Reactive Maze Solving with a Biologically-Inspired Wind Sensor

Tim Chapman
Department of Psychology, University of Stirling, Scotland, FK9 4DN, UK
t.p.chapman@stirling.ac.uk

Adam T. Hayes
MS136-93, California Institute of Technology, Pasadena, CA 91125, USA
athayes@its.caltech.edu

Mark W. Tilden
MSA454, Los Alamos National Laboratory, Los Alamos, NM 87545, USA
mwtilden@lanl.gov

Abstract

We describe the design and testing of a novel biologically-inspired wind sensor, based on the cricket's cercal filiform hair cells. The sensor is mounted on a Khepera mobile robot, and competence at following a real wind plume is shown, even at low wind speeds. Further testing demonstrates the ability of the system to perform a dynamic maze solving task, and the exploitation of the hair sensors as tactile whiskers is discussed.

1 Introduction

Robot designs can be improved by considering biological mechanisms. There is evidence that a variety of organisms can detect wind direction, and use this information to orientate and locomote upwind (positive anemotaxis). Such behaviour is known in beetles [1], moths, flies, wasps [2] and crickets [4], as well as in mammals such as rodents and larger predators including cats, dogs and apes. The primary reason appears to be related to the detection of airborne chemical odours. Chemicals released into the atmosphere are carried along in the direction of the air flow, creating a plume. Locating the source of the plume will necessarily reveal the source of the odour, be it food or a predator [3].

To measure wind direction, an animal must either sense the passing mass of air directly, or some causal effect of it, such as a surface temperature cooling. Example mechanical equivalents of these techniques include the wind vane and hot-wire anemometer respectively. Equivalent biological structures are wind sensitive hair cells and temperature sensors in the skin. A particularly well studied biological system capable of determining wind direction is the cercal system found in insects such as the cricket and cockroach [4]. The cerci, two antennae-like appendages that protrude from the rear of the animal, are covered by hundreds of filiform hair cells. These hair cells, which vary in length from 0.1–2mm, have been shown to respond to both constant wind and wind puff stimuli. In addition, the response of each hair cell exhibits a characteristic wind directionality tuning curve, typically with a unidirectional preference. From the response of these hair cells the cricket is able to calculate the direction of approach of an incoming predator, such as the digger wasp, and execute a suitably oriented escape behaviour.

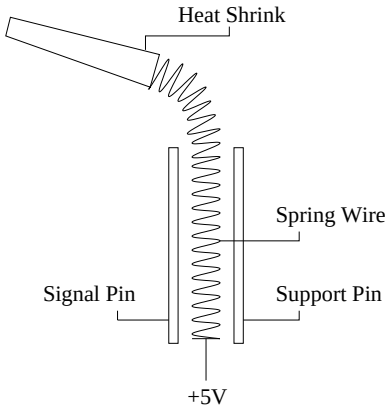

Figure 1: Diagrammatic representation of a single minimal wind sensor [5].

This paper describes the design of a neuromorphic wind-sensor inspired by the unidirectional filiform hairs of the cricket. The sensor is mounted on a Khepera mobile robot and its ability to track a wind plume under a range of conditions is demonstrated.

2 Transduction

The basic sensor device is a spring surrounded by four metal pins, as shown in Figure 1. The spring wire was 0.004ml stainless steel wound to 2mm diameter,

and the signal pins are gold-plated 0.7mm square Augat Wirewrap Headers. The spring was longer than the pins, and stretched and end-weighted so that it deflected in the presence of wind. The spring was connected to +5V, so that when it made contact with the signal pin, the pin was pulled high. The contacts were supposedly binary, but due to the bouncing knife-edge reaction of the spring against the pins, the output signals were far more stochastic than expected. The sensitivity of the sensor could be adjusted by changing the properties of the spring (length and stiffness) and adjusting the heat shrink counterweights so that the inherent stability of the sensors was always at threshold, providing a significant noise floor but also optimising sensitivity.

