

Trajectory Optimization of Unmanned Aerial Vehicles in the Electromagnetic Environment

Anvarbek Atayev*, Jörg Fliege† and Alain Zemkoho‡

June 2, 2024

Abstract

We consider a type of routing problems common in defence and security, in which we control a fleet of unmanned aerial vehicles (UAVs) that have to reach one or more target locations without being detected by an adversary. Detection can be carried out by a variety of sensors (radio receivers, cameras, personnel, etc) placed by the adversary around the target sites. We model the act of detecting a UAV from first principles by noting that sensors work by monitoring frequencies in the electromagnetic spectrum for signals or noise emitted. By this, we are able to provide a flexible and versatile nonlinear optimisation framework in which the problem is modeled as a novel trajectory optimisation problem with paths of the UAVs as continuous arcs in an Euclidean space. The flexibility of our approach is exhibited by the fact that we can easily consider various relevant objectives, among them minimising the overall probability of detection and maximising the location error that the adversary experiences when trying to locate our UAVs. Our model is also versatile enough to consider the act of *jamming*, in which one or more of our UAVs intentionally send out signals to interfere with the operations of the adversary's sensors. Numerical results show the flexibility of our framework, and that we can solve realistic instances of this problem type.

Keywords: Nonlinear Programming, OR in Defence, trajectory optimisation

1 Introduction

We consider the problem of one or more intruders that have to reach a certain site or number of sites without being detected. In this problem, sensors of various types (e. g. motion detectors, antenna, cameras, or personnel) are often used to surround the sites of interest to detect incoming vehicles or personnel. In this article we focus on unmanned aerial vehicles (UAVs) as intruders and will provide a novel mathematical programming model that is computationally tractable using state of the art optimisation tools. In particular, we will provide a novel model for the problem at hand by way of a trajectory optimisation problem in which the path used by any UAV is modelled by a continuous arc in two- or three-dimensional space.

We model the detection of a UAV from first principles by considering sensors monitoring particular frequencies in the electromagnetic domain for a potential signal emitted by the UAV. (Here, we consider both actual communication signals as well as noise from the UAV as a signal that a sensor might pick up.) In particular, we consider both the probability of detection and an error measure on the location of a UAV as well-defined quantities, the former to be minimised and the latter to be maximised by the UAV operator. We also allow particular UAVs to act as *jammers*, i. e. they temporarily degrade some sensor capabilities by creating noise and interference in relevant parts of the electromagnetic spectrum. As far as we are aware, this is the first time corresponding trajectory optimisation problems of this form have been considered.

We further allow UAVs to temporarily reduce the signal they are emitting. Accordingly, our model allows to *jointly optimise* not only trajectories but other relevant operational decisions as well. Furthermore, we

*University of Warwick, Mathematics Institute, Coventry, CV4 7AL, UK. anvar.atayev@warwick.ac.uk

†University of Southampton, Mathematical Sciences, Southampton, SO17 1BJ, UK. J.Fliege@soton.ac.uk

‡University of Southampton, Mathematical Sciences, Southampton, SO17 1BJ, UK. A.B.Zemkoho@soton.ac.uk

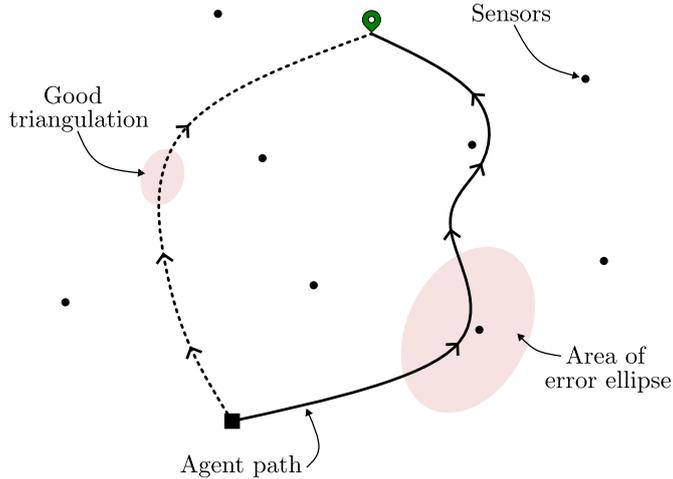


Figure 1: Choice of optimal path with respect to area of error ellipse. Solid line depicts path of minimal cumulative area of error ellipse where sensors have poorly approximated agent position. Dashed line depicts path where the sensors have obtained good triangulation of the agent position along its path, resulting in a small cumulative area of error ellipse.

investigate how to incorporate uncertainty into our computational framework under realistic conditions and provide numerical results illustrating the efficacy of our approach.

Finding optimal UAV trajectories under the risk of detection or location has recently garnered significant attention. Ruz et al. [23] describe a mixed-integer linear formulation for a trajectory optimization problem with path constraints that exclude zones of high risk of detection by radar. Pelosi et al. [21] consider a similar setup in which the UAV maximises any terrain-masking opportunities to avoid detection, while Zhang et al. [30, 31] minimises the cumulative radar cross section of the UAV. A more general approach is taken by Jun and D’Andrea [16], who use a map of the probability of threats for their path planning algorithm. All of these formulations have all nonlinear terms fully linearized, and the magnitude of the linearization error is not discussed.

Kabamba et al. [17] derive optimality conditions for a trajectory optimization problem in which one seeks to minimise the probability of being tracked by a radar. Hui et al. [15] provide a framework of joint optimization of target tracking and threat avoidance. Avoiding adversary sensors can also be seen as a case of an obstacle avoidance problem, see e. g. [7, 9, 24]. Joint optimization of trajectories and other decision has been considered only very recently; trajectories and signal strengths are considered by Zhou et al. [34], while joint strategies for trajectory optimization and jamming resp. anti-jamming have been investigated by a number of authors, see e. g. [10, 28, 29, 32, 33].

The rest of this article is structured as follows.

Section 2 describes the basic setting for the infinite-dimensional optimisation problem: we provide the underlying notation, the decision variables, the underlying vector space of the variables, and some basic constraints. Several objective functions that are relevant to the decision maker are then derived in Section 2.1. In particular, we provide complete model as an example in Section 2.4. We then discuss a variety of extensions to the basic model template in Section 3: jamming (both isotropic and anisotropic) in Section 3.1, intermittent communication in Section 3.2, and the case of several agents in Section 3.3.

In Section 4, we provide means for a numerical solution for the problem described in Sections 2–3. Note that the model template and all its extensions discussed in Section 2 is an infinite-dimensional one (we seek solutions $\mathbf{x}(t)$ for an infinite number of times t). How to appropriately discretize a given problem of this form to make it amenable to standard optimisation tools is discussed in Section 4.1. In Section 4.2 we depict various two-dimensional ‘landscapes’ in which the UAV(s) has to move. This provides for an intuitive understanding of the problem at hand. Trade-offs between different objectives are visualised in Section 4.3, while various model extensions and variants are discussed and visualised in Section 4.4. We conclude the numerical experiments with Section 4.5 where we present a numerical solution to a

simplified version of a possibly real-world scenario of an agent traversing an adversary environment with a multitude of sensors, one of which is mobile. More specifically, we also provide an example solution where the rolling horizon approach is used to take into account uncertain environment information, such as the position of sensors.

We finish this article with our conclusions, see Section 5.

2 Model development

Consider a set of sensors surrounding a site of interest, called the *target*, and a collection of UAVs or *agents* that wish to reach said target in a certain period of time. We will show from first principles that each sensor can detect a UAV with a certain probability depending on the electromagnetic emissions of the UAV and the distance between sensor and vehicle. The identification of the agent is contingent on a sufficient level of detection by one or more sensors. Furthermore, we allow for some agents to be able to perform disabling/disrupting interventions to sensors along their path to the target, whereby an agent is able to disrupt sensors for a period of time, and performing such an intervention incurs an equipment and time cost to the agent concerned.

In this section we outline an infinite-dimensional mathematical optimisation model that can be used to obtain a solution to the optimal path problem described in Section 1.

In most of what follows, we assume that the working domain is two dimensional. Considering the two-dimensional case is appropriate if all sensors are near the surface of the Earth and the UAVs stay close enough to the surface that the curvature of the Earth's surface can be neglected. It will be seen in the below that it is straightforward to extend the problem to the three-dimensional case.

For ease of notation, we will start by introducing constraints, decision variables, and objectives for the single-UAV case first. It will be seen that extending the problem formulation to the case of multiple UAVs is straightforward. Throughout this article, we fix the following notation and basic assumptions:

- Let $\tau > 0$ denote the time of flight of the agent and $T > 0$ representing the maximum time of flight, i. e. $0 \leq \tau \leq T$;
- For the given UAV, prescribe to it a parametrisation of its path via a continuous mapping $x : [0, \tau] \rightarrow \mathbb{R}^2$ where $[0, \tau]$ with $\tau \leq T$ denotes the time interval through which the UAV is travelling;
- We assume that the UAV start from a given starting point $x(0) = x_0$;
- Assume that all flight paths are restricted to lie within a bounded domain $\Omega \subset \mathbb{R}^2$, i. e. such that $x(t) \in \Omega$ for all $t \in [0, \tau]$;
- We have been given a target location $X \in \Omega$;
- Denote $y_i : S_T \rightarrow \mathbb{R}^2$ for $i = 1, \dots, n$ the locations of a collection of n sensors over a bounded time interval $S_T \subset \mathbb{R}$ with $[0, T] \subset S_T$. In the case where the sensors are stationary, we simply write $y_i = y_i(t)$ for all $t \in S_T$.

The aim is to obtain *optimal* flight paths for the given UAV from the initial position to the target position whilst minimising one or more objectives functions. For a given UAV, we aim to determine an optimal path $x : [0, \tau] \rightarrow \mathbb{R}^2$ such that $x(0) = x_0$ and $x(\tau) = X$, where, $\tau > 0$ denotes the time of flight of the agent. We assume throughout that τ is a decision variable that satisfies the maximum time of flight constraint $\tau \leq T$.

For a given $\tau \in [0, T]$, define by $\mathcal{A}_\tau(x_0; X)$ the space of twice continuously differentiable admissible paths $x : [0, \tau] \rightarrow \mathbb{R}^2$ with $x(0) = x_0$ and $x(\tau) = X$, i. e.,

$$\mathcal{A}_\tau(x_0; X) := \left\{ x : [0, \tau] \rightarrow \Omega \left| \begin{array}{l} x(0) = x_0 \\ x(\tau) = X \\ x \in \mathcal{C}^2([0, \tau]) \end{array} \right. \right\}.$$

Furthermore, for $x \in \mathcal{A}_\tau$, we have given bounds for the velocity

$$\|x'\|_\infty \leq v_{max}$$

and acceleration

$$\|x''\|_\infty \leq a_{max}$$

via some user defined constants $v_{max} > 0$ and $a_{max} > 0$, respectively, that depend on the characteristics of the engine of the UAV.

Then, the set of all our decision variables associated with a particular UAV can be written as

$$\mathcal{D} := \bigcup_{\tau \in [0, T]} \mathcal{A}_\tau(x_0, X). \quad (1)$$

This decision space is expanded upon in Sections 3.1, 3.2, and 3.3 where we introduce jamming, intermittent communication, and multiple agents, respectively.

For brevity, throughout this article we will use the notation $\mathbf{x} := x \in \mathcal{D}$, with $x : [0, \tau] \rightarrow \mathbb{R}^2$, and $\mathbf{y} := (y_1, \dots, y_n) : [0, \tau] \rightarrow \mathbb{R}^2$.

Our aim is to obtain a path for our UAV which attempts to optimally *evade* detection when travelling through an adversary environment of sensors. In addition, we also want to minimize the flight time to the target.

In particular, given a collection of sensors $\mathbf{y} = (y_1, \dots, y_n) : S_T \rightarrow \mathbb{R}^{2n}$, we want to consider some objective F that represents how well the agent "avoids" the sensors, and in the most basic scenario we want to solve a problem of the form

$$\min_{\mathbf{x} \in \mathcal{D}} F(\mathbf{x}, \mathbf{y}) \quad (2)$$

to determine an optimal path $\mathbf{x} \in \mathcal{D}$. In more complicated scenarios we also want to decide on possible jamming procedures, and variable intervals of active lines of communication. We will discuss corresponding extensions in Section 3.

2.1 General objective function framework

In what follows we will derive several novel choices of such an objective F that can all serve to compute flight paths for our UAVs that avoid being detected. Thus, in essence, our setting is one of *multiobjective optimization*. Note however that it will turn out that not all our objectives are 'in conflict' with each other; some will simply measure slightly different characteristics of a given flight path or the way sensors try to detect a UAV. It is also important to note here that for all times t , we know the location of our UAV $x(t)$ and we also know all the locations of the adversary sensors $y_j(t)$ ($j = 1, \dots, n$). The first assumption is not necessarily a strong one: many UAVs are equipped with GPS or inertial navigation devices, or provide a direct camera feed to the operator so that we can identify its location via landscape features observed. As for potentially unknown sensor locations, we provide a model extension in Subsection 3.5. Our numerical experiments use a rolling horizon or model predictive control (MPC) framework to alleviate for potentially unknown sensor locations or mobile sensors.

One key insight for modeling various objectives F is the fact that the act of trying to locate an agent or just determining the existence of an agent through receiving a signal is a process with uncertainty, as any signal is affected by some random background noise [9]. Thus, whenever a sensor measures e. g. the distance of a UAV to the sensor, the sensor realises a random variable. Correspondingly, further computations of the sensors to provide for an estimate of the location of the UAV need to be treated as an estimator of a random variable. We will use this as follows. For any given time $t \in [0, T]$ and given sensor locations $\mathbf{y}(t)$, we are interested in two scenarios:

(A) The sensor operator is aware that there is a UAV somewhere in Ω , but does not know where it is. The operator subsequently wants to locate the UAV. Given that all measurements and all subsequent estimates of position come with a stochastic error, we are interested in having the UAV at a location $x(t)$ with a large error of position. We consider three expressions $E_i(\mathbf{x}(t), \mathbf{y}(t))$ ($i = 1, 2, 3$) that measure this positional error: considering the position estimate as a random variable, we define

E_1 : the Circular Error Probable (CEP) as the radius of the circle that contains half the realisations of this random variable;

E_2 : the Area of Error Ellipse (AEE) as the area of the ellipse which contains 95% of the realisations of this random variable;

E_3 : the Geometric Dilution of Position (GDOP), which can be seen as the first derivative of CEP, i. e. as the sensitivity with respect to errors in measurements.

(B) The sensor operator is not yet aware that there is a UAV somewhere in Ω . Consequently, we are interested in locations $x(t)$ such that

(a) the probability of detection, $p(\mathbf{x}(t), \mathbf{y}(t))$, is small;

(b) the probability of false alarms that the sensor operator experiences, $q(\mathbf{x}(t), \mathbf{y}(t))$, is large.

