

Exploiting the physics: towards Doppler-based navigation with a bat-inspired mobile robot

Jose M. Carmena* and John C.T. Hallam

Institute of Perception, Action and Behaviour
Division of Informatics, University of Edinburgh
5 Forrest Hill, EH1 2QL Edinburgh, Scotland, UK

Abstract. This work shows some preliminary results in the use of Doppler-shifts for ultrasound-based navigation in mobile robots. Doppler-shifts are a rich source of information which is not exploited by commercial ultrasonic range sensors like the Polaroid. Bats, our source of inspiration, exploit the physics of echolocation and make very reliable use of Doppler-shifts for several tasks. RoBat, a 6 DOF biomimetic sonarhead mounted on a mobile robot, is the platform on which the experiments were performed. A cochlear model for extracting Doppler-shifts from the echoes is introduced, followed by two examples of Doppler-shift related tasks: a convoy controller and an obstacle avoidance task through acoustic flow. Preliminary results for both tasks are presented, showing better performance in the former one. The cosine dependence of Doppler-shifts results in a source of inaccuracies for reflector position estimation when the bearing angle is close to 0, suggesting a limited use of Doppler-shifts for tasks, such as obstacle avoidance, in which a crude estimation of the bearing angle suffices.

1 Introduction

Most research on ultrasonic sensors for mobile robots has suffered from the ‘ultrasonic sensor = Polaroid range sensor’ fallacy. This sensor [2], by extracting Time-Of-Flight (TOF) for the first echo only, discards most of the information present in the echo signal. Hence extensive post-processing has to be performed on large numbers of range readings to construct consistent and reliable environment models [11, 9] out of them.

In fact, the use of ultrasonic sensors for mobile robot navigation has been under-rated for many years. Because of the wide beam width and unwanted reflections problems, many researchers became frustrated and decided to use other sensors instead such as vision cameras and laser range sensors which are more expensive and require much more processing time. Other people, like for example [9, 7, 14, 8], made use of these ultrasonic sensors in an intelligent way. Moreover, we believe that, from the study of biological acoustic sensorimotor systems [19, 13], much more information can be extracted from these echoes. Doing so leads to improved robotic ultrasonic sensors [17].

Our inspiration comes from biosonar, *i.e.* how bats exploit the acoustics of echolocation. CF-FM bats, the kind of bats in which we are interested, are given this non-taxonomic appellation because of the time-frequency structure of their echolocation pulses, often dominated by prolonged CF (=constant frequency) signal portions, although frequency-modulated “tails” are always present as well. Allocation of energy to the CF and FM portions of a signal varies with behavioral context [24]. Well-known species of CF-FM bats are the Greater Horseshoe Bat *Rhinolophus ferrumequinum* and the Mustached Bat *Pteronotus parnellii*.

The relationship between bats and robots arises because the sensor interpretation problems of bats, while navigating in cluttered environments such as forests, are very similar to those of mobile robots provided with ultrasonic sensors when navigating in laboratories. Moreover, the constant frequency pulse emitted by narrowband echolocators (as the CF-FM bat *Rhinolophus ferrumequinum*) is analogous to the one typically emitted by robotic ultrasonic sensors in terms of bandwidth. Hence, a better understanding of how these perceptual and motor mechanisms actually work in bats could improve the design of mobile robot navigation controllers.

The idea underlying our work with respect to the use of Doppler in mobile robots navigation is that, when moving, dynamically varying acoustic images are generated. These images could correspond, at least

* Email: jose@dai.ed.ac.uk

in principle, with the images bats get while flying through their environments. Thus, in this work, we are interested in investigating whether the use of the Doppler information is feasible for indoor mobile robot navigation or not.

RoBat, a mobile platform provided with a biomimetic sonarhead on which the experiments presented in this paper are performed, is described in section 2. Section 3 defines Doppler-shifts in a robotic context (section 3.1) followed by a description of how these Doppler-shifts are extracted (section 3.2) using a cochlear model. Finally, two preliminary examples of Doppler-related tasks performed by RoBat — a convoy controller and an obstacle avoidance task through acoustic flow — are presented (section 4) and discussed (section 5).

2 RoBat: an echolocating mobile robot

The RoBat project aims to investigate bat biosonar as a biological approach to mobile robot navigation, *i.e.* tries to understand how echolocating bats perform navigation tasks such as obstacle avoidance and prey capture, and how can this be applied to ultrasonic-based navigation in mobile robots. Hence the goals of creating a tool such as RoBat are twofold: It will help engineers to better understand the relationships existing between environment features and their corresponding acoustic images in a dynamic context. At the same time, such a tool will also help biologists studying echolocation in bats better to understand what type of information is available to the bat while performing particular tasks [15].

RoBat consists of three main components: a biomimetic sonarhead (described below), a 3 DOF mobile platform and a signal processing package whose operations, performed upon the received echoes, are based upon a filterbank model of the processing performed by the mammalian cochlea [23]. These three components are all controlled and integrated into a single system by software running on a Pentium III PC (Linux).

