

# Optimal path planning using psychological profiling in drone-assisted missing person search

Jan-Hendrik Ewers<sup>1</sup> | David Anderson<sup>1</sup> | Douglas Thomson<sup>1</sup>

Aerospace Sciences Research Division,  
University of Glasgow, Glasgow, Scotland

## Correspondence

Jan-Hendrik Ewers, Aerospace Sciences  
Research Division, University of Glasgow,  
Glasgow, Scotland.

Email: [j.ewers.1@research.gla.ac.uk](mailto:j.ewers.1@research.gla.ac.uk)

## Funding information

Engineering and Physical Sciences  
Research Council, Grant/Award Number:  
EP/T517896/1-312561-05

## Abstract

Search and rescue operations are all time-sensitive and this is especially true when searching for a vulnerable missing person, such as a child or elderly person suffering dementia. Recently, Police Scotland Air Support Unit has begun the deployment of drones to assist in missing person searches with success, although the efficacy of the search relies upon the expertise of the drone operator. In this paper, several algorithms for planning the search path are compared to determine which approach has the highest probability of finding the missing person in the shortest time. In addition to this, the use of a priori psychological profile information of the subject to create a probability map of likely locations within the search area was explored. This map is then used within a nonlinear optimization to determine the optimal flight path for a given search area and subject profile. Two optimization solvers were compared; genetic algorithms, and particle swarm optimization. Finally, the most effective algorithm was used to create a coverage path for a real-life location, for which Police Scotland Air Support Unit completed multiple test flights. The generated flight paths based on the predicted intent of the lost person were found to perform statistically better than those of the expert police operators.

## KEYWORDS

optimization, search and rescue, UAV

## 1 | INTRODUCTION

Search and Rescue (SAR) of vulnerable missing persons is unfortunately a common task for the Police and other emergency services. Organizations like the Centre for Search and Rescue (CSR)<sup>1</sup> carry out research and training in areas related to SAR and while their training and published papers offer a valuable resource to the people responsible for finding lost or missing persons, they focus only on land-based search that is, directing teams of individuals. This is a slow and methodical process that would undoubtedly benefit from the assistance of an airborne surveillance platform. As a result of rapid advances in the drone sector over the last decade, multirotors have become cheaper and more accessible than ever before. Consequently, several concept evaluation trials for measuring the efficacy of incorporating drones into the search process have recently been undertaken in Scotland.<sup>2</sup>

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Advanced Control for Applications: Engineering and Industrial Systems* published by John Wiley & Sons Ltd.

Air-based SAR is a core operational requirement of the Police Scotland Air Support Unit (PS ASU), which regularly carry out SAR operations using a Eurocopter EC135 helicopter. However, operating a single helicopter platform to service the whole of Scotland introduces significant limitations. The main problem is that the transit time taken to reach the more remote locations adds a lot of lead time to the search operation in instances where seconds could be the difference between life and death. Hence, PS ASU is placing a small fleet of drones in key locations around Scotland for rapid deployment to assist in missing person search (MPS) operations.<sup>2</sup> These will not be replacing the helicopter, but are rather intended to complement its operations.

To achieve a similar spatial resolution, and thus the probability of detection, as the costlier high-resolution sensors onboard the helicopter, the sensor footprint on each drone must be significantly smaller. Therefore, if little thought is given to the path-planning algorithm used to guide the drones, their use may well be suboptimal in time-sensitive, MPS scenarios. As with all optimization problems, the efficacy of any optimal path will only be as good as the information used in the cost function construction. Unfortunately there exists no deterministic algorithm to predict the exact location of any specific missing person, but there are psychological markers extracted via forensic analysis of previous cases by organisations such as the CSR that could be exploited to assign typical behaviors to classes of subject<sup>1,3-7</sup> and so doing generate location probabilities. Our research hypothesis is that using the predicted intent of the lost person in the form of psychological profiles, better probability maps may be generated for a given search area resulting in lower time-to-finds compared to *naïve* methods.

This article is organized as follows. Section 2 will present a short overview of the current popular search patterns and algorithms. Following this, in Section 3, the models used to capture the drone dynamics, sensor characteristics, and generate trajectories are presented. Section 4 develops the nonlinear program and introduces the probability heat-map concept and illustrates how this is integrated into the final cost function. Section 5 presents the results of a Monte-Carlo analysis of a typical scenario, comparing the performance of each optimization algorithm. Finally, the most successful technique was compared against real-world flight data from PS ASU.

## 2 | CURRENT SEARCH ALGORITHMS

Search operations are a specialization or subset of the more general trajectory generation problem that has been widely investigated in the research literature.<sup>8-11</sup> Most, if not all trajectory generation methods utilise the concept of *waypoints* or *special-points-of-interest* (SPI) to define the trajectory constraints with various interpolation schemes used to connect them and form a continuous trajectory.<sup>12-15</sup> The resulting trajectories can be flown either by a pilot or autonomously. Search algorithms may then be interpreted as trajectory generation problems with additional constraints and should ideally exhibit the following properties:

- incorporate as many constraints as possible
- yields a feasible trajectory in the form of a waypoint sequence
- be computationally efficient
- be simple to use and interpret

Trajectory generation with simple search patterns is a core functionality in many ground station control software packages. Figure 1 shows example images from QGroundControl and Mission Planner displays showing the standard trajectory specification software. Waypoints specified by the user can be seen for defining a search area in Figure 1A and a specific flightpath in Figure 1B.

The most basic kind of search algorithm is known as the *parallel swaths* pattern (also known as the lawn-mower pattern)<sup>16,17</sup> due to its distinctive, regular structure. Parallel swaths are particularly well suited for flat areas with uniform probability distributions<sup>18</sup> and operations where complete coverage of the search area is the goal. Furthermore, parallel swaths is highly accessible and is the de facto coverage algorithm within popular open-source ground control software such as QGroundControl and Mission Planner 1. However, as it does not use any *a priori* knowledge of the search area, although yielding complete coverage, it is suboptimal for a time-critical missing person search task.