The sensor design utilized the signal from only one of the four available pins per spring. The potential information that is discarded is offset by an increase in the accuracy to which the sensitivity of the device can be tuned. The tuning task therefore changes from one of balancing the four non-independent pins, to minimising the separation of a single pin from the spring. The result is that, like the cricket filiform hairs, each device can be tuned for maximum sensitivity in a single direction.

Sensitivity to wind from other directions can be added easily by combining more of the basic hair devices oriented at the desired angles. In theory an arbitrary number of homogeneous or heterogeneous devices could be combined to increase the range of sensitivity of the combined sensor. The sensor used in these experiments consisted of four devices arranged at right angles to one another.

The sensor was placed on a Khepera robot equipped with a standard I/O turret, as shown in Figure 2. The outputs from the four sensors were connected to four digital inputs and read from the K-Bus using standard operating system calls. Power for the sensor was taken from the +5V and GND turret connections. The wind source used was a 12" diameter desktop fan.

3 Algorithms

3.1 Minimalist Algorithm

As in Figure 3, the hair sensors are numbered 0 to 3 clockwise from the front-left. The first algorithm that we tested was designed to set a baseline for the minimum processing required to generate the behaviour:

1. Read inputs from all hairs.

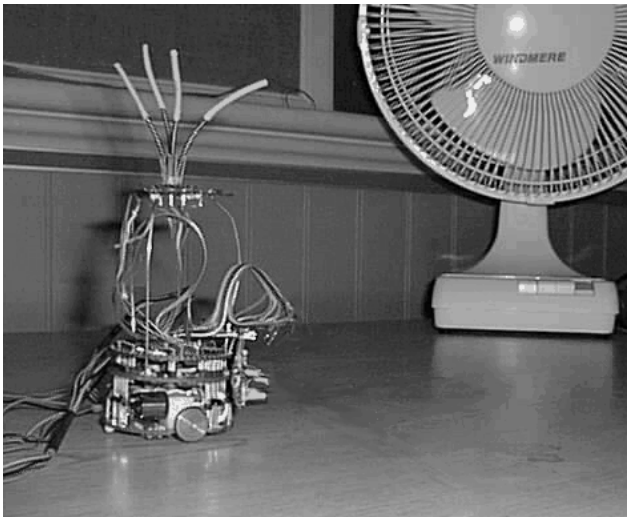

Figure 2: The combined wind sensor mounted on the Khepera miniature mobile robot. The fan used as the wind stimulus can be seen in the background. Videos of the behaviours described in this paper can be found at `ftp://ganglion.stir.ac.uk/pub/Floella-videos/`.

2. If Hair 3 (back-left) is on then turn right slowly.
3. Else if Hair 2 (back-right) is on then turn left slowly.
4. Else if Hair 1 (front-right) is on then turn left quickly.
5. Else if Hair 0 (front-left) is on then turn right quickly.
6. Else move forward slowly.
7. Go to step 1.

The algorithm was allowed to run at full speed on the Khepera's MC68331 processor. It proved to be successful at locating the wind source from a distance of 2 metres (wind speed ~1 m/s), despite the inherent asymmetry in the turning priority introduced by ordering the hair sensors with an If...else construct. The algorithm was also tested in reverse by inverting the speed and direction of turns for each of the sensors. This resulted, as expected, in the robot orienting and moving *away* from the fan. Competence level at this task was also good.

Qualitative analysis of the robot's behaviour suggested that the task competence involved using the two front sensors as 'error' signals and the two back