The 6 DOF sonarhead, as indicated in Fig. 1 (left), allows panning and tilting of the neck, and independent panning and tilting of each of the two ears (receivers). The ultrasonic transducers are Polaroid electrostatic transducers. The motors driving the different axes are standard radio-control servomotors [16].

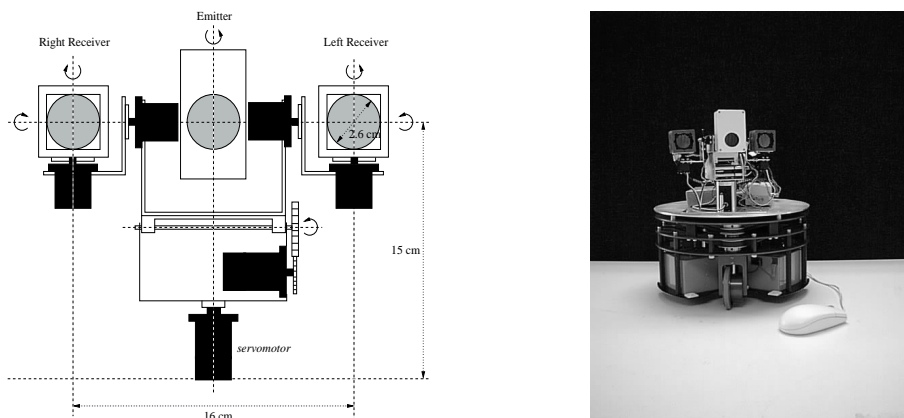


Fig. 1. Front view of the sonarhead consisting of the central emitter fixed to the head and the two receivers each independently orientable (left). RoBat, a biomimetic sonarhead mounted on a mobile platform. (right).

The transmitter is mounted at the centre of the head, moving along with the pan and tilt movements of the neck. The transmitter electronics allows the transmission of arbitrary waveforms, the latter are generated by the D/A converter board at 1 Msample/s (Gage CG1100 2 channel/12 bit/80 Msample/s D/A card) in the control computer. The ultrasonic transducer is a Polaroid electrostatic transducer: the output voltage consists of a bias voltage of 150V upon which the arbitrary waveform, amplitude 150V, is superposed. Detection and amplification of the reflected echoes is performed by the receiver electronics, which is mounted, together with the transducers, on pan/tilt servos. The output signals from the receivers are sampled at 1 Msample/s by the A/D converter board (Gage CS512 2 channel/12 bit/5 Msample/s) in the control computer. All further processing of the received data is performed on the control computer.

To provide the required mobility to the sonarhead (Fig. 1 (right)), we use a commercially available mobile platform [20]. This 3DOF platform has an on-board controller that receives motion-commands from

and sends back status information to the control computer through its serial port. The mobility offered by this robot base is used to generate dynamically varying acoustic images, i.e. acoustic flow, that correspond, at least in principle, with the images bats get while flying through their environments.

This setup when mounted on the mobile platform (Fig. 1 (right)) allows insonification of arbitrary points in space, independent of the orientation of the mobile platform. Taken together with the capability of the sonarhead to independently orient the ears it allows us to model different strategies that might be used by bats to actively explore their environment [15].

A block diagram of the modular architecture can be seen in figure 2. As seen in the figure, all the monaural modules, such as the signal processing and cue extraction, are duplicated because of the two receivers. The received echoes sampled by an A/D converter are fed into the signal processing module, whose operations are based upon a filterbank model of the processing performed by the mammalian cochlea [23]. The processed data coming from each of the monaural modules is then integrated by the binaural module, resulting in behaviours such as obstacle avoidance, motion detection — a subset of the former — and target tracking. Also, the length and call-frequency of the emitted pulses are recalculated according to the behaviour being performed by the robot at the end of each loop. All processing and interpretation of the received data is performed on the control computer, operating at a speed of 12 Hz, *i.e.* 15 complete *sense-and-act* loops per second.