Let $\mathbf{x} \in \mathcal{D}$ denote a decision parameter. If scenario (A) holds for all mission times $t \in [0, T]$, a typical example for an objective function would be

$$F(\mathbf{x}, \mathbf{y}) = - \int_0^\tau E_i(\mathbf{x}(t), \mathbf{y}(t)) dt \quad (3)$$

for a choice of $i = 1, 2, 3$, representing the maximisation of a cumulative positional error. An alternative objective would be

$$F(\mathbf{x}, \mathbf{y}) = - \min_{t \in [0, \tau]} E_i(\mathbf{x}(t), \mathbf{y}(t)), \quad (4)$$

representing the maximisation of the worst-case (smallest) positional error, for $i = 1, 2, 3$.

In scenario (B), the corresponding objectives would be

$$F(\mathbf{x}, \mathbf{y}) = \int_0^\tau p(\mathbf{x}(t), \mathbf{y}(t)) dt \quad (5)$$

and

$$F(\mathbf{x}, \mathbf{y}) = \max_{t \in [0, \tau]} p(\mathbf{x}(t), \mathbf{y}(t)),$$

where one can replace p with $-q$ in any such objective.

This approach allows one to "switch" between objectives, i. e. at the start of a mission one might want to consider objective (5) first. After a certain amount of time has elapsed, one might then assume that the UAV has been detected, and an objective of the form (3) should be used from then on. We believe that from a user's perspective, this dynamic switch from one objective to another is more important than modelling the problem as, say, a bi-objective optimisation problem.

Finally, we will also consider time flown to target as an objective, in order to penalise flight paths that let the UAV loiter for a long time in regions in which it is relatively safe from detection.

For estimators of random variables, the variance of the estimator can be used to see how well measurements can be used to deduce something about the random variable. Our aim is thus to find decision variables (mainly, paths for the UAVs) that provide for large variances in the sensor's measurements. Indeed, moving along this vein of thought, we develop our model by first providing lower bounds on the variances of various estimators of quantities of interest that the sensors may detect. Such bounds are usually called

Cramér-Rao bounds in the literature. All these lower bounds will be derived directly from properties of the electromagnetic spectrum in which the agent emits signals and noise. Taking a worst-case view, we then use these lower bounds of variances as approximations for these variances in the construction of the objectives. More specifically, we assume that the sensor operators are able to make realisations of our position with the lowest possible variance allowed by our model.

In particular, in the following we provide measures of probability of detection (to be minimised, both in the cumulative and the worst-case sense), a measure of error in position (to be maximised, likewise in the cumulative and the worst-case sense), as well as cost of equipment usage and the time of flight.

2.2 Cramér-Rao lower bounds on variances

The sensor operator will attempt to compute various estimators of certain quantities of interest, in particular the probability of existence of the UAV if existence has not yet been established, the distance to the UAV, or location of the UAV, etc. These estimators will depend on various measurements of physical quantities that the sensors might perform, for example received signal strength (RSS), angle-of-arrival (AOA), time-of-arrival (TOA), or time-difference-of-arrival (TDOA). One of the key quantities here, occurring in estimates of the quality of those estimators, is the variance of the given measurements. Cramér-Rao lower bounds (CRLBs) are bounds from below for those variances. We refer to [25] for an introduction into the subject.

In the following, we consider a sensor at location y_j and a UAV at position $x \in \mathbb{R}^2$ which is emitting an electromagnetic signal that the sensor can detect. We provide Cramér-Rao lower bounds for distance estimators from sensors to the UAV to be used in the construction of appropriate objective functions F .

Distance estimator using RSS

For the distance estimator $d_s = d_s(x)$ of the distance $\|x - y_j\|$ based on the received signal strength RSS at the sensor from the the agent, according to [13, Equation 3] we have that

$$\text{Var}(d_s) \geq \frac{(\ln 10)^2}{100\gamma^2} \|x - y_j\|^2 \text{Var}(\text{RSS}), \quad (6)$$

where $\gamma > 0$ is a pathloss coefficient, typically taking values between 1 and 5.

Remark 2.1. From equation (6), once can see that a larger pathloss, γ , results in a smaller lower bound, as the average power becomes more sensitive to the distance for larger γ . On the contrary, as the true distance between the sensor and agent, $\|x - y_j\|$ increases, the CRLB increases with it, thus deteriorating the accuracy of measurements.

We now need to express the received signal strength in terms of our decision variable, which is x . The received signal strength is a stochastic process that can be expressed as

$$\text{RSS} = \frac{c_{AG} P_T F}{\|x - y_j\|^\gamma}, \quad (7)$$

where c_{AG} is the average multiplicative gain constant, P_T is the power transmitted by the agent, and F is the stochastic process that represents fading, modelled as a zero mean Gaussian random variable with variance $\text{Var}(F)$. (Here, the fading variables F between different sensors are assumed to be uncorrelated.) See [25] for details. Accordingly,

$$\text{Var}(\text{RSS}) = \frac{c_{AG}^2 P_T^2 \text{Var}(F)}{\|x - y_j\|^{2\gamma}}$$

and

$$\text{Var}(d_s) \geq \frac{(\ln 10)^2 c_{AG}^2 P_T^2 \text{Var}(F)}{100\gamma^2} \|x - y_j\|^{2-2\gamma}. \quad (8)$$

Distance estimator using TOA

A slightly different Cramér-Rao bound can be derived for the variance of any estimator $d_t = d_t(x)$ for the

distance $\|x - y_j\|$ based on the TOA of a signal. Indeed, for this estimator, one is able to obtain the Cramér-Rao bound

$$\text{Var}(d_t) \geq \frac{c^2}{4\pi^2\beta^2} \frac{1}{t_S\beta\text{SNR}} \left(1 + \frac{1}{t_S\beta\text{SNR}}\right)$$

with c the signal velocity (often the speed of light), β the effective bandwidth, t_S the signal duration, and SNR the signal-to-noise ratio, see [20, Equation 3].

Indeed, the SNR follows the same distance law as the RSS, i. e. for some constant $\hat{c} > 0$,

$$\text{SNR} = \frac{1}{\hat{c}\|x - y_j\|^\gamma}. \quad (9)$$

Furthermore, in many common signals, the effective bandwidth and duration of a signal are related such that $t_S\beta \approx 1$. Using this approximation along with (9) leads to

$$\text{Var}(d_t) \geq \frac{c^2\hat{c}}{4\pi^2\beta^2} \|x - y_j\|^\gamma (1 + \hat{c}\|x - y_j\|^\gamma). \quad (10)$$

2.3 Some specific objective functions of interest

It is clear that we are considering a problem with several objective functions, and in the following we will derive a number of the most important ones.

In the following, and throughout this article, for $x, y \in \mathbb{R}^2$, let $\sigma(x, y)$ denote the standard deviation of the distance estimator that a sensor at position y computes for an agent at position x . For brevity, for position $x \in \mathbb{R}^2$, let $\sigma_j(x) := \sigma(x, y_j)$ ($j = 1, \dots, n$) denote the variance that the sensor at position y_j computes for the distance estimator for an agent being in position x .

2.3.1 Error on position

We want to prevent the sensor operator from computing a good estimate of the agent's position at any time. As such, for any $t \in [0, \tau]$, denote by $E(\mathbf{x}(t), \mathbf{y}(t)) > 0$ a measure for the position error that the sensors at locations $\mathbf{y}(t)$ encounter when the agent is at position $\mathbf{x}(t)$. Then we can aim to maximise the cumulative error on position or the smallest error on position respectively by setting $F(\mathbf{x}, \mathbf{y})$ in (2) as one of (3) or (4).

To arrive at computationally tractable expressions of the function E , we will use results from Torrieri [25] and derive three different proxies for the location error that the sensors are experiencing: the circular error probable (CEP), the area of error ellipses (AEE), and the geometric dilution of position (GDOP), denotes by E_1, E_2, E_3 , respectively.

Circular error probable

We first consider the CEP. Assume that a two-dimensional Gaussian random vector describes the estimate of the location of the agent produced by the sensors. The CEP is defined as the radius of the circle that has its centre at the mean of the random vector and contains half the realisations of the random vector. It can thus serve as a one-dimensional measure of uncertainty that the sensor operator experiences when trying to calculate the location of the agent, and we are thus interested in using it as a possible proxy $E(\mathbf{x}(t), \mathbf{y}(t))$. For ease of notation, we drop the dependency on $\mathbf{y}(t)$ in the remainder of this section.

By [25, Equations 89–90], if the agent is at position $x = x(t)$ for some $t \in [0, \tau]$, the CEP is given by

$$\text{CEP}(x) = \sqrt{\text{trace}(P(x))}, \quad (11)$$

where $P(x)$ is the 2×2 covariance matrix of the linear estimator of the agent location that the sensors have at their disposal. In particular, by [25, Equations 85, 87],

$$P(x) = c^2(F^\top H^\top N^{-1}(x)HF)^{-1} \quad (12)$$

$$= c^2(F^\top H^\top (HN_\epsilon(x)H^\top)^{-1}HF)^{-1}, \quad (13)$$

where c is the signal velocity (often the speed of light), H is the constant $(n-1) \times n$ matrix given by [25, Equation 82], $F \in \mathbb{R}^{n \times 2}$ is a matrix where row vector j is $r_0^\top - y_j^\top / \|r_0 - y_j\|$ given by [25, Equation 78], where r_0 is a reference point near x , and $N_\epsilon(x)$ is the $n \times n$ covariance matrix of measurement errors of the sensors.

Remark 2.2. Note that the reference point r_0 is the point around which a Taylor expansion of the measured signals is used to derive (11) and (12). The actual reference point in use within the sensor equipment might not be known (not even to the operator of the sensor), and it might or might not be possible to change the reference point during the course of a measurement campaign. In our numerical experiments we have taken a ‘worst case’ approach and assumed that the r_0 is allowed to change over time and will always be the last best known estimate of the agent location $x(t)$. That is, if we discretise the time interval $[0, \tau]$ into discrete time steps $t_0, t_1, \dots, \in [0, \tau]$, $t_{i-1} < t_i$, then at time t_i we will use $r_0 := x(t_{i-1})$. For the agent trying to avoid location, this is, in a certain sense, the worst possible choice, as it keeps the Taylor reminder that contributes further to the location error as small as possible.

Assume that the measurements of different sensors are uncorrelated, i. e. $N_\epsilon(x) = \text{diag}(\sigma_1^2(x), \dots, \sigma_n^2(x))$. Then $HN_\epsilon(x)H^\top$ is a tridiagonal symmetric positive definite matrix with main diagonal $(\sigma_1^2 + \sigma_2^2, \dots, \sigma_{n-1}^2 + \sigma_n^2)$ and subdiagonals $(-\sigma_2^2, \dots, -\sigma_{n-1}^2)$. The corresponding inverse can then be explicitly constructed, see [8], and the computation of $\text{trace}(P(x))$ is then simple numerical exercise, given the σ_j . Using a worst-case approach, we replace the variances σ_j with lower-bounds, as a smaller variance of measurements inevitably will lead to a smaller position error. For this, we replace each σ_j with the right-hand side of (8). We thus arrive at a computationally tractable expression for CEP when the sensors measure the location of the agent $x(t)$.

If instead the sensors measure arrival times, we can further assume that at time $t \in [0, \tau]$, when the agent is at location $x(t)$, we have (10). Using the worst-case approach again, by (10) we thus have

$$\sigma_j^2 = c_1 \|x(t) - y_j\|^{\gamma_j} (1 + c_2 \|x(t) - y_j\|^{\gamma_j}),$$

where $\sigma_j^2 = \sigma_j(x(t))^2$ represents the variance for each sensor j , for some constants $c_1, c_2 > 0$ and path-loss coefficients $\gamma_j > 0$ corresponding to sensor j .

To avoid difficulties in numerical gradient calculations, we propose to use as a proxy for the error in measurement

$$E_1(\mathbf{x}(t)) = \text{CEP}(\mathbf{x}(t))^2.$$

This objective is evidently sensible in terms of maximising the location error, as the square of the CEP is proportional to the radius of the ball, where 50% of the realisations of the estimate of the location reside.

Remark 2.3. If one wants to avoid implicit matrix inversion routines in the formulation of these objective functions, one can proceed as follows. As above, for $x(t)$ we consider the matrix-valued function $N_\epsilon(x(t))$ and now introduce a matrix of variables $M \in \mathbb{R}^{n \times n}$ as well as the constant identity matrices $I_k \in \mathbb{R}^{k \times k}$. We then couple $N_\epsilon(x(t))$ and M with the constraint

$$HN_\epsilon(x(t))H^\top M = I_{n-1}. \quad (14)$$

(Due to symmetry, this matrix constraint is equivalent to $n(n-1)/2$ standard nonlinear constraints.) In other words, a feasible M must fulfil $M^{-1} = HN_\epsilon H^\top$. Treating $P \in \mathbb{R}^{2 \times 2}$ as a matrix of variables, we introduce further the constraint

$$c^2(F^\top H^\top M H F)P = I_2 \quad (15)$$

using the same idea as above. These two matrix constraints now ensure that any feasible P has the form (12). With this, $\text{trace}(P)$ is readily available.

Error ellipses

As above, let P be the 2×2 covariance matrix of the linear estimator of the agent location that the sensors have computed, based on their measurements. Let $\lambda_1, \lambda_2 > 0$ be the eigenvalues of P , dependent

on $x(t)$ and $\mathbf{y}(t)$. For a 2-dimensional normal distribution with covariance P , 95% of the mass of the distribution will reside in an ellipse with half axes $2\sqrt{5.991\lambda_1}$ and $2\sqrt{5.991\lambda_2}$. Accordingly, the AEE is

$$\text{AEE}(\mathbf{x}(t)) := 23.964\pi\sqrt{\det(P(\mathbf{x}(t)))}, \quad (16)$$

which is perhaps a less crude measurement for the location error, and can likewise be used as an objective function to be maximised. We would then use as a proxy

$$E_2(\mathbf{x}(t)) = \text{AEE}(\mathbf{x}(t)). \quad (17)$$

Remark 2.4. We note that, ignoring proportionality constants, we always have

$$\sqrt{\lambda_1 + \lambda_2} \geq \sqrt{\lambda_1\lambda_2}$$

for $0 \leq \lambda_1, \lambda_2 \leq 1$, with the maximum difference occurring for $\lambda_1 = \lambda_2 = 1$. As such, in the case of small eigenvalues $\lambda_j \leq 1$ the CEP will overestimate the area of the error ellipse, by at worst 41% in case the ellipse is a circle.

