USING ACOUSTIC TOMOGRAPHY TO TRACK ECHOLOCATING BATS

D A Waters IICB, Faculty of Biological Science, University of Leeds, Leeds, UK
IJ Farr IICB Faculty of Biological Science, University of Leeds, Leeds, UK

1 INTRODUCTION]

Acoustic tomography is the process by which structures are revealed by the scattering of sound projected into a space which is produced and received by one or more transducers. The term is most often applied to the process by which the temperature and density of oceanic water is deduced through changes in the velocity of sound as revealed by difference in travel time between a transmit and receive transducer array[1]. Acoustic tomography can also be used in industry to reconstruct particle distribution or flow properties in reaction vessels or pipes, termed process tomography[2]. Both these scenarios require the active production of a known signal whose scattering within the medium can be used to infer structure. In a biological context, acoustic tomography can be used to locate animals which themselves produce an acoustic signal. The principle is that the different times of arrivals of a signal at multiple receivers (microphones, hydrophones or seismic detectors) can be used to reconstruct the location of the source. This application is complicated by the fact that the time of emission of the signal is unknown and relies on four conditions being met:

1. That the receiver locations are accurately known.
2. That the time of arrival of signal can be accurately extracted.
3. That the speed of sound does not vary from the expected value between the source and any receiver.
4. That the number and arrangement of the receivers can produce an unambiguous source location.

Satisfying **Condition 1** is usually under the experimenter's control, and with modern surveying techniques and differential GPS giving accurate locations in three dimensions to within a few centimetres, this condition can usually be met, at least in terrestrial environments. The case is much more difficult in aquatic environments where the exact location of the receiver array may not be known, and if the elements in the array are not fixed relative to one another, the relative distances between them may also vary. Movement of the array due to wave action between the signal emission and reception would also result in positional inaccuracy.
Condition 2 can be problematic where signal to noise ratios are low. An error in extracting the time of arrival at any receiver will translate into an error in source position. Usually, a cross-correlation technique is applied to the incoming signal and the peaks of the cross-correlation used as the time marker for the arrival of the signal. This technique implies that the experimenter has an appropriate signal to use as the template for the cross-correlation, which may not be true if the animal's signal is not very stereotypic, and can be computationally very expensive if the signal sequence is long. Structurally complex environments where signals may be reflected from intervening objects producing multiple path lengths and hence temporally overlapping signals at a receiver can also produce errors in the time of arrival.
Condition 3 again mainly affects aquatic systems where the velocity of sound varies with temperature, salinity and depth, although in terrestrial systems, wind direction and temperature changes such as thermals and temperature inversions can also result in velocity changes. These are especially problematic where the velocity changes depend on direction between the source and receiver giving different errors at different receivers. Additionally, if the overall actual velocity

of sound at the time of the recording varies from the velocity used in the calculations, a general error in location will result. Using the approximation that:

$$c = 331.3 + 0.606T \qquad (1)$$

Where:
c = speed of sound in ms^{-1}
T = temperature in °C

At 0 °C, c = 331.3 ms^{-1} and at 20 °C, c = 343.4 ms^{-1}. Using a five element receiver array with four receivers in an square configuration spaced 10 m apart and a central fifth receiver, a source at 50 m would result in an location error of up to 2.7 m if 330 ms^{-1} were used as the speed of sound rather then 344 ms^{-1}. Errors in location decrease as the distance from the source to the receivers decreases (algorithm developed by Spiesberger[3]). While meeting condition 3 greatly simplifies the derivation of source location, methods do exist where the speed of sound is allowed to vary between microphones[4].

