

Robotic Ecology: Tracking Small Dynamic Animals with an Autonomous Aerial Vehicle

Oliver M. Cliff,^{1*} Debbie Saunders,² Robert Fitch^{1,3}

¹Australian Centre for Field Robotics, University of Sydney

²Fenner School, Australian National University

³Centre for Autonomous Systems, University of Technology, Sydney.

*To whom correspondence should be addressed; E-mail: o.cliff@acfr.usyd.edu.au..

Understanding animal movements that underpin ecosystem processes is fundamental to ecology. Recent advances in animal tags have increased the ability to remotely locate larger species, however, this technology is not suitable for up to 70% of the world's bird and mammal species. The most widespread technique for tracking small animals is to manually locate low-power radio transmitters from the ground using hand-held equipment. Despite this labour-intensive technique being used for decades, efforts to reduce or automate this process have had limited success. Here we present a new approach for tracking small radio-tagged animals using an autonomous and lightweight aerial robot. We present a series of experimental results where we use the robot to locate migratory swift parrots (*Lathamus discolor*), a critically endangered species. The robot system combines a miniaturised sensor with novel estimation algorithms to yield unambiguous bearing- and range-based measurements with associated measures of uncertainty. We incorporate these measurements into Bayesian data fusion and information-based planning algorithms to control the position of the robot as it collects data. We report estimated positions that lie within 50 m of the true positions of the birds, which are sufficiently accurate for recapture or observation. Further, in a head-to-head race with experienced human trackers, we found our system to perform comparably, and

often faster. These results provide the first unequivocal validation of robotic systems for wildlife radio telemetry since its inception in the 1960s, and pave the way for their widespread use as human-assistive or autonomous devices.

Introduction

Conservation management of certain critically endangered species relies on understanding how these species interact with their environment. This is achieved by tagging and tracking individual animals in the wild (1–3). Aerial robot systems can access rugged areas that are difficult for humans to traverse, and thus are viewed as a potentially revolutionary tool for data collection in wildlife ecology (4, 5). However, this potential remains largely unrealised. Robot systems have yet to achieve a level of signal detection, tracking accuracy, and speed that is sufficient to legitimise their role as a replacement for human trackers.

Despite recent advances in automated wildlife telemetry tracking, very little is known about the movement of small, dynamic migratory species, of which many have reached critically endangered status. For large to medium animals, the miniaturisation of GPS tags with remote data readout has facilitated a dramatic increase in understanding the movements of a diversity of species (6, 7). Methods such as satellite telemetry have far reaching applications from investigating migration routes and wintering areas of large migratory birds (8–10) to studying the dynamics of aquatic predators (11, 12). Unfortunately, these approaches are still only suitable for about 70% of bird species and 65% of mammal species (2). In the case of smaller species that return to the same breeding areas seasonally, miniature non-transmitting data loggers can be used (2); however, retrieving this data requires relocating the animals *in situ*. Due to this challenge, VHF tracking has become one of the most useful techniques in ecology and management (13). This involves attaching small radio transmitters to animals and subsequently tracking the target species. Although scientists have been using VHF tracking since the early 1960s (14), data yielded by this approach is sparse due to the manual labour involved (2). Thus, researchers are more frequently exploiting the abundance of low-cost unmanned aerial vehicles (UAVs) equipped with visual sensors for

this type of conservation management and wildlife monitoring (4, 5). However, the capability of these systems is limited in terms of identifying individual animals and locating animals in unknown locations.

In recent years, there has been increased interest in end-to-end wildlife telemetry tracking with robotic systems (4), where the robot moves autonomously to track a live target animal. The usefulness of these systems, however, is yet to be proven in direct performance comparison to the traditional manual approach. Most notably, ongoing research is aimed at tracking radio-tagged carp in Minnesotan lakes using autonomous surface vehicles (ASVs) on the water and mobile ground vehicles (MGVs) when the lake is frozen (15–20). While this project has yielded seminal work in the field, the use of ground and surface vehicles is untenable for wildlife situated in rugged habitats. We recently validated the use of a multirotor UAV for autonomously localising live radio-tagged birds in such environments (the *Manorina Melanocephala*) (21). Here, we present the first unequivocal validation that robotic systems can perform comparably to experienced human trackers in head-to-head experiments.

The majority of research in radio tracking with an aerial vehicle focuses on isolated subsystems. Although these systems are typically motivated by the idea of tracking small animals (e.g., bird (22–25) and fish species (26, 27)), only simulations or prototypes are presented with limited field testing. Alternatively, when tracking a relatively stationary target, the observations can be considered more robust and thus attention in this field has shifted to optimising planning for single (17, 19, 20) or multi-robot systems (18, 28). The main assumption the authors make is that the sequential observations are *homoscedastic*, meaning that the uncertainty over each measurement is constant or bounded. However, with a sporadic and unpredictable live target, this assumption is violated due to the resulting wide spectrum of observation quality from noisy to precise. As we show in this paper, this induces *heteroscedastic* observations, where the uncertainty varies with every observation. Failing to distinguish between low and high quality observations can lead to overconfident measurements that cause spurious location estimates, or to highly uncertain location estimates that are of little value.

A mathematically valid observation model is also critical in planning the motion of the robot to



Figure 1: The aerial robot system is designed to track small animals with lightweight radio tags attached, e.g., swift parrots (*Lathamus discolor*) 1a, brush-tailed rock-wallabies (*Petrogale penicillata*) 1b, and noisy miners (*Manorina melanocephala*) 1c. This work demonstrates that the robot is able to track swift parrots and yield comparable performance to an expert human operator performing traditional wildlife telemetry tracking 1d. The multirotor platform 1e-1f includes a lightweight directional antenna system and payload that receives the signal strength from the tag. This data is then transmitted to a ground control station for processing and online decision making.

improve the location estimate. In robotics this general problem is known as *active perception* (29,30) and introduces a coupling between data collection and planning. The idea of passively locating transmitting radio sources has been investigated in operations research motivated by search and rescue missions where stationary distress beacons must be recovered rapidly. Hence, the task is a coverage problem solved via offline strategies with an emphasis on minimising path cost over the entire area or teleoperated by humans (31). Alternatively, when the wildlife habitat is known and bounded, sensor networks can be placed in order to precisely track an animal’s location (32, 33). In our case, we require fast, precise estimates without intervention and thus employ active strategies where the observation quality relies crucially on an appropriate sequence of viewpoints (34). Our objective is reduce uncertainty (entropy) of the target location; thus, the task of actively tracking targets falls under the *informative path planning* paradigm (35). This problem is known to be NP-hard (36) and has been studied extensively over the last decade (37) with many applications focusing specifically on UAVs (38–43). In this paper we leverage these results to obtain an approximately optimal sequence of actions by greedily selecting the most informative viewpoints at each decision step.

In this work, we present rigorous analysis and validation of a complete system for autonomous wildlife tracking. We show that this system addresses the associated theoretical and engineering challenges to a degree that is sufficient to match or surpass the performance of skilled human trackers. First, we provide mathematical derivation for our data-driven sensor model, which has previously been validated over a number of trials on real birds and stationary targets (21). This range-azimuth model is further used to predict the quality of future viewpoints in planning an approximately optimal sequence of observations. We then directly compare this system against human operators in the problem of tracking the critically endangered swift parrot (*Lathamus discolor*) species in the wild. Across eight field trials, the estimated bird locations are precise to within 50 m, which is sufficient for recapture, detailed field observation, or data readout. Moreover, the time taken to achieve these estimates is comparable to, and often faster than, experienced human trackers. This result is significant because it is the first time in

over 50 years of wildlife telemetry tracking research that a robotic system has been validated as an autonomous or human-assistive device. This milestone paves the way for the widespread use of robots in migration ecology and conservation management for small, dynamic species.

Results

In this section we report on the approaches used for sensor modelling, data fusion and decision making. Then we show, in head-to-head competition, that the autonomous system’s performance is comparable to experienced human trackers. Finally, we discuss insights into the global and local spatiotemporal movements of swift parrots that were gathered from trials.

