Principles of Echolocation
in Bats and Machines

Note 4

Acoustic Flow

How echoes move in sensor space when the sensor moves

Jan Steckel
Cosys-Lab, University of Antwerp, and Flanders Make, Belgium jan.steckel@uantwerpen.be

Version 2026-11-18 16:01

1 The World Stands Still, the Echoes Do Not

A sonar does not see the world. It sees echoes, at a range r and a bearing θ measured in its own frame: the sensor space. When the sonar moves through a world that stands still, nothing in the world moves, but every echo moves in sensor space. A pole that was straight ahead drifts to the side, a wall that was far comes closer, and when the robot turns, the whole image swings the other way. This motion of the echoes is the acoustic flow, the acoustic counterpart of the optic flow that a moving eye sees.

The key point of this note is that this flow is not arbitrary. For a static world, it is completely fixed by two things: the motion of the sensor, and where the reflectors are. A sonar that knows its own motion can therefore predict where every echo will be at the next call. Conversely, every echo that does not go where the flow says carries information: about distance, about the sensor itself, or about something in the world that is not a static point.

2 The Flow Equations

We stay in two dimensions. The sensor drives forward with speed v and turns with turn rate ω, and both the bearing θ and ω are positive to the left (counterclockwise), as in our other notes. A static point reflector sits at range r and bearing θ, so at $(x, y) = (r\cos\theta,\ r\sin\theta)$ in the frame of the sensor (x forward, y to the left).

Seen from the sensor, the point moves. The forward motion of the sensor makes it come closer along x, and the rotation of the sensor makes it rotate the other way:

$$\dot{x} = -v + \omega y, \qquad \dot{y} = -\omega x. \tag{1}$$

From $r^2 = x^2 + y^2$ and $\tan\theta = y/x$, we get $r\dot{r} = x\dot{x} + y\dot{y}$ and $r^2\dot{\theta} = x\dot{y} - y\dot{x}$, and with the two equations above, the acoustic flow:

$$\boxed{\dot{r} = -v\cos\theta, \qquad \dot{\theta} = \frac{v\sin\theta}{r} - \omega.} \tag{2}$$

Two lines, but they say a lot. The range changes only through the forward motion, and fastest for points straight ahead. The bearing has two parts that behave completely differently: a translation part, $v\sin\theta/r$, which depends on the range, and a rotation part, $-\omega$, which does not.

3 Rotation Tells Nothing, Translation Tells Distance

Figure 1 shows the flow in sensor space for three motions.

Pure rotation (Fig. 1b). When the sensor turns on the spot, every echo shifts in bearing by the same amount, $-\omega$ per second, whatever its range, and no range changes. The whole image simply slides sideways. A turning sensor therefore learns nothing about distance, but its image is fully predictable once the turn rate is known, for example from a gyroscope.

Pure translation (Fig. 1a). When the sensor drives straight, the echoes stream outward from straight ahead, where the bearing does not change: the focus of expansion. Near points sweep fast, far points slowly, because the bearing rate scales with $1/r$. This is motion parallax. A pole 5 m ahead and 1 m to the side (at $r = 5.10$ m, $\theta = 11.3°$) drifts at only 2.2°/s when the sensor drives at 1 m/s, but at 19°/s once it is 1.4 m ahead. Driving straight also conserves one quantity:

$$\frac{d}{dt}\left(r \sin\theta\right) = \dot{r}\sin\theta + r\cos\theta\,\dot{\theta} = 0 \qquad (\omega = 0). \tag{3}$$

The product $r\sin\theta$ is the lateral distance between the point and the path of the sensor, so every echo slides along its own flow line, labeled by how close it will pass (magenta in Fig. 1a). Everything whose flow line lies within half the width of the robot is on a collision course, which the robot can read directly from its sensor space, without a map.

Driving and turning (Fig. 1c). The two parts add. On the left of a left turn, there is a point where they cancel, at $\theta = 90°$ and $r = v/\omega$: the center of the turn, around which the flow swirls.

