2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

Autonomous Balloon based Adaptive Sliding Mode Control and Infinite-Horizon POMDP

Van Chung Nguyen[1], An Duy Nguyen[1], Chuong Le[1] Gaurav Srikar[1],
Thanh Nho Do[2], and Hung Manh La [1]

***Abstract*— This paper presents a novel infinite-horizon Partially Observable Markov Decision Process (POMDP) framework with adaptive sliding mode control (ASMC) for autonomous navigation of the balloons. The proposed method integrates an altitude controller designed to account for thermodynamic and real-wind field constraints with an infinite-horizon POMDP for wind-optimal navigation. First, an adaptive sliding mode control is developed to ensure the balloon's internal stability under uncertainties in pressure, external wind fields, and temperature. Subsequently, a reference strategy is formulated using the infinite-horizon POMDP to exploit wind dynamics for station-keeping. The system estimates wind direction in real time and computes actions based on these observations. Experimental results demonstrate the framework's ability to converge on efficient navigation policies while compensating for partial observability of wind dynamics. This approach is particularly suited for aerial or underwater vehicles operating in stratified flow environments, offering a computationally tractable solution for real-world deployment.**

I. Introduction

Recent advances in lighter-than-air vehicles, particularly super-pressure balloons and autonomous blimps, have attracted significant research interest due to their unique advantages for long-duration missions. These systems excel in applications ranging from atmospheric monitoring [1]–[3] to decentralized communication relays [4]–[6], leveraging their low energy consumption and persistent operational capabilities. However, while cost-effective and durable [7], their navigation remains fundamentally constrained by uncontrolled wind dynamics, which often displaces them from target areas, limiting mission efficacy. Addressing this challenge requires a dual approach: an adaptive altitude controller to maintain stability under thermodynamic, wind, and pressure uncertainties, and intelligent navigation to exploit wind patterns for station-keeping. Traditional methods often treat these aspects independently, risking suboptimal performance when atmospheric disturbances interact with navigational decisions. For the balloon, it has only one degree of freedom [8], [9]. The balloon can change its altitude by adjusting the buoyancy forces, thereby moving upward, downward, or maintaining its altitude. For lateral control of the robot, the system relies on drag forces induced by the wind. Therefore, for an autonomous balloon, it is essential to determine the appropriate altitude for navigation by following the wind. This path planning problem has recently been addressed using reinforcement learning [7], [8], [10].

While previous studies have demonstrated successful balloon navigation, several critical limitations remain. For instance, [7] and [10] proposed Q-learning-based reinforcement learning (RL) methods for high-altitude balloon navigation. In [7], the wind field and thermodynamic effects were incorporated into the balloon's motion model, whereas [10] focused solely on the balloon's movement without accounting for external disturbances. However, both approaches suffer from two key drawbacks: (1) their algorithms operate offline and require pre-training with predicted environmental data, and (2) they were validated only in simulations without field testing under real-world, uncontrolled wind dynamics. In [8], the first off-policy reinforcement learning framework was introduced for the balloon to navigate across both indoor and outdoor environments. While their approach demonstrated successful waypoint tracking, it employed simplified system dynamics that neglected the balloon's thermodynamic properties, instead depending on active propulsion for altitude regulation. Furthermore, the study did not examine the system's robustness against realistic environmental disturbances such as wind gusts or thermal variations. In contrast to previous approaches, this work introduces a novel hybrid control framework that combines: (1) an ASMC for disturbance estimation and robustness via Lyapunov stability, and (2) an infinite-horizon POMDP for optimal navigation under uncertainty. The key contributions of this work are:

- We present a novel POMDP-based method capable of real-time balloon navigation. Unlike previous approaches that rely on reinforcement learning, it often requires a trade-off between computational complexity and real-time performance. Our method estimates wind direction in real time and selects the most suitable action for the balloon based on its navigation objective.
- We reformulate the balloon model by incorporating both thermodynamic behavior and real wind-field data, and then propose an adaptive control strategy for altitude tracking. The proposed method is analyzed for stability

This work was funded by the U.S. National Science Foundation (NSF) under grants: NSF-CAREER 1846513 and NSF EPSCoR-HDRFS project, the NASA EPSCoR under grant number 80NSSC24M0141, and the U.S. Army's Engineer Research and Development Center under grant number W911NF-23-1-0186. The views, opinions, findings, and conclusions reflected in this publication are solely those of the authors and do not represent the official policy or position of the NSF and the U.S. Army.

[1]are with the Advanced Robotics and Automation (ARA) Lab, Department of Computer Science and Engineering, University of Nevada, Reno, NV 89557, USA.

[2]is with the UNSW Medical Robotics Lab, Sydney, Australia.

Corresponding author: Hung La, email: hla@unr.edu

Source code: https://github.com/aralab-unr/Autonomous-Balloon

against uncertainties and disturbances, including variations in temperature, pressure, and wind.
- We verify the effectiveness of the proposed method through both simulation and real-world experiments, demonstrating the fully autonomous capabilities of the system in real-life scenarios.
- All the code and sources will be made public for further investigation.

II. Literature review

A. Modelling

Recent studies have shown that the performance of high-altitude balloons depends on various factors, such as thermodynamic effects [11], real-wind fields [12], and initial launch conditions [13]. Over the past decades, several models have been developed to investigate the behavior of high-altitude balloons under different environmental and operational conditions. An early computer program was developed in [14] to estimate calibrated ascent rates, float altitude, and flight behavior based on initial conditions. Later, a thermal model was introduced in [15] to account for the effects of heat transfer, enabling trajectory prediction with consideration of thermal dynamics. The Navajo analysis tool was then developed in [16] to provide high-accuracy predictions for both vertical and horizontal balloon trajectories. A user-friendly software tool called BalloonAscent was later introduced in [17], featuring adaptable formulations for simulating balloon flights in extraterrestrial atmospheres. Additionally, an analysis code incorporating ballasting and valving functions was developed in [18], modeling the drag coefficient as a function of Reynolds number, Froude number, and another dimensionless parameter, which allowed for accurate fitting of flight data. The thermal behavior of balloons at float altitude was also studied by analyzing the influence of film radiation properties and cloud cover, though ascent dynamics were not considered [19]. A comprehensive model was subsequently proposed to evaluate the thermodynamic performance of scientific balloons in [20], including the temperature distribution of the balloon film. Further analysis highlighted the significant influence of initial launch conditions on balloon trajectory and performance [13]. In addition, the effects of real-wind fields on balloon behavior were explored in [12], leading to strategies for maximizing station-keeping time within a targeted area.