Condition 4 presents a considerable challenge in developing an algorithm which can be used to process the often considerable amount of data derived from field recordings. The major limitation is that the time of emission of the animal signal is unknown. If it was known, then the reconstruction of the source location becomes trivial and requires only one more receiver than the number of dimensions to localise the animal within (i.e. four receivers to localise the animal within three dimensions. If the animal is physically limited to two dimensions, then only three receivers are required). Since the time of emission is unknown this becomes a classical inverse problem where the source location needs to be established from the time of arrival of the signal at a number of receivers. With one receiver, the location of the source is totally ambiguous. Some indication of proximity may be obtained by the amplitude of the signal, and a directional receiver may give some indication of direction, but that is all. Increasing the number of receivers constrains the source location to more and more discrete regions.

Consider two receivers r(1) and r(2) at locations $r_x(i)$, $r_y(i)$ and $r_z(i)$ and detecting the time of arrival of a signal at T(1) and T(2) respectively (where *i* is the receiver number adopting the notation of Spiesberger[5]). If the signal *s* is at source location s_x, s_y, s_z and the signal emitted at time T_e, all possible solutions for s_x, s_y, s_z giving T(1) and T(2) lie on the surface of a hyperboloid.

In the simplest possible case where T(1) = T(2), then *s* is constrained to a plane normal to the axis between r(1) and r(2) and bisecting it at the mid-point. This is the same scenario as in binaural hearing, where the sound location received by two ears can lie on the surface of the 'cone of confusion'[6], or more appropriately a 'hyperboloid of confusion'. Adding a third receiver r(3) results in two intersecting hyperboloids for a given T(1), T(2) and T(3), and all possible solutions lie on the line of intersection which forms a curve which may open or closed depending on source location, the arrangement of the receivers and the consequential orientation of the hyperboloids. The addition of a further receiver r(4) describes a third hyperboloid which intersects the curve at either one point (if the curve is open) or two points (if it is closed). The solution from using four receivers may therefore be unambiguous in resolving to a single point, or ambiguous if there are two solutions. An ambiguous solution may be enough to allow source location, if for example if the animal is located on a surface and one of the two possible solutions places it underground, or if one of the possible solutions places the animal much further than the likely distance that sound can propagate over. However, to fully resolve source location, a fifth receiver r(5) is needed[7]. For every N receivers, there are N-1 hyperboloids since the difference in time of arrival of the sound between r(1) and r(2), given as τ_{21} which defines the 'focus' of the hyperboloid on the axis between r(1) and r(2), and τ_{31} defining the hyperboloid for r(1) and r(3) can give rise to τ_{23} without any extra measurement. It therefore adds no new information.

An equation for the solution of *s* using five receivers is given as[7]:

$$\mathbf{m} = \mathbf{A}^{-1}\mathbf{d}, \qquad (2)$$

Where:

$\mathbf{m} = \begin{pmatrix} s_x \\ s_y \\ s_z \\ t \end{pmatrix}$ which is the location of the source in the x, y and z dimensions, and t of emission.

$\mathbf{A} = 2\begin{pmatrix} r_x(2) & r_y(2) & r_z(2) & c^2\tau_{21} \\ r_x(3) & r_y(3) & r_z(3) & c^2\tau_{31} \\ r_x(4) & r_y(4) & r_z(4) & c^2\tau_{41} \\ r_x(5) & r_y(5) & r_z(5) & c^2\tau_{51} \end{pmatrix}$ and $\mathbf{d} = \begin{pmatrix} \|\mathbf{r}_2\|^2 - c^2\tau_{21} \\ \|\mathbf{r}_3\|^2 - c^2\tau_{31} \\ \|\mathbf{r}_4\|^2 - c^2\tau_{41} \\ \|\mathbf{r}_5\|^2 - c^2\tau_{251} \end{pmatrix}$

where $\|\mathbf{r}_2\|$ is the vector length between r(2) and r(1), and all other symbols are defined above. This solution can be easily implemented. Solutions involving higher numbers of receivers create problems since **A** becomes an *m x n* matrix and so the inverse does not exist. Instead the pseudo-inverse must be derived via singular value decomposition complicating the procedure, but ultimately yielding a solution which minimises the sum of squares of the residuals of source location[5].