The material in this section is based on the bearing-only, heuristic approach presented in previous work (21). Here, we present a full range-azimuth algorithm derived rigorously from first principles. We also provide variance analysis and proof that the objective function is *monotone submodular*, an important property that is useful in designing efficient planning algorithms.

Likelihood Functions for Observations

The most critical component of the system is the sensor model, which allows us to convert the signal received from the radio tag to an instantaneous estimate of the target’s location. An inaccurate or over-confident observation can lead to poor decision making and imprecise final location estimates.

Let $\mathbf{X} \in \mathcal{S}$ denote the vehicle location, $\mathbf{Y} \in \mathcal{S}$ the targets’ geographic coordinates within some workspace $\mathcal{S} \subset \mathbb{R}^2$, and \mathbf{Z} an observation in some measurement space \mathcal{H} . We are interested in learning the *likelihood function* $\ell(\mathbf{y}; \mathbf{x}, \mathbf{z})$, i.e., the probability of receiving the measurement \mathbf{z} at location \mathbf{x} , given the target location \mathbf{y} :

$$\ell(\mathbf{y}; \mathbf{x}, \mathbf{z}) = \Pr(\mathbf{X} = \mathbf{x}, \mathbf{Z} = \mathbf{z} \mid \mathbf{Y} = \mathbf{y}) \text{ for } \mathbf{y} \in \mathcal{S}. \quad (1)$$

Note that we could consider uncertainty in the vehicle location \mathbf{x} by including it in the measurement \mathbf{z} , however, we assume full knowledge of vehicle state in this paper. We use the convention that upper

case symbols (e.g., \mathbf{Y} and \mathbf{Z}) are random variables, lower case symbols (e.g., \mathbf{y} and \mathbf{z}) denote their realisations, and bold denotes a set. Further, true (or optimal) quantities are denoted with an asterisk (e.g., \mathbf{y}^* is the true location of the bird) and estimates are denoted with a hat (e.g., $\hat{\mathbf{y}}$ is the target estimate).

To construct our sensor models, we must determine what we are measuring and the uncertainty over these measurements. In this work we take both range and azimuth readings of the target, where both observations are assumed to be normally distributed. This results in each measurement comprising the mean and variance $\mathbf{z} = \{\mu, \sigma^2\}$. Given a *measurement function* $h : (\mathcal{S} \times \mathcal{S}) \rightarrow \mathcal{H}$ that maps the vehicle \mathbf{x} and target state \mathbf{y} to the measurement space \mathcal{H} , the Gaussian likelihood function is:

$$\ell(\mathbf{y}; \mathbf{x}, \mathbf{z}) = f(h(\mathbf{x}, \mathbf{y}); \mu, \sigma^2), \quad (2)$$

where f is the probability density function (PDF) of the normal distribution.

Observed and Expected Sensor Data

To derive our likelihood functions we first describe the raw sensor data collected and the model used for obtaining our measurement. More details on the system collecting these measurements are given in the Materials and Methods section.

The radio tag emits an on-off keyed pulse signal; this transmission is received by the payload on-board the UAV and the received signal strength indicator (RSSI) values of the signal are captured and filtered. These RSSI values are linearly related to the power received during a transmission and are used as the raw sensor data for the observation. The range and bearing likelihood functions are based on these raw values and the learned sensor model.

The antenna used on-board the UAV is a two-point phased array: a lightweight, unambiguous directional antenna designed for radio telemetry with multirotor vehicles (21). The array comprises two monopoles fed through a passive combiner circuit, which yields a radiation pattern with a front lobe and back null. In order to reduce noise and spurious readings due to multipath propagation, the UAV

remains stationary while yawing through a full rotation. During this rotation, the continuous RSSI values are filtered and sampled at a constant rate to give a scalar value g_n associated with the bearing of the n th value ϕ_n . These values are then transmitted to a base station, giving the recorded gain pattern $\mathbf{g} = (g_1, g_2, \dots, g_N)$. As a result, the random vector $\mathbf{G} = \mathbf{g}$ is a function of the vehicle \mathbf{X} and target \mathbf{Y} location. Further, let $b(\mathbf{x}, \mathbf{y})$ denote the bearing from \mathbf{x} to \mathbf{y} . The true bearing to the target from robot location \mathbf{x} is then $\theta^* = b(\mathbf{x}, \mathbf{y}^*)$. We assume that the error for each recorded RSSI value is normally distributed with unknown variance $\sigma^2(\theta^*)$ that remains constant throughout an observation, i.e., for arbitrary $g_n \in \mathbf{g}$

$$g_n = \mathbf{E}[G_n | \Theta = \theta^*] + \nu_G, \quad \nu_G \sim \mathcal{N}(0, \sigma^2(\theta^*)), \quad (3)$$

where $\sigma^2(\theta^*) = \mathbf{V}(G_n | \Theta = \theta^*)$.

We obtain the expected gain pattern $\mathbf{E}[\mathbf{G} | \Theta]$ by linear regression. Specifically, we fit the expected gain pattern to a J th-order Fourier series $\varphi : \mathbb{R} \rightarrow \mathbb{R}$, i.e., given the true bearing θ^* ,

$$\begin{aligned} \mathbf{E}[G_n | \Theta = \theta^*] &= \varphi(\phi_n + \theta^*) \\ &= a_0 + \sum_{j=1}^J a_j \cos(j(\phi_n + \theta^*)) + \sum_{j=1}^J b_j \sin(j(\phi_n + \theta^*)). \end{aligned} \quad (4)$$

From this Fourier model, we obtain the expected gain pattern $\varphi(\theta) = \mathbf{E}[\mathbf{G} | \Theta = \theta]$, where $\varphi : \mathbb{R} \rightarrow \mathbb{R}^N$ is generated by sampling the Fourier series (Eq. (4)) with a phase offset θ at N regular intervals, i.e., $\varphi(\theta) = (\varphi(\theta), \varphi(\theta + 2\pi/N), \dots, \varphi(\theta + 2\pi))$.

Given the expected and observed sensor output, φ and \mathbf{g} , the main goal of Bayesian sensor data fusion is to compute PDFs of the bearing and range to a target from the robot. Given that the likelihoods are assumed to be Gaussian, the measurement tuple $\mathbf{Z} = \{\mu(\mathbf{G}), \sigma^2(\mathbf{G})\}$. To learn the mapping from \mathbf{G} to \mathbf{Z} , we use a data-driven approach based on training experiments, described below.