B. Propulsion

As mentioned above, to navigate the balloon, it is necessary to adjust the buoyancy forces, thereby enabling it to move upward, downward, or maintain its altitude. To overcome wind fields and gravity, a propulsion system is introduced. This system generates forces that help the balloon reach the desired altitude. In [21], stratosails were used to create lateral forces that assist in steering the balloon. Similarly, tandem tethered balloons were used in [22] to generate lateral forces that support station-keeping. Another propulsion method involves propeller-based systems. In [8], [23], such systems were employed for station-keeping missions. However, they are limited by high power consumption, as maintaining vertical position requires significant energy, making long-term missions infeasible. An alternative approach is the use of electrohydrodynamic thrusters, as proposed in [24]. In this study, the thrusters received power wirelessly via microwave transmission. However, this method is complex, vulnerable to the scarcity of atmospheric particles at higher altitudes, and adds substantial external mass to the balloon, posing additional challenges for long-term missions. Another method is the use of hot-air balloons, which heat the air inside the envelope to create a pressure difference [25]. Compared to gas balloons, they are less efficient, as the density of hot air is lower than that of helium or hydrogen. Additionally, extra energy is required to regulate the air temperature inside the balloon. Moreover, altitude control can be achieved using superpressure balloons equipped with air ballasts. This type of balloon adjusts its altitude by pumping air into or out of a pressurized container. However, for continuous operation, it requires a constant power supply to run the air pump. To support long-term missions, a solar power system was successfully used to power the air pump in [26]. Finally, latex balloons that vent lift gas and deploy preloaded ballasts have been demonstrated as an energy-efficient method for altitude adjustment [27]. The combined use of venting and ballasting in this way significantly reduces the energy demands of the balloon system. The choice of system depends on the user's objectives. In this research, we propose an overall control law for propulsion using two propellers to generate buoyancy forces, designed to be robust against disturbances and uncertainties.

C. Navigation strategy

Navigation in 3D space is made possible by selecting an appropriate altitude. One of the most widely used methods for determining altitude references is reinforcement learning (RL) [7], [8], [10], [26], [28]–[30]. RL can be applied using atmospheric data and balloon flight records collected across various locations and seasons. For example, Project Loon [26] utilized a quantile regression deep Q-network (DQN) to model wind layers at different altitudes for positioning, enabling balloons to perform station-keeping at around 20 km altitude while providing internet connectivity. In [31], wind and atmospheric uncertainties were modeled as state transition probabilities. This allowed the RL problem to be framed as a sequential decision-making task, where the agent learns optimal policies by sampling experience. Among RL methods, Q-learning is particularly popular for control tasks due to its model-free nature. It has been effectively applied in several works [7], [8], [10], allowing agents to adapt actions based on environment states. Alternatively, a geometric method was proposed in [32] to determine the optimal altitude by maximizing the cosine similarity between the balloon's trajectory and wind orientation. In balloon navigation, modeling atmospheric dynamics and wind patterns is essential [33], [34]. These models provide the foundation for training learning algorithms or making real-time decisions

to select the most suitable altitude for navigation or station-keeping. In this research, instead of proposing a pre-trained model, we introduce a method that allows the balloon to learn the wind conditions in real time and navigate effectively.

III. Balloon Model

For the balloon model, we assume that the balloon is a sphere with radius R, and the volume of the balloon can be expressed as:

$$V = \frac{m_{\text{He}} R_{m_{\text{He}}} T_{\text{He}}}{P_{\text{He}}} \tag{1}$$

where $R_{m_{\text{He}}} = 2077$, $T_{\text{He}}, P_{\text{He}}$ are the specific gas constant, the temperature, and the pressure of helium. The total buoyant force and drag force produced by the balloon are given by [19]:

$$B = \rho_{atm} g V \tag{2}$$

$$D = 0.5 \rho_{atm} v_r^2 C_d S \tag{3}$$

where C_d is the drag coefficient with a value of 0.8; $\rho_{atm} = \frac{P_{atm}}{R_{atm} T_{atm}}$ is the air density, where P_{atm}, R_{atm}, and T_{atm} are the atmospheric pressure, specific gas constant, and temperature, respectively; $S = 2\pi R^2$ is the reference area; and v_r is the relative wind velocity with respect to the movement of the balloon.

$$v_r = (v_{rx}^2 + v_{ry}^2 + v_{rh}^2)^{0.5} \tag{4}$$

where v_{rx}, v_{ry} and v_{rh} are the velocity components according to the coordinate system as in [19]:

$$\begin{cases} v_{rx} = v_{wx} - v_x \\ v_{ry} = v_{wy} - v_y \\ v_{rh} = v_{wh} - v_h \end{cases} \tag{5}$$

where v_{wx}, v_{wy} and v_{wh} are the absolute velocity components of wind, v_x, v_y and v_h are the absolute velocity components of the balloon. Then, the drag force components can be expressed by:

$$\begin{cases} D_x = D v_{rx}/v_r \\ D_y = D v_{ry}/v_r \\ D_h = D v_{rh}/v_r \end{cases} \tag{6}$$

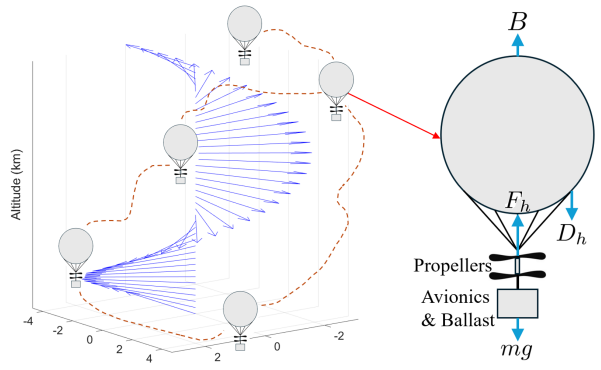

Fig. 1. Balloon diagram.