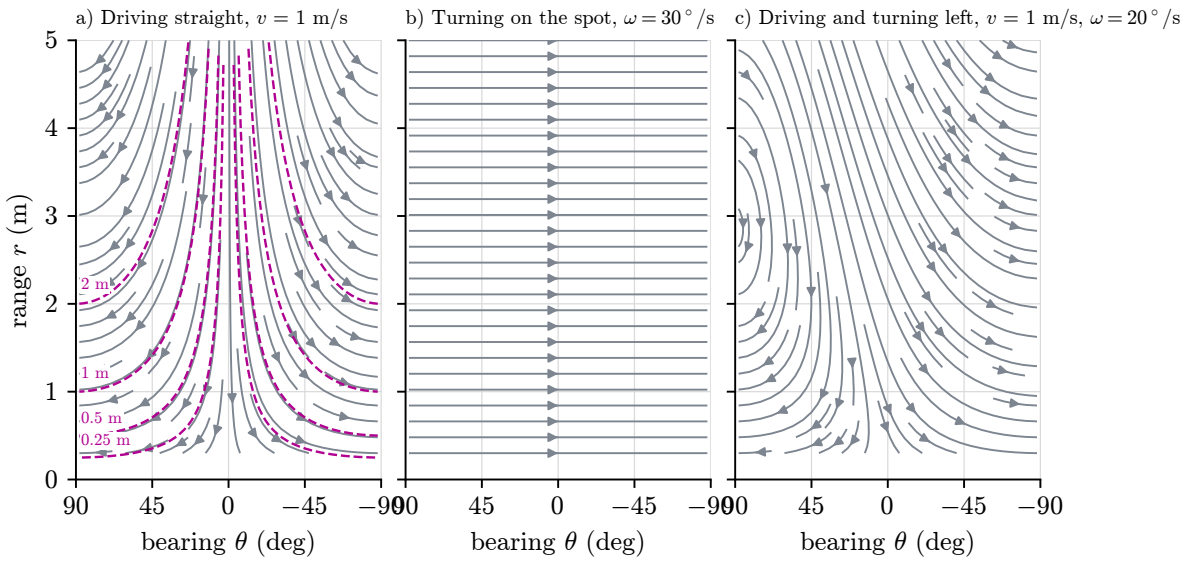


Figure 1: The acoustic flow in sensor space (bearing to the left on the left), equation (2). a) Driving straight: the echoes stream outward from straight ahead, along flow lines of constant lateral distance $r\sin\theta$ (magenta, dashed). b) Turning: all echoes shift equally in bearing. c) Both: the flow swirls around the center of the turn.

Key idea 1

In sensor space, rotation shifts every echo by the same angle and carries no information about distance. Translation moves near echoes faster than far ones (motion parallax), along flow lines of constant lateral distance.

4 Checking the Equations

Equation (2) is exact for a static point and an ideal sensor. Figure 2 checks it with a small simulation: a sensor drives a curved path at 0.8 m/s, turning at up to 18°/s, among nine static points. At every step, we compute the true range and bearing of every point (colored), and separately integrate the flow equations from the first measurement, using only v and ω (dashed). Over six seconds, the two agree to within 2.0 mm in range and 0.38° in bearing, the remaining difference being the time step of the simulation.