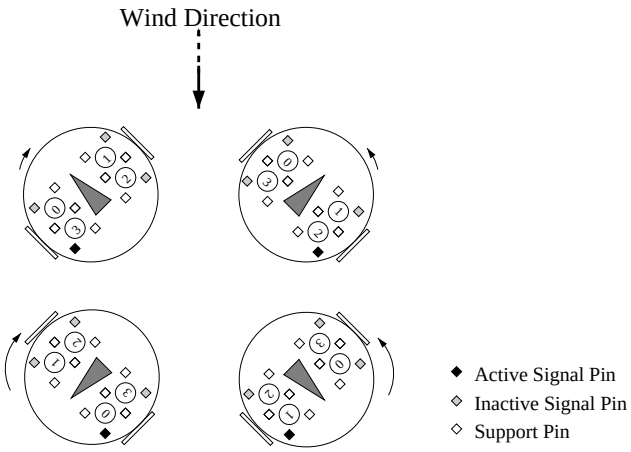

Figure 3: A minimal algorithm to perform the positive anemotaxis behaviour.

sensors for 'fine-tuning'. When the robot faced away from the wind source one or both of the front sensors would be stimulated the most, in either case causing the robot to turn rapidly until it was oriented roughly upwind. In this situation the front sensors became very unlikely to fire (the wind blowing the spring away from the active pin), instead, the back sensors were used to make fine adjustments to the robot's course.

Simple ablation studies conducted by placing an insulating material over the appropriate signal pins supported these hypotheses. As expected, removing the two front 'error' signals did not entirely prevent the robot from locating the wind source, but did increase the time taken to turn around when facing away from the wind source. Removing only a single front signal pin, however, did not significantly degrade performance as compared with the full sensor. This evidence strongly supports the notion that the informational content of the front pins is equivalent to a 'severe directional error' signal.

Removing the two back 'fine-tuning' signals had a more severe impact on performance. Without these the robot seemed to lose its ability to 'lock on' to the target, resulting in a circuitous path with many loops. Removing only a single back hair improved the situation somewhat, but performance was still poor.

3.2 Temporal Binning Algorithm

Two main problems with the minimal algorithm became apparent. Firstly, a directional bias, due to the asymmetry of algorithm. Secondly, a tendency to 'jitter' left and right, because the algorithm was executing (and thus changing the motor speeds) at many thousands of cycles per second. To correct these problems we proposed a new 'temporal binning' algorithm, as follows:

1. Create a bin counter for each hair.
2. Zero all counters.
3. Read inputs from all hairs.
4. For each hair that is on, increment the respective bin counter
5. Repeat steps 3 and 4 for a fixed number of binning cycles
6. If the combined total of the front hairs is greater than the back hairs then
   (a) Turn quickly and proportionally to the totals of the front hairs.
   (b) Else turn slowly and proportionally to the totals of the back hairs.
7. Go to step 2.

As before, the algorithm was executed at full speed on the Khepera. The bin size parameter could be adjusted from a single cycle (equivalent to the minimalist algorithm) up to many millions of cycles (lasting several minutes).

For small bin sizes the model behaved similarly to the minimalist model described above, with the added enhancement of correcting the asymmetry problem. Increasing the bin size also corrected the 'jitter' problem, with bin sizes of around 100 to 1,000 cycles resulting in the robot taking longer and smoother turns, and the ability to track the wind plume from over 4 metres (wind speed ˜.6 m/s). Bin sizes up to 10,000 cycles produced a noticeably curved path, but did not greatly reduce the ability of the robot to locate the wind source under moderate conditions. Even with bin sizes of 100,000 cycles, which allowed the robot sufficient time to perform a complete 360° turn between motor updates, the wind source could still be located, although by now the performance was much impaired. Nevertheless, these tests demonstrate the robustness of the temporal binning algorithm across bin sizes spanning six orders of magnitude and the advantages gained by integrating sensory input over an extended period on time.

3.3 Other Factors

Ego-motion Before testing the system, we expected the motion of the Khepera robot to affect the hair sensors far more than the wind itself. Surprisingly however, even though the springs were greatly affected by the movement of the robot, the wind direction could still be determined properly. This was due to integration of the signal over a long time period, relative to the vibration rate of the hairs.