Some of the most simple informed search algorithms are Reference 19 LHC\_GW\_CONV (local hill climb, global warming, convolution) and Reference 20 greedy algorithm. Both superimpose a grid over the search area and use local search<sup>21</sup> to move to the adjacent cell with the best fitness, with each using tie-breaker algorithms in the case of two



FIGURE 1 Illustration of the path-planning tools in (A) QGroundControl and (B) Mission Planner.

cells having equal values. Reference 19 used convolution kernels to determine which direction is the most promising whereas Reference 20 proposes multiple solutions with LA\_MaxMax (a greedy 1-step look ahead algorithm) being the most successful. However, Reference 20 still ran into the greedy algorithm problem of getting stuck in a mode and not exploring other areas. This is where Reference 22 algorithm is superior as it introduces the *global warming* aspect which forces exploration of the area by artificially making probability more scarce. An issue with LHC\_GW\_CONV is that it does not consider the area as a whole. It is computationally fast but potentially misses regions of high interest due to the core of it being a greedy algorithm.

To explore the full search area,<sup>22</sup> proposes the use of hierarchical heuristics using mode goodness (MG) ratio which is based on Gaussian Mixture Models (GMM). Furthermore, this work<sup>22</sup> uses a glimpse factor to differentiate between areas where a sensor might struggle, whilst most other literature does not. This allowed<sup>22</sup> to more accurately tune their algorithm to what a real-world scenario might require. Ultimately, the TopN search algorithm was found to be profoundly better than the previously mentioned LHC\_GW\_CONV due to its ability to easily break free from modes. However, LHC\_GW\_CONV did beat both TopN and Top2 in certain scenarios. A reason for this could be that LHC\_GW\_CONV was tested using single or dual-mode PDMs, which it is well suited for, whilst<sup>19</sup> may have changed aspects of the algorithm at the time had their PDMs used more information about the search area. Nonetheless, Reference 22 algorithms only managed to reach maximum efficiency of 77.36% in their realistic testing leaving 22.64% left to explore. The efficiency was defined as a path's accumulated probability along it divided by the theoretical maximum found by *teleporting* between cells with the maximum probabilities.

Another method to break free from modes within a multi-modal area whilst potentially exploring closer to 100% of the area could be optimization algorithms. Reference 23 proposes the LRH-B for the search planning whilst<sup>20</sup> proposes partially observable Markov decision process (POMDP). LRH-B primarily focuses on being faster than standard methods (like Interior Point), whereas POMDP focuses on being more optimal. However, the computational complexity of either cannot be fairly compared as<sup>23</sup> used a 2D single robot whereas<sup>20</sup> used multiple quadrotors at various altitudes. As well as this,<sup>23</sup> does not compare their paths to ones of merit, nor do they use informed PDMs, so its performance cannot be determined. POMDP, on the other hand, was extensively compared to various permutations of the aforementioned greedy algorithms (which POMDP reliably outperformed) and used PDMs based on what a real-world SAR scenario might look like.

Whilst a single optimization algorithm might excel in a specific given search area, the work of Reference 24 combines 5 different standard optimization algorithms to create the hyper-heuristic method (HHM). This works on the principle that a single optimizer cannot be ideal for every case. Thus, HHM uses a fitness function to make the Low-Level Heuristics (LLH) compete. Wang compared 15 differently sized areas, and in each scenario, the HHM out-competed the 5 individual algorithms. In particular, the larger areas proved highly difficult for the LLHs with HHM performing 34% better than the next best LLH in the largest area.

It is easy to see that relative to parallel swaths, the more information available to an algorithm the better it performs. Reference 20 greedy algorithm was trumped by Reference 19 LHC\_GW\_CONV as the global warming aspect allowed it to explore more of the area. However, Reference 20 POMDP further outperforms the greedy algorithms due to its ability to see the map as a whole. Similarly, Reference 22 extracts further heuristics from the PDM and applies objectively simple search algorithms yet outperforms LHC\_GW\_CONV.

The optimal method to generate a heuristic—the PDM—is to fit a surface across the landscape based on large amounts of historical data. This would be the most representative of where a lost person may be found, resulting in a highly accurate map for path-planning algorithms and searchers to use. However, the reality is that SAR data is sparse at best with the most complete dataset being ISRID<sup>7</sup> at 50,000 entries from less than 20 countries. High-incident areas, such as Yosemite National Park, may have more dense datasets but in general, there is not enough high-quality data to do accurate curve fitting. Thus, Reference 25 developed MapScore which scores PDMs through heuristics. Similarly, Reference 26 used transfer learning to make smaller datasets usable based on non-SAR data. Another type of historical data is the general trends in the behavior of a lost person's movement. This can be in the form of locations found,<sup>1</sup> dispersion angle,<sup>7</sup> and so forth. This type of data is usually split into various profiles with different characteristics leading to highly customized PDMs for every incident.

Due to the lack of historical data in SAR, combining models is a common occurrence. A popular method to plan a search is the Mattson model<sup>27</sup> which uses a consensus across all searchers at hand to inform where to begin the search in an objective manner. Reference 28 generalizes on this concept by showing that combining models can yield better results due to the total information available increasing. As current SAR datasets do not typically incorporate psychological data due to the difficulty in fusing the environment/geographic data with psychological profiles, there is an opportunity to significantly increase the performance of path planning algorithms through the use of intelligently designed, high-fidelity PDMs.<sup>27,28</sup>

### 3 | SYSTEM MODELLING

As the purpose of this research is to assess algorithms for autonomously creating optimal trajectories for missing person search, realistic constraints must be applied to the optimization to ensure that any computed trajectories are feasible. Physical constraints require physical models and there are plenty of dynamic UAV models in the literature.<sup>12,14,29</sup> However, as the physical model is going to be placed inside of a numerical optimization cost function, model fidelity must be minimized to reduce computational overhead while retaining only the necessary physical constraints, that is, the physical model needs to be *fit-for-purpose*.

Figure 2 illustrates the typical drone operation for search and rescue missions. Quadcopter UAVs are nonholonomic requiring the aircraft to first pitch or roll to induce either forward or lateral motion in the  $x - y$  plane. Any platform attitude change will alter the sensor footprint if the sensor is rigidly mounted to the UAV base. Fortunately, most drones use a form of platform stabilization (usually a two-axis gimbal) to isolate the camera orientation from platform angular motion. Assuming the sensor is stabilized, this motion can be neglected. No specific sensor configuration or modality was assumed to ensure the generality of the results, apart from the sensor instantaneous field-of-view (IFoV) which is shown in Figure 2 as the angle  $\theta$ . The final assumption is that the drone maintains a constant altitude  $h$  throughout the mission.




FIGURE 2 UAV configuration showing sensor instantaneous field-of-view and footprint.

