

Jointly-optimized Trajectory Generation and Camera Control for 3D Coverage Planning

Savvas Papaioannou, Panayiotis Kolios, Theocharis Theocharides,
Christos G. Panayiotou and Marios M. Polycarpou

Abstract—This work proposes a jointly optimized trajectory generation and camera control approach, enabling an autonomous agent, such as an unmanned aerial vehicle (UAV) operating in 3D environments, to plan and execute coverage trajectories that maximally cover the surface area of a 3D object of interest. Specifically, the UAV’s kinematic and camera control inputs are jointly optimized over a rolling planning horizon to achieve complete 3D coverage of the object. The proposed controller incorporates ray-tracing into the planning process to simulate the propagation of light rays, thereby determining the visible parts of the object through the UAV’s camera. This integration enables the generation of precise look-ahead coverage trajectories. The coverage planning problem is formulated as a rolling finite-horizon optimal control problem and solved using mixed-integer programming techniques. Extensive real-world and synthetic experiments validate the performance of the proposed approach.

Index Terms—Planning and control, coverage, optimization.



NOMENCLATURE

$x_t \in \mathcal{X}$	UAV state at time-step t .
$u_t \in \mathcal{U}$	UAV kinematic and camera control inputs at time-step t .
$x_{t' t}$	Predicted UAV state for time-step t' computed at time-step t .
$u_{t' t}$	Predicted control inputs for time-step t' computed at time-step t .
$\mathcal{C}_{t' t}$	Predicted UAV camera state (i.e., FOV vertices) for time-step t' .
$\mathcal{V}_{t' t}$	The convex hull of the camera FOV at the future time-step t' .
$\mathcal{L} = \{\Lambda_1, \dots, \Lambda_n\}$	The set \mathcal{L} of n light-rays $\Lambda_i, i \in \{1, \dots, n\}$.
$\mathcal{K} = \{\kappa_1, \dots, \kappa_{ \mathcal{K} }\}$	The surface of the object of interest is represented as a mesh \mathcal{K} composed of triangular facets $\kappa_i, i \in \{1, \dots, \mathcal{K} \}$.
$\bar{\mathcal{A}} = \{\varpi_1, \dots, \varpi_{ \bar{\mathcal{A}} }\}$	The environment $\mathcal{A} \subset \mathbb{R}^3$ is discretized into a 3D grid $\bar{\mathcal{A}}$ composed of $ \bar{\mathcal{A}} $ non-overlapping cells ϖ_i .
T	The length of the planning horizon in time-steps.
$\mathcal{M} = \Theta \times \Phi \times Z$	The set of all admissible camera configurations i.e., a combination of camera rotations (Θ, Φ) and camera zoom levels (Z).
$b_{\varpi, \hat{\kappa}} \in \{0, 1\}$	Decision variable indicating whether there exists a light-ray that can directly trace back to facet κ (identified by index $\hat{\kappa}$) when the UAV is located within cell ϖ (identified by index $\hat{\varpi}$).

$b_{\varpi, t' t}^{\bar{\mathcal{A}}} \in \{0, 1\}$	Decision variable indicating whether at the future time-step t' the UAV is located within cell ϖ (identified by index $\hat{\varpi}$).
$b_{\hat{\kappa}, \hat{m}, t' t}^{\mathcal{V}} \in \{0, 1\}$	Decision variable indicating whether at time-step t' the facet κ (identified by index $\hat{\kappa}$) resides within the convex-hull of the $\hat{m}_{th} \in \mathcal{M} $ camera FOV state.
$s_{\hat{m}, t' t} \in \{0, 1\}$	Decision variable that indicates whether at time-step t' the \hat{m}_{th} camera FOV state is active.
$\bar{b}_{\hat{\kappa}, \hat{m}, t' t} \in \{0, 1\}$	Decision variable indicating whether at time-step t' the facet κ (identified by index $\hat{\kappa}$) is visible through the \hat{m}_{th} camera FOV state.
$\mathcal{Q}(\kappa) \in \{0, 1\}$	UAV memory defined as a function $\mathcal{Q} : \mathcal{K} \rightarrow \{0, 1\}$, where $\mathcal{Q}(\kappa) = 1$ indicates that facet κ has been covered at some previous time-step $\tau < t$.

1 INTRODUCTION

THE interest in unmanned aerial vehicles (UAVs) has skyrocketed over the last decade. The latest advancements in robotics, automation, and artificial intelligence, coupled with the widespread adoption of small consumer aerial drones, have spurred unprecedented interest in UAV-based applications and services. Today, UAVs have the potential to be utilized across a wide range of application domains, including surveillance and security [1]–[5], automated inspection [6]–[9], and emergency response missions [10]–[15]. Automated coverage planning is one of the key functionalities that can significantly enhance the autonomy of existing unmanned aerial vehicles (UAVs), enabling them to execute fully automated missions across the diverse application scenarios mentioned earlier. In coverage path planning (CPP) [16], [17], also referred to as coverage trajectory planning, the goal is to determine a path that allows an autonomous mobile agent to fully cover a specific area or object of interest

- The authors are with the KIOS Research and Innovation Centre of Excellence (KIOS CoE) and the Department of Electrical and Computer Engineering, University of Cyprus, Nicosia, 1678, Cyprus. E-mail: {papaioannou.savvas, pkolios, ttheocharides, christosp, mpolycar}@ucy.ac.cy

while optimizing a particular mission objective, such as minimizing mission elapsed time.

Despite the extensive range of CPP approaches proposed in the literature, no dominant solution exists for autonomous UAV-based coverage planning in realistic 3D environments. Current state-of-the-art UAV-based CPP techniques [16] are primarily focused on covering planar areas (i.e., 2D terrain coverage) rather than 3D objects. Moreover, existing approaches (e.g., [18], [19]) typically assume UAVs equipped with fixed, downward-facing cameras, thereby excluding the joint optimization and control of the UAV's kinematics and camera settings during planning. This assumption simplifies the coverage planning problem, reducing it to a path-planning problem [20], [21]. Additionally, most CPP techniques achieve area coverage using simple geometric patterns (e.g., back-and-forth, zig-zag, or spiral motions) [22], which often lack the flexibility to generalize to 3D settings. In this work, we argue that a practical UAV-based coverage planning solution must: (a) operate effectively in true 3D environments (i.e., producing coverage plans for 3D objects rather than planar regions), and (b) generate coverage plans in an online fashion. Modern UAVs are typically equipped with controllable cameras featuring rotation and zoom capabilities, necessitating coverage planning techniques that optimize not only the UAV's motion plan but also the control and management of the camera for optimal coverage.

Motivated by these challenges, we propose a jointly optimized trajectory generation and camera control approach that enables a UAV agent to generate online trajectories for covering the surface area of a 3D object of interest. Specifically, we frame the coverage planning problem as a rolling finite horizon optimal control problem (FHOCPP), jointly optimizing the UAV's kinematic control inputs and its camera control inputs (e.g., camera rotation and zoom levels) to produce collision-free coverage trajectories. To facilitate the generation of look-ahead coverage trajectories over a finite planning horizon, the proposed approach incorporates ray-tracing to simulate light rays captured through the UAV's camera. This enables the anticipation of the agent's behavior over the planning horizon and the generation of trajectories that account for the parts of the object visible within a finite set of planned control inputs. The key contributions of this work are as follows:

- We propose a jointly optimized trajectory generation and camera control approach for 3D coverage planning, enabling an autonomous UAV agent to efficiently generate optimal coverage trajectories for 3D objects of interest.
- We demonstrate how ray-tracing can be integrated into the coverage planning process to identify the visible parts of the object through the agent's camera, thereby enabling the generation of look-ahead coverage trajectories.
- Finally, we illustrate how this problem can be formulated as a rolling horizon optimal control problem and subsequently solved using mixed-integer programming. The effectiveness of the proposed approach is validated through extensive synthetic and real-world experiments.

This paper is structured as follows. Section 2 discusses the related work on coverage path planning with autonomous agents. Section 3 outlines our modeling assump-

tions, Section 4 formulates the problem, and Section 5 discusses the proposed 3D coverage approach. Finally, Section 6 evaluates the proposed approach, and Section 7 concludes the paper.

2 RELATED WORK

Initial works on coverage path planning (CPP) primarily focused on ground robots operating in 2D environments [23]. Most of these approaches typically involve: (a) decomposing the area of interest into a set of non-overlapping cells, and (b) solving an optimization problem to determine the robot's path that traverses all the cells. In [24], [25], the authors propose coverage planning techniques that enable the robot to cover the region of interest while simultaneously constructing the cell decomposition. In [26], an exact cellular decomposition coverage planning approach based on the detection of natural landmarks is presented, while [27] introduces a spanning-tree-based CPP approach for a ground robot equipped with a fixed sensor. Additionally, [28] investigates the CPP problem for robots with energy constraints. More recently, a coverage planning approach for 2D obstacle-cluttered environments was proposed in [29]. This work introduces a hierarchical, hex-decomposition-based coverage planning algorithm that ensures complete coverage of the area of interest.

The coverage planning problem has also been extended to multi-robot systems [19]. For example, in [30], the authors propose a multi-robot spanning-tree coverage algorithm for the coverage of a bounded planar area, using multiple robots equipped with fixed sensors. In [31], the area of interest is divided among multiple robots, which then apply ant-colony optimization to determine the coverage path for their assigned regions. Similarly, [32] introduces a multi-robot coverage planning approach based on the capacitated arc-routing problem, constructing coverage paths while accounting for the robots' energy capacity limitations. The authors of [33] proposed a boustrophedon cellular decomposition algorithm to address the coverage problem for multiple robots under communication constraints. Meanwhile, the work in [34] employs graph neural networks to design a multi-robot coverage system.

More recently, the coverage path planning problem has been explored for UAV-based systems and applications [35]–[37]. Specifically, the coverage of polygonal 2D areas using UAVs is investigated in [38], [39] and, more recently, in [40]. In particular, [40] proposes a UAV-based CPP approach utilizing the traveling salesman problem to cover multiple 2D polygonal regions. In [41], the CPP problem is formulated as an optimal control problem for a UAV operating in a bounded 2D environment, while [42] presents a coverage planning approach tailored to cuboid-like structures. The work in [43] employs mathematical programming techniques to command a team of UAVs for area coverage in the minimum amount of time. In [44], a UAV-based cooperative CPP approach is introduced, using the simulated annealing algorithm to account for each UAV's sensing and operational capabilities, initial positions, and no-fly zones. A multi-UAV CPP approach for 2D terrain coverage is proposed in [45], aiming to minimize completion time by balancing the workload across UAVs. Additionally, [46] introduces a cell decomposition algorithm based on regular hexagons for multi-UAV area coverage. Finally, the approaches in [47]–[49] explore how the CPP problem can

be framed as a learning problem and solved using reinforcement learning.

It is important to note here that recent advancements in 3D monitoring, communications and trajectory planning using single and multiple UAVs have led to significant developments across various applications. In [50], the authors propose a UAV-based 3D monitoring framework for accurately estimating the 3D geometric features of agricultural trees, providing valuable insights for precision agriculture and crop management optimization. In [51], the authors present a UAV-based system for 3D monitoring and inspection of dams using photogrammetry for damage detection. A recent survey on UAV-based 3D mapping and monitoring is available in [52]. The work in [53] introduces a 3D monitoring and video surveillance approach using fixed-wing UAVs to track suspicious mobile targets, formulating the problem as a multi-objective optimization to balance monitoring performance and power efficiency. The study in [54] explores various strategies and methodologies for monitoring and patrolling missions with multiple UAV agents, while [55] reviews recent developments in UAV-based monitoring applications within transportation.

Regarding multi-UAV communication and trajectory planning, [56] investigates the problem of energy-efficient trajectory planning for UAV-based content coverage. A UAV-aided mobile edge computing framework is proposed in [57], where multiple UAVs with distinct trajectories fly over the target area to support user equipment on the ground. In [58], the authors design a multi-UAV-enabled wireless communication system to serve a group of ground users. The joint trajectory planning and communication problem is formulated as a mixed-integer non-convex optimization problem and solved using an efficient iterative algorithm based on block coordinate descent and successive convex optimization. Other related work [59]–[61] focuses on minimizing the completion time of UAV tasks through joint communication optimization and trajectory planning, particularly considering energy costs.

In summary, compared to existing CPP approaches, this work proposes a UAV-based coverage planning technique specifically designed for 3D environments, where both the agent and the object to be covered exist in three-dimensional space. The proposed approach jointly optimizes the UAV's kinematic and camera control inputs to generate optimal coverage trajectories. Additionally, it addresses the visibility determination problem by incorporating ray-tracing into the coverage planning process. This integration simulates the physical behavior of light rays captured by the UAV's camera, enabling the generation of look-ahead coverage trajectories over a future planning horizon. The CPP problem is formulated as a rolling finite-horizon optimal control problem (FHOCPP), which is then transformed into a mixed-integer program (MIP) and solved using existing optimization tools.

3 PRELIMINARIES

3.1 Agent Kinematics

In this work we consider a mobile UAV agent with state at time-step t , given by $x_t \in \mathbb{R}^6$, operating within a bounded 3D surveillance area $\mathcal{A} \subset \mathbb{R}^3$ according to the discrete-time dynamical model shown below:

$$x_t = \begin{bmatrix} \mathbf{I}_{3 \times 3} & \Delta t \mathbf{I}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & (1 - a) \mathbf{I}_{3 \times 3} \end{bmatrix} x_{t-1} + \begin{bmatrix} \mathbf{0}_{3 \times 3} \\ \frac{\Delta t}{m} \mathbf{I}_{3 \times 3} \end{bmatrix} f_{t-1}, \quad (1)$$

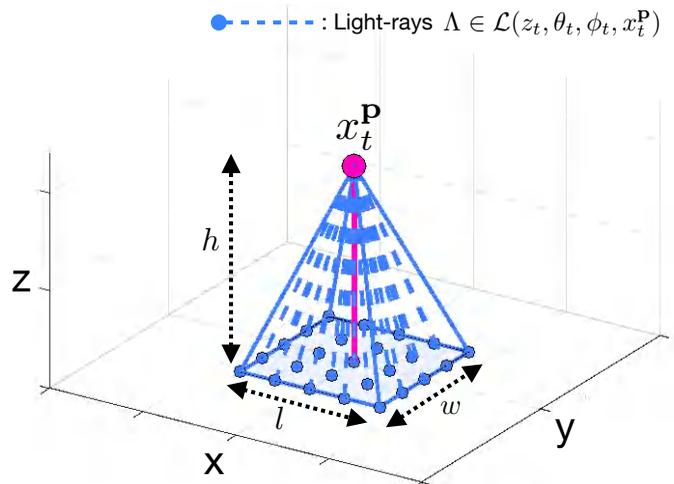


Fig. 1. The camera's FOV (with horizontal and vertical FOV angles given by φ_h and φ_v respectively) is represented by a regular right pyramid with height h , and a rectangular base of dimensions (l, w) . At each time-step t a finite set of light-rays $\mathcal{L}(z_t, \theta_t, \phi_t, x_t^P) = \{\Lambda_1, \dots, \Lambda_n\}$ enter the camera's FOV (shown with the dotted line segments) as discussed in Sec. 3.2

where the state of the agent $x_t \in \mathcal{X}$ is composed of the 3D position, and 3D velocity vectors in Cartesian coordinates, denoted hereafter as $x_t^P \in \mathbb{R}^{3,1}$, and $x_t^V \in \mathbb{R}^{3,1}$ respectively. The input vector $f_t \in \mathcal{F} \subset \mathbb{R}^{3,1}$ denotes the applied control force (i.e., the kinematic control input) which allows the agent to change its direction and speed according to the mission requirements. The parameter a is used to model the air resistance, m denotes the agent's mass, and Δt is the sampling interval. Finally, the 3-by-3 square matrices $\mathbf{I}_{3 \times 3}$ and $\mathbf{0}_{3 \times 3}$ denote the identity and zero matrices respectively. The proposed coverage controller, uses the kinematic model in Eq. (1) to construct the agent's coverage plan, which in turn is used as the reference trajectory to be tracked with a flight controller (e.g., an auto-pilot) [62]–[65] depending on the UAV's aerodynamical characteristics.