Geometric dilution of position

The GDOP can be interpreted as a scalar measure of the first derivative of the location estimate that the sensors construct, where the derivative is taken with respect to the measurement data. It thus provides a measure of sensitivity to measurements, a quantity that we would like to maximise in order to increase any error the sensors suffer. The change in measurement is given by $c \sum_{j=1}^n \sigma_j^2$. Following [25, Equations 88, 89], we can then define the GDOP as

$$\text{GDOP}(\mathbf{x}(t)) = \frac{\sqrt{\text{trace}(P(\mathbf{x}(t)))}}{c \sum_{j=1}^n \sigma_j(\mathbf{x}(t))^2}. \quad (18)$$

and we can use the same derivations as above to first express the matrix P in terms of the variances σ_j , and then use lower bounds for the variances to express GDOP in terms of $\|x(t) - y_j\|$. Indeed, we would then use as a proxy for the error in measurement

$$E_3(\mathbf{x}(t)) = \text{GDOP}(\mathbf{x}(t)).$$

Summary

In summary, for $\mathbf{x} \in \mathcal{D}$ and sensor positions \mathbf{y} , the error proxies E_i derived are given by

$$\begin{aligned} E_1(\mathbf{x}(t), \mathbf{y}(t)) &= \text{trace}(P(\mathbf{x}(t))), \\ E_2(\mathbf{x}(t), \mathbf{y}(t)) &= \sqrt{\det(P(\mathbf{x}(t)))}, \\ E_3(\mathbf{x}(t), \mathbf{y}(t)) &= \frac{\sqrt{\text{trace}(P(\mathbf{x}(t), \mathbf{y}(t)))}}{c \sum_{j=1}^n \sigma_j(\mathbf{x}(t), \mathbf{y}(t))^2}. \end{aligned}$$

2.3.2 Probability of detection

Suppose that for each sensor $j = 1, \dots, n$ we have a function $p_j(\cdot) = p_j(\cdot, y_j) : \mathbb{R}^2 \rightarrow [0, 1]$ such that $p_j(x)$ indicates the probability that sensor j detects an agent at location $x \in \mathbb{R}^2$. Denote by $p(\cdot) = p(\cdot, \mathbf{y}(t)) : \mathbb{R}^2 \rightarrow [0, 1]$ the combined probability of detection with respect to all the sensors, i. e.

$$p(x) := 1 - \prod_{j=1}^n (1 - p_j(x)). \quad (19)$$

We derive now explicit formulation for the probabilities p_j in the scenario that the sensor j treats the signal emitted by the agent as a zero mean Gaussian random process with variance $\sigma^2(x)$, in the presence of white Gaussian background noise with known variance σ_{bg}^2 . We further assume that the sensor uses a

likelihood ratio to decide if a signal has been detected, i. e. the sensor confirms the existence of the agent if, for the received signal r , we have

$$\frac{\mathbb{P}[r \mid r \text{ carries a signal}]}{\mathbb{P}[r \mid r \text{ carries no signal}]} > \gamma \quad (20)$$

for some likelihood threshold $\gamma > 0$, and a probability measure \mathbb{P} .

Finally, we assume that the sensor takes N signals to decide if the agent is present, i. e. $r \in \mathbb{R}^N$. (In practice, the number N is often dependent on the hardware in use and cannot always be easily changed. The optimal choice of N depends on the time interval over which signals are taken, background noise, and other environmental factors. Typical values range from $N = 100$ to $N = 40960$, see [11, 26] for details.) Under these assumptions, we take the logarithm in (20) and rearrange terms to arrive at a scaled threshold $\hat{\gamma}$. Then [18, Equations 5.1 – 5.3] shows that with

$$\hat{\gamma}(x) := \frac{2\sigma_{\text{bg}}^2 (\sigma^2(x) + \sigma_{\text{bg}}^2)}{\sigma^2(x)} \left(\log \gamma - \frac{N}{2} \log \left(\frac{\sigma_{\text{bg}}^2}{\sigma^2(x) + \sigma_{\text{bg}}^2} \right) \right), \quad (21)$$

we arrive at

$$p_j(x) = Q_N \left(\frac{\hat{\gamma}(x)}{\sigma^2(x) + \sigma_{\text{bg}}^2} \right), \quad (22)$$

where

$$Q_N(t) := \begin{cases} 2Q(\sqrt{t}) & \text{if } N = 1, \\ 2Q(\sqrt{t}) + \frac{e^{-t/2}}{\sqrt{\pi}} \sum_{k=1}^{N/2} \frac{(k-1)!(2t)^{k-\frac{1}{2}}}{(2k-1)!} & \text{if } N \text{ odd}, \\ e^{-t/2} \sum_{k=0}^{N/2-1} \frac{(t/2)^k}{k!} & \text{if } N \text{ even}, \end{cases} \quad (23)$$

where Q denotes the complementary cumulative distribution function of a Gaussian random variable:

$$Q(s) := \frac{1}{\sqrt{2\pi}} \int_s^\infty e^{-u^2/2} du. \quad (24)$$

Finally, we note that by (7) we can substitute

$$\sigma^2(x) = \text{const} \|x - y_j\|^{-2\gamma_j}. \quad (25)$$

with the pathloss coefficient γ_j associated with the sensor j .

These results generalise to the case where the signal is seen as a vector of uncorrelated Gaussian variables, i. e. if the sensors pool the signals received before hypothesis testing is carried out. See [18, Example 5.2]. The matter is much more complicated in the case of correlated Gaussian variables. Here, the cumulative distribution that needs to be computed belongs to the random variable $x^\top Ax$, where x is the vector of Gaussian signals and A is a positive definite matrix, dependent on the correlation matrix of x . See [18, p. 144–146] for a corresponding derivation.

Asymptotically, all functions Q_N are of the form $e^{-t}q(t)$ for some polynomial q . Setting $q = 1$, approximating the exponential via a Taylor expansion and eliminating the constant zeroth order term results in the approximation

$$p_j(x) \approx \frac{c_j}{\epsilon + \|x - y_j\|^{\gamma_j}} \quad (26)$$

with a constant $c_j > 0$, path loss exponent γ_j ($j = 1, \dots, n$), and a small constant $\epsilon > 0$ introduced to avoid numerical difficulties when x is near y . This provides a rough approximation of the model derived by Kay [18] and as explained above provides for the same qualitative behaviour of $p_j(x)$ with respect to changes in $\|x - y_j\|$. Typical values of path-loss exponents range between 1.5 and 5.

2.3.3 Probability of false alarms

To misdirect the operator of the sensors, we want to maximise the probability of false alarms that the sensor operators experience; that is a sufficiently large test statistic is achieved in (20) given that r carries no real signal. If $p_{\text{FA},j}(x)$ is the probability that sensor j raises a false alarm if the agent is at location x , the overall probability that at least one false alarm is raised is then

$$q(x) = 1 - \prod_{j=1}^n (1 - p_{\text{FA},j}(x)). \quad (27)$$

As before, consider the case that each sensor j treats the signal emitted by the agent as a zero mean Gaussian random process with variance $\sigma^2(x)$, in the presence of white Gaussian background noise with known variance σ_{bg}^2 . Assume also that the sensors take N signals to decide if the agent is present, and that the sensors use a likelihood ratio to decide if a signal has been detected, that is we follow (20). Then, by [18, Equation 5.2], we can model the probability of false alarms via

$$p_{\text{FA},j}(x) = Q_N(\hat{\gamma}(x)/\sigma_{\text{bg}}^2), \quad (28)$$

where $\hat{\gamma}(x)$ and Q_N are given by (21) and (23), respectively.

2.3.4 Time to target

We will usually also want to reach our target as soon as possible, or at a specified time $t_F \in (0, T]$. This gives rise to the objective

$$\max\{0, (\tau - t_F)\}^2, \quad (29)$$

where treating $t_F = 0$ will simply encourage paths to be shorter in time.

We recommend treating this objective as a ‘‘soft penalty’’, i. e. multiply it with a reasonably large $M > 0$ and add it to the actual objective under consideration. This allows for sufficient flexibility in treating target times t_F .

2.4 Example of a complete model formulation

Here we provide an example formulation of our UAV path planning problem, using the tools developed in the previous subsections:

$$\min_{x, \tau} \quad - \int_0^\tau E_2(x(t), \mathbf{y}(t)) dt + M \max\{0, (\tau - t_F)\}^2 \quad (30a)$$

$$\text{s.t.} \quad x \in \mathcal{D} = \bigcup_{\tau \in [0, T]} \mathcal{A}_\tau(x_0, X), \quad (30b)$$

$$\|x'\|_\infty \leq v_{\text{max}}, \quad (30c)$$

$$\|x''\|_\infty \leq a_{\text{max}}, \quad (30d)$$

where

$$E_2(x(t), \mathbf{y}(t)) = \sqrt{\det(P(x(t)))}, \quad (31)$$

$$P(x(t)) = c^2 (F^\top H^\top (H N_\epsilon(x(t)) H^\top)^{-1} H F)^{-1}, \quad (32)$$

$$N_\epsilon(x(t)) = \text{diag}(\sigma_1(x(t)), \dots, \sigma_n(x(t))), \quad (33)$$

$$\sigma_j(x(t)) = \text{const} \|x(t) - y_j\|^{2-2\gamma}. \quad (34)$$

It can be seen that we are maximising the average area of the error ellipse whilst taking care of mission times with a penalty term, and ignore all probabilities of detection. We assume that the sensors use the received signal strength of the signal emitted by the UAV to estimate the distance between sensor and UAV. Constraints (31)–(34) are meant to hold for all $t \in [0, T]$ and show how the distances $\|x(t) - y_j\|$ ultimately determine the value of E_2 . In particular, (31)–(34) can be substituted consecutively out of the problem formulation, leaving only the constraints (30b)–(30d) and a highly nonlinear objective. Constraint (34) is also meant to hold for all $j = 1, \dots, n$. The constant in (34) is the constant factor in (8).

3 Extensions to the basic model

In this section we will consider several extensions to the basic model, notably *jamming*, an electronic countermeasure in which a UAV sends out a signal in order to saturate sensors with noise; *intermittent communication*, in which a UAV stops communicating with its command centre, the case of *several UAVs*, and *regional constraints*. These extensions showcase the flexibility of our basic model: the necessary modifications are relatively straightforward, consisting of simple substitutions and the introduction of some uncomplicated constraints.

3.1 Jamming

We now discuss how *jamming* can be included in our model.

Let $\tau \in [0, T]$ and $x \in \mathcal{A}_\tau$ denote the path of an agent. We assume that the agent in question is able to interfere with the detection capabilities of the sensors (y_1, \dots, y_n) whilst travelling along its path x through the action of a limited device described by $\mathbf{j}(t)$ for $t \in [0, \tau]$.

To account for jamming action along the path in the optimisation problem (2), we introduce the space of admissible jamming functions acting over an interval $[0, \tau]$ by \mathcal{J}_τ , i.e.,

$$\mathcal{J}_\tau := \{ t \in [0, \tau] \mapsto \mathbf{j}(t) \}. \quad (35)$$

Through this, we can replace the previous decision space \mathcal{D} in (1) by

$$\mathcal{D} := \bigcup_{\tau \in [0, T]} \mathcal{A}_\tau \times \mathcal{J}_\tau. \quad (36)$$

3.1.1 Isotropic jamming

First, we assume that jamming occurs only isotropically and discuss the effects that this has on the various objectives and parameters.

Under isotropic action, we can define

$$\mathcal{J}_\tau := \{ z \mid z : [0, \tau] \rightarrow [0, 1] \}, \quad (37)$$

where $z : [0, \tau] \rightarrow [0, 1]$ acts as a measure of energy usage of the jamming device along the path.

Error on position

Let $t \in [0, \tau]$ and assume $E(\mathbf{x}(t), \mathbf{y}(t))$ measures the positional error encountered by the sensors positioned at $\mathbf{y}(t) = (y_1(t), \dots, y_n(t))$ when measuring the location of the agent that is at $x(t)$.

If $\sigma_j^2(x) = \sigma_j^2(x, y_j)$ is the variance of the measurement error of sensor j when the agent is at location x , then we further consider the case in which we have

$$E(\mathbf{x}(t), \mathbf{y}(t)) = \tilde{E}(\sigma_1^2(x(t)), \dots, \sigma_n^2(x(t)), \mathbf{y}(t)) \quad (38)$$

with a smooth function \tilde{E} . This holds for all previously considered positional error functions. If we have energy for jamming at our disposal, we can then replace each occurrence of σ_j^2 with

$$\left(1 + \frac{c_j z(t)}{\|x(t) - y_j\|^{\rho_j}} \right) \sigma_j^2, \quad (39)$$

where the term $1 + c_j z(t) / \|x(t) - y_j\|^{\rho_j}$ measures the increase in the measurement error due to the strength of our jamming signal, with $c_j > 0$ being a multiplicative factor and $\rho_j > 0$ a constant.

This straightforward modification covers changes to the circular error probability CEP, the GDOP, as well as the area of the error ellipse, according to the strength of our jamming attempts.

This approach corresponds to a linear model $1 + \text{const} \cdot z$ for the relative increase in variance, in case of a stationary agent. Indeed, results for direction of arrival models. see [27], based on Cramér-Rao bounds

indicate that the CEP itself exhibits an inverse proportional relationship with the SNR. According to this result, we have

$$\text{CEP} \approx c_1 + \frac{c_2}{\text{SNR}} = c_1 + \frac{c_3 z(t)}{\|x(t) - y_j\|^\gamma} \quad (40)$$

for suitable constants and the relevant path loss coefficient γ .

Probability of detection

Naturally, the probabilities of detection $p_j(x(t))$ increase when jamming is utilised. We can, for example, model this by replacing each occurrence of $p_j(x(t))$ by

$$((p_j(x(t)))^q + (z(t))^q)^{1/q} \approx \max\{p_j(x(t)), z(t)\} \quad (41)$$

for some reasonably large parameter $q > 1$. The smooth approximation (left) of the max term on the right facilitates the use of high-performance optimisation routines that rely on differentiable functions everywhere; the error introduced is ca. 4.7% for $q = 15$ and ca. 3.5% for $q = 20$. A more realistic model though is obtained by replacing each occurrence of $p_j(x(t))$ by

$$\min \left\{ 1, p_j(x(t)) \left(1 + \frac{\hat{c}_j z(t)}{\|x(t) - y_j\|^{\mu_j}} \right) \right\} \quad (42)$$

with some coefficients $\mu_j > 0$ and constants $\hat{c}_j > 0$. Note that here the increase in detection probability is modelled as an affine-linear function $1 + \text{const} \cdot z$ with cutoff at 1, for static agents. Models of higher fidelity can instead be constructed by using sigmoid functions, i. e. functions of the form $S(z) = 1/(1 + e^{-z})$ for the increase in probability, see [2, 6].

3.1.2 Anisotropic Jamming

We introduce here a directional model for jamming. As such, alongside the jamming energy usage function z , we define a function $\alpha : [0, \tau] \rightarrow [0, 2\pi]$, where for $t \in [0, \tau]$, $\alpha(t)$ indicates the direction the agent attempts to jam the sensors (where we have fixed the coordinate system arbitrarily).

Indeed, to include anisotropic jamming in the above model and to treat the directionality of jamming as a variable, we can set \mathcal{J}_τ in (35) to be

$$\mathcal{J}_\tau = \{ (\alpha, z) \mid (\alpha, z) : [0, \tau] \rightarrow [0, 2\pi] \times [0, 1] \}.$$

where for $t \in [0, \tau]$, $\alpha(t)$ indicates the direction the agent attempts to jam sensors, while $z(t)$ indicates the energy expenditure. In this situation, the agent employs e. g. beamforming techniques to generate a jamming signal. In the below we discuss only the scenario of one beam being formed.