### 3.1 | UAV platform model

To ensure sufficient physical constraints are included, to minimize computational complexity and impose constant altitude, a simple holonomic 2D point-mass model is sufficient fidelity to represent the UAV platform from 1.

$$\begin{bmatrix} \ddot{x} \\ \ddot{y} \end{bmatrix} = \frac{1}{m} \begin{bmatrix} F_{Tx} - d\dot{x}^2 \\ F_{Ty} - d\dot{y}^2 \end{bmatrix}, \quad (1)$$

where  $d$  is the drag coefficient.  $F_T = [F_{Tx}, F_{Ty}]$  is the thrust acting on the UAV, and  $\mathbf{x} = [x, y]$  is the UAV's position vector. Drag is added to the system to match flight through the still air. An average flight speed of  $1.5 \text{ m s}^{-1}$  was used.

The model was controlled through a basic inverse-kinematics controller such that

$$\mathbf{F}_T = m\ddot{\mathbf{x}}_{des} + \tilde{\mathbf{F}}_D, \quad (2)$$

where  $\ddot{\mathbf{x}}_{des}$  is the desired acceleration (defined in Section 3.3) and  $\tilde{\mathbf{F}}_D$  is the estimated force due to drag on the model.

The resultant force input applied to the model was limited to prevent unrealistic behavior.

$$F_T = \begin{cases} F_T, & |F_T| \leq F_{T,\max} \\ F_{T,\max}, & |F_T| > F_{T,\max} \end{cases}, \quad (3)$$

where  $F_{T,\max}$  is the maximum force value that can be applied at any time step. For the use-case of PS ASU, an equivalent value of  $3 \text{ m s}^{-2}$  was used, however during testing it was found that the acceleration never exceeded  $0.42 \text{ m s}^{-2}$  regardless.

### 3.2 | Sensor characteristics

The UAV is assumed to fly at  $h$  m above a flat ground plane. A perfect sensor, that will detect a person with 100% accuracy, is mounted below the UAV with an unobstructed vision of the ground plane, and parallel to said ground plane. A typical UAV sensor is the DJI Zenmuse Z30, which has a Field Of View (FOV) range from  $2.3^\circ$  to  $63.7^\circ$ .

The value for  $b$ , the sensor footprint diameter, is then calculated by  $b = 2 \cdot h \cdot \tan \theta$ . Where  $b$  is the footprint diameter,  $h$  is the flight altitude and  $\theta$  is the FOV. It was assumed that the normal flight altitude is from  $h = 20$  m to  $h = 30$  m. These parameters are further defined in Figure 2. To simplify the sensor model a constant FOV of  $\theta = 31^\circ$  and a constant flight altitude of  $h = 25$  m where chosen with a resultant search footprint diameter of  $b = 30$  m.

### 3.3 | Trajectory generation

A quintic polynomial scheme<sup>30</sup> with low computational complexity boundary conditions was employed. These polynomials will give adequate estimations of a valid and realistic path for the point-mass model to follow. The position, velocity, and acceleration for the quintic polynomial scheme are defined by,

$$x = c_0 t^5 + c_1 t^4 + c_2 t^3 + c_3 t^2 + c_4 t + c_5, \quad (4)$$

where the coefficients  $\mathbf{c} = \{c_i | i \in [0, \dots, 5]\}$  are found by solving the system of equations at  $x_{t=0}$  and  $x_{t=T_n}$ .  $T_n$  is the finishing time of the line segment  $n$  as calculated by

$$T_n = \frac{\{\text{distance}\}}{\{\text{mean velocity}\}} \text{ [s]}, \quad (5)$$

$x_{t=0}$  and  $x_{t=T_n}$  are the starting and finishing waypoints, respectively. Furthermore,  $\dot{x}_{t=0} = 0 \text{ m s}^{-2}$  and  $\dot{x}_{t=T_n} = 0 \text{ m s}^{-2}$ .




FIGURE 3 Geometric layout of waypoints for 7.




FIGURE 4  $|\mathbf{v}_{g,+}|$  over  $\theta : \mathbb{R} \in [0, \pi]$  with  $n : \mathbb{N} \in [1, 6]$  and  $V = 2 \text{ m s}^{-1}$ .

For the velocity vector  $\mathbf{v}_g$  at each waypoint, the waypoints before and after a waypoint are used to calculate the boundary conditions for  $\dot{\mathbf{x}}$

$$\cos(\theta) = \frac{(-\mathbf{r}_{fg}) \cdot \mathbf{r}_{gh}}{|\mathbf{r}_{fg}| |\mathbf{r}_{gh}|} \quad (6)$$

$$\mathbf{v}_{g,-} = \mathbf{v}_{g,+} = V \cdot \frac{\mathbf{r}_{fh}}{|\mathbf{r}_{fh}|} \cdot \left( \frac{1 - \cos(\theta)}{2} \right)^n, \quad (7)$$

where  $V$  is the mean desired speed. Waypoints  $f$  and  $h$  are the waypoints before and after respectively with  $g$  being the current waypoint—as seen in Figures 3 and 4.

## 4 | TRAJECTORY OPTIMIZATION

### 4.1 | Optimization algorithms

Probability Accumulation Based optimization (PABO) is the concept of using probability accumulation over an area to create a path using numerical optimization. As with most topics, there is no clear best solution as<sup>24</sup> showed. For the case

of optimum path planning, optimization is taken to mean finding the set of waypoints in a problem space that optimizes various cost parameters. In this case, the cost function  $J(\mathbf{q})$ , where  $\mathbf{q}$  is the set of control inputs, is to be maximized to find the set optimal control inputs  $\mathbf{q}^*$ . This is achieved through the numerical solution of the following optimization problem

$$\mathbf{q}^* = \arg \max_{\mathbf{q}} \int_0^T p(q(t)) dt, \quad (8)$$

s.t.

$$T \leq \{\textit{endurance limit}\} \quad (9)$$

$$\mathbf{x}_{\min} \leq \mathbf{q} \leq \mathbf{x}_{\max}, \quad (10)$$

where  $\mathbf{q}^*$  is the optimal path with  $\mathbf{q} \in \mathbb{R}^m$ ,  $m$  is the number of 2D coordinates in path  $\mathbf{q}$ , and  $p(q(t))$  is the probability *seen* by the sensor at time  $t$ .  $T$  is the total path flight time, and  $n_{\text{waypoint}} = 10$  for the experiments conducted.  $\mathbf{x}_{\min} \in [x_{\min}, y_{\min}]$ ,  $\mathbf{x}_{\max} \in [x_{\max}, y_{\max}]$  creates a rectangular constraint around the area where  $p(x, y) > 0$ .

