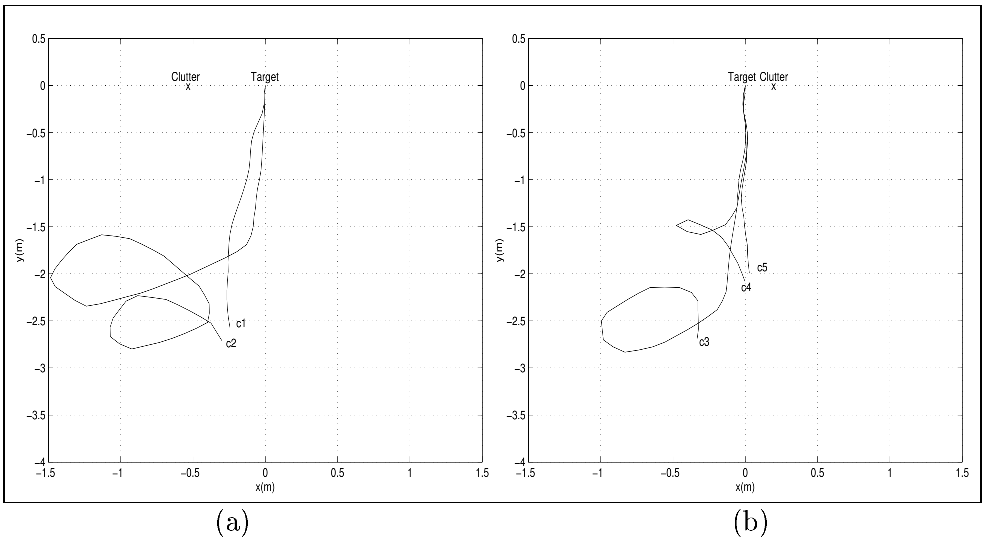

Figure 4.23:
Localisation: Target Fan in the presence of moving clutter (450 Hz).

Target Fan with moving clutter (60 Hz)

Placing a different Clutter Fan (with a rotation rate which is an integer divisor of that of the Target Fan) in the sound field has a profound effect on the ability of the localisation mechanism to detect the Target Fan. Figure 4.24 shows the overlapping spectra of the Target Fan and this Clutter Fan — *i.e.* an office fan with 3 large blades oscillating at an effective rate of 60 Hz. Using only two filter channels, it is not possible to separate these targets. Therefore, the robot localises the fan which is detected first (based upon its initial orientation). In Figure 4.25 (a), the initial orientation of the robot was toward the right side of the room and the Target Fan came into view before the Clutter Fan such that it was able to follow a sensory gradient toward the Target Fan. (In run $C2$ the Target Fan came into view immediately. In runs $C1, C3, C4, C5$, the robot started off moving into free-space and discovered the Target Fan during a search rotation at around $r = 2.3$ m.)

As shown in Figure 4.25 (b), when the robot started in an orientation wherein it first received echoes from the Clutter Fan, then it localised that fan. Moreover, because the Clutter Fan was larger and yielded stronger echoes, it was reliably detected from a larger distance. Notice also how the trajectories toward the Clutter Fan are smoother, due to the strong (*i.e.* reliable) signals that it reflected. Runs $C9$ and $C10$ show cases where, after a search rotation in which the robot spun through the direction of the Target Fan, the robot still localised the Clutter Fan because the Target Fan was out of detection range.

Binaural comparison of filter-bank energy which is weighted to favour the 51.1 kHz channels yields a localisation (or "target selection") signal whose magnitude and sign can determine the yaw necessary to direct an echolocator towards the target reflecting this signature. The use of weighted spectral energy exploited in the (closed-loop) 2D localisation experiments discussed in this chapter is the basis of the (open-loop) 3D localisation mechanisms described in Chapters 6-7 of Part III of this thesis. While it might also be possible to derive target aspect angle descriptions from the ratios of sidebands, this information did not appear useful. Likewise, the spectral broadenings might serve as an elevation cue for targets which tend to move in the same plane (*e.g.* insects which only fly right-side up, fans which sit vertically on a stand, *etc.*); however,

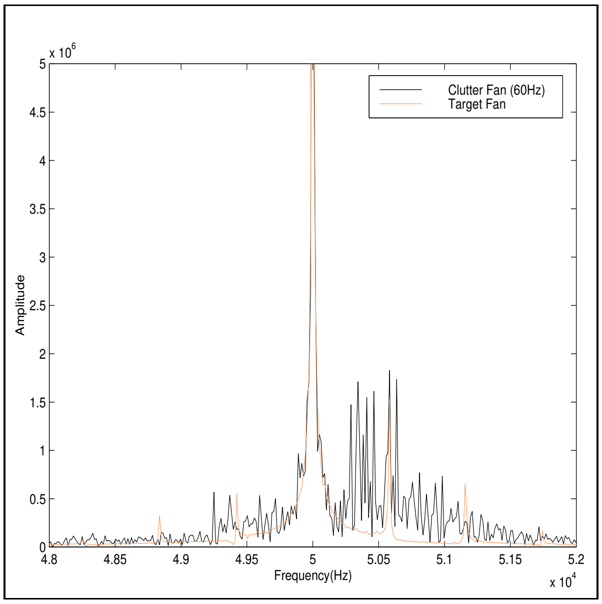

Figure 4.24:
Spectrum of echoes from Target Fan and large office fan. Target Fan at $r = 0.5$ m. Clutter Fan at $r = 1$ m.

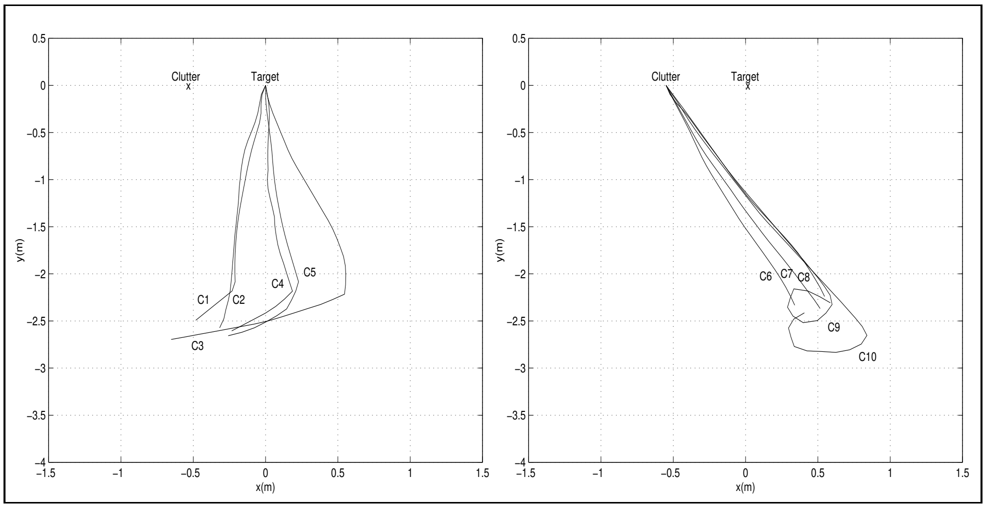

(a) (b)

Figure 4.25:
Localisation: Target Fan in the presence of moving clutter (60 Hz).

I did not want to make this assumption for more general targets. Therefore, elevation information is derived from the modulations induced by observer rather than target motion, as described in Chapters 6-8.

4.4 Discussion

How echoes are encoded in the bat's nervous system so as to highlight the essential differences in targets is unclear. It has been speculated that a bat's ability to show aspect-angle invariant selection of a target must imply model-based recognition — *i.e.* that bats integrate target descriptions collected from different orientations into a single remembered 3D "geometric model" of that target. From investigations of orientation invariant prey selection among rhinolophids, authors have drawn conclusions about the cognitive abilities of bats which are similar to those of Grinnell reported at the end of Section 3.1.2 (von der Emde and Schnitzler reported in [Moss & Schnitzler 96]):

> *"This finding suggests that bats may develop a three-dimensional representation of a fluttering insect from acoustic information contained in echoes from a single view. The echo of a fluttering insect insonified from a single orientation may thus be sufficient to construct a three-dimensional representation of the moving prey, and bats may use this representation for target classification."*

Far too little is known about neural processing in bats for one to rule out the possibility that they are able to map back from an echo(es) onto the original geometry which gave rise to it. Likewise, it is also possible that bats associate the more primitive spatio-temporal patterns in the reflections of an insect insonified from different angles into a single insect category. For example, Schuller claims that for "insect identification the bat would have to associate a given species of insect with the entire family of echoes reflected at various attitudes" [Schuller 79]. If this is true, we might say that the bat has formed a "natural category" [Herrnstein 85]. (The formation of natural categories has been used in several cases to describe the interpretation of visual signals [Herrnstein *et al.* 76, Irle & Markowitsch 87].) It is still a matter of debate as to how animals can achieve natural categorisation — *i.e.* to what extent specific local cues *vs.*

"computational process" are involved. However, it has been suggested that, because a simple parameter(s) does not characterise reflections from insects insonified at all aspect angles, high duty-cycle bats use "cognitive" processes to recognise prey.

In the present chapter, I have looked at an explanation of target cues available in echoes, with a view toward arriving at an acoustical characterisation of targets which is both unique and robust in ecologically realistic environments and sensing tasks. The results reported here demonstrate that there are invariant spectral cues in the echoes reflected from targets with periodic motion which an echolocator can use for localisation on the basis of only immediate cues. Though it remains to be seen whether insects behave as suggested by the model and whether bats can/do localise targets based only on locally available cues, the results presented here do show that target selective behaviour need not be mediated by formal recognition of targets, but, rather, can be under-pinned by a coupling between echolocator and environment to the extent that the signature signal provides a steering signal which is used directly in localisation.

In his book *Visual Explanations* [Tufte 97], Edward Tufte captures nicely this notion that design success (or, similarly, a successful explanation of design) arises out of appropriate characterisations of the invariant attributes of information:

> *"Clutter and confusion are failures of design, not attributes of information. There's no such thing as information overload."*

Part III

Structuring sound at the receiver

"In no other field of science ... does a stimulus produce so many different sensations as in the area of directional hearing."

—von Békésy

"By indirections find directions out."

—William Shakespeare (Hamlet, Act. II)

5. *Background*

The placement of two receivers on opposite sides of an acoustic perceptual system generates inter-aural intensity and timing differences which are powerful lateralisation cues. Adding pinnae to receivers creates directional cues which enable animals — using only two ears — to do more than simply lateralise a sound source/reflector; but to "project" the source out to a well defined location having attributes of front-back-above-below. In this part of the thesis, I examine intensity (Chapters 6 and 7) and spectral (Chapter 8) cues which may play a role in 3D target localisation. Although ITDs are not considered, other uses of time in target localisation — *i.e.* those whose acoustical cues do not degrade with head size — are investigated in Chapter 7. In the present chapter, I discuss the reception process in more general terms.

5.1 The role of pinnae

As the physical effects of sound propagation make echolocation a rather short range sensing modality, acoustic receivers (principally pinnae) are a central component which can maximise the angular resolution (and operating range) of an echolocation system. (Call features are another, but increasing call strength also makes the echolocator more conspicuous to potential tympanate prey [Roeder 70], conspecifics [Balcombe & Fenton 88] and predators [Fenton 80].) Although the specific acoustical cues generated by the complex pinnae of animals and the way(s) in which available cues are transformed into spatial percepts are not well understood, behavioural studies confirm their importance to bats. Distortion or displacement of the external ear reduces a bat's spatial acuity [Grinnell & Schnitzler 77, Lawrence & Simmons 82, Mogdans *et al.* 88, Jen & Chen 88], and its ability to detect a target [McCarty & Jen 83] and to avoid obstacles [Griffin 58]. Bat pinnae are well developed and, in many species, exceed the dimensions of the head [Coles & Guppy 86]. The variety of actual pinnae shapes and sizes within Microchiroptera suggest that, just as there are a variety of different approaches to calling in echolocation, there are also a variety of approaches to reception.

5.1.1 Gain

Undoubtedly, one important function of the pinnae is to serve as collectors and resonators of sound in behaviourally relevant frequency ranges. The pinna dimensions of some echolocating bats are well matched to the dominant frequency of their echolocation signals. For example, the wavelength of the fundamental component of the shallow FM search phase echolocation calls of some Molossidae is just twice the radius of the pinna mouth ($\lambda \approx$ pinna height) [Guppy & Coles 88]. These frequencies are likely to be highly amplified by the external ear.

Further examples of the role of pinnae in amplifying sound are abundant amongst gleaning bats (*i.e.* bats which hunt by listening for prey generated noises). These bats possess some of the largest pinnae of all Microchiroptera. In some species, large ears extend the gleaner's low frequency hearing range by increasing pressure gain in the ear canal. This is true of many megadermatids and nycterids which detect prey (namely, arthropods or small vertebrates) *via* the rustling sounds ($<$ 25 kHz) they generate when moving over dry substrate [Obrist *et al.* 93]. The ears of these bats amplify (gain $\approx 5 - 10$ dB) sounds below 15 kHz. By contrast, the phyllostomid *Trachops cirrhosus* — also a gleaner — uses its large ears to increase directional cues. The pinnae of *Trachops cirrhosus* produce an even $5 - 10$ dB gain over $5 - 95$ kHz — suggesting that its prey (*i.e.* frogs which make a loud, low frequency ($<$ 4 kHz) croak [Tuttle & Ryan 82]) generate sufficient energy for detection. However, the pinnae show exceptionally increased directionality around the frogs' call frequency: IIDs rise to 27 dB at 5 kHz.

5.1.2 Directionality

In many bat species, there is an increase in directionality (in particular, IID resolution) for the dominant wavelength component(s) of search-phase calls or prey generated noises. This suggests that pinnae also play an important role in generating the directional cues necessary for 3D localisation [Simmons 82, Obrist *et al.* 93]. In this context, "directionality" determines how sound originating from any given location in space is modified by the external auditory system and, therefore, provides information which may be used to localise the sound source. Thus, one might think of pinnae as acoustical magnifying glasses which, when held in front of the ear canal, increase the sensitivity of the auditory system for a particular region of the frontal sound field.

For example, the circular piston transducer used to model pinnae in these studies acts as a very simple magnifying glass which is most sensitive to sound arriving from the direction normal to the piston surface. It attenuates, or filters, sound arriving from off-axis angles due to diffraction across the surface (of radius a) of the piston as:

$$H(f,\theta) = \frac{J_1(d)}{d}$$

where $d = \frac{2\pi f\, a\, sin(\theta)}{c}$ (see Appendix A). Figure 5.1 shows the sensitivity of a piston receiver for apertures equal to (a) the radius of the Polaroid series 7000 transducer (operating at its optimum frequency 50 kHz) and (b) the radius of a typical *R. ferrumequinum* pinna (83 kHz). (Mentally rotating the 2D polar radiation plots of Figure 5.1 around their acoustic axes at 0° yields the 3-D sensitivity pattern of each symmetric receiver.)

The directionality of a binaural echolocation system is determined by both its basic transducer directionality and the spatial relationship between the transducers. Qualitatively, the transducers must be orientated in such a way that targets are likely to fall into the beam of both receivers, while maintaining a sufficient angular separation between the receivers to create the inter-aural differences necessary for localisation.

Figure 5.2 (a) shows the directionality of a binaural (two piston) auditory system in response to 83 kHz pulses arriving from a source at different angular positions along a 0.5 m radius from the centre of the head. (Here two piston receivers were

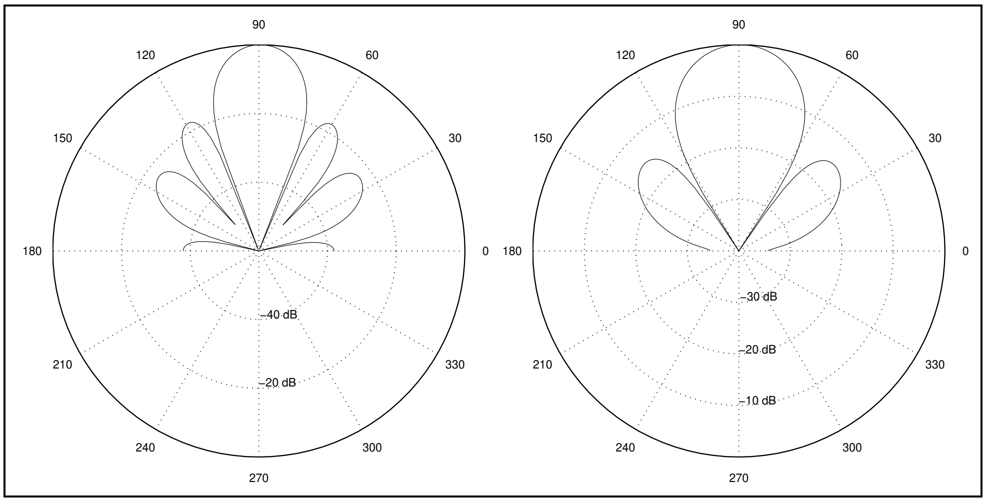

(a) (b)

Figure 5.1:
Transducer directionality. Predicted directional sensitivity of the piston receiver. (a) Polaroid series 7000 transducer: $a = 1.13$ cm, sound frequency $f = 50$ kHz. (b) *R. ferrumequinum* dimensions: $a = 4.5$ mm, sound frequency $f = 83$ kHz.

mounted on a simulated *R. ferrumequinum* head, with dimensions shown in Figure 4.1.) When receiving reflections of its own vocalisations, the emitter and receivers are equal partners in hearing such that receiver directionality must be determined in the context of the directional emission. Due to the loss of energy in the emitted pulse as it spreads away from the centre of the field, the resolution of the echolocation system as a whole is sharpened. As shown in Figure 5.2 (b), when receiving 83 kHz echoes of its own emitted sound, the sensitivity regions move inward toward the vertical mid-line and attenuation at the horizontal periphery of the sound field occurs more rapidly. In the limit that the receivers are coincident, the orientation of the receiver's acoustic axis (within the $0°$ elevation plane) would be $\theta_{azi} = \frac{\theta_{trans}+\theta_{rec}}{2}$ (for identical transmitting and receiving transducers) across all target ranges. For the simulated *R. ferrumequinum* head shown in Figure 5.2 (b), the inter-aural distance is small and the (azimuth) orientation of the acoustic axes is well approximated by $\theta_{azi} = \frac{0 \pm 15°}{2}$.

Increased directionality of the echolocating system with respect to receiver directionality has been observed in the combined echolocation systems (*i.e.* emitter plus receivers) of several bat species [Shimozawa *et al.* 74, Grinnell & Schnitzler 77, Fuzessery *et al.* 92]. Moreover, in bats, a great variety of emitting structures have

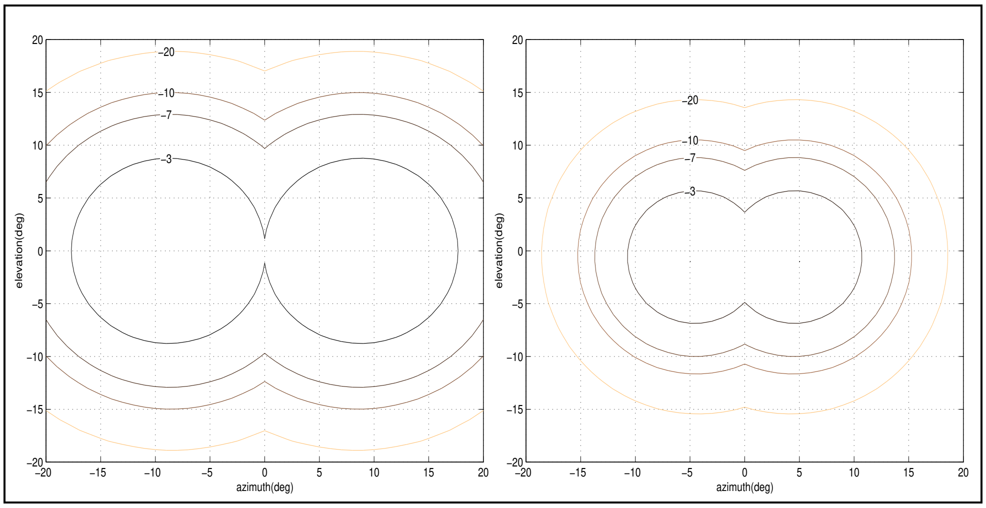

(a) (b)

Figure 5.2:
Predicted binaural directionality. Predicted binaural sensitivity of the (a) receiver and (b) echolocation system with respect to the angle of incidence of a 83 kHz sound source/target located at given angular bearing along a 0.5 m radius from the mid-line between the receivers. Shown are sensitivity values for each receiver in response to sound arriving from its own ipsilateral frontal sound field. Transducers are mounted in a common elevation plane and separated by horizontal distance of 8 mm along an axis perpendicular to the frontal sound field. As shown in Figure 4.1, the normal to the transmitter surface points directly forward and receivers are orientated at $\theta_{rec} = \pm 15°$ away from the mid-line.

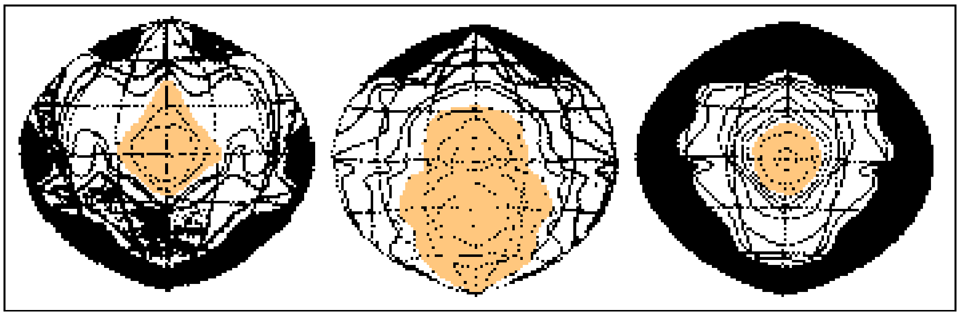

(a) (b) (c)

Figure 5.3:
Directionality of *R. ferrumequinum*. (a) Receiving system. (b) Emission. (c) Combined effects. Shaded areas show the < -10 dB intensity regions. (Data taken from [Grinnell & Schnitzler 77].)

evolved whose radiation patterns appear to complement the pinna directionality. (A brief discussion of beam-forming in high duty-cycle bats is given in Section D.1.) For example, hipposiderids and rhinolophids are nasal emitters whose fleshy appendages (*i.e.* noseleaves) decorating the muzzle play a role in shaping and directing the emitted beam. In the emission pattern of *Rhinolophus ferrumequinum*, there is a strong asymmetry along the vertical axis created by a downward pointing side-lobe which appears to compensate for a region of low hearing sensitivity in the same area due to shadowing by the snout [Schnitlzer & Grinnell 77, Grinnell & Schnitzler 77]. Figure 5.3 illustrates this phenomenon. The iso-intensity contours of the echolocation system assume a greater radial symmetry about the centre of the sound field than is seen in either receiver or emitter directionality contours. Because of this, the intensity of an echo will drop at a similar rate as a target moves in any direction away from the centre of the sound field.

In bats, the basic directionality which an animal is born into — by virtue of such anatomical factors as pinna size, shape, position and orientation atop the head, *etc.* — is dynamically altered during sensing so as to create additional directional cues. In this regard, one might distinguish two different dynamic approaches employed by species echolocating with different duty-cycles and bandwidths. The first — employed by many broad-band emitting species — exploits that fact that, due to the passive filtering properties of stationary pinnae, the acoustic axes point into different regions of space at different frequencies (Section 5.2.1). Bats echolocating with CF calls cannot "call up" different acoustical magnifying glasses, but they can create additional acoustical cues *via* dynamic reorientation of the acoustic magnifying glass created at the CF frequency (Section 5.2.2). These approaches are discussed in the following section.

5.2 Dynamic directional sensing

5.2.1 Pinna morphology

The ears of bats are complex physical structures — *e.g.* many taxa have special ridges and folds in their pinnae, and most microchiropterans have a tragus (*i.e.* a fleshy projection on the anterior edge of the ear opening) which can be deflected during echolocation [Vaughan 86]; the funnel-shaped pinnae of the Mormoopidae, each with a complex array of folds, crenellations, and leaf-like structures, are amongst the most elaborate belonging to any mammal [Kunz & Pierson 94]. Some authors have tried to model these effects using global cues. For example, Guppy and Coles claim that if the opening to the pinnae is treated as the oblique truncation of a (right) conical horn (see [Beranek 93]) diffraction effects predict receiver sensitivity lobes whose position and sharpness correspond well with those measured in *Macroderma gigas* and *Nyctophilus gouldi* [Guppy & Coles 88]. However, an accurate model of directional hearing in bats must also incorporate *local* cues involving reflections off internal structures which may play a role in determining the directionality of the ear at different frequencies [Wotton *et al.* 96].

Most of what we understand of the process of fine, or local, acoustical filtering comes from work on humans. Batteau proposed that the human auditory system uses the time delays of sound reverberations generated by reflections between various ridges and folds of the pinnae to represent vertical and horizontal directions [Batteau 67]. Each position of a sound source could be represented by a specific temporal pattern of delays because the path lengths created by the ridges and folds differ systematically according to the position of the incoming sound. (Batteau identified two different delay paths, recent studies suggest that the actual number of delays may be larger — between 8 and 16 [Chen *et al.* 92].)

Preliminary studies of the finer acoustic properties of the external ear of *Eptesicus fuscus* reveal that sounds entering the ear produce primary and secondary reflections (presumably off the tragus) which are strongly influenced by the elevation of the sound source [Lawrence & Simmons 82]. In *Eptesicus fuscus*, when the tragus is deflected, vertical discrimination performance (normally 3°) deteriorates.

These delays yield interference effects which provide spectral cues as well because the extent to which incoming sound at any particular frequency is cancelled or reinforced depends upon the delay, or phase relationship, between the original and reverberated copies converging on the cochlea. For a broad-band sound passing through the human pinnae, interference between direct and delayed versions of incoming sound produces a comb-filtered spectrum with notches and peaks at different frequencies according to the magnitude of the delay separation imposed by the external ear as a function of source location. In humans, there is at least one prominent notch in the spectrum which varies in centre frequency with vertical angle [Roffler & Butler 68, Shaw 74]. Artificial manipulation of the position of the notch leads to a perception of change in sound source elevation [Bloom 77].

Investigations in various animals have shown that there is a systematic variation in the details of spectral shape for broad-band sounds and this phenomenon is believed to be important for sound localisation in bats with fixed pinnae [Grinnell & Grinnell 65, Fuzessery *et al.* 92, Lawrence & Simmons 82, Obrist *et al.* 93]. Most bats which employ a range of frequencies for echolocation appear to hold their pinnae in a fixed position relative to the head and take advantage of the fact that, due to pinnae morphology, reflections from targets at any given point in space will give rise to different IIDs across the different frequencies. In other words, these bats actively sample different regions in space by calling up different frequencies/acoustical magnifying glasses.

For example, the pinnae of *P. parnellii* are highly directional at the second (60 kHz) and third (90 kHz) CF echolocation harmonics. They amplify 60 kHz sound when it originates from off-centre azimuth regions whose elevation coordinates lie between 0° and −30°, while they are most sensitive to 90 kHz sound arriving from the horizontal mid-line between −20° to −40° elevation [Fuzessery *et al.* 92]. (These difference areas are well within the beam of the bat's echolocation pulse.) While sound localisation using any single harmonic is ambiguous, simultaneous comparison of inter-aural intensity differences at the second and third harmonics (and possibly the fundamental) can disambiguate position in both azimuth and elevation.

5.2.2 Pinna motion

In the previous section, I discussed how the pinnae of bats act as acoustical magnifying glasses whose resolution and orientation change with frequency [Obrist *et al.* 93]. Thus, by employing a broad-band call, a stationary bat can sample several regions of 3D space within the beam of a single echo. Cues naturally occurring in the echo power spectrum will not be available to bats echolocating with a single CF call — *e.g.* many species in the families Rhinolophidae and Hipposideridae. Members of both families echolocate using the second CF harmonic and are distinguished by the lack of a tragus and by having highly mobile pinnae [Möhres 53, Schneider & Möhres 60]. Some rhinolophid and hipposiderid bats[1] independently rotate their pinnae, like two opposing antennae, during echolocation. The mobility of the pinnae may be employed in search-related tasks. However, these bats appear to investigate potential targets more systematically *via* vertical scanning motions which involve a rotation of one ear forwards while the other rotates backward. One such movement appears to be made for each CF pulse/echo — with the right and left pinnae moving along equal and oppositely signed arcs during one pulse/echo, and reversing this motion during the next pulse/echo, and so on [Schneider & Möhres 60, Griffin *et al.* 62, Pye *et al.* 62, Pye & Roberts 70, Gorlinsky & Konstantinov 78, Mogdans *et al.* 88]. Better correspondence between ear movements and pulse emission has been observed in cases where pulse emission rates are high — *e.g.* during the terminal buzz phase of echolocation, when "interest" and need for information is high [Griffin *et al.* 62]. However, other authors report good numerical correlation across a variety of sensing scenarios [Pye *et al.* 62].

The use of movement rather than bandwidth to alter the acoustic axis can be seen as another effort by high duty-cycle bats to confine analysis to the narrow window of the

[1] Bats in which this motion has been observed include Rhinolophidae: *R. fumigatus*, *R. ferrumequinum*, *R. alcyone*, *R. euryale*, *R. hipposideros*, *R. landeri*; Hipposideridae: *H. commersoni*, *Triaenops afer*, *Asellia tridens*, and *H. caffer*.

acoustic fovea. Although the specific acoustic cues generated by pinnae movements and the way in which available cues are transformed into spatial percepts are not well understood, behavioural studies have confirmed the importance of ear movements in target localisation: immobilisation of the pinnae causes disorientation and a loss of localisation acuity in *R. ferrumequinum* [Schneider & Möhres 60, Gorlinsky & Konstantinov 78, Mogdans *et al.* 88]. It is interesting to note that when deprived of the use of the muscles which move the pinnae relative to the head, *R. ferrumequinum* become disoriented and lose localisation acuity until they learn to compensate by moving the head vigorously [Schneider & Möhres 60, Mogdans *et al.* 88]. *Pteronotus parnellii* has relatively immobile pinnae, but has been observed to move its head in much this same way when emitting CF pulse [Schnitzler 70]. (Rapid movement of the ears is used by other Phyllostomidae/Mormoopidae, but the movements appear to be elicited by external sounds, unrelated to the emission of their short broad-band pulses [Pye & Roberts 70].)

Figure 5.4 shows the alterations in directionality which come about from tilting the simulated bat head (dimensions given in Figure 4.1) through $\pm 15°$. Where iso-intensity contours intersect (*e.g.* along the vertical mid-line in Figure 5.4 (b)), inter-aural intensity differences fall to 0 dB. In the orthogonal direction lies what I shall refer to as the SONAR horizon — *i.e.* the region in space wherein inter-aural intensity differences vary most sharply. In the case of a head tilting motion (Figure 5.4), the SONAR horizon is a plane containing the acoustic axes. If the echolocator uses receiver, rather than head, movement to alter the orientation of the SONAR horizon, the acoustic axes no longer lie in the same plane and the SONAR horizon is no longer strictly a planar surface. However, for the same degree of rotation, more dramatic changes in the orientation of the SONAR horizon can be brought about by employing opposing pinnae movements as shown in Figure 5.5.

In order to estimate the movement of the SONAR horizon with pinnae movement, one might make the geometrical simplification that the receivers are coincident such that the direction along which the IID gradient is the steepest is a surface with a range invariant orientation of $\theta = atan(\theta_{ele}/\theta_{azi})$. As shown in Figure 5.5, this is not a bad approximation for an echolocator with a small head. There, pinnae with azimuth orientations of $\pm 15°$ away from the vertical mid-line (as shown in Figure 4.1) have been

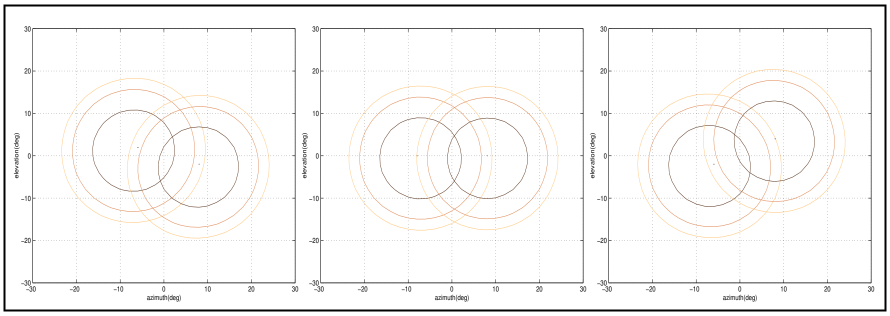

Figure 5.4:
Movement of the acoustic axes as a function of head tilt angle.(a) 15°. (b) 0°. (c) −15°. (In each figure, the $-3, -7, -10$ dB intensity contours of each receiver are shown.)

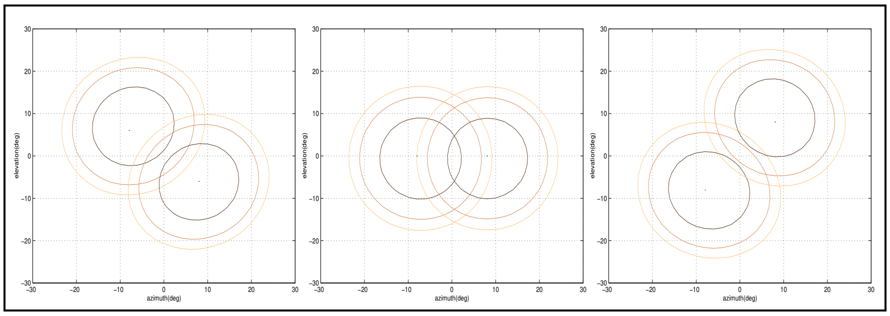

Figure 5.5:
Movement of the acoustic axes ***via*** **opposing pinna movements.** (a) 15°. (b) 0°. (c) −15°.

rotated through opposing vertical pinnae motions of $\pm 15^\circ$ which "tilts" the SONAR horizon through approximately 45°.

Tilting of the SONAR horizon so as to alter the directionality of the echolocation system is one possible function of the pinna movements. In Chapter 6, I explore the hypothesis that a CF emitting echolocator can use pinna *movement* rather than *morphology* to exploit the same localisation principle believed to be employed in broad-band echolocation — *i.e.* the use of more than one SONAR horizon along which to localise sound. By moving a single, fixed-width acoustic magnifying glass (defined by the frequency of the dominant CF harmonic) to a series of different positions during reception, rhinolophids and hipposiderids may provide themselves with additional viewing perspectives across which to sample and compare IID values in the same way that broad-band emitting bats may compare IIDs across frequencies.

In Chapter 7, I investigate the possibility that receivers may be scanned to create dynamic cues, rather than other measurement perspectives. In this regard, I explore the generation and use of timing cues. For example, receivers sweeping vertically across a target set up amplitude modulation patterns in returning echoes which change systematically with target elevation. Monaurally, target elevation is encoded in the temporal disparity between echo arrival time and echo amplitude peak time. The difference in echo amplitude peak times and intensities in the right and left ears create IID rates of change [Schnitzler 73] which also encode target elevation. Unlike ITDs, the resolvability of timing cues generated in this sensory context are determined by the sensitivity and speed of movement of the ears (*millisecond* cues), rather than the inter-aural dimensions (*microsecond* cues).

Given the fine frequency discrimination of some rhinolophids and hipposiderids, it is natural to suppose that the movement of the ears plays a role in producing local Doppler shifts which aid in target separation and (or) localisation [Pye & Roberts 70]. In Chapter 8, possible utilities of Dopper cues are investigated in simulation.

6. Investigation One: IID Maps

Auditory cues give rise to strong spatial percepts and, yet, the surface of the auditory sensor, unlike the visual sensor, is not laid out topographically with respect to space. Because of this, it is intriguing to ask whether (and, if so, where) in the auditory system a spatial map may be assembled.

6.1 Hypothesis

Motivation

Patterns of response across iso-frequency populations of neurons with particular aural preferences and different IID sensitivities may provide a substrate on which spatial maps are formed. The inferior colliculus is a likely host for such a map, as it is the first point in the auditory system that has major output to motor pathways and movement coordination systems [Casseday & Covey 96]. In *P. parnellii* there is a ventromedial population of binaural cells in the CF_2 contour of the central nucleus of the inferior colliculus in which IIDs are generated *via* a subtractive process wherein supra-threshold sounds originating in the excitatory (contralateral) ear evoke a certain discharge rate which is inhibited by sound presented to the ipsilateral ear when the ipsilateral sound intensity exceeds a certain threshold level [Wenstrup *et al.* 86, Wenstrup *et al.* 88a, Wenstrup *et al.* 88b]. Inhibitory thresholds (*i.e.* levels at which a cell's discharge rate is reduced by 50%) vary from 15 to -25 dB — encompassing much of the range that the bats would experience

[Wenstrup *et al.* 88a]. Moreover, the representation of different thresholds/IIDs by these "excitatory-inhibitory" (EI) neurons is topographically organised within this region such that neurons with high positive inhibitory thresholds (*i.e.* high, positive IIDs) are located in the dorsal EI region and there is a progressive shift to lower IIDs among EI neurons located in more ventral regions (see review in [Pollak & Park 96]).

The IIDs created by the acoustic properties of the ears (and head) combined with the topographic representation of the inhibitory thresholds within EI neurons suggests that the angular bearing of a target in space may be represented (at least partially) as the position of a border separating a region of discharging cells from a region of inhibited cells [Fuzessery & Pollak 85, Wenstrup *et al.* 88b, Pollak *et al.* 86]. Such an organisation might host the representation of space across the frontal sound field; however, not unambiguously.

We can observe the ambiguities in single-tone localisation by examining the directionality of the robotic echolocation system in terms of its IID sensitivity. In Figure 6.1 a population of cells — coding for IIDs varying from 20 dB (white) to −20 dB (black) — are laid out on an iso-frequency grid depicting the target bearing which give rise to them (based upon measured inter-aural disparities resulting from reflections of 50 kHz echolocation pulses). Across a restricted region within the SONAR horizon (*e.g.* the 0° elevation plane in which the acoustic axes lie) IID values vary monotonically. However, due to symmetries in the directionality of the echolocation system, the same IID measured within a given elevation plane may represent both a medial and lateral azimuth bearing angle — see Figure 6.2. Additional ambiguities arise along the vertical mid-line — where 0 dB IID values code for all elevations. Although the precise IID sensitivity of an auditory system is dependent upon its particular pinna/receiver and head structures, these sorts of ambiguities exist in all binaural auditory systems (*i.e.* evolved and engineered). Moreover, in all physical systems, the border between regions are fuzzy due to measurement noise. (The individual IID measurements shown in Figure 6.1 fluctuated by approximately 2 dB between measurements taken under similar circumstances.)

In the neural processing system of bats, it has been suggested that the directional emission not only increases resolution and stabilises intensities across the frontal sound

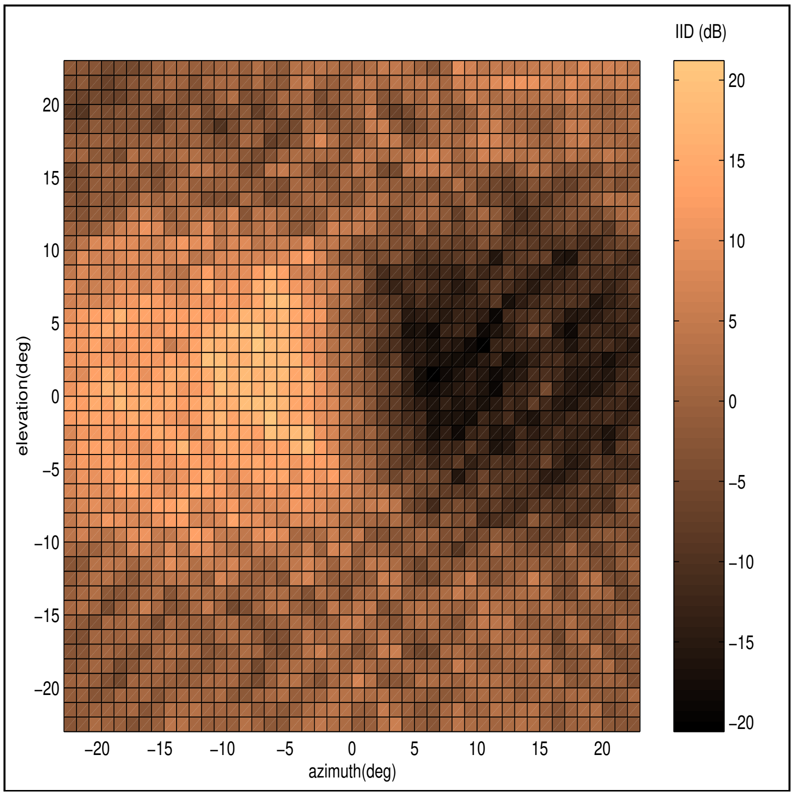

Figure 6.1:
Measured iso-frequency (50 kHz) IID map. Measured IID (*i.e.* intensity at the right receiver minus that at the left) values resulting from reflections off target located at indicated bearings in the frontal sound field ($r \approx 0.3$ m).

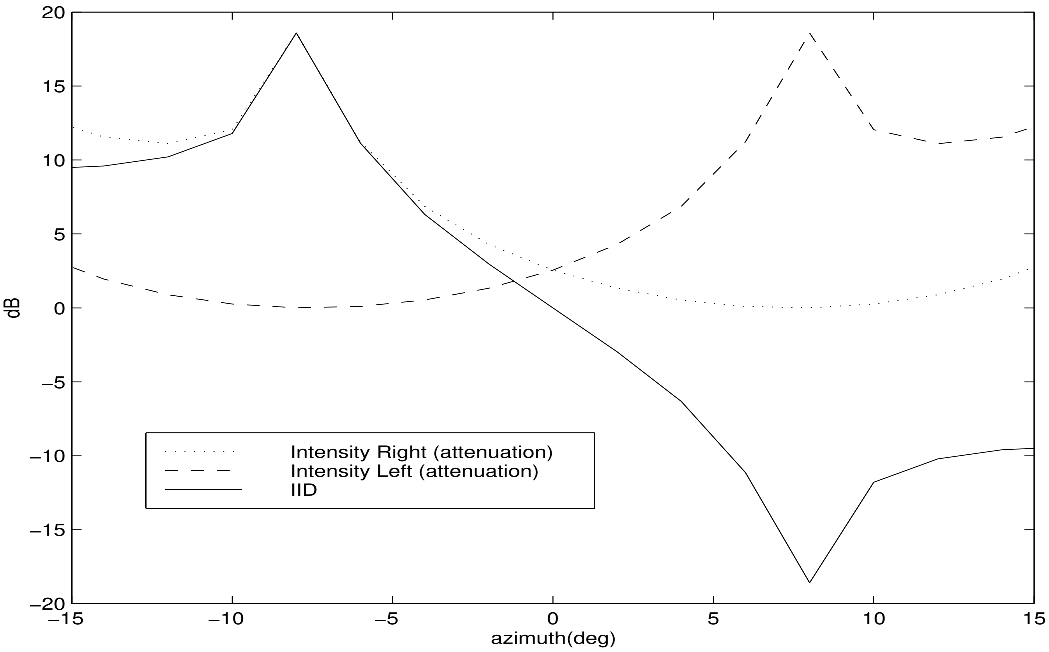

Figure 6.2:
Predicted IIDs and intensity attenuation (0° **elevation plane).** IID (right minus left) values across the sound field at 0° elevation. (Measured values agree well, but are less smooth and, therefore, illustrative.)

field, but may have the additional advantage of removing this potential medial/lateral confusion in single-harmonic localisation [Wenstrup 88]. In other words, although two different bearings give rise to the same IIDs, the absolute intensities measured at each bearing will not be the same. (See the IIDs *vs. intensity* contours of Figure 6.2.) In this way, intensity at the more lateral angle which produces the duplicate IID may be so attenuated that it is not much use in signal processing [Grinnell & Schnitzler 77, Fuzessery *et al.* 92]. Alternatively, an animal might take these both to be valid hypotheses, the truth of which can be resolved by listening along an additional viewing perspective(s).

As discussed in Chapter 5, broad-band emitting bats can create alternative measurement perspectives by sensing at different frequencies [Fuzessery 88, Pollak & Park 96]. Specifically, because bat pinnae are essentially composed of a series of acoustical magnifying glasses which point at different regions of space at different frequencies, it is argued that ambiguity in a target's bearing might be eliminated in the response accumulated across a series of iso-frequency IID maps. By contrast, the pinnae of rhinolophids and hipposiderids contain a single, fixed-width magnifying glass which has very fine resolution — directional sensitivity profiles at the dominant CF harmonic used in echolocation are amongst the highest reported for all bats (IID values up to $0.6^\circ dB^{-1}$) [Obrist *et al.* 93]. By actively moving these powerful magnifying glasses, high duty-cycle bats may collect measurements and integrate IIDs sampled from different regions of space in much the same way as broad-band emitting bats are believed to do. Thus, when viewed as a mechanism for the maintenance of maximal spatial sensitivity over a small area of the frontal auditory space, ear movements may provide these bats with versatility in maximising the spatial sensitivity of individual neurons during echolocation.

Mechanism

In this chapter, I investigate this possibility. The underlying philosophy of the experiment described below is the following. A tone can be localised (approximately) by comparing a measured IID value with a stored map of IID values. To remove ambiguities present in any one IID map comparison, the SONAR horizon can be reorientated, and the target localised again. Movement of the SONAR horizon can be achieved by tilting the head or, as discussed in Chapter 5, opposing motions of the pinnae/receivers might be employed to achieve similar effects. Note that this strategy does not require an echolocator have any particular directionality characteristics. Rather, the essential requirement is that the echolocator understand the mapping from IID to angular bearing at each SONAR horizon.

Figure 6.7 (a)-(e) (left column) show how the mapping of IIDs onto representative spatial regions changes as the SONAR horizon is rotated (in 6° increments) *via* receiver movements. Receiver motion causes the region with the steepest IID slope (*i.e.* the SONAR horizon) to tilt — increasing measurement resolution in each spatial region through which it moves.

Ambiguities due to noise and symmetries in the receiving system arise within each SONAR horizon, however they can be eliminated by combining results across alternative SONAR horizons.

6.2 Experiment

In the experiment reported here, an oscillating fan was placed at a series of known angular bearings covering the frontal sound field in 2° increments ($r \approx 0.3$ m). (The Target Fan is described in Chapter 4 and shown in Figure 6.3.) For each target bearing, the sensor head emitted a 50 kHz echolocation pulse and the receivers were stepped in opposing vertical arcs during echo reception. (See Figure 6.4.)

The results presented in Section 6.3 were obtained by stepping the receivers between extrema angles of $\pm 12^\circ$ in 6° increments. This angular excursion is used because the -20 dB measured beam-width of the Polaroid transducer (*i.e.* the level at which the first null in the directivity pattern begins to appear) is approximately 24°. Thus, for each target bearing, 5 IIDs were computed from the maximum intensity of the processed signals emerging from the 51.1 kHz channels of right and left filter-banks at each SONAR horizon. In this way, 5 IID maps were built up to describe the frontal sound field as heard from each SONAR horizon. The number of maps necessary to achieve good localisation across the frontal sound field was determined empirically (as discussed below). Less symmetrical receivers would require fewer listening perspectives.

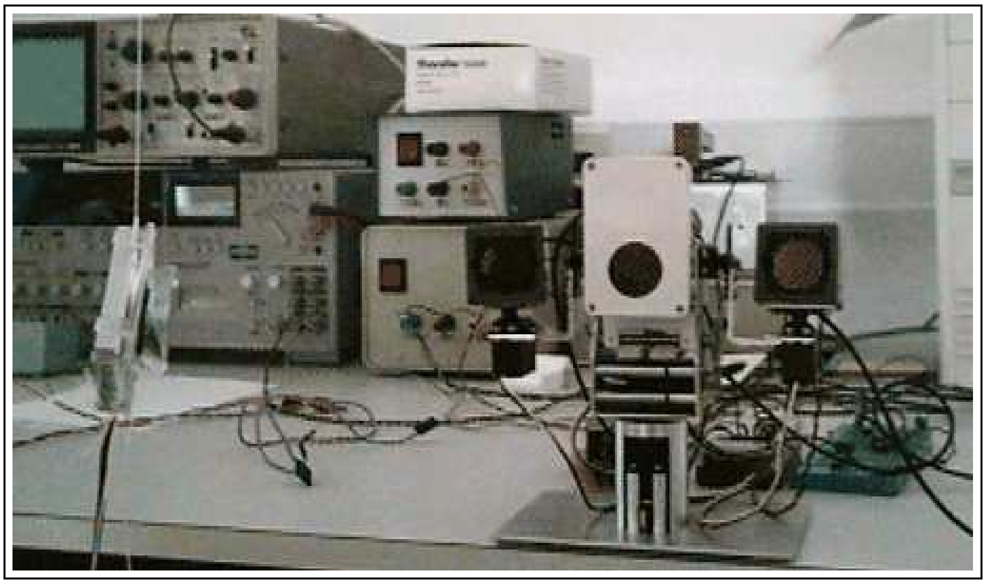

Figure 6.3:
Bionic SONAR system and fan target. The robotic sensor head and the fan target, fixed to the ceiling and floor by wires but otherwise free to rotate.