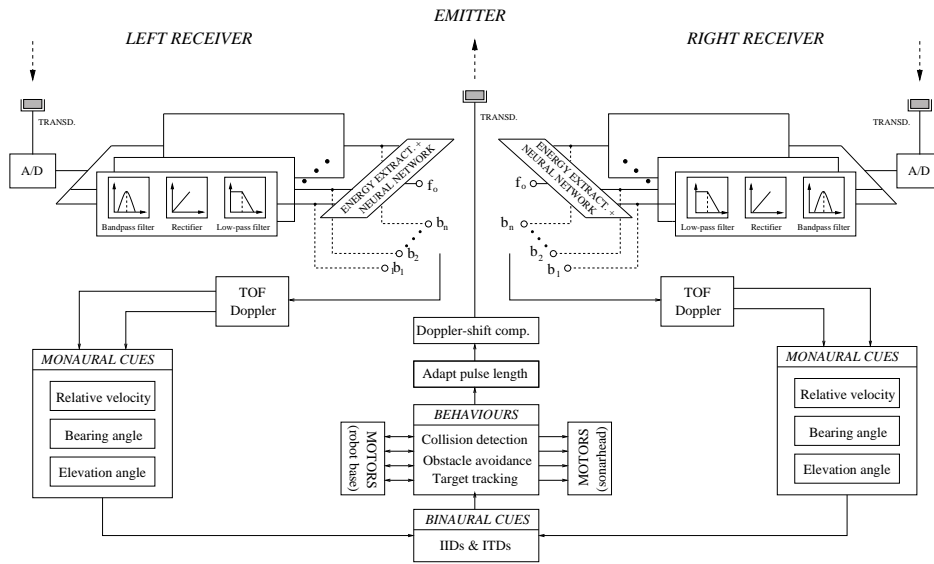


Fig. 2. RoBat's architecture.

3 Towards the use of Doppler-shifts in mobile robot navigation

3.1 Definition: Doppler-shifts in a robotic context

Next we will define and illustrate the concept of Doppler-shift in the robotic context of figure 3.

Let's consider two robots, $R1$ and $R2$, approaching each other in a straight line (relative angle $\simeq 0$) at speeds V_{R1} and V_{R2} . If f is the frequency emitted by $R1$, c the speed of sound in air (approx. 345 m/s) and φ the angle between the relative velocities of $R1$ and $R2$, we can define the Doppler-shifted f because of V_{R1} as

$$f_{R1} = f \left(1 + \frac{V_{R1}}{c} \cos \varphi \right)$$

with $\cos \varphi \simeq 1$ because of $\varphi \simeq 0$. Hence the frequency arriving at $R2$, because of V_{R2} , will be

$$f \left(1 + \frac{V_{R1}}{c} \right) \left(1 + \frac{V_{R2}}{c} \right) = f \left(1 + \frac{V_{R1} + V_{R2}}{c} + \frac{V_{R1} * V_{R2}}{c^2} \right) \simeq f \left(1 + \frac{V_{R1} + V_{R2}}{c} \right)$$




Fig. 3. Emitted Doppler-shifted call by $R1$ (left) and reflected Doppler-shifted echo by $R2$ (right).

and therefore, the Doppler-shifted frequency of the reflected echo will be

$$f_{R2} = f \left(1 + \frac{V_{R1} + V_{R2}}{c}\right) \left(1 + \frac{V_{R2}}{c}\right) \left(1 + \frac{V_{R1}}{c}\right) \simeq f \left(1 + \frac{V_{R1} + V_{R2}}{c}\right)^2 \simeq f \left(1 + 2\left(\frac{V_{R1} + V_{R2}}{c}\right)\right)$$

Defining the relative velocity between $R1$ and $R2$ as $v = V_{R1} + V_{R2}$, the Doppler-shift (δf) is then defined as

$$\delta f = 2f \frac{v}{c} \cos \varphi.$$

3.2 Doppler-shift extraction: a cochlear model

The bat’s auditory system is structured similarly to that of other mammals (see, e.g., [18]). The incident sound is directed towards the ear canal by the pinna, which is very mobile and has a highly convoluted surface in many bat species. The transduction stage located in the inner ear (cochlea) performs a joint time-frequency analysis of the incoming signal. A simple model of this analysis is a bandpass filter bank with subsequent demodulation in each channel by a combination of half-wave rectification and lowpass filtering. In the FM bats the layout of the auditory filter bank follows the general mammalian pattern of keeping filter quality constant as center frequency varies. CF-FM bats deviate considerably from this pattern by forming an auditory fovea in the frequency range where the carrier frequencies of the echoes are kept by the Doppler-shift compensation behavior of these animals [1]. Towards the center of the fovea, filter qualities rise steeply to the highest values known (maximum $Q_{10dB} \sim 400$ in *Rhinolophus ferrumequinum*); outside this frequency band the CF-FM bats follow the general mammalian pattern.

As stated in [10], “a good cochlear model should take real sound as input, in real time, and product output that resembles the signals on the cochlear nerve”. At this point, the importance of both a biologically plausible — to keep our biomimetic approach — and computationally reliable — to perform real time processing — model of the bat cochlea for extracting Doppler-shifts from the echoes is clear. Thus, in this work, the cochlear model used is a conventional digital analysis filterbank whose frequency response characteristics splits, by a set of parallel bandpass filters with a uniform distribution of center frequencies across the channels, the incoming signal into a corresponding number of sub-bands. The transduction of the movement of the membrane into neural activity, implemented biologically by the hair cells of the cochlea, is modelled fairly well by processing the outputs of these sub-bands (or channels) with a half-wave rectifier and a low-pass filter [22]. These last two steps are equivalent to a simple amplitude demodulation scheme that approximately recovers the envelope of the outputs of the bandpass filters in the filterbank. A diagram of a generic cochlear model is shown in figure 2.

