Differentiable Co-Optimization of Mass-Limited Legged Robots Across Variable Gravity

Anaïs Jeger, Leone Costi, Giacomo Acciarini, Dario Izzo
Advanced Concepts Team, European Space Agency (ESA)
ESTEC, Noordwijk, The Netherlands
anais.jeger@esa.int

***Abstract*—We present a differentiable co-design framework for studying legged locomotion across variable gravitational environments. The method performs end-to-end gradient-based optimization through rigid-body contact dynamics, using NVIDIA Warp's GPU-accelerated automatic differentiation to jointly tune morphology and gait parameters under a fixed 12 kg mass budget. By treating gravity as a first-class design variable, the framework reveals gravity-dependent trade-offs in body-leg mass allocation, limb geometry, and gait coordination. The resulting solutions expose destination-specific design trends for small, mass-limited planetary robots.**

I. Introduction

Space robotics provides a sharp motivation for morphological co-design: once a destination is selected, the robot must work within a narrow, largely irreversible design envelope. Hardware cannot be modified after launch, and on-site tuning opportunities are limited by risk, bandwidth, and operational constraints. A robot deployed on Titan ($g = 1.35\,\mathrm{m/s^2}$) operates in a qualitatively different regime than one on Venus ($g = 8.87\,\mathrm{m/s^2}$), shifting ballistic time scales, stance loading, and stability margins in ways that reorganize which morphologies and gait patterns are feasible. Sweeping gravity is therefore not an attempt to design a universal robot, but a controlled way to expose regime transitions and derive destination-specific design guidance under the same co-optimization pipeline.

Engineering practice often decouples morphology and control: a body plan is fixed first, and a controller is subsequently optimized for it [1]. This sequential approach can miss synergistic solutions that only emerge from joint optimization. Modern differentiable physics simulation addresses this directly: by propagating gradients through simulation dynamics, task performance can be differentiated with respect to both morphological and control parameters simultaneously, enabling efficient co-design over high-dimensional joint design spaces [2]–[4].

II. Method

A. Robot Model and Design Space

We represent the robot as five rigid spheres: a central body sphere of radius r_b and four leg spheres of radii $r_1, \ldots, r_4$ (Fig. 1), with total mass fixed at $M = 12\,\mathrm{kg}$. Each sphere's mass is determined by its volume and a uniform density solved from the constant-mass constraint; the body-to-leg density ratio is fixed at $\rho_b/\rho_\ell = 4.5$ to model the concentration of heavy components (batteries, electronics) in the body. This constraint enforces morphological trade-offs: enlarging the body reduces mass available for legs, and vice versa.

The 11-dimensional design vector $\boldsymbol{\theta} = [\boldsymbol{\theta}_m; \boldsymbol{\theta}_c] \in \mathbb{R}^{11}$ comprises eight morphological parameters $\boldsymbol{\theta}_m = \{r_1, \ldots, r_4, r_b, s_x, s_y, A\}$ (four leg radii, body radius, stance half-dimensions, stride amplitude) and three control parameters $\boldsymbol{\theta}_c = \{\varphi_{\mathrm{FR}}, \varphi_{\mathrm{BL}}, \varphi_{\mathrm{BR}}\}$ (inter-leg phase offsets; the front-left phase is fixed at zero). All bounded parameters are enforced via a smooth reparameterization: each physical parameter θ_j is expressed as a function of an unconstrained internal variable $\tilde{\theta}_j \in \mathbb{R}$,

$$\theta_j = \theta_j^{\min} + \tfrac{1}{2}\big(\tanh(\tilde{\theta}_j) + 1\big)\big(\theta_j^{\max} - \theta_j^{\min}\big), \tag{1}$$

so that optimization proceeds over $\tilde{\theta}_j$ without explicit projection, while θ_j remains within its box constraints and gradients stay continuous.

B. Locomotion Control

Each leg tracks a trajectory generated by a central pattern generator (CPG), a sinusoidal oscillator that produces slope-aligned elliptical foot paths:

$$x_i^t(t) = x_i^0 + A \sin\alpha \sin(\omega t + \varphi_i) \tag{2}$$

$$z_i^t(t) = z_i^0 + A \cos\alpha \cos(\omega t + \varphi_i) \tag{3}$$

where $\alpha = 25^\circ$ is the slope angle, $\omega = 2\pi \times 1.5\,\mathrm{rad/s}$ sets the stepping frequency, and φ_i is the phase offset. The leg actuator is a proportional-derivative (PD) controller ($K_p = 1500\,\mathrm{N/m}$, $K_d = 120\,\mathrm{N{\cdot}s/m}$) with a mass-proportional force cap $\|F_i^{\mathrm{act}}\| \le a_{\max} \cdot m_i$, where $a_{\max} \in \{30, 60, 120\}\,\mathrm{N/kg}$. A leg with radius below $0.07\,\mathrm{m}$ is classified as vestigial.