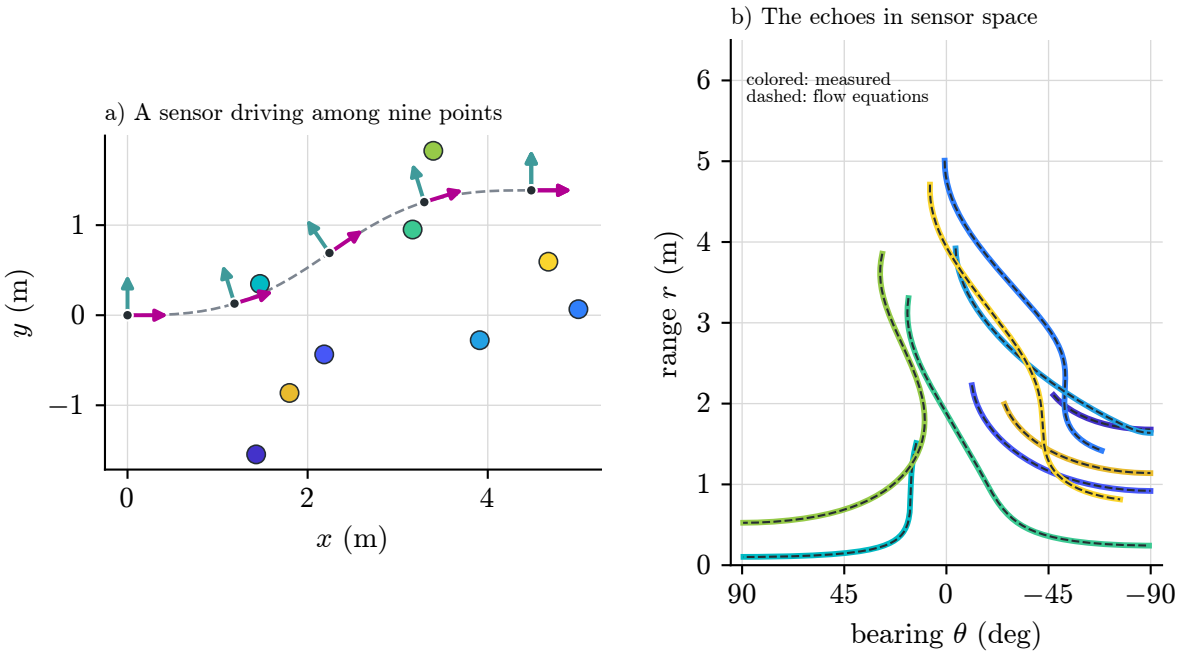


Figure 2: a) A sensor (triads: position, pointing direction in magenta, left in teal) drives a curved path among nine static points. b) The echoes of the points in sensor space while they are in the frontal field of view: measured (colored) and predicted by the flow equations from the sensor's own speed and turn rate (dashed).

5 A Sensor That Is Not at the Center of Rotation

So far, the sensor sat at the center of rotation of the robot. On a real robot, the sonar is usually mounted at the front, a distance d ahead of the axle. When the robot turns, the sensor then also swings sideways, with a speed ωd to the left, and the flow gains two terms:

$$\dot{r} = -v\cos\theta - \omega d\sin\theta, \qquad \dot{\theta} = \frac{v\sin\theta - \omega d\cos\theta}{r} - \omega. \tag{4}$$

The rotation now also produces a little translation, and with it range changes and parallax. For a point 1 m straight to the left, a robot that drives at 1 m/s while turning left at 30°/s sees the range shrink at −0.10 m/s with the sensor 20 cm ahead of the axle, against 0.00 m/s at the axle. We checked equation (4) against a numerical simulation of the same geometry, to six digits. For estimating motion (below), d is a known constant of the robot, so the equations stay linear in v and ω.

6 Time to Contact and Collisions

How long until the sensor reaches a point? Dividing the range by the speed at which it shrinks gives the time to contact,

$$\tau = \frac{r}{-\dot{r}} = \frac{r}{v\cos\theta}, \tag{5}$$

which for the pole above is 5.20 s. Lee and colleagues proposed that bats could steer by this quantity, read directly from the rate at which the delay of an echo shrinks, without knowing their speed or the distance separately.

