This is What it Is Like to Be a Bat

David Scott Farrar
Department of Cognitive Science
University of California, San Diego
May 26, 1995

Integrative Paper Committee:
David Zipser
Martin Sereno
John Batali

This is What it Is Like to Be a Bat

Cognitive Science Integrative Paper

This is What it Is Like to Be a Bat

Introduction: Nagel's Question

Nagel asks the question "What is it like to be a bat?" (Nagel, 1974). Nagel's question is partially a rhetorical device for making a philosophical point about qualia and points of view. This paper will use Nagel's question as the starting point for telling a story about how a species of animal quite alien to us perceives and conceives of its world. The "answer" given by this paper will be deliberately naive and empirical. From this perspective, asking what it is like to be a bat involves answers to scientific questions: How does the bat gather information about its world? What are some important things in a bat's life that are likely to be on its mind and reflected in its behavior? How does the bat construct and use representations of objects to form a model of the world? How are these realized in their neural correlates? Nagel's philosophical point is roughly that scientific, empirical study may be limited in what it can tell us about someone else's point of view. This paper will instead assume that empirical evidence can "answer" Nagel's question -- that it can tell us what it is like to be a bat -- and that any answer to his question must involve empirical research into bats and what they do.

Bats are interesting because they sense their world by primarily by hearing it, rather than seeing it; they can do some fairly impressive and smart things, even though most species have very impoverished visual capabilities and small, mouse-sized brains. Bats have impressive sensory specializations different from our own, complex behavior comparable to higher mammals in many respects, but are an example of a small animal with relatively tractable and accessible neural structures. Cognitive science is in the business of explaining intelligent behavior, and echolocating bats provide an excellent domain in which to study several key issues such as representation and computation. So let the story begin!

Bat Basics: The Echolocation Pulse

Almost all nocturnal bats obtain information about their environment by emitting extremely high frequency sounds (from 30-120 KHz), and listening to echoes from objects nearby. Although there are significant interspecies and individual variations in these echolocation pulses, there are two primary parts to a typical pulse. In the CF ("constant frequency") portion, the fundamental frequency and its harmonics hold steady; in the FM ("frequency modulated") portion, the fundamental frequency and its harmonics sweep downward (Figure 1). The spectrogram of a typical echolocation pulse thus looks much like a human vowel diphthong. In (Figure 1) below, the CF and FM harmonic components of both the echolocation pulse and the returning echo are shown. Keeping in mind this basic echolocation tool, we can begin to piece together a story of what the bat might be sensing from its "biosonar" pulses, and what it can derive from them.

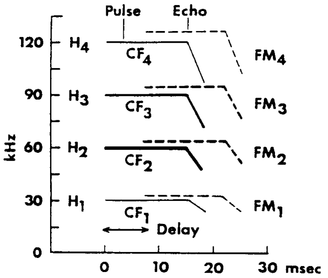

Figure 1: The Basic CF-FM Echolocation Pulse
(Adapted from Suga, 1984)

Marr's Computational Approach

In order to begin answering Nagel's question, "What is it like to be a bat?", we need a methodology, an approach to guide us. Marr's approach to understanding the problem of vision consists of developing a set of descriptions and explanations at three levels of analysis: at the level of computational theory, the algorithmic level, and at the level of mechanisms and implementations (Marr, 1982). Applied to the problem of echolocation, the approach insists that we understand the computational problems facing a bat in its nightly hunting and navigating tasks. What kinds of information are theoretically extractable from the sound waves in the bat's environment? What is to be computed via sound processing, and why? Answering the question "Why?" these computations must be performed brings us closer to the approach of the neuroethologists, who are interested not just in a description of behavior, but explanations of behavior as well.

What kinds of information are theoretically extractable from echolocation? First, note that the distance of an object is proportional to the delay between the emission of a sonar pulse and the returning echo (Figure 2). An object's relative velocity is proportional to the Doppler shift of the returning echo. The subtended angle of the object (its size divided by its distance) is proportional to the amplitude of the returning echo. The "true" size of the object can then be computed from the amplitude of the echo, some knowledge about the object's sound reflectance efficiency, and the echo delay. Internal motions due to some flexibility of an object's surface can cause fluctuations in the average Doppler shift of the whole object. The echo of an object located in a specific spatial direction will cause a complex but systematic pattern of interaural timing differences, interaural intensity differences, and spectral cues in a binaural hearer with a appropriately structured pinnas. Information about surface texture of an object (and even some the object's internal structure) can be extracted from the spectral properties of the returning echo (Suga, 1984). The bat will hear a complex interference pattern from textured surfaces, as echoes from parts of the surface with different distances overlap one another.

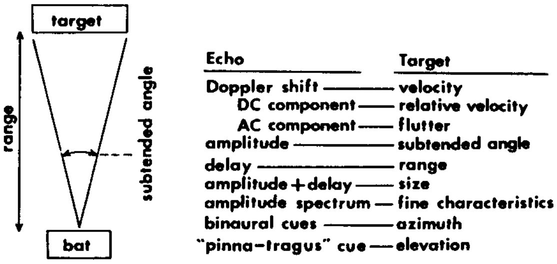

Echo	Target
Doppler shift	velocity
DC component	relative velocity
AC component	flutter
amplitude	subtended angle
delay	range
amplitude + delay	size
amplitude spectrum	fine characteristics
binaural cues	azimuth
"pinna–tragus" cue	elevation

Figure 2: Echolocation Cues
(Adapted from Suga, 1984)

What kinds of algorithms might be used to extract these kinds of information? For a quality that is mostly proportional to an information bearing parameter, measurements of that parameter can be scaled and normalized to form a simple representation of the quality. Thus, delay between the emitted sonar pulse and the echo from an object can be measured, scaled, and normalized to form a representation of object distance. There are a couple ways of performing this kind of comparison; one would be to 'timestamp' the sonar pulse and its echo, and subtract the two; another would be to systematically delay a representation of the sonar pulse, and detect a coincidence between the delayed pulse and the returning echo. The amount the pulse has to be delayed then becomes the measured echo delay. The Doppler shift of the returning echo can be measured and compared to the emitted sonar pulse, scaled, and normalized to form a representation of object velocity. Note that this requires an algorithm which can perform comparisons between frequencies, so that a shift from one frequency to another can be noted. The amplitude of the returning echo can be scaled with respect to the amplitude of the emitted pulse to form a representation of the subtended angle of the object; this requires the ability to discriminate between specific echo amplitudes (Suga, 1984).

The algorithm underlying the computation of spatial location (azimuth and elevation) is likely to involve a complex but systematic mapping between interaural time, intensity, and frequency cues. One algorithm to compute azimuth involves a comparison of interaural phase differences within a narrow frequency range; these phase differences can be compared across several frequencies to eliminate phase ambiguities and compute an actual time difference. This time difference can then be normalized to yield a measurement of sound azimuth. Another algorithm to compute both azimuth and elevation takes advantage of spectral cues which result from the directionality of the pinna. In particular, if the pinna's directionality changes as a function of frequency, a binaural comparison of the spectral components of the sonar pulse and echo should detect differences that are systematically related to the spatial location of the object that created the echo. Thus, the combination of time, intensity, and frequency difference information can yield the spatial direction of an echo (Fuzessery, 1986).

Grinnel has suggested that the Doppler shift of a sonar echo may yield information about the angle between the echo and a sonar emitter's direction of movement: objects located in front will generally be shifted by an amount proportional to the emitter's velocity, while objects located towards the periphery will generally be shifted by lesser amounts (Grinnel, 1973). Although this does not determine the direction of an auditory object, it may provide a constraint which can be combined with others to completely disambiguate spatial direction.

Other algorithms may allow the construction of a representation of surface texture or structure, or perhaps even internal object structure. How might these algorithms and representations be implemented neurally? Neural maps of several of these information bearing parameters (Suga, 1984) have been identified in the auditory midbrain and the auditory cortex of echolocating bats. Details of these maps will be discussed later in this paper. How might these algorithms and representations lead to the conception and representation of objects? The notion of an "auditory image" provides a useful intermediate step in helping us understand how representations of objects might be possible from echolocation information.

Auditory Image Perception and Scene Analysis

Auditory image perception is the determination of the sources of sounds (Yost, 1991). Auditory scene analysis is the process of extracting useful information about the surrounding environment from the sounds one hears (Bregman, 1990). The general problem is that sounds overlap one another: several sounds in the environment may happen at the same time; the intensity of one sound may obscure the intensity of others; and the sounds may be composed of overlapping frequencies. Spatial ambiguities can also arise. Understanding how a bat can separate multiple auditory "objects" from one another, and from their background, is a challenge. Although the bat can use an active sonar system to control many of the sounds it hears, and humans only passively listen to their environment, both face the same fundamental problem, and may use the some of the same methods to solve the problem. How can a bat or a human segregate and associate spectral information with objects and events in the environment?

One example is illustrated below (Figure 3). The graphs are sound spectrograms, which show how the intensity of sound changes over time and frequency. The first graph visually shows the sound energy pattern for a cat meowing; the second is for a dog barking; the third is the summation of the two. It is easy to distinguish between the two sounds when listening to their summation, even though they overlap in time and frequency.