Since the time of arrival of the signal is subject to various errors from propagation delays, errors in receiver location and errors in estimating the true time of arrival, the actual intersection of all possible hyperboloids may in fact describe a space within which the animal will be located. In this case, the addition of other receivers will increase the accuracy of localisation but with the added complexity of calculation.

A different approach to the same problem is used for echolocating bats[8], and also modified for use on cetaceans[9], in which four receiving transducers are mounted on a three spoked arm, with each arm at 120° to the next with one transducer at the end, and one in the centre. In this case, range to the animal is calculated from the differential time of arrival between the signal being received at the arm transducers and the centre, defined by:

$$R = \frac{3A^2 - c^2(\tau_{01}^2 + \tau_{02}^2 + \tau_{03}^2)}{2c(\tau_{01} + \tau_{02} + \tau_{03})} \qquad (3)$$

Where:
R = range to the source from the centre transducer.
A = the distance between the arm transducers and the centre transducer.
τ_{01}= the difference in arrival time of the signal at the centre transducer and transducer 1 on the arm of the array.
Once R is established, it is then possible to determine the location of the source in the x and y planes, as well as the time of signal production (see the Appendix in reference 9 for further details). This is a simplified version of the method presented by Aubauer[8] but gives very similar results and has been used experimentally to establish the location of bats flying down a flight tunnel[10].

One disadvantage of a system using a fixed geometry array is that range derivations are susceptible to time travel noise as the distance between individual receivers is small. Errors in the recording of time of arrival of small values such as 10μs at one transducer with A = 1m and a target range of 20 m can produce range errors in *R* of up to 0.87 m if the source is on-axis. Errors at two microphones of 10μs result in a range error of 1.67 m. As long as the source is not too far off-axis, these range errors are largely independent of the x and y location. If the timing errors reach 100μs at a single microphone, then the error can reach 6.25 m for a target at 20 m. While the ranging accuracy in the method used with an arbitrarily arranged array[7] also suffer from errors, these can be reduced by increasing the spacing between transducers.

Further information can be gained by deriving the angle of s with respect to the z plane[8] (elevation here, note change of convention for the z direction from that given in other references [9]). A microphone group would thus contribute both a range and an elevation. From this, azimuth is then derived using the calculated elevation to transform the co-ordinate system. Using multiple transducer groups then allows triangulation to the target, increasing accuracy. A similar algorithmic solution to this problem, deriving x, y and z co-ordinates from time of arrival of satellite signals can be applied to GPS positioning[11].

More recently, the development of optimisation methods to solve inverse problems has led to alternative ways to derive source location from the time of arrival at a number of receiving transducers. In one case, a stochastic search procedure to locate singing birds where predicted time delays to receiving microphones from a candidate source location were used to derive cross-correlation functions at the associated predicted delays, which were then summed[12]. The location with the maximal value then provides the estimated location. This system estimated source locations to an accuracy within 3 m over a 75 m range using eight microphones.

The application of optimization techniques allows a simplified approach to the location of acoustically active animals using arbitrarily placed receivers, or receivers matched the spatial locations of the animals. Here we present a solution to the location in 3-D space of actively foraging bats using optimisation algorithms taking data from an array of eight ultrasonic transducers.

2 APPLICATIONS TO BAT ECHOLOCATION

2.1 Practical considerations

While many acoustic tomography systems are deployed to look at cetaceans[13], some systems have been described which are specifically designed for echolocating bats[8]. Bats produce loud, regular ultrasonic echolocation calls which lend themselves to acoustic tomography since they are high amplitude and short duration. Unless the bat is flying close to clutter, the echolocation call is essentially produced in anechoic space, since excess atmospheric attenuation means that echoes from distant targets will not return to the recording microphones giving false locations. The application of such as system has led to some very useful and surprising insights into bat echolocation and behaviour[14-16].