Time to contact alone does not say whether the sensor will hit something. A point 5 m ahead and 1 m to the side has a time to contact of about 5 s as well, but it will be passed, not hit. The flow lines of Fig. 1a answer the second question: what will be hit is what lies in the collision corridor, $|r\sin\theta|$ smaller than half the width of the robot. Figure 3a shows both in sensor space. The corridor is a narrow band around straight ahead that widens at short range, and the lines of equal time to contact are arcs. Figure 3b shows how τ evolves while the sensor approaches three points. For a point straight ahead, it falls linearly with a slope of -1, reaching zero at impact, which is the quantity Lee proposed that animals control. For a point that will be passed, τ reaches a minimum and then rises again as the point moves to the side.

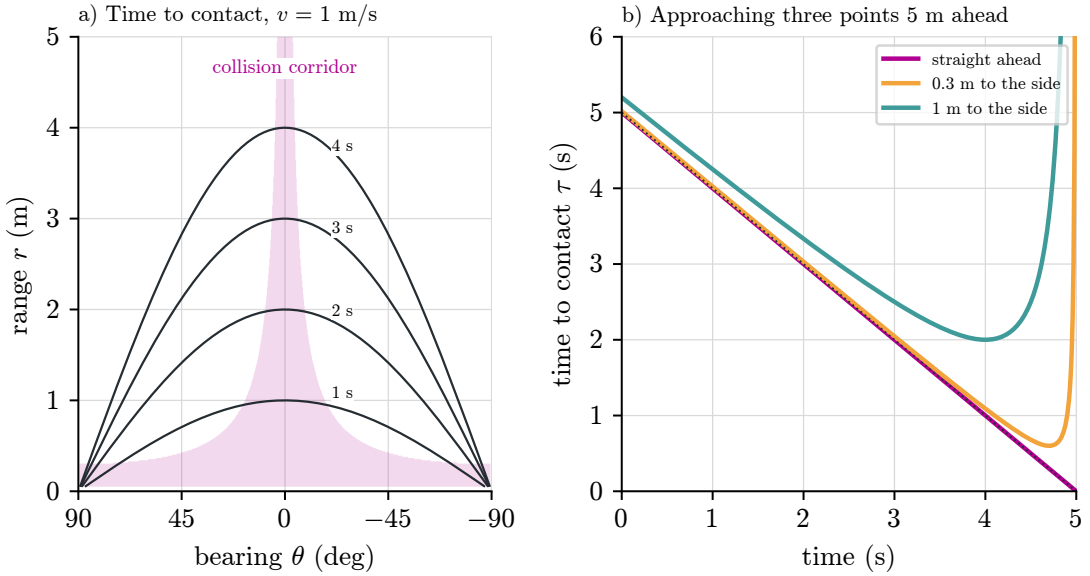


Figure 3: a) Lines of equal time to contact (black) and the collision corridor for a robot 60 cm wide (magenta) in sensor space, at 1 m/s. b) Time to contact while driving towards three points that start 5 m ahead, at 0, 0.3 and 1 m from the path. Dotted: the time left until the sensor reaches $x = 5$ m.

7 Walls Do Not Flow Like Points

The note *Points, Walls and Corners* showed that the echo of a smooth wall comes from the foot of the perpendicular, the point of the wall closest to the sensor. Its bearing is the direction of the wall's normal in the sensor frame, which changes only when the sensor turns. Its range is the distance to the wall, which changes like the range of a point at the same place. So the echo of a wall flows with

$$\dot{r}_{\text{wall}} = -v\cos\theta - \omega d\sin\theta, \qquad \dot{\theta}_{\text{wall}} = -\omega. \tag{6}$$

The range rate is that of a point, but the bearing rate lacks the parallax term $(v\sin\theta - \omega d\cos\theta)/r$. This difference is exactly what gives a wall away: for a wall 1 m to the left of a robot driving at

1 m/s, a point at the same place would sweep backward at 57°/s, while the wall echo does not move at all.