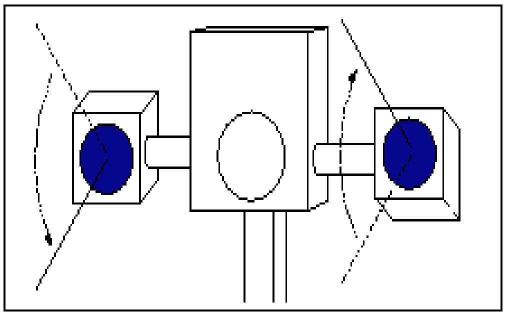

Figure 6.4: **Receiver motion.** Receivers move through opposing vertical arcs. (No pan motion introduced.)

It is possible that as few as 2 different SONAR horizons — *e.g.* measured at the start and end points of the pinna scan — could achieve the same effect as I demonstrate here.

6.3 Results

Figures 6.5 shows clearly how vertical arc scanning movements of the receivers tilt the SONAR horizon about the absolute horizon. With the receivers at $\pm 12°$, the SONAR horizon — drawn as a straight line through the sensitivity peaks of each receiver — is tilted through more than $45°$. This orientation is specific to a particular target range in a binaural system with large inter-aural distances: as the inter-aural separation increases, the effective orientation of the acoustic axes is increasingly determined by the inter-aural spacing rather than transducer orientation.

For example, considering only the azimuth orientation of the acoustic axes, the effective θ_{azi} (for receivers mounted in the same elevation plane) is given by:

$$\theta_{azi} = atan(\tan(\theta) + \frac{d}{r\cos(\theta)})$$

where θ is the bearing of a target with respect to the centre of the coordinate system (located at the emitter). Figure 6.6 shows the effect of range on IIDs for the bionic sensor head used in this investigation. For a small head, the second term in this equation drops out and θ_{azi} approaches θ (aside from any explicitly introduced, range-invariant receiver pan-angle). Thus, while widely spaced receivers can measure inter-aural timing differences more reliably, the measurement of inter-aural intensity appears to be simpler aboard a head with small inter-aural separations (wherein the orientation of the acoustic axis depends more upon the orientation of the receivers than their spacing).

In this and the following chapters, I work with targets over a limited range interval ($r \approx$ 0.3-0.5 m.) This range interval is used because it is just beyond the 2 d (where d is the inter-aural separation) boundary where IIDs stabilise and is also close enough to the sensor to provide good resolution ($< 0.5° dB^{-1}$, as shown in Figure 6.2). This resolution also corresponds well with IID resolution measured in rhinolophids and

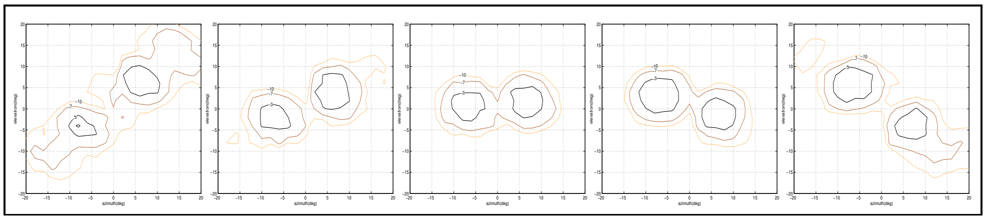

(a) (b) (c) (d) (e)

Figure 6.5:
Measured binaural directionality plots. Rotation of the SONAR horizon due to movement of the receivers through opposing motions of (right,left) receivers: (a) $(+12°, -12°)$, (b) $(+6°, -6°)$, (c) $(0°, 0°)$, and (d) $(-6°, +6°)$, (e)$(-12°, +12°)$. Intensity values are measured from reflections off a target located at indicated bearings ($r \approx 0.3$ m) in the ipsilateral frontal sound field of each receiver.

hipposiderids ($0.6° dB^{-1}$) [Obrist *et al.* 93].

The following demonstration provides an example of how the bearing of targets (located at $r \approx 0.3$ m in this case) can be determined by sensing along the 5 SONAR horizons in Figure 6.5 and comparing measurements with those stored in the corresponding IID maps shown in Figure 6.7 (a)-(e) (left column).

The middle and right columns of Figure 6.7 demonstrate how reorientating the receivers during echo reception can break the symmetry inherent in the binaural receiving system and overcome limitations due to noise and measurement resolution. In the construction of this demonstration, a test target was placed at an arbitrarily chosen bearing in the frontal sound field — *e.g.* (14° azimuth, 14° elevation). The receivers were then rotated through each of the 5 orientations and, in each receiver orientation, a new IID measurement was collected. A comparison was then made between each new IID value and those stored in the corresponding memorised IID map. Figure 6.7 (a)-(e) (middle column) shows the results of that comparison as a locus of possible target bearings whose characteristic IID values best match that of the test target.

Notice that when the SONAR horizon runs nearly perpendicular to the target (as in Figure 6.7 (a) and (b) (middle column)), a small IID is generated which makes the target appear to be located between the receivers and/or in the acoustic "blind" spot directly above or below the emitter. In these receiver orientations, there is little measurement sensitivity in the area where the target lies and, due to noise, the locus

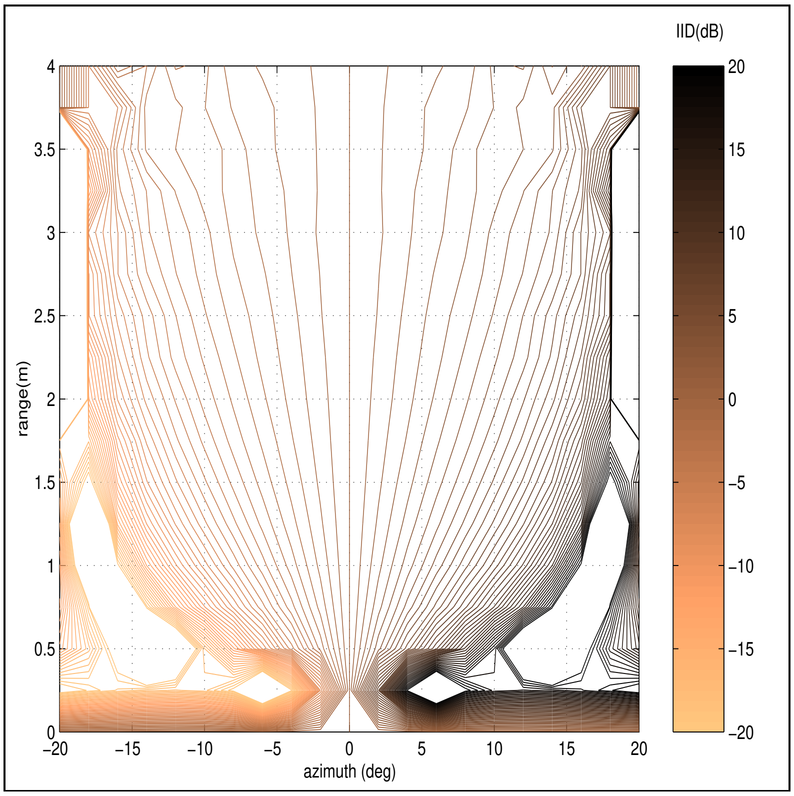

Figure 6.6:
Predicted IIDs as a function of range. Iso-IID lines (calculated for targets at 0° elevation) shows the effect of range on IID. (Sidelobes which give rise to the lateral patterns of sound create dramatic IIDs, but their overall signal levels are very low.)

of possible target bearings is wide and messy. The cluster of possibilities tightens as the SONAR horizon is rotated toward the target. Around 0° (Figure 6.7 (c) (middle column)), the target generates a ring of possibilities in which the azimuth bearing angle of the target is perceived to be in the right portion of the sound field. When the horizon starts to align itself with the target bearing (*i.e.* Figure 6.7 (d) and (e) (middle column)), target localisation is good and a small cluster of possible target bearing angles is generated.

The right column of Figure 6.7 (a)-(e) show how ambiguities may be reduced across a series of measurements by combining hypotheses generated at each SONAR horizon orientation. Each succeeding map (from (a) to (e)) takes the hypotheses of the previous map and combines them with its own by simply adding the results (*i.e.* tallying the votes at each SONAR horizon orientation). Shown are the target bearings receiving votes within 95% of the best hypothesis — except the last map (e), which shows the single hypothesis receiving the most votes. Notice how, in the first three movements, the symmetries are broken and the final movements serve to remove uncertainty due to noise and resolution limitations.

Across the whole frontal sound field (±20° azimuth, ±20° elevation), the average azimuth error is 6°, while the average elevation error is 4°. (The difference in accuracy arises because the directionality is naturally sharper in the vertical direction due to the common elevation mounting of all transducers.) Localisation accuracy varies with target bearing because sensitivity is highest in regions through which the SONAR horizon is stepped. For targets lying along a SONAR horizon, the average azimuth error is 2° and the average elevation error is less than 1°. By contrast, accuracy is low along the vertical mid-line (average error is 6° in both azimuth and elevation) because the system confuses low IID values generated along the mid-line between the receivers with those arising due to low absolute signal values in the periphery (blind spots). (Figure 6.8 shows the errors across the frontal sound field.) In this case, performance could be improved by using the biological signal processing strategy of interpreting IID values in the context of the absolute echo intensity.

Also, the use of additional measurement orientations/maps would improve results by increasing the likelihood of the SONAR horizon landing on the target. In simulation

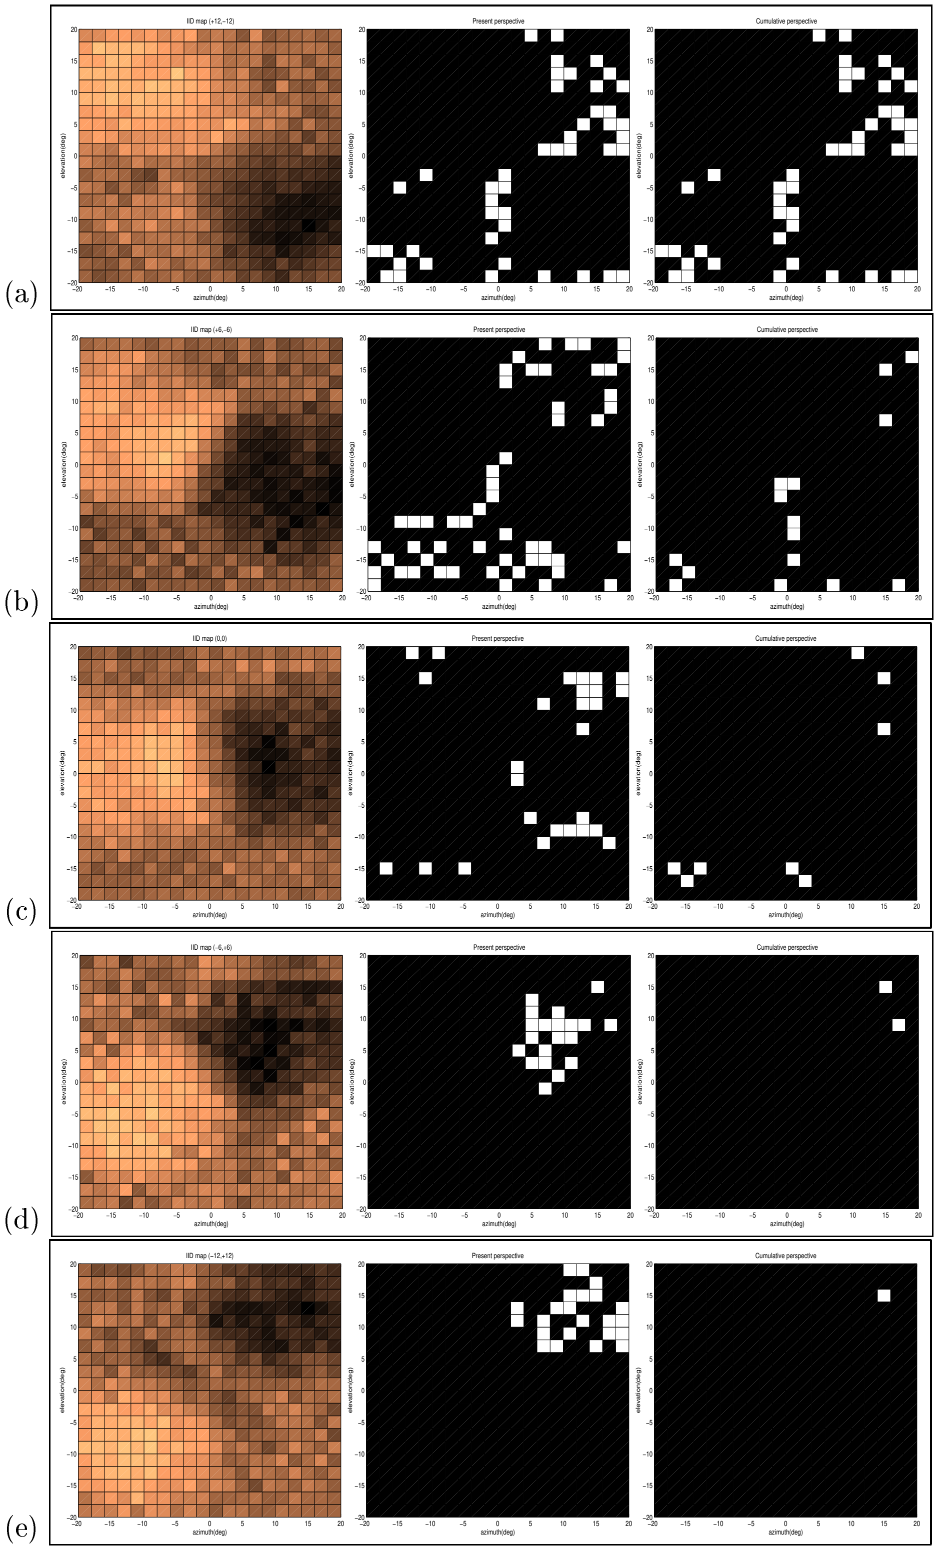

Figure 6.7:
Measured IID maps and target resolution. Left column: measured IID values across the frontal sound field for receivers held at angles (right,left) (a) $(+12°, -12°)$, (b) $(+6°, -6°)$, (c) $(0°, 0°)$, (d) $(-6°, +6°)$, (e) $(-12°, +12°)$. (Colourbar given in Figure 6.1.) Middle column: target bearing estimates shown as cells whose IID values match (*i.e.* deviate by less than 5% from) the measured IID value. True target bearing is $(14°, 14°)$). Right column: bearing resolution through cumulative tallying of votes in the middle column. (a)-(d) Shown are bearings whose IID values deviate by less than 5% from that of the measured value. (e) Bearing with the overall single best match between measured and map IID values.

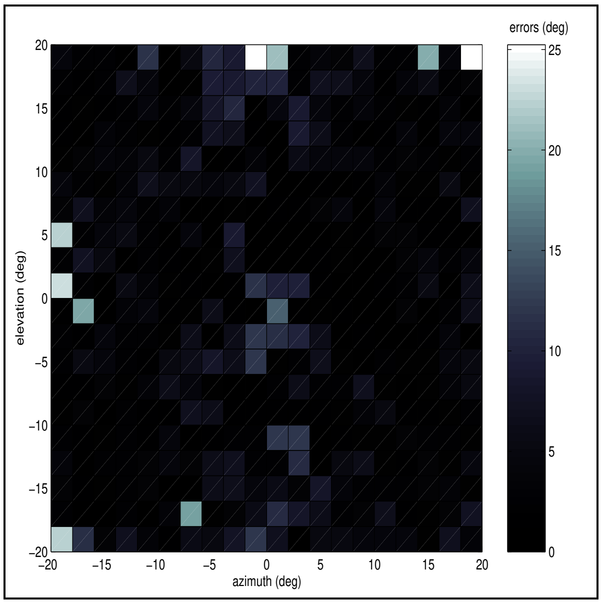

Figure 6.8: **Localisation accuracy.** Using a single ear sweep (5 static measurement orientations), target bearings were determined using the strategy depicted in Figure 6.7. Angular accuracy across the centre of the frontal sound field is shown as the root mean squared angular distance (degrees) between the true and measured bearing (across a grid corresponding to true bearing). Error values are based on a comparison of 2 sets of measurements.

work performed with Herbert Peremans [Peremans *et al.* 98a], the number of maps was manipulated and a probabilistic comparison scheme was implemented so that a reduction in entropy could be calculated during each stage of the map integration process. In this work, a probability map $P_N(A_m)$ was defined as a two dimensional array containing N cells where the value of each cell represented the probability that a target at the corresponding position has caused a particular measured IID value A_m. If the value of each cell in the probability map is denoted by $P(cell_i|A_m)$ for $i \in [1 \cdots N]$ and with A_m denoting the measured IID value, then, assuming only one target present, the cells in the probability map have to comply with:

$$\sum_{i=1}^{N} P(cell_i|A_m) = 1.$$

Employing Bayes rule, the posterior probability of a particular cell after n IID values have been measured $P(cell_i|A_m^n, A_m^{n-1}, \cdots, A_m^1)$ can be expressed as

$$\frac{P(A_m^n, A_m^{n-1}, \cdots, A_m^1|cell_i)P(cell_i)}{P(A_m^n, A_m^{n-1}, \cdots, A_m^1)}.$$

Assuming the different IID measurements to be conditionally independent, *i.e.*

$$P(A_m^n|A_m^{n-1}, \cdots, A_m^1, cell_i) = P(A_m^n|cell_i),$$

the posterior probability can be rewritten as:

$$\frac{P(A_m^n|cell_i)P(cell_i|A_m^{n-1}, \cdots, A_m^1)}{P(A_m^n|A_m^{n-1}, \cdots, A_m^1)}.$$

The term in the denominator can be seen as a normalisation term which can be implicitly calculated by normalising the posterior probability map. The first term in the numerator represents the measurement model. The posterior probability map is robust with respect to the actual measurement model chosen. An additive, zero mean, Gaussian noise was used in [Peremans *et al.* 98a]:

$$P(A_m|cell_i) = \frac{1}{\sqrt{2\pi}\sigma} \exp^{\frac{-(A_m - A_i)^2}{2\sigma^2}},$$

where σ was determined from noise measurements with the real system and the expected value A_i derived from averaging over a set of measured IID maps. However, changing the value of σ or employing a different distribution (*e.g.* the triangular of the example in Figure 6.7 or a rectangular distribution) does not alter the results significantly.

The second term in the numerator (*i.e.* the prior probability map) represents the evidence accumulated through the previous $n-1$ measurements. At startup this prior probability map is set equal to a uniform distribution $P(cell_i) = 1/N$. (If additional information about the target distribution is known, a maximum entropy approach could be used to select a more appropriate prior distribution.) Later, the posterior probability map resulting from integrating the previous measurement is used as the prior probability map for the integration of the new measurement.

The amount of ambiguity left in a particular probability map can be characterised by its entropy:

$$H(P_N(A_m)) = -\sum_{i=1}^{N} P(cell_i|A_m)\log(P(cell_i|A_m)).$$

Using five maps, the entropy decreased steadily across the comparison — suggesting that more maps could yield better accuracy. Doubling the number of maps (*i.e.* moving through receiver orientations in $3°$ increments) drove the average error to below $1°$ across the frontal sound field and nearly halved the average entropy (doubled the confidence) in the angular measurements.

As stated earlier, less regular receivers would require fewer SONAR horizons in order to resolve ambiguities in target bearing. Alternatively, sensing at just two different receiver orientations can give good results with the symmetric Polaroid transducer if several iso-frequency IID maps are employed at each SONAR horizon. Figure 6.9 shows a typical example of how the entropy and thus the position ambiguity decreases as more IID maps are combined in this way.

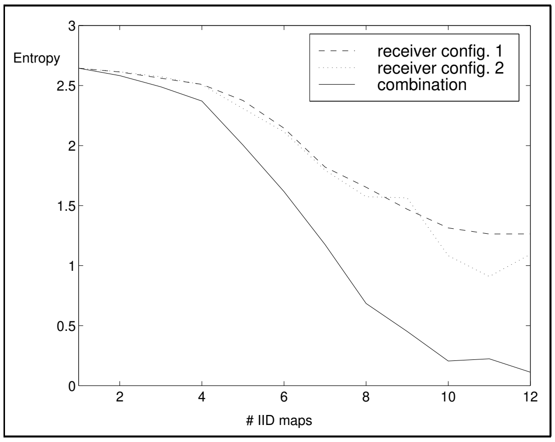

Figure 6.9: **Decrease in entropy as multiple frequency IIDs are combined with multiple SONAR horizons.** Receivers were moved into configuration 1 $(+5, -5)$ and IID measurements taken during the reception of a broadband $(30 - 90$ kHz) chirp. (IIDs extracted at 6 frequencies.) This procedure is repeated in receiver configuration 2 $(-5, +5)$ and the results combined.

6.4 Discussion

In this chapter, I investigated the hypothesis that a high duty-cycle echolocator could employ the low duty-cycle bat strategy of sampling discrete regions of space within the beam of a single echolocation pulse. Due to the fact that the acoustic axes of the ears point in different directions at different frequencies, broad-band emitting bats may sample different regions of space by altering the operating frequency of their auditory systems. By contrast, rhinolophids and hipposiderids hold that parameter largely constant and may achieve the same overall effect by actively altering the orientation of their receivers. I showed that for a binaural echolocation system with similar IID resolution to that of high duty-cycle bats, combining IID values from several viewing perspectives can be used to disambiguate target bearing — without recourse to additional coding mechanisms.

In trying to understand the representations of space which exist in the nervous system, it is tempting to look for high degrees of isomorphism between large regions of the environment and associated neural "maps". By contrast, Wehner argues that "neural maps are often distorted in the extreme, and often pay attention only to certain aspects of the outside world" [Wehner 87]. This work provides an illustration of the use of such *partial* maps which, through observer movement, can represent large areas of the behaviourally relevant frontal sound field. As stated by Wehner, neural maps used by autonomous agents must be understood in their dynamic context [Wehner 87]:

> *"The oddities of such mapping might give any die-hard inner screen theorist pause for thought. But they are less surprising when one realizes that the brain is designed to steer the animal's actions within the outside world rather than to portray this world as completely as possible. Sensory maps are not neural photographic images cast on some kind of inner neural screen, but devices shaped by particular selection pressures to pre-process sensory information in a way readily translatable into the necessary motor commands."*

The performance of the particular partial map-based localisation strategy implemented here varied with target bearing in the frontal sound field due to the fact that the SONAR horizon was moved in large steps and, therefore, only sampled a few regions of the frontal sound field. Increasing the number of maps is a cumbersome solution because a prototype IID map must be memorised for each SONAR horizon. (The precision of the motor control/proprioception systems may also limit the number of useful maps.) In the next chapter, I explore a more dynamic mechanism which involves a single tone echolocation call and continuously moving and measuring receivers.

7. Investigation Two: Temporal cues

In the nervous system of animals, various types of aurally sensitive neurons are created by combining inputs from different monaural and binaural lower auditory centres. In rhinolophids, auditory neurons sensitive to target motion have been observed [Schlegel 80, Neuweiler 80]. Although most superior olivary, lemniscal and collicular neurons faithfully encode the bearing of a target irrespective of target velocity, neurons which encode target motion by increasing discharge rates proportional to target speed and/or direction of motion have been observed. Some of these neurons only respond when the target sweeps through a narrow angular sector with a marked best angle. In other neurons, the representation of azimuth angle is lost because units extract only dynamic features.

If neurons exist to represent moving targets — as measured by stationary pinnae — then the apparent motion of a stationary target — as measured by moving pinnae — may be similarly encoded.

7.1 Hypothesis

Motivation

By scanning a target continuously with alternating pinnae movements, the amplitude modulation imposed on the envelope of a CF echo reflected from a target fixed in space changes continuously. "This change is specific for every angle of incidence of an echo and could, therefore, be used to determine target angle" [Schnitzler 73]. (Figure 7.1

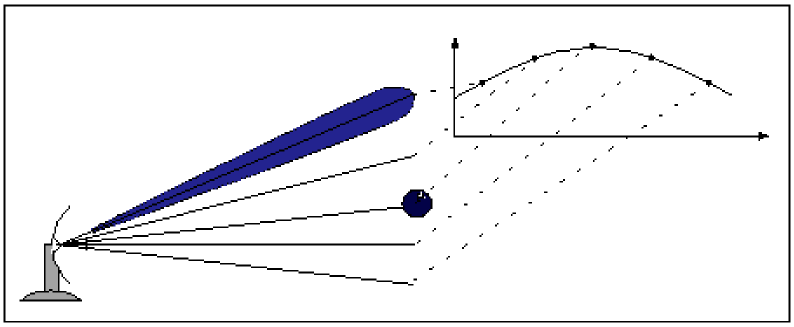

Figure 7.1: **Scanning SONAR.**

illustrates this idea.) The use of a largely vertical pinnae movements by rhinolophids and hipposiderids suggests that amplitude modulation due to apparent target motion would be likely to encode elevation, rather than azimuth. Indeed, there is some evidence to suggest that the ear movements play a greater role in localising a sound source in the vertical plane than in the horizontal one — *e.g.* in rhinolophids, vertical localisation performance is more significantly impaired by immobilization of the ears [Griffin 58, Gorlinsky & Konstantinov 78, Mogdans *et al.* 88].

Mechanism

In this chapter, both binaural and monaural elevation cues are extracted from echo envelopes modulated in this way. In the case of the latter, target elevation is encoded in the temporal disparity between echo arrival time and echo amplitude peak time (Section 7.3.2). In Section 7.3.1, the differences in echo amplitude peak times in the right and left receivers — which create *IID rates of change* [Schnitzler 73] — are employed to encode target elevation.

7.2 Experiment

As in Chapter 6, an oscillating reflecting target was placed at a series of bearing angles covering the frontal sound field in 2° increments. For each target bearing, the sensor emitted an echolocation pulse and the receivers were moved in opposing vertical arcs during echo reception. All measurements presented in Section 7.3 were collected during continuous vertical motion of the receivers which generated amplitude modulations in the returning echo envelopes. Elevation angle was derived from the modulation patterns (as described below) and azimuth is calculated independently from IIDs.

In Section 7.3.2, a procedure is described that extracts the time at which peaks occur in the amplitude modulated echoes reflected from the Fan Target located at different elevations. Peak delay times are measured relative to the arrival time of echoes in the 51.1 kHz channel of the right and left filter-banks. In this (and the following) strategy, IIDs are used to determine target azimuth. Measuring IIDs unambiguously in a map-less scenario requires that targets remain within the region of the frontal sound field wherein IIDs vary monotonically (*i.e.* $\pm 8°$ for targets at $r \approx 0.3$ m, as shown in Figure 6.2). Therefore, in measurements taken with the bionic head, targets were restricted to a 16° (azimuth) by 16° (elevation) region of the frontal sound field (and $r \approx 0.3$ m).

In Section 7.3.1, the continuous differences between right and left echo envelopes (rates of change of IIDs) are computed. Here, simulation results are presented along-side measured results. In the case of the former, the simulated ears were snapped into opposite extreme orientations (right pinna up, left pinna down) at the time of emission, and then moved — with a sinusoidal velocity profile — through opposing vertical arcs of 30° centred on the horizontal mid-line. This particular angular excursion is used because the beam-width of the main sensitivity lobe of the piston pinna model (whose dimensions (a = 4.5 mm) and frequency (83 kHz) are based on those typical of *R. ferrumequinum*) is approximately 30°. (This value corresponds well with peak to peak scan amplitudes measured at the tip of the ear of *R. ferrumequinum* which can exceed 1 cm — describing an arc of about 30° [Pye & Roberts 70].)

7.3 Results

7.3.1 An IID slope elevation sensor

Figure 7.2 illustrates the *change in IID* perceived by scanning a pair of simulated pinnae continuously during echo reception. When a target is located at a negative elevation, the right ear hears a weak reflection at the beginning of the reception interval and sound intensity increases smoothly as it is scanned downwards (*i.e.* toward the

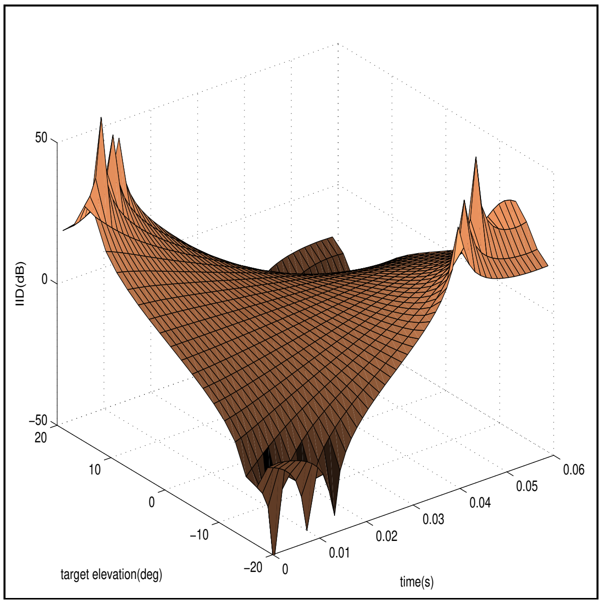

Figure 7.2:
Predicted IIDs as a function of time and target elevation. Simulated time history of IID values ("IID slopes") for ears the size of *R. ferrumequinum* moving through 30° elevation arcs. During each trial, the simulated target was located within the 0° azimuth plane at the indicated elevation. (Doppler effect created by the movement of the simulated was eliminated.)

target). Because the opposite trend is present in the left ear, the IID (right minus left) increases over time. The IID profile is steepest when the target is located in an extreme elevation location (±15°) and, as the target is brought toward the horizontal mid-line (around which the artificial pinnae scan), response peaks in each ear overlap increasingly — thereby flattening the slope. At 0° elevation, IIDs are constant across the duration of the measurement.

Monotonicity of the IID profile is broken if the target moves outside the central elevation region — as shown for target bearings $> |15|°$ elevation. Here the roughening of the slope is exaggerated by side-lobes in the piston model's polar sensitivity. Similar effects would be encountered if the target moved to the azimuthal periphery, or if the ears are driven to extreme orientations such that the main lobes of their sensitivity beams no longer overlap.

From the available sense-information, azimuth might be derived from the average IID, or "IID offset". For example, the family of slopes representing different target elevations is centred on 0 dB for targets at 0° azimuth. At other target azimuth bearings, the family of intersecting lines representing each elevation is centred at a different IID offset. Figure 7.3 shows two examples. Similar results were obtained aboard the bionic system, although the larger echolocation head operated over a more narrow region of the frontal sound field. Example IID slopes and offsets measured by the bionic system are shown in Figure 7.4. The signals shown there were extracted from the 51.1 kHz channel of the filter-bank described in Chapter 4. (In order to smooth the signal fluctuations due to amplitude and frequency modulation introduced by the target, the final stage of lowpass filtering was replaced by a narrow bandpass filter centred at the target's rotation rate.)

By rotating the acoustic axes outwards (and holding them there throughout the sweep), an echolocator can change the amount of IID offset and thereby improve target azimuth resolution. In the simulation reported here, the acoustic axes of the right and left transducers are rotated outward from the middle of the sound field by $\pm 15°$ azimuth into their respective ipsilateral sound fields (as in Figure 4.1). This yields an offset of approximately $0.7° dB^{-1}$ — see Figure 7.3. By contrast, in preliminary studies with the acoustic axes pointing directly forward, the simulator predicted that pinnae the size of *R. ferrumequinum* would only afford a $2.5° dB^{-1}$ resolution.

In the IID slope comparison scheme, the straightforward use of a single slope value to represent elevation independently of azimuth is possible in the middle of the frontal sound field. This assumption breaks down, however, as the target moves laterally such that response in one receiver predominantly determines the IID values throughout the scan. As shown in Figure 7.3, at 5° azimuth, IID slopes are similar to those for targets at 0° azimuth. Based on a first order approximation, they only deviate from those at 0° by 9%; however, when a target moves out to an azimuth angle of 10° (not shown) deviations increase to 30%.

It is possible to construct a calibration factor to correct for this systematic variation, but a simplier strategy might be to decouple azimuth and elevation sensing by employing IIDs for the former and temporal cues for the latter.

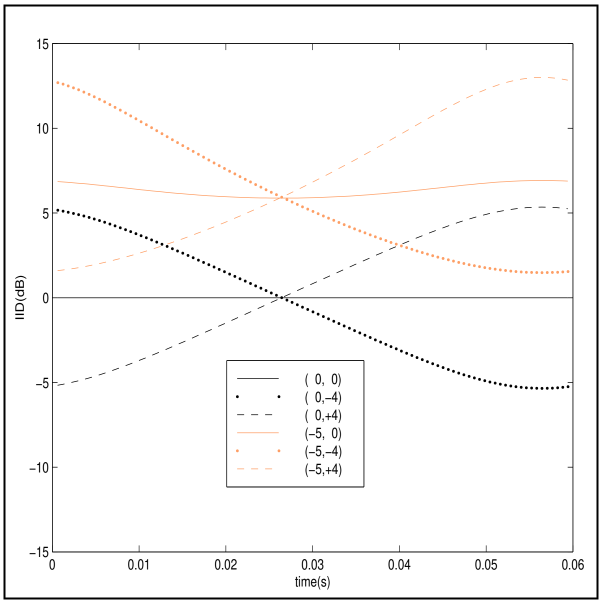

Figure 7.3:
Predicted IID profiles as a function of target bearing. Simulated IID slopes and offsets for ears the size of *R. ferrumequinum* moving through 30° elevation arcs are shown. The azimuth orientation of the acoustic axes (−15° (right), +15° (left)) is held constant throughout the scan. Target azimuth and elevation is indicated.

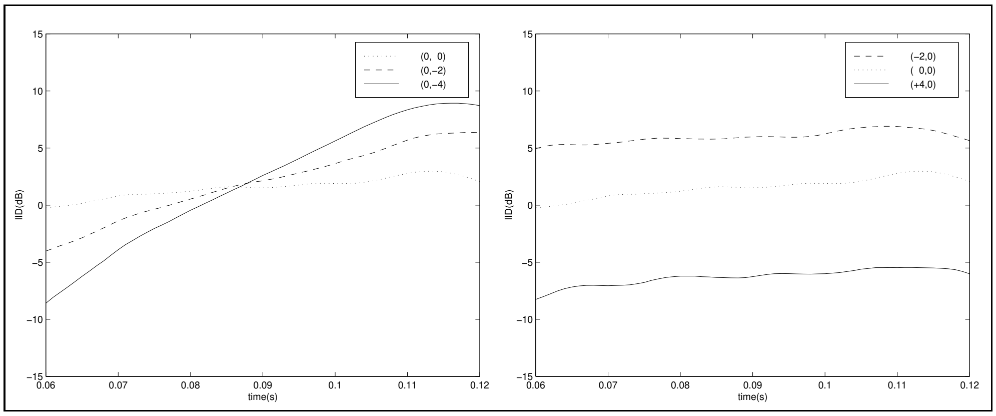

(a) (b)

Figure 7.4:
Measured IID profiles as a function of target bearing. (a) IID slopes for the bionic system. (b) IID offsets for the bionic system. Right and left receivers were scanned through 16° elevation arcs. In both panels, target azimuth and elevation are indicated ($r \approx 0.3$ m). (During scanning, no appreciable Doppler effect is induced by the receivers.)

7.3.2 Monaural peak delay

IID slopes arise because echo amplitudes peak at different times in right and left receivers moving through opposing vertical arcs. This fact suggests that an echolocator with mobile pinnae might represent target elevation *via* inter-aural (peak) timing differences (ITD_{peaks}). Alternatively, this information can be derived from monaural timing cues — *i.e.* the delay between echo arrival time and echo peak time in each ear. Although the peak echo intensity *value* reflected from a target at any particular elevation will vary with target azimuth, the *peak delay* time depends predominantly on the speed of the pinnae and the vertical angle between target and pinnae start orientation. The *difference* in monaural peak delay values encoding adjacent target elevation angles is determined by the motion of an acoustic axis and, therefore, can be understood simply as a delay-per-degree transformation (if the velocity of the receivers is sufficiently linear).

Figure 7.5 shows this transformation as the delay profile for right and left receivers (averaged across different target azimuth angles). A straight line approximation to the delay lines yields a 1° per 3 ms transformation — *i.e.* a 3 ms resolution encoding adjacent target elevation angles. This figure is based on receivers scanning over 16° in 60 ms and suggests angular resolution greater than that available to rhinolophids and hipposiderids scanning their pinnae over twice this angular extent in less than half of this time.) However, an important aspect of this mechanism is that scan rates and, therefore, elevation resolution can be manipulated explicitly.

Along the contours of Figure 7.5, the average separation between measured peak delays values encoding adjacent target elevations is 6.0 ms (per 2°). As the average variation of each measurement is nearly half of this (*i.e.* 3.3 ms, as shown *via* the error bars), it is possible to determine a target's elevation position with an accuracy of within ±2° elevation. (An obvious way to combine the monaural elevation estimates is in proportion to signal intensity, so that angular estimates are principally determined by the transducer receiving the louder, more reliable information.)

Figure 7.6 shows the variation in IID values with target azimuth (across all elevations). As IIDs change continuously throughout the scan, we adopted the simple approach of

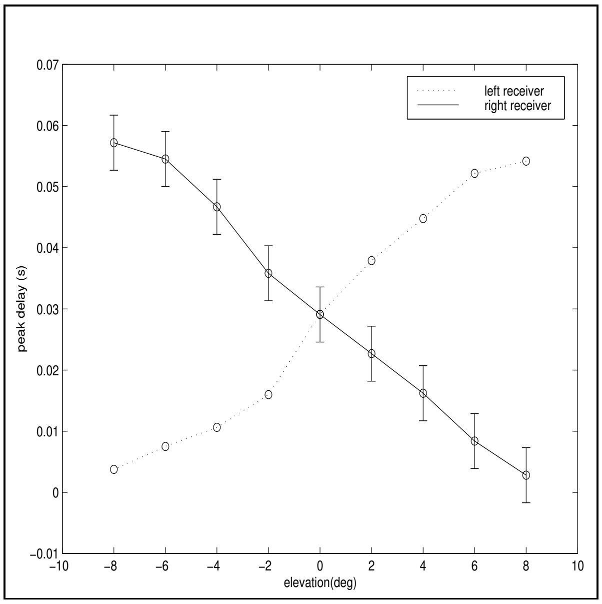

Figure 7.5:
Measured average peak delay. Peak delay in each receiver as a function of target elevation (*i.e.* values averaged across all target azimuth bearing angles). Error bars (only shown for the right receiver) indicate the average variation in each elevation measurement.

comparing the peak intensity values in each receiver. Values encoding each azimuth are averaged across all elevations, and the errorbars indicate the average variation in each measurement due to noise and differences in target elevation. As shown, this use of IIDs to determine target azimuth position independently of elevation has an accuracy of approximately $\pm 2°$ azimuth across the centre of the frontal sound field where IIDs vary monotonically. Naturally, angular resolution is highest in the front of the sound field where IIDs change most rapidly.

7.4 Discussion

Just as receiver size and shape create dramatic binaural intensity differences, receiver mobility may allow some echolocators to derive dramatic temporal cues for localisation — *i.e.* cues dependent upon pinna speed (millisecond cues) rather than inter-aural distance (microsecond cues).

Here I investigated the creation and use of various temporal cues. Across a limited

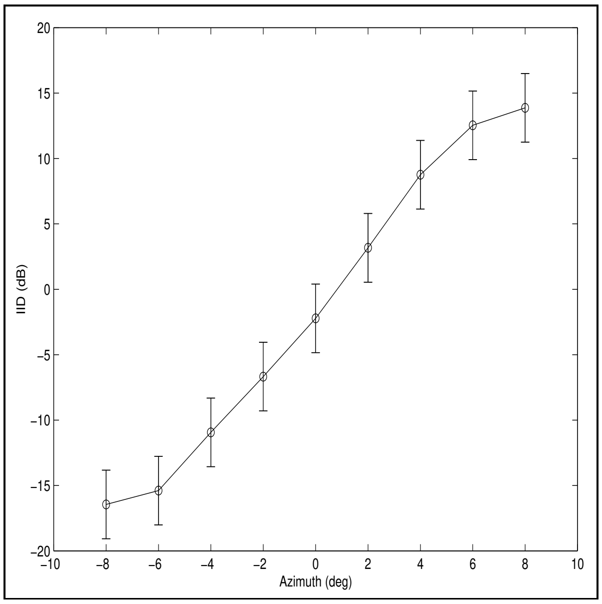

Figure 7.6: **Average IID.** Measured IID as a function of target azimuth (*i.e.* values averaged across all target elevation positions). Error bars indicate the average variance in each azimuth measurement.

region of the frontal sound field, accurate elevation measurements can be made by sweeping a pair of receivers through opposing elevation arcs and measuring the time of the peak in the amplitude modulated echo. Unlike the mechanism investigated in Chapter 6, angular errors resulting from this use of movement are consistent across the frontal sound field. Using a decoupled strategy of measuring azimuth *via* IIDs and elevation *via* monaural peak delays (with inter-aural confidence checks where possible) can provide unambiguous, high resolution 3D angular cues across a limited region of the frontal sound field. However, readings are reliable over a more limited region because targets must remain within both (i) the azimuth region where IIDs are monotonic and (ii) the elevation region between the extreme angles of the vertical scan. Outside of this region, targets can still be detected, and behavioural mechanisms may be used to bring the target into the optimised environment of the front-centre sound field for analysis [Schnitzler & Henson 80].

8. Investigation Three: Frequency cues

The relationship between pinna features and echolocation calls is striking in bats which employ pinna movements. In rhinolophids and hipposiderids the size of the ears is large with respect to the wavelength of echolocation calls (pinna mouth can be up to 3 times larger than the wavelength of the CF component of their call [Guppy & Coles 88, Obrist *et al.* 93]). The use of such high frequency harmonics for sensing does not optimise the pressure gain in the ear (predicted for a system operating closer to $ka = 1$). However, it will improve spatial resolution and target resolution (by lowering the beam-width and wavelength, respectively). Another (complementary) explanation for the exaggerated size of the ears with respect to the wavelength of the call is that use of large, mobile pinnae and high call frequencies may exaggerate Doppler cues induced by the pinnae themselves (assuming that reflections occur off the extremities of the ear). In other words, given that pinna speed will be limited by the response of motor neurons and the duration of the call (for one-to-one correlated pulses/echoes and ear movements), the evolution of large pinnae and high call frequencies provides a means for maximising Doppler cues.

8.1 Hypothesis

Motivation

The acoustic fovea of high duty-cycle bats is particularly well suited to fine frequency analysis and Doppler shifts induced by moving pinnae could be detected (although this has never been directly tested). The ability to introduce spectral broadenings *via* ear movements may facilitate the frequency domain separation of call and echo — which could be particularly important when the bat and/or target is stationary [Pye *et al.* 62]. Moreover, if bats understand how to exploit the cosine-law dependence of Doppler cues on target angle, pinna movements could be used in target localisation [Pye *et al.* 62].

Pinna speeds in rhinolophids and hipposiderids vary with circumstance. As discussed in Chapter 5, movements appear to be correlated with echolocation pulses which are emitted in groups consisting of $1 - 2$ pulses rising to approximately 20 in unfamiliar conditions. In the latter case, pulse lengths may drop to as little as 10 ms. The measured rate of pinna rotation along their opposing vertical arcs can rise to 80 Hz in *R. ferrumequinum* [Pye *et al.* 62]. This figure agrees with the independently measured speed of contraction of the superficial cervicoauricularis muscle responsible for moving the pinnae: tetanic summation was absent for muscles stimulated by 1 ms square waves up to 25 Hz, but develops progressively above 30 Hz. Pye *et al.* report that considerable amplitude of vibratory ear movement is present at 70 Hz, with the alternating component disappearing completely at about 100 Hz.

It is intriguing that the pinna movement behaviour operates over a range of pinna speeds. We saw in Chapter 7 how peak delay measurements taken by slower moving pinnae provided greater resolution (in $^{\circ}ms^{-1}$), whereas the measurement of angle *via* Doppler would clearly be facilitated by pinnae moving at a higher velocity (so as to create larger Doppler shifts). Perhaps the available cues are weighted by their magnitude such that the bat relies increasingly on Doppler cues as the pulse duration decreases and pinna speed increases. In this chapter, I fix the pulse length around its minimum which — for correlated ear movements and pulses — fixes the pinna rate at 50 Hz.

In the case of short calls, reflection of complete wing-beat cycles by fluttering targets is unlikely; however, it is possible that the Doppler shifts due to insect wing-beat and those due to pinna movement might complement each other to encode a target description. One such possibility would exploit the fact that — as pinna speeds span much the same range as target wing-beats — the pinnae might be moved so as to effectively cancel (thereby creating a stroboscopic effect) or reinforce target induced Doppler. This chapter investigates the use of pinna movements to create Doppler cues for target localisation and identification.

The work presented in this short chapter is speculative in that there is relatively little known about the way in which bat pinna induce frequency (as opposed to intensity) cues. Here I use a very simple model of Doppler shift induction by a pair of simulated piston-like receivers mounted atop pinna stalks and explore the way in which frequency cues generated by stalk motion can be used target localisation (and detection) context. Due to hardware limitations (*i.e.* limitations in the rotational speed of receivers), it was not possible to test hypothesised localisation and detection mechanisms aboard the robotic sensor.

Mechanism

The motion of the ears during reception of an echo from a stationary target creates frequency shifts which encode target angle because the magnitude of the Doppler effect depends upon how much of the ear velocity is projected onto the line-of-sight between target and receiver. In the case of a bat moving with velocities v_b (emitter) $v_{b'}$ (receiver) relative to a target, which, itself may be moving with velocity v_t relative to the emitter ($v_{t'}$ relative to the receiver), the echo frequency f_e can be calculated from the call frequency f_c:

$$f_e(t) = f_c \frac{c + v_{b'}(t)}{c - v_b(t)} \frac{c - v_t(t)}{c + v_{t'}(t)} \tag{8.1}$$

In the case of rhinolophids and hipposiderids pursuing fluttering targets, relative velocities in Equation 8.1 may have both AC components (which arise from the bat's pinna and the insect's wing motions) and DC components (due to movements of the agents as a whole). In this chapter, I will consider only the AC velocity components of $v_{b'}$ and $v_{t/t'}$ (v_b, having only a DC component, drops outs entirely). Specifically, I will consider two sensing scenarios.

In Section 8.3.1, I look at directional cues available to a stationary bat moving its pinnae in the presence of a stationary, non-fluttering target. In this case, Equation 8.1

reduces to:

$$f_e(t) = f_c \frac{c + v_{B'}(t)\ cos\ \delta_{B'T}(t)}{c} \tag{8.2}$$

where $v_{B'}$ is the absolute velocity of the pinna as it moves through its vertical arc and $\delta_{B'T}(t)$ is the Doppler angle between the vector defining the time-varying pinna velocities and the lines-of-sight from pinna to target. Both quantities vary systematically with pinna position during the sweep.

In Section 8.3.2, target motion is introduced:

$$f_e(t) = f_c \frac{c + v_{B'}(t)\ cos\ \delta_{B'T}(t)}{c} \frac{c - v_T(t)\ cos\ \delta_{TB}(t)}{c + v_T(t)\ cos\ \delta_{TB'}(t)} \tag{8.3}$$

where $v_{B'}$ is defined as above and v_T is the absolute velocity of the insect's wing which is projected onto the line-of-sight with the emitter (through the angle δ_{TB}) and receiver ($\delta_{TB'}(t)$). Note that here the line-of-sight changes with variations in both the pinna and wing positions.

8.2 Experiment

A simulated reflecting target was placed at a series of elevations throughout the frontal sound field (r = 0.5 m) and insonified by a stationary echolocator (see Appendix C for a full description of the *3D Echolocation Simulator*). In all simulations, the pinnae (and transmitter) were modelled as circular pistons (a = 4.5 mm, corresponding to the radius of the pinna mouth of an average *R. ferrumequinum*). Receivers were mounted at the "tip" of the pinnae (here pinna length was taken to be twice the width [Obrist *et al.* 93]), as it was assumed that reflections from the extremity of the pinnae are diffusely reflected into the ear canal (and that the small delay introduced is insignificant). Pinnae move through opposing vertical arcs with sinusoidally varying velocity profiles throughout the duration of echo reception.

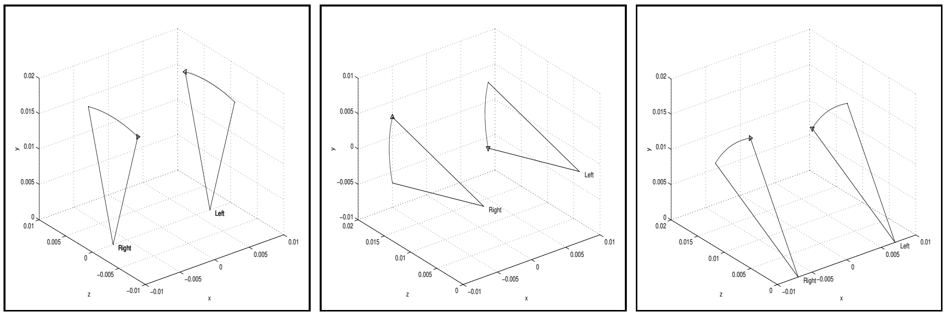

(a) (b) (c)

Figure 8.1:
Pinna position and angular excursion. (a) Rotation about 0°. (b) Rotation about 90°. (c) Rotation about 45°. (Arrows mark the end of the trajectory.)