The channels are located at centre frequencies d Hz apart across the band from $f - \frac{n}{2}d$ to $f + \frac{n}{2}d$ Hz, f being the frequency of the center of the “acoustic fovea” ($50kHz$ for our system), n the number of channels (7 for our system) and d half of the single channel $3dB$ bandwidth. The frequency range determined by n suffices for the range of Doppler-shifts received by RoBat. Because of the real time performance constraint, the type of filters chosen need to be cheap computationally — the bandpass and low-pass filters are 2nd and 1st order Butterworth respectively.

For Doppler-shift estimation, we need a higher frequency resolution than that provided by a single channel of the n channel filterbank (this is analogous to the sub-pixel resolution problem in machine vision). To perform frequency discrimination using the filterbank outputs we need to interpolate the frequency of the echo. For doing so, an artificial neural network (ANN), which will determine the frequency of the principal component of the echo, is applied to the filterbank output in the following way. First, the energy in each channel is calculated and normalised. Then, as seen in figure 4 (left), the ANN fits as a sliding window

containing channels centered on the 2 channels with highest energy. The energy of these 2 channels (at centre frequencies f and $f + d$ in the figure) feed the 2 black-filled input units of the ANN. The interpolated frequency (f_o in the figure) will be somewhere between the centres of these 2 channels.

A 6-3-1 backpropagation ANN was trained using simulated noisy signals, learning rate of 0.1 or smaller and no momentum, learning the whole acoustic fovea in 14616 iterations. The signals were made of a sinusoid of given frequency – the frequency to be estimated – added to uniformly distributed random noise.

In figure 4 (left), the output of the filterbank for a 150 Hz Doppler-shifted echo is shown. The echo is a 50150 Hz noisy sinusoid. As seen in the figure, the 2 channels (50140 and 50160 Hz) that bracket the frequency of the echo have higher energy than the rest. Thus, in this case, the ANN is centered on these 2 channels for estimating the signal frequency.

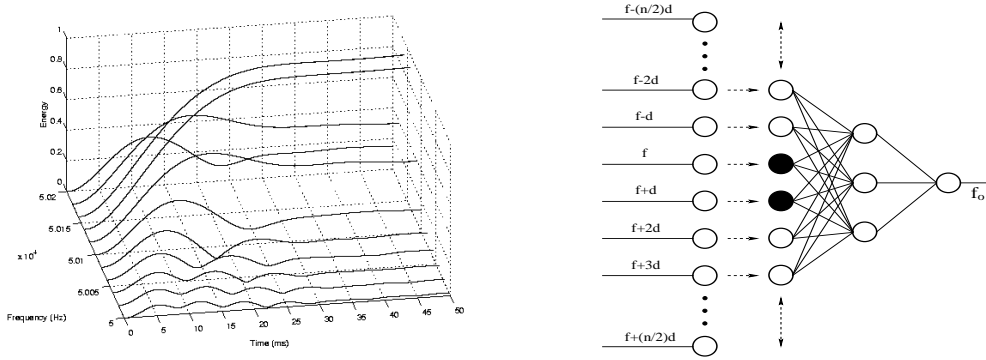


Fig. 4. Filterbank output for a 150 Hz shifted echo (left). Sliding ANN for frequency extraction (right).

An example of robot-motion-induced Doppler-shift can be seen in figure 5, in which RoBat was moving at constant velocity towards a 5 cm diameter post positioned to one side of its path at a bearing angle α_n (figure 5 (left)).




Fig. 5. RoBat driving by a post (left). Estimated *vs.* true Doppler-shifts (right).

Since Doppler-shifts depend on the cosine of the bearing angle, a maximum Doppler-shift is obtained when the robot is far from the reflector because the bearing angle between the robot and the post is close to 0° (α_1 in figure 5 (left)). This can be seen in figure 5 (right), in which the estimated Doppler-shifts while RoBat moves by the post are plotted against the mean of true Doppler-shifts. The latter are estimated by trigonometry given that the positions of both the post and the robot (through the TOF of each echo) are known, as well as the velocity of the robot, for each of the data points of a single trial. The figure shows the mean and standard deviation of 10 trials for the case of the estimated Doppler-shifts. As seen in the figure, the cochlear model performs satisfactorily well despite some fit errors near to the 90° relative position.

4 Doppler-related tasks in bats and robots

CF-FM bats detect interesting properties (*e.g.* food) in echoes, while navigating in cluttered environments, by the frequency modulation patterns produced by the wing-beat cycles of fluttering insects [21, 6, 25]. This target selection without ‘recognition’ behaviour was successfully implemented by Walker *et al.* [26], in which a previous version of RoBat was set to track a computer fan in a cluttered environment. Part of this clutter was composed of other fans oscillating at a different rate than the target fan.

In the RoBat project we are currently investigating the use of Doppler for two particular tasks: convoy navigation and obstacle avoidance, described next.