Unfortunately, the highly directional nature of echolocation calls means that in a wide-spaced transducer array, not all receivers will receive the full call amplitude, and there may be wide variation in the amplitude of signals recorded. In addition, wideband ultrasound microphones are expensive and highly sensitive to humidity, while narrowband piezo devices are sensitive but limited in bandwidth. The main limitation in acoustic tomography of bats is that in order to record full bandwidth, the sampling rate on each channel of the digital acquisition device must be at least twice the highest frequency likely to be encountered, usually at least 250 kHz per channel. For an eight receiver array, this corresponds to a total throughput of 1.5 MHz, not easily achievable in a field-portable device.

The reason that acoustic tomography is a useful tool in the study of bats is that the precise nature of emitted signals is still largely unknown due to the highly directional emission pattern[17, 18]. Such a pattern tends to under-represent higher frequencies in the bat's FM echolocation call, as well as higher harmonics since these are only recorded when the bat is directly facing the recording microphone. It is precisely these frequencies which should allow higher target resolution. By knowing the exact distance and angle of a flying bat to a calibrated recording microphone, it should be possible to identify the maximum frequency, harmonic structure, bandwidth, sweep rate and intensity of echolocation calls.

2.2 A field acoustic tomography system

The system documented here uses eight receivers, though information on location can be extracted from any subset of five transducers. Each receiver station consists of three Murata 40 kHz piezo-electric receivers orientated at 120 azimuth to one another and elevated 30 relative to the horizontal axis. The three transducers are wired in parallel to a 60 dB gain three-stage amplifier section. The disadvantage of a 40 kHz transducer is that it has a very narrow -3dB bandwidth as it is effectively tuned to 40 kHz. This means that bats which use higher or lower frequencies may be missed. In practise, many bats use frequency modulated echolocation calls that pass through the 40 kHz band, are so will be picked up. The advantage of these transducers is that they are very sensitive and have very high signal to noise ratios and are relatively unaffected by low-ultrasonic or sonic noise in the environment. Three transducers are used as the transducers are very directional, as is the transmission beam pattern of bats. Thus, to maximise the chance of recording a bat on all eight receiver stations, transducers which point in a range of directions are useful. In order to accommodate the bandwidth of recording all eight channels, the signals are low-pass filtered with a -3dB point of 5 kHz and roll off of -6 dB per octave. This generates an envelope of the recorded signal which can be digitised at much lower rates that the full bandwidth signal. The output of the eight receiver stations is digitised at 12 bit and 62.5 kHz per channel (total throughput 500 kHz) using a 6062-E National Instruments PCMCIA card under LabView running on a laptop computer.

Echolocation signals are automatically identified using a thresholding technique, and the time of the start of the echolocation call extracted at a fixed threshold level over the mean noise floor level of each channel. Only those recording stations which have recorded a signal are used, so for each echolocation call, between 5 and 8 arrival times are extracted. The array of extracted times is then passed to the optimisation procedure. This is either implemented in Microsoft Excel, or Matlab. In either case, the optimisation procedure is the same. The optimiser is supplied with the x, y and z locations of each receiver, and the time of arrival of an echolocation call. The optimisation algorithm is then supplied with an arbitrary start location and emission time proposed for the bat, usually the centre of the array at a realistic height for the foraging bat, and an emission time of 0 seconds. These candidate co-ordinates and emission times are passed through each receiver station to give the time of arrival at each station using simple 3-D geometry if that bat was at the suggested location. The difference between the actual time of arrival and the calculated time is then squared and summed across the receiver stations to give the mean square error between the suggested times and real times. The optimisation algorithm then seeks to minimise this error by altering the location and time of emission of the echolocation call over 1000 iterations. There can be an issue of local minima where an apparent optimal solution is reached, but which is in fact far away from the real location. In order to identify these, if the newly calculated location is further away that expected based on the flight speed and pulse repetition rate of the bat under study, then the optimisation process is run a second time which moves the calculated location out of the local minimum and to a more realistic location. In most cases, local minima exist and the optimization process needs to be run twice. Careful placement of transducer stations, including a range in the z axis, can minimise this problem. The version running under Microsoft Excel uses the Generalised Reduction Gradient Algorithm, while that under Matlab uses the non-linear Gauss-Newton method from the optimization toolbox. Both algorithms yield comparable results.