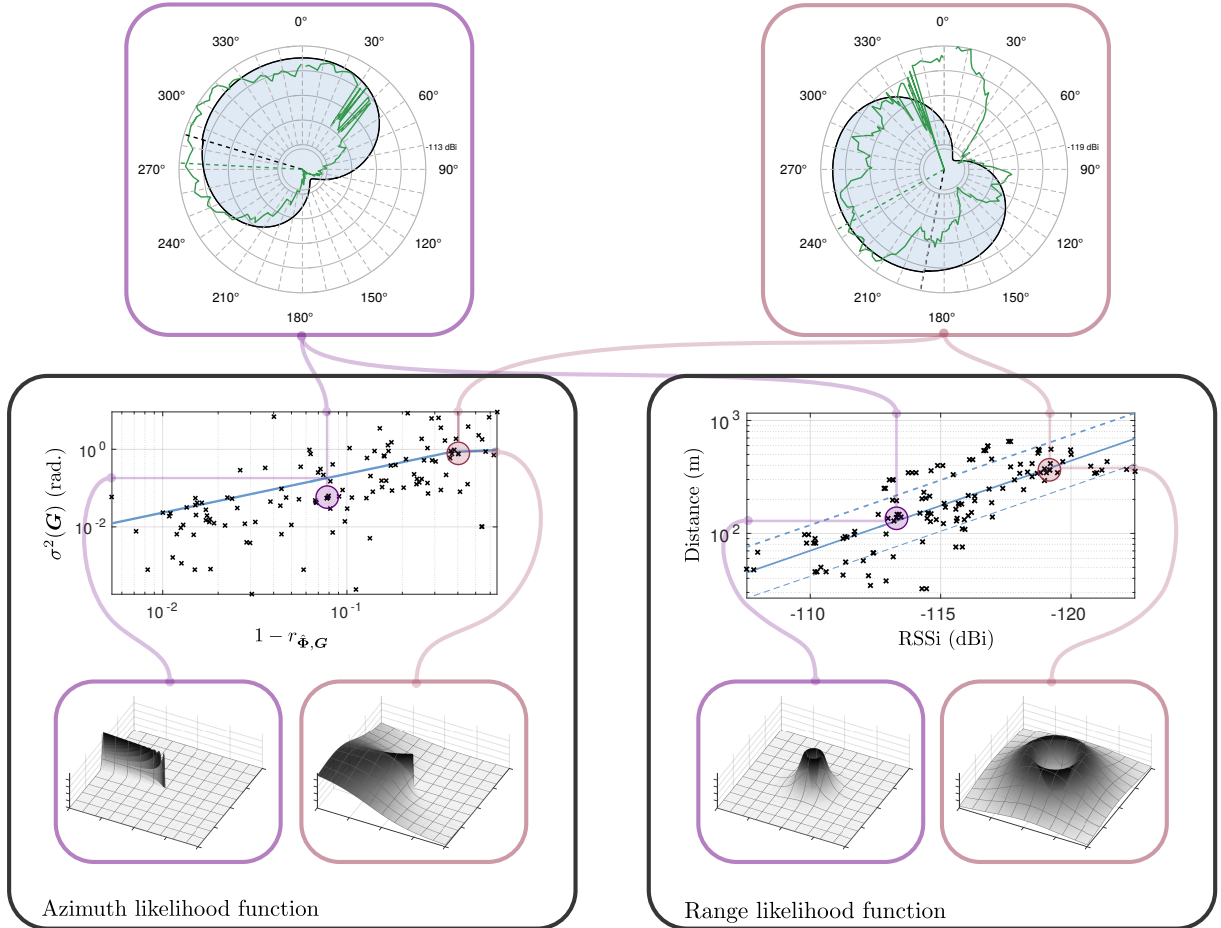


Figure 2: Obtaining range-azimuth likelihood functions from observations. The top row illustrates two example observations taken online with a stationary target. The radial plots illustrate real RSSI readings (green line) \mathbf{g} and a third-order Fourier series model $\varphi(\theta)$ of the radiation pattern (black line). The model is offset (rotated) such that it is oriented towards the true bearing to the target θ^* , and the RSSI values are offset by the maximum correlation $\mu_{\Theta}(\mathbf{g}) = \arg \max_{\theta} r_{\varphi(\theta), \mathbf{g}}$. These offsets are illustrated with dotted green and black radial lines. On the left subfigure, the maximum value correlation coefficient $r_{\varphi, \mathbf{g}}$ maps to a bearing-error $\sigma_{\Theta}^2(\mathbf{g})$, which is illustrated in the grid plots below. On the right subfigure, the maximum RSSI value g_{\max} maps to an expected range $\mu_R(\mathbf{g})$ with a fixed range-error $\hat{\sigma}_R^2$ giving the associated grid plots below.

Azimuth Likelihood Function

We model the likelihood of each azimuth measurement with a Gaussian bearing-error model (44) where $\mathbf{Z}_\Theta = \{\mu_\Theta(\mathbf{G}), \sigma_\Theta^2(\mathbf{G})\}$. That is, the difference between the true bearing to the target θ^* and the estimated bearing $\hat{\theta}$ (i.e., the bearing-error) is Gaussian distributed. Importantly, the bearing estimate $\hat{\theta} = \mu_\Theta(\mathbf{g})$ and its variance $\sigma_\Theta^2(\mathbf{g})$ are not measured directly but instead are given as functions of observation quality (i.e., the correlation coefficient, discussed below). As a result, when $\mathbf{G} = \mathbf{g}$, the bearing-error likelihood function ℓ_Θ is given by

$$\ell_\Theta(\mathbf{y}; \mathbf{x}, z_\Theta) = f(b(\mathbf{x}, \mathbf{y}); \mu_\Theta(\mathbf{g}), \sigma_\Theta^2(\mathbf{g})). \quad (5)$$

Now, given our model φ of the gain pattern, our problem becomes that of inverse regression to find the expected bearing and uncertainty. The Gaussian bearing-error assumption states

$$\hat{\theta} = \theta^* + \nu_\Theta, \quad \nu_\Theta \sim \mathcal{N}(0, \sigma_\Theta^2(\mathbf{g})), \quad (6)$$

when $\hat{\theta} = \mathbf{E}[\Theta | \mathbf{g}]$ and $\sigma_\Theta^2(\mathbf{G}) = \mathbf{V}(\Theta | \mathbf{g})$. We find the expected azimuth by minimising the sum of squares of the residuals, i.e.,

$$\hat{\theta} = \mu_\Theta(\mathbf{g}) = \arg \min_{\theta \in [0, 2\pi)} \|\mathbf{g} - \varphi(\theta)\|^2. \quad (7)$$

To infer the variance $\mathbf{V}(\Theta | \mathbf{g})$ for a given signal \mathbf{g} , we note that the collection of $\{\mathbf{G}\}$ is *heteroscedastic*, i.e., the conditional variance can change with each observation. This is shown in the scattergram in Fig. 2 where the bearing error is plotted against observation quality (correlation). We assume this unexplained variance is due to hidden causes of observation noise, such as the target animal moving during a measurement, or spurious recordings due to multipath interference. In typical regression, heteroscedasticity is considered undesirable and is reduced by introducing more regressors or non-linear transformations of the existing variables. In our case, given that this knowledge is hidden, we cannot introduce more variables and instead marginalise out this quantity to infer the conditional variance from

data. Below, we show how the *coefficient of determination* expresses the proportion of variability in our model (i.e., the heteroscedasticity is attributed to bearing error).

In the context of regression, we can obtain the *fraction of variance unexplained (FVU)* for a response variable through the coefficient of determination. In linear regression, where we have the sample variance s_g^2 as an estimate of the population variance $\mathbf{V}(G)$, the FVU is given by the sample correlation coefficient r^2 :

$$\frac{\mathbf{V}(G | \Theta = \hat{\theta})}{\mathbf{V}(G)} \simeq \frac{s_g^2 \hat{\theta}}{s_g^2} = 1 - r_{g, \varphi(\hat{\theta})}^2. \quad (8)$$

However, we are interested in the bearing variance $\mathbf{V}(\Theta | \mathbf{G})$, which we can approximate from the model variance $\mathbf{V}(\varphi(\Theta) | \mathbf{G})$ by Taylor expansion. Recall that our estimate $\hat{\theta} = \mu_{\Theta}(g_1, g_2, \dots, g_N)$ is a function of the random vector \mathbf{G} . We can approximate the variance of this mapping via a first-order Taylor expansion (45),

$$\mathbf{V}(\Theta | \mathbf{G}) \simeq \sum_{i=1}^N \sum_{j=1}^N \Sigma_{ij} \frac{\partial \mu_{\Theta}(\varphi(\Theta))}{\partial G_i} \frac{\partial \mu_{\Theta}(\varphi(\Theta))}{\partial G_j}. \quad (9)$$

Now, because the measurement \mathbf{G} comprises i.i.d. variables G_n , the covariance matrix is given by $\Sigma = \mathbf{V}(G | \Theta) \mathbf{I}_N$ where \mathbf{I}_N is the identity matrix. This gives the conditional variance in Eq. (9) as

$$\mathbf{V}(\Theta | \mathbf{G}) \simeq \mathbf{V}(G | \Theta) N \sum_{n=1}^N \left(\frac{\partial \mu_{\Theta}(\mathbf{G})}{\partial G_n} \right)^2. \quad (10)$$

Since small changes in each realisation of G will introduce small changes in μ_{Θ} , the variance in Eq. (10) is approximately linear for low noise ν_G ; however, the approximation becomes worse as ν_G becomes large. By using the coefficient of determination (Eq. (8)), we can express the variance of a given sensor reading \mathbf{g} in Eq. (10) as

$$\sigma^2(\mathbf{g}) = \mathbf{V}(\Theta | \mathbf{G} = \mathbf{g}) \simeq s_g^2 \left(1 - r_{\varphi(\hat{\theta}), \mathbf{g}}^2 \right) N \sum_{n=1}^N \left(\frac{\partial \mu_{\Theta}(\mathbf{g})}{\partial g_n} \right)^2. \quad (11)$$

Thus, $\sigma^2(\mathbf{g})$ can be expressed as a function of $s_g^2 \left(1 - r_{\varphi(\hat{\theta}), \mathbf{g}}^2 \right)$.