Error on position

Exactly as in the isotropic case, recall that $\sigma_j^2 = \text{const} \times \|x(t) - y_j\|^2$ is the variance of the measurement error of sensor j of an agent at position $x(t)$. In the same vein as before, we then replace each occurrence of σ_j^2 with

$$\sigma_j^2 \left(1 + \frac{c_j z(t)}{\|x(t) - y_j\|^{\rho_j}} \right) \times (\cos \alpha(t), \sin \alpha(t)) \cdot \frac{y_j - x(t)}{\|y_j - x(t)\|}. \quad (43)$$

Again, the term $1 + c_j z(t)/\|x(t) - y_j\|^{\rho_j}$ measures the increase in variance of measurement error by the strength of our jamming signal, while the inner product $(\cos \alpha(t), \sin \alpha(t)) \cdot (y_j - x(t))/\|y_j - x(t)\|$ is equal to one if and only if we ‘point’ our jamming device in direction $(y_j - x(t))/\|y_j - x(t)\|$. Note that for $(\cos \alpha(t), \sin \alpha(t)) \cdot (y_k - x(t)) < 0$, i. e. for sensors with an angle of more than 90 degrees from our jamming direction, the measurement error decreases: for these sensors, it has become easier to collect information on the agent’s location. If the latter effect is not intended, we can replace each occurrence of σ_j^2 instead with

$$\sigma_j^2 \left(1 + \frac{c_j z(t)}{\|x(t) - y_j\|^{\rho_j}} \right) a_j(\alpha(t), x(t), y_j), \quad (44)$$

where the function a_j is a function that measures the decay of our jamming signal away from the direction $(\cos \alpha(t), \sin \alpha(t))$. A simple example would be

$$a_j(\alpha, x, y) := \max \left\{ 0, (\cos \alpha, \sin \alpha) \cdot \frac{y - x}{\|y - x\|} \right\}^2.$$

The max term here is squared in order to make the function a differentiable when $a(\alpha, x) = 0$. With this, the jamming signal is only strengthened in the direction $(\cos \alpha(t), \sin \alpha(t))$, and all direction that have at most a 90° angle with this vector. This approach can easily be generalised to the case where beamforming is used to form a cone with apex $x = x(t)$ in which jamming takes place. As we consider here the two-dimensional case, let the cone angle be given by $\beta \in [-\pi/2, \pi/2]$. We then replace the function a_j above by

$$a_{j,\beta}(\alpha, x, y) := \frac{1}{(1 - \cos(\beta/2))^2} \times \max \left\{ 0, (\cos \alpha, \sin \alpha) \cdot \frac{y - x}{\|y - x\|} - \cos(\beta/2) \right\}^2.$$

This approach can be used to modify the calculation of CEP, GDOP, as well as the area of the concentration ellipse.

Probability of detection

Exactly as in the isotropic case, recall that $\sigma_j^2 = \text{const} \times \|x(t) - y_j\|^2$ is the variance of the measurement error of sensor j . In the same vein as before, we then replace each occurrence of σ_j^2 with

$$\sigma_j^2 \left(1 + \frac{c_j z(t)}{\|x(t) - y_j\|^{\rho_j}} \right) \times (\cos \alpha(t), \sin \alpha(t)) \cdot \frac{y_j - x(t)}{\|y_j - x(t)\|}. \quad (45)$$

Again, the term $1 + c_j z(t)/\|x(t) - y_j\|^{\rho_j}$ measures the increase in variance of measurement error by the strength of our jamming signal, while the inner product $(\cos \alpha(t), \sin \alpha(t)) \cdot (y_j - x(t))/\|y_j - x(t)\|$ is equal to one if and only if we 'point' our jamming device in direction $(\cos \alpha(t), \sin \alpha(t))$.

Remark 3.1. Jamming means that an agent spends additional energy, which is not only a limited resource, but the use of which will usually lead to increase in probability of detection. Naturally we should thus constrain the total energy usage by adding an additional constraint of the form

$$\int_0^\tau z(t) dt \leq z_{max} \quad (46)$$

with a scenario-dependent parameter z_{max} .

3.2 Intermittent communication

We now allow for intermittent communication between the agent and some central command centre. In this scenario, the agent is allowed to lower its transmission power to zero for a certain amount of time. We slightly relax this notion and allow the agent to lower its transmission power anywhere from 0% to 100%, for a certain amount of time. For a flight time $\tau \in [0, T]$, we model this reduction by a new decision variable

$$\theta : [0, \tau] \rightarrow [0, 1] \quad (47)$$

where $\theta(t)$ is a multiplicative factor applied to the transmission power at time t . Consequently, we consider the variables

$$\mathcal{I}_\tau := \{ \theta \mid \theta : [0, \tau] \rightarrow [0, 1] \}, \quad (48)$$

and as such replace the previous decision space \mathcal{D} in either (1) or (36) with

$$\mathcal{D} := \bigcup_{\tau \in [0, T]} \mathcal{A}_\tau \times \mathcal{I}_\tau$$

or

$$\mathcal{D} := \bigcup_{\tau \in [0, T]} \mathcal{A}_\tau \times \mathcal{J}_\tau \times \mathcal{I}_\tau, \quad (49)$$

respectively.

The factor $\theta(t)$ can now be introduced into the modelling of the error on position and probabilities of detection in the following way.

Error on position

To update errors on positions observed by sensor operators when transmission power of agent communications is reduced, we scale the variances of measurement errors σ_j^2 by

$$\frac{1}{(1 - \theta(t))^2} \sigma_j^2 \quad (50)$$

for $j = 1, \dots, n$. In particular, as transmission power is reduced, i.e. $\theta(t) \rightarrow 1$, the variance of the agent position estimator calculated by the sensor operators gets larger, as does the corresponding measure of error on position.

Probability of detection

In the non-logarithmic gain model we employ, we can simply replace each occurrence of $p_j(x(t))$ by $\theta(t)p_j(x(t))$. Thus, (19) becomes

$$p(x(t), \theta(t)) := 1 - \prod_{j=1}^n (1 - (1 - \theta(t))p_j(x(t))). \quad (51)$$

Remark 3.2. Operators of an agent will want to keep in contact with the agent for a certain period of time; they will thus allow the total reduction of signal energy emitted by the agent by at most a value of θ_{allow} , which we assume to be fixed. This gives rise to the additional constraint

$$\int_0^\tau \theta(t) dt \leq \theta_{allow}. \quad (52)$$

3.3 Several agents

The model described above and all its features can be extended to the case of several agents in a straightforward manner. Consider for example the case of m agents, following paths

$$x^{(1)}, \dots, x^{(m)} : [0, \tau] \rightarrow \mathbb{R}^2 \quad (53)$$

with various starting points

$$x^{(k)}(0) = x_{0,k} \quad (k = 1, \dots, m) \quad (54)$$

(some of which of course can be identical) and targets

$$x^{(k)}(T) = X_k \quad (k = 1, \dots, m). \quad (55)$$

We likewise introduce jamming variables

$$(\alpha^{(1)}, z^{(1)}), \dots, (\alpha^{(k)}, z^{(k)}) : [0, T] \rightarrow [0, 2\pi] \times [0, 1]$$

and other variables per agent as necessary. Likewise, each agent has their own objective function (to be minimised). We will associate to each objective a weight $\omega_k > 0$ ($k = 1, \dots, m$), signifying its importance, and will sum up all those objectives to one.

To include this in the optimisation problem, we replace the decision space \mathcal{D} in (1) or more generally (49) with

$$\mathcal{D} = \prod_{k=1}^m \bigcup_{\tau \in [0, \tau_k]} (\mathcal{A}_{\tau_k}(x_{0,k}, X_k) \times \mathcal{J}_{\tau_k} \times \mathcal{I}_{\tau_k}). \quad (56)$$

3.4 Regional constraints

One may also wish to prescribe regions which the agent(s) should avoid. We write such a region in \mathbb{R}^2 as a generic set of the form

$$R := \{x \in \mathbb{R}^2 \mid \rho_j(x) \leq 0, j = 1, \dots, r\} \quad (57)$$

with functions $\rho_j : \mathbb{R}^2 \rightarrow \mathbb{R}$ ($j = 1, \dots, r$), and employ it to our optimisation problem as an inequality constraint for the agent(s) path, $x : [0, \tau] \rightarrow \mathbb{R}^2$, to satisfy at each time $t \in [0, \tau]$. Note that it is straightforward to modify this constraint in case such a restriction to movement should only hold at certain times, e. g. at the start of the mission: we then consider constraint functions $\rho_j(x, t)$ depending also on time.

A special case of such forbidden regions are areas in which the probability of detection is larger than some prespecified bound p_U , $0 < p_U < 1$, or areas in which the location error is smaller than some prespecified bound $E_L > 0$. This gives rise to the constraints

$$p(x(t)) \leq p_U \quad t \in [0, \tau] \quad (58)$$

and

$$E(\mathbf{x}(t), \mathbf{y}(t)) \geq E_L \quad t \in [0, \tau]. \quad (59)$$

Some care has to be taken in choosing the values of p_U and E_L to construct R ; a too small p_U , or too large E_L could result in a $\Omega \setminus R$ that is no longer path-connected, and consequently it may not be possible for the target to be reached by the agent.

3.5 Unknown sensor locations

In case the sensor locations $\mathbf{y}(t)$ are not known by the operator of the UAV, we instead make the relaxed assumption that we believe that individual sensors lie in certain known uncertainty sets; for example we assume that at time t the location $y_j(t)$ of sensor j lies in a ball of radius $r_j(t)$ around a fixed guess \hat{y}_j of the sensor location,

$$y_j(t) \in B_j(t) := \{y \in \mathbb{R}^2 \mid \|y - \hat{y}_j\| \leq r_j(t)\}. \quad (60)$$

This gives rise to the uncertainty set

$$\mathcal{U}(t) = \prod_{j=1}^n B_j(t) \quad (61)$$

and a corresponding robust reformulation of our problem

$$\max_{\mathbf{x}} \min_{\mathbf{y}(t) \in \mathcal{U}(t)} F(\mathbf{x}, \mathbf{y}). \quad (62)$$

In the general case, it is not clear what appropriate solution approaches for problems of the form (62) are. However, in the case in which we want to minimise the overall probability of detection we note that the worst case occurs when $y_j(t)$ is as close as possible to $x(t)$, i. e. for

$$y_j(t) = y_j(x(t)) := \hat{y}_j + r_j \frac{x(t) - \hat{y}_j}{\|x(t) - \hat{y}_j\|} \quad (63)$$

as long as $x(t) \notin B_j(t)$. As such, the inner min-operator in (62) can be resolved explicitly if we also add constraints of the form $\|x(t) - \hat{y}_j\| \geq r_j(t)$, which appears to be quite reasonable in a worst-case framework. More complicated sets of uncertainty $B_j(t)$ can, at least in principle, be treated in a likewise fashion by setting $y_j(t) = \text{proj}_{B_j(t)}(x(t))$ for $x(t) \notin B_j(t)$, thus trading the min-operator against a certain amount of nonsmoothness.

3.6 Handling future uncertainty

We now briefly describe how one can use a rolling horizon approach to handle uncertain information present in the environment over which the UAV(s) traverse, whereby optimal paths and equipment usage is updated as new information is received. Due to its simplicity, the rolling horizon approach to optimisation is one of the preferred modelling approaches (see, e. g. [14]) if further information about uncertainties in the problem becomes available over the course of the solution process. The rolling horizon approach is essentially a nonlinear single-stage model predictive control (MPC) framework [12, 22], widely deployed in chemical engineering and related areas, but here used for trajectory optimisation.

We apply the rolling horizon approach with respect to the maximal time of flight interval of the UAV(s), $[0, T]$, over which we optimise and subsequently also solve optimisation problems over its subintervals. In particular, we propose the following rolling horizon approach to solve the optimisation problem (2) in the presence of uncertainties in the environment. Begin by first choosing the *scheduling horizon* $[0, T]$, and construct both the *prediction horizons* $[T_i, T]$ with $T_i < T_{i+1}$ for $i = 0, 1, \dots$ where $T_0 = 0$, and the *control horizons* $[T_i, C_i]$ with $C_i = T_{i+1}$ for $i = 0, 1, \dots$. At the i th step, beginning with $i = 0$, obtain a solution to the optimal path problem (2), $\mathbf{x}^{(i)}$ over the i th prediction horizon, $[T_i, T]$, making sure to update the boundary conditions of the problem with the change in the prediction horizon, e.g. the starting and ending positions of the agent. Within the i th control horizon, $[T_i, C_i]$, the UAV(s) move along the path $\mathbf{x}^{(i)}$. During movement within the control horizon, the UAV(s) can update their current estimate of uncertainties of their environment. Once the UAVs have reached the end of their control horizon, and thus the beginning of the next prediction horizon, make use of this updated information when computing an optimal path over the next prediction horizon, $\mathbf{x}^{(i+1)}$. Repeat this process until the target for the UAV(s) is reached within the i th control horizon.

Remark 3.3. In the rolling horizon approach algorithm outlined above, we have left several parts deliberately ambiguous, comments on which we provide below.

1. The choice of the lengths of the control horizons, i. e. at which the agent stops and updates uncertainty parameters and then solves a new optimisation problem is up to the user. An ad-hoc procedure would simply choose a fixed number J of control horizons and use $C_i := iT/K$ ($i = 1, \dots, K$). More sophisticated approaches would couple the choice of C_i with changes in the observable environment. Indeed, large changes in the environment between control horizons indicate a dynamic environment that might warrant some further care in path planning, and thus immediately proceeding control horizons $[T_{i+1}, C_{i+1}]$ should not be too large. If, on the other hand the overall environment does not seem to have changed much between time T_i and time C_i , computation time can be saved by using a comparatively large interval $[T_{i+1}, C_{i+1}]$ for an immediately proceeding control horizon.
2. In each step of the rolling horizon algorithm, we are free to choose an appropriate objective function to use within the optimisation problem (2) over the proceeding prediction horizon $[T_i, T]$.
3. Finally, we note that the actual path taken by the agent $x : [0, \tau] \rightarrow \mathbb{R}^2$ is 'stitched' together from the paths $x^{(i)} : [T_i, C_i] \rightarrow \mathbb{R}^2$ where the intervals $[T_i, C_i]$ only overlap at their endpoints, i. e. $T_{i+1} = C_i$. We ensure continuity of x by ensuring $x^{(i)}(C_i) = x^{(i+1)}(T_{i+1})$, but no further smoothness conditions are ensured. These can however easily be included into the overall approach by ensuring that corresponding constraints are added when defining the proceeding optimisation problem.
4. The rolling horizon approach can be used to account for any changes in environment, for example, newly discovered sensors, or updating the positions of moving sensors.

4 Numerical solution

In this chapter, we provide a number of illustrative results to showcase the flexibility and versatility of the mathematical model and its variants discussed in the previous chapters. More specifically, we provide means to solve the optimisation problem described by (2). Models in this chapter have been implemented both in MATLAB and the optimisation modelling language AMPL to verify the correctness

of the implementation.