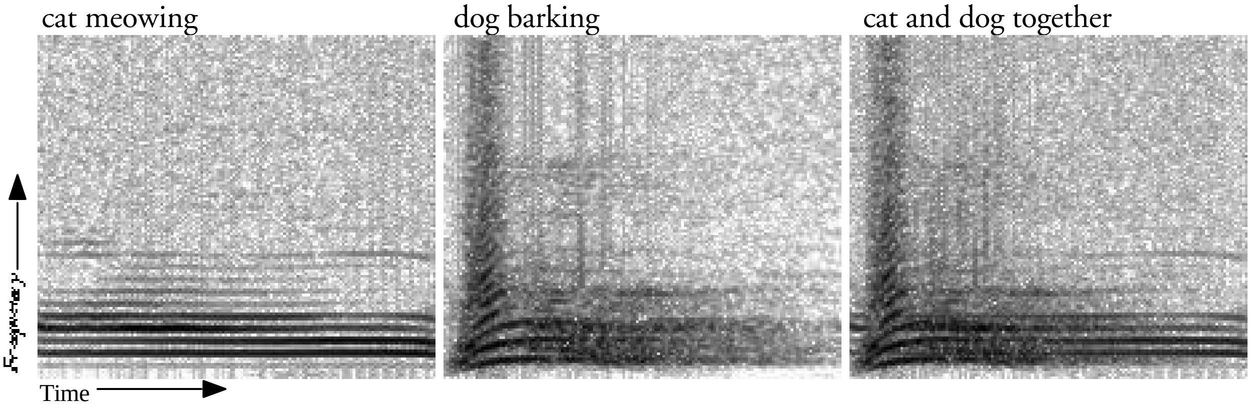

Figure 3: The Problem of Overlapping Spectra

One solution to the problem is to find a way to process sounds into some representation of the environment: to segment, segregate, combine, and integrate the sensory evidence for sounds into coherent "streams" (Bregman, 1990). The analogy between visual objects and auditory streams is quite strong. In vision, the sensory data is spatially organized intensity information from the retina; in audition, the sensory data is tonotopically organized intensity information from the cochlea. Both are changing with time. In vision, object recognition consists of dividing up the retinal data into regions, and using information about these regions and their changes to represent the object. In audition, stream segregation consists of dividing up the chochleal data into regions within the time, intensity, and frequency domains, and then using information from these regions to track and represent the stream.

Using this notion of an auditory stream, several other striking analogies with vision emerge. For example, a continuous tone may be overlapped by a brief burst of broadband noise; the auditory system has no trouble tracking the tone and the noise separately, and even if the tone is actually missing during the noise burst (Figure 4) (Bregman, 1990). This is analogous in vision to a contiguous object that is occluded by another object. Auditory streams can also be perceived from sequences of alternately low and high pitched tones, which both slowly creep upward in pitch; as the tones are played more quickly, the fact that they alternate seems to disappear, and instead one hears two simultaneous sequences of tones, one low, one high, both of which seem to rise in pitch (Figure 5) (Bregman, 1990; Yost, 1991). This is analogous in vision to the phenomenon of apparent motion, where object motion is perceived from lights which are alternately flashed on and off, as in a movie marquee. In the case of echolocating bats, there may be another kind of visual analogy. Bat echolocation pulses are discrete but rapid, much like the raster refresh of a television or a computer monitor. Humans perceive continuous motion on television and in computer games. Perhaps bats create and perceive a continuous "picture" of their environment, even though they use a discrete method to obtain information about that environment.

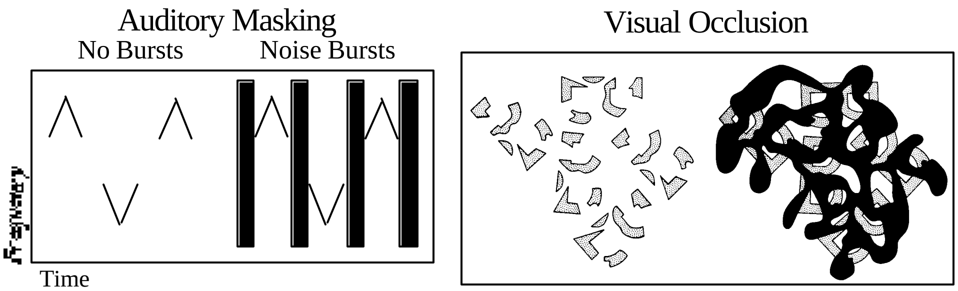

Figure 4: Auditory Masking and Visual Occlusion
(Adapted from Bregman, 1990)

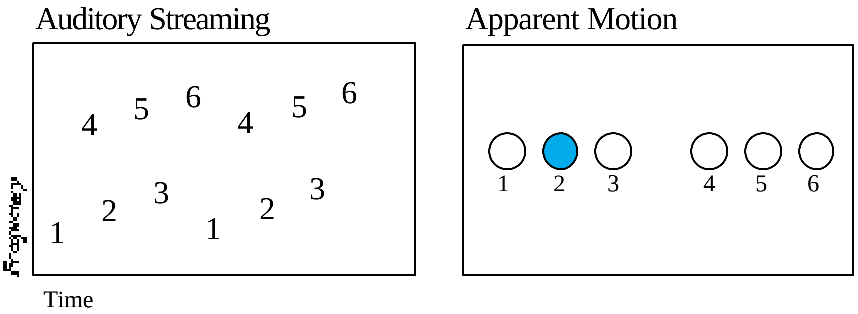

Figure 5: Auditory Streaming and Apparent Motion
(Adapted from Bregman, 1990)

Bregman identifies two kinds of stream integration, spectral and sequential. In spectral integration, the problem is how to identify which spectral components belong to the same sound source. Large differences across the spectrum are likely to be segregated into different streams, depending on how well the patterns of spectral energy are separated, and on the presence of multiple "timbres". In sequential integration, the problem is how to maintain the identity of sounds sources through time. Rapid changes in the spectrum of a sound over time will tend to cause the auditory system to segregate the sounds into different streams. Bats must be capable of both kinds of stream integration: spectral integration of the frequency components of an object echo, and sequential integration of successive echoes from the same object.

In the case of echolocating bats, the application of auditory scene analysis is complicated by the fact that bats can control the nature of the sounds they emit, and thus influence the nature of the sounds that are reflected back to them. As we shall see, bats structure their sonar pulses so as to reduce or eliminate ambiguity in the returning echoes. In Yost's account of auditory image perception, and in Bregman's account of auditory scene analysis, spatial location is but one of many factors influencing stream segregation; sounds widely separated in space, but similar in timbre or rhythm may still be lumped into the same auditory stream. For bats, spatial location may play a much more crucial role in sound stream segregation, for simple reasons: insects must be very accurately located in order to be physically caught; obstacles and spaces must be accurately located to be avoided, or used for passage and shelter. Humans can usually rely on visual cues to provide information about the detailed spatial location of events; bats have only their hearing to give them a sense of the spatial locations of objects, and it is therefore likely that their auditory systems give more weight to spatial cues.

Bregman proposes a distinction between primitive and schema-based organization of streams. Primitive stream segregation, he contends, is innate, and happens automatically, in the absence of learning and attention. It uses cues that tend to correctly segregate wide classes of sound events. Schema-based segregation divides sounds into streams using learned or "cognitive" information; he suggests that schema-based segregation requires an attentional mechanism. He claims that both kinds of segregation are used at the same time, and that sometimes they may be in tension with one another. Among echolocating bats, the distinction between primitive and schema-based stream segregation may be quite useful. For example, much of the processing of sonar pulses and their echoes appears to be done automatically, in the absence of attention; this is why some useful neural recording experiments can be done on unconscious, anesthetized bats. In particular, the processing of spatial direction, distance, and velocity are likely to contribute to the primitive segregation of streams, as each cue provides information that is likely to distinguish between multiple objects. Thus in bats, likely candidates for primitive stream segregation include differences in sound direction, distance, and velocity, while likely candidates for schema-based stream segregation are the processing of echoes from large surfaces such as walls, the spectral profiles of specific species of moths, and the communication calls of conspecifics.

Bregman's notion of primitive stream organization is comparable to Yost's notion of sound source determination through auditory image perception and analysis. Yost's general theory is that the determination of sound sources is the basis of hearing (Yost, 1991). He proposes several auditory variables which may contribute to the formation of images: spectral separation, intensity profile, harmonicity, spatial separation, temporal separation, common temporal offsets and onsets, and coherent slow temporal modulation. It is striking how appropriate most of these variables are for understanding how bats may form auditory images through echolocation. Bats are known to use spectral and intensity profiles of echoes to discriminate between multiple targets (Simmons, et. al. 1983), and can even shape the spectral and intensity profiles of their pulses to enhance tracking of a particular auditory image (Pollak, et. al. 1989). The typical echolocation pulse is composed of a somewhat damped fundamental frequency and strong harmonics, and the bat's auditory system senses the presence of the fundamental harmonic of its pulse to distinguish its pulses from those of conspecifics (Masters, et. al. 1991). Bats are quite capable of distinguishing the spatial locations of objects (Fuzessery, et. al. 1993); echo delay is one kind of temporal separation, which bats use to determine object distances (Griffin, 1958). In a typical hunting approach, a bat may overlap outgoing pulses with incoming echoes (Griffin, 1958); thus, the bat must be able to associate common temporal onsets and offsets of sonar pulses and echoes with one another, and successive pulses may refine the auditory image (Dear, et. al. 1993). Additionally, bats are known to be keenly sensitive to coherent slow temporal modulations of the CF portion of the sonar pulse, which help them to identify particular species of moths via their wingbeat patterns (Schnitzler, et. al. 1983).

