

LEGGED ROBOTS

Extreme dynamic symmetry enables omnidirectional and multifunctional robots

Jiaxun Liu^{1†}, Boxi Xia^{1†}, Boyuan Chen^{1,2,3*}

Symmetry is a central organizing principle in natural systems, yet its use as a unifying design strategy in robotics has largely remained limited to geometric form. We show that symmetry can instead be leveraged at the level of dynamic actuation capability. We introduce dynamic symmetry, the uniformity of a robot's attainable center-of-mass accelerations, and formalize it through a measure coined as dynamic isotropy. Across more than 1000 simulated morphologies, we found that higher dynamic symmetry consistently improved trajectory tracking, task success, robustness, resiliency, and energy efficiency, with the benefits becoming most pronounced as dynamic isotropy approached its theoretical limit. To study this regime systematically, we developed Argus, a family of spherical robots designed to explore the effects of increasing dynamic symmetry. Members of the Argus family vary in their actuation geometry and dynamic symmetry level while sharing a common architectural principle: radially oriented linear actuators that directly shape the robot's center-of-mass dynamics. Among them, we built a physical 20-leg Argus variant that achieved near-extreme dynamic isotropy and demonstrated orientation-invariant locomotion, agile traversal of cluttered and deformable terrain, rapid self-stabilization, and resilience to partial actuator failures. Its distributed sensing further enabled omnidirectional perception and object interaction during continuous motion. These results show that designing robots for symmetry not only in morphology but also in their attainable dynamics provides a powerful and general pathway toward agility, robustness, and multifunctionality in uncertain terrestrial and extraterrestrial environments.

INTRODUCTION

Since the term “robot” was first coined more than a century ago (1), the design of robotic systems has remained a central scientific and engineering challenge. Advances in actuation, sensing, and computation have been matched by a rich diversity of mechanical designs, many of which draw inspiration from biology (2–6). Quadrupedal robots echo the musculoskeletal architecture of dogs (7–9); humanoid robots (10–13) emulate human proportions and bipedal locomotion; articulated manipulators mimic the dexterity of arms and hands (14, 15); bioinspired underwater robots (16–18), flapping-wing drones (19, 20), and even simplified mechanisms such as parallel grippers inherit key features from their natural counterparts (21). The unifying theme in these approaches is imitation, aiming to reproduce the form and function of biological systems and then refining the designs to meet real-world constraints.

Here, we explore a different way of learning from nature: extracting and applying fundamental organizational principles that transcend specific species or morphologies. One such principle is symmetry, a pervasive feature of living systems across scales and taxa (22, 23). From the bilateral forms of vertebrates and the radial symmetry of starfish to the helical arrangements of plant seeds and the geometric regularity of viruses, symmetry is deeply entwined with both structure and perception. In animals, control laws often exploit these symmetries to simplify coordination of whole-body actuations and gain robustness under diverse environmental conditions.

Inspired by these ideas in nature, symmetry has long played a central role in how roboticists design body configurations that move

and interact with the world. Most prior work has emphasized morphological symmetry, specifically the geometric arrangement of limbs, links, or actuators. Recent studies have shown that symmetric bodies often help improve learning efficiency (24–28), simplify dynamics and control (29, 30), or enhance maneuverability (31). Examples range from legged robots with bilateral symmetry, such as biped and quadruped robots, to spherical (32–35) and tensegrity platforms (36–42) whose shapes exhibit higher-order geometric repetition. Despite this rich history, these approaches still remain grounded in the geometry of the robot's body rather than in the forces and accelerations the robot can produce.

This gap motivates a shift in perspective: from symmetry of form to symmetry of dynamic actuation capability. We refer to this idea as “dynamic symmetry,” the ability of a robot to generate forces and accelerations with uniform magnitude in all directions. Where morphological symmetry ensures that no spatial orientation is privileged, dynamic symmetry ensures that no direction of action is privileged. This can enable a fundamentally different suite of behaviors: robust omnidirectional locomotion, rapid disturbance rejection from arbitrary angles, resilient recovery from broken actuators or sensors, seamless transition between tasks, and natural multifunctionality in complex environments.