Biases Perfectly balancing the sensitivity levels of the four hair sensors proved to be almost impossible, and certainly would not have been robust had we achieved it. The result of this was the introduction of a small bias in the turning preference. In the presence of a wind stimulus the effects of this bias were insignificant, however, with no wind present the result was to cause the robot to travel in a circle. Rather than being a hindrance, however, this behaviour assisted the robot if it moved out of the wind stream. For example, with a clockwise bias, if the robot left the wind plume on the left-hand side the clockwise circle naturally returned it towards the fan. If, conversely, the robot left the wind plume on the right, the bias caused it to loop back until it returned to the stimulus further downwind. This is reminiscent of behaviours observed in real animals, such as ants following a chemical trail.

4 Maze Solving

Much of the behaviour of a real world situated agent is derived from its ability to exploit the physical characteristics of its environment through the agent-environment feedback loop [6]. Exploitation of the turning bias as a way to relocate the wind stream is one example of this. Location of a wind source through the dynamic interaction of the air flow, hair sensors and mobile base is another. We were interested in testing these properties in a more complex and challenging situation, as well as developing a more controlled test environment. Maze solving presented itself as the ideal solution.

No internal representation of the maze was used by the robot. Instead the maze was solved directly in the environment by exploiting the natural properties of air flow. This approach conferred two major benefits to the robotic system. Firstly, it simplified the control structure from a full-blown generic maze navigation algorithm to a sub-representational wind follower. Secondly, the nature of the problem solving was

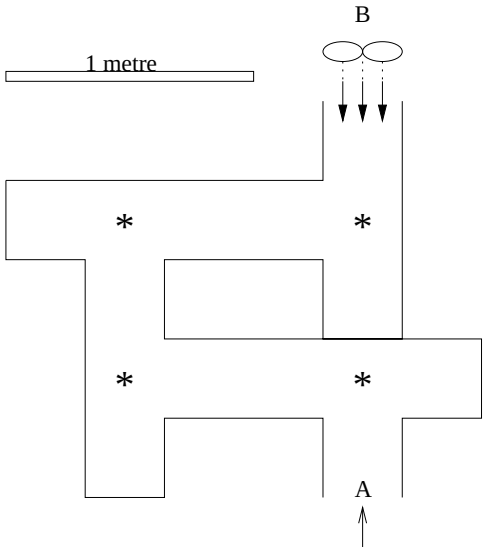

Figure 4: Layout of the 4 component T-Maze. A - entrance, B - exit, * - choice point.

inherently dynamic; the robot necessarily reacts in a sensible manner to any changes that occur in the environment, such as moved obstacles or new wind sources. Using physical sensors as matched filters for a particular stimulus in this way solves many of the difficulties associated with traditional navigation problems.

The maze design chosen for these experiments was based on a classic "T-maze" design, more commonly used for testing forced choice decisions in rats. The maze, shown in Figure 4, was based around four T components, thus providing four choice points. At each of these junctions the robot had to "make a decision" about which way to turn. This method provides an explicit mapping of the internal dynamics of the system into a behaviour that can be observed clearly and unambiguously. Correct decisions were defined as those choice points at which, on its first encounter with the junction, the robot departed along the corridor nearest to the maze exit. Departure in any other direction was counted as an incorrect decision. This provided us with a convenient 5 point error scale on which to rate the successfulness of a given trial. Time taken to complete (or fail, by leaving through the entrance) the maze was also recorded as a second measure of task performance.

The robot executed the temporal binning algorithm described in Section 3.2, with the bin size parameter set to 100 cycles. In addition to the previous algorithm, the Khepera's built-in IR obstacle proximity sensors were used to provide a simple mechanism for preventing the robot from getting stuck on the walls of

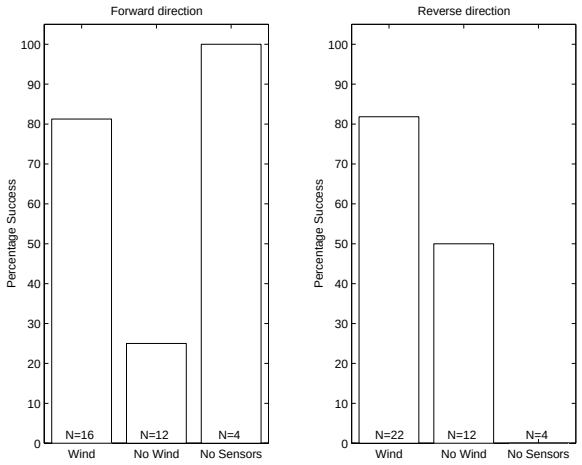

Figure 5: Percentage success rates for forward and reverse maze directions.