Computation times appear to be very modest; on a Dell laptop with intel i5 CPU at 2.6GHz, MATLAB's `fmincon` optimisation function was able to solve every optimisation problem provided in less than a minute, even with standard options. Solving the same or very similar problems in an AMPL formulation and the state-of-the-art solver KNITRO [19] resulted in seconds of computation time on the same hardware platform. Further improvements can also be achieved by appropriate code optimisation. As we aim mainly to provide illustrative results, certain parameters have been chosen arbitrarily irrespective of physical assumptions.

4.1 Discretisation

The optimisation problem (2) with various choices of decision space and objective function as described in previous sections is infinite dimensional, i.e. the vector space of variables \mathcal{D} does not have a finite basis and thus no finite representation. Even more, various constraints are also infinite dimensional. To arrive at a computationally tractable formulation with a finite number of variables and a finite number of constraints, we will need to discretise the problem. For demonstration purposes, we propose the following simple approach.

Assume that $\mathbf{x} \in \mathcal{D}$ provides a solution to (2), where $x : [0, \tau] \rightarrow \mathbb{R}^2$ denotes the path component of \mathbf{x} with travel time $\tau > 0$. Next, choose an integer $N \in \mathbb{N}$ and discretise the interval $[0, \tau]$, describing the period in which the UAV is in motion, into a finite number of points

$$0 =: t_0 < t_1 < \dots < t_{N-1} < t_N := \tau,$$

where, for simplicity, we choose $t_i := t_0 + ih$ for $i = 1, \dots, N$, with grid length $h := \tau/N$. We refer to t_0, \dots, t_N the discretisation mesh of $[0, \tau]$.

Remark 4.1. In practice, one should choose an adaptive grid based on progress within the optimisation process and estimates of local discretisation errors to improve the performance of any optimisation scheme as well as the accuracy of the solution.

To discretise the time dependent parameters in the decision space of our choosing $\mathbf{x} = (x, z, \alpha, \theta) \in \mathcal{D}$, we introduce up to $4(N + 1)$ real-valued variables

$$\begin{aligned} x_i &:= x(t_i), & \alpha_i &:= \alpha(t_i), \\ z_i &:= z(t_i), & \theta_i &:= \theta(t_i), \end{aligned}$$

for $i = 0, \dots, N$, which will be the variables of our finite-dimensional optimisation problem. As such, we denote by \mathcal{D}_h the space of h -discretised optimisation parameters.

To arrive at corresponding objective functions, we naturally replace each integral with its Riemann sum approximation, partitioned under the same discretisation mesh of $[0, \tau]$, via

$$\int_0^\tau f(t)dt \rightarrow h \sum_{i=1}^N f(t_i),$$

where we note that higher-order integral discretisations schemes can be considered as necessary.

To replace and approximate derivatives, we opt for a simple Euler discretisation scheme. In particular, we replace the velocity and acceleration values at t_i for appropriate $i = 1, \dots, N$ with

$$\begin{aligned} v(t_i) &\rightarrow v_i^h := \frac{1}{h}(x_i - x_{i-1}), \\ a(t_i) &\rightarrow a_i^h := \frac{1}{h^2}(x_{i+1} - 2x_i + x_{i-1}). \end{aligned} \tag{64}$$

We would like to emphasize here that such a straightforward discretization scheme like the explicit Euler method deployed as above can often lead to large numerical errors over time, although in our experiments described below we have not encountered any particular blowup of errors. Nevertheless, we would in

general advise to use higher order schemes (see e. g. [5]), as well as appropriate modeling tools that facilitate their usage [1, 4].

In summary, we arrive at a new discrete optimisation problem of the form

$$\min_{\mathbf{x} \in \mathcal{D}_h} F(\mathbf{x}, \mathbf{y}) \quad (65)$$

along with both the equivalent discretised, and user-defined constraints, to determine an optimal path and equipment usage $\mathbf{x} \in \mathcal{D}_h$. Furthermore, F here denotes a combination of the various metrics that can be used by the UAV to evade detection by the sensor operator, including the probability of detection (19), probability of false alarms (27), CEP (11), AEE (17), and GDOP (18), as described by Section 2.1.

4.2 Landscapes

Following the definitions of the various metrics that can be used to measure detection, or lack thereof, by the UAV(s) in determining an optimal path, in this section, we provide a description of the landscapes, colourised contour plots, formed by each metric described in Section 2.3 and observe the similarities and differences between them.

We investigate the landscapes under an example eight sensor configuration on a square grid, which we shall use for the majority of this section. Figure 2 depict the landscapes of a number of functions of interest; the probability of detection (19), probability of false alarms (27), CEP (11), AEE (17), and GDOP (18). In these figures, yellow depicts high, and blue depicts low values for the metrics. We note that, the error metrics CEP, AEE, and GDOP have very large variations, where details on their respective landscape plots can be obscured. Furthermore, such large variations can be difficult for some numerical optimisers to handle, as small variations in variables can lead to large and drastic variations in objective values. To remedy both of these flaws, in these plots and the remaining parts of this article, we replace all instances of CEP, AEE, and GDOP with their logarithm.

The probability of detection metric(s) (see Figure 2a and Figure 2b) depict high values about each sensor position which rapidly decline to low values between them. Between neighbouring sensors, one can discern ridge-like structures of relatively high probability of detection. To attain a small cumulative probability of detection, for example as in scenario (B), a UAV will wish to stay in regions of low probability of detection, i.e. in the blue shaded valleys.

In contrast, the CEP, similarly the AEE landscapes (see Figure 2d and Figure 2e) consist of a roughly polygonal valley, shaped by the outermost sensors, and with a rather flat valley floor in the inner most ring of sensors. High values of CEP and AEE are encountered near the boundary of the domain considered, seemingly where the sensors do not well triangulate those points. In addition, we observe relatively high CEP and AEE values along splines pointing away from the barycentre of the sensors, and along lines of colinearity between sensors. In the CEP and AEE environments, a UAV will wish to adhere to regions where high values can be attained to best spoof the sensor operators, as one would with in scenario (A).

It is clear that the probability of detection and error in position objectives are in some conflict with each other: paths of low overall probability need to navigate with some care between different sensors, staying away from all of them, while paths that maximise overall errors should hug the sensors and follow splines with relatively high CEP/AEE values connecting adjacent sensors. The task of finding an optimal path with either the objective *minimise probability of detection* or *maximise error* can thus be visualised as finding optimal paths in one of the two landscapes. In the following Section 4.3, we provide example numerical solutions to optimal paths along each of these environments.

It is also easy to see that the error metrics E_i are not wholly in opposition with one another, unlike with the probability of detection metric(s). Although they do present with different properties. The CEP and the AEE the most similar to one another amongst the error metrics, providing a similar landscape profile, although, we notice that the variation in the CEP occurs on slightly shorter scales, as can be seen by the balloon like regions about each sensor. Furthermore, due to the faster variation in the CEP, we observe much more clearly the splines of relatively high error along lines connecting sensors. The GDOP, although being somewhat similar to the CEP and AEE along the boundary of the sensors, is very

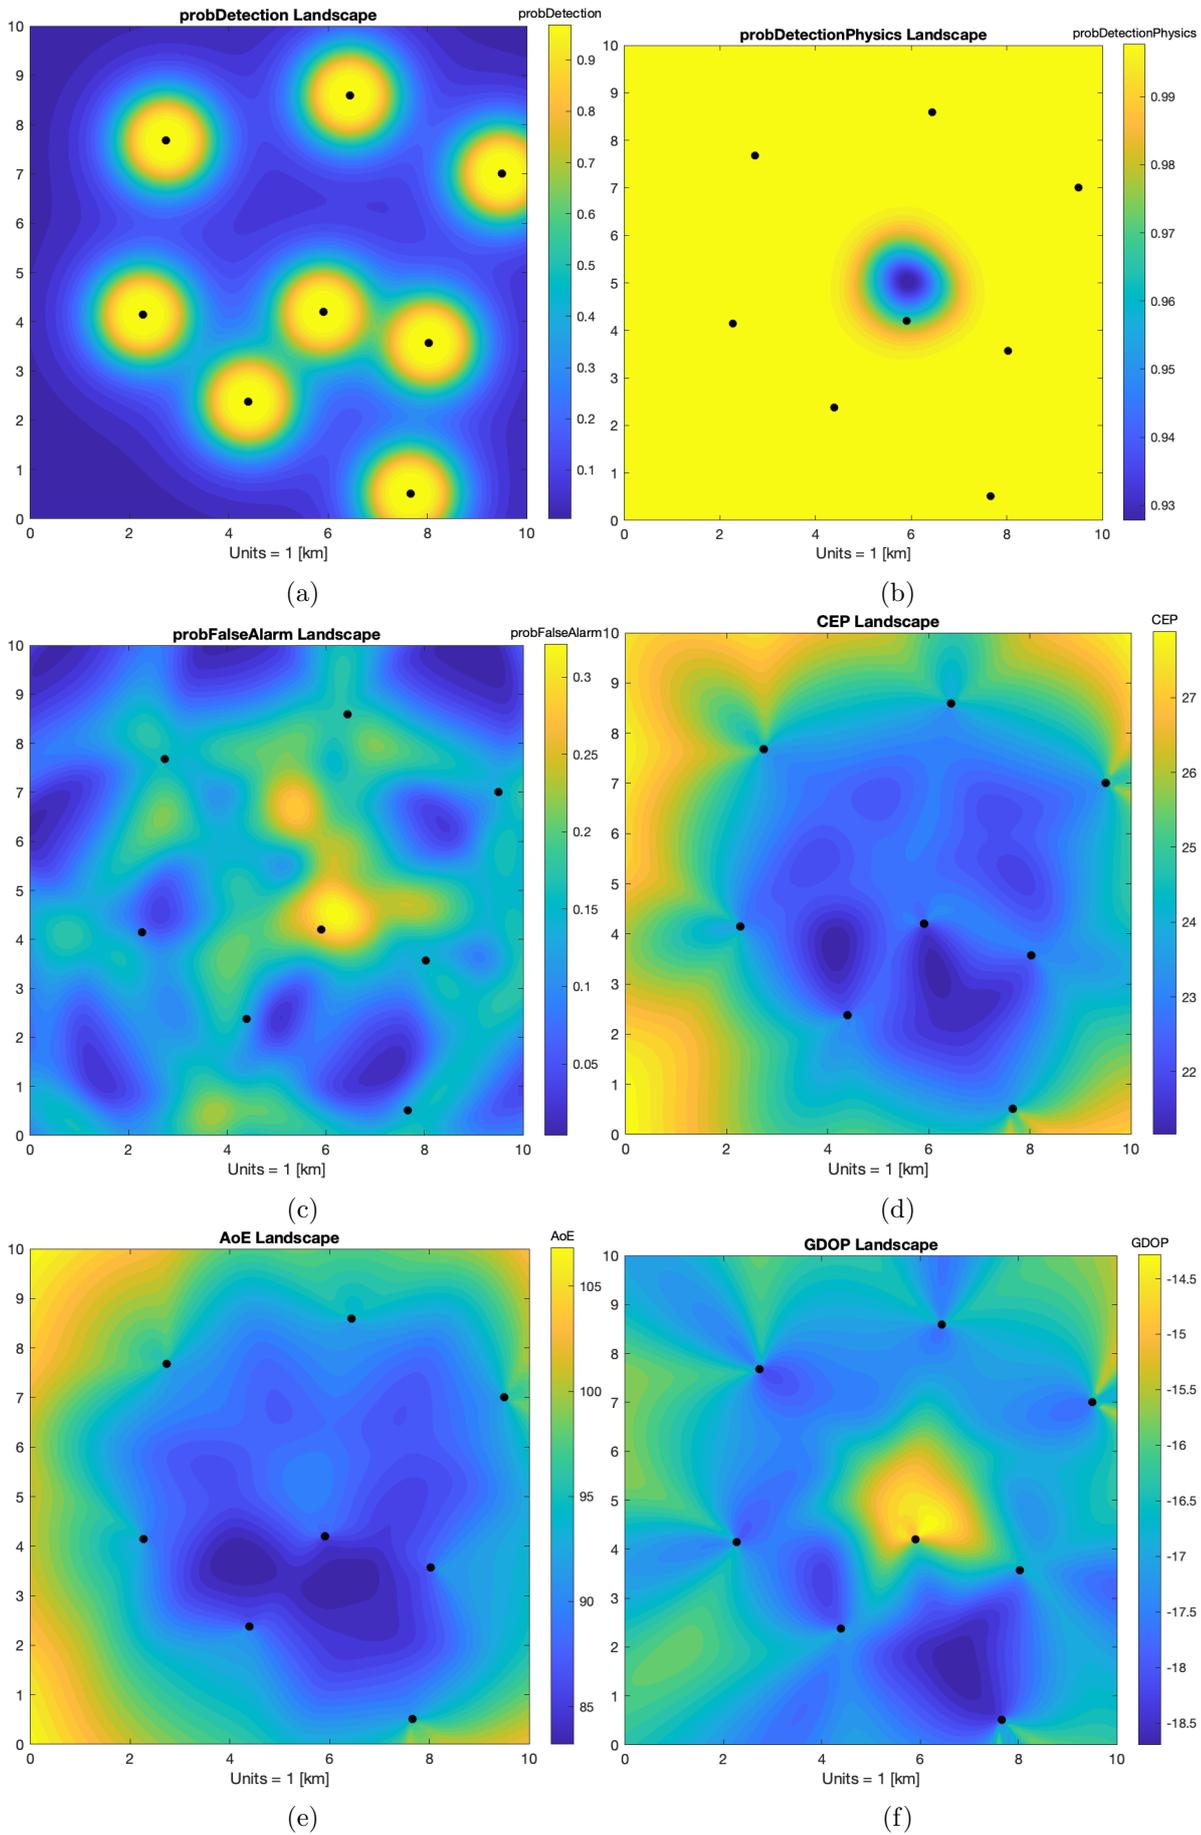


Figure 2: Landscape plots for an example sensor configuration (a) probability of detection (proxy), (b) probability of detection (physics based), (c) probability of false alarms, (d) CEP, (e) AEE, (f) GDOP. The value of each point in the plot corresponds to the value of the metric experienced by an agent at that point. Sensors are depicted as black dots.