In practice, we regress only on $\sigma^2(\mathbf{g})$, assuming the variable is a piecewise continuous function

of the explanatory variable $(1 - r_{\varphi(\hat{\theta}),\mathbf{g}})$. We can also determine azimuth $\hat{\theta} = \mu_{\Theta}(\mathbf{g})$ for each measurement \mathbf{g} by the correlation coefficient $r_{\varphi(\hat{\theta}),\mathbf{g}}$. That is, following each observation, the recorded gain pattern is correlated against the model $\varphi(\theta)$ with regular phase offsets θ and the lag that corresponds to the maximum correlation then gives the estimated angle of arrival, i.e., Eq. (7) becomes $\mu_{\Theta}(\mathbf{g}) = \arg \max_{\theta \in [0, 2\pi)} r_{\varphi(\theta),\mathbf{g}}$. This process of obtaining an azimuth observation is illustrated on the left of Fig. 2 and example likelihood functions from one trial can be seen in Fig. 3.

Range Likelihood Function

Next, we estimate the distance to the target using a Gaussian range-error model where the set $\mathbf{Z}_R = \{\mu_R(\mathbf{G}), \hat{\sigma}_R^2\}$. The range errors are assumed to be logarithmic, as discussed below. Furthermore, unlike the bearing observations, the scattergram in Fig. 2 does not indicate the noise is heteroscedastic, i.e., the variance is constant for each observation. This yields the likelihood function

$$\ell_R(\mathbf{y}; \mathbf{x}, \mathbf{z}_R) = f(\log(d(\mathbf{x}, \mathbf{y})); \mu_R(\mathbf{g}), \hat{\sigma}_R^2). \quad (12)$$

In general, range measurements in cluttered environments can be highly imprecise due to multipath interference. We anticipate the vehicle to be deployed in similar environments and estimate the variance under these conditions. Although the error in range measurements can be significant, including such observations is still useful. Because the noise is homoscedastic, we can rely on range measurements to provide an approximate location. The ability to focus on an approximate location is particularly beneficial when the search area would otherwise be expansive, such as in tracking scenarios where there is little prior knowledge of the target's location, and when bearing uncertainty is high.

We are interested in mapping the sensor output \mathbf{g} to the distance between transmitter and receiver. Due to atmospheric interactions, the signal amplitude will decrease with range. Denote $d(\mathbf{x}, \mathbf{y})$ as the Euclidean distance between our receiver \mathbf{x} and the transmitter \mathbf{y} . Then, the received power p_r is a

function of the transmitted power p_t and the attenuation per meter α (46):

$$p_r = p_t e^{\alpha d(\mathbf{x}, \mathbf{y})}. \quad (13)$$

In Eq. (13) we have assumed that p_r and p_t take into account the link budget, which characterises all gains and losses in the telecommunication system. Most of these components are fixed for a given system (e.g., transmitter and receiver losses), however, for a directional antenna, the gain relative to the average radiation intensity (the *isotropic directivity*) depends on the immediate angle of arrival ϕ_n . As a result, the RSSI values g_n are a function of the received power p_r and angle of arrival ϕ_n . The isotropic directivity is approximately constant if we take the maximum RSSI value $g_{\max} = \max_n g_n$. Thus, we use the value g_{\max} to estimate distance.

Now, let the true distance to the target be $r^* = d(\mathbf{x}, \mathbf{y}^*)$ and its estimate be a function of \mathbf{g} , i.e., $\hat{r} = \mu_R(\mathbf{g})$. From the above discussion and Eq. (13), $p_r = p(g_{\max})$ for some linear function $p : \mathbb{R} \rightarrow \mathbb{R}$. Moreover, r^* is a function of $\log p(g_{\max})$ and the Gaussian range-error assumption may be expressed as

$$\log \hat{r} = \log r^* + \nu_R, \quad \nu_R \sim \mathcal{N}(0, \sigma_R^2), \quad (14)$$

where $\sigma_R^2 = \mathbf{V}(\log R)$. We thus obtain the estimated range \hat{r} as

$$\log \hat{r} = \mu_R(\mathbf{g}) = \alpha^{-1}(\log p(g_{\max}) - \log p_t). \quad (15)$$

The function $\mu_R(\mathbf{G})$ can be fitted to a first degree polynomial function of $\log g_{\max}$. The variance σ_R^2 is estimated by the sample variance $\hat{\sigma}_R^2 = s_R^2$. The procedure for obtaining a Gaussian range-error observation is illustrated on the right of Fig. 2 and example range likelihood functions can be seen in Fig. 3.

Combined Likelihood Function

The individual likelihood functions may be combined to obtain a range-azimuth likelihood function $\ell(\mathbf{y}; \mathbf{u})$, where $\mathbf{Z} = \{\mu_\Theta(\mathbf{G}), \sigma_\Theta^2(\mathbf{G}), \mu_R(\mathbf{G}), \hat{\sigma}_R^2\}$. That is, assuming independent errors ν_Θ and ν_R ,

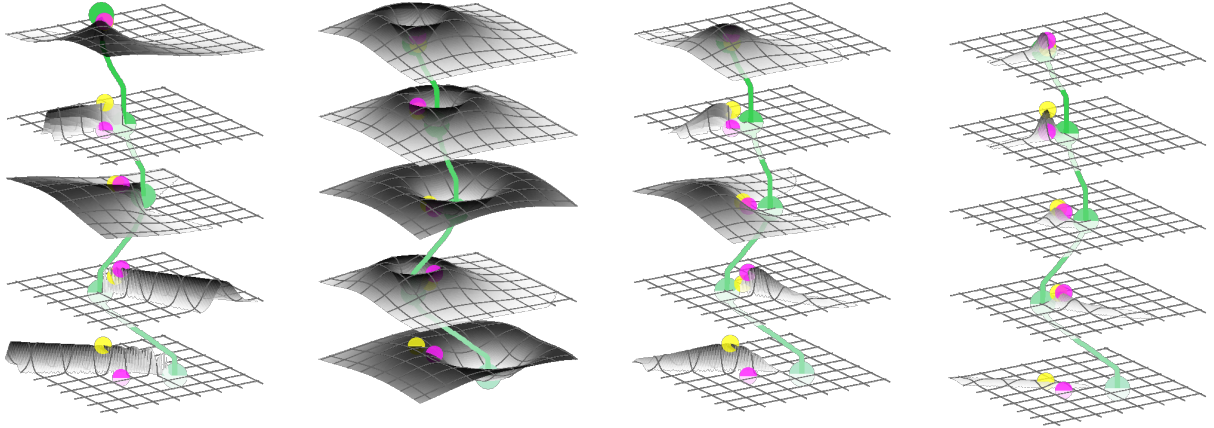


Figure 3: Bayesian data fusion to obtain target estimates. The distributions shown are spatially discrete grids over a 750 m squared area (with grid-lines every 100 m for illustrative purposes only). From left to right: the bearing-only likelihood function ℓ_{Θ} , the range-only likelihood function ℓ_R , the combined likelihood function ℓ , and the posterior belief $p(\cdot)$. In each subfigure, the first observation ($k = 1$) is shown in the lowermost grid, the last observation ($k = 4$) is uppermost, and higher probability mass is represented as darker, raised regions. The UAV location \mathbf{x}_k is indicated by a green dot, the target location \mathbf{y}_k^* in purple, and the maximum likelihood estimate $\hat{\mathbf{y}}_k$ in yellow.

the likelihood functions are multiplied pointwise (44), i.e.,

$$\ell(\mathbf{y}; \mathbf{u}) = \ell_{\Theta}(\mathbf{y}; \mathbf{x}, \mathbf{z}_{\Theta}) \circ \ell_R(\mathbf{y}; \mathbf{x}, \mathbf{z}_R). \quad (16)$$

We tested the null hypothesis that these errors are independent by computing the sample correlation coefficient. Since the errors are assumed to be normal, the hypothesis was tested via a Student's t -distribution with 95% confidence and 150 observations. The results showed a correlation of $r_{\nu_{\Theta}, \nu_R} = -0.08 \pm 0.136$, giving a confidence of less than 66% that the errors are correlated. This result further supports the heteroscedasticity assertion, i.e., that poor quality observations are not significantly correlated with distance.

