

High-Altitude Balloon Station-Keeping with First Order Model Predictive Control

Myles Pasetsky^{1*}, Jiawei Lin^{2*}, Bradley Guo^{1*}, and Sarah Dean¹

Abstract—High-altitude balloons (HABs) are common in scientific research due to their wide range of applications and low cost. Because of their nonlinear, underactuated dynamics and the partial observability of wind fields, prior work has largely relied on model-free reinforcement learning (RL) methods to design near-optimal control schemes for station-keeping. These methods often compare only against hand-crafted heuristics, dismissing model-based approaches as impractical given the system complexity and uncertain wind forecasts. We revisit this assumption about the efficacy of model-based control for station-keeping by developing *First-Order Model Predictive Control (FOMPC)*. By implementing the wind and balloon dynamics as differentiable functions in JAX, we enable gradient-based trajectory optimization for online planning. FOMPC outperforms a state-of-the-art RL policy, achieving a 24% improvement in time-within-radius (TWR) without requiring offline training, though at the cost of greater online computation per control step. Through systematic ablations of modeling assumptions and control factors, we show that online planning is effective across many configurations, including under simplified wind and dynamics models.

I. INTRODUCTION

High-altitude balloons (HAB) provide a low-cost and low-risk profile tool [1] for atmospheric and climate research [2], [3], as well as astronomical and space science observations [4], [5], among others. As horizontal control remains costly and resource-intensive, most of these balloons rely on active altitude control to take advantage of varying wind currents at different altitudes that guide the balloons laterally. However, accurate measurements of these wind currents are only available locally, making the setting partially observable. In this underactuated and partially observable setting, control effectiveness depends both on the accuracy of wind predictions and the reliability of altitude-based steering. These requirements are further complicated by internal resource limitations, which can restrict altitude control, and by external environmental factors, which may reduce the availability of diverse and favorable wind layers.

Early work on developing controllers for autonomous HABs relied on hand-designed heuristics [6], [7], which explicitly exploit the local wind column via altitude control to achieve short-term station-keeping. To approach the policy from an optimal control perspective, recent works apply model-free reinforcement learning (RL) methods to the balloon control problem [8], [9], [10], [11], [12], [13], [14], [15]. These approaches train policies in simulation to either navigate from point to point in a wind field or perform

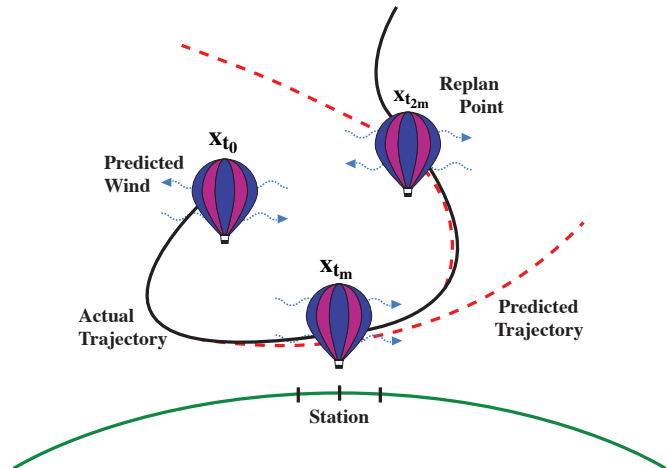


Fig. 1: Station-keeping via receding-horizon altitude control. Predicted winds generate a planned trajectory (red), which diverges from the actual trajectory (black); replanning at every m steps corrects forecast errors.

station-keeping, where the balloon remains in proximity to a ground target. Among these works, the RL agent has access only to the local predicted wind profile (typically a vertical column) at its current horizontal position and learns an approximate optimal policy through offline training in simulation.

Despite successful RL deployments in both simulation and real-world, their main performance gains are demonstrated only with respect to simple heuristic-based policies. These works are often motivated by the purported difficulty of applying model-based tools to the station-keeping problem [14], and the need for costly replanning under uncertain wind predictions [8]. Moreover, strong baselines are critical for validating empirical results in deep RL, where reported gains can be undermined by statistical uncertainty with few evaluation runs [16] and high training variance across seeds [17]. Without strong baselines, deep RL results remain difficult to interpret, as costly training, variance across seeds, and generalization failures can exaggerate apparent progress. Hence, there is a critical gap in the literature: a systematic model-based baseline for direct comparison with RL methods.

Motivated by this absence of model-based baselines, we develop First-Order Model Predictive Control (FOMPC) as the first rigorous approach to fill this gap. By first implementing balloon dynamics and wind models as differentiable

¹Cornell University {mhp58, bzg4, sdean}@cornell.edu

²University of California San Diego jal214@ucsd.edu

*Equal Contribution

functions in JAX [18], then formulating the HAB station-keeping task as an unconstrained optimization problem, we can use first-order gradient methods to optimize balloon actions. This approach explicitly incorporates both dynamics and wind predictions into long planning, enabling the controller to exploit future wind patterns rather than reacting only to current wind conditions. Frequent replanning then mitigates forecast errors from imperfect wind predictions and dynamics fidelity, allowing even simplified dynamics and wind models to outperform model-free RL policies.

We evaluate FOMPC in the Balloon Learning Environment (BLE) [19] against Google’s state-of-the-art RL controller, Perciatelli44 [20], and show that FOMPC achieves up to a 24% improvement in time-within-radius (TWR), or 1.8 extra hours in radius, with the same wind information. Through systematic ablations on planning horizon, replan interval, model fidelity, and wind models, we further identify the factors most responsible for performance gains and robustness. Our main contributions are:

- We release¹ an open-source, fully differentiable implementation of high-altitude balloon dynamics in JAX to support future research in model-based control and learning.
- We propose First-Order Model Predictive Control (FOMPC), which outperforms Perciatelli44 by 24% in time-within-radius performance, an additional 1.8 hours a day within radius, and establishes a systematic model-based baseline for station-keeping.
- We conduct a comprehensive ablation study of modeling and control parameters, identifying which factors most strongly drive performance and where simplified models remain effective. These insights provide concrete guidance for future RL and hybrid approaches.