8.3 Results

8.3.1 Stationary target

In this first example, the simulated pinnae were placed atop the notional bat head and swung through opposing arcs of 30° — see Figure 8.1 (a). The resulting motion of the right pinna is away from the frontal sound field and that of the left is into it, as indicated by the small arrows.

Figure 8.2 (a) and (b) show the variation in instantaneous echo frequency (over the duration of a 10 ms pulse) received in the right and left ears (respectively) for targets at elevations varying from ±90° (negative target elevations in black, positive in orange). For a target located at 0° elevation, the plot of the frequency deviations over time is symmetrical because the angle ($\delta_{B'T}$) between the vectors defining the line-of-sight from pinnae to target and pinna velocity is zero in the middle of the sweep. For targets lying above or below 0° elevation, the pinnae are no longer maximally aligned with the target as they move through the middle of their trajectory — which gives rise to asymmetries in the time waveforms. (Another way to look at this is that the relative velocity of the right/left pinna point into the negative/positive elevation regions of the frontal sound field at the beginning of its trajectory and this progressively delays the arrival of echoes from targets lying in the other half of the sound field.)

To investigate target angle resolveability, a detection strategy must be adopted. A

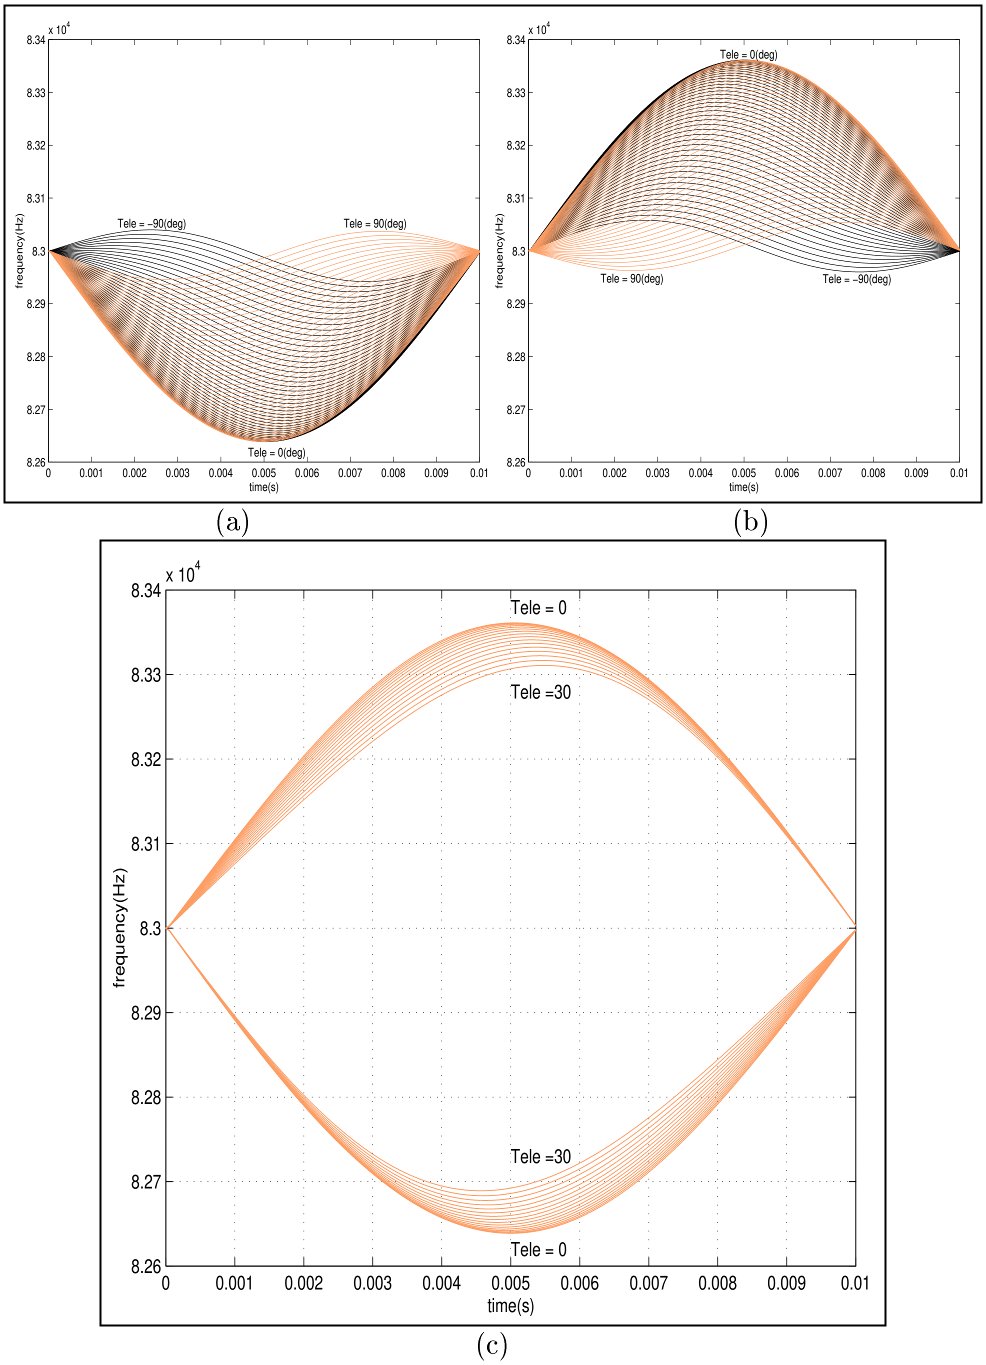

Figure 8.2:
Instantaneous frequency shift *vs.* **target elevation (pinnae moving overhead as in Figure 8.1 (a).** Echo frequency deviation over time for targets located at different elevations along the vertical mid-line ($r = 0.5$ m). (a) Right pinna. (b) Left pinna. (c) Both pinnae over a limited range of target elevation angles.

simple directional cue which might be extracted from Doppler shifted echoes is the frequency extreme. Figure 8.3 (a) shows the minimum and maximum frequencies of the right and left (respectively) echoes over all target elevations. Instantaneous frequency deviations rise to 450 Hz for targets located at 0°. As *R. ferrumequinum* is reported to distinguish frequency shifts of approximately $f = 30 - 50$ Hz [Schuller *et al.* 74], this simulation suggests that stationary targets could be detected (on the basis of Doppler shift) at all elevations. (Echo attenuation, not considered, would certainly reduce this detectability at the peripheries of the sound field.)

Another possible detection strategy might involve the use of inter-aural frequency differences (IFDs). Figure 8.3 (b) shows IFDs measured continuously over time. The IFDs change systematically with target angle due to the underlying cosine-law dependency of the Doppler effect. This dependency creates a situation wherein IFDs are greatest for targets in the middle of the sound field; however, resolution (in terms of $^\circ Hz^{-1}$) is worst in this region.

As shown in Figure 8.3 (b), around the centre of the sound field, IFDs drop 50 Hz (from their maximum at 0°) over 14°. (Figure 8.2 (c) shows the subtle variations in the Doppler effect across target elevations within the centre of the frontal sound field for left and right pinnae individually.) It is not clear what inter-aural frequency differences a bat might resolve; however, it would need to detect IFDs of < 3 Hz to resolve targets separated by 2° elevation. (Figures 8.4 shows how altering the pinna length and sweep velocity increase the absolute frequency deviations, but the differences in frequency which encode adjacent target elevation angles do not change.)

An echolocator can take a more clever approach to exploiting the cosine dependence of the Doppler effect. If there is high resolution when the pinna velocity is more orthogonal to the target position, perhaps the ears are rotated about an elevation offset angle which causes Doppler shifts to change more rapidly with target angle. Figure 8.1 (b) depicts an extreme variant of this where the pinnae are held forward (as opposed to upright) and swept through the front of the sound field. Figure 8.5 (a) and (b) show how right and left ears moving about this new offset angle experience both positive and negative Doppler shifts. Here, shifts fall towards zero when targets are centred at the mouth elevation (0°) such that a minimal frequency deviation would encode centred targets.

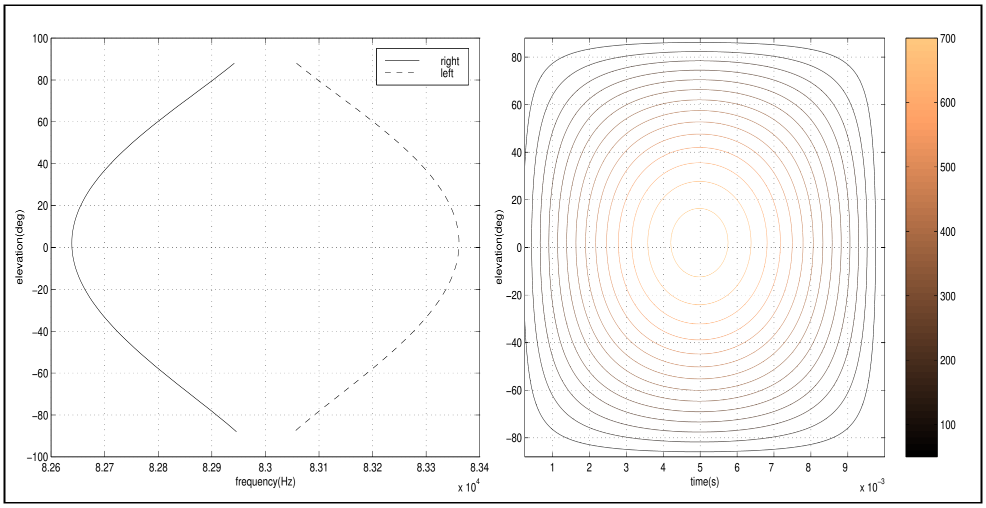

(a) (b)

Figure 8.3:
Doppler shift extrema and IFDs. (a) Maximum Doppler as a function of target elevation for right and left pinnae. (b) Inter-aural frequency differences (IFD) (left minus right) shown as a function of echo length and target elevation. Colourbars indicate IFDs (Hz). (IFD plot does not peak at 0° elevation due to the mounting of the pinnae *above* the head.)

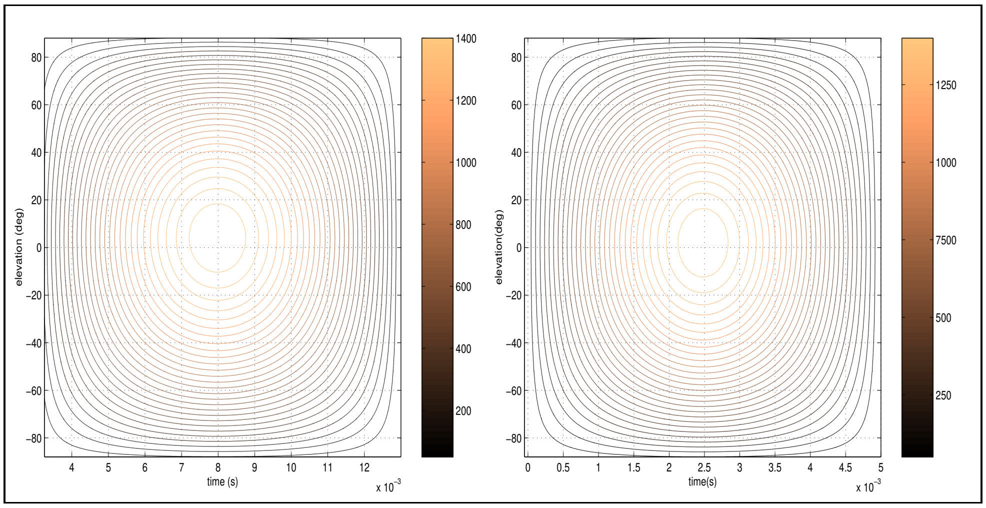

(a) (b)

Figure 8.4:
IFDs across varying ear parameters. (a) Pinna length doubled. (b) Sweep speed doubled (pulse length halved). In each simulation, all other parameters were held constant. (Colourbars indicate IFDs (Hz).)

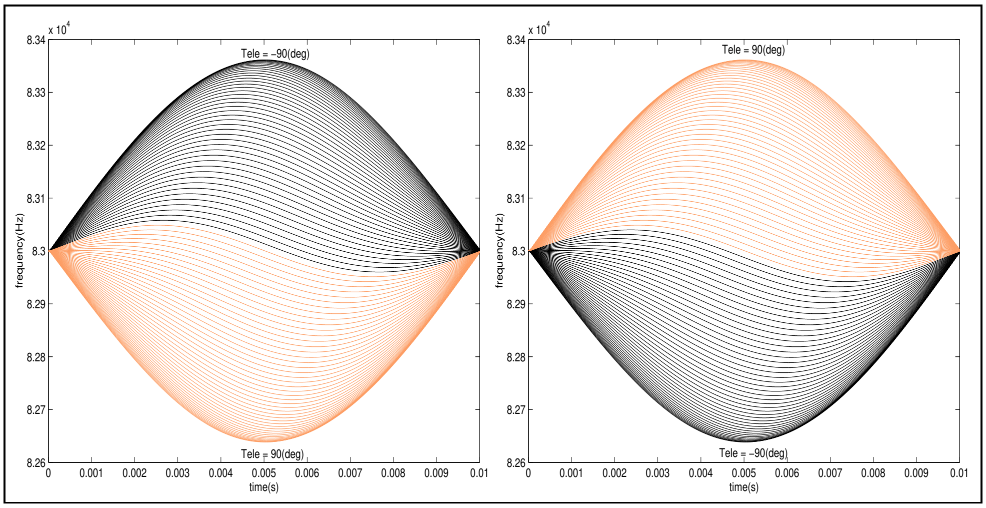

(a) (b)

Figure 8.5:
Instantaneous frequency shift *vs.* **target elevation (pinnae moving as in Figure 8.1 (b).** Echo frequency deviation over time for targets located at different elevations along the vertical mid-line ($r = 0.5$ m). (a) Right pinna. (b) Left pinna.

The IFDs corresponding to this situation are shown in Figure 8.6 (a). Targets centred at $0°$ elevation yield no IFD, while targets at just $4°$ off-axis give rise to a 50 Hz difference. This strategy is also appealing because the ambiguity in target elevation sign (due to IFD symmetry, see Figure 8.3) is removed.

Although illustrative, a $90°$ offset is unrealistic in that it is hard to imagine that pinnae could receive echoes in this configuration without introducing occlusion effects. A more realistic elevation offset — which would introduce many of the same benefits (*i.e.* symmetry resistance and high resolution) — would be $45°$, as shown in Figure 8.1 (c). Figure 8.6 (b) shows the corresponding IFDs. Figure 8.7 shows the least squares straight line approximation to the slope of the IFD curves (from $t = 0 - 5$ ms) for pinnae sweeping around elevation offsets of $0°$, $45°$ and $90°$. The relative values of the slopes are more dramatic around $90°$ and the differences change sign at $0°$; however, pinnae rotating around $45°$ induce a more dramatic Doppler effect and resolution is good through $\pm 20°$ elevation. Also, pinnae sweeping about an elevation offset of $45°$ experience no sign ambiguity.

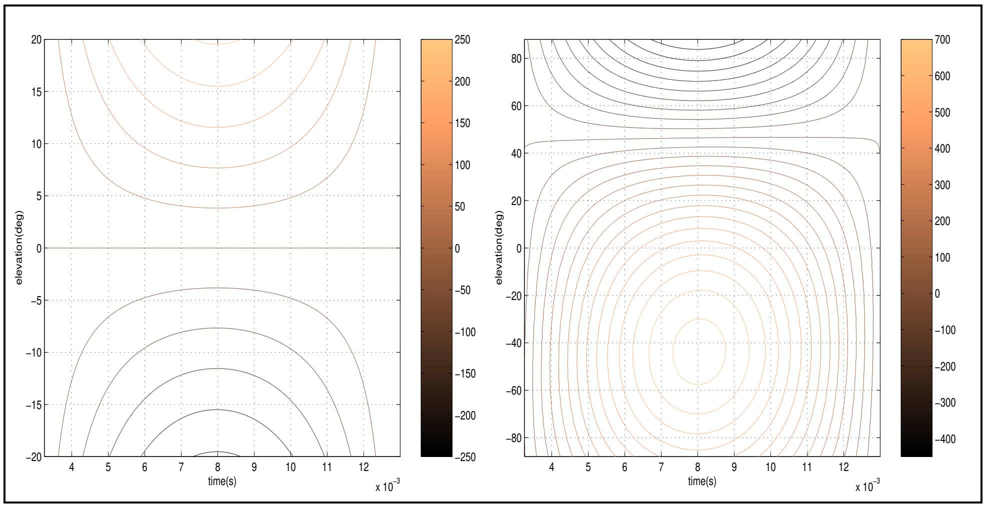

(a) (b)

Figure 8.6:
IFDs across different ear offset angles. (a) 90°, as in Figure 8.1 (b). (b) 45°, as in Figure 8.1 (c). (Colourbars indicate IFDs (Hz).)

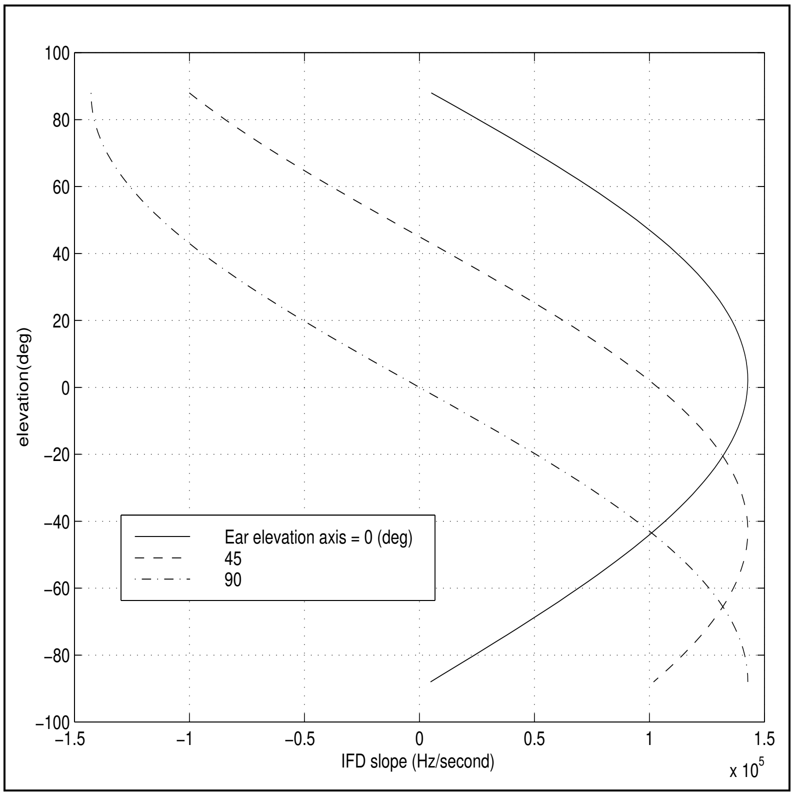

Figure 8.7: **IFD slopes for different ear offset angles.**

8.3.2 Pinna and target motion

In this section, I allow target (*i.e.* wing) as well as bat (pinnae) motion. Figure 8.8 shows the frequency deviations induced by a fluttering target (*via* 50 Hz iso-frequency deviation contours) as heard by stationary pinnae held at 45° offset elevation angles. (Target characteristics are based on those of a generic moth flapping a 16 mm long wing through a 30° vertical flap plane at a wing-beat rate of 50 Hz.)

By combining frequency cues into an IFD, directional cues are largely unaffected by target motion. We can see this by expressing echo frequency f_e differences as:

$$IFD(t) = f_c \frac{v_{b'_{left}(t)} - v_{b'_{right}(t)}}{c} \frac{1}{1 - \frac{v'_t(t)}{c}} \tag{8.4}$$

For simplicity, Equation 8.4 treats v'_t as being the same in both ears (*i.e.* $v'_t = v_{t'_{right}} = v_{t'_{left}}$). This is a reasonable approximation for closely spaced pinnae whose line of sight with the target do not change greatly during the movement. As shown by this equation, the influence of low speed target movements on the comparison of inter-aural frequency cues is so minimal that it can be safely ignored. Therefore, IFDs for fluttering targets are essentially the same as those for a stationary target located at the same elevations. (See Figure 8.8.)

By contrast, as indicated by Equation 8.3, the monaural Doppler cues are significantly altered due to the fact that target induced Doppler exaggerates or diminishes that due to ear movement. Figure 8.9 (a) and (b) show the frequency deviations in the right and left ears in response to a fluttering target positioned at elevation angles of ±20° across frontal sound field. The motion of the right pinna acts to cancel the wing induced Doppler, while the motion of the left pinna (moving opposite to the right) exaggerates it.

The insect parameters (*i.e.* flap rate, flap extent and wing length) used here were chosen to approximately compliment those of the bat pinnae. However, they fall within those of typical prey of *R. ferrumequinum* (*e.g.* several species of nocturnal moths). It is tempting to wonder whether the ears of bats, which move at similar speeds to insect wings and can be of comparable lengths, could induce stroboscopic effects such

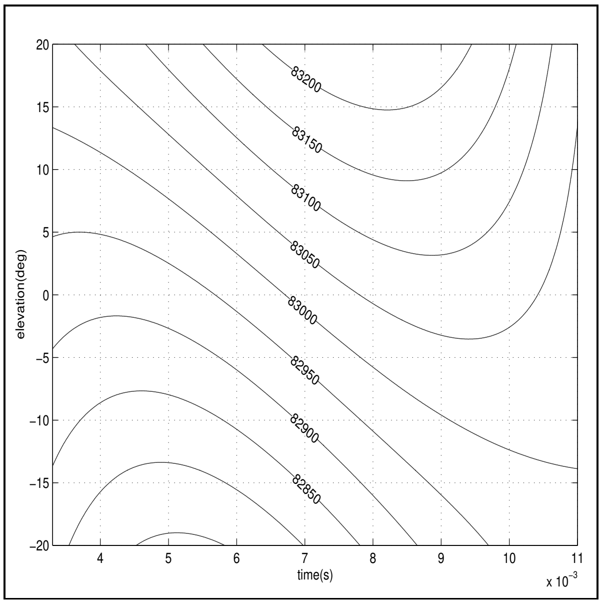

Figure 8.8:
Doppler due to fluttering target. Echo frequency deviation over time for fluttering targets at different elevations as heard in right and left stationary pinnae. Target is flapping a 16 mm long wing at a rate of 50 Hz through 30° within a vertical flap plane and comprises one reflecting point at the tip of the wing.

as those shown in Figure 8.9.

8.4 Discussion

In this brief exploratory chapter, I looked at frequency cues created by scanning simulated receivers through the vertical arc motions of rhinolophids and hipposiderids. Results shown here can only serve a qualitative exploration of the sensory space, as precise (numeric) predictions can only be validated once investigations with bats have determined how frequency cues are generated at the surface of moving pinnae and how the acoustic axes and Doppler collecting surfaces are related. Nevertheless, it is clear from this work that, due to the cosine dependence of the Doppler effect, fine target angle discrimination is likely to be difficult using frequency cues. However, frequency cues generated by movement of receivers through even small arcs (*e.g.* like those through which *R. ferrumequinum* drives its pinnae) can induce significant Doppler shifts (*i.e.* well above detection thresholds). This information could certainly

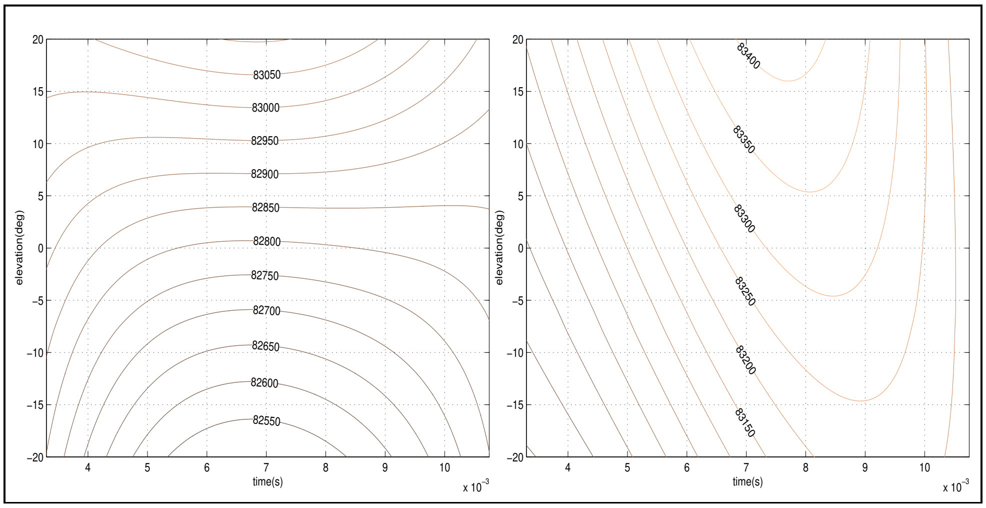

(a) (b)

Figure 8.9:
Instantaneous frequency shift *vs.* **target elevation (fluttering target).** Frequency deviations as heard by (a) right and (b) left pinnae swept through 30° about a 45° elevation offset angle.

be used to make a target more audible in situations, for example, where a stationary echolocator is searching for prey. Also, such movements could be synchronised to the movements of similarly sized targets to create stroboscopic effects which may be useful in detection and classification.

As echo frequency is also altered by the relative DC velocity between bat and target, any complete theory of Doppler cue exploitation by rhinolophids and hipposiderids must integrate cues derived from pinna movements with acoustic flow cues induced by the motion of the agents as a whole.

Part IV

Conclusions

"Twinkle, twinkle, little bat!
How I wonder what you're at!"

—sung by the Hatter at the Great Concert given by the Queen of Hearts in Lewis Carroll's *Alice's Adventures in Wonderland*

9. Evaluation: A biological explanation

In the empirical work reported in Parts II and III of this thesis, a bionic model was used to investigate the acoustic cues generated by particular types of target and receiver motion. Moreover, the sufficiency of those cues to engender a binaural echolocation system with robust and accurate spatial percepts was tested aboard this model system. In this and the following two chapters, I evaluate the results of this investigation. In the present chapter, I consider the performance of the system in terms of the biological hypotheses which it was constructed to investigate. In Chapter 10, I compare the bionic echolocation system with other artificial echolocation systems in the field of robotics. Finally, in Chapter 11, I look at the work in terms of its methodological contribution — *i.e.* I examine the bionic system as an platform from which to investigate perception.

The question of how bats localise prey can be examined from a number of different viewpoints. To this end, Ostwald *et al.* suggest that information from three different fields of study must be integrated [Ostwald *et al.* 88]:

> "a) *The spectro-temporal analysis of echoes by physical measurements and theoretical considerations gives insight into which echo parameters characterise the filter properties of a target. This shows which cues could be used by the bat.*
> b) *Systematic variations of these cues in behavioural experiments in the laboratory demonstrate whether bats can use this kind of information. Field studies indicate which parameters seem to be most important for*

bats under natural conditions.

c) Electrophysiological recordings of neuronal responses to simple and more natural auditory stimuli give hints about the processing of these echo cues."

In this computational (or "theoretical") study, my intention was (i) to investigate acoustical cues generated under conditions of realistic target and observer motion and (ii) to explore the simplest mechanisms which might transform acoustical cues contained in a single echo into target angular positions. From analysis of acoustical cues available in single echoes, I addressed the following four hypotheses.

- **Recognition is lunch (Chapter 4).** Echolocators hunting exclusively for oscillating targets may be so coupled to their environments that prey selective behaviour occurs as a result of a single localisation mechanism tuned to perform binaural comparisons of echo energy in particular frequency channels. Such a strategy is possible for a long CF echolocator hunting fluttering targets because periodically moving surfaces modulate an insonified probing beam so as to redistribute acoustic energy into particular frequency bands.

- **Partial IID maps for complete target localisation (Chapter 6).** By rotating pinnae through opposing vertical arcs, an echolocator can alter the directionality of its echolocation system (*i.e.* the region of the frontal sound field from which IIDs are measured most reliably). High duty-cycle bats could use such a strategy to realise a 3D echolocation technique which is believed to be employed by low duty-cycle bats — *i.e.* the creation of alternative viewing perspectives across which to compare IIDs. CF echolocators use pinna *movement* rather than *morphology/bandwidth* to alter the orientation of their acoustic axes. However, in both cases, echolocators who understand the mapping from 3D space onto in internal array of IID measuring neurons could disambiguate both target azimuth and elevation by using this strategy to sample in both dimensions.

- **Timing cues for elevation (Chapter 7).** Continuously scanning receivers set up amplitude cues which vary systematically with the relative angle of a target. I investigated the use of IID rates of change [Schnitzler 73] and the use of a novel timing cue: the delay between echo arrival and echo peak, to encode target elevation. In this way, pinna mobility could provide a high duty-cycle bat with a mechanism for creating dramatic temporal cues for directional sensing which, unlike inter-aural timing differences, do not degrade with head size.

- **Frequency cues for elevation (Chapter 8).** Scanning pinnae also induce frequency modulations (*via* the Doppler effect) which are related to the off-axis angular position of a target through the cosine-law. I investigated Pye's hypothesis that, using frequency cues, an echolocator with fine frequency sensitivity can partition echoes based on Doppler shift and thereby extract the angular position of a target [Pye & Roberts 70].

To assess the extent to which the results of the hypothesis testing aboard the bionic system might shed light on the mechanisms underlying target localisation in high duty-cycle bats, I describe the accuracy of the model (Section 9.1), discuss the extent to which the behaviour of the model supported the hypotheses (and why) (Section 9.2), and explore how consistent the model behaviour is with reports of behaviour of rhinolophids and hipposiderids (Section 9.3). This section concludes with a list of predictions and extensions to the model.

9.1 Model accuracy

The bionic model mimics the external auditory apparatus of a bat using functionally analogous transducers, electro-mechanical motion control circuitry and digital signal processing/filtering. The choice of the particular model components used here was constrained by the availability of hardware, descriptions of their acoustical properties (in the case of the transducers) and ease of implementation; as well as the availability of consistent and quantifiable descriptions of aspects of the bat's external and internal auditory system. Its purpose is not to serve as a detailed model of a bat head or its peripheral auditory nervous system, but as a working model which can exploit some of the same essential physics which might be used by the animal to perform target localisation.

Like any computable representation of a biological system, the required simplifications imply a number of assumptions about what are the relevant factors to represent, and what constitutes a satisfactory way to represent them. Webb argues that the "strength of the basis for these assumptions is critical for the strength of any conclusions that can be drawn from operations on the model" [Webb 93]. Wartofsky reinforces this view by advocating that the "existential commitment" of models "be supported by an analysis of the ways in which models ... *represent*: ... not simply an analysis of their structural properties, but of the relation between such properties and the purported reference of the models" [Wartofsky 79]. The following subsections revisit each of the important assumptions made in this work in an effort to quantify the ways in which various model components represent aspects of the bat's external and internal auditory system.

9.1.1 Pinna morphology

The ears of bats are complex physical structures and local effects involving reflections from pinna folds and crenellations are likely to generate important acoustic cues. Nevertheless, in this work, I ignore local effects and employ only the global directionality brought about by diffraction across the opening of the pinnae. I adopt an assumption common in pinna modelling; that the opening of the pinna experiences sound diffraction similar to a single, circular aperture [Fletcher & Thwaites 79, Fletcher & Thwaites 88, Guppy & Coles 88, Kuc 94, Andrews 95]. Across a limited range of frequencies, the directivity characteristics of the main sensitivity lobe of a circular piston (receiving in an infinite baffle) can be related to the diffraction limits imposed by the effective radius of the pinna opening in relation to sound wavelength [Guppy & Coles 88, Kuc 94, Andrews 95]. A more appropriate pinna model, which lacks the extensive side-lobes of the piston receiver, might be the conical horn. However, the piston model is a convenient choice for use in these studies because there exists an analytic expression for its directionality and inexpensive transducers with similar behaviour are readily available (*e.g.* the electrostatic Polaroid transducer [Biber *et al.* 80]). Although the directionality profile of this manufactured transducer is more symmetrical than that generated by the less regular pinnae of rhinolophids and hipposiderids, I argue that increased symmetry only makes unambiguous target localisation more difficult using the mechanisms hypothesised here. Furthermore, the mechanisms considered here do not rely on the echolocator having a *particular* directionality, rather, they simply require that the echolocator understand how the frontal sound field maps onto IIDs after filtering by the pinnae and head.

9.1.2 Directionality of the head

The radius of the Polaroid transducer used in these investigations ($a \approx 11.3$ mm) is larger than that of a typical rhinolophid or hipposiderid pinna opening and the resulting beam-width is narrower than that of bat pinnae. Moreover, due to external constraints, emitting and receiving transducers are mounted at a common elevation aboard a "head" which is approximately an order of magnitude larger than an average rhinolophid or hipposiderid head. The resulting directionality of the model sensor head is less circularly symmetrical than that of a bat (*i.e.* it is wider in the horizontal than the vertical directions). However, within the frontal sound field of both the model sensor head and those of bats echolocating at their dominant CF frequency, similar high resolution IID profiles arise: approximately $0.5°dB^{-1}$ robot (measured at $r \approx 0.3$ m), $0.7°dB^{-1}$ *R. ferrumequinum* (simulated), $0.6°dB^{-1}$ *Rhinolophus rouxi, R. clivosus, R. eloquent, Asellia tridens, Hipposideros lankadiva* (measured) [Obrist *et al.* 93]. Where differences in beam-width were significant, simulation results using a model head based on average *R. ferrumequinum* pinna dimensions ($a \approx 4.5$ mm, length ≈ 18 mm) and call frequency (83 kHz) were presented.

9.1.3 Pinna motion

The observed motion of rhinolophid and hipposiderid pinnae is not purely vertical. Qualitative observation of additional head movements, ear rotations, and changes in pinna shape have been associated with the systematic ear arc scanning behaviour [Möhres 53, Schneider & Möhres 60, Pye *et al.* 62]. These additional motions may allow the bat to change the size as well as the position of its acoustic magnifying glass. Nevertheless, here I investigated only changes in the vertical position of the acoustic axes.

9.1.4 Synchronisation of sound and movement

Ear movements appear to take place before the production of a pulse [Pye & Roberts 70]; therefore, it is speculated that movements may actually be synchronised with the arrival of (previous) echoes. In this investigation, maximal use of the receiver motion is made by ensuring that the receivers are at their extreme positions before echo reception and are moving throughout its duration.

9.1.5 Call structure

The echolocation calls of rhinolophids and hipposiderids contain multiple harmonics, and the long CF portion of each harmonic may be preceded and followed by short, approximately $10 - 30$ kHz wide FM sweeps. In many rhinolophids and hipposiderids, vocal tract processes largely strip the fundamental and third harmonics (see Section D.1 for a brief review). In this work, I assume that the bulk of the information used by these bats during echolocation is gained from the dominant CF_2 portion of the signal. (In this research, I worked with CF signals at the most efficient frequency for the transducer (50 kHz). Where simulation results are used, I employed the 83 kHz dominant CF frequency of *Rhinolophus ferrumequinum*.)

9.1.6 Information content per echo

I assume that the information available in a single echo contains complete (though, probably, imperfect) localisation information. In other words, I assume that a target's angular position can be determined from one echo measurement, but subsequent echoes may be necessary to reduce positional uncertainty. To my knowledge, the assumption that spatial information can be extracted from a single echo has not been directly tested, but would be of obvious value — allowing an echolocator to make immediate adjustments to its trajectory following the initial detection of a target. *Myotis* aims its head at an insect within 0.1 s after initial detection [Webster 67]. Assuming a search phase pulse rate of approximately 10 Hz [Schnitzler 70], it is possible that *Myotis* makes this initial orientation on the basis of a single observation.

9.1.7 Information content per ear movement

The previous assumption, in turn, implies that the pinnae need only make one sweep (rather than, for example, a full down-and-up cycle) to generate acoustic cues. This assumption is in keeping with the popular theory that bats employing broad-band sounds *simultaneously* sample target position along a number of acoustic axes within the beam of a single echo. More convincing, perhaps, is the fact that rhinolophids and hipposiderids emit pulses in groups containing as few as a single pulse/movement [Pye *et al.* 62].

9.1.8 Filtering

The sound filtering operations used to process echoes are based upon simple abstractions of those performed by the filter units in the cochlea: bandpass filtering in narrow (but overlapping) channels around the echo carrier, rectification and band/lowpass filtering of the echo envelope. Also, in the digital filter implementation, tuning curves reflect realistic Q_{10dB} values. However, these operations represent only a small subset of those actually carried out by filter neurons, *e.g.* they do not take into account the shape of the tuning curves, nor their preferences for echo intensity and different echo modulation parameters. Also, the digital filters do not encode information as latencies or firing rates, but pass filtered versions of the original time waveforms (regardless of signal level). The use of a neural model would have provided a more realistic combination of time and intensity cues and would neatly overcome the problem of eliminating low intensity, noisy signals. However, I chose to deal with timing cues explicitly.

9.1.9 Targets

It is well known that among high duty-cycle bats "oscillating targets and not specifically fluttering wings are detected and induce catching responses..." [Neuweiler *et al.* 87]. In many behavioural studies, targets such as rotating fans, propellers and vibrating loudspeakers are used to amplitude and frequency modulate call signals [Schnitzler 78, Trappe 82, von der Emde & Schnitzler 86, Nitsche 87, Ostwald *et al.* 88, Sum & Menne 88, Roverud *et al.* 91, Moss *et al.* 92].

The use of fans (and the use of the simulated insect described in Appendix C) as a model of real insect flutter is based on two broad assumptions: (i) insect flutter is truly periodic and (ii) reflecting surfaces are continuously visible throughout each flutter cycle. These assumptions (and other simplifications of the complex motion of real insects) are discussed below.

Periodicity

In realistic interactions between predators and prey, insects may employ avoidance behaviours which involve a disruption to normal flutter patterns. For example, some moths can detect a range of high frequency sound and respond by folding their wings and dropping [Surlykke 88]. This anomaly aside, is insect flutter truly periodic? Reports of wing position as measured by high speed photographs do not conclusively answer this question. Some authors have observed “abrupt changes” [Moss & Zagaeski 94] in wing-beat, while other suggest that wing motion is periodic for prolonged periods. For example, Kober and Schnitzler report that the “wing position cannot be distinguished between ... photographs although more than 100 wing-beats are between them” [Kober & Schnitzler 90]. Some of the most impressive observations come from insonification studies which found that wing-beat rate varied by less than $\pm 10\%$ (sometimes less than 5%) in one or more individuals from 40 nocturnal insect species, (7 orders) [Kober & Schnitzler 90]. Also, insect wing-beat rate is used in taxonomies as a species specific descriptor.

Continuous line-of-sight

The second assumption, *i.e.* that insect wings are continuously visible is a convention which I have adopted specifically for my present purpose. (In my simulation, insect wings are composed of one or more individual point-like scattering surfaces which each generate echoes that are visible throughout the hemisphere normal to the wing surfaces.) It is not, in general, possible to reason comprehensively about reflector visibility given the complex shape and movements of insect wings. In support of this assumption, I note that in a study of echoes reflected from fluttering insects insonified from a number of different angles, Kober and Schnitzler observed that spectral broadenings from opposites sides of fluttering wings produce Doppler shifts which are approximately equal (in magnitude) and opposite (in sign) [Kober 88, Kober & Schnitzler 90] — suggesting that wings have a strong reflective patch(es) located at a fixed distance(s) from the stalk which contribute principally to the echoes heard from either side of the wing. In demonstrating this — Schnitzler's *glint hypothesis* [Schnitzler & Flieger 83] — Kober remarked that in "the horizontal and vertical planes the spectrograms changed gradually in a systematic way when moving around the insect. ... By turning the loudspeaker in a horizontal plane around the flying insect the energy is slowly moving to higher frequencies and is symmetrically distributed at the 90 degree position" [Kober 88].

Reflected AM waveforms

The motion of insect wings is complex — involving twisting and changes in shape brought about by folding and buckling, as well as elevation and depression. These flap mechanics control the form of the modulating waveforms by determining *when* and *in which directions* wing surfaces are extended over the course of a wing-beat cycle. For example, in the case of the latter, Kober and Schnitzler characterised stroke planes by noting that when insonified from various angles within the horizontal plane, some species reflected louder glints from front/rear viewing orientations (0° or 180°) — *e.g.* this was true for beetles (*Coleoptera*) which flap their wings through a relatively small arc in a plane at a 45° angle to the horizontal [Kober 88, Kober & Schnitzler 90]. By contrast, moths (*Lepidoptera*), caddis flies (*Trichoptera*) and lace wings (*Plannipennia*) produce their strongest glints at 90° (from within the horizontal plane). Investigations of individual moth species (*Autographa gamma* and *Macroglossum stellatarm*) reveal that they move their wings in a stroke plane which is more vertical and therefore most strongly reflect side-insonifying acoustic beams. The *Diptera* studied did not conform to either of these patterns because they change the major axis of wing movement during the stroke [Pringle 57, Vanderplank 50].

The complex structure of insect wings plays a role in determining the number of glints per wing-beat. Insects with morphologically or functionally one pair of stiff wings produce simple echoes with largely one steep amplitude peak per wing-beat cycle — *e.g.* the true fly (*Diptera*) and wasp (*Hymenoptera*) [Seifert 75, Nachtigall 72]. The wings of these insects provide a single scattering area which reflects an acoustic probing beam when the wings move perpendicular to it during the up stroke. By contrast, the wings of, for example, moths (*Lepidoptera*), beetles (*Coleoptera*) and lace wings (*Pannipennia*) are twisted into a third dimension during flight and, therefore, present an array of multiple perpendicular scattering surfaces per wing-beat cycle, resulting in the production of several glints around the main one.

The artificial insects employed in the *3D Echolocation Simulator* did not mimic these complex morphologies and motions. Instead they flapped simple wings within a single plane and generated amplitude glints described by simple waveform elements (*e.g.* rectified sinusoids or sinusoids raised to various powers, pulses, saw-teeth). To serve as an accurate echo generator, the model would need to be extended to produce the fine temporal detail of lower amplitude reflections (off of wings and stationary parts of the insect body) which occur between glints. However, the present model does capture the gross scale width and extent of the amplitude peaks reflected by real insect wings by using realistic wing lengths, flap rates and percent AM modulations taken from [Kober & Schnitzler 90]). Moreover, signal processing algorithms do not rely on this simplicity; they only require that the echolocator learn the particular periodic amplitude modulating function reflected by a desired target.

The computer cooling fans used as targets in the robotic investigations represent a relatively simple reflecting surface compared to the surface of an insect's wing. However, the fan targets also produce realistically scaled glints. At the moment of glint production, reflections off fan blades (which are of comparable length to the wings of typical nocturnal species of European insect) are approximately 20 dB higher than those off stationary fan parts. Recall that insonification studies show that glints from insect wings may be 20 dB (up to 40 dB in the extreme) greater than those off stationary body parts [Kober & Schnitzler 90].

Reflected FM waveforms

As in amplitude modulation, wing morphology effects the FM waveform in complicated ways which are not investigated here. Functionally distinct pairs of fluttering surfaces may have different motions. For example, beetles have one pair of large, lift producing wings and a second pair of covering wings (the *elytra*) which simply covibrate at a smaller amplitude. Likewise, floppy wings may vibrate significantly — *e.g.* moths' wings are linked during flight, but they are large and flexible and twist in the middle of the up-stroke. These effects are likely to give rise to fine-detail (intra-glint) Doppler shifts.

An accurate model of the way in which insect wings Doppler sift echoes would also require detailed consideration of insect flight kinematics which, for many insect species, is not well understood. For simplicity, many authors (including myself) assume that, within a single wing stroke, the motion is sinusoidal. However, measurements with high speed cinematographs show that velocity of wing movements varies significantly through the cycle. Pringle claims "... the velocity of wing movement varies through the cycle in a manner far from sinusoidal and incapable of description by any simple mathematical formula" [Pringle 57]. For example, the complex motion of the fly *Glossina palpalis* involves different delays at top and bottom of the strokes and the wings are in motion for only about 12% of the total duration of the cycle [Vanderplank 50].

My simulation model only captures gross-scale features of the frequency modulating function reflected by real insects. Recall that insonification studies reveal that degrees of spectral broadening may be as much as $\pm 1 - 3$ kHz in typical modulation processes involving insect prey such as moths, beetles and flies with wing lengths ranging between $3 - 30$ mm and wing-beat rates between $20 - 100$ Hz [Kober 88, von der Emde & Menne 89, Kober & Schnitzler 90, von der Emde & Schnitzler 90]. The simulator calculates Doppler shift induced by artificial insect wings using a full 3D model of the spatial relationship between echolocator and target and, consequentially, spectral broadenings in artificial echoes contain a good analogue of frequency depths and rates of modulation introduced by real insects. Echoes reflected by the small computer cooling fans used in the robotic investigation show a modulation depth within a realistic range as well ($\pm 1 - 3$ kHz); however, their modulation rate is $5 - 20$ times faster than the wing-beat rate of typical nocturnal species of European insect. The latter fact made it relatively easy to detect individual sidebands in echoes from the fans.

9.2 Model behaviour

An important part of evaluation involves determining whether the model's behaviour constitutes an analogue of the hypothesised behaviour. In this next section, I review the performance of the system and discuss "why it works".

The purpose of any working model is not to provide a complete evaluation of a hypothesis, but to show that the proposed mechanism can produce the behaviour in question. As stated in Chapter 2, the extent to which the results can be applied to the animal system depends on the assumptions of the model; however, even general models (which rely on strong assumptions) can prove useful in understanding biological systems if the interaction of the model with its environment is thoroughly explained in terms of the underlying physical processes which gave rise to the results.

Thus, determining whether a behavioural hypothesis is realised by the behaviour of the model requires (i) systematic testing of the behavioural capacities displayed by the device and (ii) an analysis of why the system works. The latter includes the answer to such questions as: what does the agent do in a particular environment?, in which environments does the agent work?, under what conditions will it achieve its goals?, what forms of interaction require internal architectures such as memory?

In the following subsections, I address these questions concerning each hypothesised mechanism.

9.2.1 Recognition is lunch

In Chapter 4, I explored the possibility that the selection of targets with periodic motion might be explained by the operation of a single localisation mechanism performing binaural comparison of energy on the weighted output of a filter-bank. In this way, targets are preferentially localised depending upon the extent to which they reflect an echo spectrum which matches the weighting function. Such a mechanism for selecting amongst targets is possible when target echoes contain a specific, orientation invariant (*i.e.* locally available) cue.

A target with periodic motion can reflect such a cue. When a long CF carrier pulse falls on reflector surfaces moving with a characteristic periodic motion, that signal is amplitude and frequency modulated such that echo energy becomes concentrated in particular spectral (sideband) channels — the positions of which are determined by the particular structure and motion of the target. By using this target signature energy in binaural comparisons, target selective localisation can occur without the need to translate signals into intermediate geometric or associative recognition models. This localisation strategy is effective in environments containing large amounts of stationary clutter. Moving clutter (conspecifics, predators, other potential prey) will confuse an echolocator using this scheme in proportion to the extent to which reflections from these targets match the signature of the desired target.

The simple (two filter channel) target selection scheme utilised in the closed-loop fan localisation experiments (of Chapter 4) would require increasingly sophisticated filtering or weighting functions to select between targets reflecting similar signatures. For example, it cannot distinguish between similar targets moving at the same rate (or rates which are integer multiples of each other). The solution which high duty-cycle bats appear to have evolved involves the partitioning of AMFM echoes on the basis of a number of "features" which were not considered here. For example, in *R. ferrumequinum*, several auditory nuceli contain increased numbers of filter neurons tuned to pass 50 Hz echo envelopes (50 Hz is approximately the wing-beat rate of nocturnal moths that comprise a large amount of this bat's diet [Jones 90]). These units respond preferentially to particular combinations of echo intensity, frequency modulation depth, percent amplitude modulation, *etc.*

9.2.2 Partial IID maps for complete target localisation

In Chapter 6, I explored the possibility that a CF emitting echolocator can use pinna *movement* rather than *morphology* to exploit the same 3D localisation principle believed to be employed by broad-band emitting bats — *i.e.* the use of more than one SONAR horizon along which to localise sound. By moving a single, fixed-width acoustic magnifying glass (defined by the frequency of the dominant CF harmonic) to a series of different positions during reception, rhinolophids and hipposiderids may provide themselves with additional viewing perspectives across which to sample and compare IID values in the same way that broad-band emitting bats may compare excitation across different iso-frequency IID maps.

Such a scheme requires the echolocator to understand the way in which ear movements affect the position of the SONAR horizon. In this implementation, an IID map was memorised at several pinna positions and subsequent comparisons were made with reference to these internal representations. Using spherically symmetrical transducers, the number of maps required for good localisation performance makes this a rather cumbersome strategy. However, in systems with less symmetrical receivers (*e.g.* bat pinnae) or systems operating over a larger bandwidth, this scheme would require fewer memorised maps to remove ambiguity in target position.