4.1 Convoy controller

The fact that this bat modifies — increasing or decreasing — the carrier frequency of its own calls to compensate the Doppler-shift produced when the bat, the reflector or both are moving indicates the importance of the Doppler-shift information in the received echoes [24]. Inspired by this behaviour, a collision detection and smooth convoy navigation controller for a mobile robot using Doppler information was devised [5] and is described next.

In a real environment, after defining *maximum Doppler* as the the maximum Doppler-shift the robot could observe for a given velocity assuming a static reflector at 0° bearing angle, we can find at least three different Doppler-related situations:

- **Doppler ≤ 0** There is no reflector in the way or there is a moving reflector whose relative velocity with respect to our robot is zero or negative. In this case the robot can navigate safely within its perceptual range².
- **$0 < \text{Doppler} < \text{Max. Doppler}$** There is either a static reflector in the way or a moving reflector with a positive relative velocity with respect our robot but with a bearing angle sufficient to avoid a collision.
- **Doppler $\geq \text{Max. Doppler}$** There is a moving reflector in the way with a positive relative velocity with respect our robot and its bearing angle is 0 or close to 0. In this case the robot should change its path immediately for avoiding collision.

Thus assuming that a moving reflector (*e.g.* another robot) is within the robot’s perceptual range and at a small bearing angle (close to 0), a simple collision detection and convoy controller, which modifies the velocity of the robot in order to keep the Doppler shift close to 0, can be implemented as shown in figure 6.




Fig. 6. High-level block diagram of the Doppler-based collision controller.

As seen in the figure, after extracting the average Doppler-shift (D) from the individual Doppler-shifts received by each of the receivers, according to its sign and value, the velocity of the robot (V_D) is estimated. In case of D being greater than the maximum Doppler (D_{max}), *i.e.* reflector moving towards the robot, an immediate escape maneuver such as turning to one side would be adopted for avoiding collision. Next, the velocity is set to its new value (V_s) which is limited by a preset value (60 cm/s). Finally, from the readings of the motor encoders, the current velocity (V_e), which is used for both calculating the new velocity and the maximum Doppler, is estimated.

The experiment designed for assessing the performance of this controller consisted in RoBat moving along a corridor right behind ‘Gillespie’, a RWI-B21 robot moving in the same direction (figure 7 (upper-left)). Three different cases were tested: In the first case, Gillespie starts moving keeping its velocity constant at 35

² The maximum range from which the robot can accurately sense the environment. In RoBat this range is 3 m.

cm/s (upper-right graph). In the second case, Gillespie starts moving and keeps accelerating until reaching the maximum allowed velocity of 60 cm/s (lower-left graph). Finally, in the last case, Gillespie starts moving with high acceleration until reaching the maximum speed of 60 cm/s and then decelerates smoothly to 30 cm/s before stopping (lower-right graph).

The controller performed satisfactorily well. For each of the three cases, the Doppler-shift, maximum Doppler and RoBat’s velocity were estimated for every echo, as shown in the graphs of figure 7. Because of the large size and weight of Gillespie, the irregularities of the floor affect its velocity more drastically resulting in a small ripple in the velocity curve of the three graphs. This small ripple in the velocity can be clearly seen in the the Doppler-shift and maximum Doppler curves.