Bayesian Data Fusion

Given the likelihood function in Eq. (16), we can combine numerous observations to determine the most likely position of the target animal. To achieve this, we use Bayesian data fusion, assuming independent observations; this process is illustrated in Fig. 3.

By time t we have obtained observations of the target animal at a set of times $0 \leq t_1 \leq \dots \leq t_k \leq t$. The above notation is extended to include the observation number k . Denote $\mathbf{Y}_k = \mathbf{Y}(t_k)$ as the target's location at time t_k and, similarly, $\mathbf{Z}_k = \mathbf{Z}(t_k)$ is the observation taken at this time. Moreover, let $\mathbf{U}_k = \{\mathbf{X}_k, \mathbf{Z}_k\}$; denote the measurement process as $\mathbf{Z}_{1:k} = (\mathbf{Z}_1, \dots, \mathbf{Z}_k)$ and the values it takes on as $\mathbf{z}_{1:k} = (z_1, \dots, z_k)$. We are ultimately interested in knowing the probability of the target's state after all K observations, i.e., the *posterior belief* (35, 44),

$$p(t_K, \mathbf{y}_K) = \Pr(\mathbf{Y}_K = \mathbf{y}_K \mid \mathbf{U}_{1:K} = \mathbf{u}_{1:K}). \quad (17)$$

Further, we assume that the target can transition between observations such that $\mathbf{y}_k = \mathbf{y}_{k-1} + \boldsymbol{\nu}_Y$ with $\boldsymbol{\nu}_Y \sim \mathcal{N}(0, \boldsymbol{\Sigma}_Y)$ for some covariance $\boldsymbol{\Sigma}_Y$. This leads to the *transition density*

$$q(\mathbf{y}_k \mid \mathbf{y}_{k-1}) = \Pr(\mathbf{Y}_k = \mathbf{y}_k \mid \mathbf{Y}_{k-1} = \mathbf{y}_{k-1}). \quad (18)$$

Computing the posterior belief (Eq. (17)) becomes simpler if the process $(\mathbf{Y}(t))_{t \geq 0}$ is assumed to be Markovian and each observation \mathbf{Z}_k only depends on \mathbf{Y}_k , i.e., $\ell(\mathbf{y}_{1:K}; \mathbf{u}_{1:K}) = \prod_{k=1}^K \ell(\mathbf{y}_k; \mathbf{u}_k)$. As the likelihood function in Eq. (16) is defined this way, recursive Bayesian filtering (44) can be used to update the belief. That is, the posterior belief is computed as

$$p^-(t_k, \mathbf{y}_k) = \int d\mathbf{y}_{k-1} q(\mathbf{y}_k \mid \mathbf{y}_{k-1}) p(t_{k-1}, \mathbf{y}_{k-1}) \quad (19)$$

$$p(t_k, \mathbf{y}_k) = \eta \ell(\mathbf{y}_k; \mathbf{u}_k) p^-(t_k, \mathbf{y}_k), \quad (20)$$

where η is a normalisation constant such that $\int d\mathbf{y}_k p(t_k, \mathbf{y}_k) = 1$ and $\ell(\mathbf{y}_k; \mathbf{u}_k)$ is the likelihood function (Eq. (16)). The first step, Eq. (19), gives a *motion update*, and the second step, Eq. (20), gives the *information update* to obtain a new belief of the target location (44).

Early approaches to recursive Bayesian filtering focused on Gaussian implementations due to convenient analytical solutions to computing the posterior belief in Eq. (17), e.g., Kalman filters (KFs) and extensions such as the unscented and extended KFs. However, these methods are approximations

to the nonlinear, non-Gaussian Bayesian filter (shown in Eqs. (19)-(20)). Grid-based filtering allows for resolution-complete recursive estimation (44, 47) and can be computed in reasonable time over our workspace. Thus, we represent our workspace \mathcal{S} as an $I \times J$ grid in \mathbb{R}^2 .

The evolution model, Eq. (19), is functionally equivalent to Gaussian convolution. Further, given our grid-based workspace \mathcal{S} , this convolution is simply a Gaussian blur, a spatial (low-pass) filter commonly used in image processing. To efficiently implement this model, we leverage results from computer vision for convolution and use a separable Gaussian kernel of width $3|\Sigma_Y|$.

Finally, we require an estimate of the location of the target $\hat{\mathbf{y}}_k$ given the posterior $p(t_k, \mathbf{y}_k)$. Two obvious choices are the expected value of the posterior $\mathbf{E}[\mathbf{Y}_k] = \int d\mathbf{y}_k \mathbf{y}_k p(t_k, \mathbf{y}_k)$, or the maximum *a posteriori* probability (MAP) estimate $\arg \max_{\mathbf{y}_k \in \mathcal{S}} p(t_k, \mathbf{y}_k)$. The MAP estimate performed marginally better in preliminary trials; however, in practice the target does not remain stationary and so we instead maximise recursively over all posteriors:

$$\hat{\mathbf{y}}_k = \arg \max_{u \in [1, k], \mathbf{y}_u \in \mathcal{S}} p(t_u, \mathbf{y}_u).$$

In this way, the location estimate likelihood is strictly increasing.

Decision Making by Information Gain

Our overall objective is to know where our target animal is and with what certainty. Thus, the problem can be considered under the framework of *information gathering* (35). To quantify uncertainty, we use Shannon entropy, a standard measure for this purpose. The conditional (Shannon) entropy of a random variable V , given another variable W , quantifies the uncertainty over the outcomes of V in the context of W ; mathematically, this is given by $H(V | W) = \mathbf{E}[\log \Pr(V | W)]$.

In this context, we aim to choose a sequence $\mathbf{u}_{1:K} = \{\mathbf{x}_{1:K}, \mathbf{z}_{1:K}\}$ of state-measurement pairs such that the final entropy of the belief $H(\mathbf{Y}_N)$ is minimised. That is, letting $\mathcal{U} = (\mathcal{S} \times \mathcal{H})$ and fixing the

measurement space \mathcal{H} , the objective can be stated as

$$\mathbf{U}_{1:K}^* = \arg \min_{\mathbf{U}_{1:K} \subseteq \mathcal{U}} \mathbf{E}_{\mathbf{U}_{1:K}} [H(\mathbf{Y}_{1:K})], \quad (21)$$

where the posterior belief $p(t_k, \mathbf{y}_k)$ is a function of the robot position and measurements $\mathbf{u}_{1:k}$.