The differential equation governing the horizontal and vertical motion of the balloon is given by:

$$\begin{cases} \dot{x} = v_x;\ \dot{y} = v_y;\ \dot{h} = v_h \\ \dot{v}_x = D_x/(m + m_{add}) \\ \dot{v}_y = D_y/(m + m_{add}) \\ \dot{v}_h = (F_h + B + D_h - mg)/(m + m_{add}) \end{cases} \tag{7}$$

where x, y and h are the coordinate positions of the balloon, m is the gross mass of the balloon and m_{add} is the added mass.

$$m_{add} = C_V \rho_{atm} V \tag{8}$$

where the added mass coefficient $C_V = 0.5$. For the balloon, we can only control its altitude by changing the lift forces. In this research, we consider the robust altitude control for the balloon as proposed as:

$$\begin{cases} \dot{h} = v_h \\ \dot{v}_h = \frac{F_h + \frac{m_{\text{He}} R_{m_{\text{He}}} T_{\text{He}} g}{R_{atm} T_{atm}} + D_h - mg}{m + m_{add}} = \frac{F_h + D_h + C_b}{m + m_{add}} \end{cases} \tag{9}$$

Here, $C_b = \frac{m_{\text{He}} R_{m_{\text{He}}} T_{\text{He}} g}{R_{\text{atm}} T_{\text{atm}}} - mg$ is a shorthand term representing the internal parameters of the balloon. For the balloon model, we further consider thermodynamics and the effects of real wind fields as in [12], [13]. Consequently, the proposed model can be applied to all types of balloons. For the adaptive sliding mode control, we consider the sliding surface as:

$$s = e_2 + \lambda e_1 \tag{10}$$

where λ is positive number, $e_1 = h - h_r$, $e_2 = \dot{h} - \dot{h}_r$ are the control error, with h_r is the desired altitude.

Theorem 3.1: Considering the balloon with the state space (9), if the ASMC controller is proposed as:

$$\dot{\hat{D}}_h = \gamma s \tag{11}$$

$$F_h = (-K_s \text{sat}(s) - Ks + \ddot{h}_r - \lambda(\dot{h} - \dot{h}_r))(m + m_{add}) - \hat{D}_h - C_b \tag{12}$$

With $\hat{D}_h$ being the estimated wind disturbance, K_s and K as positive constants, and sat(s) denoting the saturation function of s, the sliding surface and the wind disturbance estimation error are asymptotically stable.

Proof: According to Lyapunov's stability, considering the Lyapunov candidate as:

$$V_b = \frac{1}{2} s^2 + \frac{1}{2\gamma(m + m_{add})} \tilde{D}_h^2. \tag{13}$$

with $\tilde{D}_h = D_h - \hat{D}_h$ is the wind disturbances estimation error. Differentiating V_b with respect to time t, the derivative of V_b yield is expressed as:

$$\dot{V}_b = s\dot{s} + \frac{1}{\gamma(m + m_{add})} \tilde{D}_h \dot{\tilde{D}}_h. \tag{14}$$

Substituting the control law and adaptive law in (11) into the state-space model (9), and noting the assumption that

the rate of change of the wind disturbance is dominated by the control rate, we have:

$$\dot{\tilde{D}}_h = \dot{D}_h - \dot{\hat{D}}_h \approx -\dot{\hat{D}}_h = -\gamma s, \tag{15}$$

then the derivative becomes:

$$\begin{aligned}\dot{V}_b &= s\dot{s} + \frac{1}{\gamma}\tilde{D}_h\dot{\tilde{D}}_h = s\dot{s} - \frac{1}{m+m_{\text{add}}}(D_h - \hat{D}_h)s \\ &= s\left(\ddot{h} - \ddot{h}_r + \lambda(\dot{h} - \dot{h}_r) - \frac{1}{m+m_{\text{add}}}(D_h - \hat{D}_h)\right) \\ &= s\left(\frac{F_h + \hat{D}_h + C_b}{m+m_{\text{add}}} - \ddot{h}_r + \lambda(\dot{h} - \dot{h}_r)\right) \\ &= s(-K_s\text{sat}(s) - Ks) = -K_s\text{sat}(s)s - Ks^2 \end{aligned} \tag{16}$$

Hence, the sliding surface is stable. Moreover, applying Barbalat's lemma [35], the sliding surface and the wind disturbance estimation error are asymptotically stable. ■

IV. Infinite-Horizon POMDP for Navigation

In this section, we present the formulation of an infinite-horizon POMDP for balloon navigation in the presence of unknown and time-varying wind disturbances. The balloon operates in an environment with uncertain wind conditions, and its task is to adjust its altitude to achieve station-keeping. Each altitude level corresponds to a different wind direction, which the agent must estimate and exploit to remain near a desired location.

To perform station-keeping, the balloon follows a three-step strategy:

I) Observe its current position and velocity;
II) Estimate the probability distribution over wind directions at different altitude levels;
III) Select an action based on its current belief and position to maintain its location. This is implemented by controlling the altitude through adjusting the amount of helium inside the balloon.

We model this problem as an infinite-horizon POMDP, solved via online planning. The infinite-horizon POMDP is defined by the tuple (S, A, T, R, Ω, O), where S is the set of states, A is the set of actions, T is the state transition model, $R : S \times A \to \mathbb{R}$ is the reward function, O is the set of observations, and Ω is the observation model. At each timestep, the agent selects an action $a \in A$, causing the environment to transition from state s to s' according to $T(s', a, s) = \Pr(s' \mid s, a)$. The agent receives an observation $o \in O$ through $\Omega(o, s', a) = \Pr(o \mid s', a)$ depending on the sensor model, along with a reward $r = R(s, a)$.

Since the true state is not directly observable, the agent maintains a belief state $b_t : S \to [0, 1]$, representing a probability distribution over possible states. After executing action a and receiving observation o, then belief is updated according to:

$$b'(s') = \eta \Pr(o \mid s', a) \sum_{s \in S} \Pr(s' \mid s, a) b(s),$$

where η is a normalizing constant to ensure the belief sums to one. Our objective is to find an action based on policy π that maps belief states to actions in a way that maximizes the average reward:

$$a^\pi = \lim_{T\to\infty} \frac{1}{T}\mathbb{E}^\pi \left[\sum_{t=0}^{T-1} R(s_t, a_t)\right],$$

where the expectation is taken over the trajectories generated by following policy π, this average-reward formulation, which does not employ a discount factor, is particularly suitable for persistent, ongoing tasks such as balloon station-keeping.