To quantify dynamic symmetry, we introduce a theoretical construct for rigid-body legged robots that we call dynamic isotropy. Classical isotropy theory in robotics, such as manipulability-based isotropy and Jacobian condition number analyses (43–48), is purely kinematic and applies only to end-effector velocity mappings of serial manipulators. These metrics characterize how uniformly joint motions map to Cartesian velocities, but they do not incorporate actuator force limits, mass properties, thrust directions, or the geometry of whole-body dynamic actuation capability. Although previous dynamic isotropy (46, 49) and dynamic manipulability ellipsoids (50) introduce inertia, they remain local and

¹Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC, USA. ²Department of Electrical and Computer Engineering, Duke University, Durham, NC, USA. ³Department of Computer Science, Duke University, Durham, NC, USA.

†These authors contributed equally to this work.

*Corresponding author. Email: boyuan.chen@duke.edu

end effector centric, capturing neither full-body acceleration feasibility nor global actuation symmetry.

In contrast, we define dynamic isotropy directly on the robot's attainable center-of-mass (CoM) acceleration set. It measures how isotropic the feasible acceleration space becomes as the robot's actuation directions increase in number, uniformity, and spatial distribution. This yields a global characterization of whole-body dynamics, which reflects how evenly the robot can act on the environment rather than being limited to how its joint motions map to velocity. Dynamic isotropy thus provides a principled way to study and design dynamic symmetry as an axis of robot capability.

In this article, we investigate extreme dynamic symmetry as a design principle for omnidirectional and multifunctional rigid-legged robots (Movie 1). Our central hypothesis is that, as dynamic isotropy approaches its theoretical limit, a robot gains enhanced dynamic actuation capabilities, enabling it to become an omnidirectional and multifunctional machine. We tested this hypothesis through both theoretical and empirical studies. First, we derive and analyze dynamic isotropy as a performance measure grounded in whole-body dynamics and control. We then introduce Argus as the experimental platform, a class of spherical robot designs (fig. S2) that feature one-degree-of-freedom linear actuated legs (Fig. 1A) oriented toward the robot's center and directly alter its CoM dynamics. With Argus, we conducted large-scale simulation sweeps across thousands of morphologies to investigate how increasing dynamic symmetry alters the attainable acceleration set and enhances robustness, resiliency, and task performance. Last, we realized a physical Argus variant with 20 legs and 20 cameras at its foot nodes (Fig. 1B) with near-extreme dynamic symmetry that demonstrated capabilities exceeding those of prior spherical and tensegrity platforms across locomotion, robustness, agility, resilience, and multifunctionality (Fig. 1C). Beyond locomotion, we show that extreme dynamic symmetry pairs naturally with omnidirectional perception (Fig. 1C): The robot can maintain omnidirectional perception under aggressive maneuvers. Together, dynamic isotropy and omnidirectional sensing produce a versatile, multifunctional system capable of seamless transitions between locomotion,

reorientation, environment interaction, and perception-driven tasks such as whole-body loco-manipulation. Overall, our work establishes dynamic isotropy as a theoretical framework for whole-body robotic capability and demonstrates that pushing dynamic symmetry toward its limit unlocks versatile behaviors in mobility, robustness, perception, and multifunctionality.

RESULTS

Dynamic symmetry and dynamic isotropy

Robots differ not only in appearance but also in how they can interact with the environment in terms of their dynamic actuation capabilities. We define dynamic symmetry as the uniformity of a robot's attainable CoM accelerations: A dynamically symmetric robot can accelerate itself with uniform authority in many, ideally all, directions.

Formally, we model the relationship between joint torques $\boldsymbol{\tau}$ and CoM accelerations \mathbf{a}_c as

$$\mathbf{a}_c = \mathbf{A}(\mathbf{q}) \boldsymbol{\tau} \quad (1)$$


where $\mathbf{A}(\mathbf{q})$ is a configuration-dependent matrix that incorporates actuation directions, actuator limits, and the robot's mass distribution (see the "Dynamic symmetry and dynamic isotropy" and "Theoretical analysis of dynamic isotropy for stability, robustness, and control efficiency" sections for the full derivation). For a unit direction $\mathbf{u} \in \mathbb{S}^2$, we define $a_{\max}(\mathbf{u})$ as the maximum CoM acceleration attainable along \mathbf{u} under actuator constraints. Sampling many directions on the unit sphere yields a discrete approximation of the robot's attainable acceleration set.