This algorithm, like all others employed here, may be used to localise stationary, as well as targets with AC motion (oscillatory motion, flutter, rotation, *etc.*). In the case of targets with these characteristic AC motions, clutter can be eliminated automatically by only localising echo energy in the spectral sidebands reflected by the target. To localise stationary targets, other means for extracting target signatures and comparing only the relevant intensities are needed.

Targets with DC motion would change position from map to map, such that the simple combinations of IID maps used in this work would not result in an accurate position measurements.

9.2.3 Timing cues for elevation

In Chapter 7, I showed how continuous receiver motion (and continuous measurement of IID rate of change or echo peak delays) can lead to uniform angular accuracy across the frontal sound field because the receivers visit all SONAR horizons, rather than a discrete subset of them. The measurement of IIDs in a map-less scenario restricts sensing to a narrow region of the frontal sound field wherein IIDs change monotonically. However, the monaural peak delay encoding of elevation operates over the union of the receiver beam-widths (rather than their intersection) and could be used to localise target elevation approximately so that the echolocator can bring them into the centre of the frontal sound field where more accurate measurements can be taken.

The introduction of timing cues generated by receiver motion is appealing because it provides a small echolocator with a means for overcoming the temporal resolution problems inherent in measuring target angle *via* ITDs. Although inter-aural distances are fixed, by scanning its pinnae slowly, an echolocator can increase temporal resolution. Thus, just as dramatic pinna morphologies may have evolved to increase IID resolution, pinna mobility may have evolved to provide an echolocator with mechanisms for creating dramatic temporal cues.

Targets whose elevation changes during a measurement could not be localised accurately *via* these mechanisms. Although DC motion would confuse the system, peak delays can be used to localise both targets with AC motion and stationary targets. In the case of the former, however, target oscillation rate must be faster than pinna scan rate so that the two amplitude modulations do not interfere.

9.2.4 Frequency cues for elevation

The use of Doppler shift to encode elevation might be seen as a companion strategy to the use of timing cues for this purpose. It works over the same target types but, here resolution improves with pinna speed. The difficulty in using this scheme is that the cosine dependence of the Doppler effect creates a situation wherein targets in angular positions which give rise to the largest frequency deviations are also the hardest to differentiate. An echolocator can compensate for this somewhat by rotating its pinnae about an elevation offset angle so that their velocity is more orthogonal to the target. However, the results obtained here suggest that the pinnae motion would be more effective in obtaining a frequency domain separation of call and echo, than in inducing directional cues.

Another possibility for further research is that flapping pinnae may induce frequency domain cues for target recognition. When echolocator and target are moving receptive and reflective structures of similar dimensions and at similar rates, receiver motion can be synchronised with target motion so as to create particular cancellations or exagerations of Doppler cues.

9.3 Comparison with target behaviour

There are difficulties in making comparisons with biology because animal behaviours are rarely simple to describe. In this work, means for establishing correspondence were limited from the outset by the choice of a target system of high sophistication for which detailed neuroethological explanations are not available. Many "circuits" and sensory modalities underlie echolocation in bats and the mechanisms described here (should they prove to be used by the bat) will only make up a subset of the perceptual mechanisms which these animals have evolved to overcome various interference problems and detection problems (*e.g.* those associated with rapid changes in the relative position of a target).

The hypotheses investigated here addressed the question of target selection and proposed a selective localisation routine rather than independent recognition and localisation schemes which are commonly assumed. The distinction is somewhat subtle — given that the term "recognition" is broad enough to encompass this meaning — but this distinction serves as a useful reminder that the function of a recognition system in the context of this prey selective behaviour is to sub-serve localisation: the bat need not recognise targets for which a localisation (steering toward) response is not needed and, likewise, it need not localise targets that it cannot recognise.

Consideration of these factors leads to hypotheses of whole mechanisms. Typically, explanations are more fragmented. For example, much of the psychophysical work with high duty-cycle bats is aimed assessing discrimination performance with the explicit aim of understanding how fluttering insects are "classified". The fact that oscillograms and high resolution spectrograms do not reveal an orientation-independent, species-specific cue "obvious to a human observer" [von der Emde & Schnitzler 90], has resulted in an increasing amount of interest in the possibility that bats learn to associate cues with insect "categories" or even map back from acoustical signals onto 3D geometric models of insects. These speculations take us farther away from an explanation of which cues might be used to localise a target.

At the other end of the spectrum, recently, several neurophysiologists have suggested the possibility that targets with periodic motion might be localised on the basis of a binaural comparison of low frequency echo envelope modulations. Specifically, it has been suggested that filter units with ON/OFF response in lower and mid-brain regions of the auditory nervous system might play a role in a mechanism for deriving a localisation signal from inter-aural differences in echo envelopes modulated by low frequency wing-beats. Such a mechanism may provide a way of selecting insect targets over stationary ones, but does not suggest a mechanism for selecting amongst fluttering insects.

The mechanism which I proposed in Chapter 4 suggests prey selective behaviour can arise from a simple mechanism based on binaural comparison of particular patterns of spectral energy. In the following subsections, I discuss what evidence there is for bats exploiting sideband patterns in target identification and then compare the localisation performance of the bionic model with that which has been observed in various bats.

9.3.1 Target cues

Detection of an echo from a target with periodic motion (be it rotating fan, vibrating loudspeaker or flapping wing) induces a new phase in echolocation behaviour which is characterised by a marked increase in duty-cycle in rhinolophids, hipposiderids and *Pteronotus parnellii*. Calls which capture one or more cycles of wing-beat (and which are held within the acoustic fovea *via* Doppler compensation) will contain spectral sidebands — the resolution of which will increase with signal length.

Von der Emde and Schnitzler's investigation with manipulated Cranefly *Tipula oleracea* echoes showed that both the amplitude and frequency modulating waveforms are necessary for orientation invariant recognition. These authors concluded that a high duty-cycle bat's ability to identify prey in an orientation independent fashion is not based upon "any single amplitude or frequency parameter like bandwidth or a special kind of amplitude peak .... Their method of analysis is far more complex, requiring both frequency- and amplitude information of the echo" [von der Emde & Schnitzler 90]. Strong evidence in support of this claim comes from von der Emde and Schnitzler's 2-AFC experiments with two *R. ferrumequinum* trained to distinguish phantom echoes of *Tipula oleracea* orientated at 90°. These bats were able to select echoes from the same target at novel (unseen) viewing angles and, moreover, this selection ability was not restricted to the originally presented individual *Tipula oleracea*, but included other individuals of the same species. During testing wherein the amplitude glint component of the training echo was digitally sliced out and played back to the bats, both individuals could still recognise *Tipula oleracea*. However, when the echoes were manipulated such that the glints were left in but the information between glints was replaced by a constant frequency and amplitude signal, only one bat moved to the loudspeaker playing the manipulated *Tipula oleracea* sound. Discrimination performance further deteriorated when echoes were manipulated such that either the AM or FM component was removed: in both cases, choices of the two bats fell into chance.

It would be interesting to know how the spectrum of the sound changed after digital manipulation of the sound. In most studies, this information is not given. As an exception, Schnitzler's early work included an investigation of the influence of modulation rate on modulation depth discrimination for which he published the spectra of the stimuli used [Schnitzler 78, Schnitzler & Flieger 83]. From observation of the echo spectra used in these experiments, it is clear that at $f_{mod} = 50$ Hz the first sideband starts to emerge out of the wide spectral peak centred on the carrier. In tests with bats, Schnitzler discovered that discrimination performance required a modulation depth of $\Delta f \approx 60$ Hz for modulation rates below $f = 50$ Hz (which corresponds well with Doppler shift resolution). Above this rate, the minimum modulation depth required for target perception decreased with increased rate. He concluded that the bat's "high sensitivity to frequency modulations at high oscillation rates can be explained on the basis of sideband detection" (see review in [Schnitzler & Henson 80]).

Even if the use of sideband energy patterns could be established unequivocally, this still leaves open the question of whether an insect might reflect such a cue from all aspect angles. Numerous experiments have shown that bats learn a preference for a species, and not a specific aspect angle, such that preferred targets can be recognised from novel angles not presented in the training set. However, as far as I am aware, all studies have been conducted on bats captured as adults, so one cannot logically conclude that an orientation invariant clue is present in the echoes themselves because the bats tested may have previously met, *e.g. Tipula oleracea*, and, therefore, already learnt a model/category. However, in the absence of this information, the results leave open the possibility.

It is interesting to note that local cues are believed to be used by other bat species in the discrimination of fluttering targets (though prey selectivity has not yet been addressed specifically in these studies). For example, although their calls are generally too short (*i.e.* < 10 ms) to carry a description of a complete wing-beat cycle, low duty-cycle bats are thought to use the ensuing echo variations to recognise flying insects [Griffin 58, Roeder 63, Kober & Schnitzler 90]. For example, it has been suggested that to distinguish between the rotation frequencies of a propeller, two broad-band emitting bats *Pipistrellus stenopteris* [Sum & Menne 88] and *Eptesicus fuscus* [Roverud *et al.* 91, Moss *et al.* 92] used the interference pattern between the Doppler shifted echo from the moving blades of the propeller and the unmodulated echo from stationary parts of this wing-beat simulator to determine an effective delay separation [Sum & Menne 88].

Another proposed mechanism for broad-band flutter detection is that of stroboscopic hearing [Feng *et al.* 94, Moss & Zagaeski 94]. This mechanism was suggested based on the observation that amplitude glints across a train of echoes are faithfully encoded in the discharge rates of neurons in inferior colliculus of *Myotis lucifugus* when the AM rate is *not* an integer multiple of the pulse repetition rate. In other words, when pulse repetition rate and insect flap rate are synchronised, discharge drops. Therefore, it is proposed that a broad-band emitting bat may alter the repetition rate of its emission until it matches the wing-beat frequency of the insect at which time the insect's wings will appear stationary. It is noteworthy that sound repetition rates of *Eptesicus fuscus* during target pursuit do not change continuously over time, but rather, remain stable over brief periods before increasing.

9.3.2 Localisation

Target localisation performance varies between species and between experiments involving a particular species. One remarkable (and consistent) feature of the directional hearing systems of bats is that, unlike most mammals, vertical target discrimination appears to be nearly as good (or, in some cases, as good) as horizontal. In passive listening experiments with *R. ferrumequinum*, 4° has been reported as both the horizontal and vertical localisation acuity [Konstantinov *et al.* 73, Gorlinsky & Konstantinov 78]. This symmetric localisation acuity was measured using 50 ms CF pulses at 81 kHz and the experimenters ensured that the pulses were long enough or occurred at a high enough repetition rate to allow ear movements to take place between them. When the experimenters presented the bats with FM signals (in the range of those used in the echolocation cry) the results fell to 6° azimuth and 19° elevation. This result is suggestive of an underlying decoupled azimuth and elevation localisation scheme. One could speculate that IIDs (which can work across a number of frequencies) were used to obtain the good azimuth resolution and that elevation performance dropped due to missing CF_2 modulation cues.

On the other hand, one might interpret results of pinna distortion experiments to imply the use of IID comparisons across different SONAR horizons. Pinna distortion experiments with *R. ferrumequinum* reveal a more substantial decrease in the discrimination of target elevation, than azimuth, when the pinnae are immobilised [Schneider & Möhres 60, Gorlinsky & Konstantinov 78, Mogdans *et al.* 88]. However, in these studies, the *R. ferrumequinum* which learnt to compensate for pinna immobility by moving the head vigorously recovered elevation localisation acuity [Schneider & Möhres 60, Mogdans *et al.* 88]. Mogdans *et al.* even go so far as to state that rhinolophids are still able to fly through horizontal wires after the pinna immobilisation surgery because they tilt their heads with respect to the horizontal to obtain elevation cues [Mogdans *et al.* 88].

The good performance of the monaural peak delay mechanism investigated in this work suggests that if rhinolophids and hipposiderids used such a strategy, they might show better target localisation performance in ear-plug studies. In nearly all species studied, plugging of one ear so as to create asymmetrical intensities between the ears reduces obstacle avoidance performance; however, Möhres found that plugging one ear resulted in no degradation of performance in *R. ferrumequinum* [Möhres 53]. Unfortunately, subsequent experiments with rhinolophids [Flieger & Schnitzler 73] contradict these findings, leaving the situation unclear.

9.4 Predictions

Building and observing the operation of the bionic system did suggest a number of lines of investigation for behavioural and neurophysiological studies.

Detectable differences in echo signatures

Psychophysical experiments have not been performed using phantom echoes with carefully varied AM and FM waveforms containing characteristic sideband patterns. Although parameters such as modulation rate and extent have been implicitly investigated in studies with different insect species, further studies should hold these variables constant while measuring the bats ability to discriminate synthesised echoes containing different AM and, particularly, FM waveforms.

Insonification studies

Further insonification studies with real insects — wherein echoes are subjected to different degrees of feature detecting — are needed. Nonlinear techniques could be exploited to cluster echoes which have been encoded using various time-frequency representations. (Of course, any categorisation of echoes should subsequently be explained in terms of the underlying acoustic cues.)

Monaural localisation

To my knowledge, no studies have been undertaken wherein a single pinna was immobilised. Would high duty-cycle bats this condition resort to head movements, or prefer to scan a single pinna.

Modulation-based localisation

It may be possible to induce spatial illusions in high-duty bats performing target capture by using modified playback echoes in which the degree of amplitude or frequency modulation is altered so as to alter a target's apparent bearing.

Pinna movements

Pinna movements should be studied more carefully with 3D cinematography and, perhaps, by mounting devices on the pinnae themselves to measure the extent of the movement. Bats should be presented with different targets (*e.g.* insects fluttering at different rates) so as to determine whether pinna movements — which occur over the same range of rates as insect flutter —- may be synchronised with wing-beat.

9.5 Extensions

There are a great many modifications to the current bionic system which could make it a stronger model. Several possibilities which would be natural extensions to the present system are enumerated below.

Reduced scale head architecture

A new head architecture with smaller inter-aural distances would allow for the measurement of IIDs which are less dependent upon range.

Transducers

Similarly, it would be useful to replace the Polaroid electrostatic transducers with smaller transducers — possibly one of the many piezoceramic models. Transducers with a higher operating frequency would facilitate better measurement of Doppler.

Artificial Pinnae

Mounting of horns or acoustic baffles on the transducers could allow for the focusing of echoes from particular parts of the sound field onto the transducers — thereby producing more desirable directional properties. (See preliminary work in [Papadopoulos 97, Peremans *et al.* 98b].) Likewise, these structures — if mounted above the receivers — could increase the effective speed of the receivers so that Doppler cues could be used aboard the bionic system. (Higher mounts for the receivers, *i.e.* further away from the axis of rotation, would also help in this respect.)

Filters

The current model reflects very few of the details of the known functional properties of filtering neurons. A more biologically plausible filter-bank — whose output was encoded as response latencies or firing rates — would increase the biological plausibility of the system and improve its performance in certain tasks. For example, by extracting

IID within the context of overall intensity, unreliable readings could be eliminated.

Implementation of filtering operations in analogue circuits would remove the current restriction on the number of frequency channels which can concurrently perform real-time processing.

3D target tracking

An important extension to the model is to introduce 3D target tracking mechanisms. Bats rapidly orient their head toward a target upon detection and, clearly, the results of this study show that keeping targets in the information rich centre of the frontal sound field improves the accuracy of target localisation schemes. Likewise, "fixing" a localised target in the axis of the emitted beam also facilitates scrutiny of target features by removing/stabilising projection effects [Kick & Simmons 84, Hartley & Suthers 87, Kuc 96].

Receiver rotation (*i.e.* panning) is possible with the current apparatus and could be used to overcome potential correspondence problems. Further useful directionality alterations may be brought about *via* panning as well — *e.g.* IID resolution along a single SONAR horizon might be stretched or compressed in this way.

Pulse patterning

Relatively few studies have addressed the question of how bats combine their individual measurements. Clearly, significant improvements in signal-to-noise ratios can be obtained *via* multiple pulse integration [Barrett 87].

It is intriguing to wonder what other cues bats generate by patterning their pulse emissions. The tendency to group pulses and the number of pulses per group varies between species and circumstances (*e.g.* attack phase). In almost all bats, pulse repetition rate increases in novel situations and (generally) as the bat converges on a target. In the case of FM emitting bats, it has been hypothesised that stroboscopic effects can be created when the pulse emission rate is synchronised with insect wing-beat rate. In the case of high duty-cycle bats, it is possible that varying the pulse emission rate could enable a bat to set up cyclical patterns of amplitude modulation like those of conical scanning radar [Skolnik 62]. There are many other possibilities.

Other target signatures

Although the present work focused on targets with periodic motion, the results could be extended to reflections from targets with other characteristic motions. There are many motion cues which may be used to characterise targets. In the case of insects, echoes may contain a description of the species-specific flight path and/or flight speed — *e.g.* beetles fly slowly while hawk moths can reach speeds up to 5 m/s; mosquitoes tend to fly in swarms, which places specific restrictions on their movement. Similarly, moths exhibit a host of avoidance behaviours which may distinguish them (though not necessarily make them easier to catch) [Surlykke 88].

9.6 Summary

Within Microchiroptera, bats employ an impressive range of morphological specialisations (*e.g.* noseleaves, specialised pinna structures, *etc.*), neurophysiological specialisations (*e.g.* an acoustic fovea) and behavioural specialisations (*e.g.* systematic ear movements, "compensation" of call frequency, *etc*). Just as much may be learned about the operation of radar from intercepted signals and the appearance of its aerials (as has been shown by military photo-reconnaissance), so "the faces and the ultrasounds of bats are highly significant indicators of their modes of operation" [Pye 84]. Because form and function of these specialised structures are influenced by the physics of the task they perform, physical characterisation and modelling of these sensor systems is likely to be an increasingly successful method of investigation.

10. Evaluation: An echolocation system

In addition to any value that the bionic system has as a bat model, the work described here resulted in a robotic implementation of a functional perceptual mechanism. As such, it can be compared to other artificial echolocation systems in the field of robotics. Here I do that *via* a brief overview of SONAR sensing in robotics (Section 10.1). This is followed by a review of "intelligent SONAR sensors" (Section 10.2). This review is by no means comprehensive. (Most conspicuously, it misses out a great deal of work on the use of ultrasonics in map building.) Instead, it looks at a few artificial echolocation systems which have exploited the principles of matched and dynamic filtering.

10.1 Historical

Probert has identified two factors which, in practice, have governed the selection of sensors for autonomous robots [Probert 94]. First, much of the initial work in robotics was carried out by researchers in Computer Science and Artificial Intelligence who had little interest in the development of hardware. Also, as robotics is a small field unable to provide demand for mass-produced, cheap sensors, many of the commonly used robotic sensors are based on technologies designed for other applications. Consequentially, Probert argues [Probert 94]:

> *"... the technology performs below its capability in robotics. This weakness has been turned to a strength in leading development of serious sensor modelling techniques. However the future lies in exploiting hardware tech-*

nologies as well as software models."

The use of SONAR sensing in robotics has been classically disadvantaged by this history. Initially heralded as a cheap, rugged and simple sensor for measuring range, the ultrasonic transducer was amongst the most common device used in robotic navigation systems for more than a decade. The most popular "ranging" kits — designed and manufactured by the *Polaroid Corporation* for use in camera auto-focusing systems — were sold with a signal processing circuit which extracted time-of-flight (TOF) to the first supra-threshold echo return (derived through a capacitive charge-up circuit with some automatic gain control). The perceived advantage of this system was that all the analogue electronics could be bought off-the-shelf. Researchers did not have to bother with echoes themselves: the sensor promised to require only a timing circuit to deliver range. Indeed, in many robotic applications, range was taken to be directly related to TOF. This proved to be a bad assumption in a number of acoustically smooth indoor environments.

The problem is essentially the following. Unlike light, sound wavelengths may approach or even exceed the dimensions of typical terrestrial reflecting surfaces. In indoor environments — *e.g.* the laboratories in which generations of robotic SONAR navigation/ranging systems have been developed — most surfaces are large and acoustically smooth such that reflections will be in the so-called "optical region" ($\frac{a}{\lambda} >> 1$), where a represents surface radius. In this case, reflected energy is independent of call frequency and echoes are highly directional. Such a reflection is sometimes referred to as "specular". As depicted in Figure 10.1 (a), (b) and (c), specular reflections produce relatively loud echoes which radiate in a direction determined by the angle of incidence. The physics of this situation naturally encodes object orientation; however, it can render objects invisible (*e.g.* Figure 10.1 (b)) to echolocators with limited means for re-orientating their sensors.

If one takes the TOF measured by the Polaroid transducer to be proportional to range along the direction normal to the transducer, echoes which return indirectly – after scattering off multiple intermediate surfaces — will overestimate range. Equally problematic to the ray-tracing assumption is the fact that echoes reflected from off-axis targets — which may fall into the path of the lower energy side-lobes of the insonified

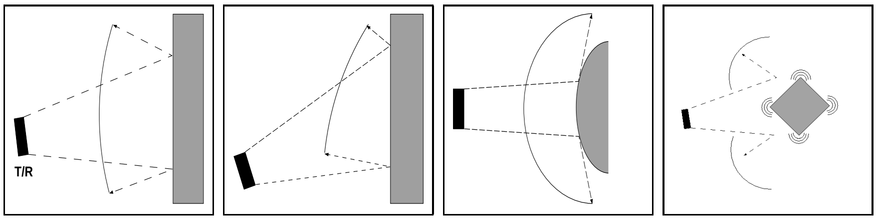

(a) (b) (c) (d)

Figure 10.1:
Surface orientation and visibility. (a) Transducer orientated at 5° with respect to the surface normal of acoustically smooth, large planar reflector. Intensity of returned echo $I \propto \frac{exp^{-2mr}\sigma_{wall}}{r^2}$, where the numerator is the product of the effective target radius (σ) and atmospheric absorption (m is an energy attenuation constant which is inversely proportional to frequency) and the denominator contains the frequency independent beam spreading term. (b) Transducer orientated at 25° degrees with respect to the surface normal. No echo received. (c) Convex, specular surface. A convex reflector introduces beam spreading such that incoming spherical waves are scattered to form new waves, centred at the point of reflection. For a cylindrical reflector, $I \propto \frac{exp^{-2mr}\sigma_{convex}}{r^3}$. (d) Composite surface producing scattering *via* reflection off planar faces (specular) and spherical facets (diffuse), and *via* diffraction from surfaces in the shadow region. $I \propto 3\,\frac{exp^{-2mr}\sigma_{sphere}}{r^4}$

probing beam (see Figure A.2) — lead to range values which underestimate the true target distance normal to the transducer.

Roboticists have employed a host of techniques for correcting *a posteriori* the errors resulting from straightforward interpretation of a range data. For example, *grid based approaches* integrate readings from multiple sensors and/or multiple viewing angles into a grid map describing the occupancy probabilities of cells in a Bayesian framework (*e.g.* [Moravec & Elfes 85]). *Feature based approaches* extract descriptions of surface type, position and/or orientation. There are a number of variants of this approach which include the application of statistical techniques (*e.g.* Kalman filtering [Leonard 90]) or qualitative criteria (*e.g.* [Mataric 90, Walker 92, Kurz 93]) to combine wheel odometry with several uncertain range measurements of the same object as seen from different viewpoints. These approaches have met with degrees of success, but other techniques are needed for exploiting the information available in the echoes themselves.

Over the last $5-10$ years this situation has been gradually improving due, in part, to the efforts of some researchers to incorporate the principles of matched and dynamic filtering in their echolocation systems. These sensory systems have come to be called "intelligent sensors". In such a system, "the sensor signal processing algorithms interpret the observed data, the interpretation being based upon the *physical principles* governing the sensor and a *model* of the environment that is being examined" [Barshan & Kuc 90]. In the next section, I consider these physical characterisations in terms of the principles of matched and dynamic filtering.

10.2 Intelligent artificial echolocation systems

10.2.1 Matched filters

Much of the work on intelligent sensors which I will discuss here has its roots in a variant of the feature based approach wherein *physical* constraints are employed to extract descriptions of surface type, position and/or orientation (*e.g.* [Hallam 84, Kuc & Siegel 87, Drumheller 87, Barshan & Kuc 90, Kuc & Viard 91, Peremans 94]). In indoor robotic environments, an emerging classification standard for target types includes *planes*, *corners* and *edges* [Barshan & Kuc 90, Leonard & Durrant-Whyte 91, Sabatini 92, Manyika & Durrant-Whyte 93, Peremans 94, Kleeman & Kuc 95]. These reflectors are composed of smooth surfaces that act as mirrors and allow one to reason about a transmitting/receiving (T/R) transducer as a separate transmitter and virtual receiver which have particular geometric relationships with targets — see Figure 10.2.

Barshan and Kuc showed that planes and corners cannot be differentiated using the amplitudes measured by a single circular (T/R) transducer. They discriminate planes and corners by exploiting the fact that the difference in sign of the virtual receiver characterising these surfaces can be detected in the pulse amplitude measurements using a two transducer array (both transmitting and receiving) — see Figure 10.2 (a) *vs.* (b). Peremans uses TOF readings to classify planes and edges using three laterally separated receivers (central transmitter) and employs sensor movement to distinguished between corners and planes [Peremans 94]. Kleeman and Kuc show that the movement of Peremans' system is equivalent to placing another transmitter at a new location and that, in principle, a sensor composed of two transmitters and two receivers is necessary and sufficient for discriminating planes, corners and edges [Kleeman & Kuc 95]. They

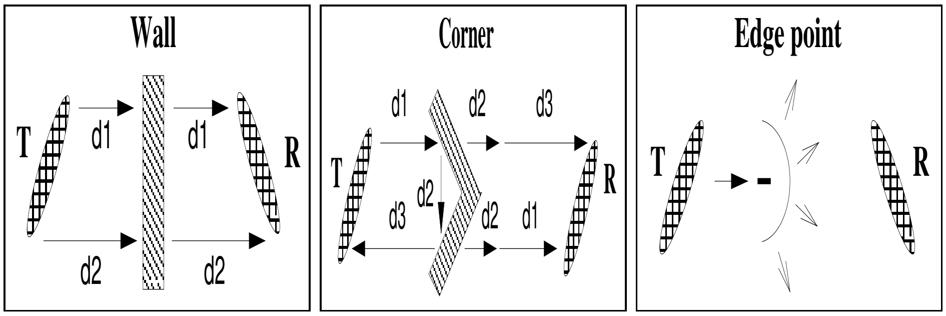

(a) (b) (c)

Figure 10.2:
Wall, corner and edge reflectors. For specular targets, a single transducer can be viewed as transmitter (T) and virtual receiver (R).

go on to build a sensor of this sort using a novel receiver array wherein transducers are closely spaced so as to minimize the potential correspondence problem of associating different receiver echoes from multiple targets.

All of these systems operate on whole echoes and derive time of flight (and from that target angle) based on carefully detected/filtered echoes. For example, Peremans obtains good range accuracy (and derives from that $\approx 1^\circ$ angular resolution) by employing a pseudo-random Barker codeword to bi-phase modulate a 50 kHz carrier pulse. The resulting waveform yields optimally low side-lobe levels in the autocorrelation function, which, in turn, decreases range ambiguity and, in the case of multiple overlapping echoes, increases range accuracy. Kuc and Kleeman use a linear filter model to generate a set of echo templates as a function of range and bearing angle such that the optimal echo arrival time can be estimated from the maximum cross-correlation of the echo with the templates. Range estimates are accurate to within 1 mm and bearing (which is derived from range) is accurate to 0.1°.

Klangbilder

The physical characterisations which gave rise to these models employing the directionality of specular reflections are appropriate in a number of indoor environments. However, this use of SONAR exploits few of the potentially available cues which echolocating animals appear to rely upon. In the case of bats, who are believed to use echolocation primarily for prey capture [Kunz & Pierson 94], targets (*i.e.* insects) are complex, moving and small (insect dimensions are approximately $3 - 30$ mm [Ostwald *et al.* 88, Kober & Schnitzler 90]) reflecting surfaces. Small targets ($\frac{a}{\lambda} << 1$) scatter acoustic waves diffusely — like light waves. In this way, surface topology (*i.e.* roughness or variations in 3D surface geometry whose dimensions approach λ) is the acoustic equivalent of colour. Echo "colouring", or *Klangbilder*, also comes about through the creation of new frequencies (*e.g.* *via* the Doppler effect). In both these cases, the spectrum can contain a description of the reflector type (its size, shape, and/or motion) and location, as discussed below. (Unlike specular reflections, which are highly directional, diffuse reflections can be seen over a wide area, see Figure 10.1 (d).)

A dramatic testimony to the richness and robustness of information available in the echo spectrum comes from human users of Kay's *Ultrasonic Mobility Aid* [Kay 61, Kay 74] — developed as an acoustic navigation aid for the blind. This device — mounted aboard a headset or spectacles — converts the high frequency reflected echoes of FM emissions down into the humanly audible range by mixing them with a copy of the transmitted signal. The frequency of the resulting beat-note is proportional to distance and its fine-structure contains target descriptions [Kay 61]. All further transformations of the signals take place in the auditory systems of users who demonstrate far superior target localisation performance using the spectral as opposed to temporal cues [Kay 61]. For example, Gissoni, a blind user, states that after several months of wearing the echolocation device, notes heard singularly or in combination can be as meaningful and informative as words spoken in your own native tongue [Gissoni 66].

Kao and colleagues built an automatic technique to recognise three types of architectural structures on the basis of beat-note structure: periodic reflectors (*e.g.* stairs), smooth and rough surfaces [Kao *et al.* 96]. Periodicity and decay rates of autocorrelated echoes could be used for the former categorisation tasks; however, rough surfaces have a strong orientation dependence and require integration of several readings taken while scanning the surface in question. In this work (and that of Bozma and Kuc [Bozma & Kuc 94]) surface roughness is modelled using a zero mean Gaussian distribution where σ represents roughness in terms of the height of echo energy deviations. Kao *et al.* describe a number of different rough surfaces as T/σ, where T is a spatial correlation constant encoding the regularity of the roughness.

Though the beat-note hypothesis is no longer used as an explanation of bat signal processing, the scrutiny of spectral cues is clearly of fundamental importance in mammalian hearing. Broadcasting a broad-band sound and listening for the filtering effects is one way to induce this information. Alternatively, a system can look for spectral broadenings in a narrow-band signal which come about due to target motion. In this respect, echoes originating from moving targets partition the world into animated foreground targets (which generally require immediate behavioural responses) and background targets (which may be thought to form a framework/envelope within which interaction with foreground targets can take place).

In this thesis I argue that a set of behaviourally relevant targets for an autonomous mobile robot must include dynamic as well as static aspects of the real world. Furthermore, I showed that acoustic motion descriptions can add descriptional richness without proportionally increasing complexity because certain sensing tasks are more naturally described in an animated acoustic framework. In the case of long CF echolocation, targets with periodic motion reflect a pattern of spectral sidebands which can be filtered to eliminate call energy, reflections from stationary clutter and reflections from objects with the "wrong" motion. In the previous chapter (Section 9.2.1), I discussed in detail why and under which circumstances this mechanism works.) It is possible that this approach will give rise to the development of reflectivity classes for a series of indoor targets with periodic motion — *e.g.* different pedal gaits, types of wheeled robot motion, *etc.* — and to characterisation of target motions of other types.

10.2.2 Dynamic filters

The principled use of transducer motion is only just beginning to be used in artificial SONAR systems, although optic flow and various types of visual-servoed motion have been employed in dynamic sensing for some time.

Some of the indoor robotic sensors described above employed motion to effectively create another transmitter or receiver (*e.g.* [Peremans 94]). Aboard a robotic system, the choice of whether to create additional measurement axes in time or space is largely task dependent. For example, Manyika and Durrant-Whyte are concerned with information rate aboard a highly mobile platform and overcome some of the inherent limitations imposed by the speed of sound *via* feature tracking — *i.e.* they exploit dynamic reorientating of the sensor to overcome the traditional "stop-look-move" cycle which is often necessary using sensors with fewer degrees of freedom [Manyika & Durrant-Whyte 93]. Their sensor consists of a pair of transducers rotating on a common axis (one transmitting, both receiving) to extract *Regions of Constant Depth* (RCDs) [Leonard 90], *i.e.* series of readings composed of regions in which the measured range is constant — thereby encoding range and bearing of wall, corner and edge targets.

Kuc has developed several "biologically inspired" target tracking systems to explore the "minimum amount of information and actuation that is needed to track an object in three dimensions" [Kuc 93]. His $ROBAT^{3D}$ consists of five wide-beam Panasonic transducers mounted in a cruciform pattern with a central transmitter flanked by a pair of receivers along both the horizontal and vertical axes. After matched filtering, the earliest arrival time detected in each pair is used to determine $sign(\theta)$ (where θ is the target angle) upon which nonlinear yaw and pitch corrections are made in the appropriate direction.

Kuc went on to build another system that adaptively changes the lateral position of its transducers in response to the echoes it receives [Kuc 96]. In this bionic system, a centre transmitter is flanked by two receivers — which each rotate on stepper motors. When echoes are received, all transducers pan to "focus" on the echo producing object using range and angular estimates derived from TOF triangulation. This has several consequences: it standardises the location of an object within the beam (thereby reducing the dependency of the reflected waveform on target aspect angle), it also maximises echo amplitude and bandwidth (by eliminating the lowpass filtering effects induced by the sweeping of a wavefront across a tilted receiver aperture). Though the system only moves in 2D, he notes that "a change in elevation of the object produces a systematic attenuation in the high-frequency components in the echoes that may prove useful for elevation determination"[Kuc 96].

Another dynamic echolocation mechanism involves the control of emission parameters — *e.g.* rate [Kuc 96] and frequency [Lindstedt & Olsson 93]. In the case of the latter, Lindstedt and Olsson probe for 3D target information (within an assembly work-cell filled with machine parts) directly by switching the frequency of an emitted pulse (heard by an array of four receivers surrounding the centrally positioned transmitter). By varying the frequency of transmission (between 40 and 200 kHz), range estimates from $1 - 2$ m down to a few centimetres can be resolved with an accuracy of 0.5 mm. Phased-array techniques are used in robotic echolocation applications. For example, Webb and Wykes obtain target angle aboard an airborne robotic echolocation system by digitally phasing the echoes collected in novel array of receivers [Webb & Wykes 96].

In this work, I've explored how moving receivers through systematic scanning motions can disambiguate the 3D position of a target in the frontal sound field. This kind of strategy can enable a CF echolocator to concentrate its signal processing resources on filters centred around its own call frequency so that locations can be sensed as fine amplitude and frequency modulations superimposed upon a long carrier. (See evaluation in Sections 9.2.2- 9.2.4.) The most obvious benefit of an acoustic localisation system which employs cues other than ITDs for target angle discrimination is that the baseline (*i.e.* inter-aural separation) can be made arbitrarily small. Also, by exploiting observer motion rather than transducer bandwidth to induce cues for 3D hearing, it is possible to manufacture ultrasonic sensor systems using one of many off-the-shelf narrow-band piezoceramic ultrasonic transducers — thereby driving the material cost of the device way down. Given the expected centimetre range accuracy and $\pm 2^{\circ}$ angular accuracy, such a sensor system would be practical for use within a robotic work-cell, aboard a automobile positioning system, an intelligent mobility aid for a blind or disabled person, *etc.* Its 3D localisation capability would be a significant improvement over existing commercially available ultrasonic sensor systems which deliver only range [Biber *et al.* 80] and 2D bearing [Kay 80].

10.3 Summary

The history of the development of SONAR systems for robots is interesting in that we can clearly see how the lack of a physical characterisation in early work led to inappropriate assumptions (*i.e.* a simple relationship between range and TOF in wide-beam systems) and poor system performance. Perhaps it is useful here to draw a distinction between *simple* and *sloppy* uses of sensory information. Animal echolocators employ a number of simplifying assumptions concerning the range of inputs for which solutions are necessary; however, these are based upon the exploitation of specific geometric constraints which are by no means sloppy. Similarly, intelligent artificial echolocation systems achieve impressive performance (and higher information rates) when they exploit simplifying assumptions appropriate to their indoor, specular environments.

Arriving at a physical characterisation of an environment, agent and task which affords the use of task-simplifying assumptions is difficult. Where possible, it seems highly beneficial to speed up the robot design process by employing biological sensing strategies. The use of target and sensor motion is clearly one area where inter-disciplinary research is likely to be fruitful.

From their work in modelling dynamic cochlear mechanisms, Mead and Lyon arrive at this same perspective [Lyon & Mead 88]:

> *"When we understand how hearing works, we will be able to build amazing machines with brain-like abilities to interpret the world through sounds (i.e. to hear). As part of our endeavour to decipher the auditory nervous system, we can use models that incorporate ideas of how that system works to engineer simple systems that hear in simple ways. The relative success of these systems then helps us to evaluate our knowledge about hearing, and helps to motivate further research."*

11. Evaluation: A methodology

AI's goal of understanding perception is shared with a number of longer established fields such as Philosophy, Psychology and Neuroscience. The essential difference is that in AI researchers are committed to computational modelling as a methodology for explicating the perceptual processes which underlie intelligent behaviour [Howe 94]. Clancy affirms this perspective: "We can generalise what AI programming is in terms of modelling methodology" [Clancey 91].

The computational approach to investigating perception is based on the perspective that a physical system carries out a transformation of information which can be described analytically due to its relation to a more abstract algorithm [Sejnowski *et al.* 88]. The computational model serves as the platform on which that algorithm, or mechanism, is implemented (and observed, tested and evaluated). Thus, the types of models which are typically built in AI can be distinguished as *physical models*. They are "physical embodiments" (*i.e.* bionic models) of the hypothesised perceptual mechanisms underlying a theory of behaviour which act dynamically to produce data [Webb 93]. It is important to note that, in robotics research, the robot itself may be the subject of study; however, when the robot is used to study animal perception, its behaviour is investigated as an explanation of a biological system's behaviour. Obviously, the robot can serve both functions; however, in interpreting the behaviour of the robot, these distinctions must be kept clear.

The primary advantages of the physical model (some of which it shares with purely mathematical and other forms of models) include the following.

1. **Alternative platform.** A physical model serves as an alternative platform on which to carry out experiments which are difficult (or impossible) to perform in living organisms.

2. **Independent levels of explanation.** A physical model makes a system with complex interacting components more accessible by allowing the investigator to examine it at a particular level of abstraction.

3. **Explicit descriptions.** A physical model provides an explicit, detailed description of a system which is accessible to formal analysis and can, therefore, be used to communicate results across the boundaries of different disciplines.

4. **Separation of *Merkwelt*.** A physical model forces a separation between the experimenter's perceptual world (*Merkwelt* [von Uexküll 21]) and that of the system under investigation.

Modelling *via* physical systems (*e.g.* bionic systems) is a relatively new approach and its methodology has not been established. In the following subsections, I discuss, in more detail, the value of this approach and provide examples from this investigation of perception *via* echolocation.

11.1 Alternative platform

There are many ways in which to investigate theories of perception. Current neuroscience technology only affords investigation of the generation and modification of behaviour in simple reflexive systems. In mammals, the relation between perception and single neurons is particularly difficult to study because the sensory capacities assessed with psychological techniques are the result of activity in many neurons from many parts of the brain [Sejnowski *et al.* 88]. Explaining relatively simple behaviours, such as stereotypical eye movements, involves complex interactions among large numbers of neurons distributed in many different brain areas [Lee *et al.* 88]. Sejnowski and colleagues state [Sejnowski *et al.* 88]:

> *"Explaining higher functions is difficult, in part, because nervous systems have many levels between the molecular and systems levels, each with its*

> *own important functions. ... Properties not found in components of a lower level can emerge from the organisation and interaction of these components at a higher level. ... The sources of such network properties are not accessible by the use of methods suited to investigating single neurons."*

Advances in neuroscience technology are improving the spatial and temporal resolution of current approaches (*e.g.* techniques have been developed for simultaneously recordings from multiple single units [Eckhorn *et al.* 88]); however other approaches to understanding perceptual mechanisms are needed.

Thus, the construction of physical models allows for styles of study which are currently impossible, impractical, unethical or too expensive to perform on animals [Moravec 66]. For example, in bats, simultaneous recordings from filter neurons with different best frequencies, while the bat is receiving echoes off a fluttering insect, cannot be performed. Therefore, it is difficult to assemble a picture of the aggregate neural response which the motion of insect wings (insonified from different viewing angles) might invoke. (The situation is more complicated when both target and bat motion is involved.) By contrast, aboard a bionic system it is possible to view the response of a filter-bank to echoes from different targets and to observe the extent to which receiver motion modifies the acoustical cues afforded by a target.

Although the lack of a neuronal signal processing model aboard the bionic system does weaken the strength of its claim to be a model of bat audition, separating out timing and intensity cues for directional sensing (as opposed to wrapping them up into a single encoding such as latency or firing rate) allows one to see the relative contribution of the cues and to identify, for example, timing cues for directional sensing which do not degrade with head size (*i.e.* peak delays, ITD_{peaks}, rates of change of IIDs). The lack of a neuronal model also highlighted the utility of neuronal thresholding mechanisms which perform binaural comparisons in the context of overall signal levels so that comparison of low level signals does not result in, for example, unreliable IIDs.

11.2 Independent levels of explanation

Rosenblueth tells us that "the best material model of a cat is another cat, or preferably the same, cat" [Rosenblueth 70]. This is true in the sense that such a model would be the most realistic representation; however, such a model would be of limited use to the computational scientist. Models must simplify [Achinstein 64].:

> *"To propose something as a* **model** *of (an)* **x** *is to suggest it as a way of representing* **x** *which provides at least some approximation of the actual situation; moreover, it is to admit the possibility of alternative representation useful for different purposes."*

The flexibility which a computational model allows for selection of alternative "representations" is both the strength and the weakness of the approach. To ensure it is the former, it is useful for models to function within a larger investigative methodology wherein the questions of "why?", "what?" and "how?" are largely decoupled [Marr 82]. For any characterisation of "why" a system is constructed in the way it is, a decoupled framework would allow us to hypothesise one or more equivalent information transformations required to duplicate the behaviour of the system ("what?") and for each of these, one or more implementations ("how?"). Here, the physical model of AI — *i.e.* the robot or computer simulation — serves as an implementational device which can take on the appropriate degree of "existential realism" [Wartofsky 79] required of the information mapping function.

It is important to note that other fields may be committed to a particular level of explanation. For example, neuroscience is clearly tied to the third ("how?") level of description, and psychophysics to the second ("what?"). However, each of these *three levels* of description should have its place in the eventual understanding of a system [Marr 82]. Moreover, by testing and investigating perceptual information transformations independently of implementation-level details (*e.g.* the full complexity of neurons), one can create a new framework in which to understand a system. A classic example of the advantage of this approach is the difficulty of understanding bird flight by only studying wings, Marr claims: "It just cannot be done. In order to understand bird flight, we have to understand aerodynamics; only then do the structure of feathers and the different shapes of birds' wings make sense" [Marr 82].

Thus, a physical model can make a system with complex interacting components more accessible by allowing the investigator to examine it at a particular level of abstraction. In this regard, designing and building physical models of animal systems facilitates investigation of a mechanism *via* "downhill synthesis" — a term coined by Braitenberg to suggest that it can be easier to put together a mechanism that produces apparently complex behaviour than to work back from the behaviour to the responsible mechanisms [Braitenberg 84].

For example, although selection of fluttering prey in acoustically cluttered environments does appear to be a complex behaviour, it does not necessarily require a complex mechanism. In designing an artificial system to localise oscillating targets, it becomes clear that an elegant way to overcome a correspondence problem is to adopt the principle that any cues which might be used to 'recognise' or classify a target, must also provide localisation cues. (This isn't the case when invoking one sensory modality, say, vibration, to recognise the presence of the snoring giant, and another, *e.g.* smell, as a directional gradient; however, when performing a task using a single sensory modality, robust integration of cues can occur at a low level.) By constructing localisation algorithms which weight the outputs of the different frequency channels of a filter-bank (and binaurally comparing them to provide a proportional steering signal), prey selective behaviour results without the need for formal target recognition. In this case, the computational perspective can allow one to apply Occam's razor to the problem of describing perceptual phenomena.

11.3 Explicit descriptions

The aim of computational analysis should be the development of formal models of sufficient explicitness, internal consistency, and completeness to enable an analytic characterisation or computer simulation. Hawkins *et al.* suggest that "ultimately, perhaps, the aim is the construction of synthetic circuits that emulate the performance characteristics of the biological system under analysis" [Hawkins *et al.* 96].

The presence of an explicit model in which the hypothesised perceptual mechanism is to account for system behaviour endows this approach with attributes analogous to scientific theories in other disciplines of science: the consistency of the assumptions may be verified; the adequacy of the underlying physical characterisation can be assessed; and these two lines of analysis can suggest extensions to the model, experiments to verify its predictions, and outline the limits of its applicability. In this way, the physical model can provide the same sort of guidance for investigation of perceptual processes as a physical theory gives in the study of physical processes. Furthermore, guidance derived from this approach may be applicable in more than one discipline [Hallam 87].

A related issue is that the way in which we describe a perceptual system fundamentally affects our understanding of it. The physical model offers some advantages in its insistence on a particular description [Hallam 87]:

> *"... a description of a perceptual system intended as a specification for a computer [or robot] program modelling that system must be complete, in the sense of there being no unspecified detail, and unambiguous. The computer [or robot] is an unforgiving and unimaginative servant, for which everything must be spelled out in full."*

Mechanisms proposed for computational investigation are no more likely to be right, but implementing them does force the explication of complete theories. Moreover, getting those mechanisms to work on an alternative platform reveals strengths and weaknesses in the hypotheses that are not always clear from verbal descriptions.

For example, in this work, implementation of an echo envelope phase comparison mechanism for locating oscillating targets highlighted the difficulty in extracting a reliable time waveform for binaural comparison. As each receiver has a different line-of-sight to the target, Doppler effects can vary the time waveforms in corresponding frequency channels of the right and left filter-banks. These effects may be subtle for a small head; however, temporal differences in the envelopes are also subtle. Binaural comparison mechanisms can be made more robust by exploiting the much more dramatic signal intensity differences in corresponding channels. This can be achieved by, for example, comparing summed energy or *via* temporal comparison processes operating on the output of a time-intensity trading filtering process. This strategy is paricularly well suited to a small echolocator, as inter-aural intensity measurements are largely range invariant for closely spaced receivers.

11.4 Separation of *Merkwelt*

As stated in the previous section, models may be right or wrong (*i.e.* they may or may not support hypotheses which account satisfactorily for the behaviour of the system they purport to explain); however, a computational model has a third possibility: it may be right, but irrelevant. In AI, this situation frequently arises when a particular physical model is too dependent upon its designer(s) to solve key parts of their perceptual problems. For example, when the agent under investigation lives in a world created by the investigator (*e.g.* the paper world of a mathematical model or a "habitat" simulated in a computer), it is difficult to ensure that the investigator has not projected some of her/his *Merkwelt* onto the system under investigation.

Ten years ago, Brooks (and others before him, *e.g.* see [Dennett 79]) identified this as a key problem in robotics [Brooks 86]:

> *"A person (even a young child) can make the right interpretation of a photo and suggest a plan of action. For AI planning systems however, the experimenter is required to abstract away most of the details to form a simple description in terms of atomic concepts such as PERSON, CHAIR, and BANANAS. ... But this abstraction is the essence of intelligence and the hard part of the problems being solved. ... The problems of recognition, spatial understanding, dealing with sensor noise, partial models, etc. are all ignored. The abstraction reduces the input data so that the program experiences the same perceptual world (Merkwelt) [von Uexküll 21] as humans."*

Remarkably, these problems are still with us today. Many AI models are based upon assumptions about sensing and actuation which are tractable, but overly simplistic. Models which do not realistically simulate sensors and actuators quickly become "irrelevant" because they either oversimplify a task (*e.g.* fail to take into account the impact of noise) or make overly difficult tasks for which details of the physics might be exploited. Webb claims: "If, as I have argued, the interaction with the environment determines the perceptual systems, then failing to deal with this interaction in model building means that one of the most fundamental problems of autonomous systems is simply not being addressed"[Webb 93].

I maintain, with Brooks, that "the best domain for trying to build a truly artificially intelligent system is a mobile robot wandering around an environment which has not been specially structured for it" [Brooks 86]. Experiments "with an actual robot ensure than an essence of reality is maintained and that no critical disabling problems have been ignored" [Brooks 89]. However, this is not to say that simulation has no value. Indeed, for some hypothesised mechanisms physical analogues cannot be built easily or without introducing too many side-effects of the particular hardware implementation. In this respect, a robotic implementation is "best" because it is an easier domain in which to maintain a rigorous conceptual separation between the agent and designer. Simulated models require that the modeller find some alternative means to establish the relationship between the model and the real world. It is not surprising to find that models which have succeeded in including realistic characterisations of the environment tend to be simulations of actual robots which incorporate experience of what is important in the real system.

In this work, the use of a physical model embodied in the environment enabled me to observe the hypothesised perceptual mechanisms under various conditions — even conditions which were not possible to formally analyse. Particularly in cases where phenomena are hard to describe formally, it may be more satisfactory to analyse mechanisms given the performance of the artificial system. Computer simulation and mathematically modelling of agent-environment interactions can not afford this perspective, as parts of the problem may be missed out when they cannot be adequately represented.