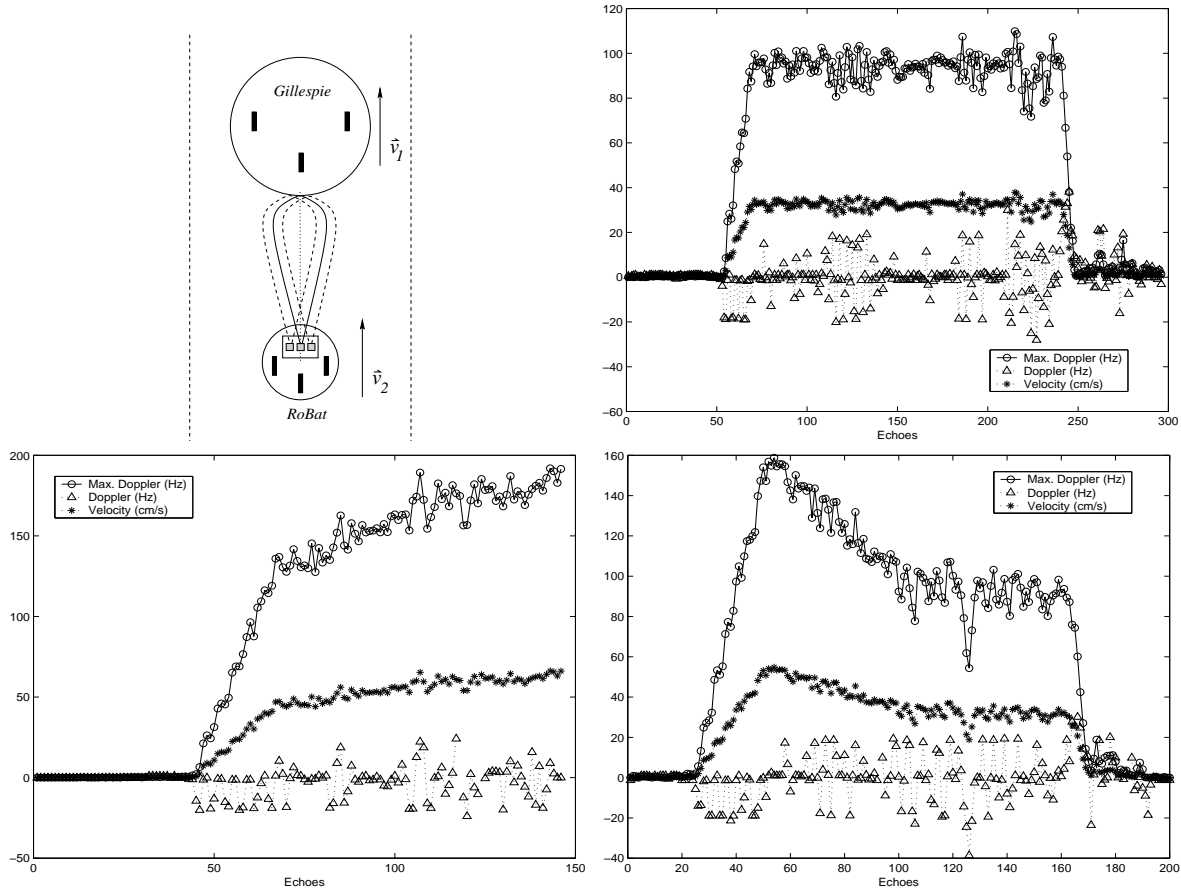


Fig.7. Measurement setup for the motion detection task (upper-left). Doppler-shifts and velocity values estimated by RoBat while Gillespie was moving at constant velocity (upper-right), accelerating (lower-left) and decelerating (lower-right).

4.2 Obstacle avoidance

Another interesting example of the use of Doppler-shifts by CF-FM bats is Müller’s acoustic flow theory for obstacle avoidance [12]. Müller bases his acoustic flow hypothesis on two perceptual dimensions — changes in Doppler-shift and sound pressure amplitude — which CF-FM bats may employ for the extraction of two-dimensional spatial information. In [3, 15], a first step towards the use of acoustic flow for mobile robot navigation was introduced, showing how robot-motion-induced Doppler-shifts can be extracted from echoes and used for calculating the bearing angle between the robot and a reflector. If we combine these Doppler-shifts with the echoes’ TOF, a crude estimation of the cartesian position of a reflector can be obtained as we will see next. We use TOF instead of changes in echo amplitude (as in [12]) because of its being more reliable in our system.

We aim to investigate whether the accuracy of this crude estimation suffices for an obstacle avoidance task in a robotic context. Acoustic flow is monaural, *i.e.* intensity differences (TOFs in our system) and Doppler-shifts can be extracted independently by each of RoBat’s ears. The way this information is combined at higher levels of the bat’s auditory system (binaurality) is nowadays a matter of investigation. In our system we simply take the averages of the two TOF and Doppler estimations.

An experiment in which RoBat had to avoid two obstacles (posts) on its path was designed. In figure 8 (left) we can see the Doppler-shift estimations of the two posts from the cochlear model. The first wiggling part of the plot is due to noisy estimations when the robot is not yet moving, having the first post in front of it. As soon as RoBat starts moving, the post is detected and RoBat starts turning until the bearing angle increases — Doppler-shift decreases — up to a desired value (in our system this was set to 50°) when the second post was detected, making the avoiding process start again. The 0 Hz Doppler-shifts between posts and after the second post are due to the lack of echoes at that particular moment. A crude position estimation of the two posts along the robot’s trajectory is shown on the right plot of figure 8.