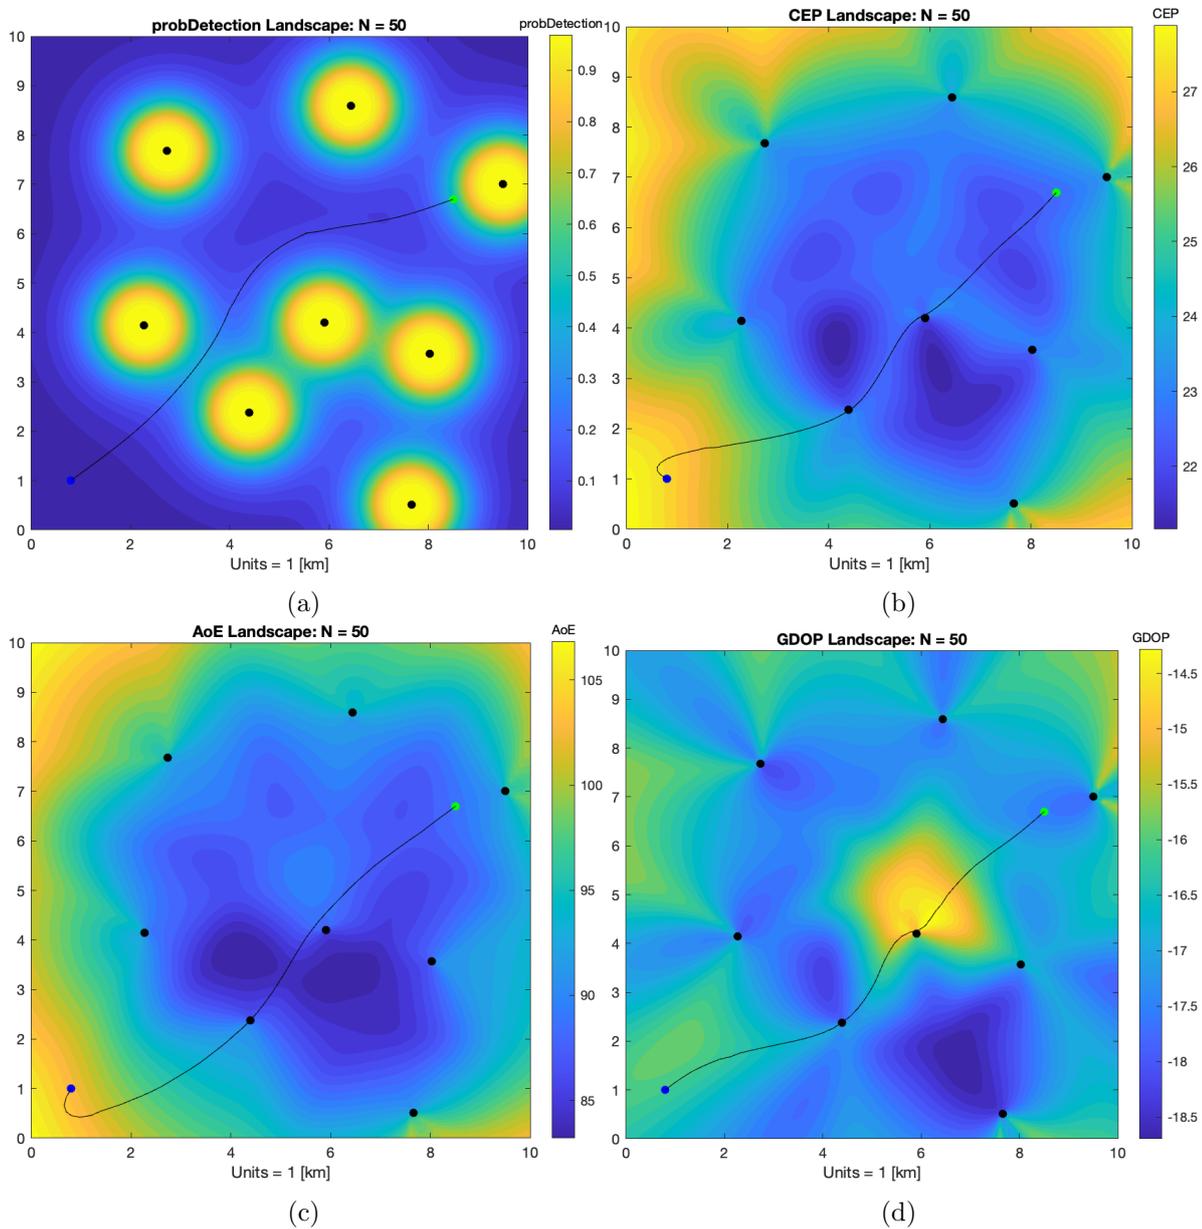


Figure 3: Optimal paths from (0.8, 1) to target location (8.5, 6.7) in (a) probability of detection, (b) CEP, (c) AEE, (d) GDOP landscapes. Optimal path is depicted as a black line, and sensors are depicted as black dots.

different at near the barycentre of the sensors. The GDOP exhibits very high values at the barycentre compared to relatively low valued in the CEP and AEE. By construction, the behaviour of the GDOP makes sense. Along regions of constant variance, the GDOP will scale very much like the CEP, although in regions of high or low variance in measurement, the GDOP relation to the CEP will be reduced, or increased, respectively. Since there is a relatively low variance at the barycentre of the sensors, due to good triangulation of position by the sensors, the GDOP will be high in this region.

4.3 Tradeoffs between different objectives

We now compute the optimal paths for a UAV under the various described landscapes; CEP, AEE, GDOP and the probability of detection. In particular, we depict solutions to the following optimisation problem

$$\begin{aligned} & \min_{\mathbf{x} \in \mathcal{D}_h} F(\mathbf{x}, \mathbf{y}) \\ \text{s.t.} \quad & v_i^h \leq v_{max} \quad i = 1, \dots, N, \end{aligned} \tag{66}$$

$$a_i^h \leq a_{max}, \quad i = 1, \dots, N, \tag{67}$$

$$x_0 = (0.8, 1), \quad x_{N+1} = (8.5, 6.7), \tag{68}$$

where F describes the cumulative cost objective(s) given by (3) and (5), \mathcal{D}_h denoting the space of real-valued variables $(x_i)_{i=0}^{N+1} \subset [0, 10]^2$, and \mathbf{y} being the positions of the eight stationary sensors in the square. Figure 3 depicts solutions to this problem under different choices of objective function.

Figure 3a provides a solution to the optimisation problem when F is chosen to denote the cumulative probability of detection, as described in (5), whereas Figure 3b, Figure 3c, and Figure 3d depict solutions under F given by (3) with error proxies chosen to be E_1, E_2, E_3 , respectively.

In the case of minimising the cumulative probability of detection, the UAV avoids sensors as much as possible whilst crossing low lying ridges instead of higher terrain features. In contrast, when maximising the CEP, AEE, and GDOP cumulative cost, the agent hugs the sensors as much as possible to utilise the large errors present directly next to the sensors.

4.4 Model variants

In this section, we extend numerical experiments and investigate the effect of setting target times, implementing jamming, and considering multiple UAVs working together to accomplish a common task.

Varied target times

Consider the same set up as before, and instead optimise over cumulative probability of detection along with an additional target time cost introduces as a penalty to the objective of the form

$$\max\{0, (\tau - t_F)\}^2, \tag{69}$$

where τ denotes the time of flight of the variable path, and τ_F denotes a target time. The UAV will now be encouraged to reach the target by the target time t_F . Indeed, the penalty can be added as a soft penalty with a large coefficient by summing with the original objective, or be included in the time constraint for the optimal paths by replacing the search space for the time of flight from $[0, T]$ to $[t_F - \varepsilon, t_F + \varepsilon]$ for some chosen (sufficiently small) slack parameter $\varepsilon > 0$.

For simplicity, we consider target arrival times to be scaled versions of the minimal possible arrival time t_M (minimal arrival time computed as the shortest distance divided by the maximum prescribed velocity of the UAV). In particular, we consider $t_F = t_M, 1.05t_M, 1.1t_M, 1.15t_M$; Figure 4b shows the corresponding optimal trajectories. In all four cases, the targeted arrival time is met, i.e. we have $\tau = t_F$ at the computed optima. When $t_F = t_M$, the optimal path is, as expected, a straight line from the start to finish, due to the velocity constraint. Indeed, it can be observed that increasing the target time allows the agent to maximise its distance to the sensors and enter valleys of low probability of detection for longer periods of time. This effect is seen more and more with larger target times.

Figure 4b shows the corresponding probability of detection along each path. The segment of the paths with the highest probability of detection can be clearly identified by the peaks, corresponding to moments

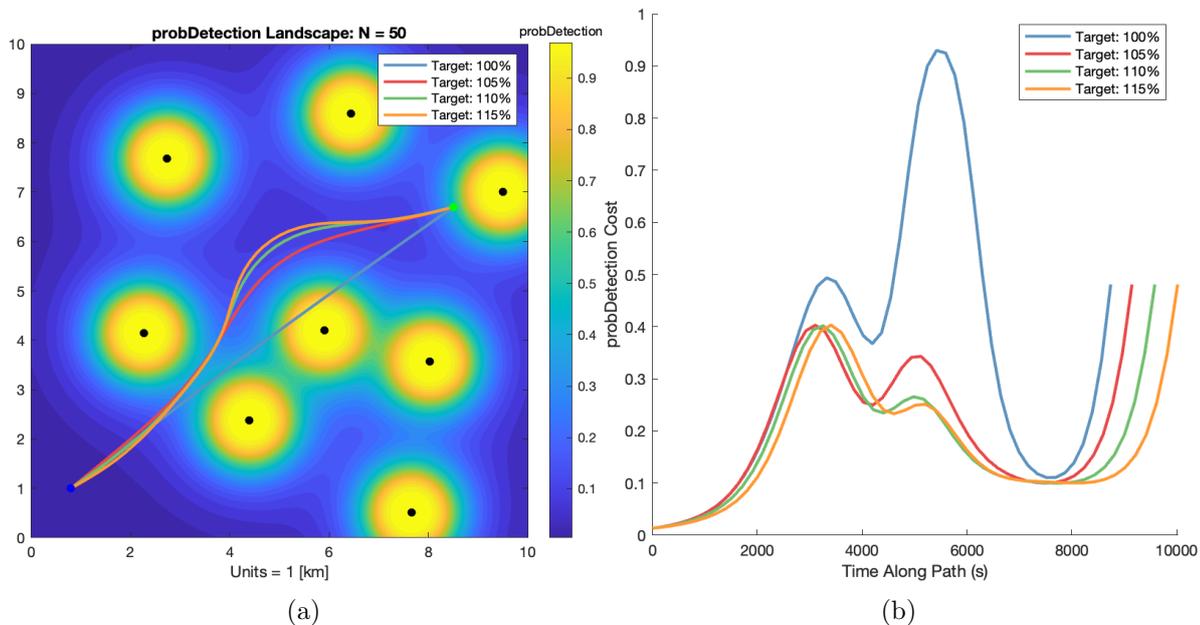


Figure 4: (a) Optimal paths (lowest cumulative probability objective) between the starting point (0.8, 1) and the target location (8.5, 6.7) in a scenario with eight sensors and subsequent increments of target time as percentages of minimal flight time. (b) Probability of detection over time for each of the optimal paths.

along the paths where the agent is close to a sensor. Also observable is the increase in the probability of detection near the target for all optimal paths. Furthermore, it can be seen that a path with a larger budget/target time spends more time in regions of low probability of detection.

Effect of jamming

We now investigate the effect of jamming action on the choice of optimal path for a sole UAV under the same conditions as described above, although with the objective being chosen as (3) with E_2 (AEE) plus a rescaled target time cost. The results of this set up are provided Figure 5.

We notice that without jamming, the optimal path hugs splines about sensors where a relative high AEE is observed. Although with jamming, the UAV is instructed to take a much more direct path through the low AEE valleys as allowed by the jamming action, which overall results in a smaller objective cost.

Regarding jamming, we observe that the optimal jamming action occurs where it is most needed. Indeed, from Figure 5, in regions of high AEE, jamming action is very small, and conversely, jamming action is high in regions of low AEE.

Multiple agents

Considering again the same sensor configuration as before, we now introduce a secondary UAV with jamming capabilities, but with the first UAV not being able to jam. We assume that both UAVs start from the same point, although the new second UAV is set to finish at a different point to the first. The aim is to determine how the two UAVs interact as to both reach their respective targets. We choose as the objective a sum of four terms, two of which relate to the cumulative AEE for each agent via (3), and another two being rescaled versions of the target time penalties, set to be at equal times. We observe the resulting optimal paths for this problem in Figure 5a and the value of the chosen error function along the path in Figure 5b.

It is easy from Figure 5b to see that the presence of a secondary agent “far” from the field within a region of high AEE can have a substantial effect on the AEE environment of an agent trying to reach a target within an adversary environment. Furthermore, it can be seen that the secondary agent enacts the most jamming when the primary agent (with no jamming capabilities) enters the lowest region of AEE, in an effort to aid the agent in spoofing the sensor operators.

Intermittent communications

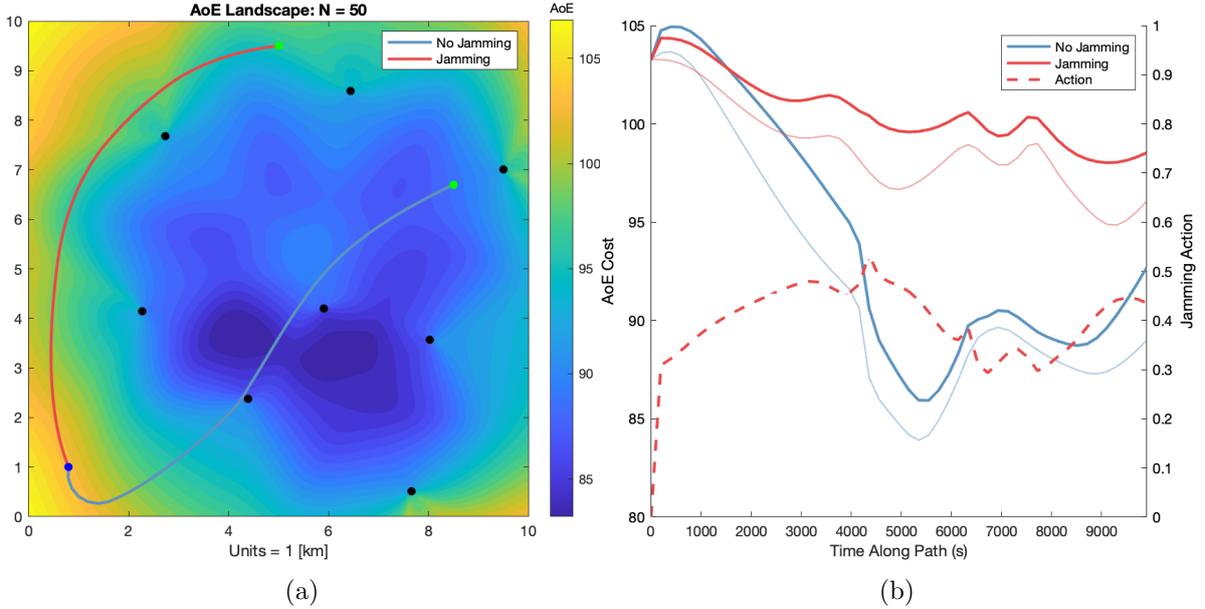


Figure 5: (a) Computed optimal paths (lowest cumulative AEE) between the starting point (0.8, 1) and the target locations (8.5, 6.7) and (5, 9.5) in a scenario with eight sensors and two agents with one having jamming capabilities (b) AEE along the optimal path(s) as a solid line, and the jamming action as a dashed line. Lighter coloured lines represent the AEE of the agents in the absence of jamming.

In this section we consider four sensors at positions (-4,10), (2, 18), (6, 8), and (1, 5) and investigate the effect of intermittent communications on the optimal paths, i. e. where the UAV reduces their emission by a certain amount for a certain amount of time; see Section 3.2 for details.

