Materials Research Proceedings 69 (2026) 846-849 https://doi.org/10.21741/9781644904251-149

# Reinforcement Learning-Based Bezier Airship Hull Optimization: A Data-Driven Framework for Shape Adaptation Under Aerodynamic Constraints

Qian ZHAO[1,a]*, and Carlo Emanuele Dionigi RIBOLDI[1,b]

[1]Department of Aerospace Science and Technology, Politecnico di Milano, via La Masa 34, 20156 Milano, Italy

[a]qian.zhao@polimi.it, [b]carlo.riboldi@polimi.it

**Keywords:** Airship Design, Bezier Curves, Reinforcement Learning, DDPG, Lift-to-Drag Ratio Optimization, Robust Design, Aerodynamic Efficiency, Data-Driven Optimization

**Abstract.** This paper proposes a data-driven reinforcement learning framework for aerodynamic optimization of airship hulls. It combines Deep Deterministic Policy Gradient (DDPG) with Bezier curve parameterization to adjust the shape of the classic airship ”Lotte” [1,2,3]. A volume-preserving constraint ensures buoyancy consistency. Different from previous studies focused solely on drag minimization, we target lift-to-drag ratio ($L/D$) and robustness [4,5]. Simulations in SILCROAD [6,7,8,9] demonstrate significant improvement in $L/D$ while preserving volume and smoothness. The approach offers a modular, scalable framework for future multi-objective extensions..

## Introduction

Airships are regaining attention for applications like surveillance, communication relay, and environmental monitoring [12]. A key design challenge is aerodynamic shape optimization of the hull, which affects drag, $L/D$, energy use, and mission range.

Traditional airship shapes often rely on empirical formulas or fixed parametric models, limiting flexibility for mission-specific designs. Prior aerodynamic optimization focused primarily on drag at constant angles, overlooking $L/D$, a critical factor for propulsion and endurance [11].

Reinforcement learning (RL) offers a flexible alternative for high-dimensional, constraint driven problems. Unlike gradient-based methods, RL learns from interaction and handles noisy evaluations.

This work introduces an RL framework integrating DDPG and Bezier parameterization for airship optimization [10]. The agent modifies control points to maximize $L/D$ while preserving volume. Evaluation is performed using SILCROAD [6,7,8,9].

Main contributions include: (1) reinforcement learning with embedded physical constraints, (2) continuous shape control via Bezier parameters, and (3) performance improvements in $L/D$ and robustness.

## Methodology

The upper hull is modeled with a Bezier curve using $n + 1$ control points $\{P_0, \ldots, P_n\}$. The curve is:

$$R(x) = \sum_{i=0}^{n} B_{i,n}(x) P_i, \qquad B_{i,n}(x) = \binom{n}{i} (1 - x)^{n-i} x^i \tag{1}$$

where x ∈ [0, 1] is the normalized axial coordinate. Endpoints $P_0$, $P_n$ are fixed at zero and the agent adjusts the $n - 1$ interior points.

Volume is maintained via scaling factor $\lambda$:

$$\lambda = \sqrt{\frac{V_0}{V'}}, \qquad P_i^{\text{scaled}} = \lambda P_i, \tag{2}$$

where $V_0$ is the volume of the original hull and $V'$ is the volume of the perturbed shape, both computed via

$$V = \pi \int_0^L R(x)^2 dx \tag{3}$$

Here, $L$ denotes the length of the hull.

*RL formulation.* The state is the set of current radii; the action is a perturbation vector on the interior control points. The reward is:

$$r = \omega_1 \left(\frac{L}{D}\right)_{avg} - \omega_2 \sigma_{\frac{L}{D}} - \omega_3 \|\Delta P\|_1 \tag{4}$$

The weights in (4) trade off performance and regularity. $\omega_1$ scales the primary objective $(L/D)_{avg}$ (suggested range 0.5–1; we use $\omega_1 = 1$ to keep the natural scale and emphasize aerodynamic merit). $\omega_2$ would penalize the variability $\sigma_{L/D}$ across $\alpha$ to promote robustness (typical 0–0.5), but is set to $\omega_2$=0 here because our solver interface returns only the mean $L/D$. $\omega_3$ regularizes shape smoothness via $\|\Delta P\|_1$ (implemented as the mean absolute first difference of control-point radii), for which 0.05–0.2 keeps the penalty one order below $(L/D)_{avg}$; we choose $\omega_3$=0.1 as it curbs oscillatory updates without suppressing beneficial adjustments. Hence, ($\omega_1$, $\omega_2$, $\omega_3$) = (1, 0, 0.1), matching the implementation $reward = AvgLD - 0.1 * smoothPenalty$.

*Learning algorithm.* Training uses MATLAB Reinforcement Learning Toolbox's rlDDPGAgent, the built-in implementation of the Deep Deterministic Policy Gradient (DDPG) algorithm—an off-policy, model-free actor–critic method for continuous-action control with a deterministic policy, experience replay, and target networks. Actor and critic networks are two-layer multilayer perceptrons. At each episode, the agent queries SILCROAD to compute the $L/D$ curve over an angle-of-attack grid spanning 0°–30°; the resulting mean (and, when used, dispersion) of $L/D$ across this range feeds the reward in (4).

## Result and Discussion

Trained over 10,000 episodes, the agent consistently outperforms the baseline.