In echolocation studies, it is difficult to reason about the acoustical visibility of targets with complex geometries and motions. Furthermore, even when a desired target signature is well understood, reasoning about the effects of clutter and secondary reflectors on the localisability of a target is difficult. In the investigation reported here, the use of only a computer simulation model would weaken the justification for extending the results of testing the hypothesis to the bat. Indeed, in these terms, the simulation model used in this research was much more difficult to build than the robotic model. It was a worthwhile exercise; however, because it facilitated the testing of explanations of observed aspects of the physical problem. These explanations, in turn, could be fed back into the simulation model to increase realism.

11.5 Summary

Clearly an exchange between biology and robotics can be fruitful for one or both fields. However, the assumption that biology is full of mechanisms which might be directly implemented is clearly wrong (just as the use of engineering system concepts, *e.g.* the ideal receiver, provide only a limited heuristic for understanding systems with other purposes and constraints). The principal idea behind the methodology presented here is that the *process* of attempting to implement physical models of biological systems can potentially contribute to our *understanding* of how perceptual systems work [Webb 93].

12. Summary of Contributions

It was once believed (by certain Greeks thinkers) that we perceived objects in our environment because they sent off little copies of themselves in all directions, and that those replicas got into the brain through sensory channels [Gibson 66]. Present day thinkers must deny the possibility that the odour of the cheese *is* the cheese, or that the sound of the bell *is* the bell. However, if no quantity enters the ear except mere mechanical vibrations, then how is it that we detect objects?

What is currently understood about perception suggests that within the flux of signals arising from a particular source there are invariant cues which an agent detects and uses to identify and maintain a spatial relationship with those sources. We do not understand how invariants are extracted — indeed there is little known about the physiology of *sense information* pick-up, as compared with the catalogues of descriptions of receptor excitation by *sense stimuli* — but this kind of information pick-up is a fundamental aspect of intelligent behaviour.

Understanding intelligent behaviour in these terms is an increasingly important topic on the research agenda in AI. Many hope that, through modelling of animal-environment interactions, a useful dialogue between biology and AI will emerge — comparable to the interaction between psychology and AI which gave rise to Cognitive Science [Webb 93]. However, few research programmes present methodologies which aim to systematically integrate ideas and techniques from both fields in a mutually comprehensible and beneficial inter-disciplinary framework.

12.1 Methodological contributions

This thesis sought to provide a demonstration of how AI methods can be applied to understanding perception through a process akin to reverse engineering. I have followed what I refer to as a "bionics" approach to investigating perception which advocates that we characterise agent-environment interactions in terms of the physical principles governing the operation of an agent's sensors (or, equivalently, actuators) throughout the execution of a specific tasks. This is somewhat of a departure from other AI approaches which study perception in terms of internal representations and control architectures — both of which can lead researchers to make strong assumptions about the capabilities of their sensors and actuators. By contrast, an important principle underlying the bionics approach is that information is not constructed in the process of perception; but, rather, perceptual systems are structured, through development and evolution, so as to register directly the information inherent in the environment. Because it is difficult to identify the task and environment factors which act to constrain and simplify this process of sense-information extraction, building models of evolved perceptual mechanisms of particular animals is fruitful way to approach the problem (both for biology and AI).

It is by no means clear whether substantial work to create detailed models of specific sensory mechanisms will contribute to a more general understanding of perception. Our present understanding would suggest that animals host a plethora of specialised mechanisms — each appropriate to the task and environment of its user. In this regard, McFarland argues: "just as there are no general-purpose animals, so there should be no general-purpose robots" [McFarland 91]. I would agree that — at least for the time being — in-depth examination of the many and varied perceptual mechanisms used by animals is just as likely to lead to the identification of general perceptual principles as the direct development of general purpose sensor systems.

12.2 Practical contributions

The problems addressed in this thesis lay at a number of different levels. They involved the characterisation of an animal sensory mechanism in terms of the underlying task-physics; the construction of novel sensor and signal processing strategies to mimic aspects of the animal morphology, neurophysiology and behaviour; and the testing and evaluation of this sensor in an appropriate task. The results revealed new dimensions of information which have not (to my knowledge) been explored in a robotic engineering context. Moreover, they provide an insight into the acoustical world of bats which is currently not available *via* more established biological investigative methodologies.

"... on the suspect of a new sense ..."

The study of how bats orient themselves and catch prey dates back to 1794 when Lazzaro Spallanzani first realized that bats do not use vision [Galambos 42] (and announced his discovery in the quotation above). Spallanzani carried out many experiments to prove this, but was unable to suggest an alternative to sight. His result was lost for several centuries because of the prevailing view of science. "Spallanzani's bat problem" remained unsolved until 1938 when Donald Griffin detected ultrasonic transmissions from bats using a high frequency microphone developed by G. W. Pierce [Pierce & Griffin 38, Griffin 58]. It is interesting to note that it was an interdisciplinary contribution, namely the application (by a pair of physicists) of the wave theory of sound to a biological problem, which re-awoke the study of echolocation in bats.

"Bat problems" of great variety remain with us today. Moreover, the picture that is emerging from decades of research is that bats are amazingly diverse creatures and detailed investigations of particular relationships that bats have developed with their environments will reveal a host of perceptual strategies currently unimagined by roboticists and biologists.

Bibliography

[Achinstein 64] P. Achinstein. Models, analogies and theories. *British Jounral for the Philosophy of Science*, 62:102–119, 1964.

[Agre & Horswill 92] P. E. Agre and I. Horswill. Cultural support for improvisation. In *Proc. 10th Nat. Conf. on A. I.* Morgan Kaufmann, 1992.

[Agre 95a] P. E. Agre. Computation and embodied activity. *Informatica*, 19(4):527–535, 1995.

[Agre 95b] P. E. Agre. Computational research on interaction and agency. *Artificial Intelligence*, 72(1-2):1–52, 1995.

[Albus 81] J. S. Albus. *Brains, Behaviour and Robotics.* McGraw Hill, 1981.

[Altes 80] R. A. Altes. Models for echolocation. In R. G. Busnel and J. F. Fish, editors, *Animal SONAR Processes and Performance (NATO ASI Series)*, pages 625–671. Plenum Press, 1980.

[Anderson & Donath 90] T. Anderson and M. Donath. Animal behaviour as a paradigm for developing robot autonomy. *Robotics and Autonomous Systems*, 6, 1990.

[Andrews 95] P. T. Andrews. Rhinolophid acoustic orientation. *Myotis*, 32-33:81–90, 1995.

[Anke 74] D. Anke. Luftschallwandler nach dem Sell-prinzip fur Frequenzen von 50 khz bis 100 khz. *Acoustica*, 30:30–39, 1974.

[Arbib & Lee 94] M. A. Arbib and H. B. Lee. Anuran visuomotor coordination for detour behaviour: From retina to motor schemas. In J-A. Meyer, H. L. Roitblat, and S. W. Wilson, editors, *From animals to animats: Proc.*

Third Int. Conf. Simulation of Adaptive Behaviour, 1994.

[Arkin 91] R. C. Arkin. Integrating behavioural, perceptual, and world knowledge in reactive navigation. In P. Maes, editor, *Designing Autonomous Agents*. M. I. T. Press, 1991.

[Balcombe & Fenton 88] J. Balcombe and M. B. Fenton. Eavesdropping by bats, the influence of echolocation call design and foraging strategies. *Ethology*, 79:158–166, 1988.

[Ballard 87] D. H. Ballard. Eye movements and spatial cognition. Dept. C.S. Techonical Report 218, University of Rochester, 1987.

[Ballard 89] D. H. Ballard. Reference frames for animate vision. In *Proc. IJCAI*, pages 1635–1641, 1989.

[Barrett 87] J. S. Barrett. Principles of modern radar. In J.K. Eaves and E.K. Reedy, editors, *Adaptive thresholding and automatic detection*, pages 368–393. Van Nostrand Reinhold, 1987.

[Barshan & Kuc 90] B. Barshan and R. Kuc. Differentiating sonar reflections from corners and planes by employing an intelligent sensor. *P.A.M.I.*, 12(6):560–569, 1990.

[Bateman & Vaughan 74] G. Bateman and T. A. Vaughan. Nightly activities of mormoopid bats. *J. Mammal.*, 55:45–65, 1974.

[Batteau 67] D. Batteau. The role of pinna in human localization. *Proc. R. Soc. Lond. Ser. B*, 168:158–180, 1967.

[Beer 90] R. D. Beer. *Intelligence as Adaptive Behaviour*. Academic Press, 1990.

[Bell & Fenton 84] G. Bell and M. B. Fenton. The use of Doppler shifted echoes as a clutter rejection system: the echolocation and feeding behavior of *Hipposideros ruber* (Chiroptera: Hipposideridae). *Behav. Ecol. Sociobiol.*, 15:109–114, 1984.

[Beranek 93] L. Beranek. *Acoustics*. Acoustical Society of America, 1993.

[Berkowitz 65] R. S. Berkowitz. *Modern Radar: Analysis, Evaluation and System Design*. John Wiley and Sons, Inc., 1965.

[Beuter 80] K. Beuter. A new concept of echo evaluation in the auditory system of bats. In R. G. Busnel and J. F. Fish, editors, *Animal SONAR Systems (NATO ASI Series)*, pages 747–761. Plenum Press, 1980.

[Biber *et al.* 80] C. Biber, S. Ellin, E. Shenk, and J. Stempec. The polaroid ultrasonic ranging system. In *Proc. 67th Convention Audio Engineering Soc.*, 1980.

[Blitz 63] J. Blitz. *Fundamentals of Ultrasonics.* Butterworths, 1963.

[Bloom 77] P. J. Bloom. Determination of monaural sensitivity changes due to the pinna by use of minimum-audible-field measurements in the lateral vertical plane. *J. Acoust. Soc. Am.*, 61:820–828, 1977.

[Bodenhamer & Pollak 83] R. D. Bodenhamer and G. D. Pollak. Response characteristics of single units in the inferior colliculus of mustache bats to sinusoidally frequency modulated signals. *J. Comp. Physiol. A*, 153:67–79, 1983.

[Bodenhammer & Pollak 83] R. D. Bodenhammer and G. D. Pollak. Response characteristics of single units in the inferior colliculus of mustache bats to sinusoidally frequency modulated signals. *J. Comp. Physiol. A*, 153:67–80, 1983.

[Boring 42] E. G. Boring. *Sensation and Perception in the History of Experimental psychology.* Appleton-Century, 1942.

[Bourgine & Varela 92] P. Bourgine and F. J. Varela. Towards a practice of autonomous systems. In *Proc. First European Conf. Artificial Life.* M. I. T. Press, 1992.

[Bozma & Kuc 91] O. Bozma and R. Kuc. Characterizing pulses from rough surfaces using ultrasound. *J. Acoust. Soc. Amer.*, 89(6):2519–2531, 1991.

[Bozma & Kuc 94] O. Bozma and R. Kuc. A physical model-based analysis of heterogenous environments. *P.A.M.I.*, 16(5):497–506, 1994.

[Braitenberg 84] V. Braitenberg. *Vehicles: Experiments in Synthetic Psychology.* M. I. T. Press, 1984.

[Brooks 86] R. A. Brooks. Achieving artificial intelligence through building robots. A. I. Memo 899, M. I. T., 1986.

[Brooks 89] R. A. Brooks. A robot that walks; Emergent behaviours from a varefully evolved network. Dept. A. I. Memo 1091, M. I. T., 1989.

[Brooks 91a] R. A. Brooks. Challenges for complete creature architectures. In J. A. Meyer and S. W. Wilson, editors, *From animals to animats: Proceedings of the First Int. Conf. on Simulation of Adaptive Behaviour.* M. I. T. Press, 1991.

[Brooks 91b] R. A. Brooks. Elephants don't play chess. In P. Maes, editor, *Designing Autonomous Agents.* M. I. T. Press, 1991.

[Brosset 66] A. Brosset. *La biologie des chiropteres.* Masson et Vie, 1966.

[Brown 89] C. M Brown. Predictive gaze control. In *Proc. Fifth Alvet Vision Conf.*, pages 103–108, 1989.

[Bruns 76a] V. Bruns. Peripheral auditory tuning for fine frequency analysis by the CF-FM bat, *Rhinolophus ferrumequinum* I. Mechanical specializations of the cochlea. *J. Comp. Physiol. A*, 106:77–86, 1976.

[Bruns 76b] V. Bruns. Peripheral auditory tuning for fine frequency analysis by the CF-FM bat, *Rhinolophus ferrumequinum* II. Frequency mapping on the cochlea. *J. Comp. Physiol. A*, 106:87–97, 1976.

[Busnel & Fish 80] R. G. Busnel and J. F. Fish. *Animal SONAR Systems (NATO ASI Series).* Plenum Press, 1980.

[Camhi 84] J. M. Camhi. *Neuroethology: Nerve cells and the natural behaviour of animals.* Sinauer Associates Inc., 1984.

[Casseday & Covey 92] J. H. Casseday and E. Covey. Frequency tuning properties of neurons in the inferior colliculus of an FM bat. *J. Comp. Neurol.*, 1992.

[Casseday & Covey 96] J. H. Casseday and E. Covey. A neuroethological theory of the operation of the inferior colliculus. *Brain, Behavior and Evolution*, 47:311–336, 1996.

[Chen *et al.* 92] J. Chen, B. D. VanVeen, and K. E. Hecox. External ear transform modelling: A beamforming approach. *J. Acoust. Soc. Am.*, 92:1933–1944, 1992.

[Clancey 91] W. J. Clancey. Situated cognition: Stepping out of representational flatland. *AICOM*, 4:109–112, 1991.

[Clark 89] A. Clark. *Microcognition.* M. I. T. Press, 1989.

[Cliff 90] D. T. Cliff. Computational neuroethology: A provisional manifesto. In *Proc. SAB90*, 1990.

[Coles & Guppy 86] R. B. Coles and A. Guppy. Biophysical aspects of directional hearing in the tammar wallaby, *Macropus eugenii*. *J. Exp. Biol.*, 121:371–394, 1986.

[Covey 93] E. Covey. The monaural nuclei of the Lateral Lemniscus: Parallel pathways from cochlear nucleus to midbrain. In M. A. Merchan, J. M. Juiz, and D. A. Godfrey, editors, *The Mammalian Cochlear Nuclei: Organization and Function*, pages 321–334. Plenum Press, 1993.

[Covey *et al.* 91] E. Covey, M. Vater, and J. H. Casseday. Binaural properties of single units in the superior olivary complex of the mustache bat. *J. Neurophys.*, 66(3):1080–1090, 1991.

[Dennett 79] D. C. Dennett. *Brainstorms: Philosophical Essays on Mind and Psychology.* The Harvester Press, 1979.

[Drumheller 87] M. Drumheller. Mobile robot localization using sonar. *P.A.M.I.*, 9(2):325–332, 1987.

[Eckhorn *et al.* 88] R. Eckhorn, R. Bauer, W. Jordan, M. Brosch, W. Kruse, M. Munk, and H. J. P. Reitboeck. Coherent oscillations - a mechanism of feature linking in the visual-cortex - multiple electrode and correlation analyses in the cat. *Biol. Cyber.*, 60:121–130, 1988.

[Eisentraut 50] M. Eisentraut. Die Ernährung der Fledermäuse. *Zoologische*, 79:115–177, 1950.

[Evans 77] E. F. Evans. Recognition of complex acoustic signals. In T. H. Bullock, editor, *Peripheral processing of complex sounds.* Abakon Verlagsgesellschaft, 1977.

[Everett 89] H. Everett. Survey of collision avoidance and ranging sensors for mobile robotics. *Robotics and Autonomous Systems*, 5(1):5–67, 1989.

[Ewing 89] A. W. Ewing. *Arthropod Bioacoustics: Neurobiology and Behaviour.* Edinburgh University Press, 1989.

[Fay & Popper 96] R. Fay and A. Popper, editors. *Hearing By Bats.* Springer-Verlag, 1996.

[Feng & Vater 85] A. S. Feng and M. Vater. Functional organization of the cochlear nucleus of rufous horseshoe bats (*Rhinolophus rouxi*): Frequencies and internal connections are arranged in slabs. *J. Comp. Neurol.*, 235:529–555, 1985.

[Feng *et al.* 94] A. S. Feng, C. J. Condon, and K. R. White. Stroboscopic hearing as a mechanism for prey discrimination in frequency-modulated bats? *J. Acoust. Soc. Am.*, 95(5):2736–2744, 1994.

[Fenton 80] M. B. Fenton. Adaptiveness and ecology of echolocation in terrestrial (aerial) systems. In R. G. Busnel and J. F. Fish, editors, *Animal SONAR Systems (NATO ASI Series)*, pages 427–446. Plenum Press, 1980.

[Fenton 84] M. B. Fenton. Echolocation: implications for the ecology and evolution of bats. *Rev. Biol.*, 59:33–53, 1984.

[Fenton 96] M. B. Fenton. Natural history and bisonar signals. In R. Fay and A. Popper, editors, *Hearing By Bats*, pages 37–86. Springer-Verlag, 1996.

[Fletcher & Thwaites 79] N. H. Fletcher and S. Thwaites. Physical models for the analysis of acoustical systems in biology. *Quarterly Rev. of Biophysics*, 12:25–65, 1979.

[Fletcher & Thwaites 88] N. H. Fletcher and S. Thwaites. Obliquely truncated simple horns: Idealized models for vertebrate pinna. *Acoustica*, 65:1194–204, 1988.

[Flieger & Schnitzler 73] E. Flieger and H-U. Schnitzler. Ortungsleistungen der Fledermaus *Rhinolophus ferrumequinum* bei ein- und beidseitiger Ohrverstopfung. *J. Comp. Physiol. A*, 82:93–102, 1973.

[Flynn & Brooks 89] A. M. Flynn and R. A. Brooks. Battling reality. Dept. A. I. Memo 1148, M. I. T., 1989.

[Forster 91] P. Forster. A transputer-based autonomous mobile robot. Dept. A.I.,Technical Report 6, University of Edinburgh, 1991.

[Franceschini *et al.* 85] N. Franceschini, C. Blanes, and L. Oufar. Appareil de mesure passif et sans contact de la vitesse d'un objet. *Patent ANVAR (France) No. 51549*, 1985.

[Franceschini *et al.* 91] N. Franceschini, J-M. Pichon, and C. Blanes. Real time visuomotor control: from flies to robots. In *Proc. IEEE Fifth Int. Conf. Advanced Robotics*, 1991.

[Fuzessery & Pollak 85] Z. M. Fuzessery and G. D. Pollak. Determinants of sound location selectivity in the bat inferior colliculus: a combined dichotic and free-field stimulation study. *J. Neurophysiol.*, 54:757–781, 1985.

[Fuzessery 88] Z. M. Fuzessery. A mechanism for horizontal and vertical target localization in the mustache bat (*Pteronotus p. parnellii*). In P. E. Nachtigall and P. W. B. Moore, editors, *Animal SONAR Processes and Performance (NATO ASI Series)*, pages 471–476. Plenum Press, 1988.

[Fuzessery *et al.* 92] Z. M. Fuzessery, D. J. Hartley, and J. J. Wenstrup. Spatial processing within the mustache bat echolocating system: possible mechanisms for optimization. *J. Comp. Physiol. A*, 170:57–71, 1992.

[Galambos 42] R. Galambos. The avoidance of obstacles by flying bats: Spallanzani's ideas (1794) and later theories. *Isis*, 34:132, 1942.

[Gibson 66] J. J. Gibson. *The Senses Considered as Perceptual Systems.* Houghton Mifflin Company, 1966.

[Gibson 79] J. J. Gibson. *The ecological approach to visual perception.* Houghton Mifflin, 1979.

[Gill 65] T. Gill. *The Doppler Effect.* Academic Press, 1965.

[Gissoni 66] F. Gissoni. My cane is twenty feet long. *The New Outlook for the Blind*, 1966.

[Goldmann & Henson 77] L. Goldmann and O. W. Henson. Prey recognition and selection by the constant frequency bat, *Pteronotus p. parnellii. Behav. Ecol. Sociobiol.*, 2:411–419, 1977.

[Gorlinsky & Konstantinov 78] I. A. Gorlinsky and A. I. Konstantinov. Auditory localization of ultrasonic source by *Rhinolophus ferrumequinum.* In R. J. Olembo, J. B. Castelino, and F. A. Mutere, editors, *Proc. of the Fourth Int. Bat Research Conf.*, pages 145–153, 1978.

[Grey-Walter 53] W. Grey-Walter. *The Living Brain.* Duckworth, 1953.

[Griffin 58] D. R. Griffin. *Listening in the Dark.* Yale University Press, 1958.

[Griffin *et al.* 60] D.R. Griffin, F.A. Webster, and C. Michael. The echolocation of flying insects by bats. *Anim. Behav.*, 8:151–154, 1960.

[Griffin *et al.* 62] D. R. Griffin, D. Dunning, D. Cahlander, and F. A. Webster. Correlated orientation sounds and ear movements of horseshoe bats. *Nature*, 196:1185–1186, 1962.

[Grinnell & Grinnell 65] A. D. Grinnell and V. Grinnell. Neural correlates of vertical localization in echolocating bat. *J. Physiol.*, 181:830–851, 1965.

[Grinnell & Schnitzler 77] A. D. Grinnell and H-U. Schnitzler. Directional sensitivity of echolocation in the horseshoe bat, *Rhinolophus ferrumequinum.* II. Behavioural directinality of hearing. *J. Comp. Physiol. A*, 116:63–76, 1977.

[Grinnell 70] A. D. Grinnell. Comparative auditory neurophysiology of neotropical bats employing different echolocation signals. *Z. Vergl. Physiol.*, 68:117–153, 1970.

[Grinnell 96] A. D. Grinnell. Hearing in bats: An overview. In R. Fay and A. Popper, editors, *Hearing By Bats*, pages 2–36. Springer-Verlag, 1996.

[Grothe 94] B. Grothe. Interaction of excitation and inhibition in processing of pure tone and amplitude-modulated stimuli in the medial superior olive of the mustache bat. *J. Neurophys.*, 71(1):706–721, 1994.

[Guppy & Coles 88] A. Guppy and R. B. Coles. Acoustical aspects of hearing and echolocation in bats. In P. E. Nachtigall and P. W. B. Moore, editors, *Animal SONAR Processes and Performance (NATO ASI Series)*, pages 289–294. Plenum Press, 1988.

[Habersetzer & Vogler 83] J. Habersetzer and B. Vogler. Discrimination of surface structured targets by the echolocating bat *Myotis myotis* during flight. *J. Comp. Physiol. A*, 152:275–282, 1983.

[Hallam 84] J. C. T. Hallam. *Intelligent automatic interpretation of active marine sonar*. Unpublished PhD thesis, University of Edinburgh, 1984.

[Hallam 87] J. C. T. Hallam. Computational descriptions for interdisciplinary research in vision. In J.L. Casti and A. Karlqvist, editors, *Real Brains, Artificial Minds*, pages 107–133. Elsevier, 1987.

[Halme 97] A. Halme. Book review of sensors for mobile robots. *IEEE Trans. Robotics and Automation*, 12(6):922–923, 1997.

[Hammond *et al.* 95] K. J. Hammond, T. M. Converse, and J. W. Grass. The stabilization of environments. *Artificial Intelligence*, 73(1-2):305–327, 1995.

[Harnad 90] S. Harnad. The symbol grounding problem. *Physica D.*, 42:335–346, 1990.

[Harnischfeger *et al.* 85] G. Harnischfeger, G. Neuweiler, and P. Schlegel. Interaural time and intensity coding in superiour olivary complex and inferior colliculus of the echolocating bat, *Molossus ater*. *J. Neurophys.*, 53(1), 1985.

[Hartley & Suthers 87] D. J. Hartley and R. A. Suthers. The sound emission pattern and the acoustical role of the noseleaf in the echolocating bat, *Carollia perspicillata*. *J. Acoust. Soc. Am.*, 82(6):1892–1899, 1987.

[Hartley & Suthers 88] D. J. Hartley and R. A. Suthers. The vocal tract acoustics of the horseshoe bat, *Rhinolophus hildebrandti*. *J. Acoust. Soc. Am.*, 84:1201–1213, 1988.

[Hartley & Suthers 90] D. J. Hartley and R. A. Suthers. Sonar pulse radiation and filtering in the mustache bat, *Pteronotus parnellii rubiginosus*. *J. Acoust. Soc. Am.*, 87(6):2756–2772, 1990.

[Hawkins *et al.* 96] H. L. Hawkins, T. A. McMullen, A. N. Popper, and R. R. Fay. *Autitory Computation*. Springer-Verlag, 1996.

[Heidegger 61] M. Heidegger. *Being and Time*. Harper and Row, 1961. (Originally published in German in 1927).

[Helmholtz 25] H. Helmholtz. *Physiological Optics*, volume 3. Optical Soc. Am., 1925.

[Henning 74] G. B. Henning. Detectability of interaural delay in high frequency waveforms. *J. Acoust. Soc. Am.*, 55:84–90, 1974.

[Henson 65] O. W. Henson. The activity and function of the middle ear muscles in echolocating bats. *J. Physiol. Lond.*, 180:871–887, 1965.

[Henson 78] M. M. Henson. The basilar membrane of the bat *Pteronotus p. parnellii*. *Am. J. Anat.*, 153:143–154, 1978.

[Herrnstein 85] R. Herrnstein. Riddles of natural categorization. *Phil. Trans. R. Soc. Lond. B*, 308:129–144, 1985.

[Herrnstein *et al.* 76] R. Herrnstein, D. Loveland, and C. Cable. Natural concepts in pigeons. *J. Exp. Psychol.*, 2:285–302, 1976.

[Hoare 85] C. A. R. Hoare. *Communicating Sequential Processes*. Prentice-Hall Int., 1985.

[Horswill 95] I. Horswill. Analysis of adaptation and environment. *Artificial Intelligence*, 73(1-2):1–30, 1995.

[Howe 94] J. Howe. Artificial Intelligence at Edinburgh University : a perspective. W.W.W. Page, Dept. of A. I., University of Edinburgh, 1994.

[Hoyle 84] G. Hoyle. The scope of neuroethology. *BBS*, 7:367–412, 1984.

[Hughes 77] A. Hughes. The topology of vision in mammals of contrasting life style: comparative optics and retinal organization. In F. Crescitelli, editor, *Handbook of Sensory Physiology*, pages 613–756. Springer-Verlag, vii/5 edition, 1977.

[Irle & Markowitsch 87] E. Irle and H. Markowitsch. Conceptualization without specific training in squirrel monkeys (*Saimiri sciureus*): A test using the non-match-to-sample procedure. *J. Comp. Psychol.*, 101:305–311, 1987.

[Jen & Chen 88] P. Jen and D. Chen. Directionality of sound pressure transformation at the pinna of echolocating bats. *Hear. Res.*, 34:101–118, 1988.

[Johnson *et al.* 74] R. Johnson, O. W. Henson, and L. Goldman. Detection of insect wing beats by the bat *Pteronotus parnellii*. *J. Acoust. Soc. Am.*, 55:53, 1874.

[Jones & Rayner 89] G. Jones and J. M. V. Rayner. Foraging behavior and echolocation of wild horseshoe bats *Rhinolophus ferrumequinum* and *R. hipposideros* (Chiroptera, Rhinolophidae. *Behav. Ecol. Sociobiol.*, 25:183–191, 1989.

[Jones 90] G. Jones. Prey selection by the greater horseshoe bat: optimum foraging by echolocation? *J. Anim. Ecol.*, 59:587–602, 1990.

[Kao *et al.* 96] G. Kao, P. Probert, and D. Lee. Object recognition with FM SONAR: An assistive device for blind and visually-impared people. In *AAAI Fall Symposium on Developing Assistive Technologies for People with Disabilities*, 1996.

[Kay 61] L. Kay. Orientation of bats and men by ultrasonic echo location. *British Communications and Electronics*, 8, 1961.

[Kay 74] L. Kay. A sonar aid to enhance spatial perception of the blind: engineering design and evaluation. *The Radio and Electronic Engineer*, 44(11):605–627, 1974.

[Kay 80] L. Kay. Air sonars with acoustical display of spatial information. In R. G. Busnel and J. F. Fish, editors, *Animal SONAR Systems (NATO ASI Series)*. Plenum Press, 1980.

[Kay 82] R. H. Kay. Hearing of modulation in sound. *Physiological Reviews*, 62(3):894–975, 1982.

[Kick & Simmons 84] S. A. Kick and J. A. Simmons. Automatic gain control in the bat's sonar receiver and the neuroethology of echolocation. *J. Neurosci.*, 4:2725–2737, 1984.

[Kim 84] D. O. Kim. Functional roles of the inner- and outer-hair-cells subsystems in the cochlea and brainstem. In *Hearing Science*. C. Berlin College-Hill, 1984.

[Kinsler & Frey 62] L. Kinsler and A. Frey. *Fundamentals of Acoustics.* John Wiley and Sons, 2 edition, 1962.

[Kirsh 95] D. Kirsh. The intelligent use of space. *Artificial Intelligence*, 73(1-2):31–68, 1995.

[Kleeman & Kuc 95] L. Kleeman and R. Kuc. Mobile robot sonar for target localization and classification. *Int. J. Rob. Research*, 14(4):295–318, 1995.

[Kober & Schnitzler 90] R. Kober and H-U. Schnitzler. Information in sonar echoes of fluttering insects available for echolocating bats. *J. Acoust. Soc. Am.*, 87:882–896, 1990.

[Kober 82] R. Kober. *Analyse der Ultraschallechos von einheimischen Insekten.* Unpublished PhD thesis, Fakultät Biologie, Tübingen, 1982.

[Kober 88] R. Kober. Echoes of fluttering insects. In P. E. Nachtigall and P. W. B. Moore, editors, *Animal SONAR Processes and Performance (NATO ASI Series)*, pages 477–481. Plenum Press, 1988.

[Koch & Segev 89] C. Koch and I Segev. *Methods in Neuronal Modeling: From Synapses to Networks.* M.I.T. Press, 1989.

[Konstantinov *et al.* 73] A. I. Konstantinov, I. M. Stosman, and I. A. Gorlinsky. Characteristics of the directionality of reception and precision of localisation of ultrasound by echolocating bats. *Ref. Dokl. IV Vses. Konf. po. Bionike*, 4:5156, 1973.

[Kössl & Vater 85] M. Kössl and M. Vater. Evoked acoustic emissions and cochlear microphonics in the mustache bat, *Pteronotus parnellii. Hear. Res.*, 19:157–170, 1985.

[Kössl & Vater 90] M. Kössl and M. Vater. Resonance phenomena in the cochlea of the mustache bat and their contribution to neural reponse characteristics in the cochlear nucleus. *J. Comp. Physiol. A*, 166:711–720, 1990.

[Kössl & Vater 96] M. Kössl and M. Vater. Cochlear structure and function in bats. In R. Fay and A. Popper, editors, *Hearing By Bats*, pages 191–234. Springer-Verlag, 1996.

[Kuc & Siegel 87] R. Kuc and M. Siegel. A physically based simulation model for acoustic sensor robot navigation. *P.A.M.I.*, 9(6):766–778, 1987.

[Kuc & Viard 91] R. Kuc and V. B. Viard. A physically based navigation strategy for sonar guided vehicles. *Int. Journal of Robotics Research*, 10(2):75–87, 1991.

[Kuc 93] R. Kuc. Three-dimensional tracking using qualitative bionic sonar. *Robotics and Autonomous Systems*, 11:213–219, 1993.

[Kuc 94] R. Kuc. Sensorimotor model of bat echolocation and prey capture. *J. Acoust. Soc. Am.*, 96(4):1965–1978, 1994.

[Kuc 96] R. Kuc. Biologically motivated adaptive sonar system. *J. Acoust. Soc. Am.*, 100(3):1849–1854, 1996.

[Kunz & Pierson 94] T. H. Kunz and E. D. Pierson. An introduction. In *Walker's Bats of the World*. Johns Hopkins University Press, 1994.

[Kurz 93] A. Kurz. Building maps based on a learned classification of ultrasonic range data. In *Proc. IFAC Workshop Intelligent Autonomous Vehicles*, 1993.

[Langton 97] C. G. Langton, editor. *Artificial Life*. M.I.T. Press, 1997.

[Lashley 51] K. S. Lashley. The problem of serial order in behaviour. In L. A. Jeffress, editor, *Cerebral Mechanisms in Behaviour: The Hixon Symposium*. Wiley, 1951.

[Lawrence & Simmons 82] B. Lawrence and J. A. Simmons. Echolocation in bats: The external ear and perception of the vertical position of targets. *Science*, 218:481–483, 1982.

[Lee 74] D. N. Lee. Visual information during locomotion. In *Perception: Essays in Honor of James J. Gibson*, pages 251–267. Cornell University Press, 1974.

[Lee 80] D. N. Lee. The optic flow field: the foundation of vision. *Phil. Trans. R. Soc. Lond. B*, 290:169–179, 1980.

[Lee 94] D. N. Lee. Perception and motor control in birds: An ecological approach. In *An eye or ear for flying*, pages 270–291. Springer-Verlag, 1994.

[Lee *et al.* 88] C. Lee, W. H. Rohrer, and D. L. Sparks. Population coding of saccadic eye-movements by neurons in the superior colliculus. *Nature*, 332:357–360, 1988.

[Lee *et al.* 92] D. N. Lee, F. R. van der Weel, T. Hitchcock, E. Matejowsky, and J. D. Pettigrew. Common principle of guidance by echolocation and vision. *J. Comp. Physiol. A*, 171:563–571, 1992.

[Lenoir *et al.* 87] M. Lenoir, J-L. Puel, and R. Pujol. Stereocilia and techorial membrane development in the rat cochlea. *Anat. Embryol.*, 175:477–487, 1987.

[Leonard & Durrant-Whyte 91] J. Leonard and H. F. Durrant-Whyte. Mobile robot localization by tracking geometric beacons. *IEEE Trans. Robotics and Automation*, 7(3):376–382, 1991.

[Leonard 90] J. Leonard. *Directed sonar sensing for mobile robot navigation.* Unpublished PhD thesis, University of Oxford, 1990.

[Lesser *et al.* 90] H. D. Lesser, W.E. O'Neil, R. D. Frisina, and R. C. Emerson. On-off units in the mustache bat inferior colliculus are selective for transients resembling "acoustic glint" from fluttering insect targets. *Exp. Brain Res.*, 82:137–148, 1990.

[Lim 86] D. J. Lim. Functional structure of the organ of corti: a review. *Hear. Res.*, 22:117–146, 1986.

[Lindstedt & Olsson 93] G. Lindstedt and G. Olsson. Using ultrasonics for sensing in a robotic environment. In *Proc. IEEE Conf. Robotics and Automation*, pages 671–676, 1993.

[Link *et al.* 86] A. Link, G. Marimuthu, and G. Neuweiler. Movement as a specific stimulus for prey catching behaviour in rhinolophid and hipposiderid bats. *J. Comp. Physiol. A*, 159:403–413, 1986.

[Long & Schnitzler 75] G. Long and H-U. Schnitzler. Behavioural audiograms from the bat *Rhinolophus ferrumequinum. J. Comp. Physiol. A*, 100:211–219, 1975.

[Loughlin & Tacer 95] P. J. Loughlin and B. Tacer. On the AM-FM decomposition of signals. Dept. E.E. Techonical Report, Univeristy of Pittsburgh, 1995.

[Lyon & Mead 88] R. F. Lyon and C. Mead. An analog electronic cochlea. *IEEE Trans. Acoustics, Speech and Signal Processing*, 36(7):1119–1134, 1988.

[Lyons 85] D. Lyons. *Robot Schemas.* Unpublished PhD thesis, University of Massachussetts, 1985.

[Maes 89] P. Maes. The dynamics of action selection. In *Proc. IJCAI*, pages 991–997, 1989.

[Maes *et al.* 96] P. Maes, M. J. Mataric, J-A. Meyer, J. Pollack, and S. W. Wilson, editors. *From animals to animats: Proc. Forth Int. Conf. Simulation of Adaptive Behaviour.* M.I.T. Press, 1996.

[Malcolm 90] C. Malcolm. Behavioural modules in robotic assembly. Dept. A. I. Research Report, University of Edinburgh, 1990.

[Malcolm 97] C. Malcolm. A hybrid behavioural/knowledge-based approach to robotic assembly. In *Proc. Evolutionary Robotics*, pages 221–256, 1997.

[Malcolm *et al.* 89] C. Malcolm, T. Smithers, and J. C. T. Hallam. An emerging paradigm in robot architecture. Dept. A. I. Research Report 447, University of Edinburgh, 1989.

[Manyika & Durrant-Whyte 93] J. M. Manyika and H. F. Durrant-Whyte. A tracking sonar system for vehicle guidance. In *Proc. of IEEE Conf. Robotics and Automation*, pages 424–429, 1993.

[Marr 82] D. Marr. *Vision*. Freeman, 1982.

[Mason 86] M. T. Mason. Mechanics and planning of manipulator pushing operations. *Int. J. Rob. Res.*, 5(3):53–71, 1986.

[Mataric 90] M. J. Mataric. Environment learning using a distributed represenation. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 402–406, 1990.

[Maturana & Varela 80] H. R. Maturana and F. J. Varela. *Autopoesis and cognition*. Reidel, 1980.

[McCarty & Jen 83] J. McCarty and P. Jen. Bats reject clutter interference for moving targets more successfully than for stationary ones. *J. Comp. Physiol.*, pages 447–454, 1983.

[McFarland 91] D. McFarland. What it means for robot behaviour to be adaptive. In J. A. Meyer and S. W. Wilson, editors, *From animals to animats: Proceedings of the First Int. Conf. on Simulation of Adaptive Behaviour*. M. I. T. Press, 1991.

[McKerrow & Hallam 90] P. J. McKerrow and J. C. T. Hallam. An introduction to the physics of echolocation. *Third National Conf. Robotics*, pages 198–209, 1990.

[Merleau-Ponty 62] M. Merleau-Ponty. *Phenomenology of Perception*. Humanities Press, 1962. (Translated from French by Colin Smith).

[Meyer 97] J-A. Meyer, editor. *Adaptive Behaviour*. M.I.T. Press, 1997.

[Miller *et al.* 60] G. A. Miller, E. Galanter, and K. H. Pribram. *Plans and the Structure of Behaviour*. Holt, 1960.

[Mills 60] A. W. Mills. Lateralization of high-frequency tones. *J. Acoust. Soc. Am.*, 32(1):132–135, 1960.

[Mogdans *et al.* 88] J. Mogdans, J. Ostwald, and H-U. Schnitzler. The role of pinna movement for the localization of vertical and horizontal wire obstacles in the greater horseshoe bat, *Rhinolophus ferrumequinum*. *J. Acoust. Soc. Am.*, 84(5):1676 – 1679, 1988.

[Mogdans *et al.* 93] J. Mogdans, H-U. Schnitzler, and J. Ostwald. Discrimination of 2-wavefront echoes by the big brown bat, *Eptesicus fuscus*: Behavioural experiments and receiver simulations. *J. Comp. Physiol. A*, 84:309–323, 1993.

[Möhres 53] F. Möhres. Uber die Ultraschallorientierung der Hufeisennasen. *Z. Vergl. Physiol*, 34:547–588, 1953.

[Möller *et al.* 78] J. Möller, G. Neuweiler, and H. Zoeller. Response characteristics of inferior colliculus neurons of the awake CF-FM bat, *Rhinolophus ferrumequinum* I. Single tone stimuli. *J. Comp. Physiol. A*, 125:217–225, 1978.

[Moravec & Elfes 85] H. P. Moravec and A. Elfes. High resolution maps from wide angle sonar. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 116–121, 1985.

[Moravec 66] H. Moravec, editor. *Computer Simulation Techniques*. John Wiley and Sons, 1966.

[Moravec 84] H. Moravec. Locomotion, vision and intelligence. *Robotics Research*, 1984.

[Morse & Ingrad 68] P. Morse and K. Ingrad. *Theoretical Acoustics*. McGraw-Hill Book Company, 1968.

[Moss & Casseday 96] E. Moss and J. H. Casseday. The lower brainstep auditory pathways. In R. Fay and A. Popper, editors, *Hearing By Bats*, pages 235–295. Springer-Verlag, 1996.

[Moss & Schnitzler 96] C. F. Moss and H-U. Schnitzler. Behavioural studies of auditory information processing. In R. Fay and A. Popper, editors, *Hearing By Bats*, pages 87–145. Springer-Verlag, 1996.

[Moss & Zagaeski 94] C. F. Moss and M. Zagaeski. Acoustic information available to bats using frequency-modulated echolocation sounds for the perception of insect prey. *J. Acoust. Soc. Am.*, 95:2745–2756, 1994.

[Moss *et al.* 92] C. F. Moss, C. Gounden, J. Booms, and J. Roach. Discrimination of target movement by the FM bat, *Eptesicus fuscus*. *Abstracts of the 15th Midwinter Research Meeting of the Society for Research in Otolaryngology*, page 142, 1992.

[Müller & Schnitzler 97] R. Müller and H-U. Schnitzler. Acoustic flow in echo amplitudes and spectra: a viable concept for obstacle avoidance? *J. Acoust. Soc. Am.*, 101(5):3137, 1997.

[Nachtigall & Moore 88] P. E. Nachtigall and P. W. B. Moore. *Animal SONAR Processes and Performance (NATO ASI Series)*. Plenum Press, 1988.

[Nachtigall 72] W. Nachtigall. Der Flug. In A. Kaestner, I. Band, and A. Teil, editors, *Lehrbuch der speziellen Zoologie*, pages 51–65. Fischer Verlag, 1972.

[Nehmzow *et al.* 90] U. Nehmzow, T. Smithers, and J. C. T. Hallam. Steps towards intelligent robots. Dept. A. I. Research Paper 502, University of Edinburgh, 1990.

[Nehmzow *et al.* 91] U. Nehmzow, T. Smithers, and J. C. T. Hallam. Location recognition in a mobile robot using self-organising feature maps. Dept. A. I. Research Paper 520, University of Edinburgh, 1991.

[Neuweiler & Vater 77] G. Neuweiler and M. Vater. Response patterns to pure tones of cochlear nucleus units in the CF-FM bats, *Rhinolophus ferrumequinum*. *J. Comp. Physiol. A*, 115:119–133, 1977.

[Neuweiler 70] G. Neuweiler. Neurophysiologische Untersuchungen zum Echoortungsystem der Grossen Hufeisennase *Rhinolophus ferrumequinum*. *Z. Vergl. Physiol*, 67:273–306, 1970.

[Neuweiler 80] G. Neuweiler. Auditory processing of echoes: Peripheral processing. In R. G. Busnel and J. F. Fish, editors, *Animal SONAR Systems (NATO ASI Series)*, pages 519–548. Plenum Press, 1980.

[Neuweiler 83] G. Neuweiler. Echolocation and adaptivity to ecological constraints. In F. Hubel and H. Markl, editors, *Neuroethology and Behavioural Psychology*, pages 280–302. Springer-Verlag, 1983.

[Neuweiler 90] G. Neuweiler. Auditory adaptations for prey capture in echolocating bats. *Physiol. Rev.*, 70(3):615–639, 1990.

[Neuweiler *et al.* 80] G. Neuweiler, V. Bruns, and G. Schuller. Ears adapted for the detection of motion, or how echolocating bats have exploited the capacities of the mammalian auditory system. *J. Acoust. Soc. Am.*, 68(3):741–753, 1980.

[Neuweiler *et al.* 87] G. Neuweiler, W. Metzner, U. Heilmann, R. Rübsamen M. Eckrich, and H. Costa. Foraging behaviour and echolocation in the rufous horseshoe bat *Rhinolophus rouxi*, of Sri Lanka. *Behav. Ecol. Sociobiol.*, 20:53–67, 1987.

[Neuweiler *et al.* 88] G. Neuweiler, A. Link, G. Marimuthu, and R. Rübsamen. Detection of prey in echocluttering environments. In P. E. Nachtigall and P. W. B. Moore, editors, *Animal SONAR Processes and Performance (NATO ASI Series)*, pages 613–617. Plenum Press, 1988.

[Nevins & Whitney 78] J. L. Nevins and D. E. Whitney. Computer controlled assembly. *Sci. Am.*, 1978.

[Newell & Simon 63] A. Newell and H. A. Simon. GPS:A program that simulates human thought. In E. A. Feigenbaum and J. Feldman, editors, *Computers and Thought*, pages 279–296. Mc-Graw Hill, 1963.

[Nitsche 87] V. Nitsche. *Das Unterscheidungsvermögen für Schlagfrequenzen mechanisch stimulierter Flügel in störungsfreier und akustisch gestörter Umgebung bei der FM-CF-FM Fledermaus Rhinolophus ferrumequinum.* Unpublished PhD thesis, Fakultät Biologie, München, 1987.

[Novick 77] A. Novick. Acoustic orientation. In W. A. Wimsatt, editor, *Biology of Bats*, volume 3, pages 73–287. Academic Press, 1977.

[Nowak 94] R. M. Nowak. *Walker's Bats of the World.* Johns Hopkins University Press, 1994.

[Obrist *et al.* 93] M. Obrist, M. B. Fenton, J. Eger, and P. Schlegel. What ears do for bats: A comparative study of pinna sound pressure transformation in chiroptera. *J. Exp. Biol.*, 180:119–152, 1993.

[O'Neill 96] W. E. O'Neill. The bat auditory cortex. In R. Fay and A. Popper, editors, *Hearing By Bats*, pages 416–480. Springer-Verlag, 1996.

[Ostwald 84] J. Ostwald. Tonotopical organization and pure tone response characteristics of single units in the auditory cortex of the greater horseshoe bat. *J. Comp. Physiol.*, 155:821–834, 1984.

[Ostwald 88] J. Ostwald. Encoding of natural insect echoes and sinusoidally modulated stimuli by neurons in the auditory cortex of the greater horseshoe bat, *Rhinolophus ferrumequinum.* In *Animal SONAR Processes and Performance (NATO ASI Series)*, pages 483–487. Plenum Press, 1988.

[Ostwald *et al.* 88] J. Ostwald, H-U. Schnitzler, and G. Schuller. Target discrimination and target classification in echolocating bats. In P. E. Nachtigall and P. W. B. Moore, editors, *Animal SONAR Processes and Performance (NATO ASI Series)*, pages 413–434. Plenum Press, 1988.

[Panter 65] P. Panter. *Modulation, Noise and Spectral Analysis: Applied to Information transmission.* McGraw-Hill Book Company, 1965.

[Papadopoulos 97] G. Papadopoulos. Evolving ears for echolocation. Unpublished M.Sc. thesis, University of Edinburgh, 1997.

[Peremans & Hallam 98] H. Peremans and J. C. T. Hallam. The spectrogram correlation and transformation receiver, revisited. *J. Acoust. Soc. Am. (accepted)*, 1998.

[Peremans 94] H. Peremans. *Tri-aural perception for mobile robots.* Unpublished PhD thesis, Universiteit Gent, 1994.

[Peremans *et al.* 97] H. Peremans, V. A. Walker, and J. C. T. Hallam. A biologically inspired sonarhead. Dept. A. I. Technical Report 44, University of Edinburgh, 1997.

[Peremans *et al.* 98a] H. Peremans, V. A. Walker, and J. C. T. Hallam. 3d object localisation with a binaural sonarhead, inspirations from biology. *IEEE Conf. Robotics and Automation (in press)*, 1998.

[Peremans *et al.* 98b] H. Peremans, V.A. Walker, G. Papadopoulos, and J.C.T. Hallam. Evolving batlike pinnae for target localisation by an echolocator. In *Proc. 2nd Int. Conf. Evolvable Systems (submitted)*, 1998.

[Pfeifer & Verschure 92] R. Pfeifer and P. Verschure. The challange of autonomous agents: pitfalls and how to avoid them. In *Proc. of the Workshop on Emergence, Situatedness, Subsumption and Symbol Grounding*, 1992.

[Pichon *et al.* 89] J-M. Pichon, C. Blanes, and N. Franceschini. Visual guidance of a mobile robot equipped with a network of self-motion sensors. In *SPIE Mobile Robots IV*, volume 1195, 1989.

[Pierce & Griffin 38] G. W. Pierce and D. R. Griffin. Experimental determination of supersonic notes emitted by bats. *J. Mammalogy*, 19, 1938.

[Pollak & Bodenhammer 81] G. D. Pollak and R. D. Bodenhammer. Specialized characteristics of single units in the inferior colliculus of the mustache bat: frequency representation, tuning and discharge patterns. *J. Neurophys.*, 45:605–620, 1981.

[Pollak & Casseday 89] G. D. Pollak and J. H. Casseday. *The Neural basis of Echolocation in Bats*. Springer-Verlag, 1989.

[Pollak & Park 96] G. D. Pollak and T. J. Park. The inferior colliculus. In R. Fay and A. Popper, editors, *Hearing By Bats*, pages 296–367. Springer-Verlag, 1996.

[Pollak & Schuller 81] G. D. Pollak and G. Schuller. Tonotopic organization and encoding features of single units in the inferior colliculus of horseshoe bats: Functional implications for prey identification. *J. Neurophysiol.*, 45:208–226, 1981.

[Pollak 88] G. D. Pollak. Time is traded for intensity in the bat's auditory system. *Hearing Research*, 36:107–124, 1988.