We introduce dynamic isotropy η as

$$\eta = \frac{a_{\min}}{a_{\max}} \quad (2)$$

where a_{\min} and a_{\max} are the minimum and maximum values of $a_{\max}(\mathbf{u})$ across all sampled directions. When $\eta \rightarrow 1$, the robot can accelerate almost equally well in every direction; when η is small, certain directions are much harder to actuate than others.

Figure 2 illustrates dynamic isotropy among different robots and joint configurations. We present the detailed description and analysis of the dynamic isotropy of different robots in the Supplementary Methods section. We also present a more detailed analysis of the robot's dynamic isotropy under different configurations during forward motion in fig. S1. These examples also demonstrate that dynamic symmetry is a distinct property that is not guaranteed by geometric symmetry alone. For each system, we visualized the distribution of attainable acceleration magnitudes over 2048 uniformly sampled directions. Robots with low dynamic isotropy exhibited lopsided, streaked acceleration clouds with large variations in magnitude, indicating pronounced directional preferences. In contrast, robots with higher



Movie 1. Overview of Argus. Overview of the omnidirectional and multifunctional capabilities of the Argus robot through extreme dynamic symmetry.

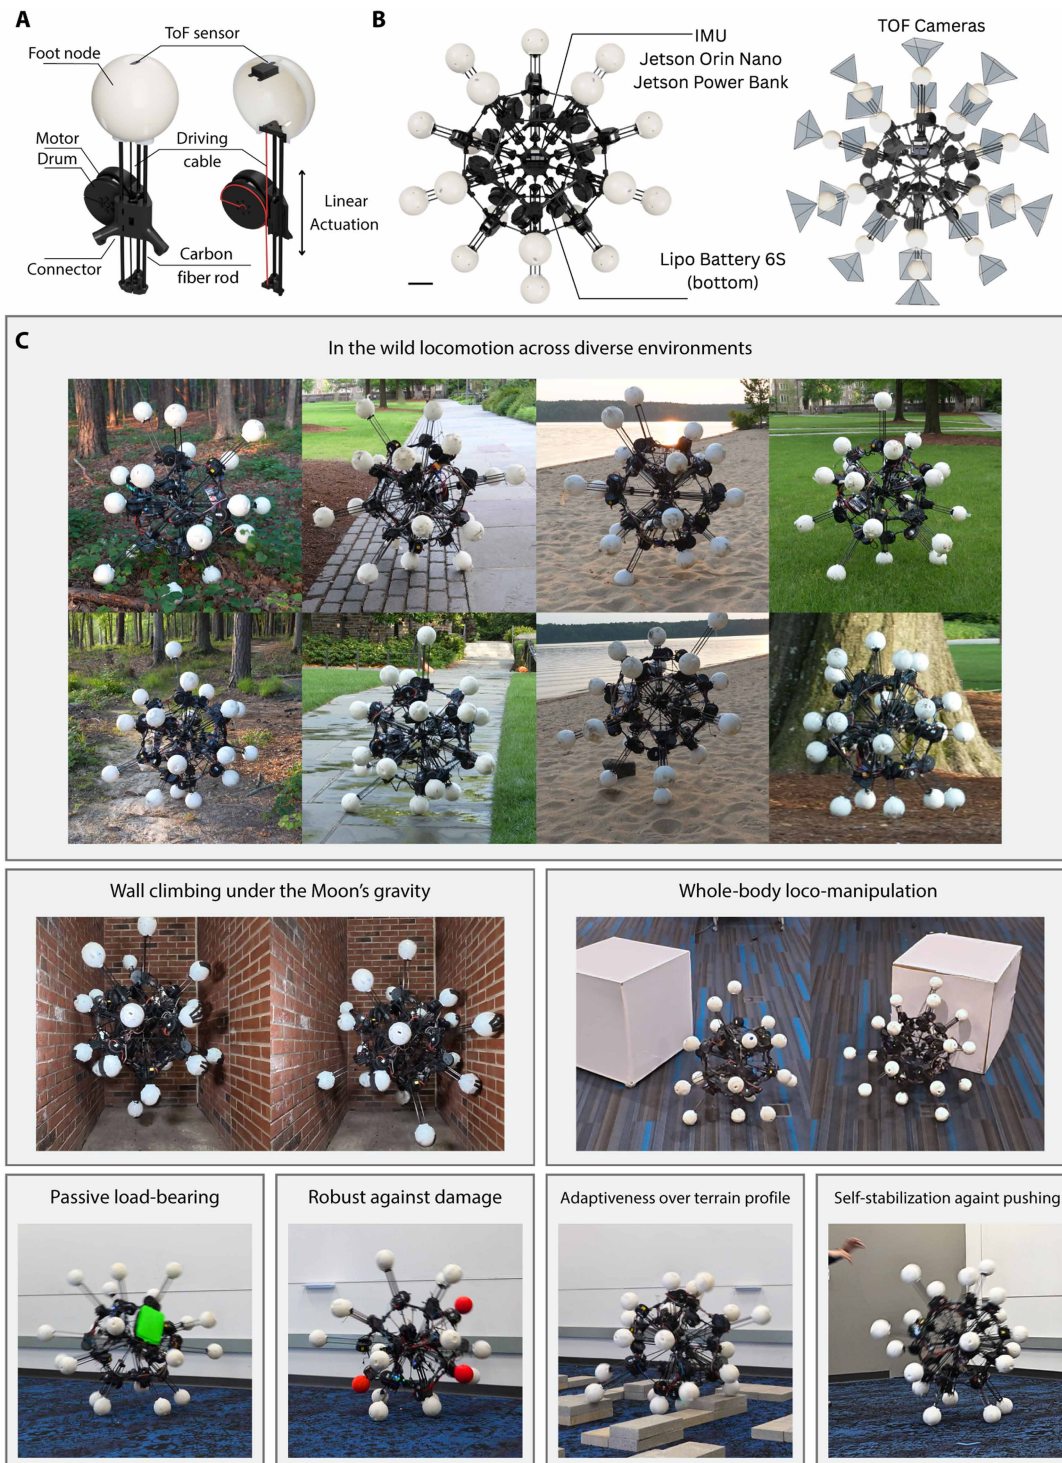


Fig. 1. Argus: A dynamically symmetric robot family and its 20-leg hardware realization. (A) Each leg module is identical and consists of a radially oriented one-degree-of-freedom linear actuator driven by a cable-drum mechanism. A ToF depth sensor is embedded in each foot node, enabling omnidirectional perception collocated with actuation. (B) The 20-leg Argus prototype arranges its actuators at the vertices of a regular dodecahedron, producing a near-spherical distribution of actuation directions and achieving near-extreme dynamic isotropy of 0.91. Computation, inertial sensing, and power modules are mounted at the center of the structure, whereas 20 outward-facing ToF cameras collectively provide omnidirectional depth measurements for perception and interaction. Scale bar, 0.1 m. (C) The physical robot operating across diverse indoor and outdoor environments—including grass, sand, bark, pavement, narrow corridors, and whole-body loco-manipulation settings—demonstrating that extreme dynamic symmetry enables robust omnidirectional locomotion and multifunctional behaviors in unstructured environments.