II. RELATED WORK

Most prior research on HAB station-keeping focuses on model-free RL, motivated by the perceived difficulty of applying model-based approaches under uncertain winds. Non-RL alternatives remain limited, typically relying on heuristic-based or data-driven methods, and rarely compared against RL. Meanwhile, model predictive control (MPC) in broader aerospace applications shows that MPC can perform well under wind disturbances and nonlinear dynamics, suggesting it to be a promising but underexplored baseline for the HAB station-keeping problem.

A. Reinforcement Learning Approaches

Most prior work focuses on reinforcement learning (RL) methods. Project Loon addressed the station-keeping problem using quantile regression deep Q-network (DQN) to maintain balloons at altitudes around 20 km for long durations while providing Internet connectivity [14]. Later, they further improved efficiency by reusing prior computation and experience to accelerate learning, demonstrated in their Balloon Learning Environment [20]. Subsequent work

applied variants of DQN [9], [12], [13] and SAC [11] to tasks including station-keeping, navigation, and no-fly zone avoidance, as well as to different balloon types such as sounding and latex balloons. Collectively, these efforts show that RL can exploit wind structure without explicit models, but do not amend the view that model-based approaches are impractical under stratospheric uncertainty.

B. Non-RL Alternatives

Non-RL alternatives remain sparse. Heuristic and model-based altitude controllers have been proposed without long-horizon planning [7], and extremum-seeking controllers learn favorable winds without explicit models [21], but neither is compared to RL. While many works argue that model-based methods are fragile under underactuation and forecast errors, studies in related flow settings show that replanning can enable robust performance [22]. Most closely related, a receding-horizon controller using short-term forecasts and optimistic optimization has been demonstrated [23], but the method relies on gradient-free optimization and is only compared against tree search, not RL.

C. Model Predictive Control

Model Predictive Control (MPC) has been widely applied in unmanned aerial systems (UAS), demonstrating robustness under wind disturbances and nonlinear dynamics. For example, constrained MPC solved via quadratic programming has stabilized quadrotor attitude and position [24], while efficient nonlinear formulations have improved real-time performance by shortening prediction horizons [25]. Adaptive nonlinear MPC has been used for highly unsteady aircraft dynamics [26], and cooperative MPC has enabled coordinated aerial-ground rendezvous with reduced computation [27]. In lighter-than-air systems, MPC has guided stratospheric airships using linearized dynamics [28] and extended to constrained trajectory tracking under unknown disturbances with Laguerre function parameterizations [29]. These successes highlight MPC’s versatility but leave open whether it can plan effectively in uncertain wind fields with the limited control authority of HABs.

III. PROBLEM FORMATION

The goal of station-keeping is to maintain the balloon’s position within a communication range of a ground station. The evaluation metric is time-within-radius (TWR), defined as the fraction of time steps during which the balloon remains within 50 km of the ground station center. We frame station-keeping as a finite-horizon discrete-time optimal control problem. Over a planning horizon with H steps, the objective is to compute the control policy π that maximizes the expected TWR of subject to the balloon dynamics:

$$\begin{aligned} \max_{\pi} \quad & \mathbb{E} \left[\frac{1}{H} \sum_{k=0}^{H-1} \mathbf{1}(x_k^2 + y_k^2 \leq 50^2) \right] \\ \text{s.t.} \quad & \mathbf{x}_{k+1} = F_{\Delta t}(\mathbf{x}_k, \mathbf{w}_k, u_k, t_k), \\ & k = 0, \dots, H-1, \quad \mathbf{x}_0 = \mathbf{x}_{\text{init}}, \\ & u_k = \pi(\mathbf{x}_k, \hat{W}), \end{aligned} \tag{1}$$

¹<https://github.com/sdean-group/balloon-learning-environment>

where $\mathbf{1}(\cdot)$ is the indicator of being within the 50 km radius. Here, \mathbf{x}_k is the balloon state, $u_k \in \mathbb{R}$ is the vertical control input, t_k is the timestamp. \hat{W} is the wind measurements, while \mathbf{w}_k is the wind at the balloon's position (Sec IV-A). $F_{\Delta t}$ is the state update (5). The expectation is over forecast uncertainty and stochastic wind realizations. This objective is consistent with prior RL evaluations, enabling direct comparison.

IV. SYSTEM DYNAMICS MODEL

We describe balloon dynamics as described in BLE [19]. These are re-implemented in JAX with an additional continuous control relaxation for differentiability. HAB dynamics can be decomposed into horizontal dynamics, governed solely by the wind, and altitude dynamics, determined by buoyancy and net mass.

Let \mathbf{x} define the state vector of the superpressure balloon.

$$\mathbf{x} = [x, y, l, n, T, E, V, p_{\text{env}}], \quad (2)$$

where (x, y) are horizontal offsets, l pressure, n ballonnet air moles, T temperature, E battery energy, V envelope volume, and p_{env} envelope superpressure.