A. Definition

For our problem, the infinite-horizon POMDP is defined by the tuple (S, A, T, R, Ω, O), where the state space is given by $S = (W_b, z_b)$, with W_b representing the wind direction and z_b denoting the altitude level. Since wind direction varies across different altitude layers, the agent must estimate the wind direction at each level to effectively perform station-keeping and remain near a desired location. The action space A consists of three discrete actions: ascend, descend, or maintain altitude. Each action changes the balloon's altitude by transitioning it to a corresponding level within the predefined altitude range.

The transition model T defines how the agent's state evolves over time. At each time step t, the agent takes an action $a \in A$, causing a transition from state s to a new state s'. The state space is discretized into a grid, where each cell center corresponds to a unique state. After the transition, the agent receives an observation $o \in O$ from the environment, providing information that helps update its belief about the current state. For the observation model, we observe the balloon's position and velocity, represented as $O = (p_b, v_b)$, where p_b and v_b denote the position and velocity, respectively. These observations are obtained from onboard sensors such as GPS or optical flow, which provide real-time measurements of the balloon's motion.

Finally, the reward function R is defined as $R(s, o, a) = -d(o' \mid (s, a), p_0)$, where o' is the resulting position after taking action a from state s, and p_0 is the target location. The function penalizes the distance between the balloon's new position o' and the desired location p_0, encouraging the agent to keep the balloon near the goal for effective station-keeping.

For the policy π used to select actions, we adopt an ϵ-greedy strategy. Actions are primarily chosen based on the one that maximizes the expected reward, while with probability ϵ, a random action is selected to encourage exploration. The exploration rate ϵ decays over time to gradually favor exploitation over exploration. Additionally, when the confidence in the belief exceeds a defined threshold-indicating a possible change in the direction of the wind, we reset ϵ to reintroduce exploration and allow the policy to adapt to the changing environment. The entire Infinite-horizon POMDP algorithm for navigating the balloon is shown from Algorithm 1 to Algorithm 3.

Algorithm 1 Infinite-horizon POMDP navigation

Input: Current position (x, y, h), velocities (v_x, v_y).
Output: Desired altitude h_{des}.
Initialize belief B_h, exploration rate ϵ, decay factor ϵ_{decay}, the set of actions, observations, area, and goal: $\mathcal{A}_{\text{actions}}, \mathcal{W}_{\text{winds}}, \mathcal{S}_{\text{area}}, \mathbf{g}$; threshold for belief rate: ε; Belief learning rate α.
for episode $= 1$ to ∞ **do**
    Measure current state s_t
    if The robot reaches the desired h_{des} for a longer time Δt **then**
        $\mathbf{B_h} \leftarrow \text{UPDATEWINDBELIEF}(\mathbf{B_h}, v_x, v_y, h)$
        if $\mathbf{B_h}$ rate $\geq \varepsilon$ **then**
            $\epsilon \leftarrow$ initial ϵ
        else
            $\epsilon \leftarrow \max(\epsilon \cdot \epsilon_{\text{decay}}, 0.001)$
        end if
        $a^* \leftarrow \text{SELECTOPTIMALACTION}(s_t, \varepsilon, \mathbf{B_h}, \mathbf{g})$
        $h_{\text{des}} \leftarrow \text{Clamp}(h_{\text{des}} + \Delta a^*, \mathcal{S}^h_{\text{area}})$
    else
        Maintain current altitude strategy
    end if
end for

Algorithm 2 Wind Belief Update

function UPDATEWINDBELIEF($\mathbf{B_h}, v_x, v_y, h$)
    $\mathbf{w}_{\text{obs}} \leftarrow (\text{sgn}(v_x), \text{sgn}(v_y))$
    for $i \leftarrow 1$ **to** $|\mathcal{W}|$ **do**
        if $\mathcal{W}_i = \mathbf{w}_{\text{obs}}$ **then**
            $B_{h,i} \leftarrow (1-\alpha)B_{z,i} + \alpha$
        else
            $B_{h,i} \leftarrow (1-\alpha)B_{z,i}$
        end if
    end for
    Normalize $\mathbf{B}[h, :]$ to probability distribution
    return $\mathbf{B_h}$
end function

Algorithm 3 Action Selection Policy

function SELECTOPTIMALACTION($s_t, \varepsilon, \mathbf{B_h}, \mathbf{g}$)
    if $\xi < \varepsilon$ **then** ▷ $\xi \sim \mathcal{U}(0, 1)$
        return random exploration action
    else
        for $h \in \mathcal{S}^h_{\text{area}}$ **do**
            $\mathbf{w}_h \leftarrow \mathcal{W}_{\arg\max \mathbf{B}[h,:]}$
            $\mathbf{s}' \leftarrow \mathbf{s} + \mathbf{w}_h$
            $d_{xy} \leftarrow \|\mathbf{s}'_{xy} - \mathbf{g}\|$
            Maintain altitude proximity cost
        end for
        return a that minimizes d_{xy}
    end if
end function

V. Results

In this section, we present results conducted in both Simulink/MATLAB and Gazebo/ROS 2 using the parameters of a real balloon to evaluate the effectiveness of the proposed method. More details about the available code can be found at `https://github.com/aralab-unr/Autonomous-Balloon`. The robot and environment parameters are listed in Table I. The effectiveness of the proposed method is demonstrated in two scenarios. First, to evaluate the stability of the controller, the balloon is required to maintain a specific altitude. Second, to test its autonomous capability, the balloon is directed to reach a desired target and then perform station-keeping. In this scenario, the balloon is guided to a target area and remains there by tracking the desired altitude provided by the infinite-horizon POMDP.