8 Estimating the Sensor's Own Motion

Equations (4) are linear in v and ω. Every echo that is followed from one call to the next gives two equations,

$$\begin{bmatrix} \dot{r}_i \\ \dot{\theta}_i \end{bmatrix} = \begin{bmatrix} -\cos\theta_i & -d\sin\theta_i \\ \sin\theta_i / r_i & -(1 + d\cos\theta_i / r_i) \end{bmatrix} \begin{bmatrix} v \\ \omega \end{bmatrix}, \tag{7}$$

so a handful of echoes of static points suffices to estimate the robot's own motion by least squares, without wheels or gyroscope. Figure 4 shows how well this works, for calls 0.1 s apart, a measurement noise of 3 mm in range and 0.5° in bearing per call, and points spread over the field of view between 0.5 and 5 m (2000 random trials per point count, $v = 0.8$ m/s, $\omega = 15°$/s, $d = 20$ cm). With three points, the speed is found to within 3.5 cm/s and the turn rate to within 3.6°/s (RMS), with twenty points to within 1.2 cm/s and 1.39°/s. The error falls roughly with the square root of the number of points.

A single wall echo among the points spoils this. Treated as a point, it attributes the missing parallax to a turn, and with three points, the turn rate comes out biased by 8.3°/s on average (1.7°/s with twenty points). The remedy follows from the previous section: the residual of a wall echo in the bearing equation is the full parallax term, 46°/s for a wall 1 m away at 0.8 m/s, far above the noise of the bearing rate (7°/s here). A robust estimator, which fits the motion to the echoes that agree and rejects the rest, therefore finds and removes the walls, and in passing classifies the echoes into points and non-points.

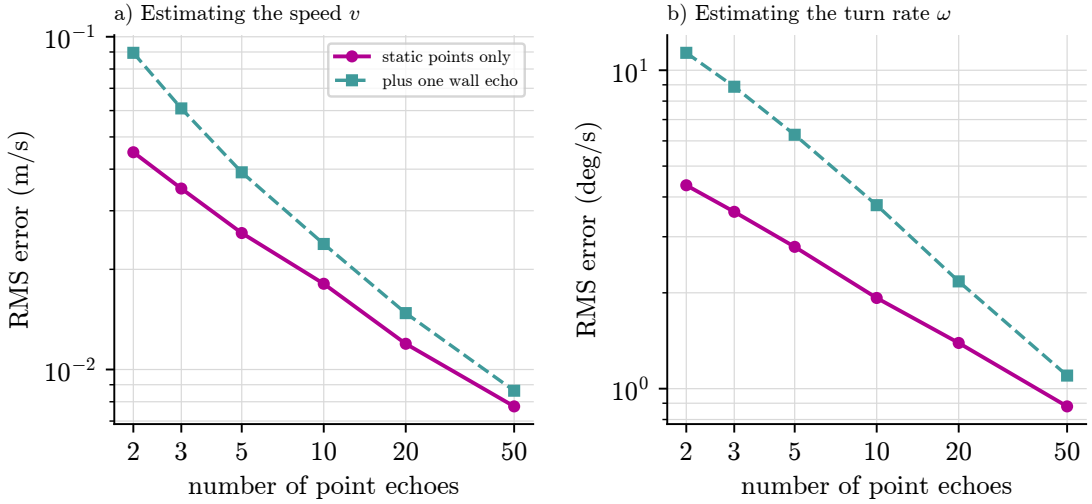


Figure 4: Estimating the sensor's own motion from the flow of n static points by least squares (calls 0.1 s apart, 3 mm and 0.5° noise per call): RMS error of a) the speed and b) the turn rate. Dashed: one wall echo added and treated as a point.