Dynamics follow

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{w}, u, t), \quad (3)$$

implemented with discrete updates computed with $\Delta t = M\delta t$ via M Euler substeps ($\delta t = 10\text{s}$, $M = 18$) as follows:

$$\begin{aligned} \mathbf{x}^{(m+1)} &= \mathbf{x}^{(m)} + \delta t f(\mathbf{x}^{(m)}, \mathbf{w}_k, u_k, t_k + m\delta t) \\ m &= 0, \dots, M-1, \quad \mathbf{x}^{(0)} = \mathbf{x}_k \end{aligned} \quad (4)$$

The final state after one simulator step is $\mathbf{x}_{k+1} = \mathbf{x}^{(M)}$. This composition defines our discrete update step:

$$\mathbf{x}_{k+1} = F_{\Delta t}(\mathbf{x}_k, \mathbf{w}_k, u_k, t_k) \quad (5)$$

A. Wind Fields in Environment

BLE generates forecast winds W_0 from a variational autoencoder (VAE) trained on the ERA5 reanalysis dataset [30]. The ground truth wind, W , is defined with added Simplex noise b to emulate forecast errors:

$$W(x, y, l, t) = W_0(x, y, l, t) + b(x, y, l, t). \quad (6)$$

We also define $\mathbf{w}_k \equiv W(x_k, y_k, l_k, t_k)$.

B. Balloon Dynamics

We decompose balloon dynamics into six components, consistent with the Balloon Learning Environment (BLE) simulator.

1) *Horizontal Dynamics*: The balloon drifts with the ambient wind at its current location and pressure level.

$$\dot{x} = \mathbf{w}_x, \quad \dot{y} = \mathbf{w}_y$$

2) *Vertical Dynamics*: Let h denote the altitude of the balloon. Vertical motion is defined by buoyant, gravity, and drag forces:

$$m\ddot{h} = F_b - F_g - F_d, \quad (7)$$

but BLE enforces steady state ($\ddot{h} = 0$):

$$0 = \rho_{\text{air}} V g - m g - \frac{1}{2} \rho_{\text{air}} C_d A |\dot{h}| \dot{h}. \quad (8)$$

We can directly solve for \dot{h} from (8). Using the U.S. Standard Atmosphere Model [31], we convert \dot{h} into a pressure differential, \dot{l} , which allows us to fully model pressure dynamics.

3) *Thermal Model*: The envelope temperature T evolves from ambient temperature, solar flux, and radiation balance, which determines the density of air inside the envelope. Warmer internal temperature reduces gas density, increases buoyancy, and thus promotes ascent; vice versa.

4) *Volume and Superpressure Model*: In (8), a larger volume V directly increases buoyancy and promotes ascent. Superpressure p_{env} limits control: once fully stretched, added gas/heat raises p_{env} not V , increasing venting when high and risking collapse when low.

5) *Altitude Control System (ACS)*: The ACS adjusts n , the moles of air, via a continuous action $u \in [-1, 1]$, with venting for $u > 0$ and pumping for $u < 0$:

$$\dot{n}(u) = \begin{cases} -c_{\text{vent}}(l, T, p_{\text{env}})u, & u > 0, \\ c_{\text{pump}}(l, T, p_{\text{env}})(-u), & u < 0, \\ 0, & u = 0, \end{cases} \quad (9)$$

with power

$$P_{\text{acs}}(u) = \begin{cases} \eta(l, p_{\text{env}})(-u), & u < 0, \\ 0, & u \geq 0. \end{cases} \quad (10)$$

This is a relaxation compared with the original discrete actions defined in BLE. It allows our algorithm to plan a continuous control input u , which is mapped into physically realistic venting or pumping rates.²

6) *Battery Dynamics*: Battery energy evolves as $\dot{E} = P_{\text{solar}}(t) - P_{\text{load}}(t) - P_{\text{acs}}(u)$, with BLE disabling downward pumping when $E/E_{\text{max}} < 0.025$ until recovery above 0.05 (considered a power safety violation).

V. MODEL PREDICTIVE CONTROL DESIGN

MPC is a control strategy that optimizes actions over a receding horizon. Typical MPC formulations attempt to simplify the planning problem into a convex surrogate. However, iterative approximations and standard nonlinear optimization techniques can be computationally prohibitive [32]. Instead, we reformulate the planning task as an unconstrained optimization problem. This allows us to leverage gradient-based optimization to solve the problem with significantly higher efficiency. We propose *First-Order MPC (FOMPC)*, which we describe by its main components:

1) *Continuous Action Construction*: To ensure our control actions, which are only defined for $u \in [-1, 1]$, can be treated as unconstrained decision variables, we impose a squashing function $u = \sigma(\tilde{u}) = \frac{2}{1 + \exp(-\tilde{u})} - 1$ to guarantee admissibility by construction.

²Re-discretizing post-optimization yields similar time-within-radius performance, supporting the validity of the continuous relaxation. Using the M Euler substeps to approximate continuous actions as multiples of $\frac{1}{M}$ by alternating discrete actions also produced comparable results.