TABLE I
Robot, Environment, and Control Parameters

Balloon parameters	Environment & control parameters
$m = 0.075(kg)$	$\lambda = 0.125, \gamma = 10.75, K_s = 0.1, K = 1.5$
$m_{He} = 0.00975(kg)$	$T_{atm} = 288 - 0.0065h, R_{atm} = 287$
$R_0 = 0.315(m)$	$P_{atm} = 101325(\frac{288-0.0065h}{288})^{5.25}$

A. Simulation results

1) Altitude tracking: For the altitude-holding task, the desired altitude is set to 10m. The balloon is launched from the origin and ascends toward the target altitude. The results are shown in Figure 2. However, due to the paper's length limitations, more details about the simulation results in Gazebo/ROS2 are provided in the attached link. In this section, we focus on the performance of the balloon in the Simulink/MATLAB simulation. To further assess our method, we compare it with other approaches by evaluating control performance using the following benchmarks: Settling time 5% (s), Overshoot (%), Integral Absolute Error (IAE), Integral Square Error (ISE), Integral of Time-weighted Absolute Error (ITAE), and Integral of Time-weighted Square Error (ITSE). The results, shown in Table II, indicate that our method outperforms all others across all metrics. Regarding estimation, the adaptive term demonstrates excellent performance in estimating disturbances under wind effects, as illustrated in Figure 2.

TABLE II
The Control Performance Benchmark

Methods	ST (s)	Overshoot	IAE	ISE	ITAE	ITSE
PID	-	29.7	119.4	321.4	4777	6591
SMC	25.4	4.69	92.8	462.3	1116	1945
H_∞	-	130.7	759.7	7152	43300	40380
ASMC	**23.9**	**0.48**	**79.9**	**400.8**	**638.4**	**1591**

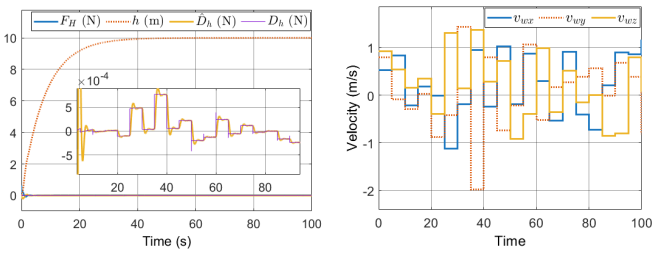

Fig. 2. Tracking performance, control law & wind velocity.

2) Station Keeping Performance: For station-keeping, we used wind data from the Global Wind Atlas [36] at latitude 39.53794° and longitude -119.761963° (Reno, Nevada). We observed wind velocities at altitudes ranging from 0 to 250 m

and used the maximum observed wind velocity as the input for the model. The wind velocity is chosen as the maximum of the mean value. The altitude is divided into five levels, and at each level, the wind direction is randomly selected from northeast (NE), southeast (SE), southwest (SW), and northwest (NW).

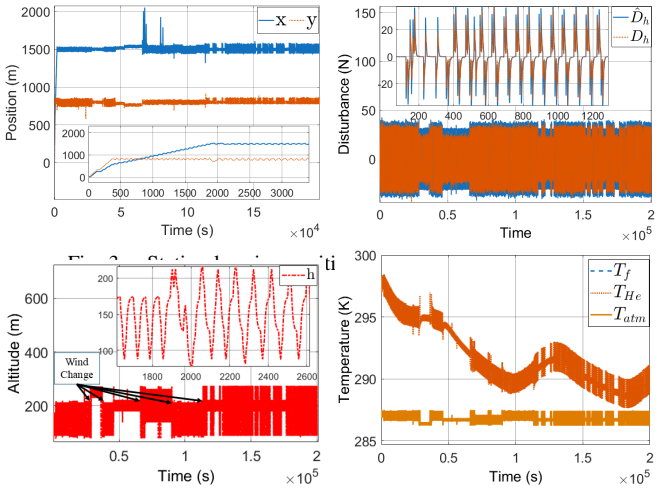

Fig. 4. Altitude & Environment temperature.

In order to maintain its position, the balloon must estimate the wind direction and choose the most suitable altitude accordingly. To further evaluate the disturbance-handling capability of the infinite-time POMDP, the wind direction is changed every 3000 interactions. The balloon is then operated for 2.35 days to demonstrate its robustness against disturbances originating from thermodynamic effects and the real wind field. The setpoint of the balloon is defined as $[x_r, y_r] = [15, 8]$ km. The results, shown from Figure 3 to Figure 4, indicate that the balloon successfully maintains its position and adapts well to environmental changes. The adaptive term also demonstrates excellent performance in estimating disturbances.

B. Experiment results

1) Experiment setup: The setup of the indoor prototype is shown in Figure 5. The balloon is driven by two DC servo motors. The control board is handled by an ESP8266 module, which sends data to the station via the TCP protocol.

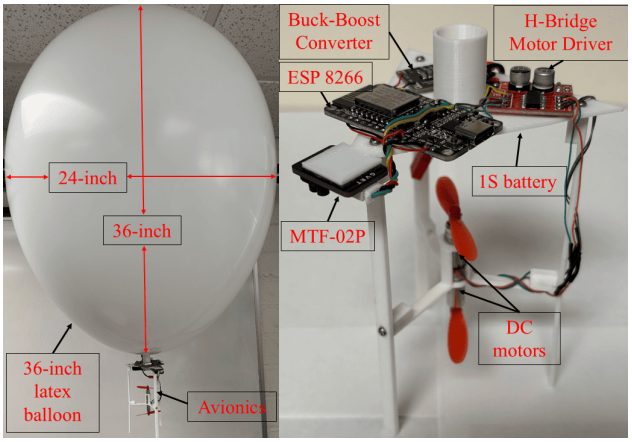

Fig. 5. The setup of Balloon's prototype.

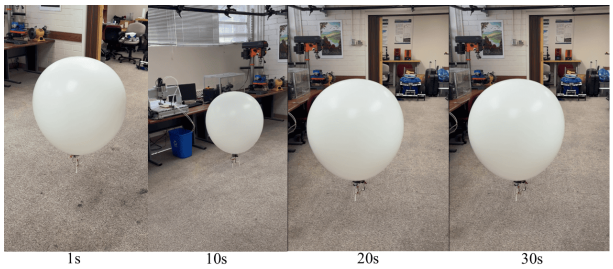

Fig. 6. Visual demonstration of balloon flying indoors.