However, it is more convenient to consider the equivalent problem of maximising the information gain of each observation. Let \mathbf{Y}_k^- be distributed according to the target belief after the motion update step, i.e., Eq. (19)). The information gained in taking the action $\mathbf{U}_k = \mathbf{u}_k$ is quantified by the mutual information $I(\mathbf{Y}_k; \mathbf{Y}_k^-)$ between the posterior and the prior belief:

$$I(\mathbf{Y}_k; \mathbf{Y}_k^-) = H(\mathbf{Y}_k) - H(\mathbf{Y}_k | \mathbf{Y}_k^-). \quad (22)$$

Decomposing Eq. (22) using the chain rule, the entropy minimisation problem defined in Eq. (21) can be expressed as

$$\mathbf{U}_{1:K}^* = \arg \max_{\mathbf{U}_{1:K} \subseteq \mathcal{U}} \mathbf{E}_{\mathbf{U}_{1:K}} \left[H(\mathbf{Y}_1) + \sum_{k=2}^K I(\mathbf{Y}_k; \mathbf{Y}_k^-) \right]. \quad (23)$$

The objective of Eq. (23) is equivalent to entropy minimisation and is, in general, non-convex and analytically intractable. However, the mutual information given in Eq. (22) is monotone submodular and thus the quality of the solution provided by a greedy algorithm is at least 63% of optimal (36). That is, given a deterministic greedy algorithm that selects the action

$$\mathbf{U}_k = \arg \max_{\mathbf{U}_k \in \mathcal{U}} \mathbf{E} [I(\mathbf{Y}_k; \mathbf{Y}_k^-)] \quad (24)$$

at each decision step, the resulting path $\hat{\mathbf{u}}_{1:K}$ is within a constant factor of optimal of the objective shown in Eq. (23), i.e.,

$$\hat{\mathbf{u}}_{1:K} \geq \left(1 - \frac{1}{e}\right) \mathbf{u}_{1:K}^*. \quad (25)$$

Furthermore, this is the most efficient algorithm to obtain such a bound unless $P = NP$ (36).

Optimising each observation \mathbf{U}_k is constrained in that only the vehicle locations $\mathbf{x}_k \subset \mathbf{u}_k$ can be selected, and consequently only the expected information gain at each sample s can be computed, i.e.,

we choose future waypoints \mathbf{x}_k such that

$$\mathbf{x}_k = \arg \max_{\mathbf{s} \in \mathcal{S}} \mathbf{E} [I(\mathbf{Y}_k; \mathbf{Y}_k^-) | \mathbf{X}_k = \mathbf{s}]. \quad (26)$$

As mentioned above, we assume independent errors in the likelihood functions shown in Eq. (16), giving $\Pr(\mathbf{Z}) = \Pr(\mathbf{Z}_\Theta) \Pr(\mathbf{Z}_R)$. However, even solving for independent priors requires inverting all possible distributions at all sample locations $\mathbf{s} \in \mathcal{S}$; this is generally intractable.

As an efficient alternative, we assume that the target location for the next observation is the maximum likelihood position after the motion update, i.e., $\mathbf{Y}_k = \hat{\mathbf{y}}_k^-$. As a result, for a fixed viewpoint \mathbf{s} , the expected range measurement $\mathbf{E}[R | \mathbf{X}_k = \mathbf{s}, \mathbf{Y}_k = \hat{\mathbf{y}}_k^-] = d(\mathbf{s}, \hat{\mathbf{y}}_k^-)$ and expected bearing measurement $\mathbf{E}[\Theta | \mathbf{X}_k = \mathbf{s}, \mathbf{Y}_k = \hat{\mathbf{y}}_k^-] = b(\mathbf{s}, \hat{\mathbf{y}}_k^-)$ to the target are given. Moreover, the expected variance $\hat{\sigma}_\Theta^2$ is given by marginalising out \mathbf{G} such that $\hat{\sigma}_\Theta^2 = \mathbf{E}[\sigma_\Theta^2(\mathbf{G})] \simeq 0.2$ radians. In this case, the expected observation is a function of the viewpoint \mathbf{s} :

$$\hat{\mathbf{z}}_k(\mathbf{s}) = \{d(\mathbf{s}, \hat{\mathbf{y}}_k^-), \hat{\sigma}_\Theta^2, b(\mathbf{s}, \hat{\mathbf{y}}_k^-), \hat{\sigma}_R^2\},$$

and the optimisation over potential viewpoints \mathbf{s} from Eq. (26) becomes

$$\mathbf{x}_k = \arg \max_{\mathbf{s} \in \mathcal{S}} I(\mathbf{Z}_k = \hat{\mathbf{z}}_k(\mathbf{s}); \mathbf{Y}_k = \hat{\mathbf{y}}_k^-). \quad (27)$$

To reduce computation time, instead of sampling every location in the workspace $\mathbf{s} \in \mathcal{S}$ as indicated in Eq. (27), we simply sample a uniformly distributed subset. Given the stochastic nature of observations this does not appear to affect the quality of the planner.

Evaluating the Performance of the System

To validate our approach, we compared the performance of the robotic system against human tracker performance in locating swift parrots in the wild. The results indicate that the robot is able to approximate the location of the target species in less time than human trackers. Moreover, the reported position estimates are obtained in less than 5 observations (approximately 10 minutes) and lie within 50m of the




Figure 4: Evaluating the performance of the robotic system through comparison with human trackers. We performed two flights at each of four trial sites (eight flights in total). The box plot shows 1) the error in the robot system’s target estimate $\hat{\mathbf{y}}_k$ for each observation k , and 2) the Euclidean distance between the human tracker’s location at time t_k and their final estimate. In two of the flights, the final location was uncertain because the bird moved during the flight and visual confirmation was lost. The case where data from these two flights are removed is labelled *Robot (certain)* (blue boxes).

true bird location.

The boxplot in Fig. 4 collates the tracking performance from eight flights at four different sites near Temora, New South Wales, Australia. At each site, we obtained the GPS trajectory of a novice and an expert tracker performing manual wildlife telemetry. Once a human tracker had established the true location \mathbf{y} for the target, the UAV began its flight trial. A flight was performed for each type of tracker (novice, expert). Thus, we obtained two tests of the robot system at each site with known true bird location.

For fair comparison, both the robot and the humans began trials from the same initial coordinates, with the target animal location unknown. This starting location was chosen such that the radio signal was strong enough to be measured by the on-board payload. In order to quantify performance, we compare the robotic tracker estimate $\hat{\mathbf{y}}_k$ with the Euclidean distance between the human tracker and the final ground truth \mathbf{y} at that time t_k . These locations are plotted in Fig. 4, where the robot, on average, takes

less time to locate the bird to within 50 m (around Observation 2, between 143–289 s).

Ecological Significance of Trials

The quantitative data from our trials provide significant insight into the movement and habitat of swift parrots. The Temora region was chosen because, based on a small number of sightings, it was assumed that numerous swift parrots had migrated to the area in the weeks leading up to the trial (see Materials and Methods). The results in this paper were obtained over a seven-day trial in the region and the posterior estimates from all flights were aggregated to yield the heatmap shown in Fig. 5. The figure shows that the flocks used two distinct areas for foraging and roosting, including sites where the species had not previously been recorded.

Swift parrots are small, critically endangered migratory birds that are dependant on highly variable winter nectar resources. As a result, the small population (less than 2000 birds) spreads across vast areas of South-Eastern Australia each year in search of suitable food. Given their small body and hence tag size, as well as their capacity for highly variable and large movements, this species has never before been successfully radio-tracked. Figure 5 provides an example of the distribution and abundance of swift parrots across their winter range in a single season, together with the location of our study site.

Discussion

In this paper, we validated that our customised aerial robot can be used to perform autonomous wildlife tracking. We presented rigorous mathematical derivation of all algorithmic components of the system, including a novel approach to computing the uncertainty of each bearing-only observation where heteroscedasticity is assumed. The resulting system performed comparably to and often better than skilled humans in tracking the critically endangered swift parrot (*Lathamus discolor*).