Bat Behavior: Navigating Through the World

We have examined some of the kinds of information about the world bats might gather from echolocation, but we have not yet explored why bats use this information. To understand this we must look at bat behavior. Behavioral studies of the bat's ability to navigate through its environment have been conducted for centuries, but electronic recording equipment capable of measuring the ultrasonic frequencies of its sonar have been available only in the past few decades. With this equipment, detailed studies of the bat's perceptual acuity have been possible. Bats are capable of making judgments of fine distances in the range of a few centimeters from several meters away (Griffin, 1958); some species are capable of passive sound localization to within 1° (Fuzessery, et. al. 1992); some are capable of detecting metal wire obstacles as narrow as 0.4mm (Griffin, 1958; Pollak, et. al. 1989); some are capable of discriminating between species of flying insects, and selectively hunt preferred species solely from the dynamic spectral pattern of the insect's echo (Schnitzler, et. al. 1983). The echolocation of bats can be a very precise method of gathering information about the world.

Griffin describes several experiments he devised to test the bat's ability to navigate through its environment, and detect changes in its environment (Griffin, 1958). In one experiment, Griffin placed a bat in a room with some simple obstacles, one of which was a moveable wall (Figure 6). The bat would normally roost in the larger, open part of the room, but would feed in the corner opposite the obstacles. After the bat had become accustomed to the room, and could navigate around the obstacles without touching them, the room was changed when the bat was feeding, so that one wall segment (AB) was moved to block the path the bat normally used (CD).

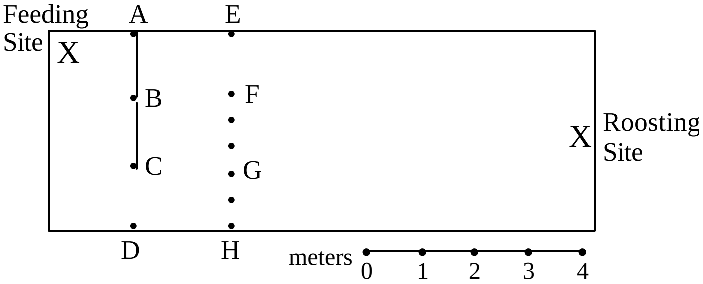

Figure 6: Moving obstacle experiment.
(Adapted from Griffin, 1958)

When it finished feeding and started to fly back to its normal roosting site, the bat ignored the new opening AB, and instead proceeded to head back to the open area along its old path, crashing into CD not just once but twice. The bat hovered about CD for a while, and eventually 'discovered' the new entrance AB. The bat did not immediately plunge into the opening, but instead tended to hover around AB a while before moving through, and back to the open area. The next feeding time, the bat proceeded from the open area to the feeding site along its old path; when it got to the new opening AB, it hovered about the entrance for several seconds before moving through. Sound detectors showed that the bat emitted echolocation pulses the whole time, even when the bat crashed into the wall CD.

There are two curious elements to this experiment. The first is why the bat blundered into the newly moved wall CD, even though it was emitting echolocation pulses, and presumably was able to sense echoes from the wall. The second is why the bat hesitated in front of the AB opening, even though there was no obstacle there, and thus no echo from the part of space the bat seemed to be examining.

What this experiment clearly demonstrates is that the bat does not navigate through its environment just from its immediate sensory impression; rather, the bat has a memory of what the world used to be like -- presumably a memory of the way the bat thinks the world is like -- and this guides the path it decides to take. It is unknown if the echoes the bat hears back from the wall CD contain enough information to locate or identify the wall, or even to warn the bat that something is "out there"; since the bat is in a familiar environment and not hunting insects, it may be shaping its echolocation pulses for purposes other than detecting new objects. It may be possible that the bat has heard an echo from the wall, but disregards the echo because the area is otherwise so familiar. Alternatively, the bat may be using its echo pulses to locate its own orientation with respect to all of the walls in the surrounding environment; it may not necessarily be attending to what is in front. Whatever the case, it is clear that the bat's behavior is guided by more than just its immediate sense of echolocation; it is also using a structured memory of its environment. When the bat collides with a wall, or hesitates when it discovers a new opening, this structured memory must be reconciled with the new echolocation information. Sometimes the world contains the structure which guides behavior, sometimes the structure which guides behavior is contained within the behaving creature.

Griffin notes from his and others' observations of bats killed from collisions with newly erected doors and buildings that bats seem to operate in two 'modes'. In the first mode, the bats are either in unfamiliar environments or hunting and seem to pay detailed attention to the echoes they receive in order to avoid obstacles or catch prey. In the second mode, the bats are usually in familiar environments, and seem to pay attention to the echoes they receive primarily to orient themselves within their environment; they otherwise seem to navigate from 'dead reckoning', and memory of the obstacles in the environment. In both modes, the bats are usually quite good at avoiding obstacles. The presence of two modes is also suggested by ear blocking experiments. Blocking both ears, of course, disables most of the bat's echolocation abilities. When only one of a bat's ears is blocked, the bat is able to navigate through rooms, avoid walls, roost, etc., but is unable to avoid small obstacles or catch prey (Griffin, 1958).

Tests of the bat's ability to navigate through its environment suggest that it is able to use its memory of the world systematically. In a second series of experiments, Griffin placed a bat in a rectangular room with a series of evenly spaced vertical metal wire obstacles down the middle (Figure 7). He prompted the bat to fly across the room by disturbing its roost on either side; he then observed if the bat collided with the wires. He tried many variations on the experiment, including different inter-wire spacings, different wire widths, orienting the wires vertically or horizontally, and wires made of different substances.

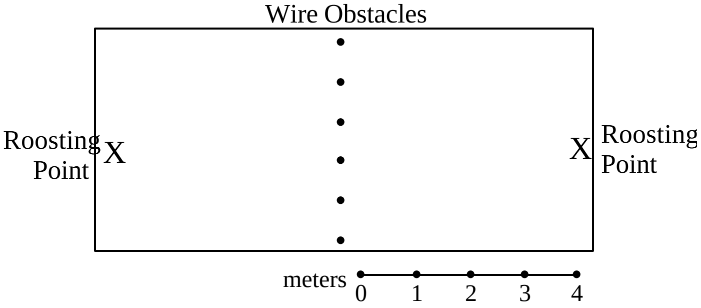

Figure 7: Wire Obstacle Experiment
(Adapted from Griffin, 1958)

He found that the bats were usually quite capable of avoiding wire obstacles, as long as the they were well-rested, i.e. not immediately awakened from sleep or hibernation. The bats were able to avoid wires as small as 0.4mm in diameter; wires made of cloth or tape; and wires oriented vertically or horizontally. The bats were able to avoid vertical wires spaced more closely than their wingspan by tilting their bodies in flight and making sharp banking turns between the wires. And after becoming habituated to this obstacle environment, when the wire obstacles were removed, the bats acted to avoid the wires, even to the point of making sharp banking turns. At all times, the bats emitted echolocation pulses. Others have extended such wire obstacle experiments by using grids of tightly spaced wires, where the bat must scrunch its wings together to fly through (Pollak, et. al. 1989). The fact that the bats have learned the spatial locations of the grid wires can be demonstrated by replacing the wires with a series of light beams and photodetectors, which the bats cannot sense: for hours or days, the bats will still retract their wings to avoid the places where the wires used to be.

These obstacle avoidance tests strongly suggest that the bat has a world model which it is using to navigate through its environment. When the bat must fly through arrays of narrow vertical wires, it can sense where the wires are in the world, and adjust its configuration or wingspan to avoid a collision. But this in itself does not demonstrate that the bat is using a model of the world; it may just be reacting to the immediate stimulus of an echo from the wire, and computing on-the-fly corrections to its path to avoid collision. No world model necessarily needs to be constructed, and the bat may not need to know or represent where it is located within its environment.

But after a while we find that the bat learns the locations of the wires, and if we remove them, the bat will still attempt to avoid the spaces where the wires used to be. The behavior of the bat cannot be explained as an immediate response to echo stimuli, because there are no wire echoes. Bats seem to have knowledge of the locations of where the wires used to be, knowledge of where their own bodies are with respect to these locations, and know how to change the configuration of their bodies to avoid these places. The fact that bats are able to manipulate this knowledge systematically -- to avoid collisions -- suggests that this knowledge is represented systematically, so that the relationships between a bat's body, obstacles, and distant objects in the world are preserved during manipulations. This is an important part of what it may mean to have a world model.

The Social World of Bats

Bats are very social animals: many species live in large colonies. Like other mammals, bats form complex social units and group structures. Given the overall acoustic specializations of many species of bats, it is not surprising that they depend to a large degree on an acoustic communication system for social interactions. The syntactic complexity of bat calls is comparable to that of other mammals, including primates (Kanwal, et. al. 1994). A strong similarity between the structure of bat social calls and the vowels and consonants of the human voice has also been noted (Kanwal, et. al. 1994; Suga, 1988). Because of its complexity and specialization, the study of the neuroacoustical system of the bat has been suggested an appropriate supplement to the study of the human auditory system for understanding the neural basis of speech perception and recognition (Suga, 1988).