To find the solution to the optimization problem, and thus the *optimal* path, two optimization algorithms were selected due to their evolutionary nature: particle swarm optimization<sup>13,31-33</sup> (PSO), and genetic algorithm<sup>34-36</sup> (GA). For PSO, the particles within the swarm (the initial values of  $q(t)$ ) are randomly initialized and maintain a record of what their individual best score has been. To avoid convergence on local minima, PSO filters out identical particles and selects the top third for the next generation. The remaining particles are discarded and replaced with new particles in different and randomly-chosen dimensions. Similarly, GA undergoes a random mutation phase at each step allowing beneficial *traits* to be passed onto the next generation whilst the population continues to explore. This ability to incorporate random mutations in both algorithms encourages exploration of the global space which makes them well-suited for exploring the problem space at hand.

## 4.2 | Probability heatmap

The Centre for Search and Rescue regularly publishes a study on the behaviors of missing persons<sup>1</sup> which is used by the likes of Police Scotland and Northumberland National Park Mountain Rescue Team. This study analyses data gathered by multiple Mountain Rescue Teams (MRT) and segments it into useful categories such as dementia, despondent (defined as “*any person who is thought to have disappeared deliberately*”<sup>1</sup>), and developmental problems. Within these categories, there are common statistics analyzed such as location found.

The “location found” is of particular interest as it gives a powerful insight into the search area (this search area is predefined by other means and how the area was determined is not important). Table 1 gives an understanding into how the data is organized. For a person with dementia, it is clear that they are more likely to be near a travel aid (i.e., road) than a body of water. This allows us to draw a PDM for an area with travel aids and water focusing the search on these travel aids, then water.

In order to do this, one can represent a higher likelihood of a person being at  $(x_1, y_1)$  than at  $(x_2, y_2)$  by placing a bivariate Gaussian at  $(x_1, y_1)$  such that

$$p(x, y, \sigma) = \frac{1}{2\pi\sigma^2} \exp \left[ \frac{-(x - x_1)^2 + (y - y_1)^2}{2\sigma^2} \right], \quad (11)$$

where  $p(x, y, \sigma)$  is the probability of finding a person at  $(x, y)$ . This is a continuous probability function and would yield a different probability for any  $x, y \in \mathbb{R}$  which is favorable for numerical optimization methods. This can then be discretized into a probability distribution map (PDM) for performance.

However, using a PDM for a numeric optimization problem poses some issues. primarily, numerical optimization methods rely on small changes to the input giving a change in the cost function. If these small changes are carried out on a discrete probability map, there may be no change in cost function and thus the optimization algorithm assumes a minimum is found, exiting prematurely. A solution is to use a continuous gaussian function to represent hotspots. This would work with a few locations and associated radii, however, once the problem is applied to the real world this becomes tedious and difficult to set up. A method of having a local expert (mountain rescue leader, police search leader, etc.) using

TABLE 1 Location found by terrain and gender for a person with dementia.<sup>1</sup>

Location	Farmland							
	Female		Male		Urban		Other	
	<i>n</i>	%	<i>n</i>	%	<i>n</i>	%	<i>n</i>	%
Travel aid	8	42	12	29	35	36	1	4
Building	5	26	8	20	29	30	2	7
Linear feature	2	11	12	29	17	17	5	21
Open ground	1	5	7	17	11	11	8	29
Trees	3	16	1	2	1	1	7	21
Water	0	0	1	2	4	4	5	18
Not recorded	0	0	0	0	1	1	0	0

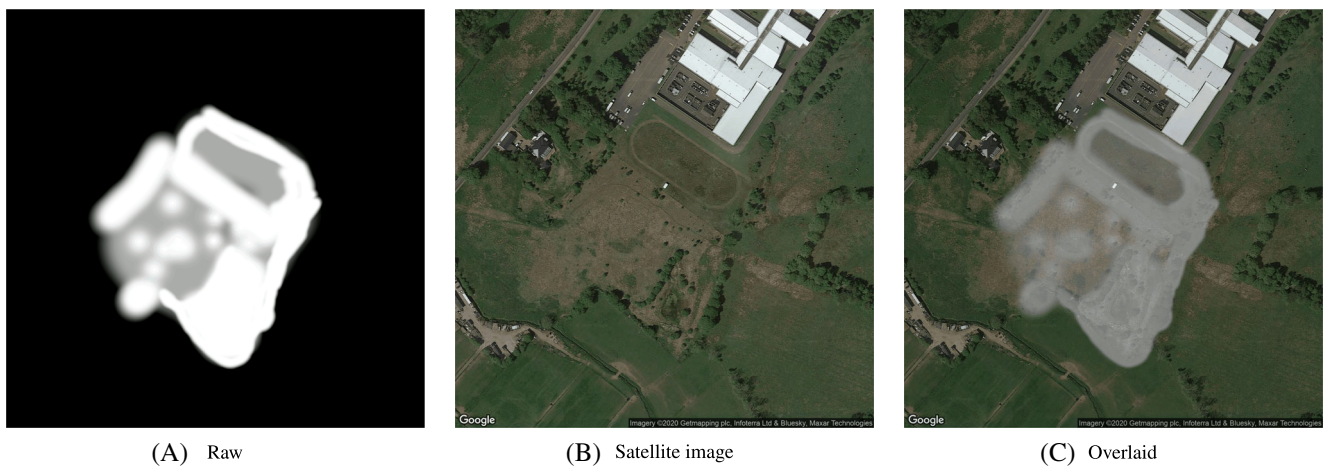


FIGURE 5 Example of how the brush method might be used out in the field by an expert to create a PDM.

*brushes*, akin to digital drawing, to manually draw hotspots on a map is more intuitive and one can imagine someone in the field with a tablet taking seconds to create a map. An example of this can be seen in Figure 5.

To solve this problem, it is possible to extrapolate a pseudo-continuous function from the discrete probability map. This is done by taking the theoretical search path and moving a search radius along it. All grid squares within this radius are completely seen (100% of the *grid square* is within the *search area*), but those at the edge will be partially seen (< 100% of the grid square is within the search area). Figure 6 shows a grid square within the  $b/2$  search radius of the path being partially seen. The *intersection polygon* is the polygon created by the common intersection area of the grid square and the search area. Its area is then normalized and the probability from this square is  $A_{int} \cdot P_{x,y}$  where  $A_{int}$  is the area of the intersection polygon and  $P_{x,y}$  is the probability of the grid square at  $(x, y)$ . This is repeated for all grid squares intersecting with the edge of the search area.