Wildlife tracking is known to be an important but difficult problem and tracking members of this species is particularly challenging due to their small size and highly dynamic movements. The ability of our system to track such animals thus exemplifies the capability of robotic wildlife trackers and the

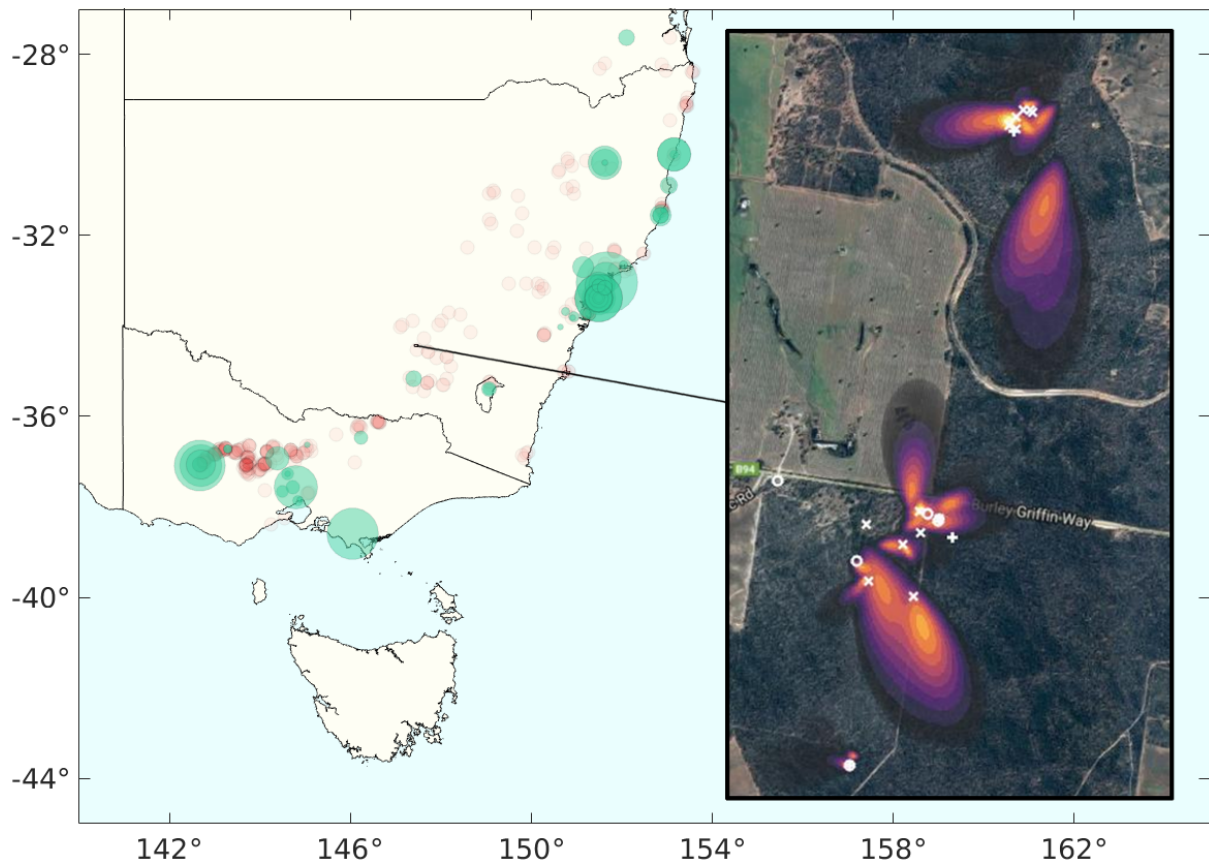


Figure 5: Recorded spatial distribution of swift parrots. Main figure: swift parrot sightings in South-East Australia for May 2007. Green circles denote sites where flocks were confirmed, red circles denote failed surveys where no birds were found. The size of the green circles indicate the size of the flock, ranging from 1 to 100 birds. Inset: heatmap illustrating aggregated posterior distributions from our trial in July 2017. The posterior distributions of all trials were normalised and aggregated to give an indication of the most likely foraging and roosting areas. White symbols indicate locations where ground truth data (confirmed by visual sightings) were available; each tag has a unique symbol: 'o', 'x', '+', or '*'.

possibility of these systems to dramatically facilitate conservation management.

Although we performed this study within a landscape that is relatively easy to traverse on foot (as discussed in Materials and Methods), the greatest benefits of this aerial tracking technology are likely to be realised within densely vegetated areas or rugged and dangerous landscapes. These situations require significantly increased amounts of time and effort on the ground to locate tagged animals (relative to flat terrain) but no additional time or effort when tracking from the air.

There are many avenues of inquiry for improving the system hardware and decision making algorithms. In particular, multi-robot extensions and long time-horizon planning with travel costs would allow for efficient search and tracking of numerous animals simultaneously. The problem of multi-robot wildlife telemetry tracking has been partially addressed by designing optimal information gathering algorithms without considering travel cost (18). However, these are yet to be used in real tracking experiments. A recent approach to information gathering for decentralised active perception allows for any general reward functions in order to perform distributed optimisation (48).

When studying fine-scale movement patterns (e.g., of highly dynamic animals) it is desirable to maintain real-time information about individual trajectories. In robotics, this general problem is known as *persistent monitoring*. Existing approaches, however, often seek to maintain information about an entire, continuous environment (49) rather than monitoring a small number of discrete features (e.g., birds). Recent approaches partition the environment into discrete spatial locations and model the likelihood of observing birds using Poisson processes (50). Extensions to this model have been made where it is assumed that the presence of a robot interferes with the animals' behaviour (51).

Materials and Methods

Robot System

The aerial robot system used in this work was originally presented in (21). We summarise key details here for convenience.

The system comprised a commercial eight-rotor UAV (Ascending Technologies' *Falcon 8*), custom antenna array, sensor payload, and ground control station (GCS) laptop running ROS (52). All algorithmic component implementations were executed on the GCS. A radio-link between the GCS and the UAV allowed for telemetry and autonomous waypoint following.

The radio signal from each radio-tagged bird was received and filtered by custom electronics on-board the UAV and transmitted to the GCS (producing the sampled signal described in Observed and Expected Sensor Data). After each observation and subsequent decision making step, the GCS communicated a new waypoint to the UAV for autonomous navigation.

Experimental Setup

Here we present the procedure used to choose the survey location and describe its terrain. We then briefly explain the manual and robotic tracking procedures.

The trials were performed on 3–7 July 2017 in Temora, New South Wales, Australia. Prior to these trials, six birds were detected by an experienced volunteer undertaking targeted surveys in the surrounding Riverina region, where swift parrots are known to migrate on a regular basis. Follow-up surveys were conducted by the authors in late June, confirming that the survey location was suitable for this trial by detecting at least 30 birds. By the end of August, at least 200 swift parrots (10% of the global population) were detected in the area. The study site is an open, grassy, Box Ironbark woodland, and thus relatively easy to traverse on foot. However, locating the birds was often complicated by logistical issues such as limited road accessibility, fence lines and different land tenures, including private property.

For the trials, several birds were captured and, subsequently, a BioTrack Pip Ag393 radio tag was taped to their back feathers. Each tag transmitted on a unique frequency that was preprogrammed into both the manual and robotic receiver systems. The tags emitted an on-off key modulated signal with a pulse width of 10 ms and a period of 1.05 s. Moreover, they were lightweight (approximately 2 g) and, subsequently, low-power (less than 1 mW) transmitters due to the small size of the species.

Manual tracking was undertaken using a Titley Australis 26k very high frequency (VHF) radio receiver system and a Yagi three element hand-held directional antenna (shown earlier in Fig. 1). The approximate location of a bird was identified by driving in an offroad-capable vehicle to different sites until a radio signal was audible from the receiver. Once a signal was detected, the tracker continued to point the antenna towards the strongest (loudest) signal while walking through the landscape. This procedure involved constant adjustment of the volume and the gain of the receiver and continued until the bird was sighted. The GPS trajectory followed by the manual tracker was recorded.

After achieving visual confirmation of a bird's location, the UAV was launched approximately from the manual tracker's starting position. The UAV trajectory and raw sensor data were recorded in real-time and later replayed to generate the figures reported. Each flight was performed at a constant altitude of 75 m (such that the canopy was cleared) and each observation took approximately 45 s to complete. For planning viewpoints, the UAV was constrained to choose locations within 300 m of the GCS (i.e., the starting position) for the pilot to maintain visual line-of-sight.