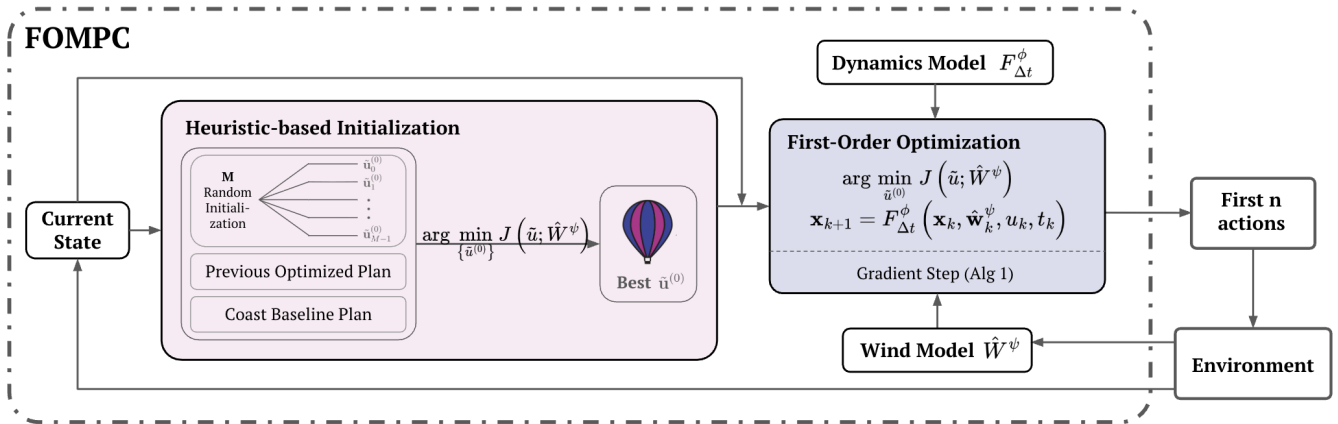


Fig. 2: Block diagram of FOMPC. Candidate plans are initialized using heuristics, then refined by first-order optimization with dynamics model $F_{\Delta t}^\phi$ and wind model \hat{W}^ψ . The lowest-cost plan provides the first n actions, after which the process repeats with the updated state.

2) *Dynamics Model*: In FOMPC, the planning dynamics $F_{\Delta t}^\phi$ need not match the full BLE dynamics $F_{\Delta t}$. We define reduced-order models by omitting selected physical states and updates. Each fidelity level $\phi \in \{\phi_0, \dots, \phi_4\}$ specifies which processes are integrated for each δt and which are clamped (Table I). The resulting dynamics are

$$\mathbf{x}_{k+1} = F_{\Delta t}^\phi(\mathbf{x}_k, \mathbf{w}_k, u_k, t_k), \quad (11)$$

where $F_{\Delta t}^\phi$ applies M sub-steps of Euler integration with the subset of updates determined by ϕ .

Level	ACS (§IV-B.5)	Vol/Superpres. (§IV-B.4)	Temp. (§IV-B.3)	Batt. (§IV-B.6)
ϕ_0	×	×	×	×
ϕ_1	✓	×	×	×
ϕ_2	✓	✓	×	×
ϕ_3	✓	✓	✓	×
ϕ_4	✓	✓	✓	✓

TABLE I: Fidelity levels and enabled submodules

When clamped, the values of states in a submodule are held to their values in \mathbf{x}_0 . As the ACS is not a state variable directly, we instead replace it with a piecewise-linear surrogate

$$\dot{n}(u) = \begin{cases} -0.012u, & u > 0, \\ 0.007(-u), & u < 0, \\ 0, & u = 0, \end{cases} \quad (12)$$

and the power usage with a similar linearization $P_{\text{acs}}(u) = 195(-u)$ when $u < 0$ (power usage remains 0 for $u \geq 0$). These constants are averages derived from typical ACS usage.

3) *Wind Model*: The balloon's horizontal dynamics depend on the true wind W , while MPC plans, at best, can have an approximate model \hat{W}^ψ . Because the forecast W_0 contains systematic errors, we construct several wind models for investigation.

We define four variants:

- (i) $\hat{W}^{\psi_0} = W_0$, where MPC can plan through the forecast.
- (ii) $\hat{W}_{\mathbf{x}_0}^{\psi_1}(x, y, l, t) = \hat{W}(x_0, y_0, l, t_0)$, where W_0 is treated as constant across (x, y, t) , forming a wind column centered around the balloon state at the start of a replan.
- (iii) $\hat{W}_{\mathbf{x}_0}^{\psi_2}(x, y, l, t) = W_0(x_0, y_0, l, t_0) + \mathcal{G}(x_0, y_0, l, t_0)$, where \mathcal{G} is a Gaussian Process trained to predict forecast noise from forecast errors $\{\mathbf{w}_{0:t} - \mathbf{w}_{0:t}\}$ sampled at the locations $\{\mathbf{x}_{0:t}\}$ traversed since previous replan, from [14].
- (iv) $\hat{W}_{\mathbf{x}_0}^{\psi_3} = \alpha \cdot W_0 + (1 - \alpha) \cdot \hat{W}_{\mathbf{x}_0}^{\psi_2}$, which is a combination of the forecast and the processed wind column to take advantage of online re-estimation of winds while still incorporating the wind variation across (x, y, t) . We set $\alpha = 0.5$.

As described in [14] and implemented in [19], the GP uses a Matérn kernel and is trained once per replan at \mathbf{x}_0 , returning estimated wind forecast errors. When combined with the original forecast, this provides an online estimate of the local wind forecast we use in \hat{W}^{ψ_2} and \hat{W}^{ψ_3} .

4) *Differentiable Cost Function Design*: To replace the non-smooth TWR objective, we use a differentiable surrogate penalizing the squared distance and low battery:

$$c(\mathbf{x}) = \left(\frac{x}{1000}\right)^2 + \left(\frac{y}{1000}\right)^2 + R^2 c_{\text{power}}(\mathbf{x}) \quad (13)$$

where $c_{\text{power}}(\mathbf{x}) = 1 - \sigma(100(\frac{E}{B} - 0.1))$, and $\sigma(z) = 1/(1 + \exp(-z))$. The c_{power} term begins penalizing usage around 10% capacity, approximating the power safety layer that activates at 2.5% and 5% (Sec. IV-B.6).

5) *Receding-horizon Planning*: At replanning instant t_n , with wind model \hat{W}^ψ , dynamics model $F_{\Delta t}^\phi$, current balloon state $\mathbf{x}(t_n)$, planning horizon H and discount factor γ , we define the optimization objective Eq. 14.