## 5 | RESULTS

### 5.1 | Experimental setup

To test the effectiveness of each algorithm, a Monte Carlo simulation was run over multiple generated paths to gather data on the time to find. Lost persons were then placed over the abstract search area using the generated PDM as seen in Figure 7.



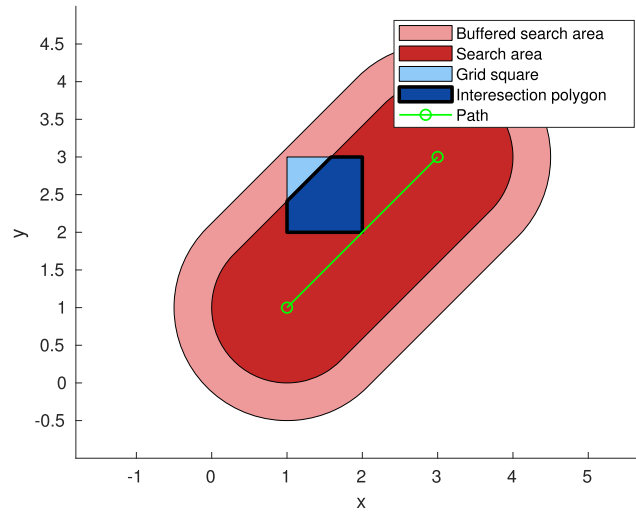


FIGURE 6 Pseudo-continuous probability map technique.




FIGURE 7 Using a PDM to place a reduced number (200) of persons.

Once the LPs have been placed, the distance comparison  $|\mathbf{p}_i - \mathbf{x}| < \frac{b}{2}$  was used to determine if the  $i$ th LP's position  $\mathbf{p}_i$  was within the search range of the UAV at position  $\mathbf{x}$ .

## 5.2 | Comparison against figures-of-merit

There were two PDMs used for the comparison of the algorithms—a 300 m × 300 m area and a 700 m × 700 m area (Figure 8A and Figure 8B respectively). The larger area is used to give an insight into the endurance capabilities of the algorithm. In this section, Figure 8A and Figure 8B will be referred to as PDM A and PDM B respectively.

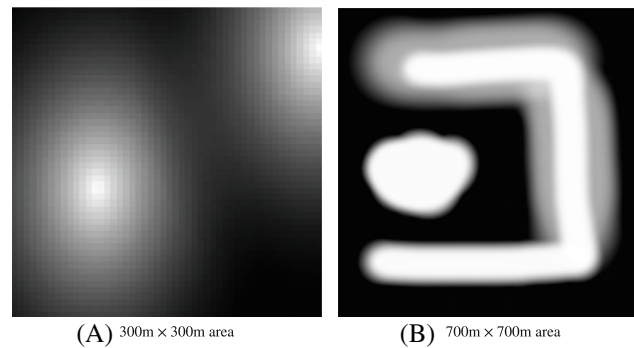


FIGURE 8 The PDMs used for this project.

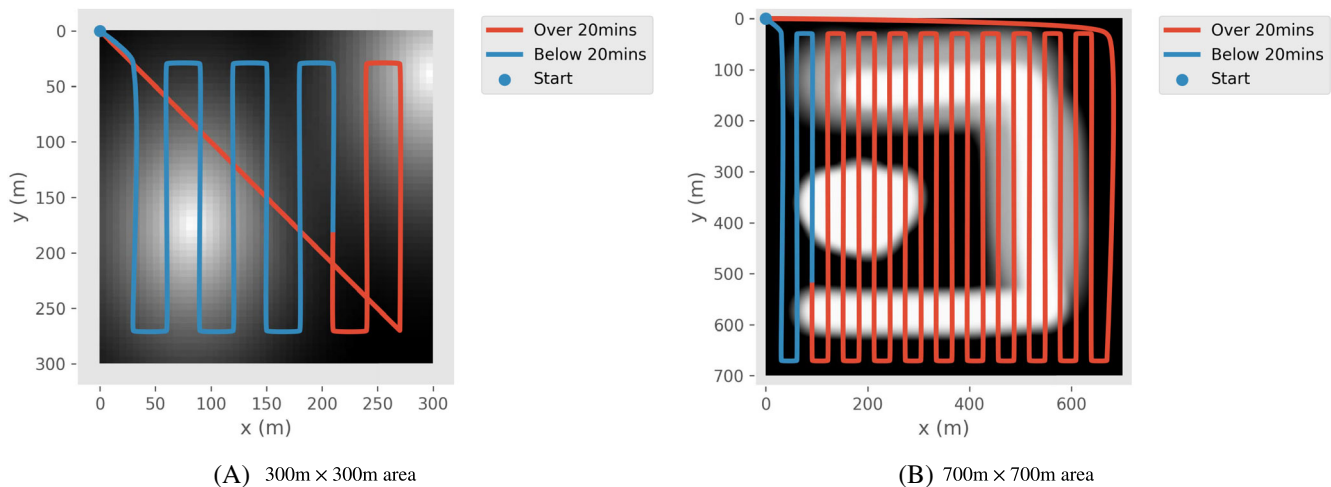


FIGURE 9 Simulated parallel swaths.

The parallel swaths path can be seen in Figure 9A. The back-and-forth path covers the whole left area of PDM A. However, due to its inability to use a priori information, the parallel swaths do not cover the right region of higher probability before the endurance limit. As well as this, the algorithm spends a substantial amount of time in the bottom left region of very low probability.

Looking at Figure 9B, it can be seen that the first sweep of PDM B completely misses all higher probability regions. The second sweep is not much better, barely gathering any probability from any of the high-probability regions.

It is easy to see that this algorithm is great for complete coverage, if either endurance is not a problem or multiple flights are possible. However, for the use case in SAR, this algorithm is unusable due to its inability to use the information at hand.

The genetic algorithm works very well for the PDM A. GA also has long parallel lines (from (50,225) to (250, 50) in Figure 10A) showing the probability accumulation aspect of the cost function in action. The path traversed in PDM A avoids low-probability regions very effectively, and it targets the regions of high probability quickly.

In Figure 10B, GA can be seen flying out to (650,250) whilst covering a lot of the high probability area and then returning home. Most importantly, it adheres to the endurance constraint.