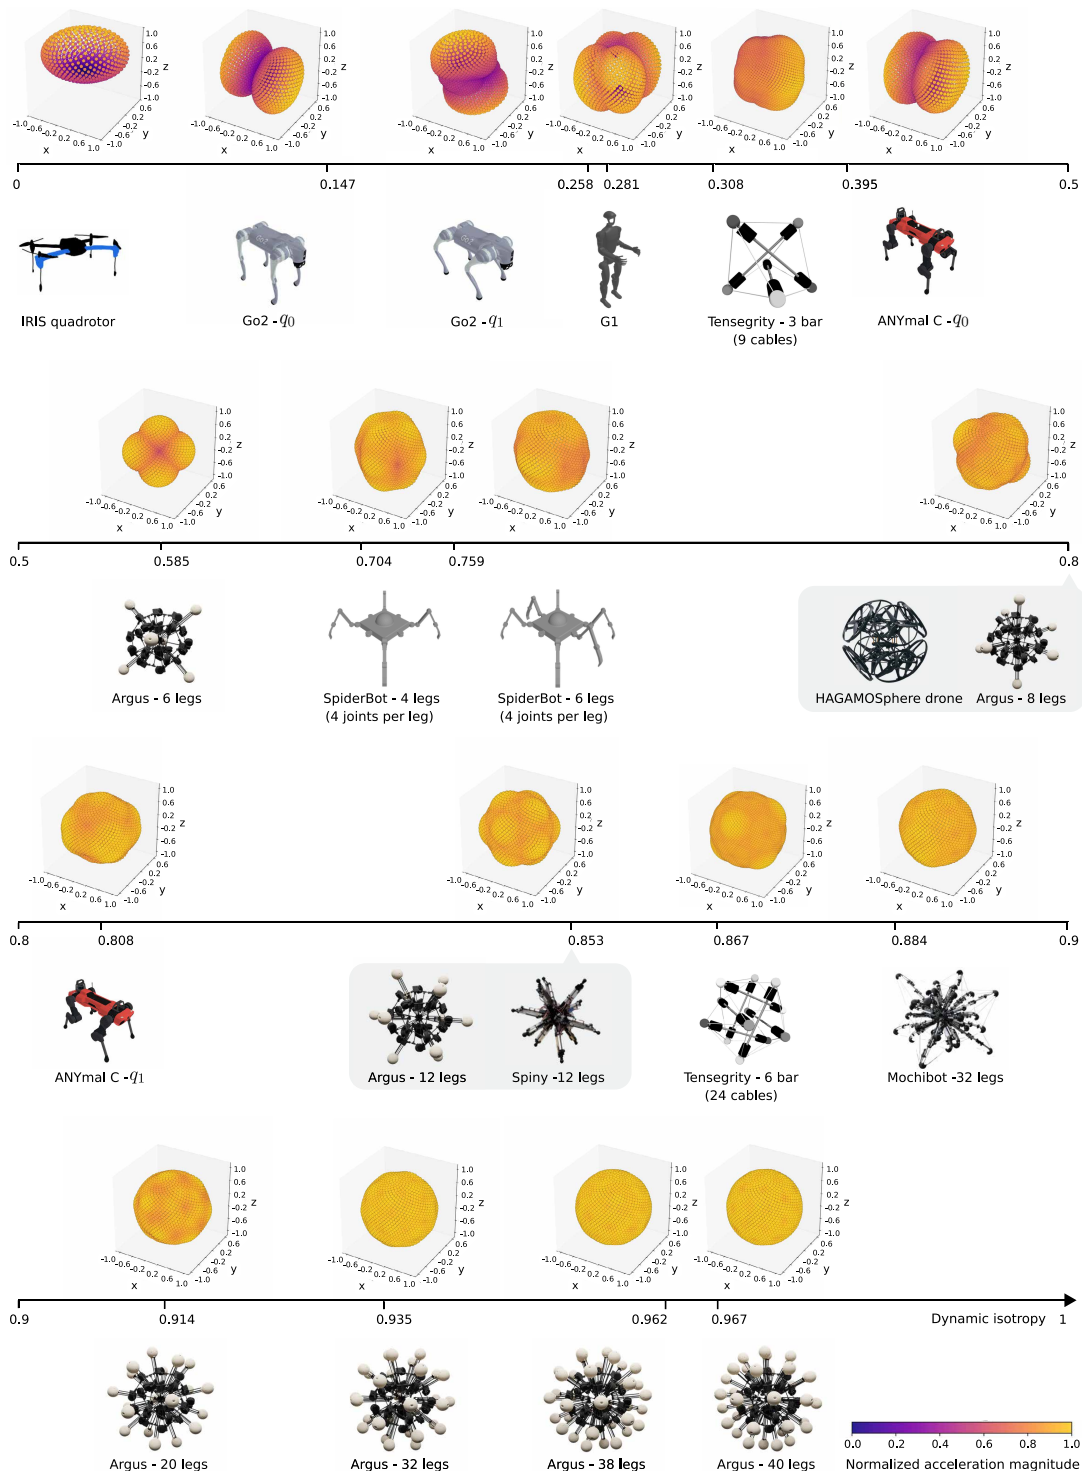


Fig. 2. Dynamic isotropy across different robots and joint configurations. These robots include the IRIS quadrotor, Unitree Go2, Unitree G1, tensegrity robots (42), ANYmal C, spider robots (62), HAGAMOSphere (research prototype, DIC Corporation), Spiny (53), and Mochibot (33). For each robot, the 3D scatter plot shows attainable CoM acceleration magnitudes sampled over 2048 uniformly distributed directions on the unit sphere. Each point lies along a direction \mathbf{u} , with radial distance and color indicating the normalized maximum acceleration magnitude $a_{\max}(\mathbf{u})$ relative to the global maximum. Systems with low dynamic isotropy exhibited highly lopsided, streaked acceleration clouds, whereas systems with high dynamic isotropy displayed nearly spherical, uniformly bright acceleration clouds. Most of the robot designs exhibited dynamic isotropy lower than 0.9. In contrast, the Argus design can achieve a higher dynamic isotropy score, approximating the maximum value of 1 in the last row.

CREDITS: IRIS QUADROTOR, UNITREE GO2 AND G1, AND TENSEGRITY ROBOTS (42); ANYMAL C, ANYBOTICS; HAGAMOSHERE, PROTOTYPE CURRENTLY UNDER RESEARCH AND DEVELOPMENT, USED WITH PERMISSION FROM DIC CORPORATION; SPIDERBOTS REPRINTED WITH PERMISSION FROM (62); SPINY REPRINTED WITH PERMISSION FROM (53); IEEE; MOCHIBOT REPRINTED WITH PERMISSION FROM (33), IEEE.

Downloaded from <https://www.science.org> on May 27, 2026

dynamic isotropy showed nearly spherical, uniformly colored clouds, reflecting an almost direction-agnostic ability to accelerate the CoM.

Certain scores may appear unintuitive, such as the low dynamic isotropy of conventional quadrotor drones. This outcome arises because dynamic isotropy quantifies the uniformity of instantaneously attainable linear accelerations of the CoM, which is different from overall maneuverability. Although quadrotors can execute agile three-dimensional (3D) motions through attitude reorientation, their coplanar rotor configuration restricts certain directions of instantaneous thrust generation. As a result, lateral acceleration authority is limited without prior reorientation, leading to reduced dynamic isotropy. As seen in our plot, designs such as spherical drones achieve higher dynamic isotropy than conventional quadrotors.

Overall, these results demonstrate that the uniformity of instantaneous attainable CoM accelerations can be consistently characterized across diverse robots, including aerial robots, tensegrity systems, legged robots, and humanoids, using dynamic isotropy as a unified theoretical framework. As discussed in the next section, we further focus on understanding the benefits for legged robotic systems operating near extreme dynamic isotropy.

Our visualizations highlight that dynamic symmetry is a property of force and acceleration capability, distinct from geometric symmetry alone. Moreover, our theoretical analysis shows that high dynamic isotropy ensures that the robot's acceleration map is well conditioned with orientation-invariant stability margins, uniform robustness to external disturbances, and balanced control efforts (see the "Theoretical analysis of dynamic isotropy for stability, robustness, and control efficiency" section in Materials and Methods).

Argus: A family of robot designs to study dynamic symmetry

To systematically explore extreme dynamic symmetry, we developed a family of spherical robots named Argus. Each Argus variant consists of one-degree-of-freedom linear legs mounted on a spherical frame and oriented radially toward the robot's CoM. Because each leg can only apply force along its own axis, the spatial distribution of leg directions directly shapes the attainable acceleration set: Adding more legs increases actuation density, and distributing them more uniformly over the sphere increases dynamic isotropy.