One of the reasons for developing this tomography array is to understand how received call structure depends on the angle of the bat to the transducer. For this, a high quality microphone is also used in vicinity of the array. This microphone is either a Pettersson D-980 solid dielectric microphone which exhibits high sensitivity but does not exhibit a linear frequency response, or a Larson Davis ¼" microphone (grid-off) which has lower sensitivity but is linear to 100 kHz and can be used to measure sound pressure level. In each case, the output from this microphone is sampled at the same time as the tomography receiver stations on the same laptop using the same hardware start trigger, giving a 1:1 correspondence in received calls. The output from this microphone is sampled at 16 bit at 1.25 MHz on a National Instruments 6251 USB data acquisition device running from the same laptop as the tomography array data capture device. If required, both microphone types can be sampled at the same time at 1 MHz per channel (there are separate A/D converters on each channel in this device). The microphone is mounted on a 4 m high stand, ensuring that it more in-line with the emission height of the bat that a ground-based microphone which can influence the structure of the recorded signal[19].

3 INITIAL RESULTS

3.1 Theoretical resolution and calibrations

The two methods of using a solution for five simultaneous equations using matrix algebra and the optimization method were tested for the same five receiver array and source locations over a range of array patterns and source locations. In both cases, errors in time of receipt of signal were added at one or two receiver stations. For five receivers in the same arrangement for both methods, a source at 15 m distant with distances between receivers of 1 m resulted in an average spatial error of 1.8 m with a 10μs timing error at one receiver station and 4.6 m with the error at two receiver stations when using the matrix algebra method. Using the optimization method, exactly the same array, source location and timing error resulted in a mean location error of 0.89 m with a 10μs timing error at one receiver station and 4.35 m for the error at two receiver stations. In both cases, the error in the z location was the largest, but in all simulations, the optimization method appears to be more robust to timing inaccuracies.

In a field calibration, an ultrasound source from an 'ultrasonic tape-measure' (a broadband impulse source) was used to simulate a bat. The array was a four-armed arrangement, with two transducers per arm, each 2 m away from each other, with one receiver in each arm elevated by 1 m. The array was able to locate the source from 8 m away to an accuracy within 10 cm.

3.2 Field Results

Initial trials during the summer of 2007 used an array spacing of 1 m, 2 m or 4 m between receiving stations. The closer 1m spacing ensured that all receiver stations detected each individual echolocation call, but resulted in greater errors in location with distance. A trace of a foraging bat (*Pipistrellus pipistrellus*) tracked by eight receiver stations is shown in Figure 1. The rapid changes in direction are real phenomena and not errors in location. In the field, wider spacing of the receiver stations mean that not all stations will receive each call, but they ensure greater accuracy with greater range. The fact that the receiver stations to not have to have a fixed geometry means that the arrangement and spacing can be adapted to individual recording situations. Loud bats using low frequency calls can be detected at greater range, and the spacing of the array can be made larger to accommodate this, meaning greater accuracy at greater range. Similarly, bats which forage close to vegetation and have quieter more directional calls can be accommodated using a more closely spaced array. One of the greatest difficulties in the field is satisfying condition 1, in knowing with sufficient accuracy the location of the receiver stations. Taking a nominal accuracy of time of arrival of the signal of 10 μs translates to knowing the receiver location to within 3.4 mm, which is certainly a challenge outside of the laboratory.