In our numerical experiments, it was always the case that the emission of the agent was reduced fully to 0% for an interval of time. We did not observe reductions of less than the full amount, nor was the time over which the reduction occurred split into more than one interval.

Figure 6 shows optimal paths for the problem of minimising detection probability for different budgets of communication dropout. A reference solution where communication is not allowed to be intermittent is depicted in blue. Parts of the path where emissions are reduced to zero are depicted in red. As it can be seen, the blue path avoids the sensor at (1, 5) and the sensor at (0, 8) by about the same distance, which is driven by the travel time budget (see earlier examples in this chapter). With a communication reduction budget of 17%, the agent now gets closer to (1, 5) for a certain amount of time, which allows keeping a wider distance to (0, 8) later. The effect is even more pronounced when the communication reduction budget is increased to 35%.

4.5 An example scenario

We present now a simplified experiment that one might encounter in the real world, a form of Scenario (B) as described in Section 2.1. In particular, we consider a UAV traversing a 14km \times 10.5km region within the Isle of Wight, on which we distribute 9 sensors (8 stationary and one mobile, denoted by y_9). We task the agent with reaching a target on the east coast of the Isle of Wight, having begun near its centre whilst optimally evading and/or spoofing detection according to some cost F . The initial positions of the sensors are known to the agent, although their future positions and any other changes in the environment are not and will be handled via the use of a rolling horizon approach, as described in Section 3.6. Furthermore, we assume that within the Isle of Wight, there are regions which the UAV(s)/agents need to avoid, such as the sea, as to depict either logistical or physical constrains; see Section 3.4.

More specifically, we assume that the agent wishes to optimise over the objective

$$F(\mathbf{x}, \mathbf{y}) = -\lambda \int_0^\tau E_2(\mathbf{x}(t), \mathbf{y}(t)) dt + M \max\{0, \tau - t_F\}^2,$$

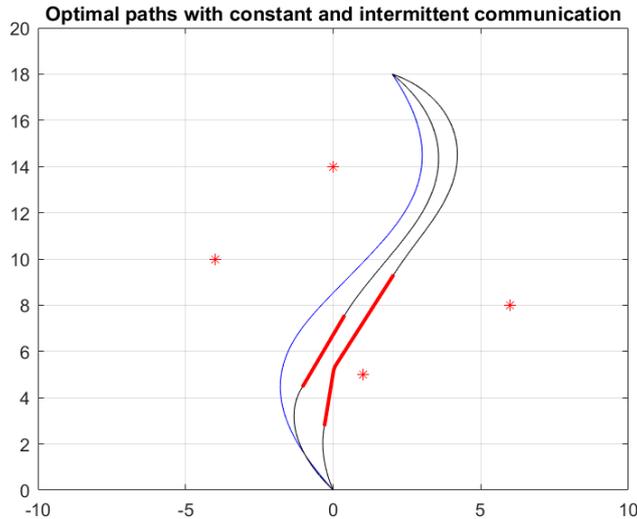


Figure 6: Optimal paths (minimise cumulative probability of detection) for intermittent communication with different budgets of communication dropout (black) and reference solution with full communication (blue). Parts of the path where emissions are reduced to zero are depicted in red.

where $\lambda = 10^{-5}$, $M = 10^{-2}$, and $t_F = 11076$ seconds, i.e. we aim to maximise the cumulative AEE (given by E_2 , see (17)) computed by the sensors whilst at the same time wishing to reach our target at a specified time t_F . Furthermore, we perform the optimisation over 18 equispaced control horizons $T_i = \frac{iT}{18}$ for $i = 1, \dots, 18$ that span a maximal time of flight interval $[0, T]$ with $T = 14000$ seconds.

See Figure 7 for a depiction of the initial configuration of sensors; the start and target location of the agent are shown in blue and green, respectively, the boundaries of the forbidden regions are shown in red, sensor positions are shown as black dots. In Figure 7a, we show the initial AEE landscape of the domain, the cumulative of which we will wish to optimise over. It is easy to see that the regions of small AEE (in blue) correspond to areas where sensors can well triangulate the position of the UAV(s). On the contrary regions showing high AEE (in yellow) occur where either the sensors have little coverage or along lines where sensors are colinear; the effects of large errors due to colinearity can be seen clearly along the line of sensors at $y \approx 8$.

We perform an optimisation to determine optimal paths and equipment usage in both scenarios where the agent is not able and is able to jam sensors.

Without jamming

The optimal path(s), when no jamming capabilities are in place, over a select set of prediction horizons are depicted in Figure 8, and the corresponding probability of detection and AEE along this path is shown in Figure 9.

In Figure 8a it is clear that the presence of the two forbidden regions provides a barrier for the optimal path as it restricts access to a region of low sensor count and thus high AEE. Although, the computed path makes use of the region between the two forbidden circles and arches southwards to spend as much time as possible in an area of relatively large AEE. In Figure 8b, we observe sensor 9 reaching the region between the two forbidden circles, thereby decreasing the local AEE values of a previously relatively high AEE value region. The optimal path computed accounts for this intrusion at this prediction horizon and now arches further southwards. As the moving sensor moves away from this region, the newly acquired shape of the path remains, with minor adjustments for when sensor 9 is particularly close to the optimal path, see Figure 8d.

Recall that the cost function used in this example solely considers the AEE of the agent and a time penalty term. Therefore, high regions of probability of detection should not be unexpected. This is precisely what is observed for this example in Figure 9a. The probability of detection of the agent peaks twice throughout the journey, at 1000 and 4000 seconds, occurring when the agent is close to sensors at

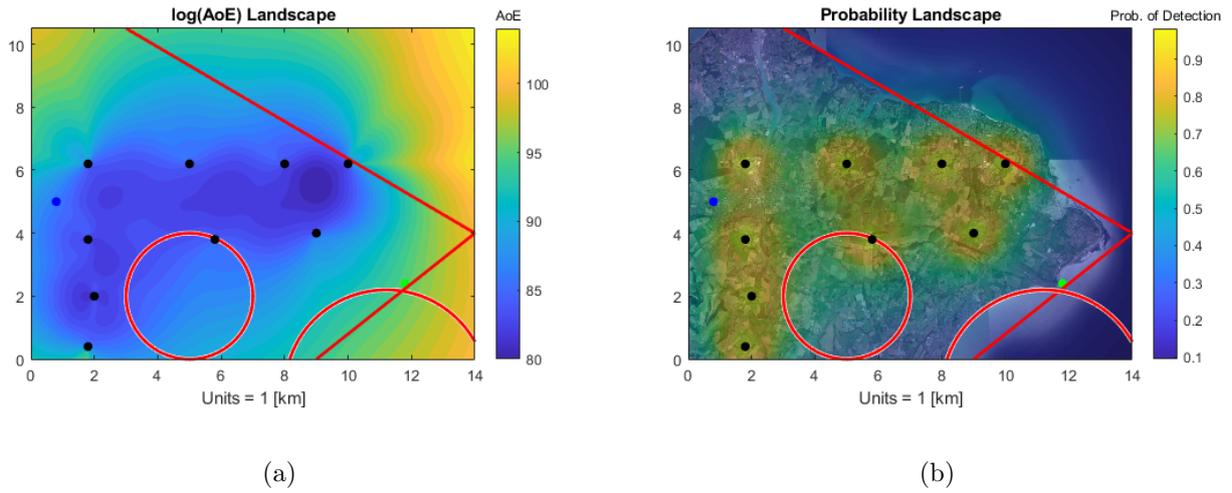


Figure 7: Initial configuration of sensors (black), along with the starting and ending positions of the agent (blue and green). Left: landscape plot of the logarithm of the area of error ellipse of the sensors. Yellow corresponds to a large area of error ellipse and thus a large error in measurement at that position, whereas blue corresponds to a small area of ellipse. Right: landscape plot of the probability of detection in the domain. The red circles denote boundaries of forbidden regions, those being inside said circles.

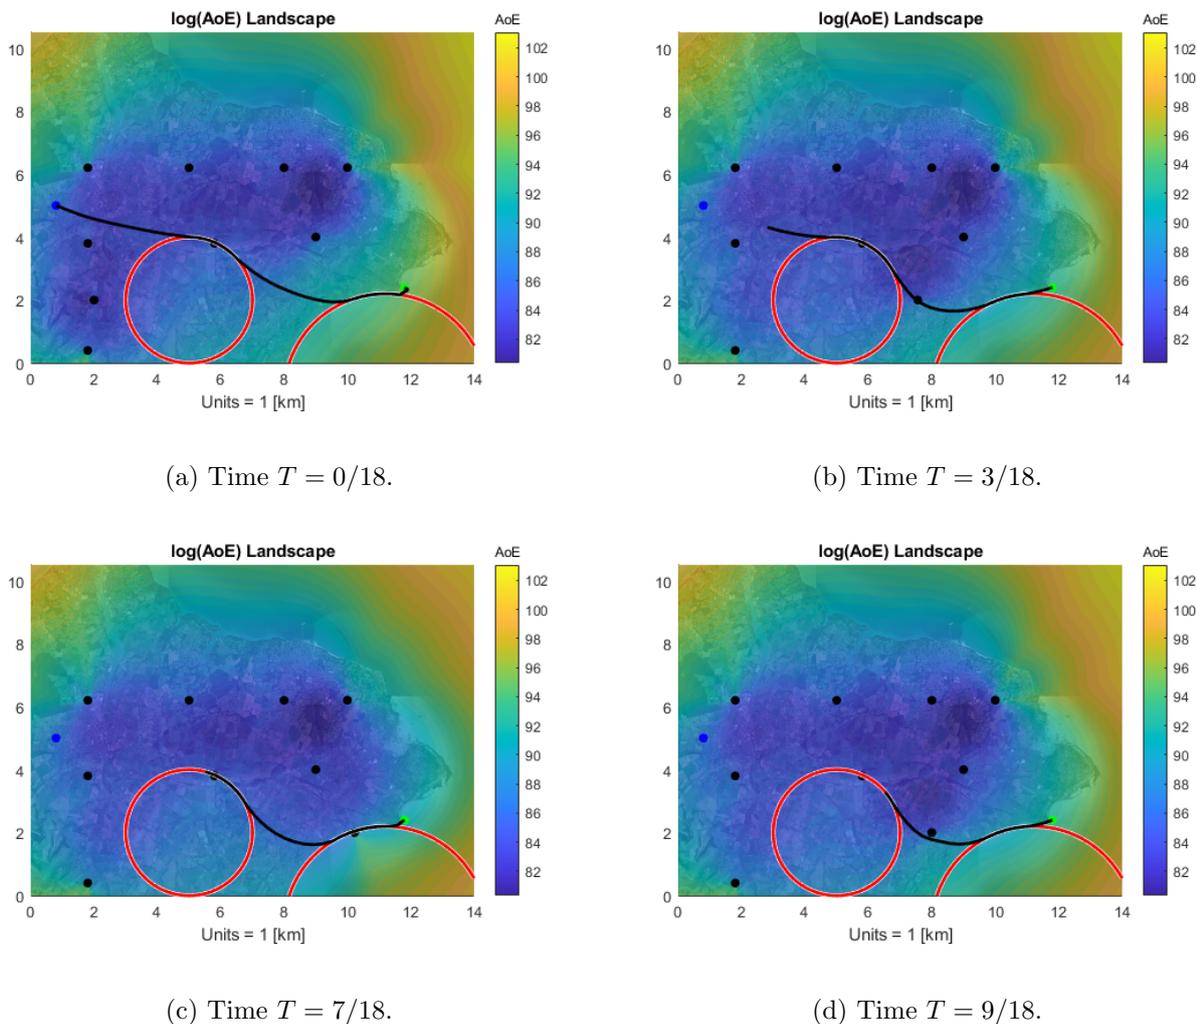


Figure 8: Optimal path computations at each control horizon for example scenario situation.

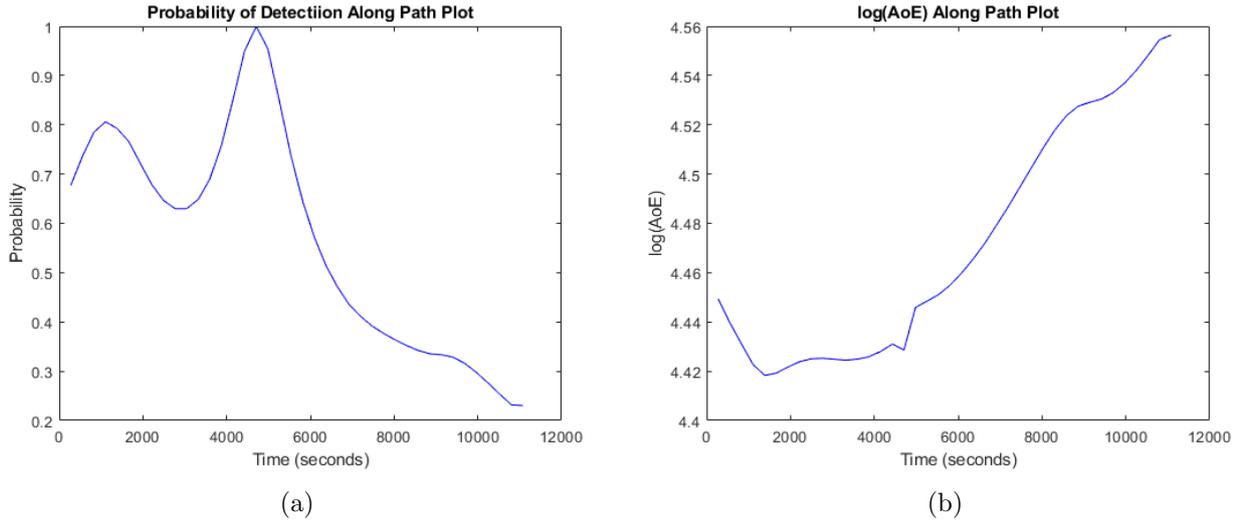


Figure 9: Plots are of quantities along the path traversed by the agent **without** jamming capabilities. (a) Probability of detection. (b) Logarithm of the AEE.

positions (2,4) and (6,4), respectively.

With jamming

The optimal path(s), when jamming capabilities are in place, over a select set of prediction horizons are depicted in Figure 10, and the corresponding probability of detection and AEE along this path is shown in Figure 11.

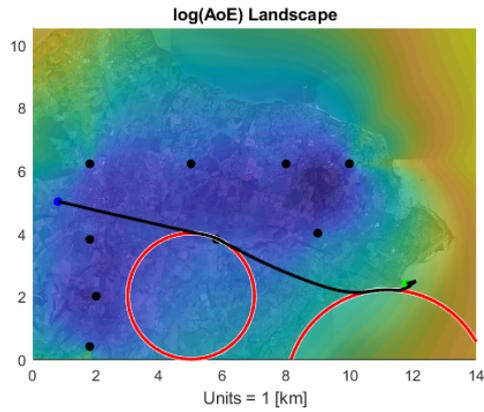
As expected jamming usage along the path allowed for a shorter and more direct trajectory of the UAV to the target, see Figure 10a. This phenomenon continues until sensor 9 has moved to the region between the two forbidden circles, where the trajectory bends southwards to avoid the newly created region of low AEE, see Figure 10b. As sensor 9 moves further eastwards, the optimal path computed recovers from the southward bend and bends northwards for a more direct path to the target, see Figure 10c, unlike the path obtained when no jamming is utilised, compare Figure 8c.