Several important questions about the neural basis of bat social behavior remain to be answered. Given that bats are inherently social animals, how is this social element reflected in the bat's view or model of the world? How does the bat act as if it were a member of a community? What does it mean for a bat to be aware of conspecifics, and how does this awareness compare to that found in other mammals, primates, and humans? Social calls in invertebrates and lower vertebrates are often accompanied by changes in the behavioral state of the listener, including hormonal changes which can last seconds, hours, or days (Heiligenberg, 1991). But there is relatively little evidence for the relationship between neural and social behavior in most animals, particularly bats. Do the same parts of a bat's mind participate when the bat is engaged in a social problem (courtship or dominance, for example) versus a navigational one? Do the same part of a bat's brain participate? And if there are different answers to these questions, what does that difference say about our explanation of what a mind may be versus what a brain can do? Before we probe questions about the bat's brain however, we need to take a look at one more important part of the way the bat hears its world.

How the Bat Shapes Its Acoustic World

Most echolocating bats have evolved a remarkable set of physiological adaptations that significantly shape the way they hear the world. The most obvious external adaptations are the outer ears and pinna: they are usually enormous, with odd ridges or pinna convolutions. These physical structures shape the bat's perception of its acoustic environment. The pinna for example is highly directional, and its directionality usually varies as a function of frequency (Fuzessery, et. al. 1992). This means that the spectral profile of echoes will be systematically altered as function of their location in the sound field. This function may be learned by the auditory midbrain or cortex in order to derive a map or representation of the spatial direction of objects (Fuzessery, 1986). The point is that the bat's own body literally shapes the way it hears the world.

Most studies of ear directionality do not include the extreme lateral or rear sound field. There are several species of bats capable of hunting insects or small rodents purely from passive localization of the rustling noises of prey. Presumably, this would require the ability to locate sound sources from all directions. How this may be achieved is not quite clear, but may involve similar processing of the pinna-altered spectral profile of the sounds (Fuzessery, 1986). An understanding of how bats perceive sounds from non-frontal directions is an important part of any explanation of how bats hear their world as an entirety.

Other important physiological adaptations help the bat to shape its acoustic environment. The vocal tract, mouth, and in some species even the nostrils are capable of directing the sonar pulse into a cone or tight beam in the direction of the bat's gaze. The horseshoe bat, for example, earned its name from a peculiar horseshoe-shaped formation on its nostrils, which acts as a very efficient directional horn; this particular bat is capable of moving its ears and using its horn as a 'flashlight' to actively scan its environment (Griffin, 1958). The mustached bat is also capable of focusing its sonar pulse into a narrow cone, so that sound energy is relatively evenly distributed across the width of the cone, but falls off dramatically away from the cone. Such intensity stabilization in the center of the acoustic field may aid in the processing of fine characteristics of targets (Fuzessery, et. al. 1992).

The extreme directionality of the ears in combination with the tight focusing of the sonar pulse makes it clear that when bats are echolocating, they gain detailed information about the world immediately in front of them. The bat's ability to alter the shape and spread of its sonar pulse has one other consequence: it has the ability to compensate for the directionality of its ears (Fuzessery, et. al. 1992). The alterable sonar pulse spread and ear directionality supplement and complement one another. Echolocation is a perceptual system which give bats rich information about the world. Now we shall take a look at how bats use this rich information to form neural representations of the objects they encounter -- in the form of dedicated cortical maps.

Representations of the World in the Auditory Cortex of the Bat:

Echo Delay Representation: Object Distance

Multiple neural maps have been identified in the primary auditory cortex of the mustached bat. In area FM-FM, neurons respond to echo delays systematically, so that neurons on one end of the area respond to short echo delays, and neurons on the other end respond to long echo delays (Figure 8) (Suga, 1984; Suga, 1990; Sullivan, 1982; Suga, et. al. 1979). Since echo delay is directly related to the distance of the object, this area can be said to map target distance. The most likely neural mechanism which could construct such a representation is a set of delay lines and coincidence detectors (Simmons, 1973). If a signal indicating the outgoing sonar pulse is delayed, and then brought together with a signal indicating the returning echo, a coincidence detector can signal the length of the echo delay, and thus the distance to the target. Since the bat needs to calculate delays on the order of several milliseconds, it is most likely that these delay lines are polysynaptic; this is in contrast to the barn owl, which uses single axons and monosynapses to delay the timing signal (Suga, 1973).

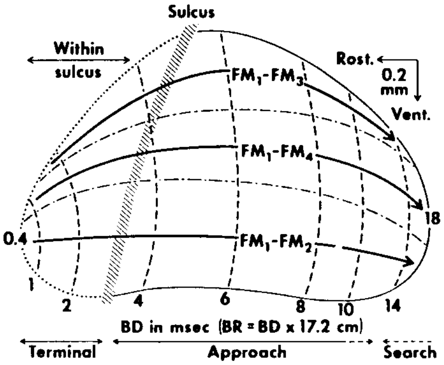

Figure 8: Echo Delay - Distance Map (Area FM-FM)
(Adapted from Suga, 1984)

Two classes of neurons have been identified in area FM-FM (O' Neill, et. al. 1979; Sullivan, 1982). The best echo delay of range-tuned neurons stays constant, while the best echo delay of tracking-neurons changes as the distance to a target becomes smaller. Both have been suggested to play a role in the formation of an auditory image (O' Neill, et. al. 1979; Dear, et. al. 1993). It is possible that groups of range-tuned neurons represent detailed information about the specific location of objects at specific distances; a group coding might account for the extremely fine acuity of bat ranging judgments (Heiligenberg, 1991). The behavior of tracking-neurons may be the result of a relatively high-level of processing which separates targets of interest from clutter in the environment.

Doppler Shift Representation: Object Velocity

In area CF-CF, neurons respond to the Doppler shifts of echoes systematically, by place-coding one sonar harmonic versus one echo harmonic (Figure 9) (Suga, 1984; Suga, 1990). Along one axis, neurons respond to increasing sonar pulse frequencies; along a roughly perpendicular axis, neurons respond to increasing returning echo frequencies. Along the diagonal axis, neurons respond to increasing frequency differences, or Doppler shifts. Since the Doppler shift of the returning echo is closely related to the velocity of the target with respect to the bat, this area can be said to map target velocity. Note that since the bat is capable of adjusting the frequencies of its sonar pulses, neurons need to be able to combine information from both auditory events to compute the true Doppler shift.

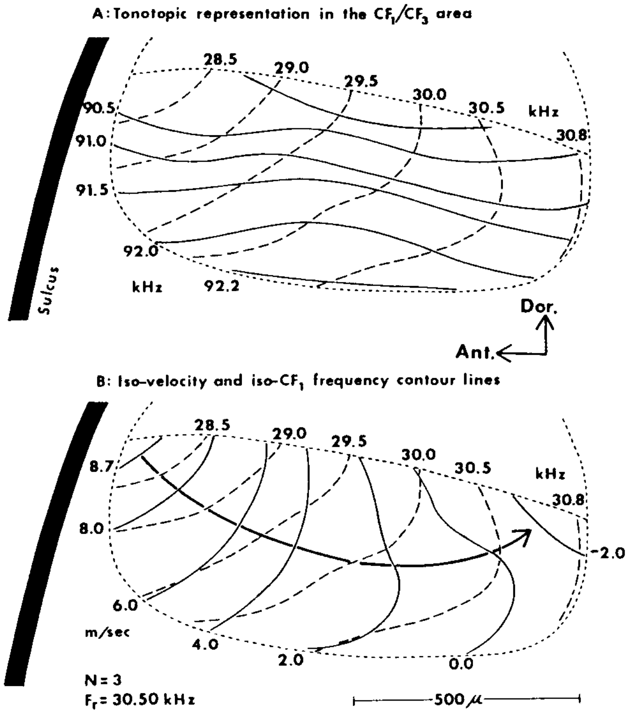

Figure 9: Doppler Shift - Velocity (Area CF-CF)
(Adapted from Suga, 1984)

A different mechanism of velocity determination has been discovered in some species of bats which do not use CF components in their sonar pulse. These FM-only bats use the pulse-to-pulse changes in echo delay to derive a measurement of target velocity (Tanaka, et. al. 1993). In these bats, neurons which are sensitive to specific target ranges may also be sensitive to specific target velocities. If a similar mechanism operates in CF-FM bats, the two velocity representations (Doppler shift, and change in echo delay) might complement one another in a way that gives CF-FM bats an adaptive advantage suitable for their particular environments (Tanaka, et. a. 1993), whereas the FM bat may not need two neural derivations of velocity to survive in its environmental niche.

Constant and Variable Doppler Shifts: Object Flutter

To a first approximation, the echo returning from a target is Doppler shifted by an amount proportional to the velocity difference between the bat and the target. But small motions due to flexibility of the object itself can also cause small Doppler shifts in the returning echo (Figure 10). Such object motions can be detected and processed in area DSCF. Although there is an overall tonotopic organization in DSCF, the area significantly overrepresents the harmonic frequencies of the CF portion of the bat's sonar pulse. Since neurons in DSCF are extremely well-tuned to particular frequencies, they have the capability of representing very small Doppler shifts (Suga, 1984). Suga reports that roughly half of DSCF neurons respond to sinusoidally varying Doppler shifts, via excitation which is phase-locked to the variable, AC component of the Doppler shift. The AC Doppler shifts ranged from 0.16% to 1.6% of the base sonar pulse frequency of 61KHz (Suga, 1984). Although the Doppler shifts of the wingbeats of flying insects are far more complicated than a sinusoidally modulated Doppler shift, it is clear that DSCF is capable of creating rich, detailed descriptions of the dynamic spectra of sonar targets. Neurons which phase lock to AC Doppler shifts have also been reported in area CF-CF, where roughly one third of the neurons seem to be sensitive to the AC component (Suga, 1984). Although the AC component of a Doppler shift does not appear to be place-coded in DSCF or CF-CF, Suga notes that other areas may respond systematically to particular wingbeat frequencies or amplitude spectra.