Substituting dynamics and control updates yields an unconstrained direct shooting formulation. We approximately

$$\begin{aligned}
\min_{\{\tilde{u}_k\}} \quad & J_n(\{\tilde{u}_k\}, \mathbf{x}_0; \hat{W}^\psi) := \sum_{k=0}^{H-1} \gamma^k c(\mathbf{x}_{k+1}) \\
\text{s.t.} \quad & \mathbf{x}_{k+1} = F_{\Delta t}^\phi(\mathbf{x}_k, \hat{\mathbf{w}}_k^\psi, u_k, t_k), \quad k = 0, \dots, H-1 \\
& u_k = \sigma(\tilde{u}_k) \mathbf{1}\left(\frac{E_k}{B} \geq 0.025\right), \quad \mathbf{x}_0 = \mathbf{x}(t_n)
\end{aligned} \tag{14}$$

solve for H actions, executing the first m before the next replanning cycle.

6) *Heuristic-based Plan Initialization*: To mitigate non-convexity, FOMPC initializes from the best of M random candidates, the previous optimized plan, and a ‘‘coast’’ baseline generated from $\text{Unif}(-0.2, 0.2)$ noise that keeps the balloon nearly stationary while remaining differentiable. Random candidates are constructed by precomputing the steps $\hat{t}(h)$ needed to reach altitudes $h \in [h_{\min}, h_{\max}]$, moving toward a sampled target altitude with near-saturated controls for $\hat{t}(h)$ steps, and adding small random perturbations for the remainder. Only the lowest cost initialization among all plans is then used in gradient descent.

7) *Gradient Descent Optimization*: We implement the first-order optimization directly on unconstrained $\tilde{\mathbf{u}}$.

Algorithm 1 First-order gradient descent

- 1: **Input**: horizon H , step size $\eta > 0$, tolerance $\epsilon > 0$, max iters S_{\max} , wind model \hat{W}^ψ , heuristic-based initialization $\tilde{\mathbf{u}}^{(0)}$
 - 2: **for** $s = 0$ to S_{\max} **do**
 - 3: $g^{(s)} \leftarrow \nabla_{\tilde{\mathbf{u}}} J_n(\tilde{\mathbf{u}}, \mathbf{x}_0; \hat{W}^\psi) \Big|_{\tilde{\mathbf{u}}=\tilde{\mathbf{u}}^{(s)}}$
 - 4: **if** $\|g^{(s)}\|_2 < \epsilon$ **then**
 - 5: **break**
 - 6: **end if**
 - 7: $\tilde{\mathbf{u}}^{(s+1)} \leftarrow \tilde{\mathbf{u}}^{(s)} - \eta \frac{g^{(s)}}{\|g^{(s)}\|_2}$
 - 8: **end for**
 - 9: **Output**: $\{u_k\}_{k=0}^{H-1} \leftarrow \{\sigma(\tilde{u}_k^{(s)})\}_{k=0}^{H-1}$
-

For notation simplicity, we define $\tilde{\mathbf{u}}^{(s)} := \{\tilde{u}_k^{(s)}\}_{k=0}^{H-1}$. We solve (14) with gradient descent, computing $\nabla_{\tilde{\mathbf{u}}} J$ by autodiff (JAX), demonstrated in Algorithm 1. We use step size $\eta = 1$, $S_{\max} = 100$, tolerance $\epsilon = 10^{-7}$.

VI. EXPERIMENTS

We explore how FOMPC performance compares to Perciatelli44 under identical conditions and how the hyperparameters of FOMPC affect its own performance and computational efficiency. Perciatelli44 is the implementation of the state-of-the-art deep distributional Q-learning (QR-DQN) station-keeping controller described in [14], [20], [33] and implemented in the Balloon Learning Environment (BLE) [19].

After running FOMPC and Perciatelli44 in BLE with the same set of conditions, we measure performance metrics such as time-within-radius (TWR). We also vary the control parameters and modeling assumptions of FOMPC and evaluate under these conditions to understand how they affect FOMPC performance.

A. Experimental Setup

We begin by describing the evaluation protocol. All experiments are conducted in BLE on the station-keeping task. Each trial consists of 960 steps, with $\Delta t = 3$ minutes of simulation time between steps, totaling 2 days of simulated flight time. Additionally, each trial is initialized with a seed, which generates a realistic balloon starting state and unique wind conditions. The time and location of the trial is randomly selected, also based on the seed, uniformly between 2011-1-1 and 2014-12-31 and between $[-10, 10]$ latitude and $[-175, 175]$ longitude. The randomness is seeded from the seed of the whole episode. After every step and at the start of each episode, an agent is given a set of features about the balloon and environment and produces an action. Each agent we evaluate is measured on the same set of randomized wind field seeds to ensure comparability. After running each seed, we collect the TWR (1), online execution time per step (i.e., time it takes an agent to calculate an action), power safety layer violations (IV-B.6), and the Perciatelli44 reward over all the steps. The Perciatelli44 reward function counts the steps within the radius and has an exponential drop-off term for steps outside the radius, with an additional power regularization term for safe power usage [19].

For FOMPC parameter ablations, each agent is evaluated on seeds $[0, 1000)$ in BLE. This range provides a broad distribution of conditions that captures the variability of real-world stratospheric winds, from featureless flows that are difficult to navigate to diverse winds that lend themselves to station-keeping. We collected a larger set of seeds $[0, 10\,000)$ for Perciatelli44 and our default FOMPC to provide a more thorough comparison. For all results, we report the mean of these metrics across trials and the 95% confidence interval computed by $\pm 1.96 \left(\frac{\text{variance}}{\text{number of trials}}\right)^{1/2}$. Seed sets $[0, 1000)$ and $[0, 10\,000)$ are provided as evaluation benchmarks in BLE [19].