Jamming availability has also allowed the agent to reach the target with a smaller cumulative probability of detection and AEE, compared to when jamming was not in place. To achieve this, the total jamming availability was fully utilised along the path, with peaks in usage at 1000 and 7000 seconds along the path, corresponding to moments when the UAV was in regions of relatively low AEE.

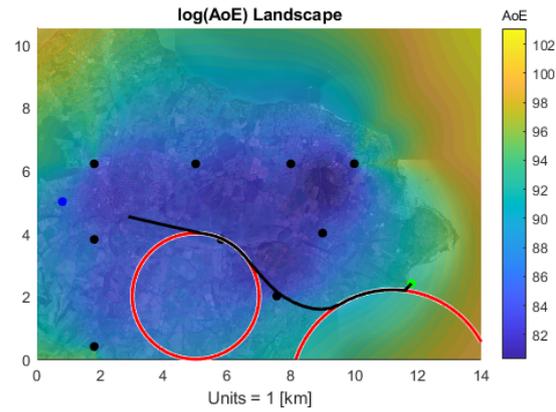
5 Conclusions

In this article, we considered the problem of evading detection and/or location of one or more agents that are tasked to traverse an adversary environment from prespecified starting and ending positions. Distributed across the area under consideration are several sensors that aim to detect and locate the agents. We have provided a flexible mathematical modelling framework by posing the problem as an *optimisation problem*. The developed model rests on a precise mathematical framework that has previously been developed in estimation theory and has widely been applied and further developed in statistical signal processing. Optimisation provides a powerful modelling paradigm that is able to provide mathematical models for a variety of situations, such as:

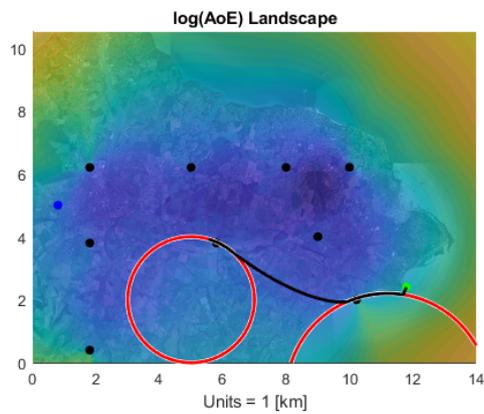
- One or more agents are under consideration, with possibly different start points and targets.
- Agents have different characteristics, e.g. different maximum velocity and capabilities.
- Several different objectives can be considered, e. g. minimise cumulative (average) probability of detection, minimise worst (smallest) probability of detection, maximise cumulative location error that the sensors experience when trying to measure the location of the agents, maximise worst (smallest) location error, minimise travel time, and others.



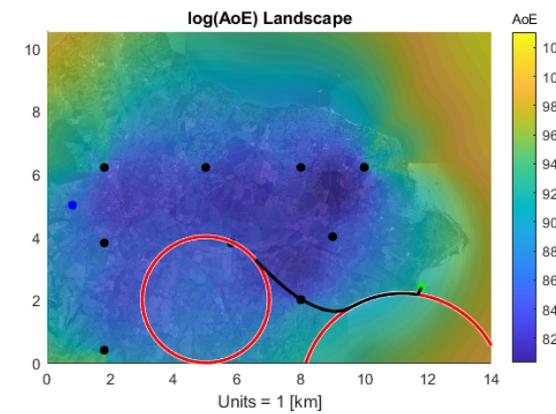
(a) Time $T = 0/18$.



(b) Time $T = 3/18$.

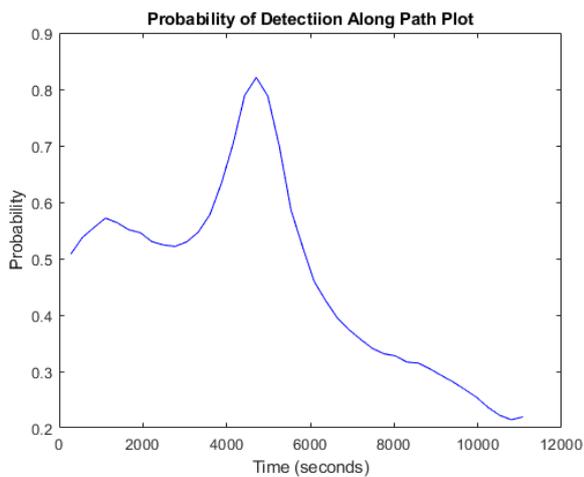


(c) Time $T = 7/18$.

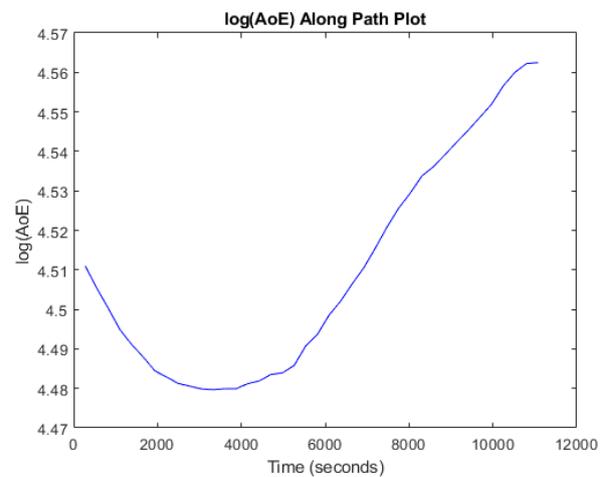


(d) Time $T = 9/18$.

Figure 10: Optimal path computations at each control horizon for the example scenario. The area of ellipse landscape shown in the background is constant in the absence of jamming.



(a)



(b)

Figure 11: Final traversed path for agent with jamming capabilities. (a) Probability of detection. (b) Logarithm of the AEE.

- We also consider the case that these objectives can change over time.
- Our framework allows for various additional constraints like energy constraints (battery power), regions forbidden to travel into, and others.
- The basic optimisation model developed can be readily extended to a number of important scenarios, in particular moving sensors, jamming (both isotropic and anisotropic), and intermittent communications.

Finally, we provide a variety of numerical test results based on a prototype implementation of the model developed. These results do not only confer the flexibility of the chosen modelling approach, but also show that the framework developed is computationally tractable and leads to realistic solutions.

We also note that several possible extensions to our work that can be made subject to further study:

- We might want to consider movement dynamics of the agents, beyond bounds on velocity and acceleration. Modelling such dynamics might become important, depending on the spatial scale of the problem.
- We might want to consider inhomogeneous terrain in which the speed of the agents in a particular direction can differ from speeds in other directions, and in which different parts of the underlying area possess different types of inhomogeneity.
- Likewise, we consider constant pathloss coefficients for the sensors only and do not model any influence of the terrain on location error or probability of detection.
- Our continuous model considers locally optimal trajectories only. This is partly alleviated by considering a hybridisation approach that merges a discrete approach with the continuous approach here. However, more work is needed to better understand the interplay of the two approaches and their advantages and disadvantages.
- In our framework, there are no actions the sensor operator can take while the agents move: for example, all sensor movements are prespecified. A more appropriate framework would be a bilevel optimisation setting, in which the sensor operator decides on the locations and action plans of the sensors first, before the UAV operator decides on their actions.
- The study performed is two dimensional in nature and thus probability and error landscapes considered are also two dimensional. Of course, extensions to the three dimensional case are expected to be straightforward.

Nevertheless, based on encouraging numerical tests, we believe that the optimisation model developed is generic, flexible, and versatile enough to warrant further investigations into its practical usage.

Declarations

Ethics approval

Not applicable.

Consent to participate

All the authors have consented to participate in this research.

Consent for publication

All the authors have given their consent for the paper to be submitted for publication.

Availability of data and material

The codes and data use for the experiments can be made available as necessary.

Funding and acknowledgements

This article is the outcome of a study first considered during the *Mathematical Challenges in the Electromagnetic Environment Study Group* event, hosted by the Newton Gateway, University of Cambridge, held between 6-8th January 2021, and subsequently reported in [3]. The support of DSTL and PA Consulting are gratefully acknowledged.

References

- [1] Andersson, J.A., Gillis, J., Horn, G., Rawlings, J.B., Diehl, M., 2019. Casadi: a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation* 11, 1–36.
- [2] Aquanallah, N., Khalfa, A., 2017. Analysis study of radar probability of detection for fluctuating and non-fluctuating targets. *Algerian Journal of Signals and Systems* 2, 12–20.
- [3] Atayev, A., Delaosa, C., 2021. Escape the Sensor. Technical Report. Challenge 2 of *Mathematical Challenges in the Electromagnetic Environment*, <https://warwick.ac.uk/fac/sci/masdoc/people/studentpages>.
- [4] Beal, L., Hill, D., Martin, R., Hedengren, J., 2018. Gekko optimization suite. *Processes* 6, 106. doi:10.3390/pr6080106.
- [5] Biegler, L., 2010. Nonlinear programming, mos-siam series on optimization. SIAM, Philadelphia .
- [6] Breloy, A., Ginolhac, G., Pascal, F., Forster, P., 2016. Robust covariance matrix estimation in heterogeneous low rank context. *IEEE Transactions on Signal Processing* 64, 5794–5806. doi:10.1109/TSP.2016.2599494.
- [7] Chand, B.N., Mahalakshmi, P., Naidu, V., 2017. Sense and avoid technology in unmanned aerial vehicles: A review, in: 2017 International Conference on Electrical, Electronics, Communication, Computer, and Optimization Techniques (ICEECCOT), IEEE. pp. 512–517.
- [8] Da Fonseca, C., Petronilho, J., 2001. Explicit inverses of some tridiagonal matrices. *Linear Algebra and its Applications* 325, 7–21.
- [9] Dieumegard, P., Cafieri, S., Delahaye, D., Hansman, R.J., 2023. Rotorcraft low-noise trajectories design: black-box optimization using surrogates. *Optimization and Engineering* 24, 2475–2512.
- [10] Duo, B., Wu, Q., Yuan, X., Zhang, R., 2020. Anti-jamming 3d trajectory design for uav-enabled wireless sensor networks under probabilistic los channel. *IEEE Transactions on Vehicular Technology* 69, 16288–16293.
- [11] Elias, F.G., Fernández, E.M., 2021. An analysis of energy detector based on improved approximations of the chi-square distributions. *EURASIP Journal on Wireless Communications and Networking* 2021, 1–18.
- [12] Findeisen, R., Allgöwer, F., Biegler, L.T., 2007. Assessment and future directions of nonlinear model predictive control. volume 358. Springer.
- [13] Gezici, S., 2008. A survey on wireless position estimation. *Wireless Personal Communications* 44, 263–282.
- [14] Hoch, B., Liers, F., 2023. An integrated rolling horizon and adaptive-refinement approach for disjoint trajectories optimization. *Optimization and Engineering* 24, 1017–1055.
- [15] Hui, C., Fengqi, W., Chongzhao, H., 2022. Uav path planning strategy based on threat avoidance in multiple extended target tracking optimization. *Journal of Radars* 11, 1–12.
- [16] Jun, M., D’Andrea, R., 2003. Path planning for unmanned aerial vehicles in uncertain and adversarial environments, in: *Cooperative control: models, applications and algorithms*. Springer, pp. 95–110.
- [17] Kabamba, P.T., Meerkov, S.M., Zeitz III, F.H., 2006. Optimal path planning for unmanned combat aerial vehicles to defeat radar tracking. *Journal of Guidance, Control, and Dynamics* 29, 279–288.
- [18] Kay, S.M., 1993. *Fundamentals of statistical signal processing, Volume II: Detection Theory*. Prentice Hall PTR.
- [19] KNITRO, (Accessed on 7 October 2023). Software. <https://www.artelys.com/solvers/knitro/>.
- [20] Lanzisera, S., Pister, K.S., 2008. Burst mode two-way ranging with cramer-rao bound noise performance, in: 2008 Global Telecommunications Conference, IEEE. pp. 1–5.

- [21] Pelosi, M., Kopp, C., Brown, M., 2012. Range-limited uav trajectory using terrain masking under radar detection risk. *Applied Artificial Intelligence* 26, 743–759.
- [22] Rawlings, J.B., Mayne, D.Q., Diehl, M., et al., 2017. *Model predictive control: theory, computation, and design*. volume 2. Nob Hill Publishing Madison, WI.
- [23] Ruz, J.J., Arévalo, O., Jesús, M., Pajares, G., 2006. Using milp for uavs trajectory optimization under radar detection risk, in: *2006 Conference on Emerging Technologies and Factory Automation*, IEEE. pp. 957–960.
- [24] Ruz, J.J., Pajares, G., Jesus, M., Arevalo, O., 2009. *UAV trajectory planning for static and dynamic environments*. INTECH Open Access Publisher.
- [25] Torrieri, D.J., 1984. Statistical theory of passive location systems. *IEEE Transactions on Aerospace and Electronic Systems* AES-20, 183–198. doi:10.1109/TAES.1984.310439.
- [26] Waghmare, R.G., Nalbalwar, S.L., Das, A., 2012. Transient signal detection on the basis of energy and zero crossing detectors. *Procedia Engineering* 30, 129–134.
- [27] Wang, D., Yu, H., Wu, Z., Wang, C., 2017. Performance analysis of the direct position determination method in the presence of array model errors. *Sensors* 17, 1550. doi:10.3390/s17071550.
- [28] Wu, Y., Guan, X., Yang, W., Wu, Q., 2021. Uav swarm communication under malicious jamming: joint trajectory and clustering design. *IEEE Wireless Communications Letters* 10, 2264–2268.
- [29] Xu, Y., Ren, G., Chen, J., Zhang, X., Jia, L., Feng, Z., Xu, Y., 2019. Joint power and trajectory optimization in uav anti-jamming communication networks, in: *2019 IEEE International Conference on Communications*, IEEE. pp. 1–5.
- [30] Zhang, Z., Wu, J., Dai, J., He, C., 2020. Rapid penetration path planning method for stealth uav in complex environment with bb threats. *International Journal of Aerospace Engineering* 2020, 1–15.
- [31] Zhang, Z., Wu, J., Dai, J., He, C., 2022. Optimal path planning with modified a-star algorithm for stealth unmanned aerial vehicles in 3d network radar environment. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering* 236, 72–81.
- [32] Zhong, C., Yao, J., Xu, J., 2018. Secure uav communication with cooperative jamming and trajectory control. *IEEE Communications Letters* 23, 286–289.
- [33] Zhou, X., Wu, Q., Yan, S., Shu, F., Li, J., 2019a. Uav-enabled secure communications: Joint trajectory and transmit power optimization. *IEEE Transactions on Vehicular Technology* 68, 4069–4073.
- [34] Zhou, X., Yan, S., Hu, J., Sun, J., Li, J., Shu, F., 2019b. Joint optimization of a uav’s trajectory and transmit power for covert communications. *IEEE Transactions on Signal Processing* 67, 4276–4290.