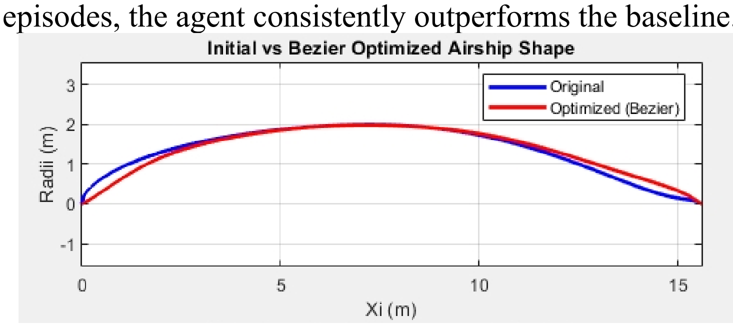


*Figure 1: Original and optimized Bezier-based hull shapes.*

Figure 1 shows original vs optimized shapes. The optimized hull has a smoother profile with gradual
tapering.

The shape preserves volume:

$$V_0 = V_{final} = 107.5273\, m^3, \quad \lambda = 1.16$$

Figure 2 shows the $L/D$ improvement. Average values:

$$\left(\frac{L}{D}\right)_{\text{initial}} = 4.7766, \quad \left(\frac{L}{D}\right)_{\text{optimized}} = 5.3044, \qquad \text{Gain} \ = +11.05\%$$

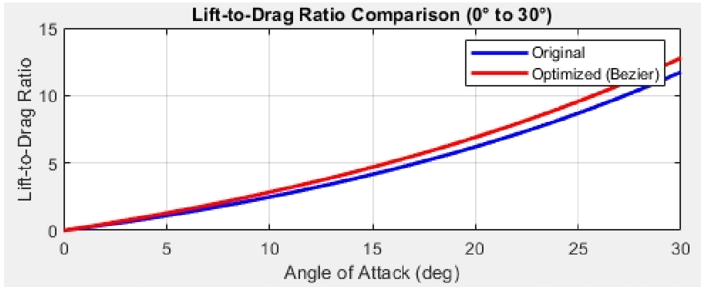


*Figure 2: Lift-to-drag ratio comparison (0°–30° angle of Attack).*

Perturbation tests showed reduced standard deviation in $L/D$. Shape smoothness was retained via $\|\Delta P\|_1$.

**Conclusion**

This study presented a reinforcement-learning-based framework for airship shape optimization. The DDPG agent adjusted Bezier control points while preserving volume. Evaluation in SILCROAD confirmed an 11.05% $L/D$ improvement with retained robustness and smoothness [13]. The framework is modular, supporting future extensions such as tail design and stability control.

**References**

[1] A. K. Shakya, G. Pillai, and S. Chakrabarty, “Reinforcement learning algorithms: A brief survey,” *Expert Systems with Applications*, vol. 231, 2023, Art. no. 120495. https://doi.org/10.1016/j.eswa.2023.120495.

[2] P. Alfeld, M. Neamtu, and L. L. Schumaker, “Circular Bernstein–Bezier polynomials,” in *Mathematical Methods for Curves and Surfaces*, 1995, pp. 11–20.

[3] T. P. Lillicrap et al., “Continuous control with deep reinforcement learning,” arXivpreprint arXiv:1509.02971, 2015

[4] R. T. Jones, Classical Aerodynamic Theory. National Aeronautics and Space Administration, 1979.

[5] S. P. Jones and J. D. DeLaurier, “Aerodynamic estimation techniques for aerostats and airships,” Journal of Aircraft, vol. 20, no. 2, pp. 120–126, 1982.

[6] C. E. D. Riboldi, M. Belan, S. Cacciola, R. Terenzi, S. Trovato, D. Usuelli, and G.Familiari, “Preliminary sizing of a low-altitude airship including ion-plasma thrusters,”in Proc. 34th Congress of the International Council of the Aeronautical Sciences (ICAS),Florence, Italy, Sept. 9–13, 2024.

[7] C. E. D. Riboldi and A. Rolando, “Layout analysis and optimization of airships with thrust-based stability augmentation,” Aerospace, vol. 9, p. 393, 2022.

[8] C. E. D. Riboldi and A. Rolando, “Thrust-based stabilization and guidance for airships without thrust-vectoring,” Aerospace, vol. 10, 2023.

[9] C. E. D. Riboldi, A. Rolando, and G. Regazzoni, “On the feasibility of a launcher-deployable high-altitude airship: Effects of design constraints in an optimal sizing framework,” Aerospace, vol. 9, pp. 1–37, 2022.

[10] A. Kornienko, System Identification Approach for Determining Flight Dynamical Characteristics of an Airship from Flight Data. Ph.D. dissertation, University of Stuttgart, 2006

[11] T. Lutz and S. Wagner, "Drag reduction and shape optimization of airship bodies,"Journal of Aircraft, vol. 35, no. 3, pp. 345–351, 1998.

[12] C. E. D. Riboldi and L. Fanchini, "Assessing the Technical–Economic Feasibility of Low-Altitude Unmanned Airships: Methodology and Comparative Case Studies," Aerospace,vol. 12, no. 3, p. 244, 2025.

[13] C. E. D. Riboldi and M. Tomasoni, "Formation Flight of Fixed-Wing UAVs: Dynamic Modeling, Guidance Design, and Testing in Realistic Scenarios," Aerospace, vol. 12, no. 3, pp. 1–49, 2025.