[Pollak *et al.* 86] G. D. Pollak, J. J. Wenstrup, and Z. M. Fuzessery. Auditory processing in the mustache bat's inferior colliculus. *TINS*, pages 556–561, 1986.

[Pringle 57] J. W. S. Pringle. *Insect Flight*. Cambridge University Press, 1957.

[Probert 94] P. Probert. Low cost range sensor for reactive navigation. In S. Cameron and P. Probert, editors, *Advanced Guided Vehicles: Aspects of the Oxford AGV Project*, pages 61–83. World Scientific, 1994.

[Pye & Roberts 70] J. D. Pye and L. Roberts. Ear movements in a hipposiderid bat. *Nature*, 225:285–286, 1970.

[Pye 67] J. D. Pye. Theories of sonar systems and their application to biological organisms (discussion). In *Animal SONAR Systems*, pages 1121–1136. Lab. Physiol. Acoust. CNRS Jouy-en-Josas, 1967.

[Pye 80] J. D. Pye. Adaptiveness of echolocation signals in bats, flexibility in behaviour and evolution. *Trends Neurosci.*, pages 232–235, October 1980.

[Pye 84] J. D. Pye. The bionic bat: Evolution and wave-physics. *Proc. Roy. Instrn. Gt. Brtn.*, 56:175–203, 1984.

[Pye *et al.* 62] J. D. Pye, M. Flinn, and A. Pye. Correlated orientation sounds and ear movements of horseshoe bats. *Nature*, 196:1186–1188, 1962.

[Raibert 86] M. H. Raibert. Running with symmetry. *Int. Journal of Robotics Research*, 5:3–19, 1986.

[Raibert 89] M. H. Raibert. Dynamically stable legged locomotion. Dept. A. I. Technical Report 1179, M. I. T., 1989.

[Rasnow *et al.* 88] B. Rasnow, C. Assad, M. E. Nelson, and J. M. Bower. Simulation and measurement of the electric fields generated by weakly electric fish. In D. S. Touretzky, editor, *Advances in Neural Information Processing*. Morgan Kaufman, 1988.

[Rayner 91] J. Rayner. Complexity and a coupled system: flight, echolocation and evolution in bats. In N. Schmidt-Kittler and K. Vogel, editors, *Constructional Morphology and Evolution*, pages 173–191. Springer-Verlag, 1991.

[Reid 65] T. Reid. *Essays on the Intellectual Powers of Man.* M.I.T. Press, 1965. Originally published in 1785.

[Robinson 96] M. F. Robinson. A relationship between echolocation calls and noseleaf widths in bats of the genera *Rhinolophus* and *Hipposideros*. *J. Zool. Lond.*, 239:389–393, 1996.

[Roeder 63] K. D. Roeder. Echoes of ultrasonic pulses from flying moths. *Biol. Bull.*, 124:200–210, 1963.

[Roeder 70] K. D. Roeder. Episodes in insect brains. *Am. Sci.*, 58:378–389, 1970.

[Roffler & Butler 68] S. K. Roffler and R. A. Butler. Localization of tonal stimuli in the vertical plane. *J. Acoust. Soc. Am.*, 51:1260–1266, 1968.

[Rosenblueth 70] A. Rosenblueth. *Philosophy of Science.* M.I.T. Press, 1970.

[Roverud *et al.* 91] R. Roverud, V. Nitsche, and G. Neuweiler. Discrimination of wingbeat motion by bats correlated with echolocation sound pattern. *J. Comp. Physiol. A*, 168:259–263, 1991.

[Sabatini 92] A. M. Sabatini. Active hearing for external imaging batsed on an ultrasonic transducer array. In *Proc. IROS*, pages 829–836, 1992.

[Saillant *et al.* 93] P. A. Saillant, J. A. Simmons, and S. A. Dear. A computational model of echo processing and acoustic imagin ing FM echolocating bats: The SCAT receiver. *J. Acoust. Soc. Am.*, 94(5):2691–2712, 1993.

[Samuel 83] K. Samuel. 11th R. A. Fisher Memorial Lecture. Hosted by the Royal Society, 20, April 1983.

[Schildberger 85] K. Schildberger. Recognition of temporal patterns by identified auditory neurons in the cricket brain. In K. Kalmring and N. Elsner, editors, *Acoustic and Vibrational Communication in Insects*. Verlag, 1985.

[Schlegel 77] P. Schlegel. Directional coding by binaural brainstem units of the CF-FM bat, *Rhinolophus ferrumequinum*. *J. Comp. Physiol. A*, 118:327–352, 1977.

[Schlegel 80] P. Schlegel. Single brain stem unit responses to binaural stimuli simulating moving sounds in *Rhinolophus ferrumequinum*. In R. G. Busnel and J. F. Fish, editors, *Animal SONAR Systems (NATO ASI Series)*, pages 973–975. Plenum Press, 1980.

[Schmidt 88] S. Schmidt. Discrimination of target surface structure in the echolocating bat, *Megaderma lyra*. In P. E. Nachtigall and P. W. B. Moore, editors, *Animal SONAR Processes and Performance (NATO ASI Series)*, pages 507–512. Plenum Press, 1988.

[Schmidt 92] S. Schmidt. Perception of structured phantom targets in the echolocating bat, *Megaderma lyra*. *J. Acoust. Soc. Am.*, 91:2203–2223, 1992.

[Schneider & Möhres 60] H. Schneider and F. Möhres. Die Ohrbewegungen der Hufeisenfledermäuse (Chiroptera, Rhinolophidae) und der Mechanismus des Bildhörens. *Z. Vergl. Physiol.*, 44:1–40, 1960.

[Schnitlzer & Grinnell 77] H-U. Schnitlzer and A. D. Grinnell. Directional sensitivity of echolocation in the horseshoe bat, *Rhinolophus ferrumequinum*. I. Directionality of sound emission. *J. Comp. Physiol. A*, 116:51–61, 1977.

[Schnitzler & Flieger 83] H-U. Schnitzler and E. Flieger. Detection of oscillating target movements by echolocation in the greater horseshoe bat. *J. Comp. Physiol. A.*, 135:385–392, 1983.

[Schnitzler & Henson 80] H-U. Schnitzler and O. W. Henson. Performance of airborne animal sonar systems: I. Microchiroptera. In R. G. Busnel and J. F. Fish, editors, *Animal SONAR Systems (NATO ASI Series)*, pages 109–181. Plenum Press, 1980.

[Schnitzler & Ostwald 83] H-U. Schnitzler and J. Ostwald. Adaptation for the detection of fluttering insects by echolocation in horseshoe bats. In J. Ewert, R. Capranica, and D. Ingle, editors, *Advances in Vertebrate Neuroethology*, pages 801–827. Plenum Press, 1983.

[Schnitzler 68] H-U. Schnitzler. Die Ultraschall-ortungslaute der Hufeisen-fledermäuse (Chiroptera-Rhinolophidae) in verschiedenen orientierungssituationen. *Z. Verlag. Physiol.*, 57:376–408, 1968.

[Schnitzler 70] H-U. Schnitzler. Echoortung bei der Fledermaus *chilonycteris rubiginosa*. *Z. Vergl. Physiol.*, 68:25–39, 1970.

[Schnitzler 73] H-U. Schnitzler. Die Echoortung der Fledermäuse und ihre hörphysiologischen Grundlagen. *Fortschr. Zool.*, 21:136–189, 1973.

[Schnitzler 78] H-U. Schnitzler. Die Detektion von Bewegungen durch Echoortung bei Fledermaüsen. *Verh. Dtsch. Zool. Ges.*, pages 16–33, 1978.

[Schnitzler *et al.* 83] H-U. Schnitzler, D. Menne, R. Kober, and K. Heblich. The acoustical image of fluttering insects in echolocating bats. In F. Huber and H. Markl, editors, *Neurophysiology and Behavioural Physiology*, pages 235–249. Springer-Verlag, 1983.

[Schnitzler *et al.* 94] H-U. Schnitzler, E. Kalko, I. Kaipf, and A. D. Grinnell. Fishing annd echolocation behaviour of the greater bulldog bat Noctilio Leporinus, in the field. *Behav. Ecol. Sociobiol.*, 35:327–345, 1994.

[Schuller & Pollak 79] G. Schuller and G. D. Pollak. Disproportionate frequency representation in the inferior colliculus of horsehoe bats: evidence for an "acoustic fovea". *J. Comp. Physiol. A*, 132:47–54, 1979.

[Schuller 72] G. Schuller. Echoortung bei *Rhinolophus ferrumequinum* mit frequenzmodulierten Lauten. *J. Comp. Physiol. A*, 77:306–331, 1972.

[Schuller 77] G. Schuller. Echo delay and overlap with emitted orientation sounds and Doppler-shift compensation in the bat, *Rhinolophus ferrumequinum. J. Comp. Physiol. A.*, 114:103–114, 1977.

[Schuller 79] G. Schuller. Coding of small sinusoidal frequency and amplitude modulations in the inferior colliculus of the CF-FM bat, *Rhinolophus ferrumequinum. Exp. Brain Res.*, 34:117–132, 1979.

[Schuller 80] G. Schuller. Hearing characteristics and Doppler shift compensation in South Indian CF-FM bats. *J. Comp. Physiol. A*, 139:349–356, 1980.

[Schuller 84] G. Schuller. Natural ultrasonic echoes from wingbeating insects are encoded by collicular neurons in the CF-FM bat,*Rhinolophus ferrumequinum. J. Comp. Physiol. A*, 155:121–128, 1984.

[Schuller *et al.* 74] G. Schuller, K. Beutler, and H-U. Schnitzler. Response to frequency shifted artificial echoes in the bat *Rhinolophus ferrumequinum. J. Comp. Physiol. A.*, 89:275–286, 1974.

[Schuller *et al.* 75] G. Schuller, K. Beuter, and R. Rübsamen. Dynamic properties of the compensation system of Doppler shifts in the bat, *Rhinolophus ferrumequinum. J. Comp. Physiol. A.*, 97:113–125, 1975.

[Schweizer 81] H. Schweizer. The connections of the inferior colliculus and the organization of the brainstem auditory system in the greater horseshoe bat (*Rhinolophus ferrumequinum*). *J. Comp. Physiol. A*, 201:25–49, 1981.

[Seifert 75] G. Seifert. *Entomologisches Praktikum.* Georg Thieme Verlag, 1975.

[Sejnowski *et al.* 88] T. J. Sejnowski, C. Koch, and P. S. Churchland. Computational neuroscience. *Science*, 241:1299–1306, 1988.

[Shamma *et al.* 86] S. A. Shamma, R. Chadwick, J. Wilbur, J. Rinzel, and K. Moorish. A biophysical model of cochlear processing: Intensity dependence of pure tone responses. *J. Acoust. Soc. Am.*, 80:133–145, 1986.

[Shaw 74] E. Shaw. The external ear. In W. Keidel and W. Neff, editors, *Handbook of Sensory Physiology.* Springer-Verlag, 1974.

[Shimozawa *et al.* 74] T. Shimozawa, N. Suga, P. Hendler, and S. Schuetze. Directional sensitivity of echolocation system in bats producing frequency modulated signals. *J. Exp. Biol.*, 60:53–69, 1974.

[Siebert 68] W. M. Siebert. Recognising patterns. In *Stimulus transformations in the peripheral auditory system.* M.I.T. Press, 1968.

[Simmons & Stein 80] J. A. Simmons and R. Stein. Acoustic imaging in bat sonar: Echolocation signals and the evolution of echolocation. *J. Comp. Physiol. A*, 135:61–84, 1980.

[Simmons 82] J. A. Simmons. The external ears as receiving antennae in echolocating bats. *J. Acoust. Soc. Am.*, 72:41–42, 1982.

[Simmons *et al.* 74] J.A. Simmons, W. A. Lavender, B. A. Lavender, C. A. Dorshow, S. W. Kiefer, R. Livingston, A. C. Scallet, and D. E. Crowley. Target structure and echo spectral discrimination by echolocating bats. *Science*, 186:1130–1132, 1974.

[Simmons *et al.* 90a] J. A. Simmons, M. Ferragamo, C. F. Moss, S. B. Stevenson, and R. A. Altes. Discrimination of jittered sonar echoes by the echolocating bat *Eptesicus fuscus*: the shape of target images in echolocation. *J. Comp. Physiol. A.*, 167:589–616, 1990.

[Simmons *et al.* 90b] J. A. Simmons, C. F. Moss, and M. Ferragamo. Convergence of temporal and spectral information into acoustic images of complex sonar targets perceived by the echolocating bat, *Eptesicus fuscus. J. Comp. Physiol. A*, 166:449–470, 1990.

[Simon 57] H. A. Simon. *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization.* Macmillan, 2 edition, 1957.

[Skolnik 62] M. Skolnik. *Introduction to Radar Systems.* McGraw-Hill Book Company, 1962.

[Sokolov & Makarov 71] B. V. Sokolov and A. K. Makarov. Direction of the ultrasonic radiation and the role of the nasal leaf in *Rhinolophus ferrumequinum. J. Biol. Sci.*, 7:37–44, 1971.

[Stremler 90] F. Stremler. *Introduction to Communication Systems.* Assison-Wesley Publishing Company, 3 edition, 1990.

[Strother & Mogus 70] G. K. Strother and M. Mogus. Acoustical beam patterns for bats: Some theoretical considerations. *J. Acoust. Soc. Am.*, 48:1430–1432, 1970.

[Suga & Jen 75] N. Suga and P. Jen. Peripheral control of acoustic signals in the auditory system of echolocating bats. *J. Exp. Biol.*, 62:277–311, 1975.

[Suga & Jen 77] N. Suga and P. Jen. Further studies on the peripheral auditory system of CF-FM bats specialized for fine frequency analysis of Doppler shifted echoes. *J. Exp. Biol.*, 69:207–232, 1977.

[Suga & Schlegel 72] N. Suga and P. Schlegel. Neural attenuation of responses to emitted sounds in echolocating bats. *Science*, 177:82–84, 1972.

[Suga & Shimozawa 74] N. Suga and T. Shimozawa. Site of neural attenuation of responses to self-vocalized sounds in echolocating bats. *Science*, 183:1212–1213, 1974.

[Suga 90a] N. Suga. Cortical computational maps for auditory imaging. *Neural Networks*, 3(1):3–21, 1990.

[Suga 90b] N. Suga. Cortical computational maps for auditory imaging. *Neural Networks*, 3:3–21, 1990.

[Suga *et al.* 75] N. Suga, J. A. Simmons, and P. Jen. Peripheral specialization for fine analysis of Doppler-shifted echoes in the auditory system of the CF-FM bat *Pteronotus parnellii*. *J. Exp. Biol.*, 63:161–192, 1975.

[Suga *et al.* 76] N. Suga, G. Neuweiler, and J. Möller. Peripheral auditory tuning for fine frequency analysis by the CF-FM bat, *Rhinolophus ferrumequinum* IV. Properties of peripheral auditory neurons. *J. Comp. Physiol. A*, 106:111–125, 1976.

[Sum & Menne 88] Y. Sum and D. Menne. Discrimination of fluttering targets by the FM bat *Pipistrellus stenopterus*. *J. Comp. Physiol. A*, 163:349–354, 1988.

[Surlykke 88] A. Surlykke. Interaction between echolocating bats and their prey. In P. E. Nachtigall and P. W. B. Moore, editors, *Animal SONAR Processes and Performance (NATO ASI Series)*, pages 551–566. Plenum Press, 1988.

[Suthers & Fattu 73] R. A. Suthers and J. M. Fattu. Mechanisms of sound production by echolocating bats. *Am. Zool.*, 13:1215–1226, 1973.

[Suthers 88] R. A. Suthers. The production of echolocation signals by bats and birds. In P. E. Nachtigall and P. W. B. Moore, editors, *Animal SONAR Processes and Performance (NATO ASI Series)*, pages 23–46. Plenum Press, 1988.

[Trappe & Schnitzler 82] M. Trappe and H-U. Schnitzler. Doppler-shift compensation in insect-catching horseshoe bats. *Naturwissenschaften*, 69:193–196, 1982.

[Trappe 82] M. Trappe. *Verhalten und Echoortung der Frossen Hufeisennase beim Beutefang.* Unpublished PhD thesis, Tübingen, 1982.

[Tufte 97] E. A. Tufte. *Visual explanations: images and quantities, evidence and narrative.* Graphics Press, 1997.

[Tuttle & Ryan 82] M. D. Tuttle and M. J. Ryan. The role of synchronized calling, ambient light and ambient noise, in anti-bat-predator behaviour of a tree frog. *Behav. Ecol. Sociobiol.*, 11:125–131, 1982.

[Vakman 72] D. Vakman. On the definition of concepts of amplitude, phase and instantaneous frequency of a signal. *Trans. Radio Eng. and Electron. Phy.*, pages 754–759, 1972.

[Vanderplank 50] F. L. Vanderplank. Air-speed/wing-tip speed ratios of insect flight. *Nature*, 165:806–807, 1950.

[Varela *et al.* 91] F. J. Varela, E. Thompson, and E. Rosch. *The Embodied Mind: Cognitive Science and Human Experience.* M. I. T. Press, 1991.

[Vater & Lenoir 92] M. Vater and M. Lenoir. Ultrastructure of the horseshoe bat's organ of corti. I. Scanning electron microscopy. *J. Comp. Neurol.*, 318:367–379, 1992.

[Vater 82] M. Vater. Single unit responses in cochlear nucleus of horseshoe bats to sinusoidal frequency and amplitude modulated signals. *J. Comp. Physiol. A*, 149:369–388, 1982.

[Vater *et al.* 85] M. Vater, A. S. Feng, and M. Betz. An HPR study of the frequency-place map of the horseshoe bat cochlea: morphological correlates of the sharp tuning to a narrow frequency band. *J. Comp. Physiol. A*, 157:671–686, 1985.

[Vater *et al.* 92] M. Vater, M. Lenoir, and R. Pujol. Ultrastructure of the horseshoe bat's organ of corti. I. Transmission electron microscopy. *J. Comp. Neurol.*, 318:380–391, 1992.

[Vaughan 86] T. A. Vaughan. *Mammalogy.* Saunders, 3 edition, 1986.

[von der Emde & Menne 89] G. von der Emde and D. Menne. Discrimination of insect wingbeat frequencies by the bat *Rhinolophus ferrumequinum. J. Comp. Physiol. A*, 164:663–671, 1989.

[von der Emde & Schnitzler 86] G. von der Emde and H-U. Schnitzler. Fluttering target detection in hipposiderid bats. *J. Comp. Physiol. A*, 14:43–55, 1986.

[von der Emde & Schnitzler 90] G. von der Emde and H-U. Schnitzler. Classification of insects by echolocating greater horseshoe bats. *J. Comp. Physiol. A*, 167:423–430, 1990.

[von der Emde 88] G. von der Emde. Greater horseshoe bats learn to discriminate simulated echoes of insects flutteing with different wingbeat rates. In P. E. Nachtigall and P. W. B. Moore, editors, *Animal SONAR Processes and Performance (NATO ASI Series)*, pages 495–500. Plenum Press, 1988.

[von Uexküll 21] J. von Uexküll. *Ummwelt und Innenwelt der Tiere.* Berlin, 1921.

[Walker 92] V. A. Walker. Bee-haviour in a mobile robot. Unpublished M.Sc. thesis, Department of Artificial Intelligence, University of Edinburgh, 1992.

[Walls 42] G. L. Walls. *The vertebrate eye and its adaptive radiation.* Hafner, 1942.

[Wartofsky 79] M. W. Wartofsky. Models: Representation and scientific understanding. In R. S. Cohen, D. Davidson, G. Nuchelmans, and W. C. Salmon, editors, *The model muddle: Proposals for an immodest realism*, volume 129, pages 1–11. D. Reidel Publishing Co., 1979.

[Webb & Wykes 96] P. Webb and C. Wykes. High-resolution beam forming for ultrasonic arrays. *I.E.E.E Trans. Robotics and Automation*, 12:138–146, 1996.

[Webb 93] B. H. Webb. *Perception in real and artificial insects: A robotic investigation of cricket phonotaxis.* Unpublished PhD thesis, University of Edinburgh, 1993.

[Webb 95] B. Webb. Using robots to model animals: a cricket test. *Robotics and Autonomous Systems*, 16:117–134, 1995.

[Webster & Brazier 65] F. A. Webster and O. G. Brazier. Experimental studies on target detection, evaluation, and interception by echolocating bats. TDR AMRL-TR-65-172, Aerospace Medical Division USAF Systems Command, 1965.

[Webster 67] F. A. Webster. Interception performance of echolocating bats in the presence of interference. In *Animal SONAR Systems*, pages 673–713. Lab. Physiol. Acoust. CNRS Jouy-en-Josas, 1967.

[Wehner 81] R. Wehner. Spatial vision in arthropods. In H. Autrum, editor, *Handbook of Sensory Physiology*, pages 287–616. Springer-Verlag, vii/6c edition, 1981.

[Wehner 87] R. Wehner. Matched filters — Neural models of the external world. *J. Comp. Physiol. A*, 161:511–531, 1987.

[Wenstrup 88] J. J. Wenstrup. Binaural neurones in the mustache bat's inferior colliculus: physiology, functional organization, and behavioural implications. In P. E. Nachtigall and P. W. B. Moore, editors, *Animal SONAR Processes and Performance (NATO ASI Series)*, pages 329–333. Plenum Press, 1988.

[Wenstrup *et al.* 86] J. J. Wenstrup, L. S. Ross, and G. D. Pollak. Binaural response organization within a frequency-band representation of the inferior colliculus: implications for sound localization. *J. Neurosci.*, 6:962–973, 1986.

[Wenstrup *et al.* 88a] J. J. Wenstrup, Z. M. Fuzessery, and G. D. Pollak. Binaural neurons in the mustache bat's inferior colliculus I. Responses of 60 khz E-I units to dichotic sound stimulation. *J. Neurophysiol.*, 60:1369–1383, 1988.

[Wenstrup *et al.* 88b] J. J. Wenstrup, Z. M. Fuzessery, and G. D. Pollak. Binaural neurons in the mustache bat's inferior colliculus II. Determinants of spatial responses among 60-khz EI neurons. *J. Neurophysiol.*, 60:1384–1404, 1988.

[Williamson 96] M. W. Williamson. Postural primitives: Interactive behaviour for a humanoid robot arm. In P. Maes, M. J. Mataric, J-A. Meyer, J. Pollack, and S. W. Wilson, editors, *From animals to animats: Proc. Forth Int. Conf. Simulation of Adaptive Behaviour*, 1996.

[Winograd & Flores 86] T. Winograd and F. Flores. *Understanding computers and cognition: A new foundation for design.* Ablex Publishing Co., 1986.

[Wotton *et al.* 96] J. M. Wotton, T. Haresign, and J. A. Simmons. Spatially dependent acoustic cues generated by the external ear of the big brown bat, *Eptesicus fuscus. J. Acoust. Soc. Am.*, 98(3):1423–1445, 1996.

Part V

Appendices

"The machines that are first invented to perform any particular movement are always the most complex, and succeeding artists generally discover that, with fewer wheels the same effects may be more easily produced."

—Adam Smith

A. The Polaroid Transducer

In the transduction process considered here, signals are generated by applying an electrical signal to an electrostatic transducer. This produces a vibration at the transducer's surface which, in turn, produces acoustic pressure waves. In the case of a linear, reversible transducer — such as the Polaroid series 7000 transducer used in this work [Biber *et al.* 80] — reciprocity dictates that the directivity of the transmitter is exactly that of the receiver [Anke 74] . Therefore, the filter characteristics presented for the transmitting transducer are identical to those of the receiver.

A.1 Generating acoustic energy

The Polaroid series 7000 electrostatic transducer operates according to the Sell principle [Anke 74]. As shown in Figure A.1, a gold plated plastic foil serves as one plate of a capacitor which, when a bias voltage is applied to the transducer, is attracted to a grooved backplate. This configuration creates multiple, largely independent, ring-shaped capacitors — one for each groove.

From this model, Anke derives a transfer function $H(f)_{tr}H(f)_{rec} = \frac{V_{out}}{V_{in}}$ based on the assumption that the circular regions of the membrane move like a concentrated mass with a concentrated stiffness. That is, the radial dependencies are ignored so that the velocity of these parts of the membrane can be taken as a constant.

The measured transfer function can be characterised as a bandpass filter with resonant frequency $f_{res} = 54$ kHz and centralised bandwidth $B_c = 9.9$ kHz [Peremans 94], where

$$f_0 = \frac{\int_0^\infty f|X(f)|^2 df}{\int_0^\infty |X(f)|^2 df} \tag{A.1}$$

and

$$B_c^2 = \frac{\int_0^\infty (f - f_0)^2 f|X(f)|^2 df}{\int_0^\infty |X(f)|^2 df} \tag{A.2}$$

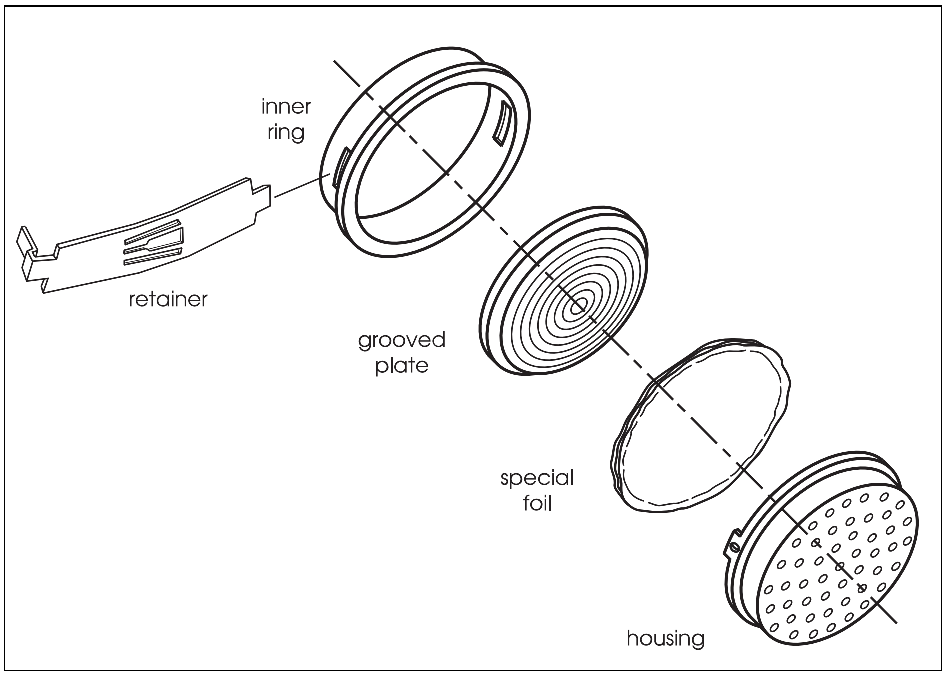

Figure A.1: **The Polaroid transducer [Biber** *et al.* **80].**

This measured transfer function also agrees with the widely used approximation which defines the impulse response of an aligned transducer pair as

$$h_{tr/rec} = exp(-\frac{t^2}{2\sigma^2} sin(2\pi f_0 t)) \tag{A.3}$$

where σ is inversely proportional to the centralised bandwidth of the filter [Bozma & Kuc 91].

This transfer function agrees with those derived empirically up to the first resonant frequency. Higher resonant frequencies — corresponding to higher order vibrations — are not accounted for by the model and, although a bias voltage is applied to linearise the transducer's response, nonlinear effects are still present due to high signal amplitudes at transmit time. However, the second harmonic of the signal contains approximately 0.29% of the energy of the signal [Peremans 94] and, therefore, the assumption that the transducer can be modelled by a linear transformation is reasonable.

A.2 Radiating acoustic energy

This transfer function describes the filtering characteristics of an aligned transmitter/receiver pair — *i.e.* a receiver (or reflector) whose surface normal is aligned with the transmitter's normal axis. However, additional filtering is introduced during off-axis sensing. To account for this, an additional transfer function is required to describe

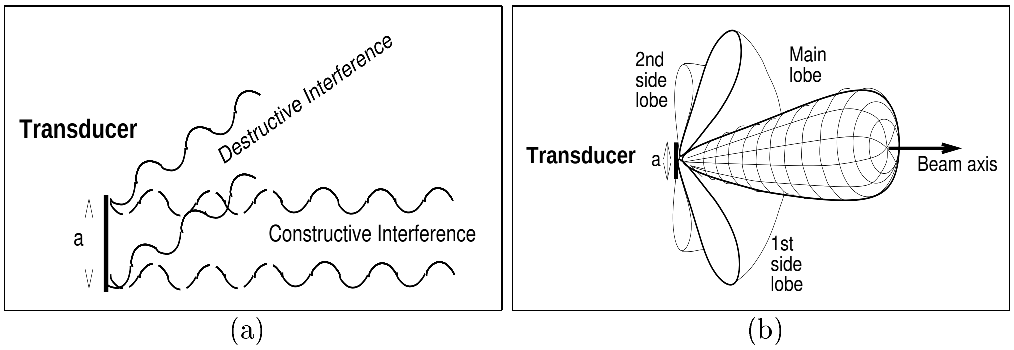

Figure A.2:
A monochromatic Polaroid pulse. Far field behaviour of the Polaroid transducer. Normal to the transducer (0°), waves from the top and bottom halves of the transducer reinforce to produce a region of high intensity. At some angles relative to this axis, waves from the top and bottom halves of the transducer are 180° out of phase and destructively interfere to produce intensity minima [McKerrow & Hallam 90]. (a) Wave interference patterns. (b) Resultant lobe pattern (side lobes cut away).

the angle dependent filtering effects.

The directionality of the Polaroid transducer comes about through the phenomena of diffraction. As shown in Figure A.2, interference of waves generated at different parts of the finite extent of the transducer surface yield a non-uniform radiation pattern. Employing the principle of superposition, the radiation pattern of a spatially extended transducer can be determined analytically from the total acoustic field produced by an appropriate collection of simple sources.

An easily integrable (and, therefore, widely used) extended transducer model is that of a rigid circular piston mounted flush with an infinite baffle, and vibrating with simple harmonic motion $u = U_0 cos\omega t$. For a piston of radius a, generating a continuous tone into a homogeneous, non-attenuating medium, the pressure in the far field[1] (approximately $r > \frac{a^2}{\lambda}$) is given by

$$p = Re\{\frac{j\rho ck}{2\pi r} U e^{j(\omega t - kr)} D(\theta)\} \tag{A.4}$$

where $k = \frac{2\pi}{\lambda}$ is the wavelength constant and ρ is density [Kinsler & Frey 62, Morse & Ingrad 68]. The directivity term $D(\theta)$ (Equation A.4) is found by a geometrical approximation in which the differences in pressure phase at a point in the field $P_{ref}(r, \theta)$ — due to path length differences between sound produced by individual surface elements — are integrated over the radiating surface (see Figure A.3) as

$$D(\theta) = \int_0^a \sigma d\sigma \int_0^{2\pi} e^{jk\sigma sin\theta cos\psi} d\psi = 2\pi a^2 [\frac{J_1(kasin\theta)}{kasin\theta}] \tag{A.5}$$

[1] In the near field, the beam is contained within a cylinder of diameter $2a$, where a is the transducer radius.

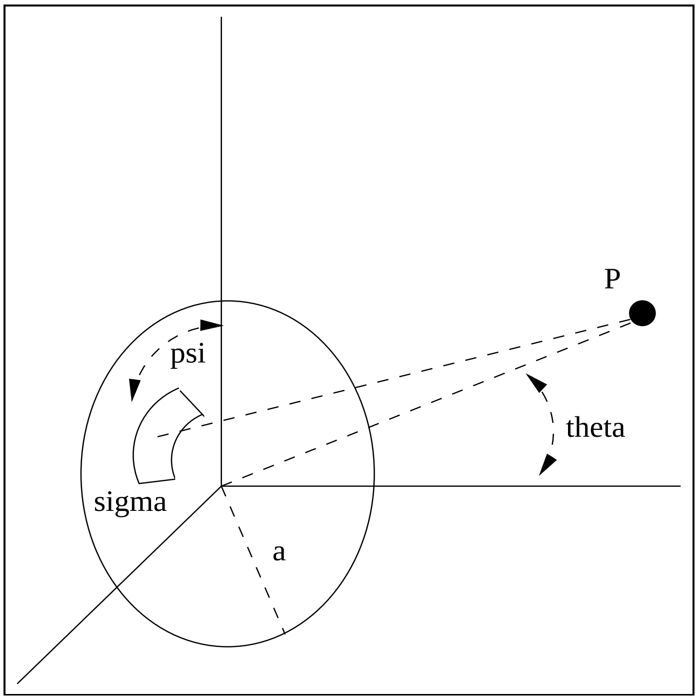

Figure A.3: **Coordinate system used in deriving the radiation characteristics of a flat circular piston [Kinsler & Frey 62].**

The directivity term defines pressure minima in the radiation pattern at well defined viewing angles. They occur for θ given by:

$$0 = \frac{J_1(kasin(\theta))}{kasin(\theta)} \tag{A.6}$$

For the Polaroid series 7000 transducer used in these studies ($a = 1.13$ cm), the first three minima occur at angles $\theta = 22°$, $42°$, and $77°$ at $f = 50$ kHz. (Figure 5.1 depicts this polar sensitivity pattern.) Since the pressure radiated by a piston source is a function of polar angle as well as distance from the source, the wave is not spherically symmetrical. However, it still has the characteristic property of a diverging spherical wave: pressure is inversely proportional to the radial distance from the center of the source. The phase of the pressure on any spherical wavefront is also the same at all points included within the major lobe.

Transduction in a Sell type transducer, where vibration of the surface occurs in accordance with the multiple capacitor model, is well approximated by the piston model [Peremans 94]. Therefore, the latter, simpler expression is used in these studies.

B. The Bionic Hardware Model

The contents of this Appendix is taken from a [Peremans *et al.* 97].

B.1 Introduction

The sonarhead is designed to make it possible to test various theories proposed to explain particular measurement strategies of biological in-air sonar systems. Since we believe that the success of those sonar systems is to a large extent due to their use of active sensing strategies we have tried to include as many relevant degrees of freedom in our sonarhead as possible. As a result, the designer of an experiment using the sonarhead can control the 6 DOF pose of the sonarhead, as well as the parameters defining the call sequence, i.e. repetition rate, duration of call, shape of transmit pulse. The sonarhead consists of a hardware and a software component, both components are described in this text.

B.2 Hardware description

The hardware of the system consists of three components: the sensorhead, the transmitter/receiver module and the signal processing module. We describe each of these three components in more detail below.

B.2.1 Sensorhead

The sensorhead has 6 DOF's as indicated in Fig. B.1, allowing panning and tilting of the neck, and independent panning and tilting of each of the two ears. The motors driving the different axes are standard radio-control servo motors. The control signals, i.e. pulse-width modulated signals, are generated by a transputer. This approach makes it possible for one transputer plus a simple motor interface module, consisting of a link adaptor and a buffer, to control up to 8 different motor axes simultaneously, as shown in Fig. B.2. The schematics of the motor interface module are discussed in Section B.5.

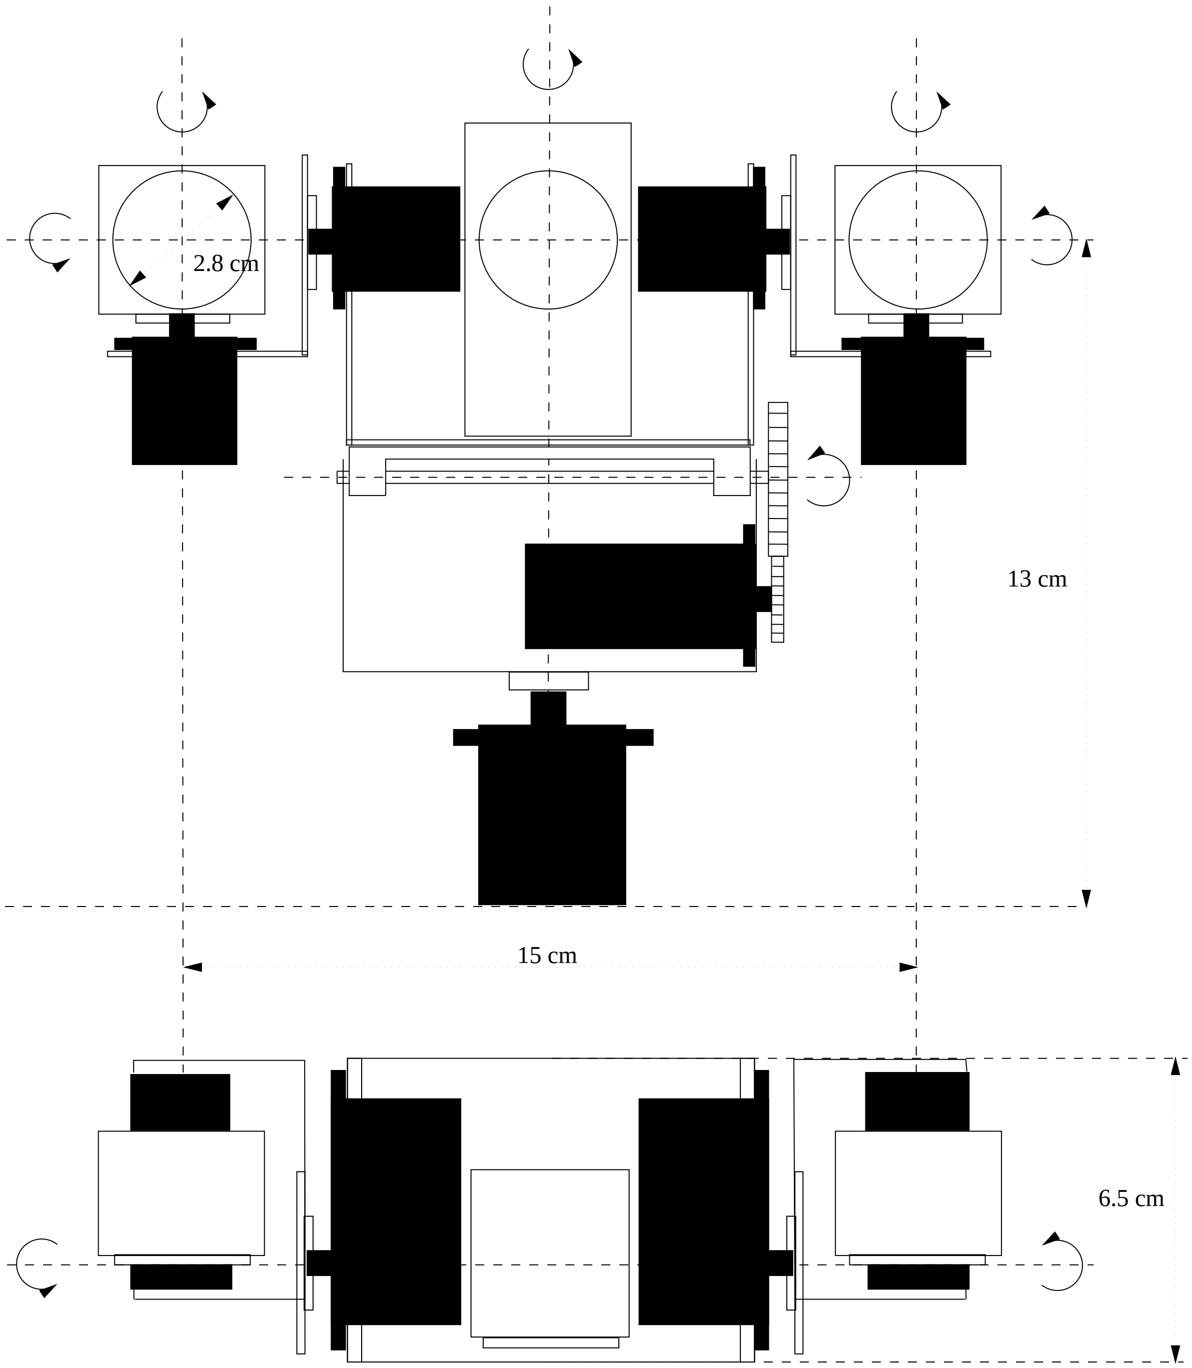

Figure B.1: The ultrasonic sensor consisting of the central transmitter fixed to the head and the two receivers each independently orientable.

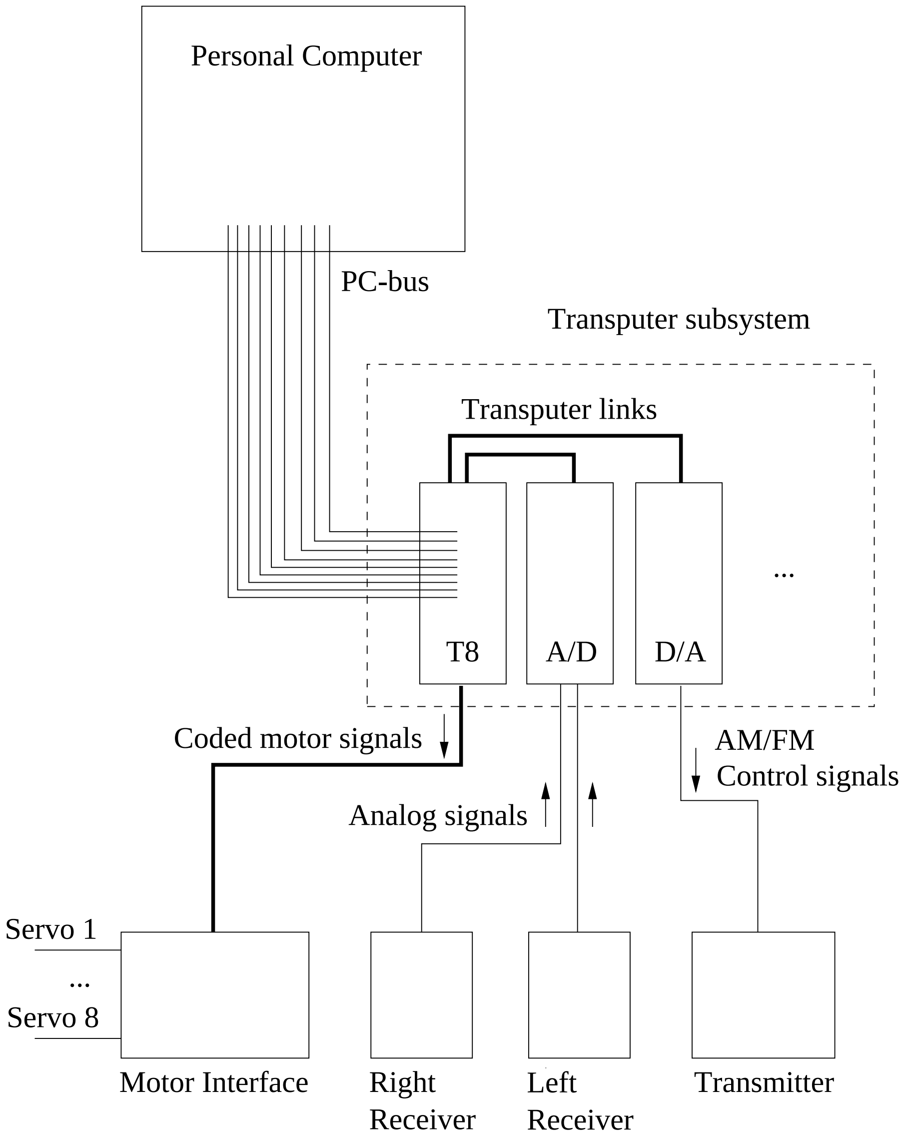

Figure B.2: The sonarhead's control system.

The positions of the servo motors are set by the width of the pulse applied to their control pin. These pulses have to be repeated every 20 ms to keep the servomotors locked in place. How this is accomplished by the program running on the transputer is explained in Section B.3.1.

B.2.2 Transmitter/receiver module

The transmitter module is split in two parts: the AM/FM modulator + power amplifier and the actual transmitter consisting of the transformer and transducer. The latter part is mounted at the centre of the head, moving along with the pan and tilt movements of the neck. It has inputs for the low voltage signal from the power amplifier and for the DC bias (max. 200V) signal. The transformer and DC bias stage, transducer dependent, is separated from the signal generation and amplification electronics, generic, to make it easier to switch between different transducers. The first part of the transmitter module has two inputs for 'fm' and 'am' modulations to be imposed on a carrier wave. In our system the two modulation signals are generated by a two-channel, transputer controlled, D/A converter. Tables B.1 and B.2 in Section B.4 show the conversion from input voltage level to modulation parameters. By setting an internal jumper the 'fm' input can also be used to specify the transmit signal directly. In this case, the signal generation part of the transmitter module is bypassed and the signal at the 'fm' input is fed directly into the power amplifier.

Detection and amplification of the reflected echoes at the ears is performed by the receiver modules, which are mounted, together with the transducers, on the pan/tilt servo's. Each receiver module has inputs for a 12/15V power supply and for a DC bias (max. 200V) signal. The output signals from the receivers are sampled at 200 kHz by a two-channel, transputer controlled, A/D converter. Note that a small extra circuit connects the sample clocks of the D/A and A/D converters in such a way that both clocks are in phase and the sample rate of the D/A converter is half that of the A/D converter. This circuit can be easily removed if so desired. All further processing of the received data is performed on a transputer based multiprocessor.

The schematics of the transmitter/receiver module and the clock synchronisation circuit are discussed in Section B.4. The processes running on the transputers supporting the transmitter/receiver module are described in Section B.3.2.

B.2.3 Signal processing module

The signal processing operations performed upon the received signals are based upon a simple model of the processing performed by the mammalian cochlea, Fig. B.3. These operations are all executed in software, there is no hardware support for them. First, the signals are filtered by a filter-bank consisting of bandpass filters centered at frequencies within the frequency regions of interest. Depending on the emitted signal used those might span the entire frequency range of an 'fm' cry, or they might be concentrated in a few regions of interest around the harmonics of a single 'cf' cry. The filters have been implemented by 5th order Butterworth bandpass filters with Q values ranging from 5 to 150. Next, the outputs of the filter-bank are half wave rectified and

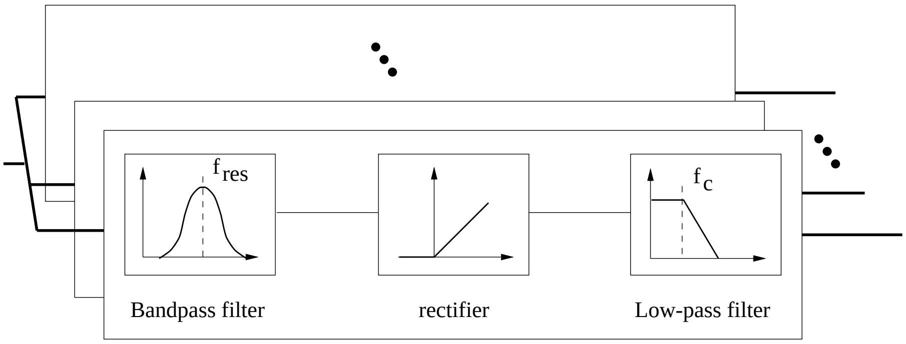

Figure B.3: The cochlear model.

smoothed. A 2nd order Butterworth low-pass filter is used with a cut-off frequency at 2000 Hz. We show an example of the outputs of the model at one ear for an 'fm' cry reflected by a stationary planar obstacle in Fig. B.4.

B.3 Software description

The sonarhead is controlled by software running on a transputer network (a set of TRAM modules), the different TRAM modules and their interconnections are shown in Fig. B.5. The file containing the program for the C004 switch is called `benhope.wir`, it is reproduced in Section B.6. This diagram also shows the names given to the different TRAM's (upper left corner) as well as the processes that run on each of them. The connections are labelled with the respective link numbers. The detailed configuration information is contained in the file `benhope.cfs` reproduced in Section B.7.

In Fig. B.6 we show what files are used to define the different processes that run on the transputer network. Below we give a more detailed description of those processes.

B.3.1 Controlling the pose of the sonarhead

As stated before the positions of the servo motors are set by the width of the pulse applied to their control pins. These pulses have to be repeated every 20 ms to keep the servo motors locked in place. The desired pulse-widths are controlled by the transputer by switching on and off the different bits of a control byte. Each bit of the control byte corresponds with one motor axis as shown in Fig. B.7.