3.2 Agent Camera Model

We assume that the UAV agent carries onboard a gimballed camera with rotation and zoom functionalities, which is used for viewing and sensing the surrounding environment. The shape of the camera's field-of-view (FOV) is represented in this work by a regular right pyramid which is composed of four triangular lateral faces and a rectangular base. The camera's optical center is assumed to be located at the apex which is positioned directly above the centroid of the FOV base. The FOV footprint and range is determined in this work by the parameter set (l, w, h) , where the pair of parameters (l, w) defines the size of the FOV's rectangular base, and h (i.e., the pyramid height) is the camera's observation range, as depicted in Fig. 1. Subsequently, the vertices of the camera's FOV for an agent centered at the origin of the 3D cartesian coordinate frame, assuming that the camera is facing downwards, is given by the 3-by-5 matrix \mathcal{C}_0 as:

$$\mathcal{C}_0 = \begin{bmatrix} -l/2 & l/2 & l/2 & -l/2 & 0 \\ w/2 & w/2 & -w/2 & -w/2 & 0 \\ -h & -h & -h & -h & 0 \end{bmatrix}. \quad (2)$$

As previously mentioned the agent's camera is equipped with optical zoom capabilities, which can be used to alter the

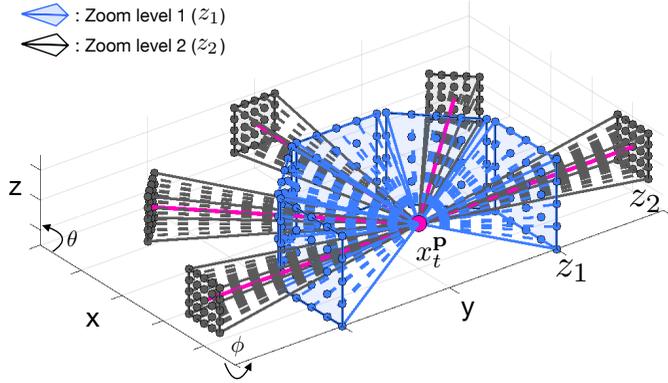


Fig. 2. The figure demonstrates various configurations of the camera's FOV at time t with respect to the control inputs z_t, θ_t, ϕ_t (i.e., zoom-level and rotation angles). In this figure the $z_t \in \{z_1, z_2\}$, $\theta_t \in \{\frac{\pi}{2}\}$, and $\phi_t \in \{0, \frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}, \pi\}$ as shown.

FOV characteristics. Specifically, the optical zoom-in functionality increases the camera's observation range (i.e., distant objects can become observable through magnification), however this operation reduces the FOV footprint size. A particular zoom-level z which belongs to the finite set of all supported zoom-levels that can be applied to the camera i.e., $z \in Z = \{z_1, \dots, z_{|Z|} | z_i \in \mathbb{R}, z_i \geq 1\}$ (the notation $|\cdot|$ denotes the set cardinality), alters the FOV parameters as follows:

$$h' = h \times z, \text{ and, } (l', w') = (l, w) \times z^{-1}, \quad (3)$$

where (l', w', h') are the modified FOV parameters after applying the optical zoom level z . Consequently, the FOV matrix \mathcal{C}_0 is in fact a function of the applied zoom-level z i.e., $\mathcal{C}_0(z)$. A zoom-level with value of $z = 1$ does not alter the FOV characteristics, and thus describes the original FOV state without zoom. Zoom-levels with values $z > 1$, increase the observation range h , and scale down the FOV footprint size (l, w) , as shown in Eq. (3).

Additionally, the camera's FOV can be rotated in 3D space by commanding the onboard camera controller to execute sequentially two elemental rotations i.e., one rotation by angle $\theta \in [0, \pi)$ around the y -axis, followed by a rotation $\phi \in [0, 2\pi)$ around the z -axis. Subsequently, at time t the agent can rotate and point the FOV of its camera anywhere inside the 3D surveillance space, as shown in Fig. 2, by applying the following geometric transformation:

$$\tilde{\mathcal{C}}_t(z_t, \theta_t, \phi_t)_i = R_\phi(\phi_t)R_\theta(\theta_t)\mathcal{C}_0(z_t)_i, \forall i \in \{1, \dots, 5\} \quad (4)$$

where $\mathcal{C}_0(z_t)_i$ is the i th column of the matrix $\mathcal{C}_0(z_t)$, $\tilde{\mathcal{C}}_t(z_t, \theta_t, \phi_t)_i$ is the rotated FOV vertex, and the rotation angles θ_t and ϕ_t are the camera controller's input signals at time t . The 3D rotation matrices $R_\theta(\theta_t)$ and $R_\phi(\phi_t)$ are further given by:

$$R_\theta(\theta_t) = \begin{bmatrix} \cos(\theta_t) & 0 & \sin(\theta_t) \\ 0 & 1 & 0 \\ -\sin(\theta_t) & 0 & \cos(\theta_t) \end{bmatrix}, \quad (5)$$

$$R_\phi(\phi_t) = \begin{bmatrix} \cos(\phi_t) & -\sin(\phi_t) & 0 \\ \sin(\phi_t) & \cos(\phi_t) & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (6)$$

We should mention that the two rotation angles θ_t and ϕ_t take their values from the finite sets of admissible rotation

angles $\Theta = \{\theta_1, \dots, \theta_{|\Theta|}\}$ and $\Phi = \{\phi_1, \dots, \phi_{|\Phi|}\}$, respectively (i.e., $\theta_t \in \Theta$ and $\phi_t \in \Phi$). To summarize, the agent's camera pose (i.e., position and orientation) in 3D space, also referred to as the camera's state, can be computed as:

$$\mathcal{C}_t(z_t, \theta_t, \phi_t, x_t^P) = \tilde{\mathcal{C}}_t(z_t, \theta_t, \phi_t) + x_t^P, \quad (7)$$

where x_t^P is the agent's position at time-step t according to Eq. (1). Finally, we consider that at each time instance t a finite-set of (straight) light-rays (indicating the propagation of light) enters the camera's optical center, causing matter to be imaged (and eventually observed). The set of light-rays captured through the agent's camera with state $\mathcal{C}_t(z_t, \theta_t, \phi_t, x_t^P)$ is denoted in this work as $\mathcal{L}(z_t, \theta_t, \phi_t, x_t^P) = \{\Lambda_1, \dots, \Lambda_n\}$, where $\Lambda_i, i = 1, \dots, n$ denotes the individual light-ray in the set which is further given by the line-segment

$$\Lambda_i = \{\lambda_i + d(x_t^P - \lambda_i) : d \in [0, 1]\}, \quad (8)$$

where x_t^P is the light-ray's end point which enters the camera's optical center at time t (denoted by the agent's position), λ_i is a fixed point on the camera's FOV base denoting the ray's origin, and d is a scalar. This is illustrated in Fig. 1.

4 PROBLEM FORMULATION

We can define the agent's joint control signal u_t at time t , to be composed of the agent's kinematic control input f_t , as well as the camera zoom-level z_t , and rotation angles (θ_t, ϕ_t) as:

$$u_t = \{f_t, z_t, \theta_t, \phi_t\}, \quad (9)$$

where we define $u_t^f = f_t$, $u_t^z = z_t$, $u_t^\theta = \theta_t$, and $u_t^\phi = \phi_t$. It is easy to see that, given a known agent state \bar{x} at time $t = 0$ i.e., $x_0 = \bar{x}$, the agent's coverage trajectory (i.e., a sequence of kinematic and camera FOV states) over a finite planning horizon of length \tilde{T} time-steps can be designed to meet specific criteria and goals (i.e., mission objectives and constraints) by suitably choosing the sequence of control inputs u_t inside the planning horizon. In particular, the agent's coverage trajectory over the planning horizon $t \in \{1, \dots, \tilde{T}\}$, can be obtained as:

$$x_t = A^t x_0 + \sum_{\tau=0}^{t-1} A^{t-\tau-1} B u_\tau^f, \quad (10a)$$

$$\mathcal{V}_{1:\tilde{T}} = \bigcup_{t=1}^{\tilde{T}} \mathcal{V}_t(u_t^z, u_t^\theta, u_t^\phi, x_t^P) \quad (10b)$$

where Eq. (10a) computes agent's sequence of kinematic states, with the matrices A , and B to contain the agent's motion model parameters according to Eq. (1) i.e.,

$$A = \begin{bmatrix} \mathbf{I}_{3 \times 3} & \Delta t \mathbf{I}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & (1-a) \mathbf{I}_{3 \times 3} \end{bmatrix}, B = \begin{bmatrix} \mathbf{0}_{3 \times 3} \\ \frac{\Delta t}{m} \mathbf{I}_{3 \times 3} \end{bmatrix}, \quad (11)$$

and \mathcal{V}_t denotes the camera's FOV convex hull generated by the FOV vertices \mathcal{C}_t as computed in Eq. (7) i.e., $\mathcal{V}_t(z_t, \theta_t, \phi_t, x_t^P) = \Delta(\mathcal{C}_t(z_t, \theta_t, \phi_t, x_t^P))$, where Δ denotes the convex hull operator. In essence, $\mathcal{V}_{1:\tilde{T}}$ computes the total area that the agent covered (i.e., observed) with its camera during its mission. Throughout the rest of the paper and depending on the context, both notations \mathcal{C}_t and \mathcal{V}_t will be used interchangeably to refer to the agent's camera state (also referred to as the FOV configuration).

Suppose now that we are given an arbitrary bounded convex object of interest $\mathcal{W} \in \mathcal{A}$, with total surface area determined by its boundary $\partial\mathcal{W}$. The problem tackled in this work can be stated as follows: *Given a sufficiently large planning horizon of length \tilde{T} time-steps, find the joint control inputs $u_t, \forall t$ over the planning horizon, which optimize the coverage objective i.e., \mathcal{G} , and allow the mobile agent to maximally cover with its camera the surface area $\partial\mathcal{W}$ of the object of interest.*

The aforementioned coverage problem is formulated in a high-level form as the optimal control problem depicted in Problem (P1). In particular, in (P1) we are interested in the sequence of control inputs $u_0, u_1, \dots, u_{\tilde{T}-1}$ which: a) optimize a given mission-specific coverage objective denoted hereafter as \mathcal{G} (e.g., minimizing the mission elapsed time), and b) allow the agent to maximally cover the surface area of the object of interest in 3D.

The constraints in Eq. (12b)-(12c) are due to the agent's kinematic model as described in Sec. 3.1. Subsequently, the constraint in Eq. (12d) implements the coverage functionality, which essentially states that the total surface area $\partial\mathcal{W}$ of the object of interest must reside inside the collective area covered by the agent's camera FOV during the mission, and finally, the constraint in Eq. (12e) makes sure that the agent avoids collisions with various obstacles $\xi \in \Xi$ in its path, including the object of interest. Lastly the constraints in Eq. (12f) make sure that the agent's state and control inputs respect the desired operational limits.

Problem (P1): Coverage Problem

$$\arg \min_{u_0, u_1, \dots, u_{\tilde{T}-1}} \mathcal{G}(x, u) \quad (12a)$$

subject to: $t \in \{1, \dots, \tilde{T}\}$

$$x_t = A^t x_0 + \sum_{\tau=0}^{t-1} A^{t-\tau-1} B u_\tau^f \quad \forall t \quad (12b)$$

$$x_0 = \bar{x} \quad (12c)$$

$$\partial\mathcal{W} \in \mathcal{V}_{1:\tilde{T}} \quad (12d)$$

$$x_0, x_t^p \notin \xi, \quad \forall \xi \in \Xi \quad \forall t \quad (12e)$$

$$x_0, x_t \in \mathcal{X}, \quad u_t \in \mathcal{U} \quad \forall t \quad (12f)$$

5 JOINTLY-OPTIMIZED TRAJECTORY GENERATION AND CAMERA CONTROL FOR 3D COVERAGE PLANNING

The coverage planning problem discussed in the previous section and shown in problem (P1) is quite challenging to be solved efficiently. In particular, observe that a feasible solution to this problem is directly coupled with the length of the planning horizon \tilde{T} . If \tilde{T} is too short, then no feasible solution may exist, while if \tilde{T} is too long then the computational complexity increases unnecessarily.

In order to bypass this problem, (P1) is transformed into a rolling finite horizon optimal control problem (FHOC), where at each time step: a) the current state of the agent $x_{t|t}$ is used as the initial state, and b) the agent's control inputs $u_{t+\tau|t}, \tau \in \{0, \dots, T-1\}$ are computed inside an arbitrary shorter finite horizon $T < \tilde{T}$. The first control input in the sequence $u_{t|t}$ is then applied to the agent, the agent moves to its new state, and the optimization problem is repeated for the next time step. This optimization procedure

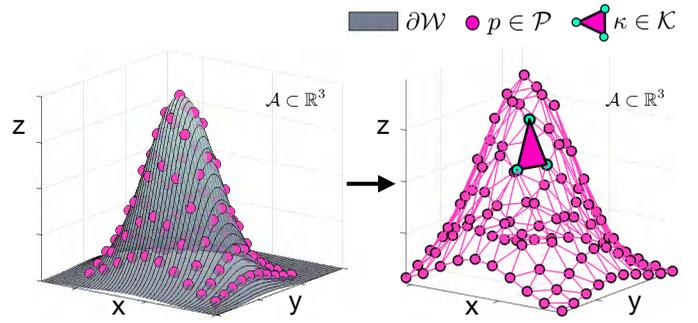


Fig. 3. A 3D point-cloud representation $\mathcal{P} = \{p_1, \dots, p_{|\mathcal{P}|}\}$ of the object's surface $\partial\mathcal{W}$ is extracted, and then triangulated to form a triangle mesh \mathcal{K} consisting of triangular facets $\kappa \in \mathcal{K}$ which need to be observed through the agent's camera.

is performed iteratively until the total surface area of the object of interest is covered. The notation $x_{t'|t}$ is used here to denote the future predicted agent state at time-step t' , which is computed at time-step t .

In this section we will also show how we have incorporated a ray-tracing based procedure into the proposed 3D coverage controller, in order to allow the determination of the visible parts of the scene through the agent's camera, and generate look-ahead coverage trajectories during the optimization.

5.1 Object of Interest

It is assumed that a 3D point-cloud representation of the object of interest is readily available prior to the coverage planning mission. Such representation can be acquired through a 3D scene reconstruction step [66], [67] where multiple calibrated images are collected from the object of interest, from which a 3D point-cloud representation $\mathcal{P} = \{p_1, \dots, p_{|\mathcal{P}|}\} \in \partial\mathcal{W}$, $p_i \in \mathbb{R}^3$ of the object's boundary can be extracted (i.e., points $p \in \mathcal{P}$ belong to the surface area on the object's boundary). The proposed approach takes as input the generated 3D point-cloud \mathcal{P} , which is then triangulated by subdividing the object's surface into a finite set \mathcal{K} of triangular facets $\kappa \in \mathbb{R}^{3 \times 3}$. In essence, \mathcal{P} is partitioned into a finite number of non-overlapping triangular facets $\kappa \in \mathcal{K}$ as shown in Fig. 3. Therefore, the constraint shown in Eq. (12d) i.e., the coverage constraint becomes:

$$\exists \mathcal{V}_\tau \in \mathcal{V}_{1:\tilde{T}} : \kappa \in \mathcal{V}_\tau, \quad \forall \kappa \in \mathcal{K}, \quad (13)$$

where with slight abuse of notation the expression $\kappa \in \mathcal{V}_\tau$ is used here to denote that the facet κ must reside within the convex-hull of the agent's camera FOV at time-step τ , and therefore we are interested in finding the agent's trajectory (i.e., kinematic and camera states) which results in the coverage of all triangle facets $\kappa \in \mathcal{K}$ which compose the object's surface area. Although, in this work we used Delaunay triangulation [68], [69] to create the triangle mesh \mathcal{K} , depending on the application scenario alternative triangulation methods can be used as well [70], [71].

5.2 Determining Visibility

Although, the constraint in Eq. (13) can be used to construct the agent's trajectory which covers all facets $\kappa \in \mathcal{K}$ (given a sufficiently large planning horizon \tilde{T}), it does not take

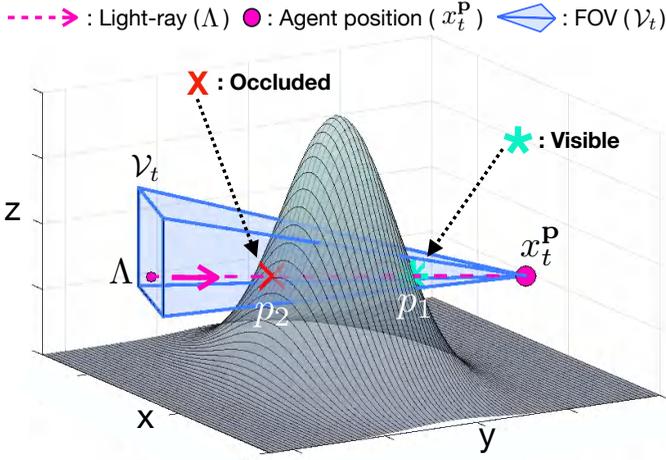


Fig. 4. The figure shows that although the two points p_1 and p_2 (marked with * and \times respectively) reside inside the volume \mathcal{V}_t covered by the agent's camera FOV, only point p_1 is visible. Specifically, in the camera configuration shown above there is no light-ray $\lambda \in \mathcal{L}(z_t, \theta_t, \phi_t, x_t^P)$ which traces back to point p_2 as all light-rays are blocked by the object of interest as shown.

into account the visibility problem i.e., determining whether some facet κ which resides inside the agent's camera FOV (i.e., $\kappa \in \mathcal{V}_t$) is actually visible. An illustrative example of this problem is illustrated in Fig. 4, which demonstrates the propagation of light-rays Λ through the agent's camera. As shown, in this example the two points on the object's boundary p_1 and p_2 , both reside inside the agent's camera FOV \mathcal{V}_t at time-step t , however only point p_1 is visible given the agent location x_t^P and camera state \mathcal{V}_t . In fact, as illustrated in the figure point p_2 resides on the occluded side of the object's surface, and as a result there is no light-ray that can be traced back to point p_2 as all light-rays are blocked by the object's body as shown.