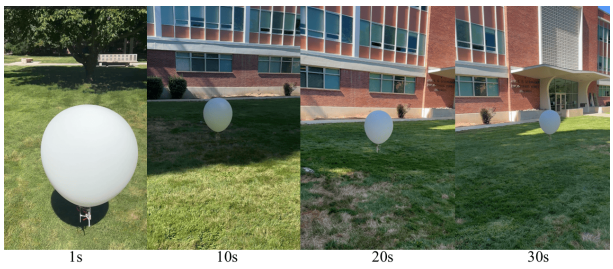

Fig. 7. Visual demonstration of balloon flying outdoors.

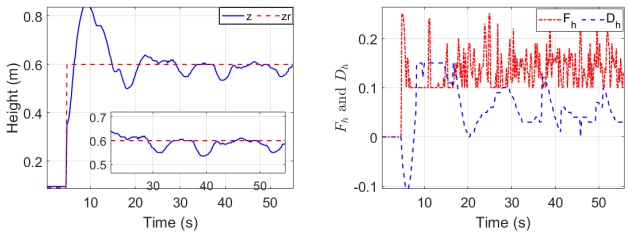

Fig. 8. Height, Control Law, and Estimated Disturbance (Indoor).

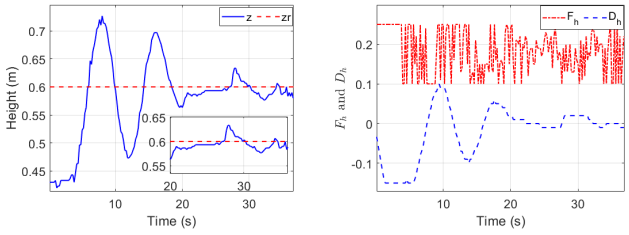

Fig. 9. Height, Control Law, and Estimated Disturbance (Outdoor).

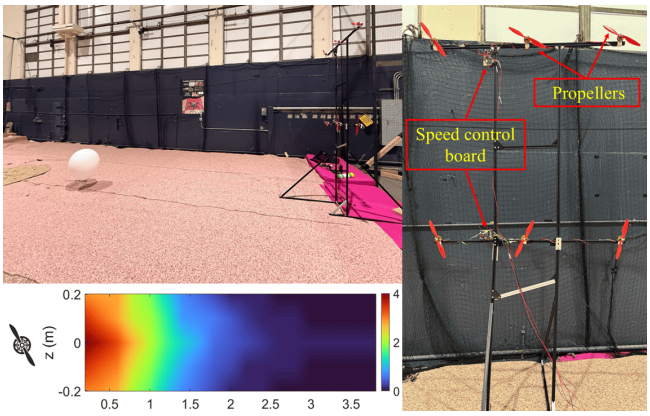

Fig. 10. Indoor environment setup with the custom wind generator, the wind velocity is measured by an anemometer and bilinear interpolation.

The balloon's position data for the indoor prototype is collected using the MTF-02P Optical & Range Sensor. All components are lifted by a 36-inch latex balloon. The total weight of the balloon is 75 grams. This is necessary to ensure that the balloon can float while carrying all the required components.