The program that drives the servo motors consists of two loops. During the first loop a timer (which uses the internal transputer clock) generates pulses with the required pulse-width, resolution 1 μs, in sequential order, i.e. no two motors are simultaneously active, to avoid crosstalk. Next during the second loop, commands from the outside world are read in. This loop gets executed until the total duration of both loops is equal to 20 ms, i.e. the period of the overall servo motor control loop. The new positions, if new commands were received, or the old positions, if no new commands were received,

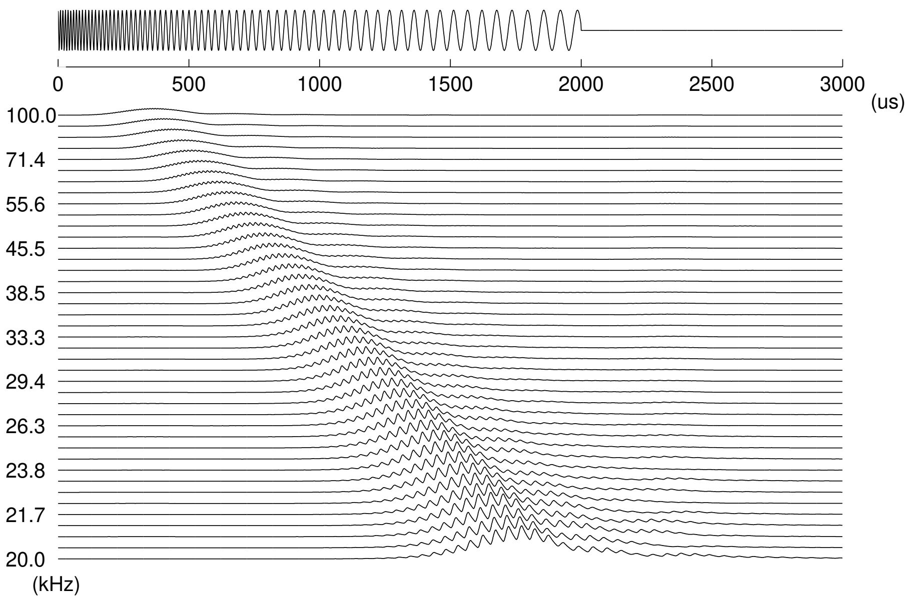

Figure B.4: Outputs of the cochlear model for an 'fm' cry.

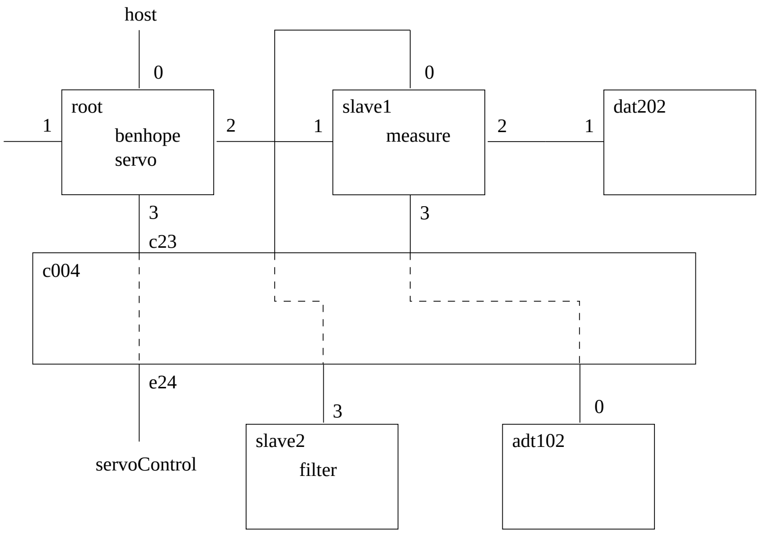

Figure B.5: The transputer network that controls the sonarhead.

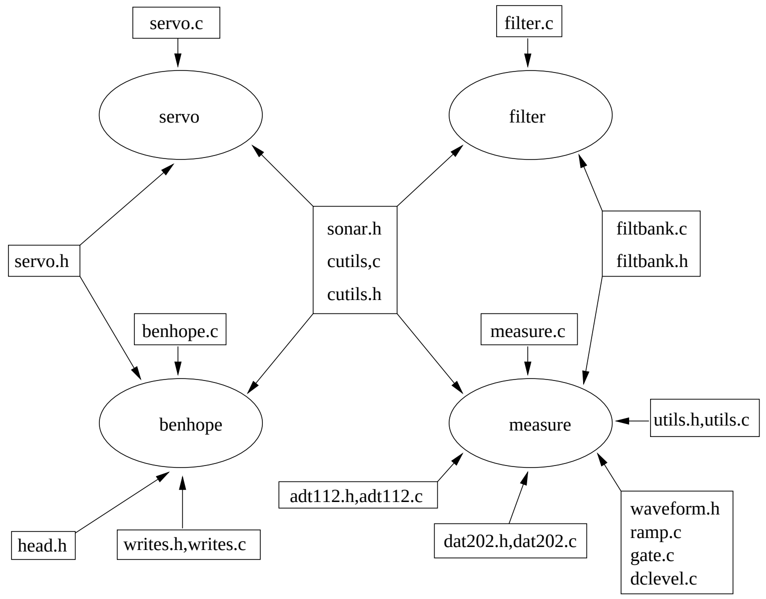

Figure B.6: The software that runs on the transputer network.

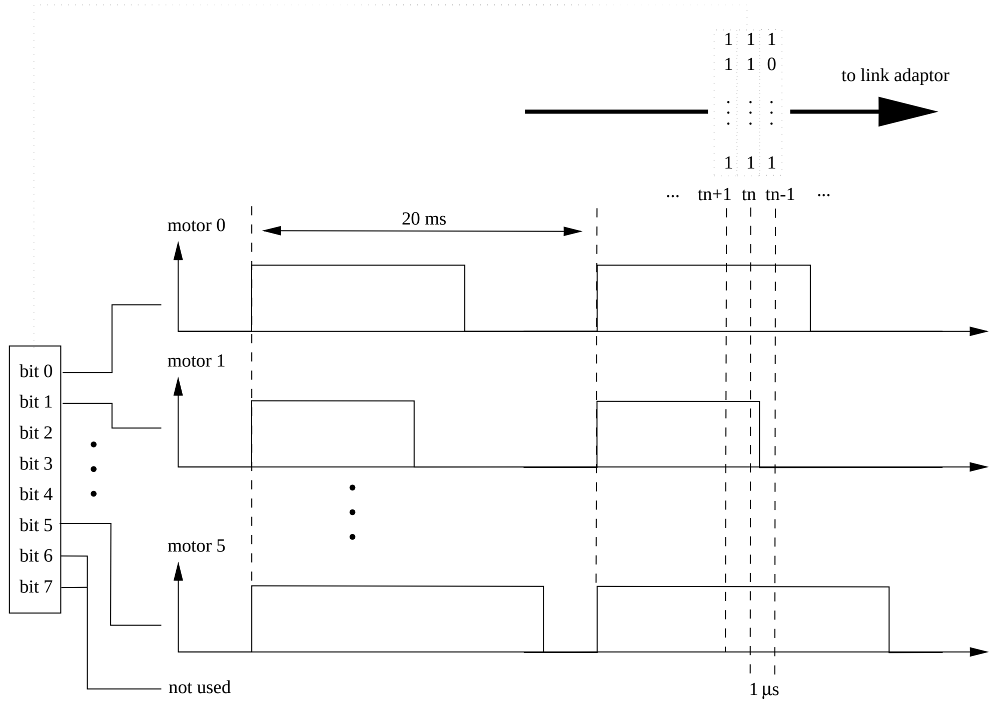

Figure B.7: The different bits of the control byte control the different servo motors.

are then used to calculate a new set of pulse-widths and a new control cycle is started. To achieve the resolution of 1 μs the program has to be executed as a high priority process on the transputer, therefore the process is put on hold regularly during the second loop to allow other processes that run on the same transputer to be executed. If sufficient resources are available it is advised to run the control program on a separate transputer as it takes up 90 % of the CPU-time to control six servo motors.

The code describing the actual servo control is in files `servo.c` and `servo.h`; the routines called in the user program to communicate with this servo control program are in `head.h` which has to be included in the user program. The basic routine to get information from the servo control program is `GetAngle(ServoMotor)` which, when given a servo motor number, returns the current position of that servo in radians. The correspondence between the different degrees of freedom of the head and the `ServoMotor` values are defined in `servo.h`. Note that this position is the one the servo motor should be at, not necessarily the one it is actually at, as there is no feedback from the servo motors to the servo control program. The value of this angle lies in the interval $[-\pi/2, \pi/2]$ and 0 is defined to coincide with the position of the servo motor when the sonar head is in its home position. Note that because of the construction of the head the neck's tilt range was restricted from $[-\pi/2, \pi/2]$ to $[-\pi/6, \pi/4]$. The basic routines to send information to the servo control program are `MoveAbs(ServoMotor,Angle)` and `MoveRel(ServoMotor,Angle)` which, when given a servo motor number and an angle in radians, send the appropriate commands to the servo control program. This results in a move to an absolute position or one relative to the current position, respectively.

The file `head.h` also contains a number of more high level routines that allow the sensors to look at specific target points, specified in cartesian or spherical coordinates, with or without panning the head. Note that these were written to drive a sonarhead that had only a panning neck so none of them use the tilt capability of the neck.

B.3.2 Controlling the transmitter/receiver modules

The processes controlling the transmitter module allow the user to specify the shape of the transmit pulse in two ways: by specifying duration, instantaneous amplitude and instantaneous frequency, or by specifying the sampled waveform (sample rate $\leq$ 100 kHz). The latter can be read in from file. As mentioned before, a jumper on the transmitter/receiver module has to be set to switch from one mode to the other. In either mode, the control signals, i.e. AM/FM modulation signals or digital waveform, are written by the transputer to the D/A board which converts them into electrical signals that can be applied directly to the inputs of the transmitter electronics. Section B.4 contains the conversion tables. The output signals from the two receivers are applied to the two inputs of the A/D board (max. sample rate = 200 kHz). The converted data are written to a buffer on the A/D board and read out by the transputer at its own pace. The transputer reads the data, under DMA control, into a large array.

The message type used to communicate between the main program and the process controlling the A/D and D/A converter boards is `data_t` defined in `sonar.h`. This message consists of different fields whose meaning varies with the command type as specified by the first field, i.e. the `type` field,

```
typedef struct {
        char type;
        char frequency_code;
        int frequency;
        int no_channels;
        int no_samples;
        char waveform;
        float v_initial;
        float v_final;
        float v_off;
        float duration;
        int servo_id;
        int servo_value;
        char contents[MAX_DATA_CONTENTS];
} data_t;
```

The possible values for the `type` field are:

- 'a' to set the A/D conversion parameters: number of samples (=`no_samples`) where the total number (=both channels) has to be smaller than 120000, number of channels (=`no_channels`) and measurement result type (=`waveform`) which can be either 'p' (=processed) or 'r' (=raw measurement data);
- 'e' to set the transmit pulse to the waveform specified in the `contents` field;
- 'd' to set the two channels of the D/A board and hence define the transmit pulse by selecting either a constant voltage `data_buff.waveform='d'`, additional parameters are D/A channel used (=`no_channels`) and voltage level (=`v_initial`), or a rectangular voltage `data_buff.waveform='g'`, additional parameters are D/A channel used (=`no_channels`), ON voltage (=`v_initial`), OFF voltage (=`v_final`) and duration in multiples of the D/A clockperiod (=`duration`), or a ramp voltage `data_buff.waveform='r'`, additional parameters are D/A channel used (=`no_channels`), initial ON voltage (=`v_initial`), final ON voltage (=`v_final`), OFF voltage (=`v_off`) and duration in multiples of the D/A clockperiod (=`duration`);
- 'f' to stop reading commands and start the actual measurement.

The A/D and D/A process, as defined in `measure.c`, consists of three parts. The first part is the command reading loop. This loop reads in commands from the main program to configure the A/D and the D/A boards. The command reading loop stops when it is given an 'f' (=fire) command. At that point, the second part is executed performing an actual measurement by creating two parallel processes: one steering the D/A board to generate a transmit pulse as set up by the previous commands and the other controlling the A/D board to capture the programmed amount of data at the programmed sample frequency. Finally, the third part starts up the signal processing module, i.e. the cochlear transform, and returns the processed results if a 'p' type measurement was requested or the raw sampled data if an 'r' type measurement was requested.

B.3.3 The cochlear processing module

The signal processing procedure `process()` is defined in the file `filtbank.c`. As explained before, it processes the raw measured data with a set of 5th order Butterworth bandpass filters. In the current implementation, it then half wave rectifies the outputs of each of those and processes the rectified signals with a first order butterworth lowpass filter. Finally, the first peak is detected in each of the frequency channels and the amplitudes and positions of those peaks are returned.

Note that different kinds of information could be extracted from the outputs of the set of bandpass filters if necessary. Contrary to the other parts of the program described above, the signal processing is still very much a topic of ongoing research and hence changes with respect to the scheme described here should be expected.

B.4 Transmitter/receiver module

The schematic of the receiver is shown in Fig. B.8. The receiver consists of a first amplification stage built around a current mirror. The current variations induced by an arriving soundwave are translated, `gain = 24000`, into voltage variations at the output of the first stage. The `LM394` is used to provide a matched transistor pair for the current mirror. The second amplification stage is built around a standard op-amp and has `gain = 11`.

The schematic of the transmitter module is shown in Fig. B.9. The first part of the transmitter module, built around the `8038`, is responsible for appropriately processing the modulation information provided through the 'fm' and 'am' inputs. The VCO is configured to generate a sinusoidal signal whose frequency is determined by the voltage applied to the 'fm' input, marked `CH0` on Fig. B.9. The little circuit around the FET `ZVN4206` uses the voltage applied to the 'am' input (marked `CH1`) to determine the amplitude of the sinusoidal signal generated by the VCO. Since this circuit works with small signals only, the output voltage of the VCO is divided first. The conversion from input voltages, applied to the 'fm' and 'am' inputs of the transmitter module, to output signal parameters are given in Tables B.1 and B.2.

The second part of the transmitter module, built around the `TL082CP`, is responsible for amplifying the 'am' and 'fm' modulated output of the VCO before applying it to the transducer. The first amplification stage is a voltage amplifier with `gain=12`. The second stage is predominantly a power amplifier, voltage `gain=6`, driving the transducer through a transformer. The two stage amplification is introduced for stability reasons. Note that the output of the first amplification stage can be replaced by the 'fm' input as an input for the final amplification stage by changing the setting of a jumper. This allows arbitrary signals to be applied to the transducer through the second amplifier stage. The power amplifier is a standard class AB amplifier. The 220 ohm pot allows regulation of the amount of quiescent current that flows through the transistor pair. The latter are a complementary pair of Darlington transistors. Because of the large currents flowing through them they have to be cooled; they are thermally connected to the metal housing. The series resistors of 2 ohm prevent thermal runaway of the

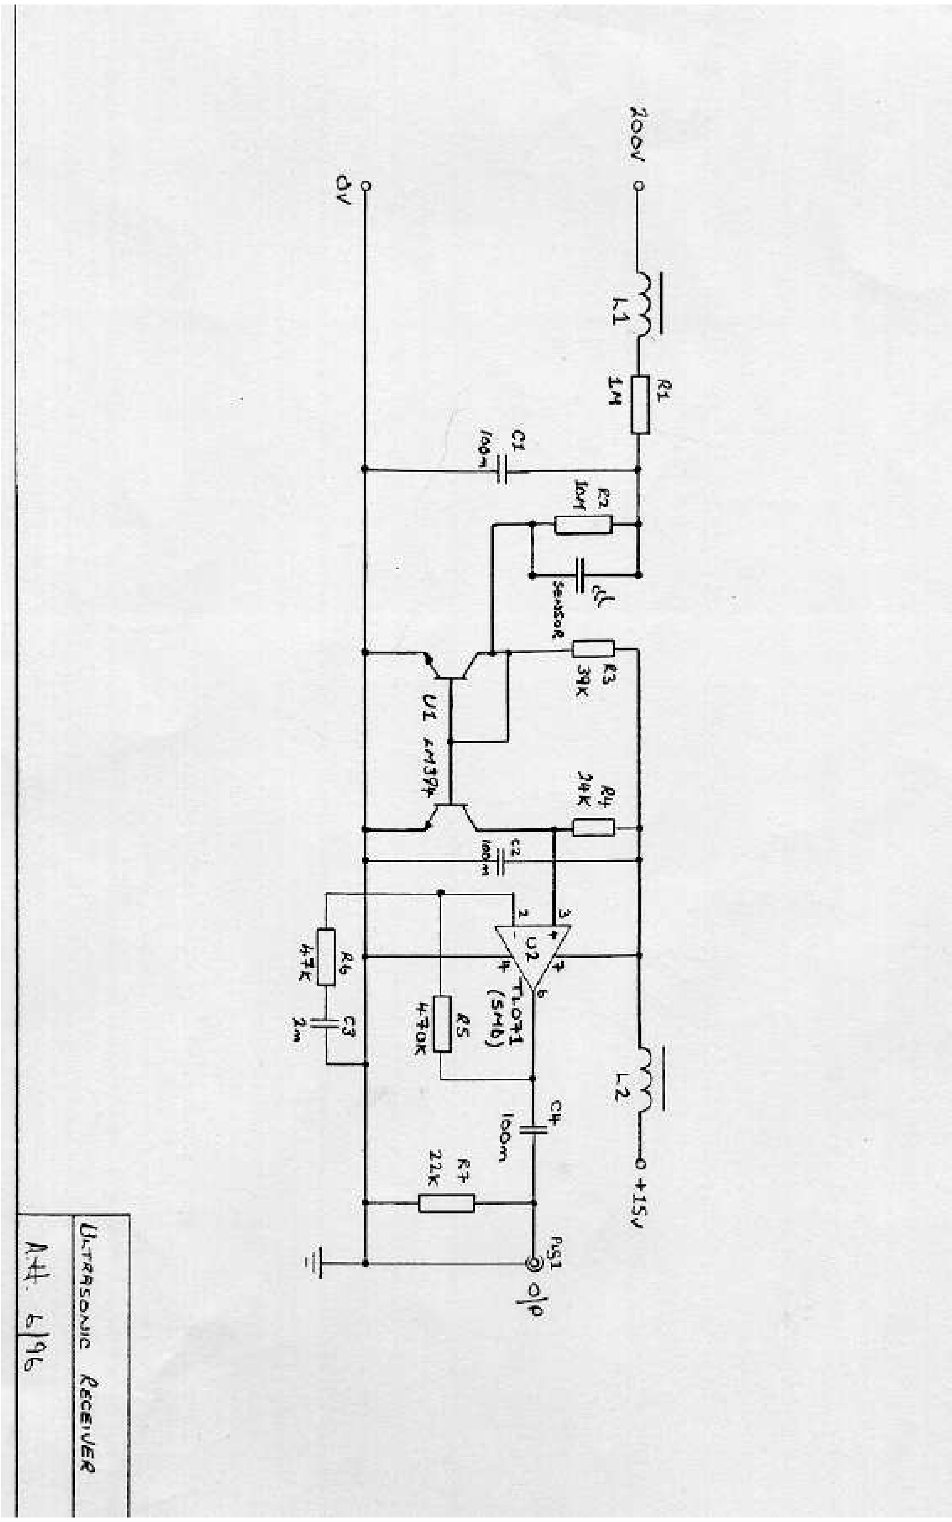

Figure B.8: The receiver.

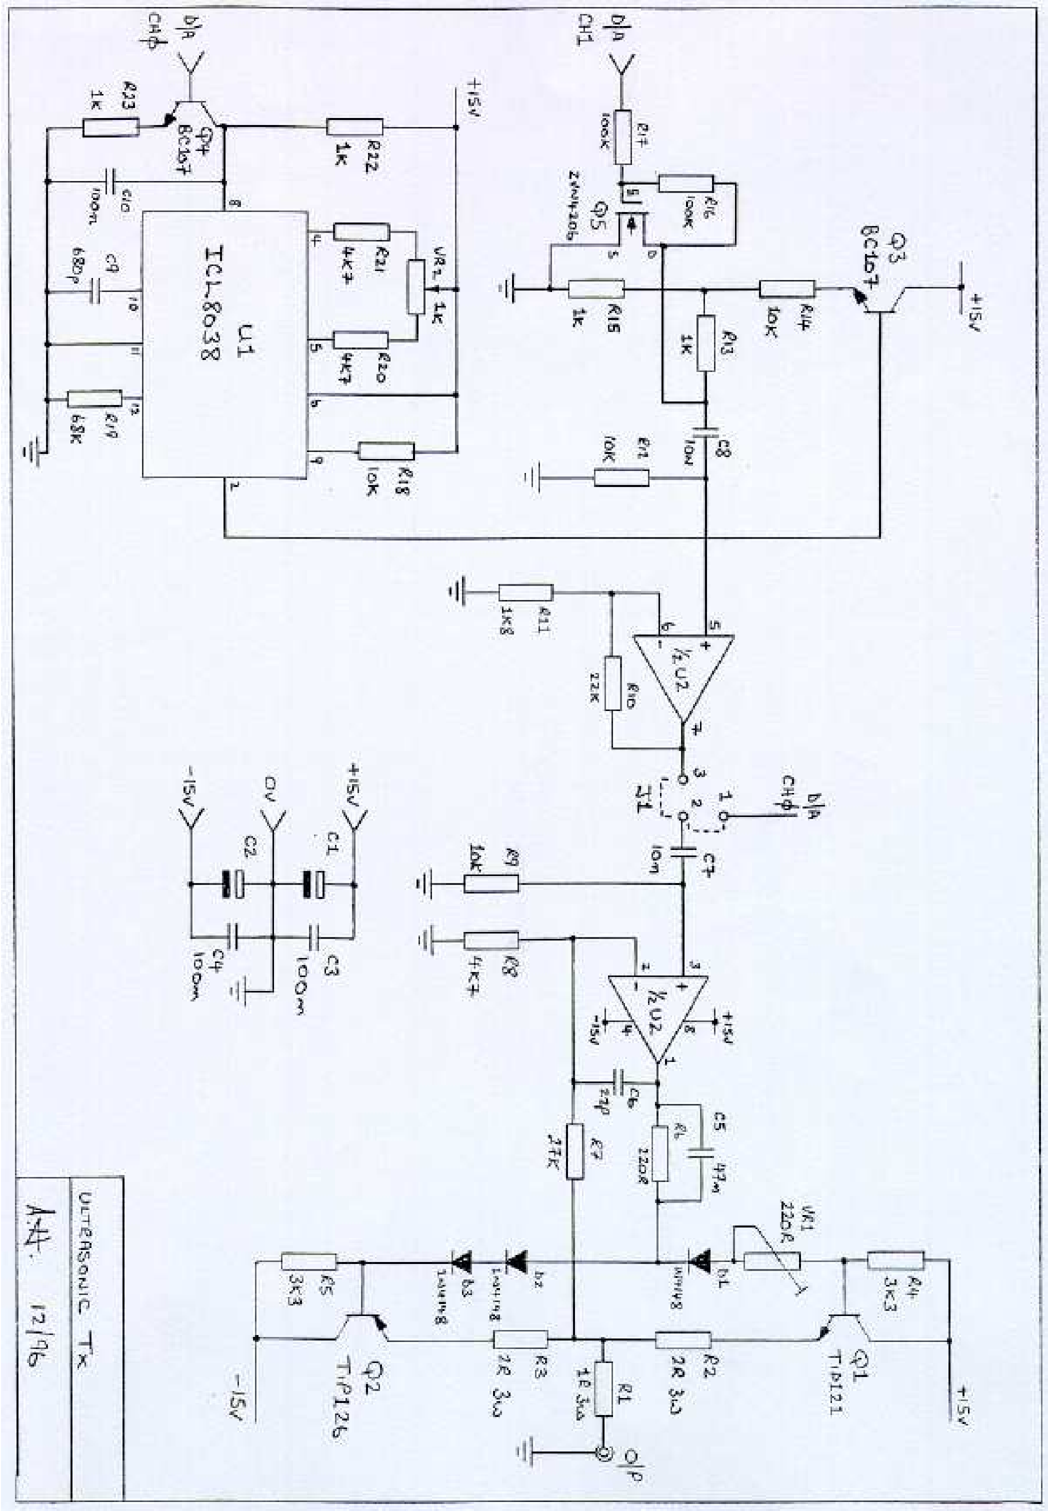

Figure B.9: The transmitter.

input (V)	frequency (kHz)
1.5	26
1.65	30
1.8	34
2.0	39
2.4	50
2.8	60
3.2	70
3.5	79
4.0	92
4.3	100
5.0	118

(a)

input (V)	output(Vpp)
1.5	120
1.2	90
1.0	70
0.8	60
0.6	50
0.4	30
0.2	18
0.0	0

(b)

Table B.1: Conversion from ‘fm’ input voltage level to (a) frequency and (b) output amplitude.

input (V)	output (Vpp)
0.0-3.3	120
3.6	110
3.9	100
4.2	80
4.35	60
4.5	40
4.65	20
4.8	11
5.0	5

Table B.2: Conversion from ‘am’ input voltage level to output amplitude.

transistors.

The final part of the transmitter/receiver module is the transformer connecting the transducer to the power amplifier. The transformer consists of 20 turns on the primary and 200 turns on the secondary. The schematic of the transducer conditioner circuit is shown in Fig. B.10.

The schematic of the DC bias voltage power supply is shown in Fig. B.11.

The schematic of the clock synchronisation circuit which is just a clock divider is shown in Fig. B.12. Note that the input to this circuit is not the A/D sample clock itself but a fixed 5 MHz clock on the A/D board from which the A/D clock is derived. Hence, we get a fixed frequency (= 100kHz) D/A sample clock synchronous to the A/D sample clock.

B.5 Motor interface module

The schematic is shown in Fig. B.13. The `C011` link adapter is used in Mode 1, converting between a link and two independent, handshaking, byte-wide interfaces. The 20MHz signal from the oscillator is divided by 4 by the `7493` to generate the 5MHz clock required by the `C011`. By connecting `IValid` to the ground and `Qack` to `QValid` the `C011` is set up to accept values from the transputer connected to it through its link at the maximal rate of the channel. This rate is set to 20Mbits/sec by connecting `SepIQ` to `ClockIn`. The values read in are kept constant at the outputs `Q0-Q7` of the `C011`. They are changed by writing a new value to the `C011`. These outputs are buffered by the `74LS244` before being fed to the output connector.

Note that the power-on reset does not work effectively, necessitating a manual reset. To do this push the reset button, after power-on!

B.6 Transputer network wiring file

```
y
s1 3 t s9 0 .
s1 0 t s8 3 .
c23 t e24 .
```

B.7 Transputer software configuration file

```
/***********************************************************************/
/* File        benhope.CFS                                             */
/* Version                                                             */
```

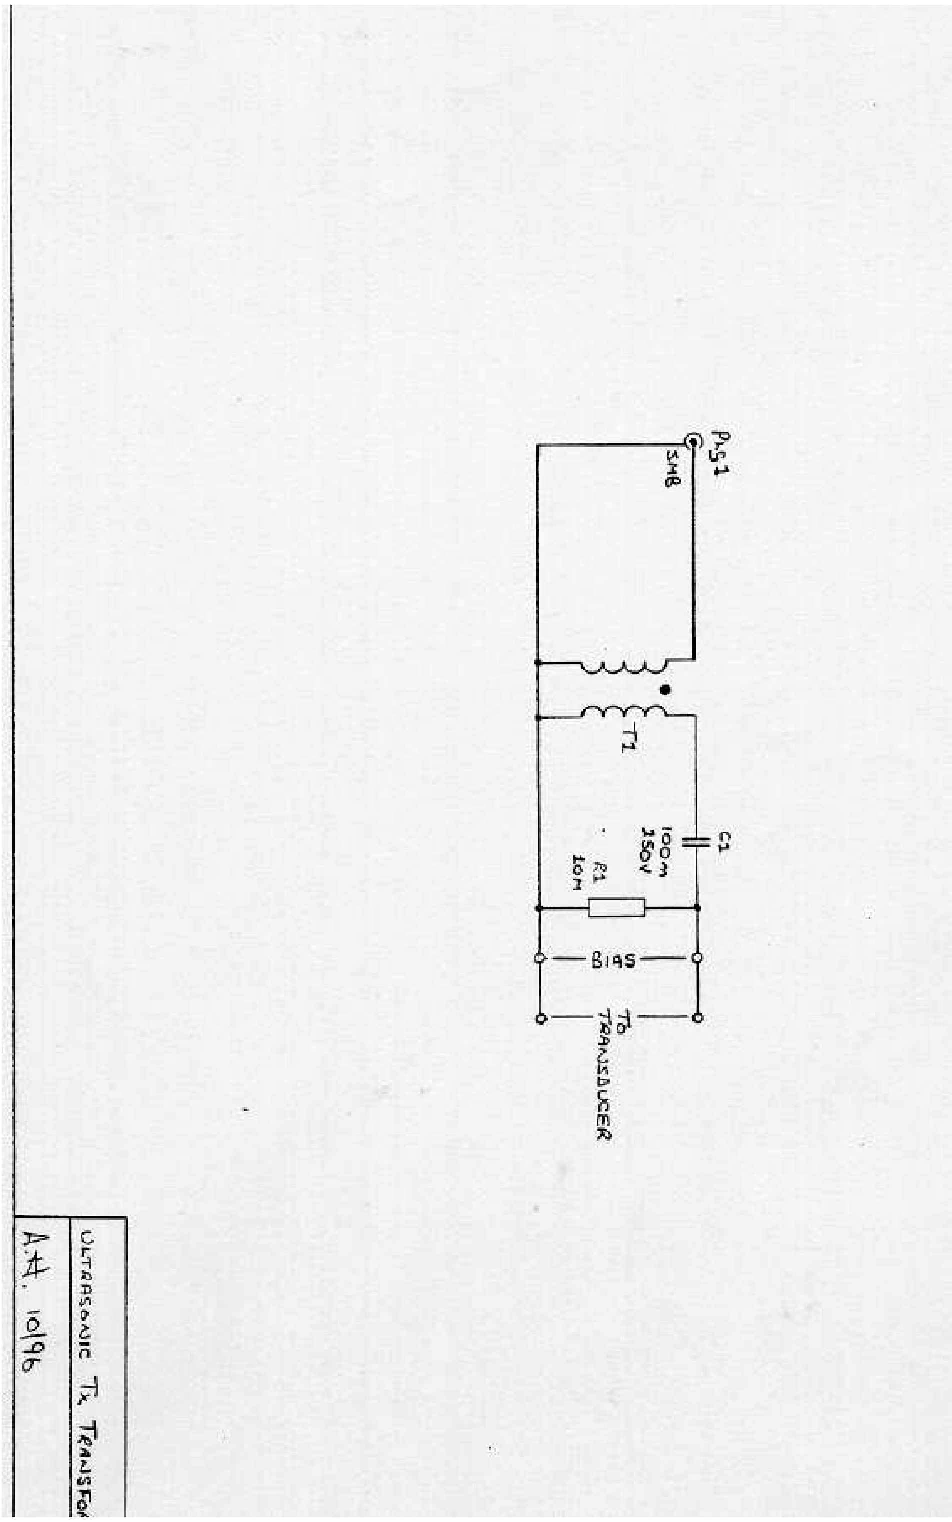

Figure B.10: The transducer conditioner circuit.

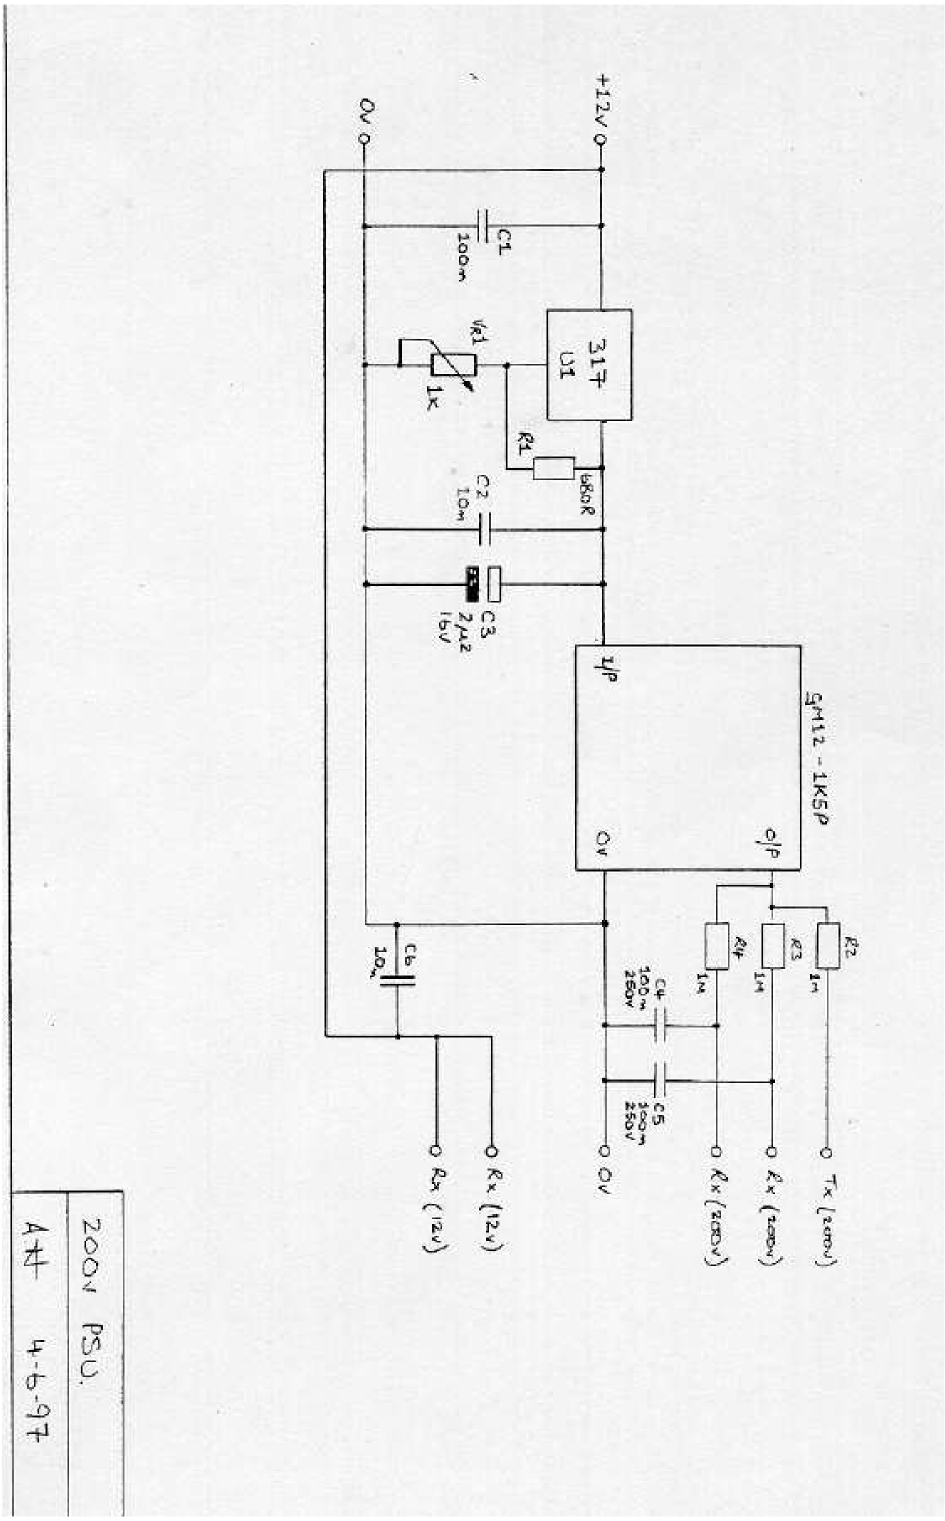

Figure B.11: The DC bias power supply.

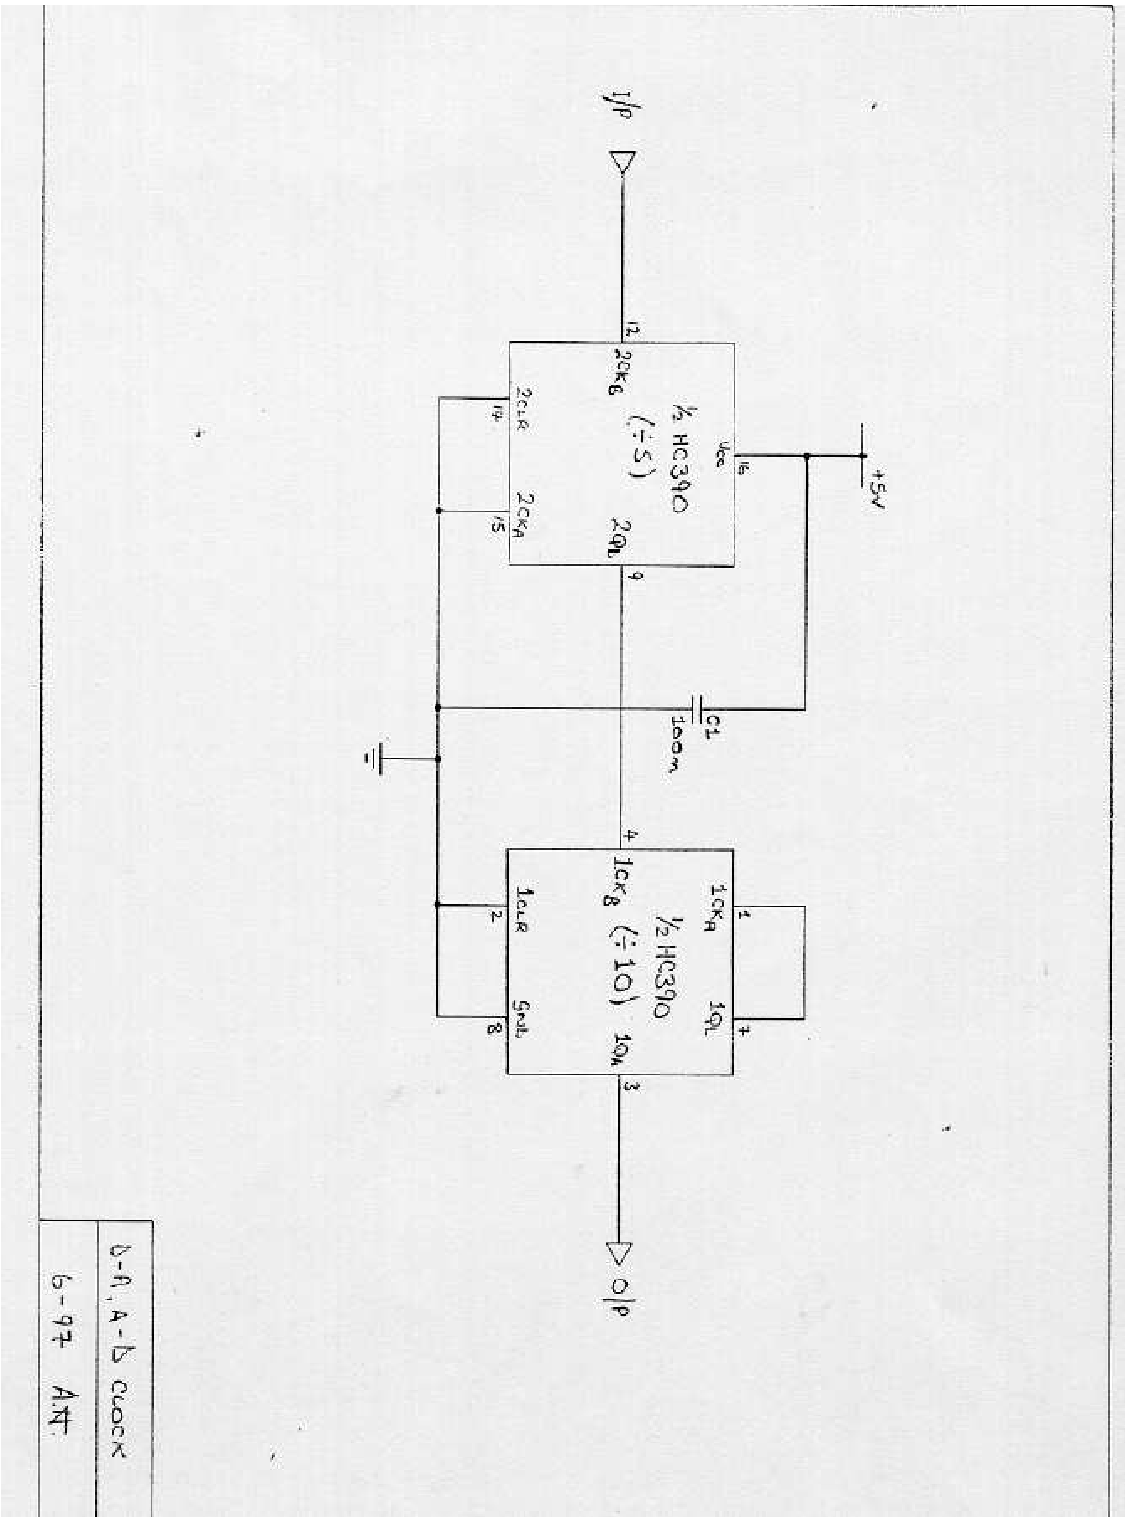

Figure B.12: The clock synchronisation circuit.

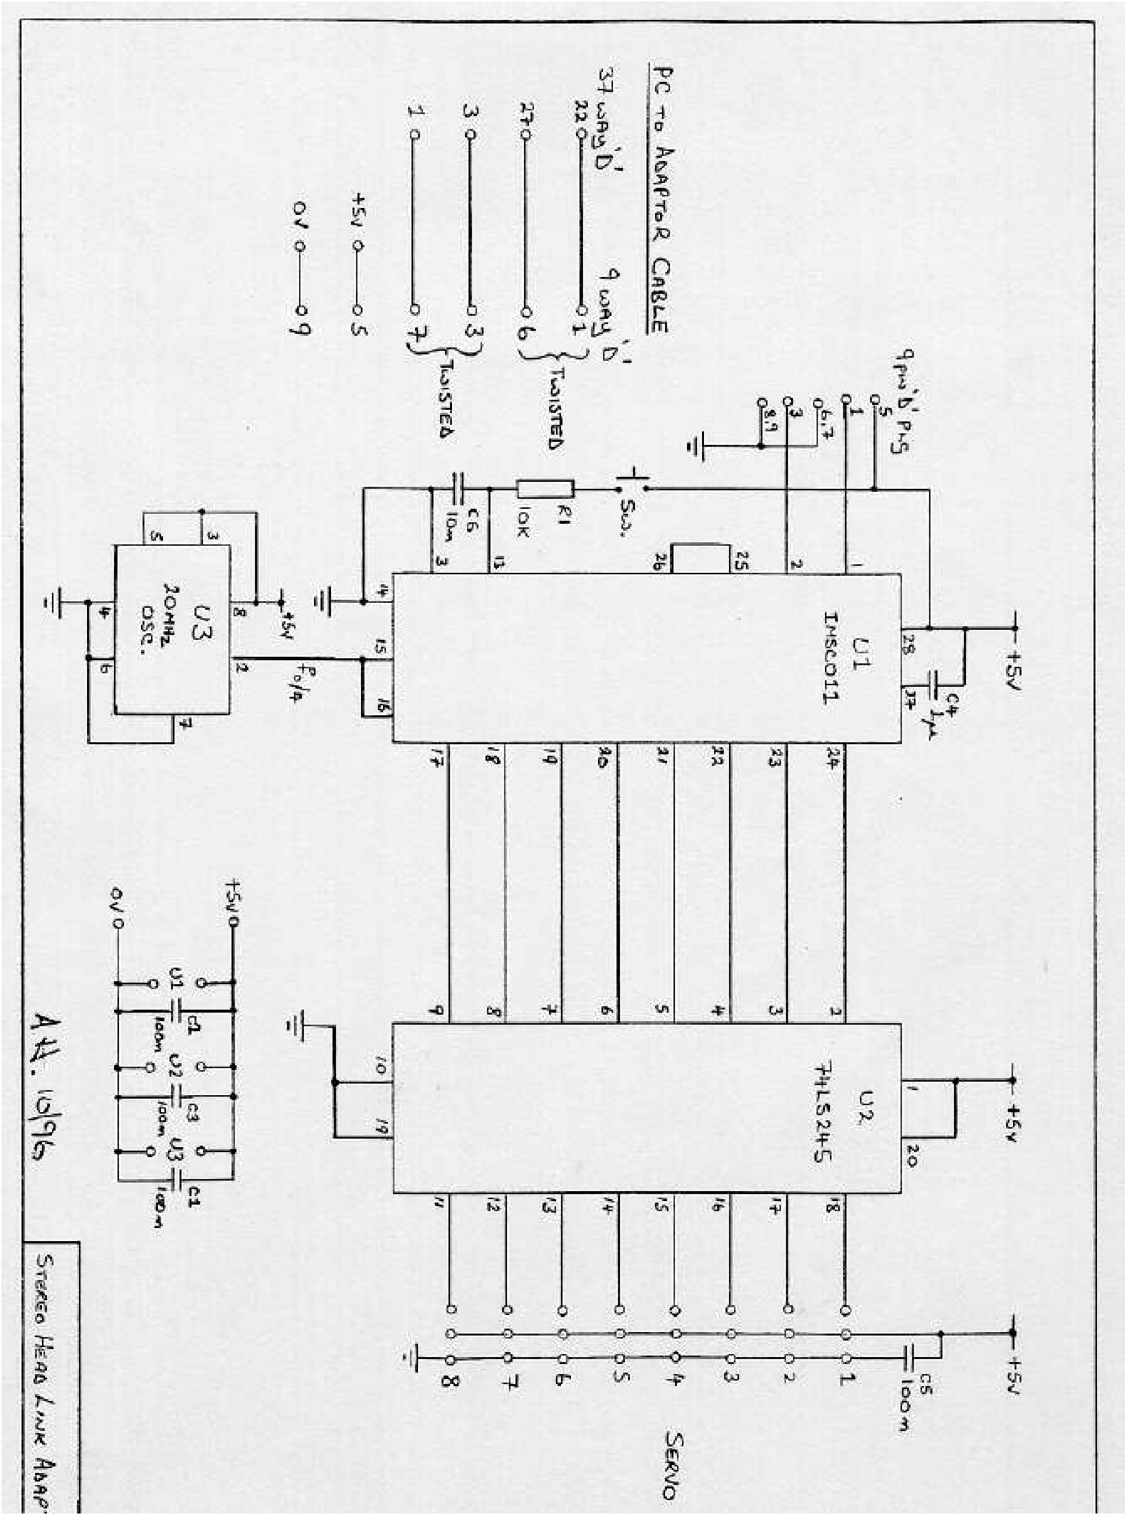

Figure B.13: The motorinterface module.

```
/* Last Updated      11/06/97                                        */
/* Compiler     Inmos 'C' Toolset                                    */
/* Description:                                                      */
/*********************************************************************/

/*Hardware Description*/
T805 (memory = 1M) root;
T800 (memory = 8M) slave1;
T800 (memory = 8M) slave2;
edge adt102;
edge dat202;
edge servoControl;
edge F68k;

connect root.link[0] to host;
connect root.link[1] to F68k;
connect root.link[2] to slave1.link[1];
connect root.link[3] to servoControl;
connect slave1.link[0] to slave2.link[3];
connect slave1.link[2] to dat202;
connect slave1.link[3] to adt102;

val i 0;

input from_server;
output to_server;

input from_adt102;
output to_adt102;

input from_dat202;
output to_dat202;

output servoOutput;

input F68kInput;
output F68kOutput;

/*Software Description*/

process (stacksize = 450K, heapsize = 80K,
         interface (input fs, output ts,
                    input from_head, output to_head,
                    input from_measure, output to_measure,
                    input F68kIn, output F68kOut)) benhope;

process (stacksize = 800K, heapsize = 3200K,
         interface (input from_control,
```

```
                    output to_control,
                    input adt102_in,
                    output adt102_out,
                    input dat202_in,
                    output dat202_out,
                    input from_filter,
                    output to_filter)) measure;

process (stacksize = 800K, heapsize = 3200K,
         interface (input from_measfilt,
                    output to_measfilt)) filter;

process (stacksize = 50K, heapsize = 50K,
         interface (input fromTalker,
                    output toTalker,
                    output toServo)) servo;
servo (priority=HIGH);

/* Mapping description */

connect from_server to benhope.fs;
connect to_server to benhope.ts;

connect from_adt102 to measure.adt102_in;
connect to_adt102 to measure.adt102_out;

connect from_dat202 to measure.dat202_in;
connect to_dat202 to measure.dat202_out;

connect servo.toServo to servoOutput;

connect servo.fromTalker to benhope.to_head;
connect servo.toTalker to benhope.from_head;

connect benhope.to_measure to measure.from_control;
connect benhope.from_measure to measure.to_control;

connect benhope.F68kIn to F68kInput;
connect benhope.F68kOut to F68kOutput;

connect measure.from_filter to filter.to_measfilt;
connect measure.to_filter to filter.from_measfilt;

use "benhope.c9x" for benhope;
use "measure.c8x" for measure;
use "filter.c8x" for filter;
use "servo.c9x" for servo;
```

```
place benhope on root;
place measure on slave1;
place filter on slave2;
place servo on root;

place from_server on host;
place to_server on host;

place from_dat202 on dat202;
place to_dat202 on dat202;

place from_adt102 on adt102;
place to_adt102 on adt102;

place servoOutput on servoControl;

place F68kInput on F68k;
place F68kOutput on F68k;
```

C. The Bionic Software Model

The *3D Echolocation Simulator* used in this investigation was built by Herbert Peremans and myself in the summer of 1996. Our original intent was to employ the simulator to create echoes from simple point-like reflectors under well-understood conditions of noise and clutter so that we could test signal processing algorithms before deploying them on the robotic sensorhead. However, the *3D Echolocation Simulator* has developed into an investigative platform itself which can host experiments that are impractical (or impossible) given the limitations of existing hardware. For example, we can introduce new call structures; change the size of the transducers and the morphology of sensor head; create oscillating targets with different geometries and motions; ... all in the time that it takes to compile and run the software (*i.e.* a few seconds on an Ultra).