Intuitively, a specific part of the object's surface, represented by the triangular facet $\kappa \in \mathcal{K}$, is visible through the agent's camera (given a particular camera pose $\mathcal{C}_t(z_t, \theta_t, \phi_t, x_t^P)$) when: a) κ resides within the camera's FOV as discussed in the previous section (i.e., $\kappa \in \mathcal{V}_t(z_t, \theta_t, \phi_t, x_t^P)$), and b) there exists a light-ray $\Lambda \in \mathcal{L}(z_t, \theta_t, \phi_t, x_t^P)$ which is not blocked and traces back to facet κ . Conversely, when no light-ray can be traced back to κ , indicates that the specific part of the object is not visible. This event might be attributed to an obstruction which blocks the propagation of light-rays and thus rendering that specific part of the object's surface unobservable. In essence, a facet is visible only when it resides inside the agent's camera FOV, and there is no occlusion intercepting the relevant light rays.

More precisely, the notion of visibility can now be described as follows: The part of the object's surface $\kappa \in \mathcal{K}$ is visible through the agent's camera with state $\mathcal{V}_t(z_t, \theta_t, \phi_t, x_t^P)$ at time t when:

$$\exists \Lambda \in \mathcal{L}(z_t, \theta_t, \phi_t, x_t^P) : \Lambda \oplus \mathcal{K} = \kappa, \quad (14)$$

where the ray-tracing operator \oplus , returns the facet $\kappa \in \mathcal{K}$ which intersects last with the light-ray Λ ; otherwise it returns \emptyset if no facet $\kappa \in \mathcal{K}$ interests with Λ . In essence $\Lambda \oplus \mathcal{K}$ computes the intersections of Λ with all the facets in the set \mathcal{K} , and returns (if exists) the one facet κ which Λ intersects just before entering the camera's optical center. The result of this operation provides the visible facet κ (i.e., the visible

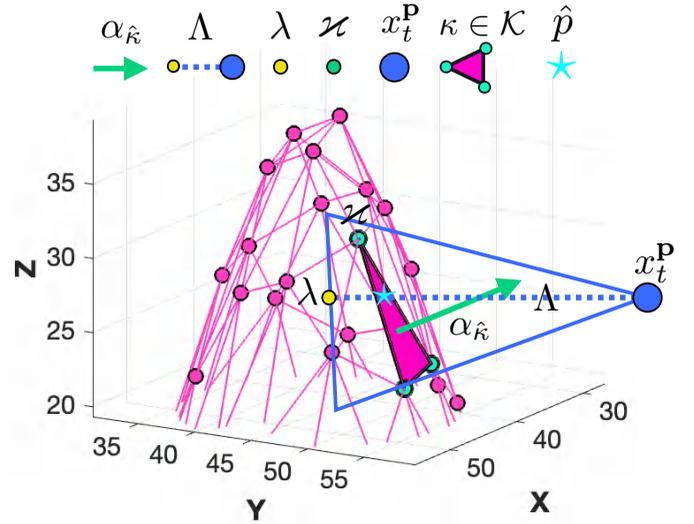


Fig. 5. The figure illustrates the ray-tracing procedure discussed in Sec. 5.2 which is used in order to determine the visibility of some facet $\kappa \in \mathcal{K}$.

part of the object's surface), since the intersecting light-ray Λ traces back to κ , allowing the surface area captured by κ to be imaged through the camera.

Let us denote the equation of the plane which contains the triangular facet $\kappa \in \mathcal{K}$ as:

$$\alpha_{\hat{\kappa}}^\top \cdot x = \beta_{\hat{\kappa}}, \quad (15)$$

where $\hat{\kappa} \in \{1, \dots, |\mathcal{K}|\}$ is the index pointing to facet $\kappa \in \mathcal{K}$, $\alpha_{\hat{\kappa}} \in \mathbb{R}^3$ is the outward normal vector to the plane containing κ , $x \in \mathbb{R}^3$, $\beta_{\hat{\kappa}} = \alpha_{\hat{\kappa}}^\top \cdot \varkappa$, $\varkappa \in \mathbb{R}^3$ is a vertex of κ , and the notation $a^\top \cdot b$ denotes the dot product of the 3D column vectors a and b .

That said, the operation $\Lambda \oplus \kappa$ first finds the intersection point (if exists) between the light-ray Λ which is given by Eq. (8) (i.e., $\Lambda = \lambda + d(x_t^P - \lambda)$), and the plane which contains the triangular facet κ given by Eq. (15) as follows:

$$\alpha_{\hat{\kappa}}^\top \cdot [\lambda + d(x_t^P - \lambda)] = \alpha_{\hat{\kappa}}^\top \cdot \varkappa \implies \quad (16a)$$

$$d = \frac{\alpha_{\hat{\kappa}}^\top \cdot (\varkappa - \lambda)}{\alpha_{\hat{\kappa}}^\top \cdot (x_t^P - \lambda)}, \quad (16b)$$

where Eq. (16a) is the result of the substitution of Λ for x in Eq. (15), and then in Eq. (16b) we solve for d . As a reminder in Eq. (16b) $\alpha_{\hat{\kappa}}$ is the normal vector to the plane which contains facet κ , \varkappa is a vertex of κ , λ is a point on the light-ray Λ (i.e., the origin of the ray), and finally x_t^P is the ray's end-point given by the agent's location. In essence the vector $x_t^P - \lambda$ denotes the direction of the light-ray. An illustrative example is shown in Fig. 5. Consequently, if the denominator of Eq. (16b) is equal to zero, the light-ray Λ and the facet κ are parallel which results in either no visibility (i.e., $\alpha_{\hat{\kappa}}^\top \cdot (\varkappa - \lambda) \neq 0$) or distorted view (i.e., $\alpha_{\hat{\kappa}}^\top \cdot (\varkappa - \lambda) = 0$). Therefore, we are interested in the scenario where $\alpha_{\hat{\kappa}}^\top \cdot (x_t^P - \lambda) \neq 0$ and thus there exists a single point intersection between the light-ray Λ and the plane which contains the facet κ . In such scenario, the light-ray Λ can be traced back to κ which can render it visible through the agent's camera. However, in order for this to happen two conditions must simultaneously be satisfied: a) the value of d must lie within the interval $d \in [0, 1]$, otherwise the point of intersection is outside the range of the camera's FOV, and b) the point of intersection $\hat{p} = \lambda + \hat{d}(x_t^P - \lambda)$ (where $\hat{d} \in [0, 1]$) is the

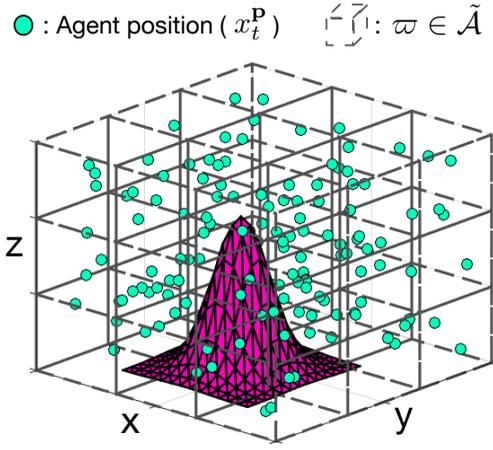


Fig. 6. A set of visibility constraints $b_{\hat{\kappa}, \hat{\kappa}}$ is learned through ray-tracing for each cell $\varpi \in \tilde{\mathcal{A}}$ and facet $\kappa \in \mathcal{K}$. Within each cell ϖ , we randomly sample the agent's position x^P , along with various combinations of the control inputs $z \in Z$, $\theta \in \Theta$, $\phi \in \Phi$ to identify the visible parts of the object via the ray-tracing procedure discussed in Sec. 5.2.

solution of Eq. (16b)) must reside within the convex hull of the triangular facet κ i.e., $\hat{p} \in \Delta(\kappa)$.

To summarize, the visibility determination process via ray-tracing discussed above must be evaluated at each time-step for all pairwise combinations of light-rays $\Lambda \in \mathcal{L}(z_t, \theta_t, \phi_t, x_t^P)$, and triangular facets $\kappa \in \mathcal{K}$, inside the planning horizon. It is easy to see that this is a computationally expensive task, which needs to be performed at each time-step. Additionally, we should mention here that, the implementation of this ray-tracing procedure introduces non-convex and non-linear constraints which are very challenging to be handled efficiently during optimization. To see this, observe that the solution to Eq. (16b) depends on the unknown control inputs (i.e., z_t, θ_t, ϕ_t) and state of the agent (i.e., x_t^P), thus rendering Eq. (16b) non-convex and non-linear.

Instead of directly implementing the ray-tracing process discussed above, in this work we employ an alternative procedure which allows the integration of ray-tracing into a mixed integer program (MIP), which can be solved using standard branch-and-bound techniques [72]. To accomplish this, the ray-tracing functionality is evaluated and learned off-line on a discretized representation of the environment. Subsequently, visibility determination is approximated with a set of binary decision variables, which can easily be incorporated into a mixed integer program, thus allowing the ray-tracing functionality to be utilized during optimization.

To achieve this, the environment \mathcal{A} is first discretized in space, to form a 3D grid $\tilde{\mathcal{A}}$ composed of a finite number of non-overlapping cells $\tilde{\mathcal{A}} = \{\varpi_1, \dots, \varpi_{|\tilde{\mathcal{A}}|}\}$, such that $\bigcup_{i=1}^{|\tilde{\mathcal{A}}|} \varpi_i = \tilde{\mathcal{A}}$. Then, within each cell $\varpi \in \tilde{\mathcal{A}}$, we randomly sample the agent's position x^P , as shown in Fig. 6, and we run the ray-tracing procedure discussed above for various joint combinations of the control inputs $z \in Z$, $\theta \in \Theta$, $\phi \in \Phi$, identifying in this way the visible parts of the object, and learning a set of state-dependent binary visibility constraints which can be embedded in a mixed integer program, and used during optimization.

Let $\tilde{\mathcal{L}}_{\hat{\kappa}}$ denote the set of light-rays that have been obtained from the application of N random combinations of the control inputs $z \in Z, \theta \in \Theta, \phi \in \Phi$, and agent states

x^P which have been uniformly sampled within the cell ϖ i.e.,

$$\tilde{\mathcal{L}}_{\hat{\kappa}} = \bigcup_{i=1}^N \{\Lambda : \Lambda \in \mathcal{L}_{\hat{\kappa}}^i(z_i, \theta_i, \phi_i, x_i^P)\}, \quad (17)$$

where $\hat{\kappa} \in \{1, \dots, |\tilde{\mathcal{A}}|\}$ is the index pointing to cell $\varpi \in \tilde{\mathcal{A}}$, $\mathcal{L}_{\hat{\kappa}}^i(z_i, \theta_i, \phi_i, x_i^P)$ is the set of light-rays given by the camera pose obtained with the set of control inputs (z_i, θ_i, ϕ_i) , for the agent location x_i^P sampled within cell ϖ . We then learn the ray-tracing constraints as follows:

$$b_{\hat{\kappa}, \hat{\kappa}} = 1 \iff \exists \Lambda \in \tilde{\mathcal{L}}_{\hat{\kappa}} : \Lambda \oplus \mathcal{K} = \kappa. \quad (18)$$

Thus in order to determine the visibility of facet κ , we utilize the binary variable $b_{\hat{\kappa}, \hat{\kappa}}$ which is activated when there exists a light-ray Λ which traces back to facet κ when the agent resides within the cell ϖ . In the next section we show how we have utilized the learned ray-tracing constraints in Eq. (18) in the proposed 3D coverage controller using mixed integer mathematical programming.

Problem (P2): 3D Coverage Controller

$$\arg \min_{u_{t|t}, u_{t+1|t}, \dots, u_{t+T-1|t}} \mathcal{G}, \quad (19a)$$

subject to: $\tau \in \{0, \dots, T-1\}$

$$x_{t+\tau+1|t} = Ax_{t+\tau|t} + Bu_{t+\tau|t}^f, \quad \forall \tau \quad (19b)$$

$$x_{t|t} = x_{t|t-1}, \quad (19c)$$

$$\tilde{\mathcal{C}}_{\hat{m}} = R_\phi(u^\phi)R_\theta(u^\theta)\mathcal{C}_0(u^z), \quad \forall \hat{m} \quad (19d)$$

$$\mathcal{C}_{\hat{m}, t+\tau+1|t} = \tilde{\mathcal{C}}_{\hat{m}} + x_{t+\tau+1|t}^P, \quad \forall \hat{m}, \tau \quad (19e)$$

$$\sum_{\hat{m}=1, \dots, |\mathcal{M}|} s_{\hat{m}, t+\tau+1|t} = 1, \quad \forall \tau \quad (19f)$$

$$\mathcal{V}_{\hat{m}, t+\tau+1|t} = \Delta(\mathcal{C}_{\hat{m}, t+\tau+1|t}), \quad \forall \hat{m}, \tau \quad (19g)$$

$$b_{\hat{\kappa}, \hat{m}, t+\tau+1|t}^\mathcal{V} = 1 \iff \kappa \in \mathcal{V}_{\hat{m}, t+\tau+1|t}, \quad \forall \hat{\kappa}, \hat{m}, \tau \quad (19h)$$

$$b_{\hat{\kappa}, \hat{\kappa}, t+\tau+1|t}^\mathcal{A} = 1 \iff x_{t+\tau+1|t}^P \in \Delta(\varpi), \quad \forall \hat{\kappa}, \tau \quad (19i)$$

$$\bar{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t} = s_{\hat{m}, t+\tau+1|t} \wedge \left(b_{\hat{\kappa}, \hat{m}, t+\tau+1|t}^\mathcal{V} \wedge [b_{\hat{\kappa}, \hat{\kappa}, t+\tau+1|t}^\mathcal{A} \wedge b_{\hat{\kappa}, \hat{\kappa}}] \right), \quad \forall \hat{\kappa}, \hat{m}, \tau \quad (19j)$$

$$\hat{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t} \leq \bar{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t} + \mathcal{Q}(\kappa), \quad \forall \hat{\kappa}, \hat{m}, \tau \quad (19k)$$

$$\sum_{\tau} \sum_{\hat{m}} \hat{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t} \leq 1, \quad \forall \hat{\kappa} \quad (19l)$$

$$x_{t+\tau+1|t}^P \notin \Delta(\xi), \quad \forall \xi \in \Xi \quad (19m)$$

$$x_{t+\tau+1|t} \in \mathcal{X}, \quad u_{t+\tau|t} \in \mathcal{U},$$

$$s_{\hat{m}, t+\tau+1|t}, b_{\hat{\kappa}, \hat{m}, t+\tau+1|t}^\mathcal{V}, b_{\hat{\kappa}, \hat{\kappa}, t+\tau+1|t}^\mathcal{A} \in \{0, 1\},$$

$$\bar{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t}, \hat{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t}, \mathcal{Q}(\kappa) \in \{0, 1\},$$

$$\hat{m} \in [1, \dots, |\mathcal{M}|], \quad \hat{\kappa} \in [1, \dots, |\mathcal{K}|], \quad \hat{\kappa} \in [1, \dots, |\tilde{\mathcal{A}}|].$$

5.3 Autonomous 3D Coverage Control

The proposed 3D coverage controller is shown in Problem (P2), where we have transformed the coverage problem discussed in Sec. 4 into a rolling finite horizon optimal control problem, formulated as a mixed integer program (MIP). More specifically, the controller shown in Problem (P2) finds the agent's joint control inputs (i.e., kinematic and camera control inputs) $u_{t+\tau|t}, \tau \in \{0, \dots, T-1\}$ over the rolling horizon T , such that the coverage objective \mathcal{G} is

optimized, subject to coverage-and-visibility constraints. As we will discuss next in more detail, Problem (P2) also makes sure that the duplication of work is minimized (i.e., since the coverage mission cannot be completed within the short planning horizon T , the state of the mission is being tracked and saved into the agent's memory, therefore allowing the mission to progress by avoiding the duplication of work), and that the agent avoids collisions with obstacles and the object of interest.