Each agent receives features from the environment before computing an action. The feature set for FOMPC is the current state \mathbf{x}_{t_n} and wind model \hat{W}^ψ , which may be initialized with state \mathbf{x}_{t_n} . The feature set for Perciatelli44 is the Gaussian process (GP)-refined wind column, $\hat{W}_{\mathbf{x}_{t_n}}^{\psi/2}$ evaluated at a pre-defined set of reachable pressure levels. This set of wind values is concatenated with the balloon state \mathbf{x}_{t_n} in one feature vector.

B. Overall Performance (FOMPC vs Perciatelli44)

Metric	MPC	Perciatelli44
Reward	513.27 \pm 5.36	425.89 \pm 5.25
Time Within Radius (%)	39.43 \pm 0.65	31.78 \pm 0.60
Execution Time per Step (s)	0.131 \pm 0.000	0.00110 \pm 0.000
Power Safety Violations	8.70 \pm 0.55	22.14 \pm 0.84

TABLE II: comparison between FOMPC and Perciatelli44. Values are reported as mean \pm 95% confidence interval over 10k trials.

We define a default FOMPC by selecting hyperparameters with high performance and reasonable computational cost.




Fig. 3: Performance–efficiency trade-offs in FOMPC. Each plot reports mean time-within-radius relative to Perciatelli44 (error bars show 95% confidence intervals) as a function of online execution time per step. Markers denote parameter settings; bolded values indicate defaults (12h horizon, 72m replan interval, 100 initializations, highest dynamics fidelity, GP-Column wind model). Each plotted settings varies one parameter at a time from this default.

These are a 12-hour planning horizon ($H = 240$ steps), 72-minute replanning interval ($n = 24$ steps), 100 random initializations, ϕ_4 dynamics fidelity level, and ψ_2 wind model, the GP-refined wind column.

To test whether a model-based planner can achieve station-keeping performance competitive with RL policies, we compare default FOMPC to Perciatelli44.

Table II shows that our FOMPC consistently outperforms Perciatelli44 in terms of RL reward and TWR, while also reducing power safety violations. This improvement comes at the expense of increase in computation time per control step. However, this time per step is still much less than the replan interval of 72 minutes between when new actions are needed.

C. Ablation Studies

Next, we investigate how individual FOMPC hyperparameters influence TWR performance and execution time per step. In each ablation, we vary a single parameter while holding all other parameters at the default setting. This controlled setup isolates the effect of each design choice. We consider five factors: planning horizon, replan interval, number of initializations, dynamics fidelity, and wind model.

1) *Planning Horizon:* We vary the planning horizon across five values: 90 minutes, 3 hours, 6 hours, 12 hours (default), and 48 hours ($H = 30, 60, 120, 240, 480$). Longer horizons allow the controller to anticipate wind patterns further into the future, but also increase computational cost. As shown in Fig. 3a, TWR improves steadily up to a horizon of 12 hours, but drops off when extended to 24 hours, indicating that overly long horizons can degrade performance despite higher computation. All but the shortest horizon outperform Perciatelli44.

2) *Replan Interval:* We vary the replan interval across five values: 576 minutes, 288 minutes, 144 minutes, 72 minutes (default), and 36 minutes ($m = 192, 96, 48, 24, 12$). Shorter replan intervals allow the controller to adapt more frequently to changing wind conditions, but also increase computational

cost. As shown in Fig. 3a, performance improves as the replan interval decreases, but gains reflect diminishing returns at higher computational cost. All but the longest replan outperform Perciatelli44.

3) *Number of Initializations:* We vary the number of random optimization initializations across six settings: 1, 5, 25, 100 (default), 250. Considering more initializations improves the chances of avoiding poor local minima in the optimization process but increases runtime. Fig. 3b shows a performance improvement with diminishing returns at higher computational costs. Even a single random initialization outperforms Perciatelli44.

4) *Dynamics Fidelity:* We vary the level of dynamics fidelity across five settings: $\phi_0, \phi_1, \phi_2, \phi_3, \phi_4$ (default) (See D). Lower-fidelity dynamics reduce computational cost by simplifying the balloon dynamics, but may degrade performance if important physical effects are omitted. As shown in Fig. 3c, TWR improves as fidelity increases, but there are diminishing returns. Even the simplest fidelity level outperforms Perciatelli44.

5) *Wind Model:* We vary the wind model across four models: Forecast, Column, GP-Column (default), GP-Column + Forecast ($\psi = \psi_0, \psi_1, \psi_2$ (default), ψ_3 See V-3). Fig. 3c reveals the importance of online estimation of forecast errors. FOMPC with the GP-informed wind column \hat{W}^{ψ_2} outperforms that with the full forecast \hat{W}^{ψ_0} despite the latter’s richer spatiotemporal information, while the combined model \hat{W}^{ψ_3} (marginally) outperforms both. This benefit comes at the cost of additional online planning time to fit the Gaussian Process. We also note that the timing difference between column and forecast models is negligible due to their similar implementations in JAX. All wind models demonstrate improvement over Perciatelli44.

D. Discussion

Our experiments highlight that FOMPC, a model-based controller, outperforms the state-of-the-art RL controller on TWR while more closely adhering to power safety

constraints. Many hyperparameter configurations still surpass Perciatelli44, with performance degrading only under extremely constrained horizons or replan intervals. Even lightweight dynamics and limited wind information retain their performance. These gains come at the cost of higher per-step computation, but this is practical in the balloon domain, where actions are required only every few minutes rather than at millisecond frequencies typical in other MPC applications. It is also the case that many balloons are controlled remotely with action sequences computed elsewhere and sent to the balloon via satellite communication. In addition, unlike Perciatelli44, FOMPC requires no training data and incurs no upfront training cost. However, further experimentation (Fig. 4) shows Perciatelli44 is more robust to unreliable forecasts; conversely, FOMPC overfits and degrades under high-wind noise.