C. Optimization Objective

The loss function balances forward progression with stability and efficiency:

$$\mathcal{L} = -d_x - \lambda_h\,\bar{z} + \lambda_u\,\overline{e_b^2} + \lambda_v\,\overline{v_{b,z}^2} + \lambda_E\,\bar{E} \tag{4}$$

where $d_x = x_{\mathrm{CoM}}(T) - x_{\mathrm{CoM}}(0)$ is net uphill displacement, $\bar{z}$ is mean center-of-mass height (encourages not sinking into the slope), $\overline{e_b^2}$ penalizes deviation of body height from a $0.34\,\mathrm{m}$ reference (upright posture), $\overline{v_{b,z}^2}$ suppresses vertical bouncing, and $\bar{E} = \sum_i \|F_i^{\mathrm{act}}\|\|v_i\|$ is a mean actuator power proxy. Weights are $\lambda_h = 0.3$, $\lambda_u = 0.4$, $\lambda_v = 0.10$, $\lambda_E = 0.005$.

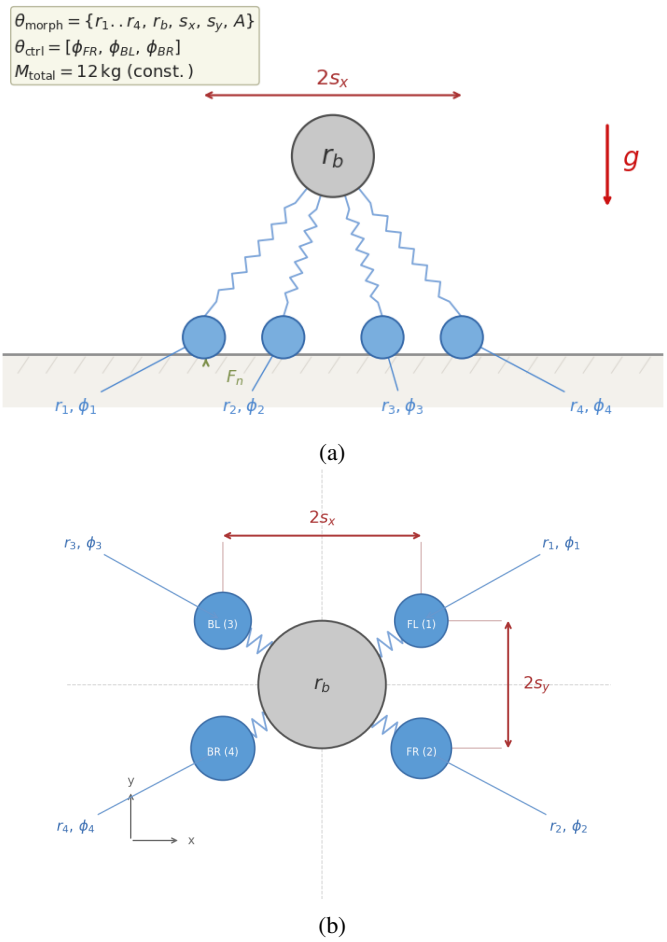

Fig. 1: Schematic of the quadrupedal robot model: (a) side-view showing the body sphere (radius r_b) connected to four leg spheres (radii r_i) via virtual spring-damper actuators on a 25° inclined surface and (b) top-view morphology parameterization showing stance half-distances (s_x, s_y) and leg indexing.

D. Optimization Protocol

We use stochastic gradient descent with separate learning rates for morphology ($\eta_m = 0.02$) and control ($\eta_c = 0.015$) parameters, with a phase gradient multiplier $\alpha_p = 15$ to compensate for the smaller magnitude of phase gradients. Gaussian exploration noise decays linearly over the first 30% of iterations. The simulator uses NVIDIA Warp [4] for forward dynamics ($T = 2000$ steps at $\Delta t = 1/480\,s$, subdivided into 4 substeps; 4.17 s simulated time), with gradients computed via reverse-mode automatic differentiation through the full trajectory.

The benchmark spans nine gravitational environments at three force caps, with 20 or 50 randomly initialized seeds per gravity level (Table I), yielding 900 total trials.

A survivor is a seed whose optimized robot achieves $d_x > 0.5\,m$ without excessive body-ground interpenetration or limb self-collision; the survival rate is the fraction of such seeds per gravity level.