5.3.1 Coverage Objective

Let us assume that each part of the object of interest i.e., every facet $\kappa \in \mathcal{K}$, is associated with a binary variable $\hat{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t} \in \{0, 1\}$ which is defined as below:

$$\hat{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t} = \begin{cases} 1, & \text{if } \exists m, \tau : \kappa \in \mathcal{V}_{m, \tau} \wedge \mathcal{J}_{\mathcal{V}_{m, \tau}}(\kappa) \neq \emptyset \\ 0, & \text{Otherwise} \end{cases} \quad (20)$$

where $\tau \in \{0, \dots, T-1\}$, $m \in \mathcal{M}$, and \mathcal{M} represents the set of all possible joint combinations of admissible camera control inputs (zoom-level and rotation angles) given by:

$$\mathcal{M} = Z \times \Theta \times \Phi, \quad (21)$$

where \times denotes the Cartesian product on sets, $|\mathcal{M}|$ is the total number of possible camera FOV configurations, and $\hat{m} \in \{1, \dots, |\mathcal{M}|\}$ is the index of the FOV configuration m . Thus $\mathcal{V}_{\hat{m}, \tau}$ denotes the convex hull of the FOV configuration $m \in \mathcal{M}$ at time-step $t+\tau+1|t$. The function $\mathcal{J}_{\mathcal{V}_{\hat{m}, \tau}}(\kappa)$ determines the visibility of facet κ with respect to the FOV state $\mathcal{V}_{\hat{m}, \tau}$, by implementing the ray-tracing procedure discussed in Sec. 5.2. When facet κ is visible at time-step τ through the $m \in \mathcal{M}$ FOV configuration $\mathcal{V}_{\hat{m}, \tau}$, the function $\mathcal{J}_{\mathcal{V}_{\hat{m}, \tau}}(\kappa)$ returns the light-ray that traces back to κ , otherwise it returns the empty set when κ is not visible (we will discuss how the functionality of $\mathcal{J}_{\mathcal{V}_{\hat{m}, \tau}}(\kappa)$ was integrated into the proposed approach later in this section).

Therefore, the binary variable $\hat{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t}$ determines whether facet κ is covered and is visible through the m FOV configuration at time-step $t+\tau+1|t$. We should point out here that $\mathcal{V}_{\hat{m}, \tau}$ depends on the applied control inputs $u_{t+\tau|t}$, and thus the activation of $\hat{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t}$ can be optimized for coverage by appropriately selecting the agent's control inputs inside the planning horizon. Based on that, we can now define the coverage objective function \mathcal{G} as follows:

$$\mathcal{G} = \omega \mathcal{D}(x_{t+1|t}^{\text{P}}, \kappa^*) - \sum_{\hat{\kappa}=1}^{|\mathcal{K}|} \sum_{\hat{m}=1}^{|\mathcal{M}|} \sum_{\tau=0}^{T-1} \gamma(\tau) \hat{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t}, \quad (22)$$

where $\mathcal{D}(x_{t+1|t}^{\text{P}}, \kappa^*) = \|x_{t+1|t}^{\text{P}} - (\delta \alpha_{\hat{\kappa}} + \kappa^*)\|_2^2$ is the squared Euclidian distance between the predicted agent position $x_{t+1|t}^{\text{P}}$ and the point $\delta \alpha_{\hat{\kappa}} + \kappa^*$, where $\kappa^* \in \mathbb{R}^3$ is the centroid of the nearest unobserved facet κ , $\alpha_{\hat{\kappa}}$ is the unit normal vector to the plane containing facet κ , and $\delta \in \mathbb{R}$ is a user defined positive scalar. The parameter $\omega \in \mathbb{R}_+$ is a positive tuning weight, and as already discussed the binary variable $\hat{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t}$ is used to track the coverage events. In essence we are trying to activate $\hat{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t}$ inside the planning horizon for as many facets κ as possible. This essentially maximizes the total surface area covered inside the planning horizon. The time-dependend function $\gamma(\tau)$ can be used here to penalize the parts of the object that are covered later in the horizon, thus encouraging the agent to cover as many facets

as quickly as possible. Observe that, the minimization of \mathcal{G} is taking place inside the finite and short planning horizon T , which results in partial coverage of the object of interest. For this reason, \mathcal{G} is optimized for subsequent time-steps t for the duration of the mission, and until the object of interest is fully covered. The function $\mathcal{D}(\cdot)$ is particularly useful for the mission progress (i.e., it drives the agent towards the facets that remain to be covered) in the event where although not all facets have been covered, the remaining facets cannot be reached for coverage within the planning horizon (i.e., the second term of Eq. (22) is 0). Next, we show in detail how a mixed integer program (MIP) can be formulated to optimize the 3D coverage objective function of Eq. (22), in a rolling horizon fashion.

5.3.2 Coverage Constraints

Model predictive kinematic control is achieved with the constraints shown in Eq. (19c) and Eq. (19d), where the kinematic control inputs $u_{t+\tau|t}^f$ are optimized inside the planning horizon according to the agent's kinematic model as discussed in Sec. 3.1.

The $m \in \mathcal{M}$ camera configuration (in terms of FOV vertices) is then given in constraint Eq. (19d) by the variable $\tilde{C}_{\hat{m}}$, by applying Eq. (4) with the \hat{m}_{th} combination of camera control inputs (i.e., zoom-level and rotation angles respectively). Observe that the operation in Eq. (19d) is precomputed for all $|\mathcal{M}|$ possible camera FOV states, and thus $\tilde{C}_{\hat{m}}$ simply selects the camera configuration m . Subsequently, the camera FOV pose $\mathcal{C}_{\hat{m}, t+\tau+1|t}$ is obtained by translating $\tilde{C}_{\hat{m}}$ to the predicted agent location $x_{t+\tau+1|t}^{\text{P}}$ via the constraint shown in Eq. (19e), where we have applied Eq. (7) as discussed in more detail in Sec. 3.2.

The binary variable $s_{\hat{m}, t+\tau+1|t}$ is used in Eq. (19f) to indicate the active camera state at time-step $t+\tau+1|t$. More specifically, at each time-step only one camera state should be active. Therefore, in order to avoid the situation where more than one camera states are activated in the same time-step, we use the constraint in Eq. (19f). The decision variable, s in Eq. (19f) can be thought as a 2D matrix with T columns and $|\mathcal{M}|$ rows, with the restriction that the sum of each column must be equal to one.

Next, the constraint in Eq. (19g) computes the convex hull $\mathcal{V}_{\hat{m}, t+\tau+1|t}$ of the FOV vertices defined by the \hat{m}_{th} camera state $\mathcal{C}_{\hat{m}, t+\tau+1|t}$ at time-step $t+\tau+1|t$. The volume enclosed within the camera FOV $\mathcal{C}_{\hat{m}, t+\tau+1|t}$ is given by the system of linear inequalities:

$$\alpha_{\mathcal{C}_{\hat{m}, t+\tau+1|t}, i}^{\top} \cdot x \leq \beta_{\mathcal{C}_{\hat{m}, t+\tau+1|t}, i}, \quad \forall i = [1, \dots, 5], \quad (23)$$

where the equation $\alpha_{\mathcal{C}_{\hat{m}, t+\tau+1|t}, i}^{\top} \cdot x = \beta_{\mathcal{C}_{\hat{m}, t+\tau+1|t}, i}$ describes the equation of the plane which contains face $i \in \{1, \dots, 5\}$ of the camera's FOV. The variable $\alpha_{\mathcal{C}_{\hat{m}, t+\tau+1|t}, i}^{\top}$ is the outward normal vector to the plane which contains face $i \in \{1, \dots, 5\}$ of the camera's FOV with state $\mathcal{C}_{\hat{m}, t+\tau+1|t}$, $x \in \mathbb{R}^3$, and $\beta_{\mathcal{C}_{\hat{m}, t+\tau+1|t}, i}$ is a constant. Therefore, any point $x \in \mathbb{R}^3$ which satisfies all 5 inequalities shown in Eq. (23) resides within the camera's FOV i.e., belongs to the convex hull defined by the camera vertices $\mathcal{C}_{\hat{m}, t+\tau+1|t}$.

The constraint in Eq. (19h) uses the binary variable $b_{\hat{\kappa}, \hat{m}, t+\tau+1|t}^{\text{V}}$ to decide whether at time-step $t+\tau+1|t$, the triangular facet $\kappa \in \mathcal{K}$ resides within the \hat{m}_{th} configuration of

the camera FOV defined by $\mathcal{V}_{\hat{m},t+\tau+1|t}$. This is implemented as follows:

$$\alpha_{i,\hat{m},\tau}^\top \cdot \kappa^c + \hat{b}_{i,\hat{k},\hat{m},\tau}^\mathcal{V} (M - \beta_{i,\hat{m},\tau}) \leq M, \quad \forall i, \hat{k}, \hat{m}, \tau, \quad (24a)$$

$$5b_{\hat{k},\hat{m},\tau}^\mathcal{V} - \sum_{i=1}^5 \hat{b}_{i,\hat{k},\hat{m},\tau}^\mathcal{V} \leq 0, \quad \forall \hat{k}, \hat{m}, \tau, \quad (24b)$$

where we have used τ as the short notation for $t+\tau+1|t$, and thus $\alpha_{\hat{m},t+\tau+1|t,i}^\top$ is abbreviated as $\alpha_{i,\hat{m},\tau}^\top$ (i.e., the outward normal vector to the plane which contains face i of the $m \in \mathcal{M}$ camera FOV at time-step τ), and $\beta_{\hat{m},t+\tau+1|t,i}$ has been shortened to $\beta_{i,\hat{m},\tau}$ as shown in Eq. (24a). Moreover, in Eq. (24a) the triangular facet κ has been approximated by its center of mass κ^c , and M is an arbitrary large positive constant, which is used to ensure that Eq. (24a) is valid for any value of the binary decision variable $\hat{b}_{i,\hat{k},\hat{m},\tau}^\mathcal{V}$. When κ^c resides within $\mathcal{V}_{\hat{m},t+\tau+1|t}$, the binary variable $\hat{b}_{i,\hat{k},\hat{m},\tau}^\mathcal{V}$ is activated for all faces i.e., $\hat{b}_{i,\hat{k},\hat{m},\tau}^\mathcal{V} = 1, \forall i = \{1, \dots, 5\}$ in order to satisfy the constraint imposed by Eq. (24a). In such scenario, the binary variable $b_{\hat{k},\hat{m},t+\tau+1|t}^\mathcal{V}$ (abbreviated as $b_{\hat{k},\hat{m},\tau}^\mathcal{V}$) in Eq. (24b) is activated to satisfy the constraint, which indicates that the \hat{m}_{th} camera FOV contains facet κ at time-step $t + \tau + 1|t$. When $b_{\hat{k},\hat{m},\tau}^\mathcal{V} = 0$, then Eq. (24b) is also valid as shown.

Similarly, the constraint in Eq. (19i) uses the binary variable $b_{\hat{\omega},t+\tau+1|t}^{\hat{A}}$ to determine if the agent's location $x_{t+\tau+1|t}^{\mathbf{p}}$ at time-step τ resides within the cell $\hat{\omega} \in \hat{\mathcal{A}}$ as shown below:

$$\alpha_{j,\hat{\omega}}^\top \cdot x_\tau^{\mathbf{p}} + \hat{b}_{i,\hat{\omega},\tau}^{\hat{A}} (M - \beta_{j,\hat{\omega}}) \leq M, \quad \forall j, \hat{\omega}, \tau, \quad (25a)$$

$$6b_{\hat{\omega},\tau}^{\hat{A}} - \sum_{j=1}^6 \hat{b}_{i,\hat{\omega},\tau}^{\hat{A}} \leq 0, \quad \forall \hat{\omega}, \tau, \quad (25b)$$

where $j \in \{1, \dots, 6\}$ indicates the faces of the rectangular-shaped cell $\hat{\omega} \in \hat{\mathcal{A}}$, the time-step $t + \tau + 1|t$ is abbreviated as τ , the parameters $(\alpha_{j,\hat{\omega}}^\top, \beta_{j,\hat{\omega}})$ define the coefficients of the equation of the plane which contains the j_{th} face of the $\hat{\omega}_{\text{th}}$ cell of the grid i.e., $\alpha_{j,\hat{\omega}}^\top \cdot x = \beta_{j,\hat{\omega}}$, and again M is a large positive constant. Therefore, when the agent resides within some cell $\hat{\omega}$, the binary variable $\hat{b}_{i,\hat{\omega},\tau}^{\hat{A}}$ is activated for all $j = \{1, \dots, 6\}$ faces, and subsequently the binary variable $b_{\hat{\omega},\tau}^{\hat{A}}$ is activated, thus indicating the presence of the agent within cell $\hat{\omega}$ at time-step τ i.e., $b_{\hat{\omega},\tau}^{\hat{A}} = 1$.

Subsequently, the coverage-and-visibility constraint shown in Eq. (20), is implemented in Eq. (19j) as the logical conjunction shown below:

$$\bar{b}_{\hat{k},\hat{m},t+\tau+1|t} = s_{\hat{m},t+\tau+1|t} \wedge \left(b_{\hat{k},\hat{m},t+\tau+1|t}^\mathcal{V} \wedge [b_{\hat{\omega},t+\tau+1|t}^{\hat{A}} \wedge b_{\hat{\omega},\hat{k}}^{\hat{A}}] \right), \quad \forall \hat{k}, \hat{m}, \hat{\omega}, \tau,$$

where the binary variable $\bar{b}_{\hat{k},\hat{m},t+\tau+1|t}$ is activated when the centroid κ^c of the facet κ , is covered and is visible through the m FOV configuration of the agent's camera at time-step $t + \tau + 1|t$. In order for this to happen, κ^c must first reside within the m FOV configuration at time-step $t + \tau + 1|t$ as indicated by the binary variable $b_{\hat{k},\hat{m},t+\tau+1|t}^\mathcal{V}$, and at the same time the agent must reside within some cell $\hat{\omega}$ as indicated by the binary variable $b_{\hat{\omega},t+\tau+1|t}^{\hat{A}}$ from which facet κ is visible (i.e., there exists a light-ray captured by the agent's camera with the m configuration which traces back to κ) as indicated by the binary variable $b_{\hat{\omega},\hat{k}}^{\hat{A}}$ which was

learned off-line as discussed in Sec. 5.2. Finally, the binary variable $s_{\hat{m},t+\tau+1|t}$ makes sure that at time-step $t + \tau + 1|t$ only the m out of the $|\mathcal{M}|$ FOV configurations is active.

The constraint shown in Eq. (19k) is used to ensure that the mission progress (i.e., the facets which have been covered) is being saved and tracked, thus allowing the agent to avoid the duplication of work. This is essential since in most scenarios due to the short planning horizon only partial coverage trajectories would be obtained. This necessitates the implementation of some form of memory in where the mission progress would be saved. The agent's memory is represented in this work with the function $\mathcal{Q} : \mathcal{K} \rightarrow \{0, 1\}$ which is defined as follows:

$$\mathcal{Q}(\kappa) = \begin{cases} 1, & \text{if } \exists t' \leq t : \kappa \text{ is covered.} \\ 0, & \text{o.w} \end{cases} \quad (26)$$

which is activated to indicate that facet κ has been covered at some time-step $t' \leq t$, where t denotes the current time-step. The constraint in Eq. (19k) discourages the agent from covering facets which have been already covered, thus avoiding the duplication of work. This is because the maximization of the binary variable $\hat{b}_{\hat{k},\hat{m},t+\tau+1|t}$ can occur either through $\bar{b}_{\hat{k},\hat{m},t+\tau+1|t}$ or through $\mathcal{Q}(\kappa)$. Thus, if κ has been observed in the past (i.e., $\mathcal{Q}(\kappa) = 1$), then there is no incentive to generate a plan which covers κ inside the current planning horizon (i.e., there is no need to activate $\bar{b}_{\hat{k},\hat{m},t+\tau+1|t}$), since $\hat{b}_{\hat{k},\hat{m},t+\tau+1|t}$ has been already maximized through $\mathcal{Q}(\kappa)$. The next constraint shown in Eq. (19l) makes sure that the agent does not plan a coverage trajectory which results in the coverage of the same facet more than once during the same planning horizon.