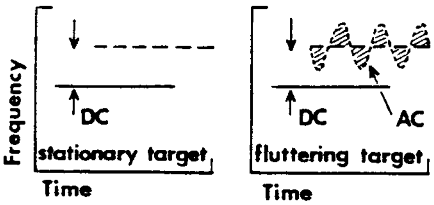

Figure 10: Constant and Variable Doppler Shifts
(Adapted from Suga, 1984)

Amplitopic Representation: Subtended Angle and Object Size

Although there is no systematic representation of amplitude at the auditory periphery, the cortex does contain a map of echo amplitude versus detailed frequency in area DSCF -- the area which contains the auditory fovea. Specifically, the best amplitude response of neurons in the DSCF auditory "fovea" varies circularly, with low amplitudes along one radius of the circle, and increasing amplitudes along more distant radii (Figure 11). Amplitopic representation of this kind is created from a combination of lateral inhibition and minimum thresholds for excitation (Suga and Manabe, 1982). Suga also suggests that lateral inhibition may provide amplitude and spectrum contrast enhancement (Suga, 1984).

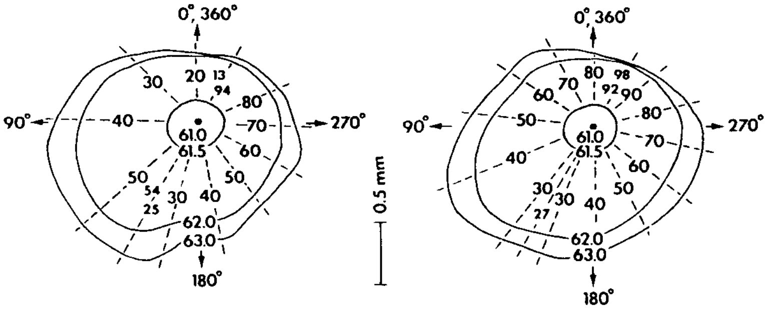

Figure 11: Amplitude - Subtended Angle (Area DSCF)
(Adapted from Suga, 1984)

One notable absence is a lack of a systematic representation of object size, rather than just the subtended angle of the object. Object size can be computed by multiplying subtended angle (amplitude) times distance (echo delay), and then combining this with some knowledge about how efficiently the object reflects sound. We might expect to find either connections between neural maps of amplitude and distance, or common projections to a third area which might represent object size. So far, no such projections or areas have been found. However, (Suga, 1984) have noted that most neurons in area FM-FM respond to not only specific echo delays/target distances, but also specific echo amplitudes. Thus these neurons may also represent specific object sizes at specific distances. Although variations in responses to echo amplitude spectra seem to occur perpendicular to the target-range axis of area FM-FM, no systematic representation of object size along this axis has been reported. It is not yet clear how knowledge of the object's sound reflectance efficiency is incorporated into this computation.

Binaural Interaction: Spatial Location

Neurons sensitive to interaural intensity differences of a particular harmonic component of the bat's sonar pulse have been found in the inferior colliculus. These neurons are likely to be part of the neural basis for the bat's keen ability to locate the spatial directions of sounds (Fuzessery, et. al. 1984). Other researchers have noted that some neurons in area DSCF may also be sensitive to spatial location. Most neurons in area DSCF respond to binaural stimuli, and can be classified into two categories. One group is excited by contralateral and ipsilateral stimuli (E-E neurons), and one is excited by contralateral and inhibited by ipsilateral stimuli (I-E neurons). (Suga, 1984) report that these neurons are segregated into two functional areas. E-E neurons tend to be located in areas where low echo amplitudes are also represented; these neurons also seem to be poorly directional, and thus integrate information about weak echoes from a wide area of space (which may enhance their detection). I-E neurons tend to be located in areas where high echo amplitudes are also represented; these neurons seem to be horizontally directional, and thus better suited to provide information about the spatial location of the target. Suga notes, however, that there does not appear to be any systematic representation of spatial location in DSCF.

Best Azimuth and Changing Azimuth: Spatial Location and Spatial Motion

Two areas of nonprimary auditory cortex have been identified as containing neurons with sensitivity to spatial location. In area DM, neurons respond relatively systematically to echoes which come from different horizontal directions (Figure 12) (Suga, 1984). The "best-azimuths" of these neurons vary radialy, so that directions in front of the bat are represented at the center of area DM, and increasingly contralateral directions are represented towards the edge of area DM. Suga reports that the "best elevations" of these neurons are not arranged systematically. He also cautions that the spatial selectivity of DM neurons is most apparent when measured by counting neural impulses, rather than by measuring the minimum threshold for response -- a kind of neural coding more often found in the auditory periphery than the central nervous system. In area VP, directions appear to be represented less systematically, but VP neurons have an additional interesting response property: these neurons respond better to the echoes of objects which are moving in the horizontal direction than they do to the echoes of stationary objects. No neurons sensitive to changes in elevation (vertical direction) or echo delay (depth) have been found in this area, however.

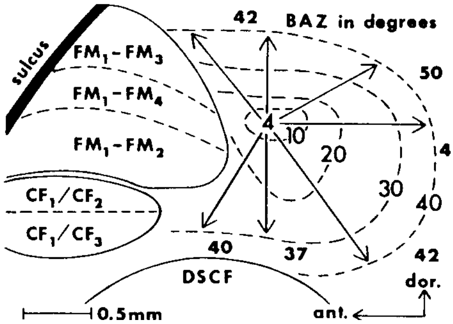

Figure 12: Horizontal Spatial Location (Area DM)
(Adapted from Suga, 1984)

Anatomical Evidence: Segregated Neural Pathways

Information is neurally segregated at a very early stage in the auditory system of echolocating bats: frequency, size, distance, and direction information are processed in separate pathways (Suga, 1990). The neural pathways share some common principles. They tend to compute and represent important information bearing parameters (Suga, 1992). The selectivity of individual neurons is improved as frequency information is processed through successive layers in each pathway. But each pathway is also specialized. The spatial direction pathway, for example, integrates information from both ears to create representations of interaural intensity differences, and from these, neural maps of spatial location. The range map pathway uses polysynaptic delay lines and coincidence detectors to compute the time difference between the emitted pulse and the returning echo. The velocity map pathway uses lateral inhibition to provide exceptionally fine-tuned neural selectivity for Doppler shift detection between the emitted pulse and the returning echo. The existence of multiple, segregated pathways in bats is consistent with studies of the sensory systems of many other animals; Heiligenberg terms this the "separation of task-specific sensory channels" (Heiligenberg, 1991). They appear to be likely neural correlates of Bregman's primitive stream segregation.

Depending on the pathway, FM, CF, or amplitude selectivity is added to some neurons; these selectivities are compounded in later stages, so that neurons are sensitive to combinations of important echo components (Suga, 1988). For example, some neurons in the distance map pathway respond only to the presence of a loud FM_1 component of a sonar pulse. Others respond only to the presence of the quieter FM_2, FM_3, or FM_4 component of an echo. Neurons in later stages of the pathway will respond only to the combination of a delayed FM_1 component of a sonar pulse and the FM_2 component of an echo. If they respond to these components only when they are separated by a certain, specific time delay, they then have the capability of representing the distance of a target. Similar single- and combination-sensitive neurons can be found in the velocity map pathway. This is one standard approach to understanding the function of neural mechanisms: to chart out the response of single neurons to successively more complex combinations of information bearing stimuli (Suga, 1992). The neural implementation of Marr's increasingly abstract representations may take the form of such increasingly complex combination-sensitive neuron response properties.

Reciprocal connections between the bat's frontal and auditory cortices have been identified. These connections may provide a neural basis for the bat's ability to remember its spatial environment(s). In combination with reciprocal connections between auditory cortex and the superior colliculus, they may also provide a neural basis for the bat's ability to coordinate its motor abilities in order to scan and navigate through its environment. (Pollak and Casseday, 1989).

Lesion Studies

Although neural maps of echo delay, Doppler shift, etc. have been discovered in the cortex, this does not prove that they are directly involved in the bat's behavior, or its world model. Lesion studies of these cortical maps provide the essential link demonstrating the necessity of these cortical maps for normal behavior (Riquimaroux, 1991). When cortical maps representing Doppler shifts (DSCF) or echo delays (FM-FM) were selectively deactivated by inhibitory neurotransmitter antagonist injections, bats trained on frequency shift or echo delay discrimination tasks were impaired. When area DSCF was deactivated, bats were unable to make fine frequency shift judgments, but were fully able to make echo delay judgments; conversely, inactivation of area FM-FM disrupted judgments of echo delay, but not frequency shift judgments. Whether or not these cortical areas are necessarily a part of the bat's world model remains to be proven, however. Some bats with bilateral cortical ablations are still capable of wire obstacle avoidance tasks (Sullivan, 1982). Further research is necessary to specify exactly which parts of the brain are necessary for high-level navigation skills and the bat's model of the world.