The ablations show that incorporating learned forecast corrections into the wind model yield the largest gains in TWR. Interestingly, spatiotemporal forecast data alone offered less benefit for long-horizon planning, suggesting that compensating for forecast bias is more impactful than merely extending wind information.

To better understand RL’s failure modes, we first compare performance under the Perciatelli44 reward. As shown in Table II, FOMPC achieves higher scores than Perciatelli44 even on this reward, despite optimizing a different surrogate for TWR. This indicates that the reward used by Perciatelli44 is a reasonable proxy for TWR, and its lower performance cannot be attributed to an objective mismatch. We also note that \hat{W}^{ψ_2} contains the same information as Perciatelli44, yet our model-based controller outperforms Perciatelli44, meaning the measurement-informed wind column is sufficient for high-performing TWR. Two remaining hypotheses for the Perciatelli44 performance are 1) a limited policy class or 2) under-optimization. Our experiments do not let us form a conclusion on whether the limitations stem primarily from the expressiveness of the Q-network architecture or from inadequate training.

VII. CONCLUSION

In this work, we revisited the problem of high-altitude balloon station-keeping through the lens of MPC, due to its absence in the literature as a baseline for RL and its potential to generalize across new dynamics and unseen wind conditions. By formulating the balloon dynamics as differentiable functions in JAX and applying FOMPC, we established a model-based controller that consistently outperforms the state-of-the-art RL controller Perciatelli44 in the BLE. Our experiments demonstrate that planning over a horizon, even when restricted to the same wind information, leads to more effective control decisions, yielding significant gains in TWR while also reducing safety violations.

Through systematic ablations of planning horizon, replan interval, initialization strategy, dynamics fidelity, and wind models, we identified which factors contribute most to performance and where computational efficiency can be traded against accuracy. These insights highlight that with the same

wind information, MPC can outperform Perciatelli44 in both Perciatelli44 reward and time within radius. It is possible that Perciatelli44’s QR-DQN structure is not able to represent long-horizon strategies or adapt to wind conditions in the way FOMPC can. Alternatively, its performance may be limited by under-optimization, as training may not fully converge given the complexity of the environment.

To investigate, future work can use MPC as a teacher for the Perciatelli44 policy network. If imitation learning enables the policy to match MPC performance, this would indicate that the network class can represent long-horizon solutions but was under-optimized. If not, it would suggest that this type of Q-network is insufficiently expressive for the task. More broadly, this points to opportunities for hybrid approaches where reinforcement learning can distill or imitate MPC, combining the generalization ability and runtime efficiency of policies with the foresight of model-based planning.

There are additional opportunities for future work on pure MPC approaches, such as addressing robustness to uncertainty and runtime. In additional experiments presented in the appendix, we observe that FOMPC remains sensitive to extreme noise (Fig. 4). Model Predictive Path Integral (MPPI) [34] is a scalable alternative which can address computation time via hardware parallelization, but it currently fails to stay in safe regions (Table III).

Our differentiable balloon model and strong MPC baselines thus provide a foundation for advancing navigation in the stratosphere and other underactuated flow-driven systems.

VIII. ACKNOWLEDGMENTS

This work was partly funded by NSF CCF 2312774, NSF OAC-2311521, NSF IIS-2442137, a gift to the LinkedIn-Cornell Bowers CIS Strategic Partnership, the AI2050 Early Career Fellowship program at Schmidt Sciences, and a 2030 Project Fast Grant from the Cornell Atkinson Center for Sustainability.

REFERENCES

- [1] V. Picouet, “Stratospheric balloon payloads for astronomy: the challenge of coping with rising complexity,” *arXiv preprint arXiv:2501.16047*, 2025.
- [2] J. B. Habeck, J. Flaten, and G. V. Candler, “High-altitude balloon measurements of atmospheric particulates,” in *AIAA Scitech 2020 Forum*, 2020, p. 1794.
- [3] 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.
- [4] M. D. Ross and M. L. Lewis, “The strato-lab balloon system for high altitude research,” *J Aviat Med*, vol. 29, no. 5, pp. 375–385, 1958.
- [5] S. Klomchitharoen, P. Wechakarn, T. Tangwattanasirikun, N. Smerwong, P. Netraphathompornkij, T. Chatmeeboon, N. Nangsue, V. Thitasirivit, K. Kaweewongsunthorn, S. Piyanopharoj *et al.*, “High-altitude balloon platform for studying the biological response of living organisms exposed to near-space environments,” *Heliyon*, vol. 10, no. 6, 2024.
- [6] H. Du, M. Lv, L. Zhang, W. Zhu, Y. Wu, and J. Li, “Energy management strategy design and station-keeping strategy optimization for high altitude balloon with altitude control system,” *Aerospace Science and Technology*, vol. 93, p. 105342, 2019.