For Bayesian data fusion, the workspace $\mathcal{S} \subset \mathbb{R}^2$ was discretised into a square, 300×300 grid, i.e., $I = 300$ and $J = 300$. Each cell represented a 5×5 m area and thus the workspace extended 750 m in all cardinal directions from the GCS. We assumed a uniform prior on the target location and evolution model covariance $\Sigma_Y = \sigma_Y I_2$, where $\sigma_Y = 20$ m.

Acknowledgements

This work was supported in part by the Riverina Local Land Services (project number RV01221), the Australian Centre for Field Robotics, and the New South Wales State Government. Special thanks to Jeremy Randle, Richard Allen, and Matthew Truman for assistance with flight trials. We also thank Graham Brooker and David Johnson for ongoing technical advice on radio signal processing.

Figure 1b is entitled 'Petrogale penicillata with radio tracking collar' by Hugh McGregor and available at <http://journals.plos.org/plosone/article/figure?id=10.1371/journal.pone.0063017.g001> under a Cre-

ative Commons Attribution 4.0 (full terms at <https://creativecommons.org/licenses/by/4.0/legalcode>).

Figure 1c was reproduced from (21) with permission of the corresponding author.

References

1. R. Nathan, *et al.*, *Proceedings of the National Academy of Sciences* **105**, 19052 (2008).
2. R. Kays, M. C. Crofoot, W. Jetz, M. Wikelski, *Science* **348**, aaa2478 (2015).
3. N. E. Hussey, *et al.*, *Science* **348**, 1255642 (2015).
4. D. Chabot, D. M. Bird, *Journal of Unmanned Vehicle Systems* **3**, 137 (2015).
5. L. F. Gonzalez, *et al.*, *Sensors* **16**, 97 (2016).
6. K. Zhao, R. Jurdak, *Scientific Reports* **6**, 31967 (2016).
7. S. Panigada, *et al.*, *Scientific Reports* **7**, 3365 (2017).
8. P. Berthold, M. Kaatz, U. Querner, *Journal of Ornithology* **145**, 356 (2004).
9. A. Kessler, N. Batbayar, T. Natsagdorj, D. Batsuur, A. Smith, *Journal of Avian Biology* **44**, 311 (2013).
10. S. Oppel, A. N. Powell, D. L. Dickson, *The Condor* **110**, 296 (2008).
11. I. Priede, *Fisheries Research* **2**, 201 (1984).
12. B. A. Block, *et al.*, *Nature* **475**, 86 (2011).
13. E. S. Bridge, *et al.*, *Bioscience* **61**, 689 (2011).
14. R. D. Lord, F. C. Bellrose, W. W. Cochran, *Science* **137**, 39 (1962).
15. P. Tokekar, D. Bhaduria, A. Studenski, V. Isler, *Journal of Field Robotics* **27**, 779 (2010).

16. P. Tokekar, E. Branson, J. Vander Hook, V. Isler, *IEEE Robotics and Automation Magazine* **20**, 33 (2013).
17. J. V. Hook, P. Tokekar, V. Isler, *Journal of Field Robotics* **31**, 296 (2014).
18. J. Vander Hook, P. Tokekar, V. Isler, *IEEE Transactions on Robotics* **31**, 864 (2015).
19. H. Bayram, J. Vander Hook, V. Isler, *IEEE Robotics and Automation Letters* **1**, 369 (2016).
20. N. Noori, A. Renzaglia, J. Vander Hook, V. Isler, *IEEE Transactions on Robotics* **32**, 261 (2016).
21. O. M. Cliff, R. Fitch, S. Sukkarieh, D. L. Saunders, R. Heinsohn, *Proceedings of Robotics: Science and Systems (RSS)* (2015).
22. A. Posch, S. Sukkarieh, *Proceedings of the ARAA Australasian Conference on Robotics and Automation (ACRA)* (2009), pp. 107–112.
23. F. Korner, R. Speck, A. H. Goktogan, S. Sukkarieh, *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots (IROS)* (2010), pp. 107–112.
24. M. Leonardo, A. M. Jensen, C. Coopmans, M. McKee, Y. Chen, *Proceedings of the ASME International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (IDETC/CIE)* (2013), p. V004T08A056.
25. G. A. M. Dos Santos, *et al.*, *Proceedings of the IEEE International Conference on Mobile Ad Hoc and Sensor Systems (MASS)* (2014), pp. 761–766.
26. A. Jensen, Y. Chen, *Proceedings of the IEEE International Conference on Unmanned Aircraft Systems (ICUAS)* (2013), pp. 1144–1149.
27. A. M. Jensen, D. K. Geller, Y. Chen, *Journal of Intelligent and Robotic Systems* **74**, 287 (2014).

28. H. Bayram, N. Stefas, K. S. Engin, V. Isler, *Proceedings of the IEEE International Symposium on Multi-Robot and Multi-Agent Systems (MRS)* (2017), pp. 97–103.
29. R. Bajcsy, *Proceedings of the IEEE* **76**, 966 (1988).
30. R. Bajcsy, Y. Aloimonos, J. K. Tsotsos, *Autonomous Robots* **42**, 177 (2017).
31. Y. Liu, G. Nejat, *Journal of Intelligent and Robotic Systems* **72**, 147 (2013).
32. F. Caballero, L. Merino, I. Maza, A. Ollero, *Proceedings of the IEEE Conference on Robotics and Automation (ICRA)* (2008), pp. 596–601.
33. R. Kays, *et al.*, *The Computer Journal* **54**, 1931 (2011).
34. T. Patten, M. Zillich, R. Fitch, M. Vincze, S. Sukkarieh, *IEEE Robotics and Automation Letters* **1**, 73 (2016).
35. S. Thrun, W. Burgard, D. Fox, *Probabilistic Robotics* (The MIT Press, 2005).
36. A. Krause, C. Guestrin, *ACM Transactions on Intelligent Systems and Technology* **2**, 32 (2011).
37. B. Grocholsky, Information-theoretic control of multiple sensor platforms, Ph.D. thesis, The University of Sydney (2002).
38. G. M. Hoffmann, C. J. Tomlin, *IEEE Transactions on Automation and Control* **55**, 32 (2010).
39. A. T. Klesh, P. T. Kabamba, A. R. Girard, *Proceedings of the IEEE American Control Conference (ACC)* (2008), pp. 1991–1996.
40. J. Nguyen, N. Lawrance, R. Fitch, S. Sukkarieh, *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)* (2013), pp. 3825–3831.
41. A. Ryan, J. K. Hedrick, *Robotics and Autonomous Systems* **58**, 574 (2010).

42. P. Skoglar, J. Nygard, M. Ulvklo, *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2006), pp. 2436–2442.
43. J. Tisdale, Z. Kim, J. Hedrick, *IEEE Robotics and Automation Magazine* **16**, 35 (2009).
44. L. D. Stone, R. L. Streit, T. L. Corwin, K. L. Bell, *Bayesian Multiple Target Tracking* (Artech House, 2013).
45. A. H.-S. Ang, W. H. Tang, *Probability Concepts in Engineering Planning and Design*, vol. I (John Wiley & Sons, Inc., 1984).
46. G. Brooker, *Introduction to sensors for ranging and imaging* (SciTech Publishing, Inc., 2009).
47. M. S. Arulampalam, S. Maskell, N. Gordon, *IEEE Transactions on Signal Processing* **50**, 174 (2002).
48. G. Best, O. Cliff, T. Patten, R. R. Mettu, R. Fitch, *The International Journal of Robotics Research* (2018).
49. S. Garg, N. Ayanian, *Proceedings of Robotics: Science and Systems* (2014).
50. J. Yu, S. Karaman, D. Rus, *IEEE Transactions on Robotics* **31**, 521 (2015).
51. B. Hefferan, O. M. Cliff, R. Fitch, *Proceedings of the ARAA Australasian Conference on Robotics and Automation (ACRA)* (2016).
52. M. Quigley, *et al.*, *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Workshop on Open Source Software* (2009).