The particle swarm path yields similar results to GA. Particle swarm spends too much time in the low probability regions of PDM A (Figure 9A). However, it does not stick to the PDM limits like GA. Furthermore, particle swarm sticks to the high probability regions very effectively in PDM A. Nonetheless, it does not adhere to the endurance constraint nearly as well as GA for PDM A.

Like GA, the PDM B performance of particle swarm is very targeted to high probability areas. It may not return home exactly within the endurance limit, but it follows a realistic path (Figure 11).




FIGURE 10 PABO GA.

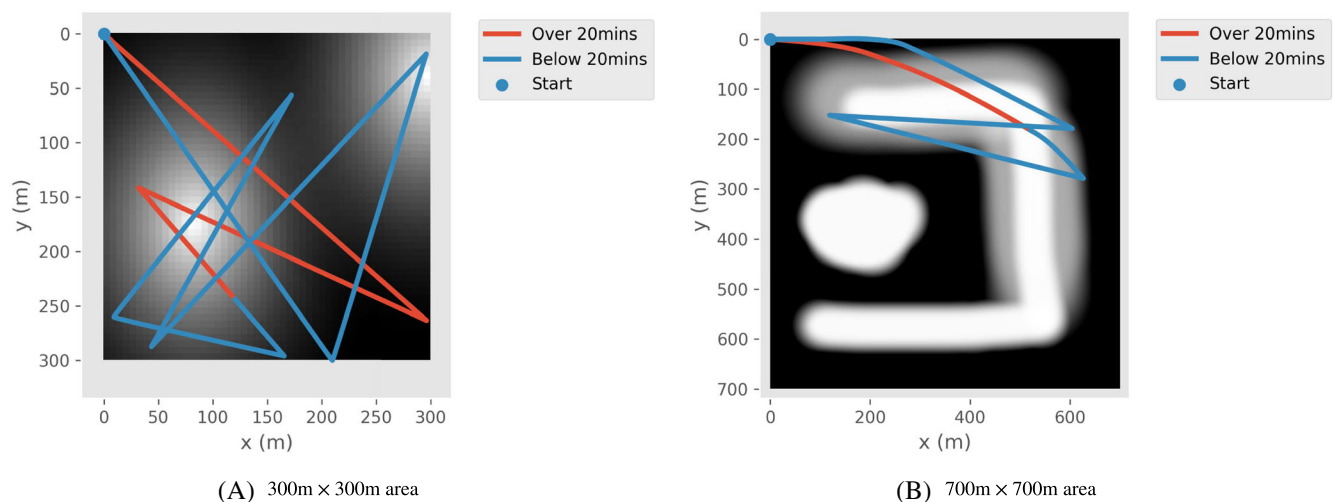


FIGURE 11 PABO particle swarm.

The below data was gathered by creating new paths for each PDM multiple times and combining the time-to-find data to give a general insight into the performance. This is preferential over a single attempt, as it prevents a single excellent path for an algorithm to influence the conclusion. SAR object placements were random given the PDM and thence all algorithms were treated the same.

Inspecting Figure 12A, it can be seen that parallel swaths find the least amount of people at any given time than the other algorithms, up until the very end. This is expected as it gives almost perfect coverage of the area but it does not use the *a priori* information of the PDM like the optimization-based algorithms.

Similarly, for PDM B in Figure 12B, the parallel swaths algorithm is outperformed by all others. The optimization-based algorithms accumulate more probability at all times.

From these results, it is evident that increasing the information available to a search planning algorithm massively increases its performance. The paths over PDM B in particular show this as the parallel swaths will search all areas no matter how irrelevant they are.

### 5.3 | Comparison of PABO against human pilots

PS ASU performed experimental flights to collect data for analysis. In total 3 test flights were performed for this project: two were piloted by trained officers based on their training and experience, and one attempted to follow a parallel swaths

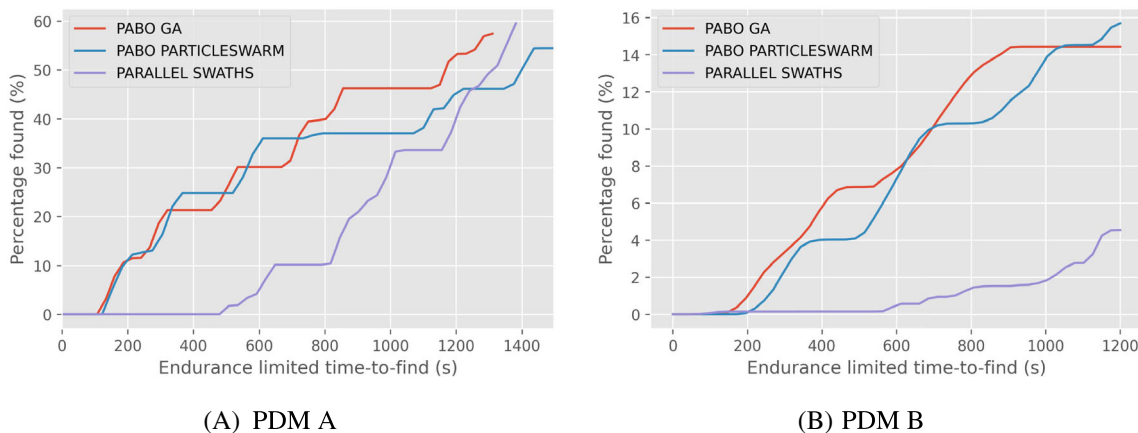


FIGURE 12 Percentage found up until time.



FIGURE 13 PDM C.

path. These will be referred to as *PS ASU 1*, *PS ASU 2*, and *Attempted Parallel Swaths* henceforth. The area used for these flights can be seen in Figure 13A and the flights were designed to be as close to a real-case SAR scenario as possible in a simulated environment.

PABO GA was selected to be pitched against the pilots as it outperformed the particle swarms optimization across both tests. The simulated flight for PABO GA was given an input value for an average flight speed of  $1.2\text{m s}^{-1}$  which was derived from the provided flight data. The PDM for this area can be seen in Figure 13B where emphasis was put on the tree lines, the dry pond, and linear features as per Table 1. Note that PS ASU could not fly out with their grounds which limited the search area. The PDM was created with this in mind.

Figure 14 shows the aforementioned paths, as well as PABO GA, superimposed onto Figure 13A along with the PDM associated with said area. From this it can be seen that PS ASU 1 and PABO GA share a common take-off location (this was intentional), however the others take-off around 100m further east, and 50m further south. This should not affect performance but will be noted nonetheless.