Finally, the constraint shown in Eq. (19m) makes sure that the agent avoids collisions with the obstacles in the environment, and with the object of interest. To implement this functionality we require that the agent's position $x_{t+\tau+1|t}^{\mathbf{p}}$ lies at all times outside the convex hull of all obstacles $\xi \in \Xi$ found in the environment. The procedure followed here first uses triangulation or tessellation techniques to decompose the surface area of the convex obstacle ξ into a finite number of faces $f_{\hat{\xi},i}, i = \{1, \dots, n^{\hat{\xi}}\}$, where each face $f_{\hat{\xi},i}$ exists on the plane with equation given by:

$$\alpha_{\hat{\xi},i}^\top \cdot x = \beta_{\hat{\xi},i}, \quad \forall i \in \{1, \dots, n^{\hat{\xi}}\}. \quad (27)$$

where as before $\alpha_{\hat{\xi},i}$ is the normal vector to the plane which contains the i_{th} face, $\beta_{\hat{\xi},i}$ is a constant, $x \in \mathbb{R}^3$, and $\hat{\xi} \in \{1, \dots, |\Xi|\}$ is the index of obstacle $\xi \in \Xi$. Subsequently, the i_{th} plane divides the environment into two half-spaces, and the intersection of all $n^{\hat{\xi}}$ negative half-spaces forms the convex-hull of the obstacle. Therefore we use the set of constraints shown in Eq. (28a) - Eq. (28b), and the binary variable $b_{t+\tau+1|t,i}^{\hat{\xi}}$ to make sure that a collisions with the obstacle ξ is avoided at all time-steps $t + \tau + 1|t$ inside the planning horizon.

$$\alpha_{\hat{\xi},i}^\top \cdot x_{t+\tau+1|t}^{\mathbf{p}} + Mb_{t+\tau+1|t,i}^{\hat{\xi}} > \beta_{\hat{\xi},i}, \quad \forall \hat{\xi}, \tau, i \quad (28a)$$

$$\sum_{i=1}^{n^{\hat{\xi}}} b_{t+\tau+1|t,i}^{\hat{\xi}} < n^{\hat{\xi}}, \quad \forall \hat{\xi}, \tau, \quad (28b)$$

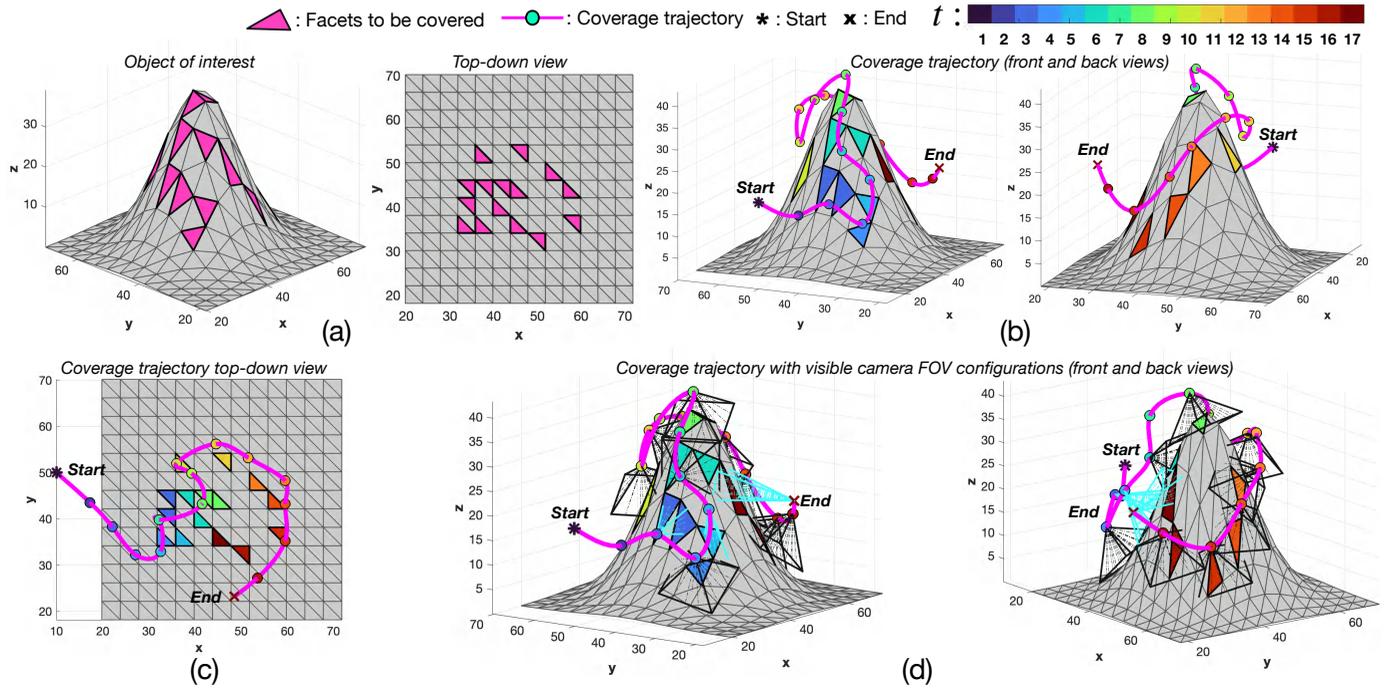


Fig. 7. The figure illustrates a simulated 3D coverage planning scenario executed with the proposed approach.

Specifically, observe from the constraint in Eq. (28a) that when the agent resides within the convex hull of obstacle ξ at time-step $t + \tau + 1|t$, then $\alpha_{\xi,i}^T \cdot x_{t+\tau+1|t}^P \leq \beta_{\xi,i}, \forall i$, and the binary variable $b_{t+\tau+1|t,i}^{\xi} = 1, \forall i$ is activated in order to satisfy the constraint (where M is a large positive constant). Thus, the agent resides outside the convex hull of the obstacle at time-step $t + \tau + 1|t$, and the collision is avoided when $\exists i \in \{1, \dots, n^{\xi}\} : b_{t+\tau+1|t,i}^{\xi} = 0$. For this reason, the constraint in Eq. (28b) makes sure that the number of times $b_{t+\tau+1|t,i}^{\xi}$ is activated at each time-step inside the planning horizon is always less than n^{ξ} as shown in Eq. (28b).

To summarize, observe that the minimization of the coverage objective \mathcal{G} shown in Eq. (22) drives the agent to select the optimal control inputs $u_{t+\tau|t}, \tau \in \{0, \dots, T-1\}$ inside the planning horizon which enable the activation of the binary variables $\hat{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t}$, which in turn indicate the coverage of certain parts of the object of interest. Full coverage is then reached by solving the problem iteratively over multiple time-steps t within the duration of the mission. This is achieved by saving the mission progress inside the agent's memory, which also allows the agent to minimize the duplication of work, and complete its mission in subsequent optimization steps. Consequently, the coverage mission is terminated when all facets $\kappa \in \mathcal{K}$ are covered as indicated by the terminal condition $\sum_{\kappa \in \mathcal{K}} Q(\kappa) = |\mathcal{K}|$.

6 EVALUATION

The evaluation of the proposed 3D coverage approach is divided into two parts. In the first part we conduct a thorough analysis of the proposed approach using synthetic tests, showcasing its performance in a variety of simulated scenarios. In particular, we show the behavior of the proposed controller in a simulated 3D coverage scenario, demonstrating the role of ray-tracing during the coverage mission, and

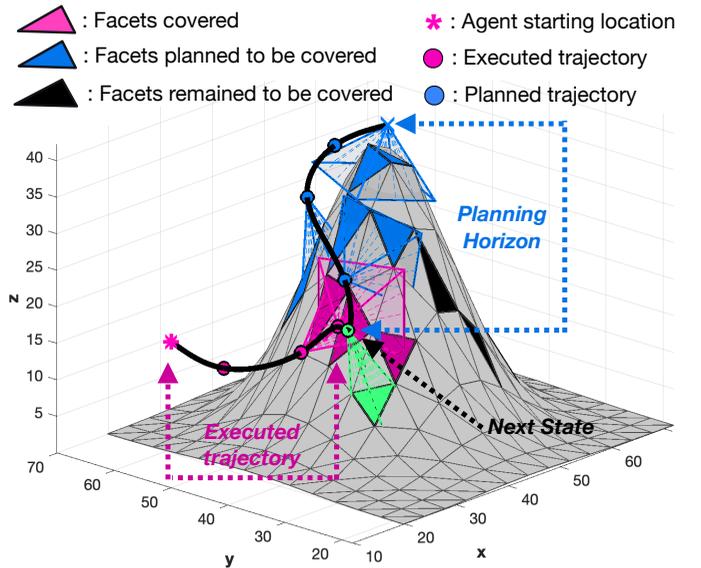


Fig. 8. The figure shows the planned coverage trajectory over a horizon length of 5 time-steps, for time instances $t + \tau + 1|t, t = 4, \tau \in [0, \dots, 4]$.

then we investigate the agent's coverage trajectory using various optimization objectives. In the second part of the evaluation, we illustrate the performance of the proposed approach in a real-world coverage mission.

6.1 Synthetic Experiments

For the evaluation of the proposed approach we have used the following simulation setup based on the specifications of the DJI Mavic Enterprise drone which was utilized in our field tests. The agent's kinematic model parameters $\Delta t, \alpha$, and m have been set to 1s, 0.2, and 1.1kg respectively. The agent's velocity x_t^V and motion control input f_t are bounded in each dimension inside the intervals $[-15, 15]$ m/s and $[-10, 10]$ N respectively. The parameter set (l, w, h) of the

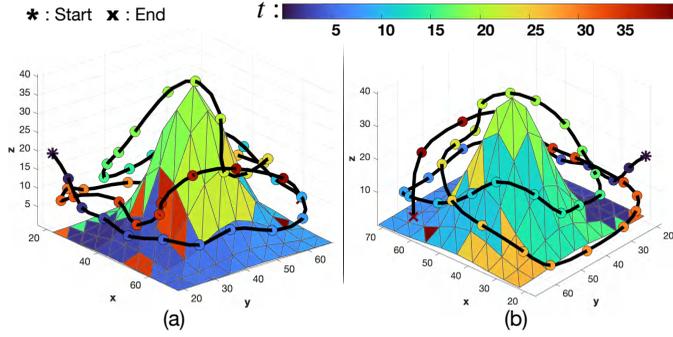


Fig. 9. The figure illustrates the agent's trajectory which was used to cover the total surface area of the object of interest. (a) Front view of the object. (b) Back view of the object.

camera FOV has been set as $(9.5, 9.5, 8)$ m, with $|Z| = 2$ zoom-levels i.e., $Z = \{1, 2\}$, and $n = 50$ light-rays. The rotation angles θ and ϕ take their values from the sets $\Theta = \{30, 90, 150\}$ deg, and $\Phi = \{30, 105, 180, 255, 330\}$ deg respectively. This enables the camera FOV to take one out of $|M| = 30$ possible configurations (i.e., 15 configurations per zoom-level). The volume of the 3D environment \mathcal{A} is equal to 100^3 m, which was uniformly discretized into $|\tilde{\mathcal{A}}| = 1000$ non-overlapping cells in order to determine the visibility constraints with the procedure described in Sec. 5.2, with $N = 100$. The object of interest to be covered is described by the Gaussian function $f(x, y) = A \exp\left(-\left(\frac{(x-x_o)^2}{2\sigma_x^2} + \frac{(y-y_o)^2}{2\sigma_y^2}\right)\right)$, with $(x_o, y_o) = (45, 45)$, $\sigma_x^2 = \sigma_y^2 = 80$, and $A = 40$, triangulated into $|\mathcal{K}| = 338$ triangular facets from a point-cloud of $|\mathcal{P}| = 200$ points. The planning horizon in the following experiments is set as $T = 5$, and the mission time is 100 time-steps. The term $\gamma(\tau)$ in the objective function of Eq. (22) is given by $\gamma(\tau) = \exp(T - \tau)$, $\tau \in \{0, \dots, T-1\}$, $\delta = 10$ m, and $\omega = 0.1$. Finally we should mention that we have used the Gurobi MIQP solver [73] to solve the optimization problem (P2), on a Mac Studio M1 Ultra.

We start the evaluation of the proposed approach with an illustrative example shown in Fig. 7. In order to simplify the analysis and for visual clarity, in this scenario we require that only a subset of facets $\tilde{\mathcal{K}} \subseteq \mathcal{K}$ need to be covered. For this reason, we have randomly sampled $|\tilde{\mathcal{K}}| = 15$ facets from the object's surface as shown in Fig. 7(a) with purple color. Therefore, given the facets $\tilde{\mathcal{K}}$, in this scenario we seek to find the agent's optimal control inputs (i.e., kinematic and camera inputs) which minimize the coverage objective in Eq. (19a), and respect the constraints shown in Eq. (19b) - Eq. (19m). Then, Problem (P2) is executed iteratively until all facets are covered. Figure 7(b) shows the resulting trajectory that the agent has executed in order to cover all facets. The agent starts at $(x, y, z) = (10, 50, 20)$ and finishes at $(x, y, z) = (48, 23, 22)$, with the start and final location shown as * and \times respectively. The facets to be covered and the agent states (shown as \circ) are color-coded according to the coverage elapsed time as shown in Fig. 7(b). Then, Fig. 7(c) shows the coverage trajectory in top-down view, and Fig. 7(d) shows the camera FOV configurations which have been used during the coverage mission. For visual clarity, we only show the FOV configurations which have resulted in a coverage event. When there is no coverage event the camera's FOV takes its initial state i.e., facing downwards. The zoom functionality is depicted with the

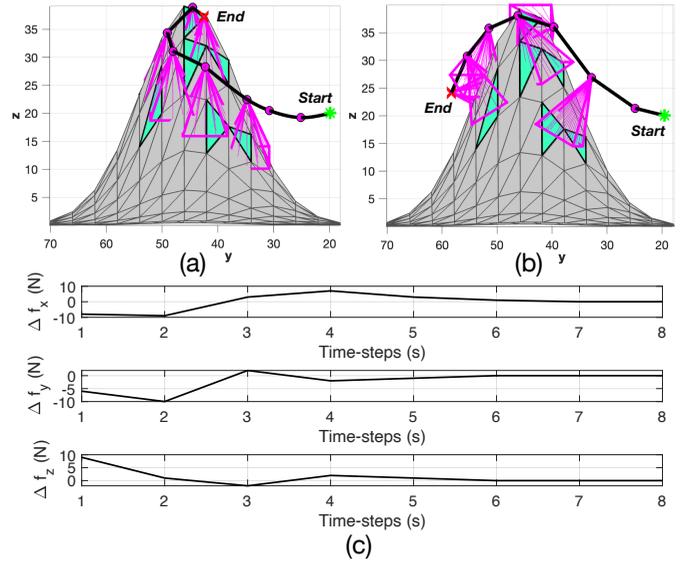


Fig. 10. The effect of different objectives on the coverage trajectory. (a) minimizing the camera gimbal utilization, (b)(c) minimizing the deviations between consecutive motion control inputs.

camera FOVs colored in cyan, as shown in Fig. 7(d). In this scenario the coverage mission was executed in 17 time-steps, no duplication of work occurred, and all facets have been covered according to their visibility constraints.

Figure 8 shows the controller's output at time-step 4. More specifically, the figure shows the agent's planned coverage trajectory during the horizon T of length 5 time-steps i.e., $x_{t+\tau+1|t}$, $t = 4$, $\tau \in \{0, \dots, 4\}$ marked with blue circles, and the predicted FOV configurations which are selected in order to maximize coverage which are also shown in blue color along with the covered facets. The agent's executed trajectory is $x_{1:4}$, and executed FOV states are shown in purple color along with the covered facets. Finally the agent's next state is shown in green color. The FOV states which result in no coverage are omitted from this figure for visual clarity.

Next, Fig. 9 shows the coverage trajectory which was executed for covering the total surface area of the object of interest. Fig. 9(a) shows the front view of the object, whereas Fig. 9(b) shows the back view. The agent states (shown in \circ) and the facets are color-coded according to the coverage time as shown in the figure. In this experiment, total coverage is achieved in 44 time-steps.