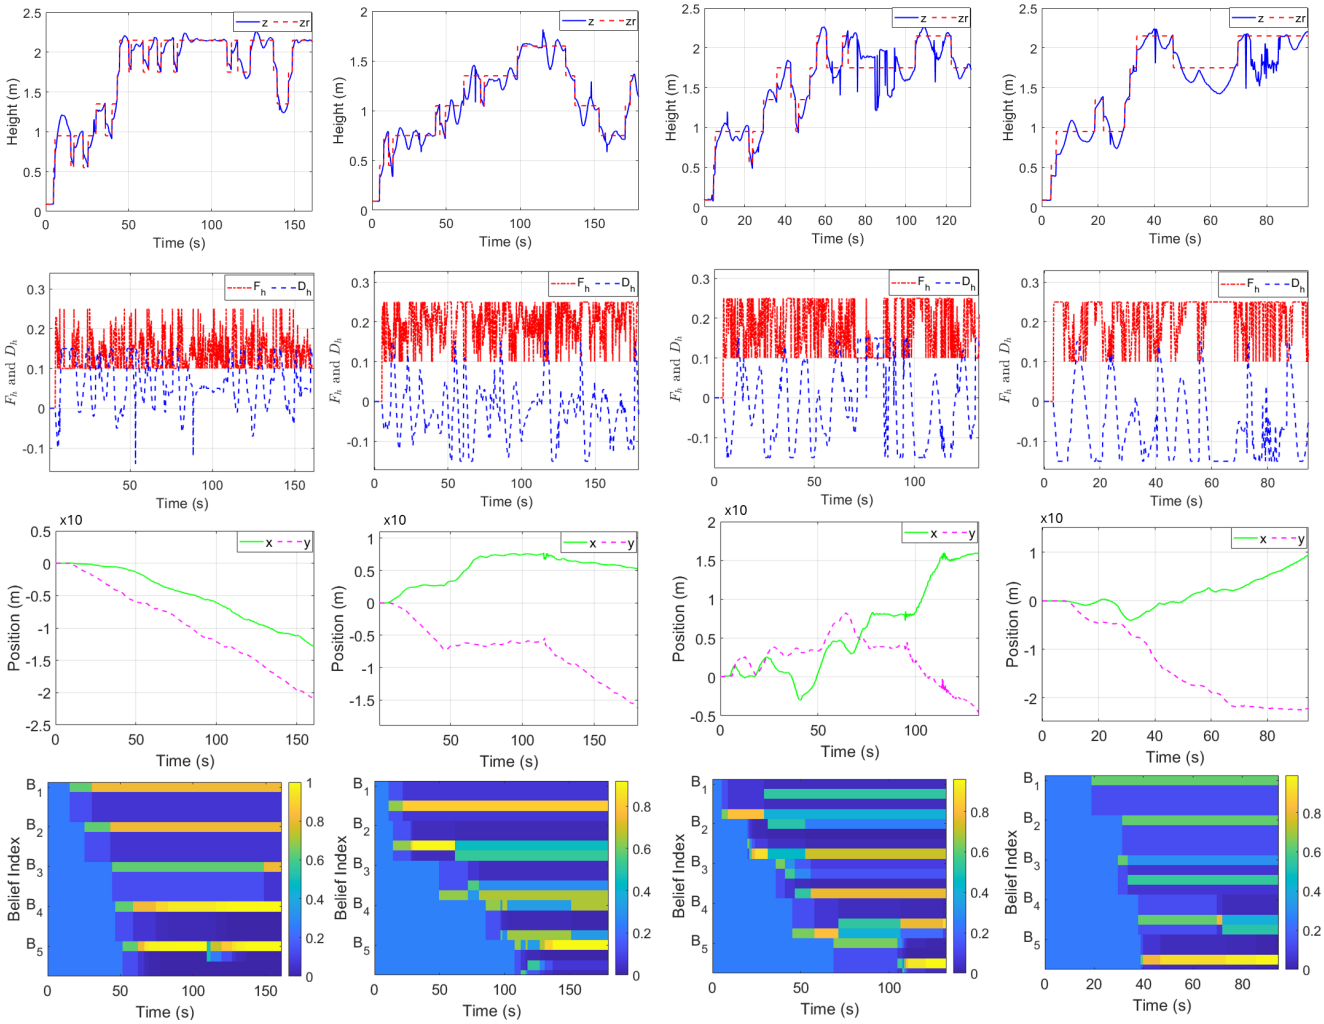

Fig. 11. Experimental results on station keeping for the balloon. The first two rows show data from indoor flights, while the last two rows correspond to outdoor flights. The results demonstrate the tracking performance in terms of height, belief values, control input, and disturbance estimation.

2) Altitude tracking results: For the altitude tracking results, we evaluated the balloon's ability to follow a specified desired altitude. The target altitude was set to 0.6 m. We test the tracking ability of the balloon in both indoor and outdoor environments. The results are shown in Figure 6 - Figure 9, the results indicate that the balloon can effectively track the desired height.

3) Station keeping results: For the station-keeping experiments, we used a custom wind generator capable of providing wind speeds up to 4 m/s and conducted the tests in a 4.5×9 m indoor space. The details of the experimental setup are shown in Figure 10. The altitude range was divided into five discrete levels, each corresponding to one belief index, with four possible wind types. The results, shown in Figures 11, demonstrate that the belief index at each altitude level converges to the true generated wind, highlighting the effectiveness of the proposed method. Moreover, although the balloon exhibits some position drift over time, it is still able to maintain station. This drift occurs because the optical flow sensor operates in the body frame and does not account for rotational motion. As a result, even when the balloon remains stationary in space, its measured position can drift. However, the wind estimator converges regardless of the balloon's orientation, further validating the robustness of the approach.

VI. Conclusion

This paper introduces a robust and computationally efficient infinite-horizon POMDP framework for the autonomous navigation of balloons in stratified flow environments. By integrating an adaptive altitude controller with a wind-adaptive, POMDP-based decision-making strategy, the proposed method ensures both system stability and long-term navigation performance under uncertainty. The adaptive controller effectively mitigates environmental disturbances such as pressure fluctuations, wind variability, and thermal constraints, while the infinite-horizon POMDP enables the balloon to exploit favorable wind currents through real-time estimation and adaptive planning. Experimental results confirm the framework's ability to generate effective navigation policies despite partial observability of wind fields. This work presents a prototype that can scale to high-endurance aerial deployments, demonstrating that robust control and

decision-making under uncertainty can be unified in a practical solution. For future work, we plan to incorporate wind profile estimation and visual inspection capabilities by equipping the balloon with a camera. We will also conduct outdoor experiments to support weather forecasting applications.

REFERENCES

[1] K. Adams, A. Broad, D. Ruiz-García, and A. R. Davis, "Continuous wildlife monitoring using blimps as an aerial platform: a case study observing marine megafauna," *Australian Zoologist*, vol. 40, no. 3, pp. 407–415, 2020.

[2] D. Vignelles, T. Roberts, E. Carboni, E. Ilyinskaya, M. Pfeffer, P. D. Waldhauserova, A. Schmidt, G. Berthet, F. Jegou, J.-B. Renard *et al.*, "Balloon-borne measurement of the aerosol size distribution from an icelandic flood basalt eruption," *Earth and Planetary Science Letters*, vol. 453, pp. 252–259, 2016.

[3] M. Tironi and M. Valderrama, "The militarization of the urban sky in santiago de chile: the vision multiple of a video-surveillance system of aerostatic balloons," *Urban Geography*, vol. 42, no. 2, pp. 161–180, 2021.

[4] A. Doi, Y. Kono, K. Kimura, S. Nakahara, T. Oyama, N. Okada, Y. Satou, K. Yamashita, N. Matsumoto, M. Baba *et al.*, "A balloon-borne very long baseline interferometry experiment in the stratosphere: Systems design and developments," *Advances in Space Research*, vol. 63, no. 1, pp. 779–793, 2019.

[5] I. Vandermeulen, M. Guay, and P. J. McLellan, "Distributed control of high-altitude balloon formation by extremum-seeking control," *IEEE Transactions on Control Systems Technology*, vol. 26, no. 3, pp. 857–873, 2017.

[6] S. Alsamhi, M. S. Ansari, L. Zhao, S. N. Van, S. Gupta, A. A. Alammari, A. H. Saber, M. Y. Hebah, M. A. A. Alasali, H. M. Aljabali *et al.*, "Tethered balloon technology for green communication in smart cities and healthy environment," in *2019 First International Conference of Intelligent Computing and Engineering (ICOICE)*. IEEE, 2019, pp. 1–7.

[7] Z. Xu, Y. Liu, H. Du, and M. Lv, "Station-keeping for high-altitude balloon with reinforcement learning," *Advances in Space Research*, vol. 70, no. 3, pp. 733–751, 2022.

[8] S. L. Jeger, N. Lawrance, F. Achermann, O. Pang, M. Kovac, and R. Y. Siegwart, "Reinforcement learning for outdoor balloon navigation: A successful controller for an autonomous balloon," *IEEE Robotics & Automation Magazine*, vol. 31, no. 2, pp. 26–38, 2024.

[9] T. Harry, M. Guay, and S. Wang, "Autonomous navigation and station-keeping of high-altitude balloon using extremum seeking control," in *2025 IEEE Aerospace Conference*. IEEE, 2025, pp. 1–7.

[10] J. Saunders, L. Prenevost, Ö. Şimşek, A. Hunter, and W. Li, "Resource-constrained station-keeping for latex balloons using reinforcement learning," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 1102–1109.