9 In Three Dimensions

In three dimensions, a sensor moves with a velocity $\boldsymbol{v}$ and rotates with an angular velocity $\boldsymbol{\Omega}$, both in its own frame. A static point at range r in the direction of the unit vector $\boldsymbol{u}$ then flows as

$$\dot{r} = -\boldsymbol{v} \cdot \boldsymbol{u}, \qquad \dot{\boldsymbol{u}} = -\frac{\boldsymbol{v} - (\boldsymbol{v} \cdot \boldsymbol{u})\,\boldsymbol{u}}{r} - \boldsymbol{\Omega} \times \boldsymbol{u}, \tag{8}$$

which reduces to equation (2) in the horizontal plane. The structure is the same: the translation part carries the range, the rotation part does not. What is new is the way rotations act. A yaw (pan) shifts all echoes in azimuth, a pitch (tilt) shifts them in elevation, and a roll rotates the whole image around the straight-ahead direction. For a sonar with a fixed directivity, these out-of-plane rotations are what makes echoes from above and below behave differently, which is why, in our self-supervised HRTF work, the elevation part of the directivity could only be learned when the head tilted or rolled.

10 What the Flow Is Good For

Separating self-motion from the world. A sonar that knows its own speed and turn rate (from wheel odometry and a gyroscope, or from the vestibular system of a bat) can predict where every echo of a static world will be at the next call. The other way around, every echo that does not follow the prediction stands out. It may be a moving object, such as a flying insect. It may also be a reflector that is not a point at all, such as a wall, whose echo lacks the parallax of a point. And it may be an artifact of the sensor, such as a sidelobe, which sticks to the reflector that causes it instead of following its own flow.

Distance from motion. A sonar measures range directly, but the flow gives a second, independent estimate: solving the bearing equation for the range gives $r = v \sin\theta/(\dot{\theta} + \omega)$. Two independent ways to the same number are the basis of self-calibration. In our self-supervised HRTF work, for example, a sonar learns its own directivity because only the right directivity makes the echoes flow as the equations predict.

Steering without a map. Animals that see use optic flow to steer: honeybees fly through the middle of a narrow tunnel by balancing the image speed on their left and right eyes. The acoustic version works for a sonar. In our robot controllers, simple behaviors each look at their own part of the flow (obstacles on a collision course, the flow line of a wall to follow, the opening of a junction), and together they guide a robot through corridors without a map. Bats seem to do something similar: in flight-tunnel experiments, the way bats flew depended on how echoic the walls were, even at the same distance, which suggests that echo-acoustic flow guides their flight.

Key idea 2

The flow of a static world is fixed by the sensor's own motion. Knowing that motion turns every call into a prediction of the next one: echoes that follow it reveal distance and structure, echoes that do not reveal moving objects, non-point reflectors and sensor artifacts.

Exercise

1. Show that for pure rotation, r is constant and $\theta(t) = \theta_0 - \omega t$. Why can a sonar that only turns not estimate distances from its echoes?
2. A robot drives straight at 0.5 m/s. An echo at $\theta = 30°$ drifts at $10°$/s. How far away is the reflector, and how close will the robot pass it?
3. Where in sensor space is the bearing flow zero when the sensor drives and turns at the same time? Sketch the curve for $v = 1$ m/s and $\omega = 20°$/s.
4. A robot drives straight towards a wall that makes an angle of $60°$ with its path. How do the range and the bearing of the wall echo change? Would a robust motion estimator still recognize it as a non-point?
5. Why does the error of the motion estimate in Fig. 4 fall roughly with the square root of the number of points?

Further Reading

- H. Peremans and J. Steckel, "Acoustic flow for robot motion control," IEEE ICRA, 316–321, 2014.
- D. N. Lee, J. A. Simmons, P. A. Saillant and F. Bouffard, "Steering by echolocation: a paradigm of ecological acoustics," J. Comp. Physiol. A, 176, 1995.
- K. Kugler, W. Greiter, H. Luksch, U. Firzlaff and L. Wiegrebe, "Echo-acoustic flow affects flight in bats," J. Exp. Biol., 219, 2016.
- M. V. Srinivasan, M. Lehrer, W. H. Kirchner and S. W. Zhang, "Range perception through apparent image speed in freely flying honeybees," Visual Neuroscience, 6, 519–535, 1991.