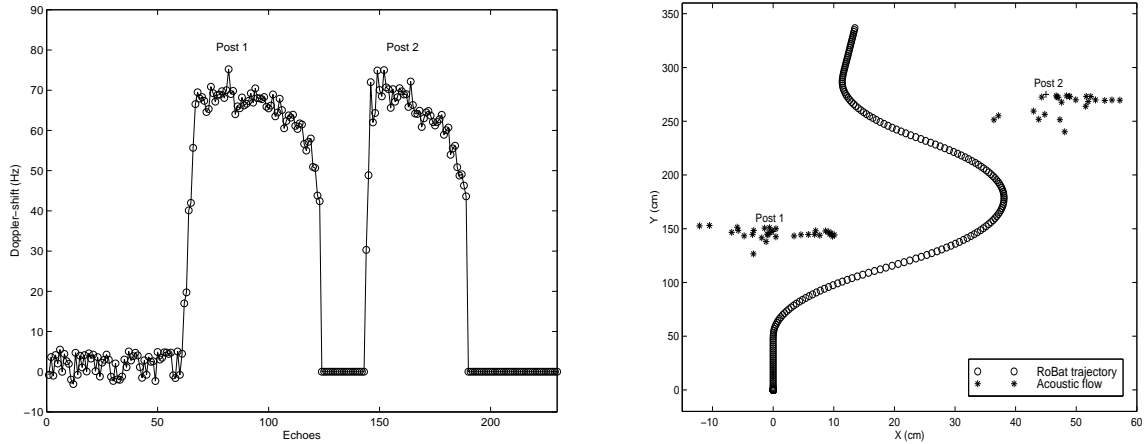


Fig. 8. Estimated Doppler-shifts of the two posts (left). Posts position estimations along RoBat’s trajectory through acoustic flow (right).

5 Discussion

The real world is dynamic: things move and change their position. We believe that for efficient navigation, an agent must take into account other agents moving around in the same environment, and likewise this should be done at the lowest level and, therefore, in the simplest way, before going to higher levels of abstraction.

CF-FM bats, the creatures from which we get the inspiration for our work, exploit the physics of echolocation, making use of Doppler-shifts for several tasks. Since, surely, bats do not perform Fourier transforms for frequency discrimination, we use a software model of the mammalian cochlea instead, consisting of a filterbank with its outputs connected to an ANN for solving the sub-resolution problem.

In the first example of a Doppler-shift related task we saw how a simple and efficient motion controller for collision detection and smooth convoy navigation can be devised. Since Doppler-shifts depend on the cosine of the bearing angle between the robot and the reflector, an angle close to 0° for situations in which a moving reflector — *e.g.* another robot — is moving in the same direction within its perceptual range can be assumed. In such situations, the controller showed how the relative velocity of the other agent can be estimated, resulting in smooth navigation while keeping the relative velocity close to 0, as seen in the results of the experiments using RoBat and Gillespie.

In the second example, RoBat was set to perform a simple obstacle avoidance task by means of acoustic flow, *i.e.* estimating the Cartesian coordinates of the obstacles from the TOF and bearing angle — calculated from the Doppler-shifts — of the received echoes along RoBat’s trajectory. The aim is to investigate whether

the crude position estimation provided by the acoustic flow suffices for an obstacle avoidance task in a robotic context. One factor to bear in mind is the difference in Doppler-shift resolution between real bats and RoBat. This is because of the difference in their velocities when navigating; in the bat this can go up to 5 m/s whereas for indoor robots 0.5 m/s is quite a fast velocity (any velocity higher than 1 m/s would be dangerous in our opinion). In addition, the 83 kHz carrier frequency of the CF-FM bat is higher than RoBat's 50 kHz . As a result, the range of Doppler-shifts obtained by RoBat is much smaller than by the real bat and, therefore, the bearing angle estimation will be worse. The effect of this poor resolution was shown in the non-consistent position estimation of the two posts of figure 8. Nevertheless, it suffices for avoiding the posts as seen by the trajectory performed by the robot. These results match with Müller's hypothesis [12] in which acoustic flow provides crude localisation information which suffices for tasks like obstacle avoidance in which lower accuracy is needed than in other tasks, *e.g.* prey capture.

Another interesting aspect to mention is that because of the cosine dependence of Doppler, bats are very sensitive to motion in the periphery (bearing angle close to 90°) and not so sensitive in the frontal plane (bearing angle close to 0°). This results in a better Doppler-shift estimation in the periphery, *i.e.* when the relative angle between the bat and the reflector is close to 90° . This affects the acoustic flow task in terms of a poor position estimation of the reflector until the bearing angle starts to increase.

The conclusion of this work is the following: Doppler-shifts are a rich source of information which is not exploited by commercial ultrasonic range sensors like the Polaroid. Why not use it? Then we have demonstrated that this is possible in a simple robotic context.

Moreover, Doppler is monaural. Therefore, when the robot is moving, both the TOF and bearing angle of a reflector can be extracted with a single sensor. This, extrapolated to the Polaroid range sensor, could be imagined as a sensor giving 2 numbers (TOF, Doppler) instead of one (TOF) from each echo. A conventional sonar-ring covering the whole body of a commercial robot would then provide Doppler information of the robot's environment. An improved cochlear model for Doppler-shift estimation (currently under investigation [4]) would improve the precision in calculating the bearing angle and therefore could allow for reducing the 30° uncertainty given by the main lobe of the commercial ultrasonic transducer's radiation pattern.

5.1 Further work

The next step is to use an improved cochlear model in more complex and realistic environments, in which the presence of more dominant frequencies in the echoes should cause the model fail because of the cross modulation terms introduced. The model presented assumes a single dominant component in the echo. The natural question here is whether this assumption holds: is it possible to sense the world in terms of a single main frequency component at a time — considering short time windows along an echo — or not?

A long term goal of the RoBat project is to integrate in RoBat's architecture (figure 2) these simple behaviours together with a target tracking behaviour in order to emulate real bat-like obstacle avoidance while engaged in a prey pursuit task.