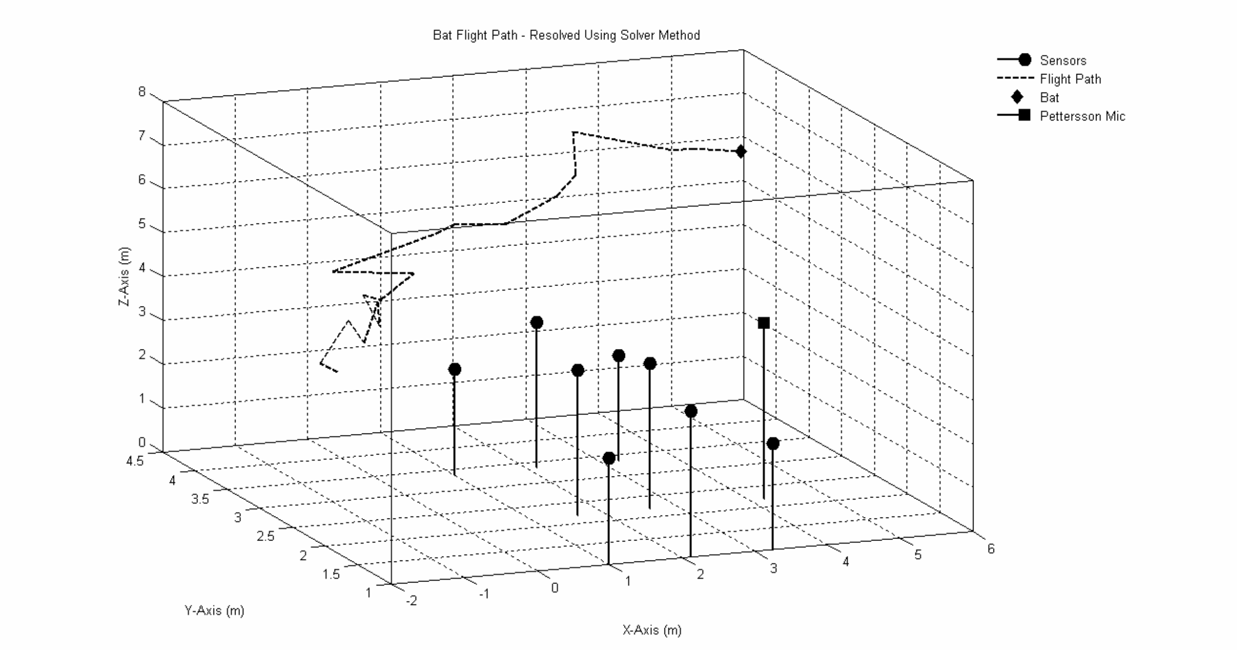

Figure 1. Track of a foraging *Pipistrellus pipistrellus* as revealed by acoustic tomography.

3.3 Future work

Simulations are currently underway to optimize the receiver placement. A model has been developed in Matlab where the space around the array has been divided into a 3-D square mesh. A source location is then placed at each vertex in turn, and the source location solved for that array arrangement. This identifies the regions of space where source locations cannot be derived from the optimization method. To date, as along as there is some variation in the source location in the x, y and z planes, the only source locations that do not have solutions using the optimization method lie in small regions within the boundaries of the array. Development of this technique will allow us to identify the best arrangement of receiver stations to track echolocating bats. Bats will be recorded throughout the summer of 2008 to measure echolocation source levels, signal structure and directionality.

4 ACKNOWLEDGEMENTS

This work is part of the BIAS (Biologically Inspired Acoustic Systems) consortium, funded by the EPSRC.

5 REFERENCES

1. J.K. Lewis, J. Rudzinsky, S. Rajan, P.J. Stein, and A. Vandiver, 'Model-oriented ocean tomography using higher frequency, bottom-mounted hydrophones'. J. Acoust. Soc. Am., 117(6), p. 3539-3554. (2005).
2. B.S. Hoyle, 'A schema for generic process tomography sensors'. Ieee Sensors Journal, 5(2), p. 117-124. (2005).