Inspecting PS ASU 1, it can be seen that the path follows the edge of the PDM. PS ASU 2 seemed to find the bottom-right of the PDM very interesting which corresponds to a grouping of trees, something that the CSR outlines as being an area of higher probability for finding missing persons.<sup>1</sup> Looking at PABO GA, it is clear that a major problem with down-sampling



FIGURE 14 Police Scotland and PABO GA paths.

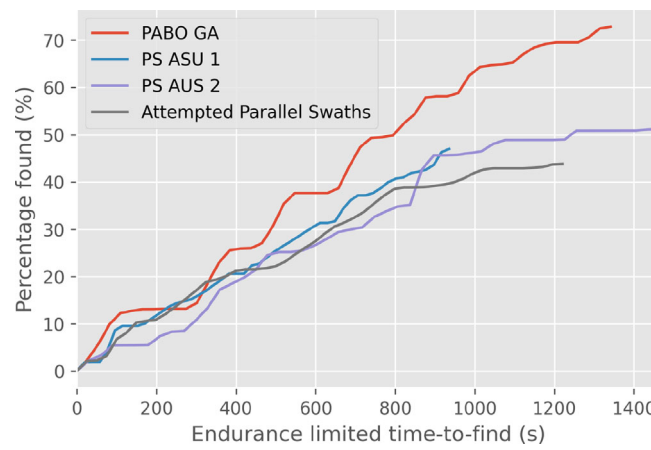


FIGURE 15 Percentage found up until time.

the PDM to  $b$  sized pixels is that the cost function believes it is in an area of higher probability, but once converted back into a higher resolution PDM it is grossly off the mark. This can be fixed by not down-sampling the PDM but would require much longer computing times or higher computational power. However, ignoring this, the path does explore the area very well. Focusing on locations like the group of trees in the bottom-right similarly to PS ASU 2.

From Figure 15 the effectiveness of the paths at finding a person can be seen. PABO GA performs well from  $t = 0$ s by staying ahead of the real-world paths for almost the whole flight time, reaching a maximum of  $\approx 73\%$ .

## 6 | CONCLUSION AND FUTURE WORK

A benchmark search algorithm and two cost-function-based algorithms were implemented. A 2D point-mass model was used along with a quintic polynomial trajectory generation scheme to simulate a SAR flight. Two different PDMs were created to compare the three algorithms through Monte Carlo simulations.

It was shown that the PABO GA was found to be the most effective algorithm overall. Parallel swaths, the benchmark, were by far the least effective overall PDMs. PS ASU officers conducted multiple flights to which PABO GA was compared. The PS ASU officers are highly trained in SAR and piloted the UAV similar to what the PDM created from the CSR report<sup>1</sup> would predict. This resulted in PABO GA outperforming all PS ASU SAR scenario flight paths.

Overall, this study has shown the viability of merging multiple information streams that are available to an optimization-based search algorithm to greatly increase its performance over an area. The psychological profile of an LP is a core part of this information, as is the physical landscape at hand. With a simple algorithm like PABO GA outperforming trained experts, it is clear that the future of SAR is data-driven.

This work is a keystone result to build from. Future work will explore automatically generating a PDM based on psychological, and geographical information. This has a potential application for machine learning which will also need to be explored.

As well as increasing the state of the art of PDM generation using ML, the application of a Deep Q-Network will be explored to create a generalized search planning agent for all PDMs. If the PDM generation is advanced enough, reinforcement learning will have the massive amounts of data it needs to closely predict a lost person's intent.

## ACKNOWLEDGMENTS

I thank PC Stuart Neilson and PC Steven Preece of the Police Scotland Air Support Unit for their generous help in providing insight into SAR and performing real-world flight testing. It would have been impossible to get this sort of data as a student otherwise during the COVID-19 pandemic.

## DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.

## ORCID

Jan-Hendrik Ewers  <https://orcid.org/0000-0002-8498-1410>

David Anderson  <https://orcid.org/0000-0001-6301-9532>

Douglas Thomson  <https://orcid.org/0000-0003-3588-5710>

## REFERENCES

- Perkins D, Roberts P, Feeney G, Mrt P. *The U.K. Missing Person Behaviour Study*. CSR; 2011.
- Skeleton L, Smith C. Remotely Piloted Aircraft System (RPAS) Evaluation Report. Tech. Rep., Scottish Police Authority; SPA Policing Performance Committee. 2020.
- Stevenson O, Woolnough P, Parr H. Geographies of Missing People: Improving Police Education and Response to Missing Persons. Tech. Rep., Scottish Institute for Policing Research; Edinburgh Napier University. 2014.
- Police G. Missing Persons: Understanding, Planning, Responding. Tech. Rep., Grampian Police; 4 Burnside St, Rothes, Aberlour AB38 7AA. 2007.
- UK Missing Persons Unit. iFIND. 2016.
- Heth CD, Cornell EH. Characteristics of Travel by Persons Lost in Albertan Wilderness Areas. *J Environ Psychol*. 1998;18(3):223-235. doi:10.1006/jevp.1998.0093
- Koester RJ. *Lost Person Behavior: A Search and Rescue Guide on Where to Look for Land, Air, and Water*. dbS Productions; 2008.
- LaValle SM. Rapidly-exploring random trees: A new tool for path planning. *Annu Res Rep*. 1998. <https://msl.cs.illinois.edu/~lavalle/papers/Lav98c.pdf>
- Hall J, Anderson D. Reactive route selection from pre-calculated trajectories – application to micro-UAV path planning. *Aeronaut J*. 2011;115(1172):635-640. doi:10.1017/S0001924000006321
- Girard A, Howell A, Hedrick J. *Border Patrol and Surveillance Missions Using Multiple Unmanned Air Vehicles*. Vol 1. IEEE; 2004:620-625.
- Weldon WT, Hupy J. Investigating methods for integrating unmanned aerial systems in search and rescue operations. *Drones*. 2020;4(3):38. doi:10.3390/drones4030038
- Beard RW, McLain TW. *Small Unmanned Aircraft: Theory and Practice*. Princeton University Press; 2012.
- Brown A, Anderson D. Trajectory optimization for high-altitude long-endurance UAV maritime radar surveillance. *IEEE Trans Aerosp Electron Syst*. 2020;56(3):2406-2421. doi:10.1109/TAES.2019.2949384
- Tsourdos A, White B, Shanmugavel M. *Cooperative Path Planning of Unmanned Aerial Vehicles*. 1st ed. Wiley; 2010.
- Ireland ML, Anderson D. Optimisation of trajectories for wireless power transmission to a quadrotor aerial robot. *J Intell Robot Syst*. 2019;95(2):567-584. doi:10.1007/s10846-018-0824-6
- Arkin EM, Fekete SP, Mitchell JS. Approximation algorithms for lawn mowing and milling. *Comput Geom*. 2000;17(1-2):25-50. doi:10.1016/S0925-7721(00)00015-8
- Koester R, Cooper DC, Frost JR, Robe RQ. Sweep Width Estimation for Ground Search and Rescue. 2004.
- Goodrich MA, Morse BS, Gerhardt D, et al. Supporting wilderness search and rescue using a camera-equipped mini UAV: UAV-enabled WiSAR. *J Field Robotics*. 2008;25(1-2):89-110. doi:10.1002/rob.20226
- Lin L, Goodrich MA. UAV intelligent path planning for wilderness search and rescue. *IEEE/RSJ Int Conf Intell Robots Syst*. 2009;0(1):709-714. doi:10.1109/IROS.2009.5354455
- Waharte S, Trigoni N. *Supporting Search and Rescue Operations with UAVs*. University of Oxford; 2010:142-147.
- Skiena SS. *The Algorithm Design Manual*. Springer London; 2008.
- Lin L, Goodrich MA. Hierarchical heuristic search using a Gaussian mixture model for UAV coverage planning. *IEEE Trans Cybern*. 2014;44(12):2532-2544. doi:10.1109/TCYB.2014.2309898
- Subramanian A, Alimo R, Bewley T, Gill P. *A Probabilistic Path Planning Framework for Optimizing Feasible Trajectories of Autonomous Search Vehicles Leveraging the Projected-Search Reduced Hessian Method*. AIAA. American Institute of Aeronautics and Astronautics; 2020.
- Wang Y, Zhang MX, Zheng YJ. A hyper-heuristic method for UAV search planning. In: Tan Y, Takagi H, Shi Y, Niu B, eds. *Advances in Swarm Intelligence*. Vol 10386. Springer International Publishing; 2017:454-464.