Acknowledgments

Jose M. Carmena is supported by the European Union TMR Network SMART2 under contract number FMRX-CT96-0052. Facilities for this work were provided by the University of Edinburgh.

References

1. O. Behrend, M. Kössl, and G. Schuller. Binaural influences on doppler shift compensation of the horseshoe bat *Rhinolophus rouxi*. *J Comp Physiol A*, 185:529–538, 1999.
2. C. Biber, S. Ellin, E. Shenk, and J. Stempeck. The Polaroid ultrasonic ranging system. In *67th Convention of the Audio Engineering Society*, New York, October 1980.
3. J.M. Carmena and J.C.T Hallam. Estimating Doppler-shift with a coarse cochlear filterbank. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 1, pages 221–226, 2000.
4. J.M. Carmena and J.C.T Hallam. A comparison of methods for estimating Doppler-shift using bat-inspired cochlear filterbank models. In *Fourth European workshop on advanced mobile robots (EUROBOT'01)*, 2001. (submitted).

5. J.M. Carmena and J.C.T. Hallam. Doppler-based motion controller for an echolocating mobile robot. In *Towards Intelligent Mobile Robots. Technical Report Series, Department of Computer Science, Manchester University*, 2001.
6. R. Kober and H.-U. Schnitzler. Information in sonar echoes of fluttering insects available for echolocating bats. *J. Acoust. Soc. Amer.*, 87:882–896, 1990.
7. R. Kuc. Three-dimensional tracking using qualitative bionic sonar. *Robotics and Autonomous Systems*, 11:213–219, 1993.
8. D. Lee. *The Map-Building and Exploration Strategies of a Simple Sonar-Equipped Mobile Robot*. Distinguished Dissertations in Computer Science. Cambridge University Press, 1996.
9. J.J. Leonard and H. F. Durrant-Whyte. *Directed Sonar Sensing for Mobile Robot Navigation*. Kluwer Academic Publishers, 1992.
10. R.F. Lyon and C. Mead. An analog electronic cochlea. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 36(7):1119–1134, 1988.
11. H. P. Moravec and A. Elfes. High resolution maps from wide angle sonar. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 116–121, 1985.
12. R. Müller and H.-U. Schnitzler. Acoustic flow perception in cf-bats: Properties of the available cues. *J. Acoust. Soc. Amer.*, 105(5):2958–2966, 1999.
13. P. Nachtigall and P. Moore, editors. *Animal SONAR Processes and Performance (NATO ASI Series)*. Plenum Press, 1988.
14. H. Peremans. *Tri-aural perception for mobile robots*. PhD thesis, University of Gent, 1994.
15. H. Peremans, R. Müller, J.M. Carmena, and J.C.T. Hallam. A biomimetic platform to study perception in bats. In G.T. McKee and P.S. Schenker, editors, *Proceedings of SPIE: Sensor Fusion and Decentralized Control in Robotics Systems III*, volume 4196, pages 168–179, 2000.
16. H. Peremans, A. Walker, and J.C.T. Hallam. A biologically inspired sonarhead. Technical Report 44, Dep. of Artificial Intelligence, U. of Edinburgh, 1997.
17. H. Peremans, A. Walker, and J.C.T. Hallam. 3D object localisation with a bionic sonarhead: inspirations from biology. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, volume 4, pages 2795–2800, Leuven, Belgium, May 1998.
18. J. O. Pickles. *An Introduction to the physiology of hearing*. Academic Press, London, 1982.
19. A. Popper and R. Fay, editors. *Hearing by Bats*. Springer Verlag, 1995.
20. Real World Interface, Inc. *B12 Base Manual (version 2.4)*, 1994.
21. H.-U. Schnitzler and J. Ostwald. *Adaptations for the detection of fluttering insects by echolocation in horseshoe bats*, volume 56, pages 801–827. NATO ASI Series, 1981.
22. M.R. Schroeder and J.L. Hall. Model for mechanical to neural transduction in the auditory receptor. *J. Acoust. Soc. Amer.*, 55:1055–1060, 1974.
23. M. Slaney. An efficient implementation of the Patterson-Holdsworth auditory filter bank. Technical Report 35, Apple Computer Inc., 1993.
24. B. Tian and H.-U. Schnitzler. Echolocation signals of the greater horseshoe bat (*Rhinolophus ferrumequinum*) in transfer flight and during landing. *J. Acoust. Soc. Am.*, 101(4):2347–64, 1997.
25. V. A. Walker. *One tone, two ears, three dimensions: An investigation of qualitative echolocation strategies in synthetic bats and real robots*. PhD thesis, University of Edinburgh, 1997.
26. V. A. Walker, H. Peremans, and J. C. T. Hallam. Good vibrations: Exploiting reflector motion to partition an acoustic environment. *Robotics and Autonomous Systems*, 24(1-2):43–55, August 1998.