Figure 10 shows that various secondary objectives can be incorporated into the proposed 3D coverage controller in order to achieve specific sub-goals. For instance we can design a sub-objective that aims in minimizing the rotations of the onboard camera (i.e., minimize the gimbal utilization), which for instance can increase the device's lifespan. Effectively, we can define this objective as $\mathcal{G}_{\text{FOV}} = \sum_{\tau=1}^{T-1} \sum_{\tilde{m}=1}^{|\mathcal{M}|} (s_{\tilde{m}, t+\tau+1|t} - s_{\tilde{m}, t+\tau|t})$, where as discussed in Sec. 5.3, the binary variable $s_{\tilde{m}, t+\tau+1|t}$ indicates the active FOV state (one out of $|\mathcal{M}|$) at time-step $t + \tau + 1|t$. This particular sub-objective can now be combined with the main cost function \mathcal{G} as a minimization of a weighted multi-objective cost i.e.,:

$$\min (\mathcal{G} + \hat{\omega} \mathcal{G}_{\text{FOV}}) \quad (29)$$

where the parameter $\hat{\omega}$ controls the emphasis given to the secondary objective. Figure 10(a) shows effect of this sub-

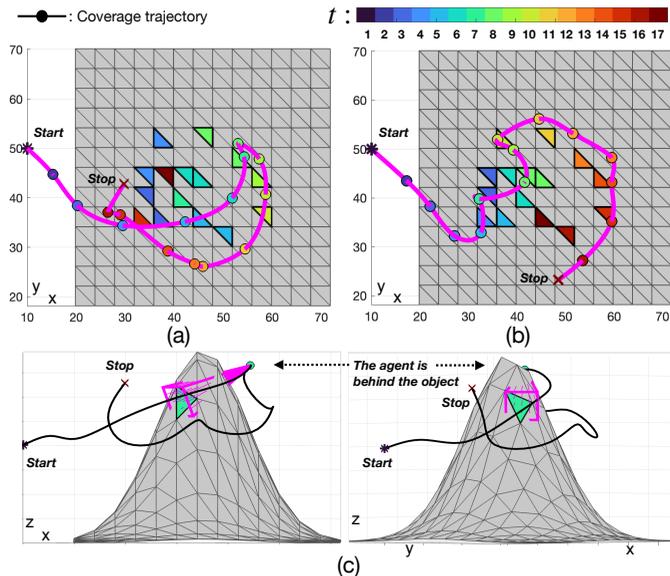


Fig. 11. The effect of ray-tracing and visibility constraints on the coverage planning performance. (a) trajectory without ray-tracing, (b) trajectory with ray-tracing, (c) trajectory without ray-tracing: illustrative example of an erroneous coverage event which occurs at $t = 7$, when visibility constraints are not enabled. The green facet is erroneously marked as covered, although it is not actually visible from the agent's location shown with a green circle.

objective i.e., Eq. (29), on the agent's coverage trajectory over a planning horizon of 8 time-steps, with $\hat{\omega} = 10$. As illustrated in the figure, the agent managed to cover all green-colored facets without rotating its camera's FOV. Effectively, we can say that this coverage trajectory emulates the coverage trajectory obtained with a UAV equipped with a fixed and uncontrollable camera sensor. Observe now from Fig. 10(b) how the coverage trajectory for the same scenario changes when we minimize the control effort which is defined here as the sum of squared deviations between consecutive kinematic control inputs i.e., $\mathcal{G}_{\text{motion}} = \sum_{\tau=1}^{T-1} \|u_{t+\tau|t}^f - u_{t+\tau-1|t}^f\|_2^2$. The coverage multi-objective cost function which now becomes $\min(\mathcal{G} + \hat{\omega}\mathcal{G}_{\text{motion}})$, is particular useful in designing energy efficient coverage trajectories which minimize sudden changes in the direction and speed of the agent, by driving $\mathcal{G}_{\text{motion}} \rightarrow 0$ as shown in Fig. 10(c). In Fig. 10(b)(c) $\hat{\omega}$ has been set to 0.1, and the horizon length is equal to $T = 8$ time-steps. In both scenarios shown in Fig. 10 the agent starts from the same location, and is tasked to cover the same facets which are shown in green color. To summarize, we have shown how different secondary objectives can be incorporated into the main coverage cost function in order to capture various mission specifications. These secondary objectives can be weighted based on their priorities in order to meet the mission requirements.

The majority of coverage planning approaches in the literature primarily focus on 2D environments [16], [39]–[41]. Approaches that address 3D environments often impose restrictive assumptions on the geometry of objects to be covered, such as cuboid-like structures [8], [22], [42]. In contrast, the proposed approach addresses the challenge of coverage planning in 3D environments for objects of arbitrary shapes, which can be represented as triangular meshes. In the proposed approach the agent operates in 3D space, and the object of interest for coverage is also fully

defined in 3D. Additionally, many existing works simplify the coverage planning problem to path planning [23], [45], [74] and also neglect the joint optimization of kinematic and camera control inputs [18], [29], [75]. In contrast, this work addresses the integrated problem of planning and control, and designs a predictive controller that simultaneously optimizes both kinematic and camera control inputs. Furthermore, the proposed approach incorporates visibility determination constraints directly into the control loop, enabling long-horizon planning and optimal solutions. In contrast, most existing approaches address the problem using metaheuristic methods which do not guarantee finding an optimal solution, lack visibility determination, and are limited to myopic planning [9], [17], [22], [44], [76], [77].

To evaluate the performance of the proposed approach in comparison to existing methods, we focus on two key aspects: (a) emphasizing the significance of visibility determination by comparing coverage performance with and without incorporating visibility constraints, effectively illustrating how methods lacking visibility consideration are likely to perform, and (b) examining the influence of planning horizon length on coverage performance, thereby demonstrating the effectiveness of the proposed approach over methods that rely on myopic planning. Finally, (c), we discuss the computational complexity of the proposed approach.

(a) Visibility determination: To assess the importance of visibility determination, we evaluate the proposed approach both with and without the visibility constraints derived via ray tracing. Specifically, we remove the binary variables $b_{\hat{\omega}, \hat{\kappa}}$ and $b_{\hat{\omega}, t+\tau+1|t}^A$ from Eq. (19j), resulting in the following modified coverage constraint: $\bar{b}_{\hat{\kappa}, \hat{m}, t+\tau+1|t} = s_{\hat{m}, t+\tau+1|t} \wedge b_{\hat{\kappa}, \hat{m}, t+\tau+1|t}^V$. The results demonstrate that, although this modified constraint directs the agent toward covering the facets of the object of interest, it does not account for whether a facet κ , located within the agent's camera FOV, is actually visible (i.e., whether a light ray can trace back to κ). Figure 11 illustrates the coverage trajectories of the object of interest with and without visibility constraints, highlighting the impact of incorporating visibility determination. Specifically, Fig. 11(a) provides a top-down view of the agent's coverage trajectory without visibility constraints (i.e., no ray-tracing), while Fig. 11(b) depicts the coverage trajectory obtained with visibility constraints enabled. In Fig. 11(a), the agent is unable to determine whether a particular facet κ within its camera field of view (FOV) is actually visible. As a result, many facets that are not visible but merely fall within the camera FOV are erroneously marked as covered, leading to an inaccurate coverage trajectory, as shown in the figure. In this specific scenario, the agent's camera FOV spans across the object's body. Without visibility constraints, the agent cannot differentiate which facets belong to the foreground (i.e., the visible surface area of the object), as illustrated in Fig. 11(c).

To further assess the importance of visibility determination in identifying visible facets for coverage, a Monte Carlo simulation comprising 100 trials was performed. In each trial, the agent was randomly initialized within the simulated environment depicted in Fig. 7, and the proposed controller was executed both with and without visibility constraints. Coverage facets were randomly selected from the set \mathcal{K} , with the number of facets sampled uniformly

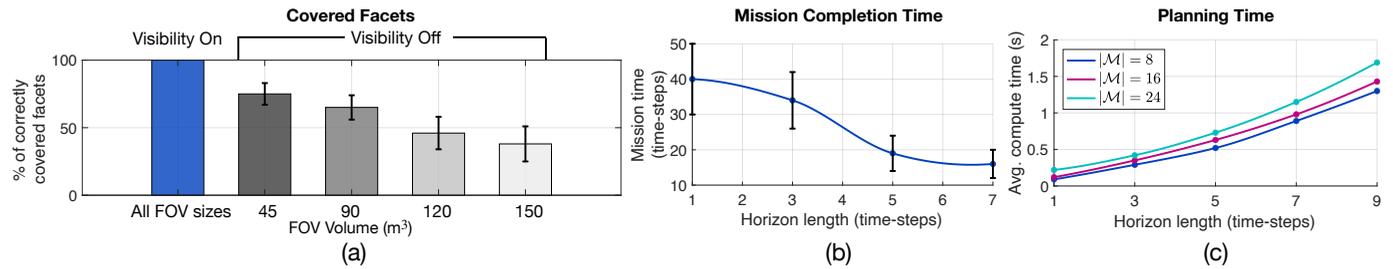


Fig. 12. The figure shows: (a) The percentage of covered facets as a function of the FOV size, with visibility determination turned on and off, (b) the time required to complete a coverage mission (in time-steps) as a function of the horizon length, and (c) the computational time complexity analysis showing the average time required to compute a coverage plan as a function of the horizon length and the number of camera control inputs.

from the range [10, 20]. The results, shown in Fig. 12(a), for various field-of-view (FOV) sizes (measured in volume), demonstrate that when visibility determination is enabled, the controller achieves full coverage of all facets, irrespective of FOV size, without any degradation in performance. This outcome is attributed to the integration of ray-tracing in the proposed method, which ensures accurate identification of visible facets. Conversely, disabling visibility determination leads to a noticeable performance decline as the FOV size increases. This reduction occurs because larger FOVs encompass facets within the convex hull of the FOV that are not genuinely visible. In contrast, smaller FOVs necessitate closer proximity to the object of interest, inherently aligning the agent’s view with the visible facets. Larger FOVs, however, may include a broader set of facets, some of which remain occluded despite being within the FOV. These findings highlight the critical role of visibility determination as a fundamental component in effective coverage planning, a feature that is currently absent in many existing approaches.

(b) Planning horizon: Figure 12(b) illustrates the relationship between the planning horizon length and the average mission completion time. These results were obtained through a Monte Carlo simulation consisting of 100 trials, where the agent was randomly initialized within the environment depicted in Fig. 7. The objective in each trial was to cover 15 facets randomly selected from the object’s surface mesh. The figure reveals that for a planning horizon of $T = 1$ (i.e., myopic planning), the average mission completion time is approximately 40 time-steps. However, as the planning horizon length increases, the mission completion time decreases significantly, stabilizing at around 16 time-steps for $T = 6$ or greater, i.e., 60% improvement in mission completion time. While this behavior is inherently dependent on the specific scenario and problem, it clearly demonstrates the importance of long-horizon planning for enhancing coverage performance. Longer planning horizons enable the agent to anticipate future control inputs, thereby optimizing the mission objective, as defined by Eq. (22). Conversely, myopic approaches, which lack the ability to predict and incorporate future states into the decision-making process, result in reduced performance and longer mission durations.

(c) Computational complexity: The coverage planning problem is formulated as a rolling finite-horizon optimal control problem (FHOCPP) and solved using mixed-integer quadratic programming (MIQP). It is important to emphasize that the proposed approach is optimal in the sense that it computes the control inputs that minimize the ob-

jective function (i.e., Eq. (22)). This is because Problem (P2) is an MIQP with a convex objective function and linear constraints, ensuring the existence of a unique global minimum for any continuous relaxation of the problem. MIQP solvers, utilizing branch-and-bound or branch-and-cut algorithms, systematically explore subsets of the feasible space by enumerating feasible integer solutions. These solvers leverage convexity to prune suboptimal branches, guaranteeing convergence to the global minimum—unlike existing approaches that rely on heuristics. While modern solvers are highly effective at solving MIQPs, combining theoretical guarantees with heuristic accelerations and exploiting problem structure and convexity to efficiently find the optimal solution, the computational complexity of MIQP problems can still grow exponentially in the worst case with respect to the number of decision variables. To evaluate this, we measured the average time required for an agent to compute a coverage plan during the mission depicted in Fig. 7 (where the agent is randomly initialized in the environment and tasked with covering n facets randomly sampled from the range [10, 20]).

Figure 12(c) illustrates the average time required to execute one iteration of Problem (P2) (i.e., generating a coverage trajectory) over 100 random trials, as a function of the planning horizon length (T) and the number of camera control inputs ($|\mathcal{M}|$). As shown, the computational complexity increases as the planning horizon length and the number of camera control inputs grow. The results confirm the theoretical properties of the problem while demonstrating that certain problem sizes can be solved efficiently to optimality. Additionally, the results presented in Fig. 12(b) and Fig. 12(c) offer valuable insights into the trade-offs between performance and computational complexity. These findings can be leveraged to fine-tune the controller, enabling it to meet specific mission requirements effectively. It is worth noting that recent advances in optimization methods, such as adaptive neighborhood search techniques [78] and machine learning-based approaches [79], enable large MIQP problems to be solved efficiently and in real-time. We should note here that in real-world settings, it is crucial to ensure that the time required for computing a coverage plan is either guaranteed, or can be accurately estimated. For this reason, future work will focus on this aspect, as well as on developing strategies to handle computation failures or timeouts encountered during optimization.

Finally, we compare the proposed approach with the sampling-based view-planning methodology i.e., a popular variant of the coverage path planning problem [80]. Sampling-based view-planning methods usually utilize a two-stage approach. First, candidate viewpoints are sam-

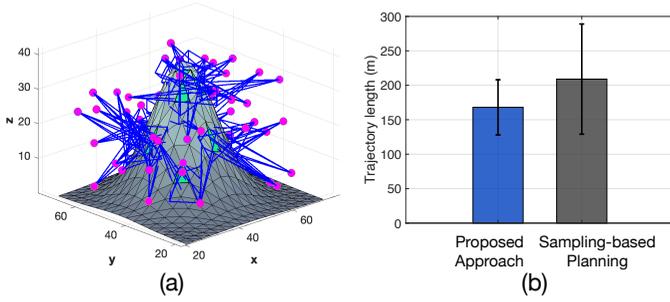


Fig. 13. (a) Illustrative example of sampled viewpoint configurations, (b) Performance comparison of the proposed approach with sampling-based view planning methods.

pled within a constrained space around the target object (e.g., Fig. 13(a)). In the second stage, a subset of these viewpoints is selected by solving an approximate version of the set covering problem (SCP) using heuristics such as greedy search, genetic algorithms, or hybrid methods. The view planning problem is nominally formulated as the following constrained optimization problem:

$$\begin{aligned}
 \min \quad & \sum_{j=1}^N \psi_j \\
 \text{s.t.} \quad & \Omega \cdot \psi \geq \mathbf{1}, \\
 & d(j, l) \leq \Delta_t \|\mathbf{v}_{\max}\|, \quad \forall j \neq l, \\
 & \psi_j \in \{0, 1\}, \quad \forall j, \\
 & \Omega(i, j) \in \{0, 1\}, \quad \forall i, j, \\
 & i \in \{1, \dots, M\}, j, l \in \{1, \dots, N\}.
 \end{aligned} \quad (30)$$

In Eq.(30), ψ_j is a binary decision variable indicating whether the j th viewpoint is selected. The binary visibility matrix $\Omega \in \{0, 1\}^{M \times N}$ defines visibility constraints, where M is the number of object surface facets requiring coverage, and N is the number of sampled viewpoints. If $\Omega(i, j) = 1$, facet i is visible and covered by viewpoint j . The constraint on the distance $d(j, l)$ ensures a minimum distance between connected viewpoints i and j , determined by the UAV's allowed maximum velocity \mathbf{v}_{\max} .

To evaluate our approach against the sampling-based view-planning technique i.e., Eq. (30), we conducted 40 Monte Carlo simulations, randomly selecting 10 facets from the object's surface (using the simulation setup described and shown in Fig. 7) and varying the UAV's initial position, measuring the length of the coverage trajectory as an indication of the coverage performance. Unlike the proposed approach, which directly optimizes the UAV trajectory, view-planning methods generate a coverage path (i.e., a sequence of connected viewpoints). In order to compare the two approaches, we convert the coverage path P , generated by the view-planning approach, into a feasible UAV trajectory by solving the following receding horizon control problem over a rolling finite horizon of length H time-steps:

$$\begin{aligned}
 \min \quad & \sum_{t=0}^{H-1} \|x_t^p - \hat{P}_t\|^2 \\
 \text{s.t.} \quad & x_{t+1} = Ax_t + Bf_t, \quad \forall t \in \{0, \dots, H-1\}, \\
 & x_0 = x_{\text{init}}, \\
 & x_t \in \mathcal{X}, f_t \in \mathcal{F}, \quad \forall t \in \{0, \dots, H-1\}.
 \end{aligned} \quad (31)$$

In essence we want to find the UAV control inputs $\{f_0, \dots, f_{H-1}\}$ in a rolling horizon fashion (we have used

$H = 5$) that optimally track the reference path \hat{P}_t . Here, \hat{P}_t is a continuous reference path obtained from P via spline interpolation [81]. Figure 13(a) illustrates an example of the sampled viewpoint configurations (i.e., UAV positions and camera orientations) generated in the first stage of the view-planning approach. In each simulation, we sample 200 valid viewpoints around the object of interest. The coverage path is then computed by solving Eq. (30) using a greedy local search strategy [82]. Figure 13(b) illustrates the average trajectory length obtained using the two approaches, demonstrating that the UAV coverage trajectory generated by the proposed approach is approximately 25% more efficient (i.e., shorter in terms of distance) compared to that of the view planning approach. The proposed method jointly optimizes the UAV's motion and camera control inputs under visibility constraints by solving an optimal control problem, whereas the view planning approach approximates this behavior through sampling.