Neural Representations During an Echolocation Pulse

We can begin to construct a story of what happens in the bat's auditory cortex during the hunt for a moth (Figure 13). When a bat emits an echolocation pulse, it protects its sensitive hearing first by blocking part of the middle ear during the pulse, and then immediately releasing the blockage at the end of the pulse. In combination with a neural mechanism of attenuation to self-stimulation, the response can be reduced by as much as 35 to 40 dB (Suga, et. al. 1974). As echoes from nearby objects return, they are matched with delayed versions of the outgoing pulse, so that the amount of delay can be measured from the FM-FM sweeps. This measurement is dynamically place-coded in cortical area FM-FM; at first, this coding is rather approximate. As echoes from objects that are farther away arrive, their distance measurements are also registered, and the representation of nearby object distances are refined (Dear, et. al. 1993). By the end of the returning echo, the accuracy of the distance representation has stabilized. The frequency of the fundamental harmonic of the outgoing pulse is also matched with the dominant harmonic of the returning echo to form a place-coded map of velocity in area CF-CF.

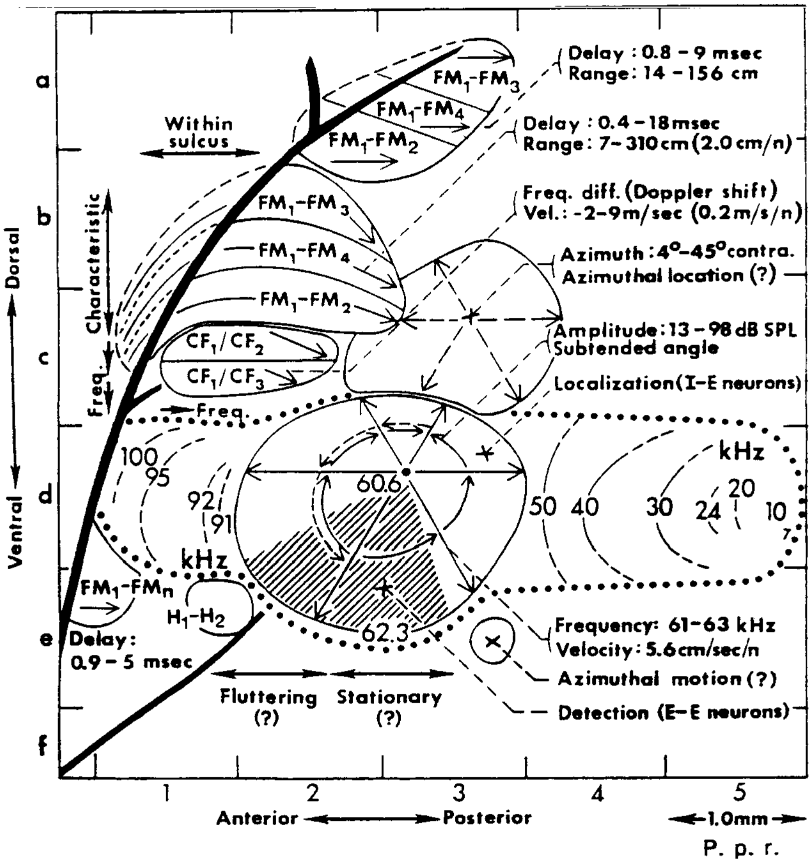

Figure 13: The Bat Auditory Cortex
(Adapted from Suga, 1984)

The spectral pattern of the echo is relayed to area DSCF, which contains disproportionate representations of the harmonic frequencies of the sonar pulse. Area DSCF is likely to perform several functions (Suga, 1984). It has the capability of detecting small Doppler shifts in the echo; this may provide the bat with the information it needs to successfully perform Doppler shift compensation on the next sonar pulse, to keep the dominant harmonics within the overrepresented areas in DSCF. It also has the capability of detecting Doppler shift fluctuations, which may signal internal object motions such as the wingbeats of insects. Since the neurons of area DSCF are extremely well-tuned to narrow frequency ranges, even slightly different echo spectral patterns may cause significantly different patterns of activation in area DSCF. This would be necessary not only to differentiate the object's texture or structure, but also to determine its spatial location.

Bats accumulate information about the world around them in rapid, successive bursts, like a series of flashes or snapshots; this abrupt, discrete method of obtaining information is likely to have significant effects on the way bats construct neural representations. For example, representations generated from information accumulated from one sonar pulse are likely to influence the representations generated in the next. If a bat detects the presence of an object 4.3 meters away while it is flying, it is likely that its neural mechanisms will be facilitated to detect the presence of an object at roughly 4.0-4.5 meters on the next sonar pulse. For some kinds of feature computations (such as velocity computation from change in target distance), this pulse-to-pulse transfer of information is absolutely necessary. Exactly how large a role these pulse-by-pulse mechanisms play is not clear, but they are likely to be an important part of the way the bat shapes its conception of the world, and are an interesting neural correlate to Bregman's proposal of sequential stream integration.

The mechanisms underlying neural representations are also likely to impact the ability of the bat to construct a notion of its world. There are comparative values to place coding a feature (a clearly explicit representation) versus representing a feature as a dynamic pattern of activation -- a less explicit representation that might need some "interpretation" in order to be used by other brain areas. Place coding allows a certain degree of hyperacuity -- the ability to sense differences or changes in a physical variable (such as distance) that are smaller than the minimum range of individual sensory receptors (Heiligenberg, 1991). Evolution has selected for some explicit, place-coded neural representations, but not others. The issue becomes, to what degree does the measurement of some environmental feature need to be made neurally explicit?

What is on a Bat's Mind

The bat's echolocation system provides a rich set of acoustic cues about the environment and enables the construction of rich representations of the world within the bat's brain. When putting these pieces together, we are faced with several questions. How are the states of the world and the states of the bat's brain related during its life in the wild? For example, what does the state of the bat's cortex look like as it is hunting a moth or attempting to find a place to roost for the night? We might think we are beginning to understand how the echolocation system allows the world to impose structure on the bat's brain; but given some structure in the bat's brain, what can we infer about the state of the world? Can we say that this structure determines "what is on a bat's mind"?

One way to discover what is on a bat's mind is to ask the question: What is interesting to a bat? We have partial answers to this question. The bat alters its sonar pulse when it is focused on a target, such as a fluttering moth; it compensates for the moth's motions in the air, for the moth's distance, and it makes fine discriminations based on the moth's behavior and spectral "texture". The bat is aware of the spatial location of obstacles, and makes efforts to avoid them. The bat is aware of and actively seeks spaces or holes where it can take shelter. The bat is aware of conspecifics, and can communicate with them. These are all things that bats are interested in -- they are what is on the bat's mind. Our task is to understand how representations of them may also come to exist as structures in the bat's brain.

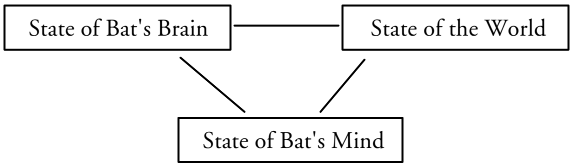

What is 'Nothing' to a Bat?

Bats are hunters and must chase after insects in order to catch them; when flying, they must also detect obstacles in order to avoid colliding with them. Bats must have the ability to identify the spatial location of objects in the environment. But bats do more than avoid object collisions when flying; they actually find safe passages -- empty spaces -- through which they can fly, and holes where they can roost and shelter. Thus, bats must also have the ability to identify the spatial location of holes in the environment. We have the beginnings of an understanding of how the bat might represent objects in a world model; the next question is, how does the bat represent spaces in the world model? What is the bat's concept of 'nothing' like? Framing the problem in Nagel's terms, is 'nothing' a qualia for a bat?

Given that we know that bats represent some aspects of objects in their brains explicitly, there are at least two answers to these questions. One possibility is that bats represent empty spaces explicitly, through place codes -- neural activity which labels a section of the world as being 'empty'. Another possibility is that bats represent empty spaces implicitly, by labeling all the objects in the environment that they can detect, and then computing the spaces via a mental subtraction operation. The right question may be, do bats ever need to make the locations of spaces explicit? Actual neural representations so far appear to be distributed and object explicit -- neural activity tends to represent the existence of information about an object. If information about the spaces in the environment is used by other parts of the brain, the information needs to be computed, extracted, or transformed. It is still possible that other regions of the brain may make the location of spaces explicit.

Do Bats Have a Notion of Shape?

Having a notion of shape requires the ability to conceive of spatial relationships between the parts of an object (Marr, 1982). The visual system does this automatically, by providing a representation of these spatial relationships from the very beginning, on the retina. The cochlea, however, cannot provide direct spatial information in this way. At best, spatial location information must be extracted from the spectrum of a sound and directionality cues from the pinnas. Although we can see how an object's location might be computed, it is not clear how object shapes can be systematically determined using this method. The problem in determining shape from sound is that the locations of multiple, nearby object points must be determined from their respective echoes, these echoes will tend to severely overlap one another. The situation is even more complicated if there are echoes from more than one object.