TABLE I: Gravitational Environments

Environment	g ($m\,s^{-2}$)	Earth ratio	Seeds
Titan	1.35	0.14	20
Moon	1.62	0.17	20
g2.5	2.50	0.25	20
Mars	3.71	0.38	20
g5.0	5.00	0.51	20
g6.5	6.50	0.66	50
g8.0	8.00	0.82	50
Venus	8.87	0.90	50
Earth	9.81	1.00	50

III. Results

A. Body Mass Redistribution

The body mass fraction decreases monotonically with gravity across all three force caps. At 60 N/kg, the mean body radius among survivors drops from 26.8 cm (Titan) to 18.7 cm (Earth), corresponding to a body mass fraction fall from 84% to 50%. At 120 N/kg the shallower range (84% to 67%) reflects that stronger actuators reduce the imperative to grow legs. The proposed design tool thus increases the practically available force by delocalizing mass away from the body, as visualized in Fig. 2.

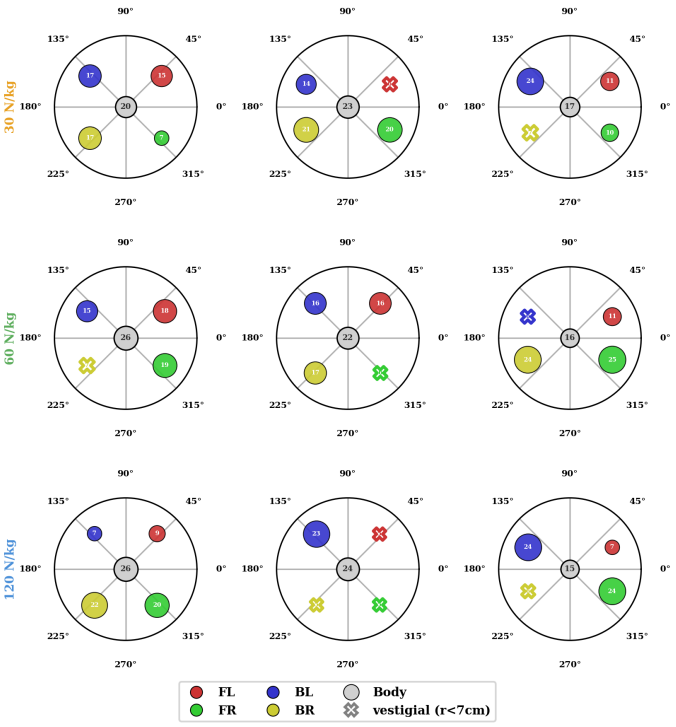

Fig. 2: Best-run morphology (top-down view, sphere size $\propto$ radius) for Moon, Mars, and Earth (left to right), at 30, 60, and 120 $N\,kg^{-1}$ (top to bottom). Hollow crosses mark vestigial legs ($r < 7$ cm).

B. Emergent Vestigial Legs

Not all legs benefit equally from mass redistribution. While some legs increase in size, others shrink to become vestigial

(Fig. 3). The appearance of vestigial legs is strongly gravity-dependent, reaching a plateau at around 8 m/s^2. Contact-load statistics confirm these legs become dynamically inactive: they contribute negligible ground reaction force during locomotion, suggesting an intentional optimizer strategy in which mass is freed and reallocated to the remaining limbs, concentrating effective actuation and support.

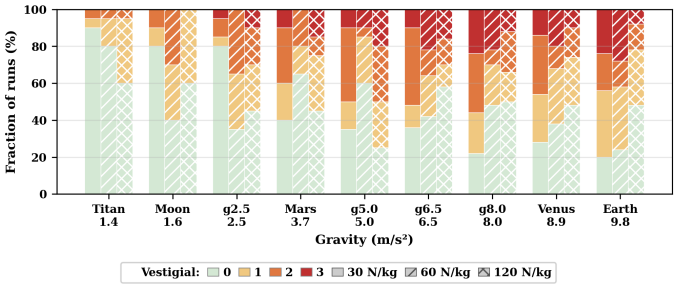

Fig. 3: Fraction of runs exhibiting vestigial legs (radius <7 cm) as a function of environmental gravity and force cap.

C. Gravity-Driven Gait Transitions

We classify gaits by inter-leg phase relationships into three classes: *walk* (sequential ground contacts), *trot/pace* (diagonal or lateral pairs in antiphase), and *aerial* (phase clustering indicating simultaneous flight phases where all feet leave the ground). Classification is performed post-hoc from optimized phase offsets.

At 30 N/kg, aerial patterns account for 65% of Moon survivors, declining to 0% at Earth where walk coordination takes over entirely (Fig. 4). At 120 N/kg, aerial gaits persist above 24% even at Earth, reflecting the optimizer's continued ability to generate ground-clearance forces with stronger actuators. This transition from aerial to ground-bound gaits under increasing gravity is consistent with partial-gravity human locomotion studies [5], [6].