- [7] 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.
- [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, 2023.
- [9] 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.
- [10] Y. He, K. Guo, C. Wang, K. Fu, and J. Zheng, "Path planning for autonomous balloon navigation with reinforcement learning," *Electronics*, vol. 14, no. 1, p. 204, 2025.
- [11] 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.
- [12] M. Gannetti, M. Gemignani, and S. Marcuccio, "Navigation of sounding balloons with deep reinforcement learning," in *2023 IEEE 10th International Workshop on Metrology for AeroSpace (MetroAeroSpace)*. IEEE, 2023, pp. 591–596.
- [13] 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.
- [14] 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.
- [15] K. Garg, "Wind based navigation for zero-pressure stratospheric balloons using reinforcement learning," *Acta Astronautica*, 2019.
- [16] R. Agarwal, M. Schwarzer, P. S. Castro, A. C. Courville, and M. Bellemare, "Deep reinforcement learning at the edge of the statistical precipice," *Advances in neural information processing systems*, vol. 34, pp. 29 304–29 320, 2021.
- [17] H. Mania, A. Guy, and B. Recht, "Simple random search provides a competitive approach to reinforcement learning," *arXiv preprint arXiv:1803.07055*, 2018.
- [18] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, "JAX: composable transformations of Python+NumPy programs," 2018. [Online]. Available: <http://github.com/google/jax-ml/jax>
- [19] J. Greaves, S. Candido, V. Dumoulin, R. Goroshin, S. S. Ponda, M. G. Bellemare, and P. S. Castro, "Balloon Learning Environment," 12 2021. [Online]. Available: <https://github.com/google/balloon-learning-environment>
- [20] R. Agarwal, M. Schwarzer, P. S. Castro, A. C. Courville, and M. Bellemare, "Reincarnating reinforcement learning: Reusing prior computation to accelerate progress," *Advances in neural information processing systems*, vol. 35, pp. 28 955–28 971, 2022.
- [21] 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.
- [22] M. Wiggert, M. Doshi, P. F. Lermusiaux, and C. J. Tomlin, "Navigating underactuated agents by hitchhiking forecast flows," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 2417–2424.
- [23] Y. Fan, X. Deng, X. Yang, Y. Long, and F. Bai, "Station-keeping control of stratospheric balloons based on simultaneous optimistic optimization in dynamic wind," *Electronics*, vol. 13, no. 20, p. 4032, 2024.
- [24] M. Aliyari, W.-K. Wong, Y. Bouteraa, S. Najafinia, A. Fekih, and S. Mobayen, "Design and implementation of a constrained model predictive control approach for unmanned aerial vehicles," *IEEE Access*, vol. 10, pp. 91 750–91 762, 2022.
- [25] M. A. Gomma, O. De Silva, G. K. Mann, and R. G. Gosine, "Computationally efficient stability-based nonlinear model predictive control design for quadrotor aerial vehicles," *IEEE Transactions on Control Systems Technology*, vol. 31, no. 2, pp. 615–630, 2022.
- [26] G. A. Garcia, S. S. Keshmiri, and T. Stastny, "Robust and adaptive nonlinear model predictive controller for unsteady and highly nonlinear unmanned aircraft," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 4, pp. 1620–1627, 2014.
- [27] L. Persson, T. Muskardin, and B. Wahlberg, "Cooperative rendezvous of ground vehicle and aerial vehicle using model predictive control," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, 2017, pp. 2819–2824.
- [28] J. Zhang, X. Yang, X. Deng, and H. Lin, "Trajectory control method of stratospheric airships based on model predictive control in wind field," *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, vol. 233, no. 2, pp. 418–425, 2019.
- [29] J. Yuan, M. Zhu, X. Guo, and W. Lou, "Trajectory tracking control for a stratospheric airship subject to constraints and unknown disturbances," *IEEE Access*, vol. 8, pp. 31 453–31 470, 2020.
- [30] H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, A. Simmons, C. Soci, S. Abdalla, X. Abellan, G. Balsamo, P. Bechtold, G. Biavati, J.-R. Bidlot, M. Bonavita, G. De Chiara, P. Dahlgren, D. Dee, M. Diamantakis, R. Dragani, J. Flemming, J. Forbes, M. Fuentes, A. Geer, L. Haimberger, S. Healy, R. J. Hogan, E. Hólm, M. Janisková, S. Keeley, P. Laloyaux, P. Lopez, C. Lupu, G. Radnoti, P. de Rosnay, I. Rozum, F. Vamborg, S. Villaume, and J.-N. Thépaut, "Complete era5: Fifth generation of cmwf atmospheric reanalyses of the global climate," Copernicus Climate Change Service (C3S) Data Store (CDS), 2017, accessed on 01-04-2021.
- [31] U. S. C. on Extension to the Standard Atmosphere, *US standard atmosphere, 1976*. National Oceanic and Atmospheric Administration, 1976.
- [32] J. Zhang, A. Jordana, and L. Righetti, "Accelerated gradient descent for high frequency model predictive control," *arXiv preprint arXiv:2409.18327*, 2024.
- [33] W. Dabney, M. Rowland, M. G. Bellemare, and R. Munos, "Distributional reinforcement learning with quantile regression," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018. [Online]. Available: <https://doi.org/10.48550/arXiv.1710.10044>
- [34] G. Williams, A. Aldrich, and E. Theodorou, "Model predictive path integral control using covariance variable importance sampling," *arXiv preprint arXiv:1509.01149*, 2015. [Online]. Available: <https://arxiv.org/abs/1509.01149>

APPENDIX

Robustness of FOMPC and Perciatelli44 is evaluated against wind field $W = W_0 + \mu \cdot b$ (Eq. 6), where μ is the noise scale. For each μ , we evaluate on seeds $[0, 100)$. Fig. 4 shows the time-within-radius difference (95% CI). FOMPC performance degrades relative to Perciatelli44, with a crossover occurring at $\mu \approx 3$, triple the typical noise level.




Fig. 4: TWR sensitivity of FOMPC and Perciatelli44 to scaled wind magnitude.

We also test MPPI [34]. It is competitive across seeds $[0, 1000)$, but susceptible to altitude-related early termination (over-pressurization). Enhanced exploration or safety-aware sampling would likely resolve these edge-case failures.

Metric	MPC	Perciatelli44	MPPI
TWR (%)	37.10 ± 2.02	29.55 ± 1.84	30.86 ± 1.87
Time / Step (s)	0.132 ± 0.00	0.001 ± 0.00	0.089 ± 0.00

TABLE III. Comparison of FOMPC, Perciatelli44, and MPPI.