25. Sava E, Twardy C, Koester R, Sonwalkar M. Evaluating lost person behavior models. *Trans GIS*. 2016;20(1):38-53. doi:10.1111/tgis.12143
26. Šerić L, Pinjušić T, Topić K, Blažević T. Lost person search area prediction based on regression and transfer learning models. *IJGI*. 2021;10(2):80. doi:10.3390/ijgi10020080
27. Mattson RJ. Establishing Search Priorities. In: Haley KB, Stone LD, eds. *Search Theory and Applications NATO Conference Series*. Springer US; 1980:93-97.
28. Surowiecki J. *The Wisdom of Crowds: Why the Many Are Smarter than the Few*. Abacus. repr ed; 2014.
29. Vargas A, Ireland M, Anderson D. *System Identification of Multi-Rotor UAV's Using Echo State Networks*. University of Glasgow; 2015.
30. Biagiotti L, Melchiorri C. *Trajectory Planning for Automatic Machines and Robots*. Springer; 2009.
31. Shi Y, Eberhart R. *A Modified Particle Swarm Optimizer*. IEEE; 1998:69-73.
32. Kiranyaz S, Ince T, Gabbouj M. *Particle Swarm Optimization*. Vol 15. Springer Berlin Heidelberg; 2013:45-82.
33. Phung MD, Quach CH, Dinh TH, Ha Q. Enhanced discrete particle swarm optimization path planning for UAV vision-based surface inspection. *Autom Constr*. 2017;81:25-33. doi:10.1016/j.autcon.2017.04.013
34. Mitchell M. *An Introduction to Genetic Algorithms*. The MIT Press; 1998.
35. Whitley D. A genetic algorithm tutorial. *Stat Comput*. 1994;4(2):65-85. doi:10.1007/BF00175354
36. Hong L, Wang Y, Du Y, Chen X, Zheng Y. UAV search-and-rescue planning using an adaptive memetic algorithm. *Front Inform Technol Electron Eng*. 2011;22(11):1477-1491. doi:10.1631/FITEE.2000632

## AUTHOR BIOGRAPHIES




**Jan-Hendrik Ewers** is pursuing a PhD in Machine Learning Driven Search and Rescue at the University of Glasgow. He graduated with a Masters of Engineering with Honours of the First Class in Aerospace Systems from the same university in 2021. Mr Ewers has practical industry experience, having interned at BAE Systems Air in 2019 and worked as a systems engineer at Gibson Robotics from 2020 to 2021. In 2021, he received the British Aerospace Engineering Systems Prize for the best industrially relevant final year project. Throughout his undergraduate degree at the University of Glasgow, he was placed on the Engineering Excellence List every year.



**David Anderson** was awarded a first class honours degree in aeronautical engineering from the University of Glasgow in 1994. He then completed his PhD study (thesis entitled “active control of turbulence induced helicopter vibration”), also from Glasgow University in 1997. Following this he was employed as a senior sightline control engineer by SELEX Galileo Sensors & Airborne Systems (then BAE SYSTEMS) in Edinburgh where he designed the precision pointing & stabilisation controllers on a number of airborne electro-optic systems. In 2001 he moved to industrial systems & control to take a position as a principal technical consultant. This role was principally one of technology transfer with Dr Anderson providing research support to the Industrial Control Centre at the University of Strathclyde and then “re-interpreting” these results for industrial

clients in the aerospace, defence and automotive sectors. In 2004 he re-joined the University of Glasgow as a lecturer in aerospace systems and was promoted to senior lecturer in 2011.



**Douglas Thomson** graduated from the University of Glasgow in 1983 with a 1st Class Honours Degree in Aeronautical Engineering. He then obtained a PhD in Aeronautical Engineering from Glasgow. After 2 years as a research assistant funded by the Royal Aerospace Establishment, he obtained a Royal Society University Research Fellowship allowing him to conduct independent research in the area of inverse simulation. In 1993 he joined the Department of Aerospace Engineering as senior lecturer in flight dynamics, and in 2000 began a period of 7 years as Head of this Department. In 2010, after the restructuring of the Faculty of Engineering into the School of Engineering, Dr Thomson was appointed as Head of the Aerospace Sciences Research Division. He is also currently

the Chief Adviser of Studies (Engineering) within the School.

**How to cite this article:** Ewers J-H, Anderson D, Thomson D. Optimal path planning using psychological profiling in drone-assisted missing person search. *Adv Control Appl*. 2023;e167. doi: 10.1002/adc2.167