Bats are capable of making discriminations between objects with different shapes, even if the objects are of identical sizes and reflect the same amount of sound energy (Schnitzler, et. al. 1979). One simple neural network model of this ability has been proposed (Dror, et. al. 1995). A second glance at Griffin's original wire obstacle avoidance experiments demonstrates that bats are not only able to sense a difference between obstacles of different shapes, but they make use of this difference (Griffin, 1958). Bats were able to navigate through rooms with both vertical and horizontal wire obstacles. The difference is important, because bats must be able to tell if there is enough horizontal clearance for their wings in order to be able to avoid collisions. They must be able to tell if the shape of the space between the wires is larger or smaller than their wingspan. This study suggests that bats do have a notion of shape, although exactly how detailed or powerful this notion may be is unclear.

The issue needs further study on three fronts. First, we must better understand how object shapes might be computable from a spectral analysis of echoes. Second, we need tighter limits on how powerful the bat's ability to sense object shapes might be. And third, we need to understand the neural basis for this ability, in particular how information within the bat's brain is manipulated to form representations of spatial relationships.

How do We Go From Features to Objects

One problem that remains unanswered by the evidence presented so far involves the construction of object representations from representations of their features. One way to look at this problem is to ask: How are multiple representations linked together to form a coherent representation of an object? Another way of posing the problem is to ask: is it really necessary for representations of object features to be directly connected with one another for the animal to act as if it can represent objects and the world?

In vision, we might imagine that the problem is relatively simple, as most early representations are distributed and retinotopic, and thus connections between them may also be retinotopic. Object features can be associated with one another easily because they occupy the same locations in the retinotopic map; multiple objects can be represented at different locations on the map. Thus a relatively simple retinotopic projection can do most of the work of keeping the representations of an object's features associated with one another. In hearing and in echolocation in particular, the neural maps are distributed; each map can represent information from multiple objects. But these maps are usually not spatiotopic. How are the maps linked together so that information which pertains to the same object is associated together, but kept distinguishable from the information for other objects? In auditory maps, tonotopy often seems to be the homologue to spatiotopy; perhaps the maps are linked together tonotopically. It is then an open question if this would be sufficient to associate an object's features together, but keep multiple objects distinguishable. The overlapping spectrum problem suggests that this may not be enough; rather, the dynamic spectral pattern may need to be further parsed. This may be accomplished by a mechanism akin to Bregman's schema-based stream analysis.

A bat is not born with a model of its environment, nor does it develop one fixed model of the world it lives in. Rather, it is born with the capacity to construct a model of the environment out of the sensory data it gathers. This model, embodied in the bat's brain and its multiple representations and states, presumably helps the bat to hunt, navigate, and survive. How flexible is this model-building capacity, and what role do representations of features play in final product? Presumably, the individual feature representations are ontogenetically fixed; all bats of a particular species will have more or less the same set of subcortical and cortical maps. How flexible is the bat at putting these maps together? How do the limitations of the combinatorics of the underlying feature representations constrain the way the bat perceives and constructs a model of its world? How do we reconcile increasingly high-level descriptions of the bat's brain with increasingly detailed descriptions of the bat's mind (Figure 14)?

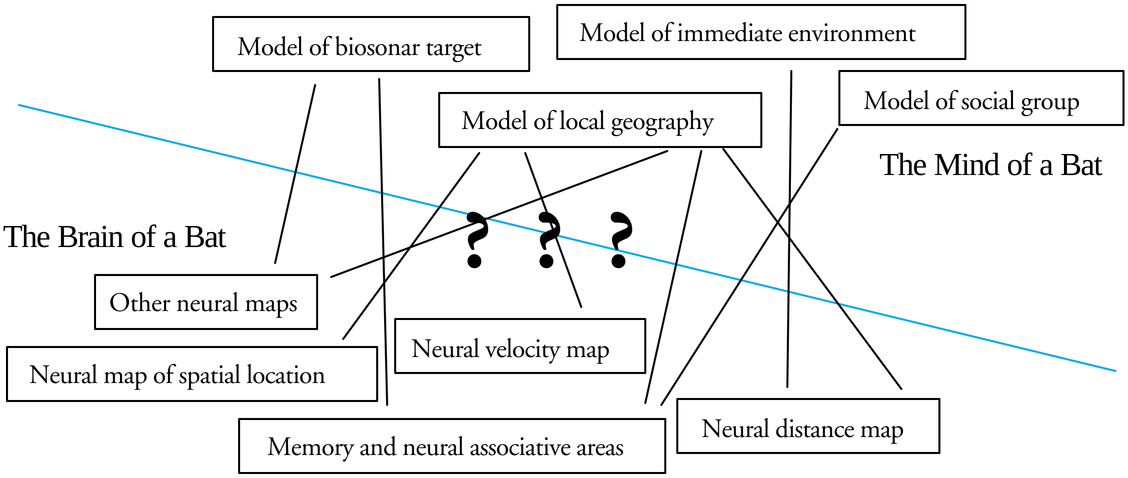

Figure 14: State of the Bat's Brain versus State of the Bat's Mind

Conclusion: The Bat Has a Mind of Its Own, But We Can Know What's on Its Mind

Because of the evidence described above and other important studies, we are now in a position to say something about what the bat's view of the world may be like: what it is like to be a bat. Part of this story is inspired by a critique of Nagel's claims (Akins, 1993). In her critique, Akins describes a fictional 'film' of what it is like to be a bat: in this film, the various sensory impressions of the bat have been translated into analogous visual stimuli so that humans can see them. Intensity of light corresponds to intensity of sound; the colors of objects correspond to the spectral properties of echoes; echo delay and distance information is converted into stereoscopic cues to mimic depth, and so on. The film is played in a special 360° movie theater, so that we have the chance to be exposed to all of the sensory information that the bat gathers: so that we can "see" everything the bat hears. Knowing some of what is important to the bat, let's push the analogy even farther, and see if it helps us understand the bat's point of view.

Although Akins uses this fictional film as a straw man to undermine Nagel's claim, the film is an excellent starting point for an understanding of the auditory images the bat is perceiving, simply because the analogies are so striking. For example, a typical movie theater film is not a continuous event; rather, it is a long sequence of frames, which when played at a fast enough frame rate (roughly 30 Hz) appear to blend together into a continuous stream of action. Bat echolocation is strikingly analogous: the sonar pulses are discrete events which occur at comparable rates (10 Hz to 100Hz). The information contained in each successive echo is neurally merged to compute representations of distance and velocity which persist from pulse to pulse. Given that we know there are auditory illusions analogous to apparent motion, it becomes quite easy to see how bats might develop a continuous perception of their environment from the discrete information available to them.

Another analogy that is often made between the visual system and bat echolocation is between the fovea of the retina and the tonotopic auditory "fovea", the overrepresented range of frequencies in the bat's auditory system. Saccades in the visual system bring objects of interest into the densest region of retinal photoreceptors; bats which perform Doppler shift compensation adjust the frequencies of their sonar pulses to compensate for the Doppler shifts of targets, and thus keep target echoes within the densely represented region of the cochlea. But there is another candidate for the foveal analogy. Bat sonar pulses are highly directional; when they probe their environment with a pulse, most of the sound energy is directed at whatever is in front of them. Due to physiological adaptations in the pinna, bat hearing is also highly directional; the ears are exquisitely sensitive to echoes immediately in front of them. In combination, the physiological adaptations give the bat another, spatial auditory fovea, which can then be moved around to point at objects of interest in the environment. And indeed, many species of bats which focus their sound into a narrow beam scan their environment by moving their heads and ears back and forth, much like the saccade of an eye.

By studying the information bats gather from sonar and understanding some of the neural representations they can compute, we can extend Akin's fictional film into a "virtual reality" experience and assign more definite meanings to the fanciful display of color and light. The individual frames of the movie can be interpreted by us as individual sonar pulses, and the feel of continuity is similarly analogous. The movie camera's lens might seem to focus on one object, and this focus could be interpreted as the bat's Doppler shift compensation for this object. Bat head and ear motions, which change the directionality of the whole echolocation system, could be interpreted as saccades of our eyes.

The directionality of the sonar system makes the movie seem dark but flashy, as objects in the center of view are brightly illuminated and in focus, while objects in the background are dark and out of focus. As the bat scans its environment, the sensation is much like a person searching with a flashlight on a dark night. The flashlight seems to have a certain tint, which is the unique spectral shape of the bat's voice. Sometimes we may see a bright, glaring flash with a slightly different tint; this is the sonar pulse of another bat flying nearby. We may also see sequences of lively, rapidly changing colors from a particular spot nearby; these sequences sometimes repeat, but are complex, and are the communication calls of another bat.

Different areas of the visual field seem to have different tints, as the bat uses spectral cues to derive information about spatial directions. When a moth happens to appear in the beam, its wings glint and flash; it also has a particular "color" and visual texture, which is related to the spectral shape of the moth's echoes. As the bat approaches the moth, the intensity of the flashlight beam seems to wane somewhat, but the moth's illumination and focus remain constant; the bat is reducing the amplitude of its sonar pulses to keep the amplitude of the returning echoes constant, and compensating for the Doppler shifts of the moth's echoes.

When the bat is navigating through a cluttered environment, such as a forest or around buildings, we see somewhat vague and unfocused blobs in different parts of the visual field, with tints and hues comparable to the spectral properties of the echoes of the surfaces of leaves, walls, and other obstacles. In summary, this extended visual analogy is an excellent way of trying to understand the bat's auditory perspective; the key is that we need to know the acoustic correlates of each of the visual features we see, so that we can begin to synthesize a "hybrid" interpretation of these translated sensory impressions. We also need to know what role these acoustic correlates play in the bat's daily life. And the only way to find that out is by empirical investigation.