[11] H. Franco and H. Cathey Jr, "Thermal performance modeling of nasa's scientific balloons," *Advances in Space Research*, vol. 33, no. 10, pp. 1717–1721, 2004.

[12] H. Du, J. Li, W. Zhu, Z. Qu, L. Zhang, and M. Lv, "Flight performance simulation and station-keeping endurance analysis for stratospheric super-pressure balloon in real wind field," *Aerospace Science and Technology*, vol. 86, pp. 1–10, 2019.

[13] Y. Zhang and D. Liu, "Influences of initial launch conditions on flight performance of high altitude balloon ascending process," *Advances in Space Research*, vol. 56, no. 4, pp. 605–618, 2015.

[14] F. Kreith and J. F. Kreider, *Numerical prediction of the performance of high altitude balloons*. Atmospheric Technology Division, National Center for Atmospheric Research, 1974, no. 65.

[15] L. A. Carlson and W. J. Horn, "New thermal and trajectory model for high-altitude balloons," *Journal of Aircraft*, vol. 20, no. 6, pp. 500–507, 1983.

[16] A. Pankine, R. S. Schlaifer, and M. Heun, "Advanced balloon performance simulation and analysis tool," in *AIAA's 3rd Annual Aviation Technology, Integration, and Operations (ATIO) Forum*, 2003, p. 6741.

[17] R. Farley, "Balloonascent: 3-d simulation tool for the ascent and float of high-altitude balloons," in *AIAA 5th ATIO and16th lighter-than-air sys tech. and balloon systems conferences*, 2005, p. 7412.

[18] G. Morani, R. Palumbo, G. Cuciniello, F. Corraro, and M. Russo, "Method for prediction and optimization of a stratospheric balloon ascent trajectory," *Journal of Spacecraft and Rockets*, vol. 46, no. 1, pp. 126–133, 2009.

[19] Q. Dai, X. Fang, X. Li, and L. Tian, "Performance simulation of high altitude scientific balloons," *Advances in Space Research*, vol. 49, no. 6, pp. 1045–1052, 2012.

[20] Q. Liu, Z. Wu, M. Zhu, and W. Xu, "A comprehensive numerical model investigating the thermal-dynamic performance of scientific balloon," *Advances in Space Research*, vol. 53, no. 2, pp. 325–338, 2014.

[21] S. S. Ramesh, J. Ma, K.-M. Lim, H. P. Lee, and B. C. Khoo, "Numerical evaluation of station-keeping strategies for stratospheric balloons," *Aerospace Science and Technology*, vol. 80, pp. 288–300, 2018.

[22] D. Zhang, H. Luo, Y. Cui, X. Zeng, and S. Wang, "Tandem, long-duration, ultra-high-altitude tethered balloon and its system characteristics," *Advances in Space Research*, vol. 66, no. 10, pp. 2446–2465, 2020.

[23] J. Wang, X. Meng, and C. Li, "Recovery trajectory optimization of the solar-powered stratospheric airship for the station-keeping mission," *Acta Astronautica*, vol. 178, pp. 159–177, 2021.

[24] E. van Wynsberghe and A. Turak, "Station-keeping of a high-altitude balloon with electric propulsion and wireless power transmission: A concept study," *Acta Astronautica*, vol. 128, pp. 616–627, 2016.

[25] I. Edmonds, "Hot air balloon engine," *Renewable Energy*, vol. 34, no. 4, pp. 1100–1105, 2009.

[26] M. G. Bellemare, S. Candido, P. S. Castro, J. Gong, M. C. Machado, S. Moitra, S. S. Ponda, and Z. Wang, "Autonomous navigation of stratospheric balloons using reinforcement learning," *Nature*, vol. 588, no. 7836, pp. 77–82, 2020.

[27] A. Sushko, A. Tedjarati, J. Creus-Costa, S. Maldonado, K. Marshland, and M. Pavone, "Low cost, high endurance, altitude-controlled latex balloon for near-space research (valbal)," in *2017 IEEE Aerospace Conference*. IEEE, 2017, pp. 1–9.

[28] T. K. Schuler, C. Prasad, G. Kiselev, and D. Sofge, "Seasonal station-keeping of short duration high altitude balloons using deep reinforcement learning," *arXiv preprint arXiv:2502.05014*, 2025.

[29] J. Saunders, S. Saeedi, A. Hartshorne, B. Xu, Ö. Şimşek, A. Hunter, and W. Li, "Identifying optimal launch sites of high-altitude latex-balloons using bayesian optimisation for the task of station-keeping," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 13 128–13 135.

[30] J. Liu, X. Li, Y. Hu, and B. Zhang, "High-altitude balloons no-fly zone avoidance based reinforcement learning," in *2024 IEEE International Conference on Real-time Computing and Robotics (RCAR)*. IEEE, 2024, pp. 490–495.

[31] N. Fathpour, L. Blackmore, Y. Kuwata, C. Assad, M. T. Wolf, C. Newman, A. Elfes, and K. Reh, "Feasibility studies on guidance and global path planning for wind-assisted montgolfière in titan," *IEEE Systems Journal*, vol. 8, no. 4, pp. 1112–1125, 2013.

[32] H. Du, M. Lv, J. Li, W. Zhu, L. Zhang, and Y. Wu, "Station-keeping performance analysis for high altitude balloon with altitude control system," *Aerospace Science and Technology*, vol. 92, pp. 644–652, 2019.

[33] T. Fields, M. Heninger, J. LaCombe, and E. Wang, "In-flight landing location predictions using ascent wind data for high altitude balloons," in *AIAA balloon systems (BAL) conference*, 2013, p. 1294.

[34] A. Sóbester, H. Czerski, N. Zapponi, and I. Castro, "High-altitude gas balloon trajectory prediction: A monte carlo model," *AIAA journal*, vol. 52, no. 4, pp. 832–842, 2014.

[35] J.-J. E. Slotine, W. Li *et al.*, *Applied nonlinear control*. Prentice hall Englewood Cliffs, NJ, 1991, vol. 199, no. 1.

[36] Global Wind Atlas, "Global wind atlas: Energydata.info," 2025. [Online]. Available: https://globalwindatlas.info/en/