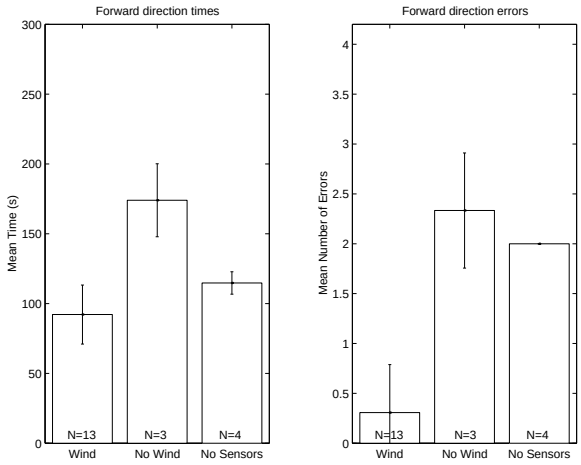

Figure 6: Mean times and errors taken to solve the maze in the forward direction. Error bars show standard deviations.

the maze. The avoidance algorithm was: if the combined output from the IR sensors is greater than a threshold value then turn away from the side with the larger activation value.

Three different wind conditions were tested: 'wind', 'no wind' and 'no sensors'. In the 'wind' condition the sensors were active and the fan was positioned at the exit to the maze and set on either 'high' or 'low' (50% of the trials each). In the 'no wind' condition the wind stimulus was switched off in all trials, although the sensors were still active. In the 'no sensors' condition both the fan and the hair sensors were inactive. Two maze direction conditions were also tested. The first was as shown in Figure 4, the second was identical except that the entrance and the exit of the maze (points A and B) were reversed. This condition introduced an interesting challenge for the robot because the appropriate moment to turn is no longer forced. Instead the decision must be made spontaneously whilst navigating along the corridor.

5 Results

For each condition, time taken to reach the exit, and the number of errors made during the journey were recorded. Trials where the robot failed to solve the maze (by exiting through the entrance) scored maximum errors and were removed from the timing data.

Forward Direction The robot solved the maze on more than 80% of trials in the forward direction (Figure 5). The timing data (Figure 6) shows that the 'wind' group performed the best, having a clear advantage over the 'no wind' group. The 'no sensors' group timings were anomalous: the times were both relatively fast and remarkably consistent, suggesting a high level of performance. However, inspection of the error graph shows that the performance was at chance level. Closer inspection of these trials showed that in every case the robot followed *exactly* the same path, making exactly the same errors, a result confirmed by the 100% success rate (Figure 5). The results can therefore be attributed to the high quality of the Khepera's engineering combined with the fact that a T-Maze acts as funnel towards its exit.

The error data supports the timing data. The mean number of errors for the 'wind' group was close to zero, indicating reliable, error-free performance. Comparatively, the 'no wind' and 'no sensors' groups were performing at around chance level.

Reverse Direction In this more difficult condition the 'wind' group performed only slightly worse. The success rate was still over 80% but more errors were made (Figures 5 and 7). The 'no wind' group, however, performed surprisingly well considering that the maze no longer funnels the robot towards the exit. Figure 5 shows that the 'no wind' group actually succeeded more often in the reverse direction. The 'no sensors' group failed on every trial (by travelling to the first dead-end, performing a 180° turn to avoid the wall, and proceeding back out of the maze) so no data were available for the time analysis and maxi-