The naive answer to Nagel's question is the summary and story of what it is like to be a bat: the evidence presented in this paper and in the scientific literature in general. The serious answer to

Nagel's question is that science and empirical study <u>can</u> tell us what it is like to be something or someone other than we are; we just have to make the effort to go and gather data, perform tests, and ask questions. There is no utility to the objection that there is some ineffable quality which keeps us from knowing the last essential quality of a bat's point of view; through empirical investigation, we can come arbitrarily close to understanding this point of view, and perhaps even the functioning of another's mind, undaunted by problems of qualia. The glass is not half-empty; it is really half-full.

References

Akins, K. (1993). A Bat Without Qualities? *Consciousness*. Oxford: Blackwell Ltd.

Altes, Richard A. (1989). An interpretation of cortical maps in echolocating bats. *The Journal of the Acoustical Society of America*. 85(2): 934-942.

Bregman, A. (1990). *Auditory Scene Analysis*. Cambridge: MIT Press.

Chelazzi, L., Miller, E. K., Duncan, J. Desimone, R. (1993) A neural basis for visual search in inferior temporal cortex. *Nature*. 363: 345-347.

Dear, S., Simmons, J., Fritz, J. (1993) A possible neuronal basis for representation of acoustic scenes in auditory cortex of the big brown bat. *Nature*. 364: 620-623.

Dror, I., Zagaeski, M., Moss, C. (1995). Three-Dimensional Target Recognition via Sonar: A Neural Network Model. *Neural Networks*. 8(1): 149-160.

Evans, E. (1992). Auditory processing of complex sounds: an overview. *Philosophical Transactions of the Royal Society of London*. 336: 295-306.

Fei Jin, Z., Srinivasan, M. (1990). Neural gradient models for the measurement of image velocity. *Visual Neuroscience*. 5: 261-271.

Feng, A., Condon, C., White, K. (1994). Stroboscopic hearing as a mechanism for prey discrimination in frequency-modulated bats? *The Journal of the Acoustical Society of America*. 95(5): 2736-2744.

Fuzessery, Z. (1986). Speculations on the Role of Frequency in Sound Localization. *Brain, Behavior, and Evolution*. 28(1-3): 95-108.

Fuzessery, Z., Pollak, G. (1984). Neural Mechanisms of Sound Localization in an Echolocating Bat. *Science*, 255: 725-728.

Fuzessery, Z., Buttenhoff, P., Andrews, B., Kennedy, J. (1993) Passive sound localization of prey by the pallid bat (*Antrozous p. pallidus*). *The Journal of Comparative Physiology*. 171(6): 767-777.

Fuzessery, Z., Hartley, D, Wenstrup, J. (1992) Spatial processing within the mustache bat echolocation system: possible mechanisms for optimization. *The Journal of Comparative Physiology*. 170(1): 57-71.

Gibson, J. (1966) *The Senses Considered as Perceptual Systems*. Houghton Mifflin Press. Boston, MA.

Griffin, D. (1958). *Listening In The Dark: The Acoustic Orientation of Bats and Men*. New Haven: Yale University Press.

Grinnell, A. (1973). Neural processing mechanisms in echolocating bats, correlated with differences in emitted sounds. *The Journal of the Acoustical Society of America*. 54(1): 147-156.

Hartley, D. (1990). Phased-array beam scanning as a possible aid to horizontal localization in horseshoe bats. *The Journal of the Acoustical Society of America*. 88(6): 2889-2891.

Heiligenberg, W. (1991). The Neural Basis of Behavior: A Neuroethological View. *Annual Review of Neuroscience*. 14: 247-267.

Kanwal, J. Matsumura, S., Ohlemiller, K., Suga, N. (1994). Analysis of acoustic elements and syntax in communication sounds emitted by mustached bats. *The Journal of the Acoustical Society of America*. 96(3): 1229-1254.

Konishi, M. (1991) Deciphering the Brain's Codes. *Neural Computation*. 3: 1-18.

Lin, Z., Chittajallu, S., Kayalar, S., Wong, D., Yurtseven, H. (1991). Modeling constant best delay-sensitive neurons and tracking neurons in the auditory cortex of the FM bat with a back-propagation neural network. *IEEE Conference on Neural Networks for Ocean Engineering*. 123-132.

Marr, D. (1982) *Vision*. New York: W. H. Freeman and Company.

Masters, W., Jacobs, S. (1991). The structure of echolocation sounds used by the big brown bat *Eptesicus fuscus*: Some consequences for echo processing. *The Journal of the Acoustical Society of America*. 89(3): 1402-1413.

Meijer, P. (1992). An Experimental System for Auditory Image Representations. *IEEE Transacations on Biomedical Engineering*. 39(2): 112-121.

Metzner, W. (1989). A possible neuronal basis for Doppler-shift compensation in echo-locating horseshoe bats. *Nature*. 341: 529-532.

Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*.

Olsen, J., Suga, N. (1991). Combination-Sensitive Neurons in the Medial Geniculate Body of the Mustached Bat: Encoding of Target Range Information. *Journal of Neurophysiology*. 65(6): 1275-1296.

O'Neill, W., Suga, N. (1978). Target Range-Sensitive Neurons in the Auditory Cortex of the Mustache Bat. *Science*. 203: 69-73.

O'Neill, W., Suga, N. (1979) Neural Axis Representing Target Range in the Auditory Cortex of the Mustache Bat. *Science*. 206: 351-353.

Pollak, G. D., Casseday, J. H. (1989). *The Neural Basis of Echolocation in Bats*. Berlin: Springer-Verlag.

Riquimaroux, H., Gaioni, S., Suga, N. (1991). Cortical Computational Maps Control Auditory Perception. *Science*. 251: 565-568.

Saillant, P., Simmons, J., Dear, S. (1993). A computational model of echo processing and acoustic imaging in frequency-modulated echolocating bats: The spectrogram correlation and transformation receiver. *The Journal of the Acoustical Society of America*. 94(5): 2691-2712.

Schnitzler, H., Henson, O. (1979). Performance of airborne animal sonar systems. In: Busnel, R-G., Fish, J. (eds). Animal Sonar Systems. New York: Plenum Press.

Schnitzler, H., Ostwald, J., (1983). Adaptations For The Detection of Fluttering Insects By Echolocation In Horseshoe Bats. In: Ewert, J. P., Capranica, R. R., Ingle, D. J. (eds). *Advances in vertebrate neuroethology*. New York: Plenum Press. 801.

Simmons, A. (1973). The resolution of target range by echolocating bats. *The Journal of the Acoustical Society of America*. 54(1): 157-173.

Simmons, J., Kick, S., Lawrence, B. (1983) Localization With Biosonar Signals In Bats. In: Ewert, J. P., Capranica, R. R., Ingle, D. J. (eds). *Advances in vertebrate neuroethology*. New York: Plenum Press. 801.

Sterbing, S., Schmidt, U., Rübsamen, R. (1994). The postnatal development of frequency-place code and tuning characteristics in the auditory midbrain of the phyllostomid bat, Carollia perspicillata. Hearing Research. 76: 133-146.

Suga, N. (1984). The Extent to Which Biosonar Information is Represented in the Bat Auditory Cortex. In: Edelman, G. M., Gall, W. E., Cowan, W. M. (eds). *Dyamic aspects of neocortical function*. New York: Wiley and Sons. 315-374.

Suga, N. (1988). Auditory Neuroethology and Speech Processing: Complex-Sound Processing by Combination-Sensitive Neurons. *Auditory Function: Neurobiological Bases of Hearing*. New York: John Wiley & Sons. 679-720.

Suga, N. (1990). Cortical Computational Maps for Auditory Imaging. *Neural Networks*. 3: 3-21.

Suga, N. (1992). Philosophy and stimulus design for neuroethology of complex-sound processing. *Philosophical Transactions of the Royal Society of London*. 336: 423-428.

Suga, N., and Schlegel, P. (1973). Coding and processing in the auditory systems of FM-signal -producing bats. *The Journal of the Acoustical Society of America*. 54(1): 174-190.

Suga, N., and Shimozawa, T. (1974). Site of Neural Attenuation of Responses to Self-Volcalized Sounds in Echolocating Bats. *Science*. 183: 1211-1213.

Sullivan, W. E. (1982) Neural Representation of Target Distance in Auditory Cortex of the Echolocating Bat *Myotis lucifugus*. *Journal of Neurophysiology*. 48(4): 1011-1031.

Sullivan, W. E. (1982) Possible Neural Mechanisms of Target Distance Coding in Auditory System of the Echolocating Bat *Myotis lucifugus*. 48(8): 1033-1047.

Sullivan, W. E. (1986). Processing of Acoustic Temporal Patterns in Barn Owls and Echolocating Bats: Similar Mechanisms for the Generation of Neural Place Representations of Auditory Space. *Brain, Behavior, and Evolution*. 28: 109-121.

Tanaka, H., Wong, D. (1993) The influence of temporal pattern of stimulation on delay tuning of neurons in the auditory cortex of the FM bat, *Myotis lucifugus*. *Hearing Research*. 66: 58-66.

Yost, W. (1991). Auditory image perception and analysis: The basis for hearing. *Hearing Research*. 56: 8-18.

Zhang, S., Xinde, S., Jen, P. (1987). Anatomical study of neural projections to the superior colliculus of the big brown bat, *Eptesicus fuscus*. *Brain Research*. 416: 375-380.