Because the software grew up alongside a robotic system, we have been able to incorporate the physics/acoustics necessary to support our observations and, therefore, the simulator delivers reasonably realistic results. In the following sections, I give an overview of the simulator environment (Section C.1) and then discuss in more detail how amplitude (Section C.2) and frequency (Section C.3) alterations are made to a sound pulse as it propagates within the simulator environment.

C.1 Overview

The simulator takes a description of a single echolocator (*i.e.* an emitter and a pair of receivers, as described in Section C.1.1); the amplitude and frequency characteristics of a pulse; and a list of reflecting targets (comprised of one or more point-like reflectors, as described in Section C.1.2). During each *SensingInterval* (*i.e.* typically, the time for a call to go out and an echo to return), a pulse(s) is propagated through the environment (as a single ray). During its out- and in-bound journey, the amplitude and frequency characteristics of the pulse are modified by environmental effects (beam spreading and frequency-dependent atmospheric absorption) and effects introduced by reflecting surfaces (scattering, Doppler shift, and directionally-dependent filtering).

While reflectors remain within the forward hemisphere defined with respect to the transmitter, each reflecting point produces one echo. All delayed, attenuated and, possibly, Doppler shifted echoes are added together to produce the resultant echo for that *SensingInterval*. No secondary bounces of sound between targets before returning to the echolocator are considered.

The echolocator and targets are described below.

C.1.1 Echolocator

The echolocator is composed of a single transmitter flanked by two receivers (see Figure C.1). The present system uses circular piston-like apertures for transmitting and receiving sound. This model was chosen because it is a good approximation of the behaviour of the Polaroid electrostatic transducers (see Appendix A); however, this model could easily be replaced with more biologically-plausible transducer models. For example, the transmitter might be replaced by a dipole emitter (to mimic the nasal emitters of rhinolophids and hipposiderids). Likewise, the receivers might be modelled as horns. Although not used in this thesis work, artificial pinnae have been added to the receivers in the form of a series of reflectors with particular orientations and positions with respect to the surface of each receiving transducer. Each of these pinnae surfaces collect energy from particular regions of the frontal sound field and reflect it onto the transducer [Papadopoulos 97, Peremans *et al.* 98b].

Once transducer type(s) has been selected, transducer sizes, start positions and start orientations must be specified for both the transmitter and receivers. The receiving surfaces may be defined to lie at a distance, *EarLen*, from the notional surface of the head — thereby approximating the physical situation wherein sound may be diffusely reflected from the tip of a pinnae into an ear canal. (In this case, the receivers can be re-orientated during a *SensingInterval* by moving them along arcs or driving them through a conical scanning motion with a particular angular extent and speed.)

The echolocator as a whole also has a start position, start orientation and constant speed which takes effect from the beginning of each *SensingInterval*.

C.1.2 Target

The targets are modelled as collections of one or more point-like reflectors (see Figure C.2) which take a 3D *Position* and, for oscillating targets (*e.g.* model insects), a *FlapRate*. In the case of the latter, the model insect is centred at the specified position and two reflecting points are placed at the tips of the wings. (Additional reflecting points can be defined to lie anywhere between $TargetPosition$ and $TargetPosition{+}WingLength$ and their speed automatically scales with distance from the non-oscillating body centre.) The insect may take on any 3D *Orientation*, but the wings move only in a 2D flap plane (orthogonal to the insect) through arcs of a given angular extent, see Figure C.2, explosion. (No twisting or wing rotation is currently defined.)

Various wing velocity profiles can be defined. Typically a sinusoidally varying wing velocity is used wherein the wings are stalled at the upper and lower extremities of their motion and move most rapidly through the middle of their angular excursion. In

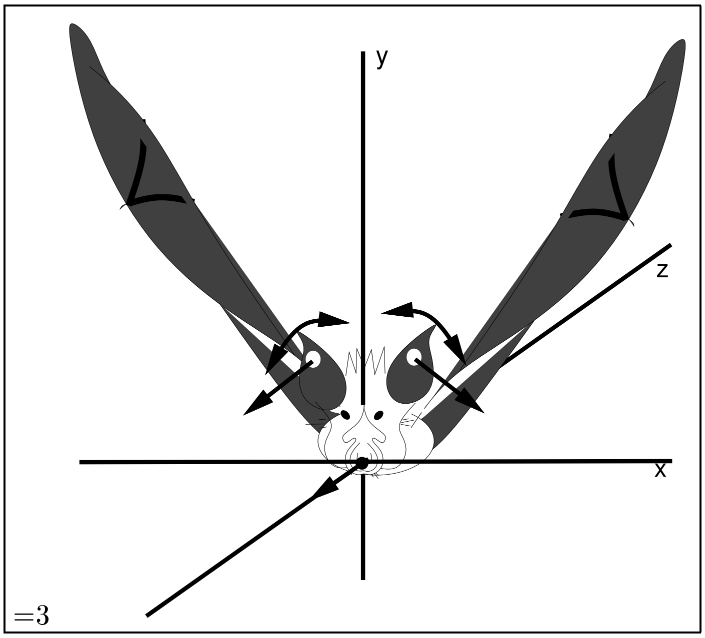

Figure C.1: **Simulated bat.**

addition to this flutter, or "AC", velocity, the reflectors may also have a "DC" velocity which is held for each $SensingInterval$.

A number of assumptions had to be made in arriving at this characterisation of an insect. In Section 9.1.9, I review what is known about the reflective properties of real insect targets.

C.2 Amplitude

The calculation of signal attenuation is made twice: once during the out-bound journey (pulse transmission, propagation and reflection) and then again for the in-bound journey (echo propagation and reception). In the case of the latter, four physical phenomena diminish the sound pressure level (SPL) of the call.

- *Spreading.* The SPL of the spherical emit pulse attenuates as $1/r^2$, where r is target range.
- *Absorption.* Frequency dependent absorptivity effects decrease the SPL (in dBs) as $(0.038\, f_c - 0.3)\, r$, where f_c is measured in kHz.
- *Directivity.* Because the intensity of the emitted pulse varies across the beamwidth (see Figure A.2), the intensity reaching the target depends upon its position (θ) relative to the emitter. As discussed in Appendix A, the directivity pattern of the circular pistons used to model emitter and receivers can be described by Equation A.5:

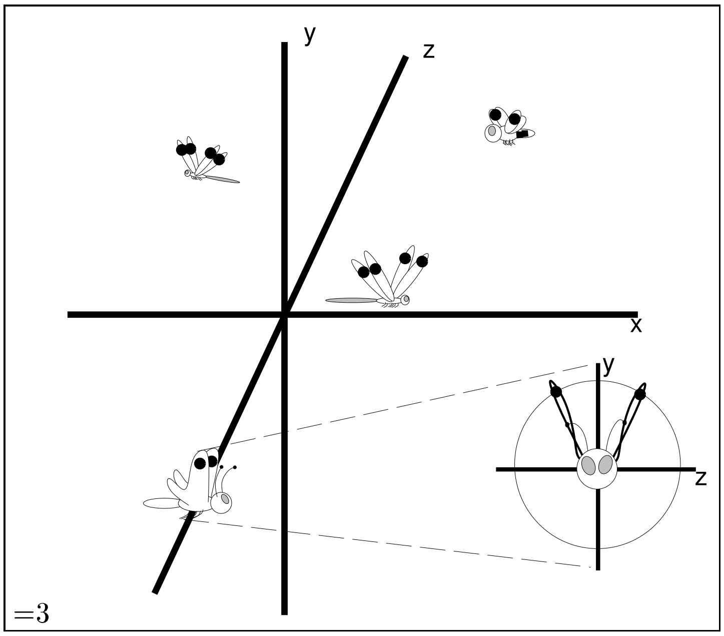

Figure C.2: **Simulated targets.**

$$D(\theta) = 2\pi a^2 [\frac{J_1(kasin\theta)}{kasin\theta}] \tag{C.1}$$

Therefore, signals are attenuated as $20 log_{10} \frac{D(\theta)}{D(0)}$. Because target and/or echolocator are moving during the simulation, at each sampling period, the angle $\theta(t)$ is defined as the angle between the normal to the transducer at $CallTime$ and the position of the target at $CallArrivalTime$.

- *Scattering.* Acoustic energy is absorbed at the surface of targets. Analytic descriptions of scatter at the surface of complex targets are hard to come by and absorptivity coefficients for various surfaces are typically measured and catalogued. Measurements of scattering at the surface of insect wings (fluttering and non-fluttering) have been published by Kober and Schnitzler [Kober & Schnitzler 90]. This data is given as attenuation of a pure tone during reflection off insects as a function of call frequency and wing length. Several of Kober and Schnitzler's measurements are used in the insect model contained in this simulator. Specifically, we use their measurements in a scatter, or "glinting", function which scales the magnitude of a reflection between the maximum (*i.e.* fluttering) and minimum (*i.e.* non-fluttering) reflection values reported for a given insect. A maximum strength reflection occurs when the angle between emitter axis and the normal to the wing (θ_n) is 0°; reflection strength decreases (sinusoidally) to its minimum value at $\theta_n = 90°$. (In this regard, one might argue that the reflectors used in the simulator are more specular than diffuse, but this is the only instance in which the target surfaces have an orientation.)

During its in-bound journey, the echo is attenuated due to spreading, absorption and directivity effects. In the case of the latter, directivity depends upon the angle $\theta(t)$ which is measured between $TargetPosition(CallArrivalTime)$ and $ReceiverPosition(EchoArrivalTime)$.

C.3 Frequency

In its most general form, the Doppler effect is defined as "the change in the apparent time interval between two events which arises from the [relative] motion of the observer together with the finite velocity of transmission of information about the events"[Gill 65]. In this simulation where emitter velocity v_B, receiver velocity v'_B and target velocity v_T are allowed, echo frequency is derived from call frequency according to the Equation C.2:

$$f_e = f_c \frac{c + v'_B\, cos(\delta_{B'T})}{c - v_B\, cos(\delta_{BT})} \frac{c - v_T\, cos(\delta_{TB})}{c + v_T\, cos(\delta_{TB'})} \tag{C.2}$$

where the Doppler angles $\delta_{B'T|BT|TB|TB'}$ denotes the angle between the line-of-sight from emitter/receiver to target, and the direction of the corresponding velocity vector.

As in the case of echo amplitudes, first the reflected f_r and then the echo f_e frequencies are calculated. In the case of the former, the target's velocity (at the arrival time of one sample of the sound pulse $v_T(CallArrivalTime)$) and the emitter's velocity (at the time that the corresponding pulse sample was emitted $v_B(CallTime)$) are projected onto the line-of-sight between the two surfaces ($p_T(CallArrivalTime) - p_B(CallTime)$).

Upon reflection, the procedure is reversed to calculate the echo frequencies heard in each ear. Therefore, at each sampling point, the receiver velocity at the echo arrival time $v'_B(EchoArrivalTime)$ and the target velocity at the time the corresponding sound pulse was reflected $v_T(CallArrivalTime)$ are projected onto the line-of-sight $p_T(CallArrivalTime) - p'_B(EchoArrivalTime)$.

D. Biological Background

In may latitudes, the air is filled with insects which many birds prey on during the day. However, visually orientating birds give up hunting when darkness falls and the richness of the night sky is exploited by a mammalian group of skilled fliers: bats [Neuweiler 90]. Birds never developed a substitute for vision; however, bats successfully use echolocation to catch prey in complete darkness. Indeed, insectivorous bats are the primary consumers of nocturnal insects (bats may consume up to 100% of their own body weight per night) and are thought to play a major role in regulating nocturnal insect populations and in transporting nutrients across the landscape [Kunz & Pierson 94].

Traditionally bats are placed in the order Chiroptera (*cheir* "wing" *pteron* "handed"), with living species arranged in two sub-orders: Megachiroptera and Microchiroptera. (They are the only mammals capable of true flight: flying squirrels and lemurs are gliders and lack the morphological and physiological adaptations which enable true flight.) Most of the approximately 150 species of Megachiroptera do not echolocate. The one exception, *Rousettus aegyptiacus*, echolocates using clicking sounds generated by the tongue. *R. aegyptiacus*, like other megachiropteran bats, has a plant-based diet and the available data suggests that *R. aegyptiacus* uses orientation sound to find its way home in the darkness of cave hollows. All of the species of Microchiroptera studied to date use vocalisations produced in the larynx to echolocate [Suthers 88]. The classification of all bats into one order Chiroptera implies that they are monophyletic, which allows two explanations for the appearance of the echolocation trait: (1) it was an ancestral trait lost in Megachiroptera and re-acquired in one species or (2) it evolved independently in both the Mega and Micro bat sub-orders. There has been substantial support raised over the past 100 years for the alternative that bats are diphyletic. Comparative analysis of cervical vertebrae and visual pathways suggest that Megachiroptera may share an immediate common ancestor with the Primate order and Microchiroptera are more closely related to Insectivora. (See review in [Rayner 91].)

This Appendix is intended to provide some background information describing the sound production mechanisms employed by high duty-cycle bats (Section D.1) and to give an overview of the major auditory nuclei discussed in the main body of the thesis (Section D.2).

D.1 The faces of echolocation

If one applied Edgar Allen Poe's aesthetic: "There is no exquisite beauty without some exaggeration of proportion" to bats, microchiropteran species might rate — along with elephants and giraffes — as some of the more magnificent creatures on the planet. Many taxa have evolved fleshy or bristly ornaments around with the nose and mouth which play role in directing and shaping their calls. Some of the more exotic structures resemble leaves (*e.g.* in Phyllostomidae, Megadermatidae, Hipposideridae, see Figure D.2) and horseshoes (Rhinolophidae, see Figure D.1). Conservationists argue that the Western worlds' fear and loathing of bats might be alleviated through a better understanding of the function of microchiroperan faces (in the same way that the trunk of an elephant is an endearing feature).

The mechanisms by which echolocators produce (and beam-form) their signals have generally received less attention than has the reception and processing of reflected signals by the auditory system. A complete picture of sound production in bats requires a knowledge of source characteristics, sub- and supra-glottal acoustics and the radiation function [Hartley & Suthers 90]. Ultrasonic vocalisations are produced by vibration of fine *vocal membranes* — *i.e.* very thin vocal membranes which lie along the vocal folds and which are unique to echolocating bats [Suthers 88]. However, the laryngeal source in bats operates in a manner roughly similar to the human larynx in that sound is produced by the action of Bernoulli forces generated when the membranes are adducted into a stream of respiratory air [Griffin 58, Suthers & Fattu 73, Suthers 88, Hartley & Suthers 90]. The source spectrum is then filtered during transmission through the supra-glottal vocal track and (possibly) radiation through the mouth or nostrils.

The high duty-cycle bats producing long CF calls are particularly interesting from the point of view of vocalisation because one or more harmonic components are missing from their calls. For example, *Rhinolophus hildebrandti* achieves a pure second harmonic emitted pulse by means of vocal track resonances that strip the fundamental and third harmonic from the sounds generated at the larynx [Hartley & Suthers 88]. Furthermore, the sublaryngeam chambers characteristic of Rhinolophidae support a tracheal standing wave system which improves the efficiency of their vocalisations [Suthers 88]. The underlying mechanism may involve the return of back propagated sound to the larynx with a phase so related to the vocal membrane motion that it increases the effective transglottal pressure [Suthers 88].

Factors which affect the directionality of the emission differ between the orally and nasally emitting Microchiroptera. In bats that emit echolocation sounds orally, the buccal cavity probably acts as a horn, giving impedance matching and, also, directing the sound by increasing the effective dimensions of the source [Strother & Mogus 70]. In nasally emitting bats, the noseleaf may play a similar role because the flanging of the nostril that it introduces may reduce internal reflection from the end of the vocal tract (whose diameter is much smaller than the wavelength of echolocation cries). Nasal emitters also appear to make use of interference between their two nostril sources to shape sound in the horizontal dimension. One can see this by modelling the nasal emitters as a pair of isotropic sources, vibrating in phase, separated by a distance b

$$D(\theta) = 2exp[j\frac{sin(\theta)\pi b}{\lambda}] + exp[-j\frac{sin(\theta)\pi b}{\lambda}]$$

In *Rhinolophus ferrumequinum*, the nostrils are separated by a half-wavelength at CF_2 [Möhres 53, Sokolov & Makarov 71, Schnitlzer & Grinnell 77] which causes destructive interference laterally and constructive interference forward — thereby automatically reducing lateral sound emission relative to forward sound emission. A highly significant relationship, independent of both body size and genus, has recently been found to exist between noseleaf width and the wavelength of the dominant echolocation frequency emitted by 14 rhinolophid and hipposiderid species [Robinson 96].

The nostrils do not behave as isotropic dipole sources in the vertical direction and the noseleaf (a typical feature of nasal emitters) is generally thought to contribute to the directionality of the system by supplementing the horizontal beam-focusing effect of nostril interference. This phenomenon has been demonstrated in *R. ferrumequinum* [Strother & Mogus 70, Sokolov & Makarov 71, Schnitlzer & Grinnell 77]. Likewise, removal of the lancet (*i.e.* a pointed, erect structure located above the nostrils and attached only by its base) doubles the initial half-amplitude vertical beam-width in *Carollia perspicillata* — leaving sound directivity in the horizontal dimension unaffected [Hartley & Suthers 87]. Some species can rotate the noseleaf structure several degrees around a vertical axis and move the tip of the lancet backwards and forwards, possibly allowing beam scanning without moving the head. Varying the phase of the emission between the two nostrils has also been suggested as a possible beam scanning mechanism.

The following three subsections briefly describe the noseleaves and other anatomical features of the Old World and New World species of high duty-cycle bats.

Rhinolophidae ("Horseshoe bats")

This Old World family consists of 69 species, 1 genus (taxonomy according to that used in [Nowak 94][1]). Rhinolophidae are found in tropical and temperate regions from Europe to Japan. (Species native to temperate regions hibernate in winter.) They live in groups and roost in damp, dark places, such as caves. Rhinolophidae usually hunt insects within 6 meters of the ground, and will also feed on the ground.

Most species are approximately $3.5 - 11$ cm long (without the $2.5 - 4.5$ cm tail) and weigh $5 - 30$ grams. They are commonly referred to as "horseshoe bats" due to their large, complex noseleaf — whose lower section resembles a horseshoe. The shape and arrangement of noseleaf varies somewhat between species, but, typically, it covers the upper lip, surrounds the nostril, and has a central notch in the lower edge. Both the horseshoe and lancet are flattened from front to back. The sella, located between the lancet and the horseshoe, is flattened from side to side; it is connected at its base by means of folds and ridges.

The ears of rhinolophids are large, mobile and lack a tragus.

Hipposideridae ("Old World leaf-nosed bats")

[1] Taxonomies differ amongst authors.

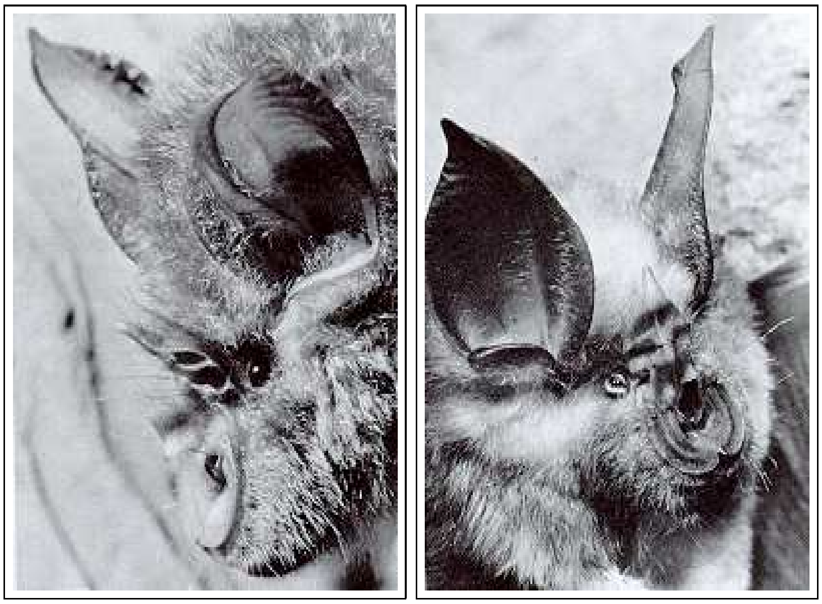

Figure D.1: **Rhinolophidae:** *Rhinolophus ferrumequinum.* (Photos from *Walkers Bats of the World* [Nowak 94].)

Hipposideridae are closely related to the Rhinolophidae and are sometimes included within that family as the sub-family Hipposiderinae. These bats are found in the tropics from Africa to Australia. Most of the 63 species (9 families) are said to be gregarious and shelter in caves or similar roosts. They feed on insects which they catch in flight. Hipposideridae are *circa* $3 - 11$ cm; the tail is either entirely lacking or, when present, measures up to 6 cm. They are characterised by a muzzle with an elaborate leaf-like outgrowth of skin.

This noseleaf consists of an anterior horseshoe-shaped part (which sometimes contains smaller accessory leaflets) and an erect transverse leaf. The latter corresponds to the lancet in the noseleaf of the Rhinolophidae and is usually divided into three cell-like parts, the apices of which may be produced into points. The noseleaves of Hipposideridae lack a sella. The ears are well developed and lack a tragus.

Mormoopidae

The 8 species (2 genera) of this New World family were formerly considered to be a sub-family of the Phyllostomidae, with the name Chilonycterinae. Mormoopidae are generally restricted to tropical habitats below 3000 m. They are gregarious cave dwellers (sometimes rooting in very large colonies) and eat exclusively insects.

Mormoopids differ most markedly from the phyllostomid bats in that they do not possess a noseleaf. Instead, the lips have been expanded and ornamented with various flaps and folds that form a "funnel" into the oral cavity when the mouth is opened. Short, bristle-like hairs surround this funnel and may act to direct airflow toward the scoop-like mouth. The nostrils have been incorporated into the expanded upper lip; above and between them are various bumps and ridges that form a sort of nasal plate.

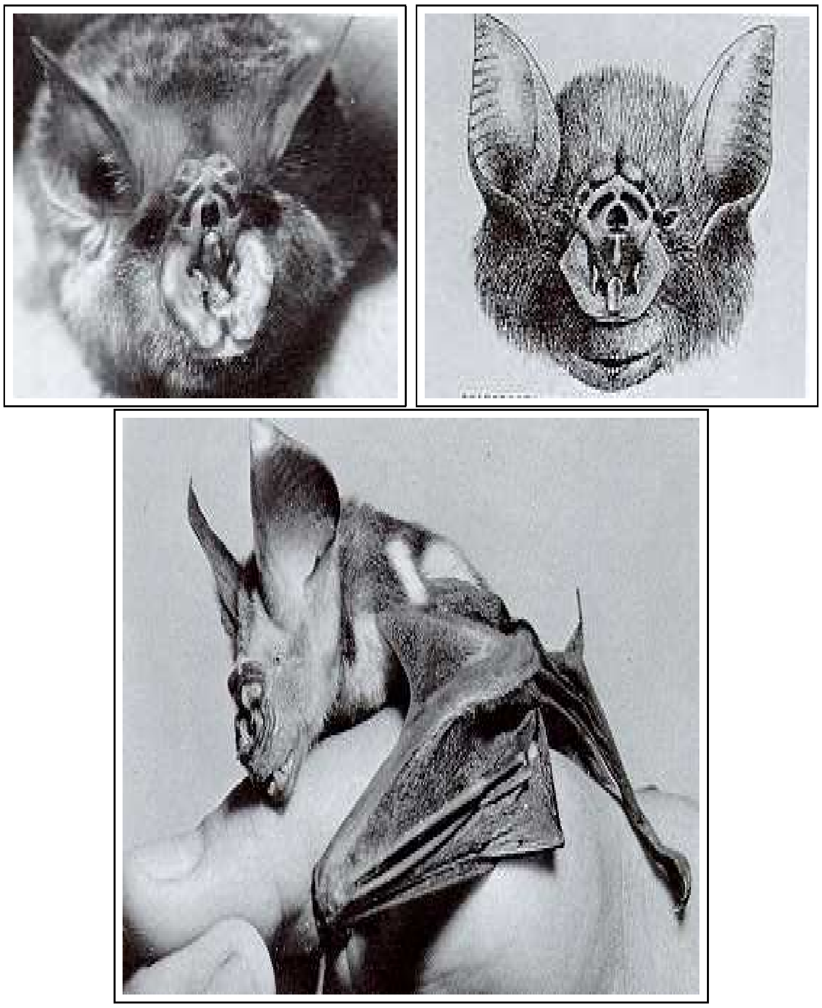

Figure D.2:
Hipposideridae: *Rhinonycteris aurantius* **(top) and** *Hipposideros diadema***.** (Line drawing from the *Catalogue of the Chiroptera* in the collection of the British Museum. Photos from *Walkers Bats of the World* [Nowak 94].)

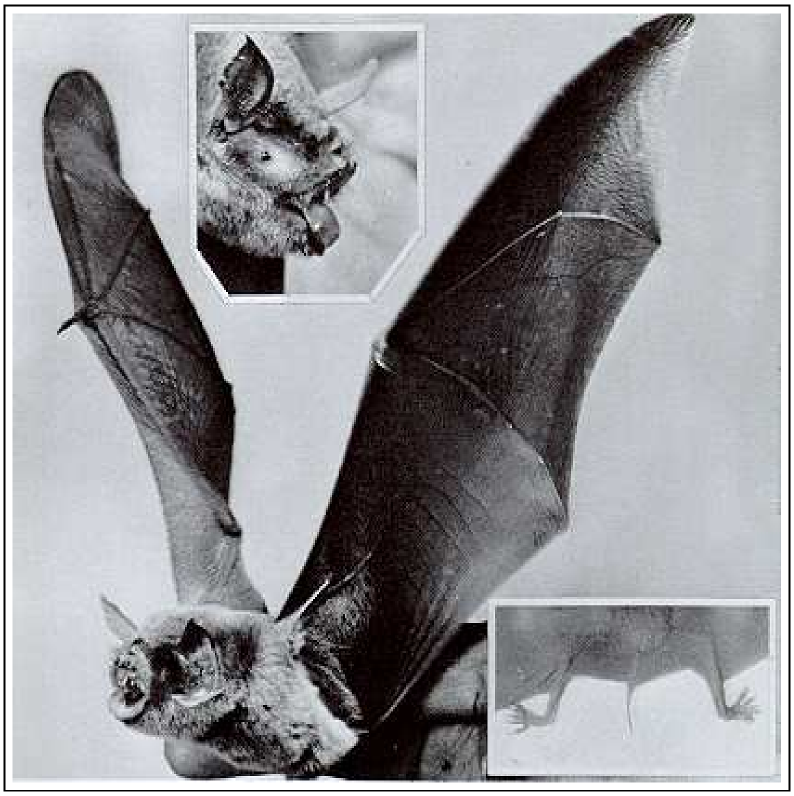

Figure D.3:
Mormoopidae: *Pteronotus parnellii.* (Photos from *Walkers Bats of the World* [Nowak 94].)

The tragus of the Mormoopidae is different from that of any other group of bats. It varies in the different genera from a simple lanceolate structure to one with a secondary fold or skin that lies at a right angle to the main longitudinal axis of the structure. This secondary fold is barely more than a pocket-like structure in the cranial edge of the tragus in *Pteronotus parnellii.*

D.2 Aspects of the bat auditory system

The bat auditory nervous system contains the same basic elements as that of other mammals: the same nuclei, cell types, synaptic connections and pharmacology. (See wiring digram of Figure D.4.) The requirements of echolocation have led to a refinement of general mammalian mechanisms as opposed to the evolution of novel ones. Here I briefly review the specialisations in structure, organisation and physiology that are critical to the echolocation mechanisms discussed in this thesis. (A one-stop, up-to-date collection of review papers covering various auditory nuclei in echolocating bats can be found in [Fay & Popper 96].)

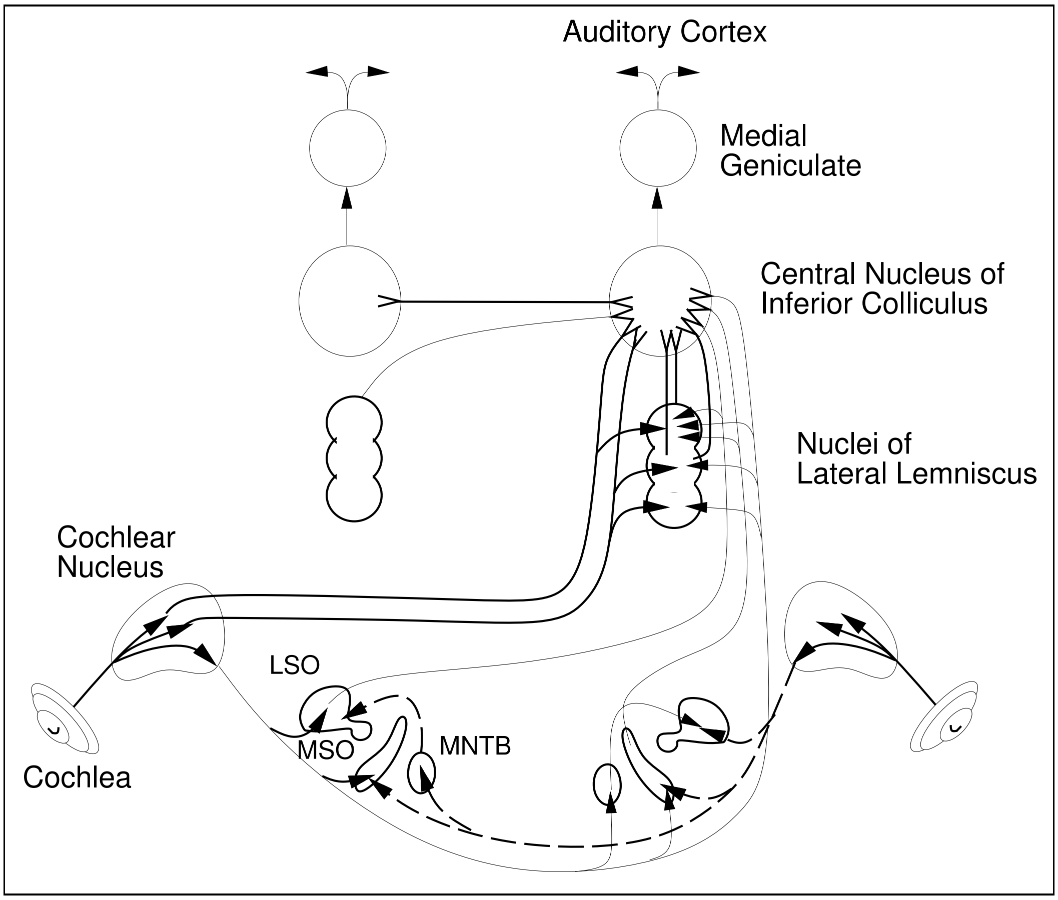

Figure D.4:
Wiring diagram showing principal connections of the mammalian auditory system (from [Grinnell 96]). LSO, lateral superior olive; MSO, medial superior olive; MNTB, medial nucleus of the trapezoid.

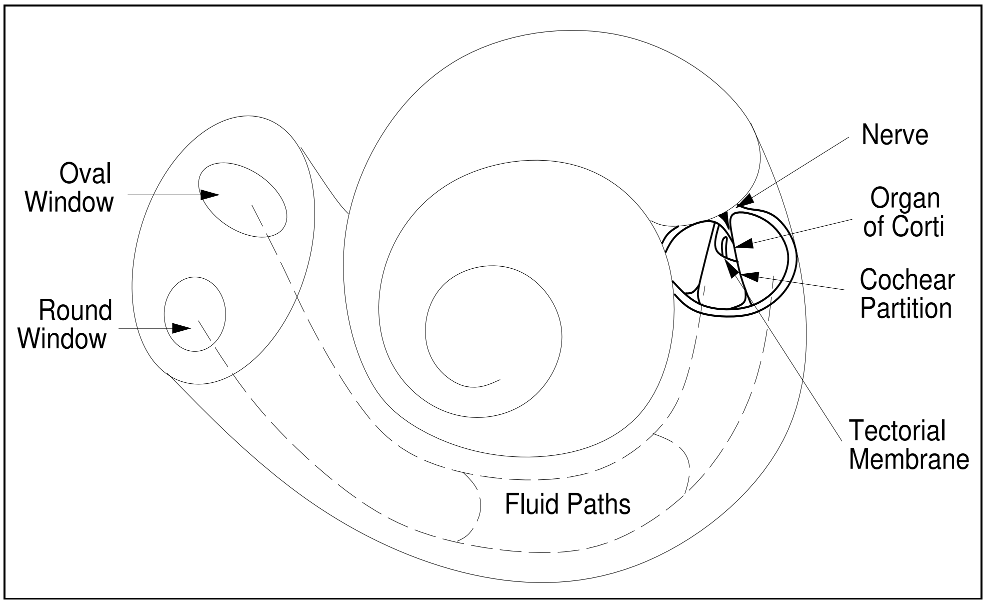

Figure D.5:
Artist's conception of the mammalian cochlea (from [Lyon & Mead 88]). Cut-away shows cross section of a cochlear duct. Dashed lines indicate the fluid paths from the input at the oval window and back to the round window for pressure relief.

Cochlea

The development of mechanisms to analyse low-level signals at many different frequencies is a relatively recent step in animal evolution, and highly developed cochleae are only found in vertebrates. As a specific mammalian adaptation, the middle and inner ear have an extended sensitivity in the high-frequency range for which cochlear macro-mechanical (*i.e.* fluid mechanical) and micro-mechanical (*i.e.* structural) specialisations have developed. Moreover, the cochlea hosts dynamic processes which enhance frequency tuning and allow it to compress the large dynamic range of acoustic pressure variations that enter the ear into the much smaller dynamic range that can be processed by the auditory neurons.

The mammalian cochlea is a spiraling, fluid filled tunnel where acoustic signals are converted into the neural code carried by the auditory nerve (see Figure D.5). The mechanical-neural transducer of the hearing system is the *hair cell*, which evolved from the flow-sensing lateral line organ of fishes and also is used in the internal equilibrium system of the inner ear [Lyon & Mead 88]. This cell detects bending motion of its hairs (*i.e. cilia*) and responds by changing an internal voltage and by releasing a neurotransmitter. Other cochlear structures perform the first level of sound separation, so that each hair cell transducer processes a uniquely filtered version of the sound entering the ear. These functions are performed as described below.

Hearing begins at the external ear, which gathers sound focuses it onto the eardrum (*tympanic membrane*). Acoustic signals enter the ear as mechanical waves which increase air pressure in front of the eardrum and push it inward, thereby moving the three bones (the *malleus*, *incus* and *stapes*) of the middle ear. The foot-plate of the stapes covers the *oval window* of the cochlea, and the movement of the stapes initiates a pressure wave in the cochlear fluid which propagates in a dispersive manner along the fluid-filled spaces of the cochlea. Thus, unlike the visual system, where input from different sources is separated early and distributed to higher processing stations along parallel pathways, all input to the cochlea converges in the middle ear to induce movement at the membrane of the *oval window*. (Figure D.5 shows the fluid path. Notice that fluid movement is enabled by the presence of the *round window* into which the fluid bulges at the end of its round-trip.)

Movement of the fluid distorts the *cochlear partition* (*i.e.* the *basilar membrane* (BM) and associated structures) — located in the turn of the cochlea. Distortion of the cochlear partition takes the form of a travelling wave which undergoes filtering as it travels from base to apex such that high frequencies attenuate rapidly near the base and the cut-off frequency gradually lowers with distance from the base to apex. The changing physical properties (*i.e.* the thickness and stiffness along the BM) control the velocity of propagation of this wave. In most mammalian cochleae, the velocity of propagation is nearly an exponential function of the distance along the membrane — starting high at the thick and stiff basal end, and decreasing toward the thinner, more flexible apical end. Lyon and Mead offer an intuitively appealing way to think about this [Lyon & Mead 88]:

> *"When the sound is a sinusoid of a given frequency, it vibrates the basilar membrane sinusoidally. The wave travels very fast at the basal end, so that it has a long wavelength. A constant energy per unit time is being put in, but, as the wave slows down and the wavelength gets shorter, the energy per unit time builds up, so the basilar membrane gets more and more stretched by the wave. For any given frequency, there is a region beyond which the wave can no longer propagate efficiently. The energy is dissipated in the membrane and its associated detection machinery."*

The upward and downward motion of the basilar membrane laterally displaces the *tectoral membrane* (TM) located directly above it. (The TM is often considered as a rigid beam providing stiffness and mass for shearing displacement or, alternatively, as a second resonator superimposed on the BM resonance.) The resulting shearing action is sensed and amplified by a row of *inner hair cells* (IHCs) which sit along the edge of the BM (in a structure known as the *Organ of Corti*). (Actually, it is a series of fine hairs (*stereocilia*) — protruding from the inner hair cells themselves — which perform the transduction of shearing motion into ion current.) In full, the shearing action causes fluid to travel back and forth across the tops of these fans of cilia, bending them to and fro. When they are bent in one direction hair cells stimulate the primary auditory neurons to fire. When bent in the opposite direction, no firing takes place. Note that by virtue of this firing property and their position along the BM, each hair cell produces a half-wave rectified version of an essentially bandpass filtered version of the original sound. Nerve fibres that originate in the inner hair cells carry this information into the central auditory system.

In addition to a row of IHCs, mammalian cochleae usually possess three rows of *outer hairs cells* (OHCs) which run parallel to the former along the length of the BM. Outer hair cells send relatively few afferent nerve fibres back into the auditory system, but, a large number of efferent fibres from higher auditory centres synapse on these cells. If not inhibited from higher centres, the outer hair cells will provide positive feedback in that if they are bent they respond by pushing harder in the same direction. These dynamic processes can put enough energy back into the BM that it will actually oscillate under some conditions. It is argued that the outer hair cells are used more like muscles which reduce the inherent damping of the basilar membrane when sound input is otherwise to weak to be detected [Kim 84, Lyon & Mead 88]; however, the role that they play in hearing is still debated.

In bats, the BM is longer (relative to body weight) than in other mammals and the organisation of the anchoring system for the BM reflects adaptations for high-frequency hearing [Kössl & Vater 96]. The claim that cochlear size is more strongly determined by functional requirements than taxonomic relationships is further realised by the non-taxonomic group of high duty-cycle bats whose cochleae are larger than other less specialised bats. Cochlear specialisations are indeed dramatic in high duty-cycle bats. In *R. ferrumequinum* [Bruns 76a, Bruns 76b] and *P. parnellii* [Henson 78], for example, an abrupt decrease in BM thickness occurs just basal to the place of representation of the second CF echolocation harmonic (CF_2). Just apical to this discontinuity, the stiffness gradient is exceptionally light — giving rise to an expanded frequency mapping. In both species, the expended representation around the CF_2 coincides with dense innervation from the auditory nerve. Within this specialised region, nicknamed the "acoustic fovea" [Schuller & Pollak 79], neurons are narrowly tuned (peak measured $Q_{10dB} = 400 - 500$).

Over-representation does not appear to be necessarily related to enhanced tuning. For example, in *P. parnellii* the sites of CF_1 and CF_3 along the BM show enhanced tuning without the morphological specialisations comparable to CF_2 region. Moreover, the expanded regions extend a few kilohertz (both apically and basally) into frequency bands where neural tuning is poor. In rhinolophids, narrow tuning has been attributed to hydro-mechanical specialisations of the cochlea [Bruns 76a, Bruns 76b]. In *P. parnellii*, there is a prominent resonance [Suga *et al.* 75] and strong evoked acoustic emission (OAE) [Kössl & Vater 85] approximately $400 - 900$ Hz above CF_2 which appears to contribute to the narrow tuning. It is possible that these two unrelated species achieve similar filtering precision by different cochlear mechanical adaptations [Neuweiler 80].

The picture becomes even more interesting (and complicated) when hipposiderids are considered. For example the acoustic fovea of *H. speros* is not centred around the CF_2 component emitted by hand held bats, but skewed towards lower frequencies encompassing about 2/3 of the frequency range of the final FM sweep [Neuweiler *et al.* 87]. Thus the acoustic fovea is broader than that of rhinolophids and Neuweiler suggests that minimal duration of the call may correlate with the sharpness of the filter across various species [Neuweiler *et al.* 80].

The hair cell transducers of echolocating bats deviate in significant ways from the general mammalian scheme in order to facilitate high gain and high frequency hearing. First, the size of the OHCs and their stereocilia is decreased, which may guarantee a high-speed micro-mechanical amplification of low-level signals by active OHC processes. Furthermore, imprints of the IHC and OHC stereocilia have been observed throughout the cochlea (*e.g.* in *R. rouxi* [Vater & Lenoir 92, Vater *et al.* 92]). While it is widely accepted that the OHC stereocilia in mammalian cochleae are firmly embedded in the subsurface of the TM, the mode of linkage of the IHC stereocilia is disputed. In some mammals, no linkage has been observed (*e.g.* in the chinchilla [Lim 86]) and in others, the imprints of the stereocilia in the subsurface of the TM are confined to the basal cochlea (*e.g.* in the rat [Lenoir *et al.* 87]). It is speculated that the extended linkage of the IHC in bats is relevant for IHC stimulation at high frequencies.

Brainstem auditory nuclei

The mammalian auditory system is a series of parallel frequency tuned pathways, beginning with the tonotopic distribution of eighth nerve endings along the BM. This organisation is maintained as a tonotopic map in (almost) every division of every nucleus en route to the cortex. In addition, information in each frequency tuned pathway is duplicated and sent through multiple channels which extract different features of the signals.

As shown in Figure D.4, information from the auditory nerve is sent to three different divisions of the cochlear nucleus which, in turn, send parallel projections to other brainstem nuclei.

One of the first places signals arrive is in the superior olivary complex (SOC). The SOC is large and well developed in bats and appears to contain the same elemental structures as non-echolocating mammals. However, there is some controversy as to whether the large and specialised medial superior olive (MSO) can be considered as an analogue of the general mammalian MSO. Relative to other mammals, wherein the MSO receives approximately equal excitatory input from the contra- and ipsi-lateral cochlear nuclei, that of bats receives primarily contralateral input. (Monaurality of the MSO is most striking in high duty-cycle bats.) The typical mammalian MSO appears to provide a measure of target azimuth in the excitation of subpopulations of neurons tuned to different delay disparities. In bats, it is believed that the extraction of binaural timing has shifted to higher cell populations [Covey *et al.* 91] and that the MSO may have evolved a new function associated with the demands localising high frequency sound reflected from flying prey in a three dimensional world. The speculated roles relevant to this thesis are discussed below.

- **Prey signature identification.** MSO neurons may encode information about echo envelope modulations. Nearly all MSO neurons are capable of following AM with ON/OFF responses: approximately one third of MSO neurons exhibit bandpass selectivity for low AM rates while the remainder have low-pass filter characteristics with upper limits between 100 and 500 Hz (with most between 200 and 300 Hz). The upper AM filter limits for MSO neurons are considerably lower than those found at the level of the cochlear nucleus which show synchronised firing up to 1000 Hz (see Section 4.1 or [Vater 82]). (Grothe describes how this bandpass filtering may be created by timing relationships between excitatory and inhibitory signals arriving from the contralateral ear [Grothe 94].) An as-yet-untested hypothesis is that MSO neurons are selective for the wing-beat rates of insect prey and that their activity could signal the presence of prey worth pursuing [Moss & Casseday 96].

- **Prey signature-based localisation.** Some MSO neurons (exclusively in the CF_2 range) elicit a phasic ON-OFF response to sharp envelope transients. It has been suggested that glints off fluttering wings may provide sharp transients of this sort and that this signal provides a timing marker — one in each MSO — which could come together (in the DNLL or IC, for example) to provide localisation cues based upon inter-aural envelope comparisons [Lesser *et al.* 90, Covey *et al.* 91].
- **Elevation cues.** The MSO might serve in target elevation localisation by transmitting spectral cues to the inferior colliculus where they would be integrated with inter-aural intensity differences encoding target azimuth.

The mammalian lateral superior olve (LSO) receives excitatory input from the ipsilateral CN and inhibitory input from the contralateral CN so as to perform the binaural intensity comparisons characteristic of the mammalian LSO. The LSO itself does not appear to contain a "map" of the frontal sound field in that the inter-aural level differences which neurons in an individual LSO are capable of encoding do not cover the necessary range of values which might be generated from targets located at arbitrary azimuth angles. Moreover, there is no evidence of a systematic topographic map of IIDs in the CF_2 region of the LSO [Covey *et al.* 91]. Thus, the map may be formed in the ICc as a result of convergence from two LSOs. However, the switching of response between those two LSOs (representing IIDs in one hemisphere) as a target moves across a vertical mid-line could provide a steering signal to pre-motor area.

The complex and hypertrophied nuclei of the lateral lemniscus are an important station for processing temporal information. Neurons in the dorsal nucleus (DNLL) show strong facilitation to the second of two similar sounds and the "best delay" intervals encoded by this population are within the range needed for echolocation [Covey 93]. In the ventral nucleus of the lateral lemniscus (VNLL) neurons receive input predominantly from cochlear nucleus axons and respond typically by firing just one spike at the onset of a sound. Thus, this population appears ideal for detecting wavefront arrival times based upon response latencies which are remarkably invariant to changes in frequency and intensity [Covey *et al.* 91]. Other parts of the nuclei of the lateral lemniscus follow rapid frequency and amplitude changes with more sustained firings — see review in [Moss & Casseday 96].

Inferior Colliculus

In echolocating bats, the IC is relatively huge — this, and its position near the dorsal surface of the skull in most species, makes it an intensely studied nucleus. All information travelling through the brainstem converges at the IC. The central nucleus (ICc) receives input from at least 10 lower centres as well as descending input from higher ones.

Most neurons receive binaural input (either excitatory input from both ears (EE) or excitatory input from one ear and inhibitory input from the other (EI)) and have narrower tuning curves than those found at lower centres. A large portion of EI neurons have *closed* tuning curves and, thus, appear to be dedicated to localising faint echoes from, perhaps, prey. Moreover, many IC neurons are *level tolerant* — *i.e.* tuning curves of many IC neurons can have nearly vertical cut-offs ($Q_{10dB} \approx Q_{20dB} \approx Q_{30dB}$), as opposed to the "V" shaped tuning curve which characterises neurons in lower auditory centres. This enables them to respond to particular frequencies independently of intensity.

These neuronal features are built up through different combinations of inhibition and excitation. Indeed, it has been suggested that the response of many neurons in the IC is temporally bounded by an inhibitory frame that defines a window in which excitatory responses can occur [Casseday & Covey 96]. Moreover, this time-frame may be considerably delayed with respect to the actual stimulus time of arrival — *e.g.* neuronal latencies varying from 5 ms to more than 30 ms in the IC (as compared with $3-6$ ms at the lateral lemniscus) [Casseday & Covey 92]. Casseday and Covey argue that this slowed rate of response of IC neurons is related to motor actions — *i.e.* the IC matches its rate of output to a rate that is appropriate for initiating and controlling movement [Casseday & Covey 96].

Auditory Cortex

The auditory cortices (ACs) of only a few bat species (two high duty-cycle, and two low duty-cycle) have been studied in detail. Differences are significant and they tend to be discussed separately. Suga and his colleagues have carried out some remarkably elegant studies of AC neurons in *P. parnellii* which have shown that the AC of high duty-cycle bats contain several specialised areas which perform different types of processing [Suga 90b, O'Neill 96]. As in other centres, the acoustic fovea covers approximately one third of the tonotopic portion of the auditory cortex. In what Suga *et al.* refer to as the DSCF area (*i.e.* the foveal portion of the tonotopic map dealing with Doppler-Shifted Constant-Frequency signals), neurons with best frequencies (BFs) at and just above CF_2 are further tonotopically organised into concentric rings. In this architecture, lower frequency neurons lie at the centre and BF's increase progressively along radial directions. Superimposed on this concentric high resolution frequency map is a map of amplitude — arranged radially with neurons having best amplitudes ranging between approximately $10-90$ dB. Finally, if we view the DSCF surface as a pie, a rather large piece receives EE inputs while the remainder receives EI inputs. The result is that the DSCF area has axes representing either target velocity or subtended target angle and is divided into areas (pie pieces) which appear suited to either target detection or localisation.

In *P. parnellii*, there are also AC regions dedicated to performing comparisons of harmonics. For example, in the so-called CF/CF region, neurons receive inputs from lower centres tuned to combinations of CF_1/CF_2 and CF_1/CF_3. These neurons are extremely narrowly tuned, level tolerant, and laid out so that the BF of neurons tuned around each harmonics increases gradually along orthogonal axes. Each cluster of combination sensitive neurons represent target velocities varying from -2 to -9 ms by virtue of the combination of Doppler shifted echo frequencies to which they respond most strongly. There are also clusters of cortical neurons selective for combinations of FM_1 with FM_2, FM_3 and FM_4. These units are laid out according to the most effective delay of the second signal with respect to the first — with best delays ranging $0.4-18$ ms, covering most of the range of delays of potential targets.

Suga *et al.* claim that this "heteroharmonic sensitivity" is one mechanism by which the SONAR information is protected from masking by conspecifics [Suga 90b]. The AC of *R. ferrumequinum* shows, qualitatively, similar organisation to that of *P. parnellii*;

however, the harmonically related combinations of signals are not as apparent.