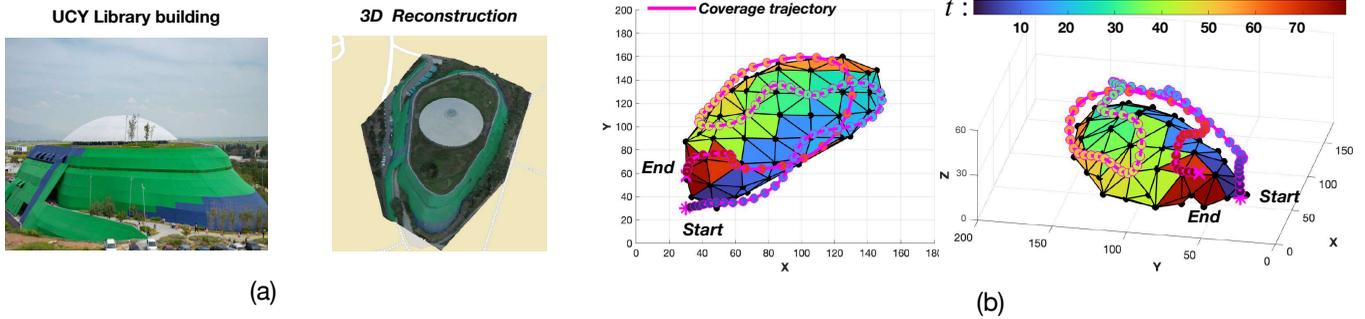
6.2 Real-world Experiment

In the second part of the evaluation, we demonstrate the performance of the proposed approach in a real-world coverage mission. More specifically, the main objective here is to compute the 3D coverage plan for the Library building of the University of Cyprus. The library building, shown in Fig. 14(a), is an oval-like structure with major and minor axis diameters of about 170m and 85m respectively, and a height of about 30m. As discussed in Sec. 5.1, as a first step we have used 3D reconstruction techniques to build a 3D model of the library structure i.e., the library building was first reconstructed as a point-cloud using Open3D [83], as shown in Fig. 14(a), and then converted into a triangle mesh using Delaunay triangulation. As a result, a total of 82 triangles were used to model the surface of the library building as shown in Fig. 14(b).

For our experimental evaluation we have used the DJI Mavic Enterprise consumer drone, equipped with a 12MP gimballed camera system, with horizontal and vertical rotation angles in the ranges $\Phi = \{300, 315, 0, 45, 60\}$ deg, and $\Theta = \{0, 30, 90\}$ deg respectively, and a square FOV footprint, at a height of approximately 18m, equal to 20m-by-20m. The zoom functionality was not utilized for this test i.e., $Z = \{1\}$, and the rest of the parameters have been set as discussed in Sec. 6.1. Figure 14(b) shows the generated coverage trajectory obtained with the proposed technique. In this figure the UAV's trajectory is color-coded based on the mission elapsed time. Similarly, the object's facets are colored based on the time-step which have been observed by the UAV (i.e., based on the coverage-time).

As illustrated in the figure the generated trajectory enabled the coverage of the total surface area of the structure (the camera FOV configurations are omitted for visual clarity). Subsequently in order to verify the generated coverage plan in practice, we have utilized our automated drone tasking platform [84] to command the UAV to execute the coverage trajectory generated with the proposed approach. Figure 14(c) shows an instance of the executed coverage plan during flight-time. Specifically, the figure shows a snapshot of the live-feed obtained from our mission control platform while executing the coverage trajectory. Figure 14(d) shows in top-down view of the executed coverage trajectory in red. The green waypoints shown in the figure correspond to the latitude-longitude coordinates of the

3D coverage planning in simulation



3D coverage plan execution in real-world settings

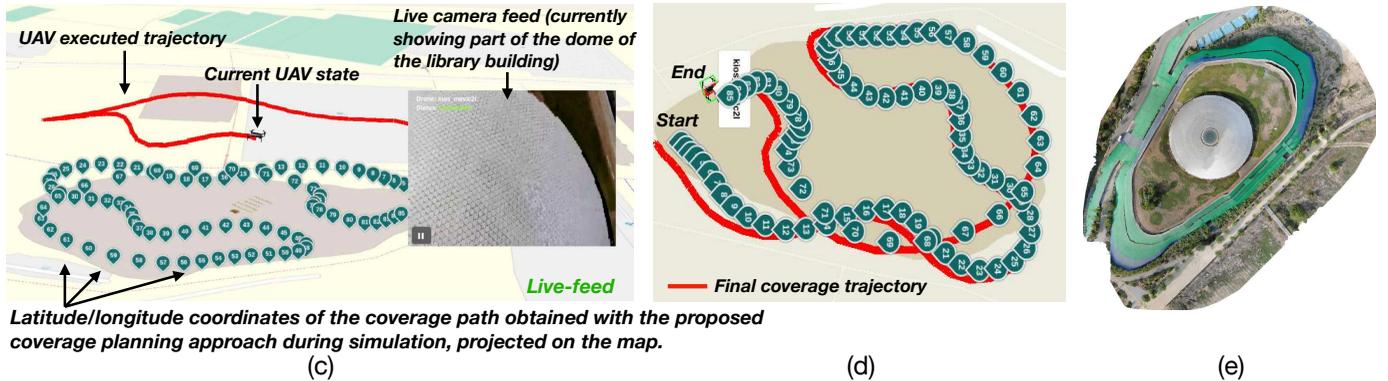


Fig. 14. Real-world 3D coverage mission. (a) Object of interest, (b) Coverage trajectory obtained with the proposed approach, (c)-(e) The generated coverage plan executed in real-world settings.

planned trajectory computed with the proposed approach during simulation. Finally, Fig. 14(e) shows the building's orthophotomap generated from the images acquired during the real-world mission, indicating the covered surface area of the structure i.e., the proposed approach achieves full coverage by minimizing the cost function shown in Eq. (22).

Implementation details and lessons learned: The UAV was operated using the DJI Mobile Software Development Kit (SDK), which facilitated the creation of a custom Android mobile application installed on a smartphone paired with the UAV. This smartphone acted as an intermediary by wirelessly connecting to the UAV's remote controller, enabling the transmission of commands to the UAV and the reception of telemetry data such as GPS coordinates, altitude, battery status, speed, and sensor readings. The proposed coverage controller was implemented and executed on a ground control station (GCS) using Matlab. The GCS computed the UAV trajectories and transmitted them to the drone via the mobile application on the Android smartphone connected to the remote controller. Telemetry data from the UAV was also relayed back to the GCS through this setup. Communication between the GCS and the UAV was facilitated by a VPN server over a wireless network, allowing data exchange between the two. Mission monitoring was conducted live using a custom-built UAV tasking platform [84].

While this prototype setup demonstrated the effectiveness of the proposed approach, it highlighted several areas for future improvement. GPS-based positioning introduced localization errors and inconsistencies, leading to distorted coverage plans in certain scenarios. Additionally, network latency resulted in communication delays that impacted the drone's operation, and environmental disturbances caused deviations from the planned trajectories. Future work will

aim to address these challenges by enhancing the approach to incorporate stochastic and robust predictive control methods for managing environmental uncertainties, integrating the controller with the UAV's onboard systems to improve real-time responsiveness, and scaling the methodology to support multiple UAV agents.

7 CONCLUSION

In this work we propose a jointly-optimized trajectory generation and camera control approach for automated coverage planning in 3D environments. The proposed coverage planning approach integrates ray-tracing into the coverage planning process thus allowing an autonomous mobile agent to determine the visible parts of the object of interest; and generate look-ahead coverage trajectories by jointly optimizing its kinematic and camera control inputs over a rolling finite planning horizon. We show how the coverage planning problem can be formulated as a finite horizon optimal control problem, and then solved using mixed integer programming. Finally, the performance of the proposed approach is demonstrated through extensive synthetic and real-world experiments.

ACKNOWLEDGMENTS

The project is implemented under the Border Management and Visa Policy Instrument (BMVI) and is co-financed by the European Union and the Republic of Cyprus (GA:BMVI/2021-2022/SA/1.2.1/015), and supported by the GLIMPSE project EXCELLENCE/0421/0586 co-financed by the European Regional Development Fund and the Republic of Cyprus through the Research and Innovation Foundation's RESTART 2016-2020 Programme for Research,

Technological Development and Innovation, and by the European Union's Horizon 2020 research and innovation programme under grant agreement No 739551 (KIOS CoE).

REFERENCES

- [1] S. Papaioannou, P. Kolios, C. G. Panayiotou, and M. M. Polycarpou, "Cooperative simultaneous tracking and jamming for disabling a rogue drone," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, pp. 7919–7926.
- [2] S. Papaioannou, P. Kolios, and G. Ellinas, "Downing a rogue drone with a team of aerial radio signal jammers," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, pp. 2555–2562.
- [3] D. Saha, D. Pattanayak, and P. S. Mandal, "Surveillance of uneven surface with self-organizing unmanned aerial vehicles," *IEEE Transactions on Mobile Computing*, vol. 21, no. 4, pp. 1449–1462, 2022.
- [4] W. Wang, H. Dai, C. Dong, F. Xiao, J. Zheng, X. Cheng, G. Chen, and X. Fu, "Deployment of unmanned aerial vehicles for anisotropic monitoring tasks," *IEEE Transactions on Mobile Computing*, vol. 21, no. 2, pp. 495–513, 2022.
- [5] S. Papaioannou, P. Kolios, and G. Ellinas, "Distributed estimation and control for jamming an aerial target with multiple agents," *IEEE Transactions on Mobile Computing*, vol. 22, no. 12, pp. 7203–7217, 2023.
- [6] P. Ramon-Soria, M. Perez-Jimenez, B. Arrue, and A. Ollero, "Planning system for integrated autonomous infrastructure inspection using uavs," in *2019 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2019, pp. 313–320.
- [7] N. Bartolini, A. Coletta, G. Maselli, and A. Khalifeh, "A multi-trip task assignment for early target inspection in squads of aerial drones," *IEEE Transactions on Mobile Computing*, vol. 20, no. 11, pp. 3099–3116, 2021.
- [8] S. Papaioannou, P. Kolios, T. Theocharides, C. G. Panayiotou, and M. M. Polycarpou, "Distributed search planning in 3-D environments with a dynamically varying number of agents," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 53, no. 7, pp. 4117–4130, 2023.
- [9] A. Zacharia, S. Papaioannou, P. Kolios, and C. Panayiotou, "Distributed control for 3D inspection using multi-UAV systems," in *2023 31st Mediterranean Conference on Control and Automation (MED)*, 2023, pp. 164–169.
- [10] S. Papaioannou, P. Kolios, T. Theocharides, C. G. Panayiotou, and M. M. Polycarpou, "A cooperative multiagent probabilistic framework for search and track missions," *IEEE Transactions on Control of Network Systems*, vol. 8, no. 2, pp. 847–858, 2021.
- [11] —, "Decentralized search and track with multiple autonomous agents," in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 909–915.
- [12] —, "3D Trajectory Planning for UAV-based Search Missions: An Integrated Assessment and Search Planning Approach," in *2021 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2021, pp. 517–526.
- [13] —, "Towards Automated 3D Search Planning for Emergency Response Missions," *Journal of Intelligent & Robotic Systems*, vol. 103, no. 1, pp. 1–19, 2021.
- [14] S. Papaioannou, S. Kim, C. Laoudias, P. Kolios, S. Kim, T. Theocharides, C. Panayiotou, and M. Polycarpou, "Coordinated crlb-based control for tracking multiple first responders in 3d environments," in *2020 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2020, pp. 1475–1484.
- [15] S. Papaioannou, P. Kolios, C. G. Panayiotou, and M. M. Polycarpou, "Synergising human-like responses and machine intelligence for planning in disaster response," in *2024 International Joint Conference on Neural Networks (IJCNN)*, 2024, pp. 1–8.
- [16] T. M. Cabreira, L. B. Brisolaro, and P. R. Ferreira Jr, "Survey on coverage path planning with unmanned aerial vehicles," *Drones*, vol. 3, no. 1, p. 4, 2019.
- [17] A. Khan, I. Noreen, and Z. Habib, "On complete coverage path planning algorithms for non-holonomic mobile robots: Survey and challenges," *Journal of Information Science & Engineering*, vol. 33, no. 1, 2017.
- [18] M. Torres, D. A. Pelta, J. L. Verdegay, and J. C. Torres, "Coverage path planning with unmanned aerial vehicles for 3d terrain reconstruction," *Expert Systems with Applications*, vol. 55, pp. 441–451, 2016.
- [19] R. Almadhoun, T. Taha, L. Seneviratne, and Y. Zweiri, "A survey on multi-robot coverage path planning for model reconstruction and mapping," *SN Applied Sciences*, vol. 1, no. 8, pp. 1–24, 2019.
- [20] P. Fazli, A. Davoodi, P. Pasquier, and A. K. Mackworth, "Complete and robust cooperative robot area coverage with limited range," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 5577–5582.
- [21] A. Nedjati, G. Izbirak, B. Vizvari, and J. Arkat, "Complete coverage path planning for a multi-uav response system in post-earthquake assessment," *Robotics*, vol. 5, no. 4, p. 26, 2016.
- [22] C. S. Tan, R. Mohd-Mokhtar, and M. R. Arshad, "A comprehensive review of coverage path planning in robotics using classical and heuristic algorithms," *IEEE Access*, vol. 9, pp. 119 310–119 342, 2021.
- [23] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1258–1276, 2013.
- [24] E. U. Acar and H. Choset, "Sensor-based coverage of unknown environments: Incremental construction of morse decompositions," *The International Journal of Robotics Research*, vol. 21, no. 4, pp. 345–366, 2002.
- [25] D. Batsaikhan, A. Janchiv, and S.-G. Lee, "Sensor-based incremental boustrophedon decomposition for coverage path planning of a mobile robot," in *Intelligent Autonomous Systems 12*. Springer, 2013, pp. 621–628.
- [26] S. C. Wong and B. A. MacDonald, "A topological coverage algorithm for mobile robots," in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, vol. 2. IEEE, 2003, pp. 1685–1690.
- [27] N. Agmon, N. Hazon, and G. A. Kaminka, "Constructing spanning trees for efficient multi-robot coverage," in *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006*. IEEE, 2006, pp. 1698–1703.
- [28] M. Wei and V. Isler, "Coverage path planning under the energy constraint," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 368–373.
- [29] X. Kan, H. Teng, and K. Karydis, "Online exploration and coverage planning in unknown obstacle-cluttered environments," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5969–5976, 2020.
- [30] N. Hazon, F. Miel, and G. A. Kaminka, "Towards robust online multi-robot coverage," in *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006*. IEEE, 2006, pp. 1710–1715.
- [31] C. Gao, Y. Kou, Z. Li, A. Xu, Y. Li, and Y. Chang, "Optimal multirobot coverage path planning: ideal-shaped spanning tree," *Mathematical Problems in Engineering*, vol. 2018, 2018.
- [32] A. Sipahioglu, G. Kirlik, O. Parlaktuna, and A. Yazici, "Energy constrained multi-robot sensor-based coverage path planning using capacitated arc routing approach," *Robotics and Autonomous Systems*, vol. 58, no. 5, pp. 529–538, 2010.
- [33] I. Rekleitis, A. P. New, E. S. Rankin, and H. Choset, "Efficient boustrophedon multi-robot coverage: an algorithmic approach," *Annals of Mathematics and Artificial Intelligence*, vol. 52, no. 2, pp. 109–142, 2008.
- [34] E. Tolstaya, J. Paulos, V. Kumar, and A. Ribeiro, "Multi-robot coverage and exploration using spatial graph neural networks," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 8944–8950.
- [35] A. Anastasiou, A. Zacharia, S. Papaioannou, P. Kolios, C. G. Panayiotou, and M. M. Polycarpou, "Automated real-time inspection in indoor and outdoor 3d environments with cooperative aerial robots," in *2024 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2024, pp. 496–504.
- [36] S. Papaioannou, C. Vitale, P. Kolios, C. G. Panayiotou, and M. M. Polycarpou, "Hierarchical fault-tolerant coverage control for an autonomous aerial agent," *IFAC-PapersOnLine*, vol. 58, no. 4, pp. 532–537, 2024. 12th IFAC Symposium on Fault Detection, Supervision and Safety for Technical Processes SAFEPROCESS 2024.
- [37] S. Papaioannou, P. Kolios, T. Theocharides, C. G. Panayiotou, and M. M. Polycarpou, "Unscented optimal control for 3d coverage planning with an autonomous uav agent," in *2023 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2023, pp. 703–712.
- [38] Y.-S. Jiao, X.-M. Wang, H. Chen, and Y. Li, "Research on the coverage path planning of uavs for polygon areas," in *2010 5th IEEE Conference on Industrial Electronics and Applications*, 2010, pp. 1467–1472.
- [39] J. Araujo, P. Sujit, and J. Sousa, "Multiple uav area decomposition and coverage," in *2013 IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA)*, 2013, pp. 30–37.
- [40] J. Xie, L. R. Garcia Carrillo, and L. Jin, "Path planning for uav to cover multiple separated convex polygonal regions," *IEEE Access*, vol. 8, pp. 51 770–51 785, 2020.