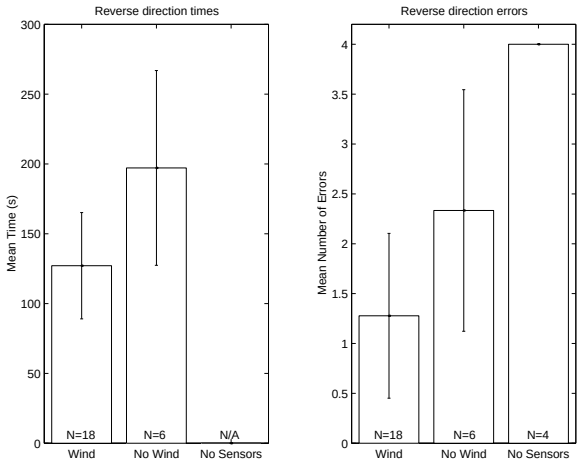

Figure 7: Mean times and errors taken to solve the maze in the reverse direction. Error bars show standard deviations. Note that the 'no sensors' group failed on every trial.

mum errors were scored in every trial.

We were interested (and concerned) as to how the 'no wind' group was performing so well in the reverse maze condition *without* a wind stimulus. Perhaps there was a latent flow of air through the maze? Close inspection, however, revealed that this was not the case. Rather, the robot was using a combination of hair sensor bias and mechanoreception. In short, the initial and final turns (both right-hand) were made due to the clockwise turning bias described in Section 3.3. The second and third left-hand turns were navigated by the left hair sensors (0 and 3 in Figure 3) brushing along the inner wall of the maze. The effect of this was to support the hairs slightly, keeping them from making contact with their respective signal pins and, therefore, the robot from making a right-hand turn. Without this balancing force, the movement of the right hair sensors induced from ego-motion alone was enough to produce regular left turns. The result was that the robot tracked its way efficiently along the left-hand wall and around the two corners, using its wind sensors as whiskers.

6 Conclusions

We have described the design and preliminary testing of a biologically-inspired wind sensor. The sensor was mounted on a mobile robot and we demonstrated its ability to track a wind plume using only a minimal control algorithm. An improved algorithm employing temporal binning techniques to allow better extraction of the signal from ego-motion induced noise, by integration of sensory input over time, enhanced the range and performance of the system.

We demonstrated the utility and robustness of the system by applying it to a maze solving task. The robot performed well above chance level under both the easier (forward) and more difficult (reverse) conditions. Interesting results were found when the wind stimulus was removed, as the robot was still able to solve the maze under the reverse condition by exploiting its wind sensors as impromptu tactile whiskers and feeling its way along the walls. Further investigation of this phenomena is in progress, along with a more precise characterisation of the sensor properties.

The sensor has applications in a range of wind sensing tasks. Its digital output simplifies interfacing issues, the power usage is low, and the sensor itself is physically robust. Planned uses of the sensor include the location of unexploded ordinance by chemical plume following and as a tool for robotic modelling of biology.

Acknowledgements

This work was supported by and carried out at the 1999 Telluride Workshop on Neuromorphic Engineering.

References

[1] H. Bohm, "Dynamic properties of orientation to turbulent air current by walking carrion beetles", *J. Expt. Biol.,* Vol. 198, pp. 1995-2005, 1995.

[2] R. T. Carde, "Odour plumes and odour-mediated flight in insects", *Ciba Found. Symp.*, Vol. 200, pp. 54-70, 1996.

[3] J. Murlis, J. S. Elkington and R. T. Carde, "Odor plumes and how insects use them", *Annl. Rev. Ent.*, Vol. 37, pp. 505-532, 1992.

[4] W. Gnatzy and R. Hustert, "Mechanoreceptors in Behaviour". In *Cricket Behaviour and Neurobiology*, Cornell: London, 1989.

[5] M.W. Tilden, "BEAM Robotics", *Algorithm 2.2 magazine*, pp 15-19, 1991.

[6] R. Brooks and L. Steen "Building brains for bodies", *MIT AI Lab memo 1439*, 1993.