3. J.L. Spiesberger, 'Locating animals from their sounds and tomography of the atmosphere: Experimental demonstration'. J. Acoust. Soc. Am., 106(2), p. 837-846. (1999).
4. J.L. Spiesberger, 'Geometry of locating sounds from differences in travel time: Isodiachrons'. J. Acoust. Soc. Am., 116(5), p. 3168-3177. (2004).
5. J.L. Spiesberger and K.M. Fristrup, 'Passive localization of calling animals and sensing of their acoustic environment using acoustic tomography'. Am. Nat., 135(1), p. 107-153. (1990).
6. E. Paulus, 'Sound localization cues of binaural hearing'. Laryngo-Rhino-Otologie, 82(4), p. 240-248. (2003).
7. J.L. Spiesberger, 'Hyperbolic location errors due to insufficient numbers of receivers'. J. Acoust. Soc. Am., 109(6), p. 3076-3079. (2001).
8. R. Aubauer, *Korrelationsverfahren zur flugbahnverfolgung echoortender fledermause*. PhD Thesis. Darmstadt, Germany. (1995).
9. W.W.L. Au and D.L. Herzing, 'Echolocation signals of wild atlantic spotted dolphin (stenella frontalis)'. J. Acoust. Soc. Am.,, 113(1), p. 598-604. (2003).
10. R.A. Holland, D.A. Waters, and J.M.V. Rayner, 'Echolocation signal structure in the megachiropteran bat rousettus aegyptiacus geoffroy 1810'. J. Exp. Biol., 207(25), p. 4361-4369. (2004).
11. R. Bucher and D. Misra, 'A synthesizable vhdl model of the exact solution for three-dimensional hyperbolic positioning system'. Vlsi Design, 15(2), p. 507-520. (2002).
12. D.J. Mennill, J.M. Burt, K.M. Fristrup, and S.L. Vehrencamp, 'Accuracy of an acoustic location system for monitoring the position of duetting songbirds in tropical forest'. J. Acoust. Soc. Am., 119(5), p. 2832-2839. (2006).
13. E.M. Nosal and L.N. Frazer, 'Sperm whale three-dimensional track, swim orientation, beam pattern, and click levels observed on bottom-mounted hydrophones'. J. Acoust. Soc. Am., 122(4), p. 1969-1978. (2007).
14. M.W. Holderied and O. von Helversen, 'Echolocation range and wingbeat period match in aerial-hawking bats'. Proc.Roy. Soc. B., 270(1530), p. 2293-2299. (2003).
15. M.W. Holderied, G. Jones, and O. Von Helversen, 'Flight and echolocation behaviour of whiskered bats commuting along a hedgerow: Range-dependent sonar signal design, doppler tolerance and evidence for 'acoustic focussing''. J. Exp. Biol., 209(10), p. 1816-1826. (2006).
16. J. Schul, F. Matt, and O. von Helversen, 'Listening for bats: The hearing range of the bushcricket phaneroptera falcata for bat echolocation calls measured in the field'. Proc. Roy. Soc. B., 267(1454), p. 1711-1715. (2000).
17. D. Henze and W.E. Oneill, 'The emission pattern of vocalizations and directionality of the sonar system in the echolocating bat, pteronotus-parnelli'. J. Acoust. Soc. Am., 89(5), p. 2430-2434. (1991).
18. K. Ghose, C.F. Moss, and T.K. Horiuchi, 'Flying big brown bats emit a beam with two lobes in the vertical plane'. J. Acoust. Soc. Am., 122(6), p. 3717-3724. (2007).
19. M.E. Jensen and L.A. Miller, 'Echolocation signals of the bat eptesicus serotinus recorded using a vertical microphone array: Effect of flight altitude on searching signals'. Behav. Ecol. Sociobiol., 47(1-2), p. 60-69. (1999).