- [41] S. Papaioannou, P. Kolios, T. Theocharides, C. G. Panayiotou, and M. M. Polycarpou, "Integrated guidance and gimbal control for coverage planning with visibility constraints," *IEEE Transactions on Aerospace and Electronic Systems*, doi:10.1109/TAES.2022.3199196, pp. 1–15, 2022.
- [42] —, "UAV-based receding horizon control for 3D inspection planning," in *2022 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2022, pp. 1121–1130.
- [43] G. S. Avellar, G. A. Pereira, L. C. Pimenta, and P. Iscold, "Multi-uav routing for area coverage and remote sensing with minimum time," *Sensors*, vol. 15, no. 11, pp. 27 783–27 803, 2015.
- [44] S. D. Apostolidis, P. C. Kapoutsis, A. C. Kapoutsis, and E. B. Kosmatopoulos, "Cooperative multi-uav coverage mission planning platform for remote sensing applications," *Autonomous Robots*, vol. 46, no. 2, pp. 373–400, 2022.
- [45] L. Collins, P. Ghassemi, E. T. Esfahani, D. Doermann, K. Dantu, and S. Chowdhury, "Scalable coverage path planning of multi-robot teams for monitoring non-convex areas," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 7393–7399.
- [46] H. I. Perez-Imaz, P. A. Rezeck, D. G. Macharet, and M. F. Campos, "Multi-robot 3d coverage path planning for first responders teams," in *2016 IEEE International Conference on Automation Science and Engineering (CASE)*. IEEE, 2016, pp. 1374–1379.
- [47] K. G. S. Apuroop, A. V. Le, M. R. Elara, and B. J. Sheu, "Reinforcement learning-based complete area coverage path planning for a modified htrihex robot," *Sensors*, vol. 21, no. 4, p. 1067, 2021.
- [48] S. Papaioannou, P. Kolios, T. Theocharides, C. G. Panayiotou, and M. M. Polycarpou, "Integrated ray-tracing and coverage planning control using reinforcement learning," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 7200–7207.
- [49] A. K. Lakshmanan, R. E. Mohan, B. Ramalingam, A. V. Le, P. Veerajagadeshwar, K. Tiwari, and M. Ilyas, "Complete coverage path planning using reinforcement learning for tetromino based cleaning and maintenance robot," *Automation in Construction*, vol. 112, p. 103078, 2020.
- [50] J. Torres-Sánchez, F. López-Granados, N. Serrano, O. Arquero, and J. M. Peña, "High-throughput 3-D monitoring of agricultural-tree plantations with unmanned aerial vehicle (UAV) technology," *PloS one*, vol. 10, no. 6, p. e0130479, 2015.
- [51] S. Zhao, F. Kang, J. Li, and C. Ma, "Structural health monitoring and inspection of dams based on UAV photogrammetry with image 3d reconstruction," *Automation in Construction*, vol. 130, p. 103832, 2021.
- [52] S. Jiang, W. Jiang, and L. Wang, "Unmanned aerial vehicle-based photogrammetric 3D mapping: A survey of techniques, applications, and challenges," *IEEE Geoscience and Remote Sensing Magazine*, vol. 10, no. 2, pp. 135–171, 2021.
- [53] S. Hu, W. Ni, X. Wang, and A. Jamalipour, "Disguised tailing and video surveillance with solar-powered fixed-wing unmanned aerial vehicle," *IEEE Transactions on Vehicular Technology*, vol. 71, no. 5, pp. 5507–5518, 2022.
- [54] K. S. Kappel, T. M. Cabreira, J. L. Marins, L. B. de Brisolará, and P. R. Ferreira, "Strategies for patrolling missions with multiple UAVs," *Journal of Intelligent & Robotic Systems*, vol. 99, pp. 499–515, 2020.
- [55] F. Outay, H. A. Mengash, and M. Adnan, "Applications of unmanned aerial vehicle (UAV) in road safety, traffic and highway infrastructure management: Recent advances and challenges," *Transportation research part A: policy and practice*, vol. 141, pp. 116–129, 2020.
- [56] C. Zhao, J. Liu, M. Sheng, W. Teng, Y. Zheng, and J. Li, "Multi-UAV trajectory planning for energy-efficient content coverage: A decentralized learning-based approach," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 10, pp. 3193–3207, 2021.
- [57] L. Wang, K. Wang, C. Pan, W. Xu, N. Aslam, and L. Hanzo, "Multi-agent deep reinforcement learning-based trajectory planning for multi-UAV assisted mobile edge computing," *IEEE Transactions on Cognitive Communications and Networking*, vol. 7, no. 1, pp. 73–84, 2020.
- [58] Q. Wu, Y. Zeng, and R. Zhang, "Joint trajectory and communication design for multi-UAV enabled wireless networks," *IEEE Transactions on Wireless Communications*, vol. 17, no. 3, pp. 2109–2121, 2018.
- [59] J. Gu, H. Wang, G. Ding, Y. Xu, Z. Xue, and H. Zhou, "Energy-Constrained Completion Time Minimization in UAV-Enabled Internet of Things," *IEEE Internet of Things Journal*, vol. 7, no. 6, pp. 5491–5503, 2020.
- [60] X. Gao, X. Zhu, and L. Zhai, "Minimization of Aerial Cost and Mission Completion Time in Multi-UAV-Enabled IoT Networks," *IEEE Transactions on Communications*, vol. 71, no. 9, pp. 5335–5347, 2023.
- [61] W. Du, T. Wang, H. Zhang, Y. Dong, and Y. Li, "Joint Resource Allocation and Trajectory Optimization for Completion Time Minimization for Energy-Constrained UAV Communications," *IEEE Transactions on Vehicular Technology*, vol. 72, no. 4, pp. 4568–4579, 2023.
- [62] S. Baldi, S. Roy, and K. Yang, "Towards adaptive autopilots for fixed-wing unmanned aerial vehicles," in *2020 59th IEEE Conference on Decision and Control (CDC)*, 2020, pp. 4724–4729.
- [63] G. H. Elkaim, F. A. P. Lie, and D. Gebre-Egziabher, "Principles of guidance, navigation, and control of UAVs," *Handbook of unmanned aerial vehicles*, pp. 347–380, 2015.
- [64] G. Garcia and S. Keshmiri, "Nonlinear model predictive controller for navigation, guidance and control of a fixed-wing UAV," in *AIAA guidance, navigation, and control conference*, 2011, p. 6310.
- [65] Z. Yang, F. Lin, and B. M. Chen, "Survey of autopilot for multi-rotor unmanned aerial vehicles," in *IECON 2016 - 42nd Annual Conference of the IEEE Industrial Electronics Society*, 2016, pp. 6122–6127.
- [66] P. Labatut, J.-P. Pons, and R. Keriven, "Efficient multi-view reconstruction of large-scale scenes using interest points, delaunay triangulation and graph cuts," in *2007 IEEE 11th international conference on computer vision*. IEEE, 2007, pp. 1–8.
- [67] A. Dai and M. Nießner, "Scan2mesh: From unstructured range scans to 3d meshes," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5574–5583.
- [68] D.-T. Lee and B. J. Schachter, "Two algorithms for constructing a delaunay triangulation," *International Journal of Computer & Information Sciences*, vol. 9, no. 3, pp. 219–242, 1980.
- [69] P. Wang, Z. Wang, S. Xin, X. Gao, W. Wang, and C. Tu, "Restricted delaunay triangulation for explicit surface reconstruction," *ACM Transactions on Graphics (TOG)*, 2022.
- [70] N. Sharp and M. Ovsjanikov, "Pointtrinet: Learned triangulation of 3d point sets," in *European Conference on Computer Vision*. Springer, 2020, pp. 762–778.
- [71] Y. Liu, W. Wang, B. Lévy, F. Sun, D.-M. Yan, L. Lu, and C. Yang, "On centroidal voronoi tessellation-energy smoothness and fast computation," *ACM Transactions on Graphics (ToG)*, vol. 28, no. 4, pp. 1–17, 2009.
- [72] T. Achterberg and R. Wunderling, "Mixed integer programming: Analyzing 12 years of progress," in *Facets of combinatorial optimization*. Springer, 2013, pp. 449–481.
- [73] L. Gurobi Optimization, "Gurobi optimizer reference manual, 2022," URL <http://www.gurobi.com>, vol. 19.
- [74] M. Jones, S. Djahel, and K. Welsh, "Path-planning for unmanned aerial vehicles with environment complexity considerations: A survey," *ACM Comput. Surv.*, vol. 55, no. 11, feb 2023.
- [75] T. Elmokadem, "Distributed coverage control of quadrotor multi-UAV systems for precision agriculture," *IFAC-PapersOnLine*, vol. 52, no. 30, pp. 251–256, 2019.
- [76] W. Jing, D. Deng, Y. Wu, and K. Shimada, "Multi-UAV coverage path planning for the inspection of large and complex structures," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 1480–1486.
- [77] S. Huang, R. S. H. Teo, and W. L. Leong, "Review of coverage control of multi unmanned aerial vehicles," in *2017 11th Asian Control Conference (ASCC)*, 2017, pp. 228–232.
- [78] G. Hendel, "Adaptive large neighborhood search for mixed integer programming," *Mathematical Programming Computation*, pp. 1–37, 2021.
- [79] D. Bertsimas and B. Stellato, "Online mixed-integer optimization in milliseconds," *INFORMS Journal on Computing*, vol. 34, no. 4, pp. 2229–2248, 2022.
- [80] M. Maboudi, M. Homaei, S. Song, S. Malihi, M. Saadatseresht, and M. Gerke, "A review on viewpoints and path planning for uav-based 3-d reconstruction," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 16, pp. 5026–5048, 2023.
- [81] H. Wang, J. Kearney, and K. Atkinson, "Arc-length parameterized spline curves for real-time simulation," in *Proc. 5th International Conference on Curves and Surfaces*, vol. 387396, 2002.
- [82] W. Jing, J. Polden, P. Y. Tao, W. Lin, and K. Shimada, "View planning for 3d shape reconstruction of buildings with unmanned aerial vehicles," in *2016 14th International Conference on Control, Automation, Robotics and Vision (ICARCV)*. IEEE, 2016, pp. 1–6.
- [83] Q.-Y. Zhou, J. Park, and V. Koltun, "Open3D: A modern library for 3D data processing," *arXiv:1801.09847*, 2018.
- [84] M. Terzi, P. Kolios, C. Panayiotou, and T. Theocharides, "A unified framework for reliable multi-drone tasking in emergency response missions," in *2019 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2019, pp. 819–827.



Savvas Papaioannou is a Senior Research Scientist at the KIOS Research and Innovation Center of Excellence (KIOS CoE) at the University of Cyprus (UCY). He earned his B.Sc. degree in Electronic and Computer Engineering from the Technical University of Crete (Greece) in 2011, his M.Sc. degree in Electrical Engineering from Yale University (USA) in 2013, and his DPhil (Ph.D.) in Computer Science from the University of Oxford (UK) in 2017. Dr. Papaioannou's research interests focus on the modeling, estimation, control, and optimization of autonomous intelligent systems. His work spans multi-agent systems, intelligent unmanned aerial vehicles (UAVs), multi-target tracking, artificial intelligence, and their applications. His work has received several awards, including the Best Paper Award at the 10th IEEE Symposium on Embedded Systems for Real-time Multimedia (ESTIMedia) in 2012, the First Prize in the Cooperative Aerial Robots Inspection Challenge at the 62nd IEEE Conference on Decision and Control (CDC) in 2023, and the Runner-up Award (2nd Prize) in the Multi-Robot Perception and Navigation Challenges in Logistics and Inspection Tasks competition at the 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). Dr. Papaioannou serves as a reviewer for numerous IEEE and ACM conferences and journals. Additionally, he has been actively involved in the organization of international conferences, including the 2018 European Control Conference (ECC), the 2022 IFAC Symposium on Fault Detection, Supervision, and Safety for Technical Processes (Safeprocess), the 2023 Mediterranean Conference on Control and Automation (MED), and the 6th International Conference on Control and Fault-Tolerant Systems (SysTol) in 2025.



Panayiotis Kolios is currently a Research Assistant Professor at the KIOS Research and Innovation Centre of Excellence of the University of Cyprus. He received his BEng and PhD degree in Telecommunications Engineering from King's College London in 2008 and 2011, respectively. Before joining the KIOS CoE, he worked at the Department of Communications and Internet Studies at the Cyprus University of Technology and the Department of Computer Science of the University of Cyprus (UCY). His

work focuses on both basic and applied research on networked intelligent systems. Some examples of systems that fall into the latter category include intelligent transportation systems, autonomous unmanned aerial systems and the plethora of cyber-physical systems that arise within the Internet of Things. Particular emphasis is given to emergency management in which natural disasters, technological faults and man-made attacks could cause disruptions that need to be effectively handled. Tools used include graph theoretic approaches, algorithmic development, mathematical and dynamic programming, as well as combinatorial optimization.



Theocharis Theocharides is an Associate Professor in the Department of Electrical and Computer Engineering, at the University of Cyprus and a faculty member of the KIOS Research and Innovation Center of Excellence where he serves as the Research Director. Theocharis received his Ph.D. in Computer Engineering from Penn State University, working in the areas of low-power computer architectures and reliable system design, where he was honored with the Robert M. Owens Memorial Scholarship, in May

2005. He has been with the Electrical and Computer Engineering department at the University of Cyprus since 2006, where he directs the Embedded and Application-Specific Systems-on-Chip Laboratory. His research focuses on the design, development, implementation, and deployment of low-power and reliable on-chip application-specific architectures, low-power VLSI design, real-time embedded systems design and exploration of energy-reliability trade-offs for Systems on Chip and Embedded Systems. His focus lies on acceleration of computer vision and artificial intelligence algorithms in hardware, geared towards edge computing, and in utilizing reconfigurable hardware towards self-aware, evolvable edge computing systems. His research has been funded by several National and European agencies and the industry, and he is currently involved in over ten funded ongoing research projects. He serves on several organizing and technical program committees of various conferences (currently serving as the Application Track Chair for the DATE Conference), is a Senior Member of the IEEE and a member of the ACM. He is currently an Associate Editor for the ACM Transactions on Emerging Technologies in Computer Systems, IEEE Consumer Electronics magazine, IET's Computers and Digital Techniques, the ETRI journal and Springer Nature Computer Science. He also serves on the Editorial Board of IEEE Design & Test magazine.



Christos Panayiotou is a Professor with the Electrical and Computer Engineering (ECE) Department at the University of Cyprus (UCY). He is also the Acting Director of the KIOS Research and Innovation Center of Excellence for which he is also a founding member. Christos has received a B.Sc. and a Ph.D. degree in Electrical and Computer Engineering from the University of Massachusetts at Amherst, in 1994 and 1999 respectively. He also received an MBA from the Isenberg School of Management, at the aforementioned university in 1999. Before joining the University of Cyprus in 2002, he was a Research Associate at the Center for Information and System Engineering (CISE) and the Manufacturing Engineering Department at Boston University (1999 - 2002). His research interests include modeling, control, optimization and performance evaluation of intelligent transportation systems, cyber-physical systems, discrete event and hybrid systems, event detection and localization, fault diagnosis, machine learning, wireless, ad hoc and sensor networks, resource allocation, and intelligent buildings. Christos has published more than 340 papers in international refereed journals and conferences and is the recipient of the 2014 Best Paper Award for the journal Building and Environment (Elsevier). He is an Associate Editor for the IEEE Transactions of Intelligent Transportation Systems, and the Journal of Discrete Event Dynamical Systems. He had a long tenure on the Conference Editorial Board of the IEEE Control Systems Society (2001-2023) and served as an Associate Editor of the IEEE Transactions on Control Systems Technology during 2016-2020 and of the European Journal of Control during 2014-2021. He held several positions in organizing committees and technical program committees of numerous international conferences, including General Chair of the 31st Mediterranean Conference on Control and Automation (MED2023), General Chair of the 23rd European Working Group on Transportation (EWGT2020), and General Co-Chair of the 2018 European Control Conference (ECC2018).



Marios M. Polycarpou is a Professor of Electrical and Computer Engineering and the Chair of the Advisory Board of the KIOS Research and Innovation Center of Excellence at the University of Cyprus. He is also a Founding Member of the Cyprus Academy of Sciences, Letters, and Arts, an Honorary Professor of Imperial College London, and a Member of Academia Europaea (The Academy of Europe). He received the B.A degree in Computer Science and the B.Sc. in Electrical Engineering, both from Rice

University, USA in 1987, and the M.S. and Ph.D. degrees in Electrical Engineering from the University of Southern California, in 1989 and 1992 respectively. His teaching and research interests are in intelligent systems and networks, adaptive and learning control systems, fault diagnosis, machine learning, and critical infrastructure systems. Prof. Polycarpou is the recipient of the 2023 IEEE Frank Rosenblatt Technical Field Award and the 2016 IEEE Neural Networks Pioneer Award. He is a Fellow of IEEE and IFAC. He served as the President of the IEEE Computational Intelligence Society (2012-2013), as the President of the European Control Association (2017-2019), and as the Editor-in-Chief of the IEEE Transactions on Neural Networks and Learning Systems (2004-2010). Prof. Polycarpou currently serves on the Editorial Boards of the Proceedings of the IEEE and the Annual Reviews in Control. His research work has been funded by several agencies and industry in Europe and the United States, including the prestigious European Research Council (ERC) Advanced Grant, the ERC Synergy Grant and the EU-Widening Teaming program.