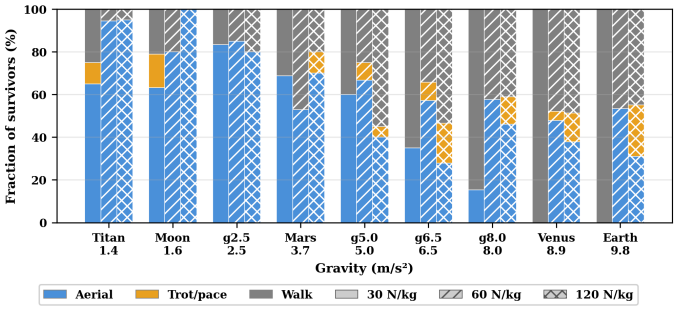

Fig. 4: Fraction of survivors exhibiting each gait class across gravity levels and three actuator force caps (30, 60, 120 N/kg).

D. Morphology–Control Coupling

Gradient analysis shows that morphology-related gradient magnitudes are comparable to, and in some cases larger than, those of the control parameters across all gravity levels (Fig. 5). This indicates that the loss is sensitive to morphological variation, justifying the treatment of morphology as a variable to be optimized jointly with control rather than fixed a priori. In contrast, the stance parameters (s_x, s_y) exhibit near-zero gradients in all environments, consistent with the sagittal-plane restriction of the locomotion task.

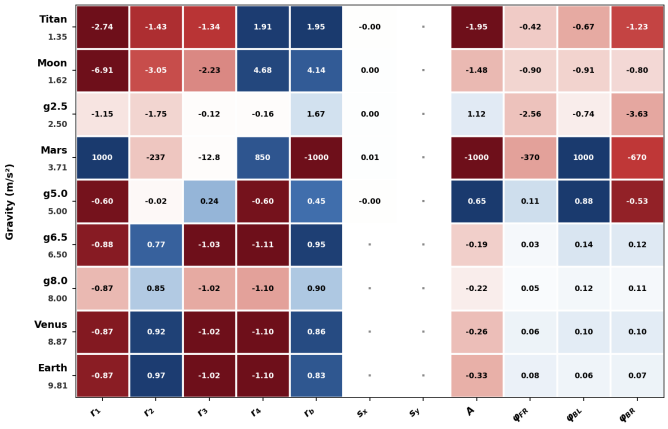

Fig. 5: Gradient magnitudes of the loss with respect to morphology (r_1–r_4, r_b, s_x, s_y, A) and control parameters (φ_{FR}, φ_{BL}, φ_{BR}) at epoch 1, across nine gravity levels. Each row is normalized by the 95$^{\mathrm{th}}$ percentile of its absolute gradient magnitude; color encodes sign and relative magnitude (red: negative, blue: positive, white: near zero).

IV. Conclusion

By embedding morphology parameters and gait coordination directly inside a long-horizon rigid-body contact simulation, gradient-based co-optimization discovers qualitatively distinct locomotion solutions across environments without manual switching between design and control phases. The experiments uncover an emergent vestigial-leg mechanism under high gravity and tight force caps, and gait statistics shift from aerial coordination to walking as loading increases, reinforcing that gravity and actuator authority jointly reshape the optimal locomotion regime.

The five-sphere model intentionally abstracts away joint kinematics to isolate the effect of mass distribution, gait coordination, and gravity on locomotion. The mass-proportional force cap is a first-order approximation of motor scaling that does not capture non-linear effects such as torque density variation with motor size and gearing, and the evaluation is limited to uphill displacement on a single slope. Bridging to articulated hardware, i.e. mapping the optimized mass distributions and phase relationships onto kinematic chains with realistic joint limits and motor torque curves, remains the subject of future work.

REFERENCES

[1] A. Rosendo, M. von Atzigen, and F. Iida, "The trade-off between morphology and control in the co-optimized design of robots," *PLOS ONE*, vol. 12, no. 10, Art. no. e0186107, Oct. 2017, doi: 10.1371/journal.pone.0186107.

[2] Y. Hu, L. Anderson, T.-M. Li, Q. Sun, N. Carr, J. Ragan-Kelley, and F. Durand, "DiffTaichi: Differentiable programming for physical simulation," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2020.

[3] C. D. Freeman, E. Frey, A. Raichuk, S. Girgin, I. Mordatch, and O. Bachem, "Brax—A differentiable physics engine for large scale rigid body simulation," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS) Datasets and Benchmarks Track*, 2021.

[4] NVIDIA, "Warp: A Python framework for high-performance GPU simulation and graphics," 2022. [Online]. Available: https://github.com/NVIDIA/warp

[5] D. J. Newman and H. L. Alexander, "Human locomotion and workload for simulated lunar and Martian environments," *Acta Astronaut.*, vol. 29, no. 8, pp. 613–620, Aug. 1993, doi: 10.1016/0094-5765(93)90078-B.

[6] M. Keime *et al.*, "How about running on Mars? Influence of sensorimotor coherence on running and spatial perception in simulated reduced gravity," *Front. Physiol.*, vol. 14, Art. no. 1201253, Jul. 2023, doi: 10.3389/fphys.2023.1201253.