We generated Argus morphologies by varying both the number of legs and their placement. Two complementary construction strategies span a wide range of dynamic isotropy: Thomson energy (51) optimization, which distributes legs on a sphere by minimizing the sum of inverse pairwise distances, producing near-uniform geometric layouts, and random sampling over a sphere to explore asymmetric morphologies. Across these designs, dynamic isotropy values range from 0.32 to 0.97 (fig. S2), enabling us to systematically study dynamic symmetry from low to near-extreme regimes.

Figure 3 compares symmetric Argus variants with different numbers of legs. As the leg count increases from 6 to 40, the attainable acceleration clouds become progressively more spherical and uniform, and the dynamic isotropy score η rises accordingly. However, this increase exhibits a plateau: Beyond ~16 to 22 legs, additional actuators yield only marginal gains in dynamic isotropy. This trend suggests a region of diminishing returns, where the robot is already close to the theoretical isotropy limit imposed by its actuation model. Although configurations with more legs have higher redundancy, their dynamic isotropy remains relatively similar, suggesting that dynamic isotropy captures effects not reflected by redundancy alone

and provides an additional criterion for selecting an appropriate leg count to achieve strong performance.

For the physical prototype, manufacturability and space for electronics impose additional constraints. Polyhedral shapes with near-extreme dynamic symmetry are ideal for physical robot construction because the edges can provide stiff support for each actuator and the faces can house electronics and wiring. We therefore built a 20-leg Argus robot by finding the closest dodecahedral configuration with a near-extreme dynamic isotropy score of 0.91. The design is modular using mounting identical cable-driven linear actuators at the 20 vertices of a dodecahedron (see the "Mechanical design" section in Materials and Methods).

Dynamic symmetry versus performance in large-scale simulations

Our key hypothesis is that robots with higher dynamic symmetry should perform better across diverse tasks, especially as isotropy approaches its theoretical limit. To test this, we conducted large-scale robot learning experiments in simulation, evaluating Argus variants in locomotion and interaction tasks under varying dynamic isotropy.

We considered four representative tasks: velocity tracking on flat ground, locomotion with 10% of legs disabled, locomotion while carrying payloads between 0 and 20 kg, and traversal over discrete obstacle terrain with obstacles up to 10 cm high. For each morphology, we trained a locomotion policy using deep reinforcement learning and evaluated trajectory tracking error, task success rate, and cost of transport (COT) over 5-s trajectories at a commanded speed of $v_{\text{cmd}} = 0.8$ m/s (see the "Policy training, tasks, and evaluations" section in Materials and Methods).

We first analyzed 18 symmetric Argus variants with leg counts between 6 and 40, each constructed via Thomson energy minimization to yield highly uniform geometric layouts. As shown in Fig. 4, performance improved with the number of legs across all four tasks: Tracking errors decreased, success rates increased, and COT dropped. Notably, the onset and plateau of these performance gains closely tracked the corresponding rise and saturation of dynamic isotropy in Fig. 3. Beyond around 16 to 22 legs, both isotropy and task performance showed diminishing returns, indicating that pushing isotropy toward its limit offers substantial benefits up to the point where the attainable acceleration field is already nearly isotropic.

Spatial redundancy has been studied as a key mechanism for reliable legged robot locomotion. To study the relationship between our dynamic symmetry with spatial redundancy (in our case, the number of legs of Argus variants), we next generated 512 morphological variants for each of the 12-, 20-, and 32-leg configurations, for a total of 1536 Argus variants. Leg directions were sampled to span a broad range of dynamic isotropy values between 0.25 and 0.93 while keeping the number of legs fixed. Representative configurations are visualized in fig. S2.

Figure 5 shows the resulting performance as a function of dynamic isotropy. Across all three leg counts and all four tasks, higher isotropy consistently correlated with lower tracking errors, higher success rates, and reduced COT. The fully symmetric Argus designs (labeled points) lie near the Pareto frontier of this dynamic isotropy-performance space, confirming that extreme dynamic symmetry brings a functional advantage.

We further analyzed the role of redundancy by comparing performance trends across leg counts (fig. S3). In flat-ground and payload-carrying tasks, additional legs provided benefits in the mid-range of

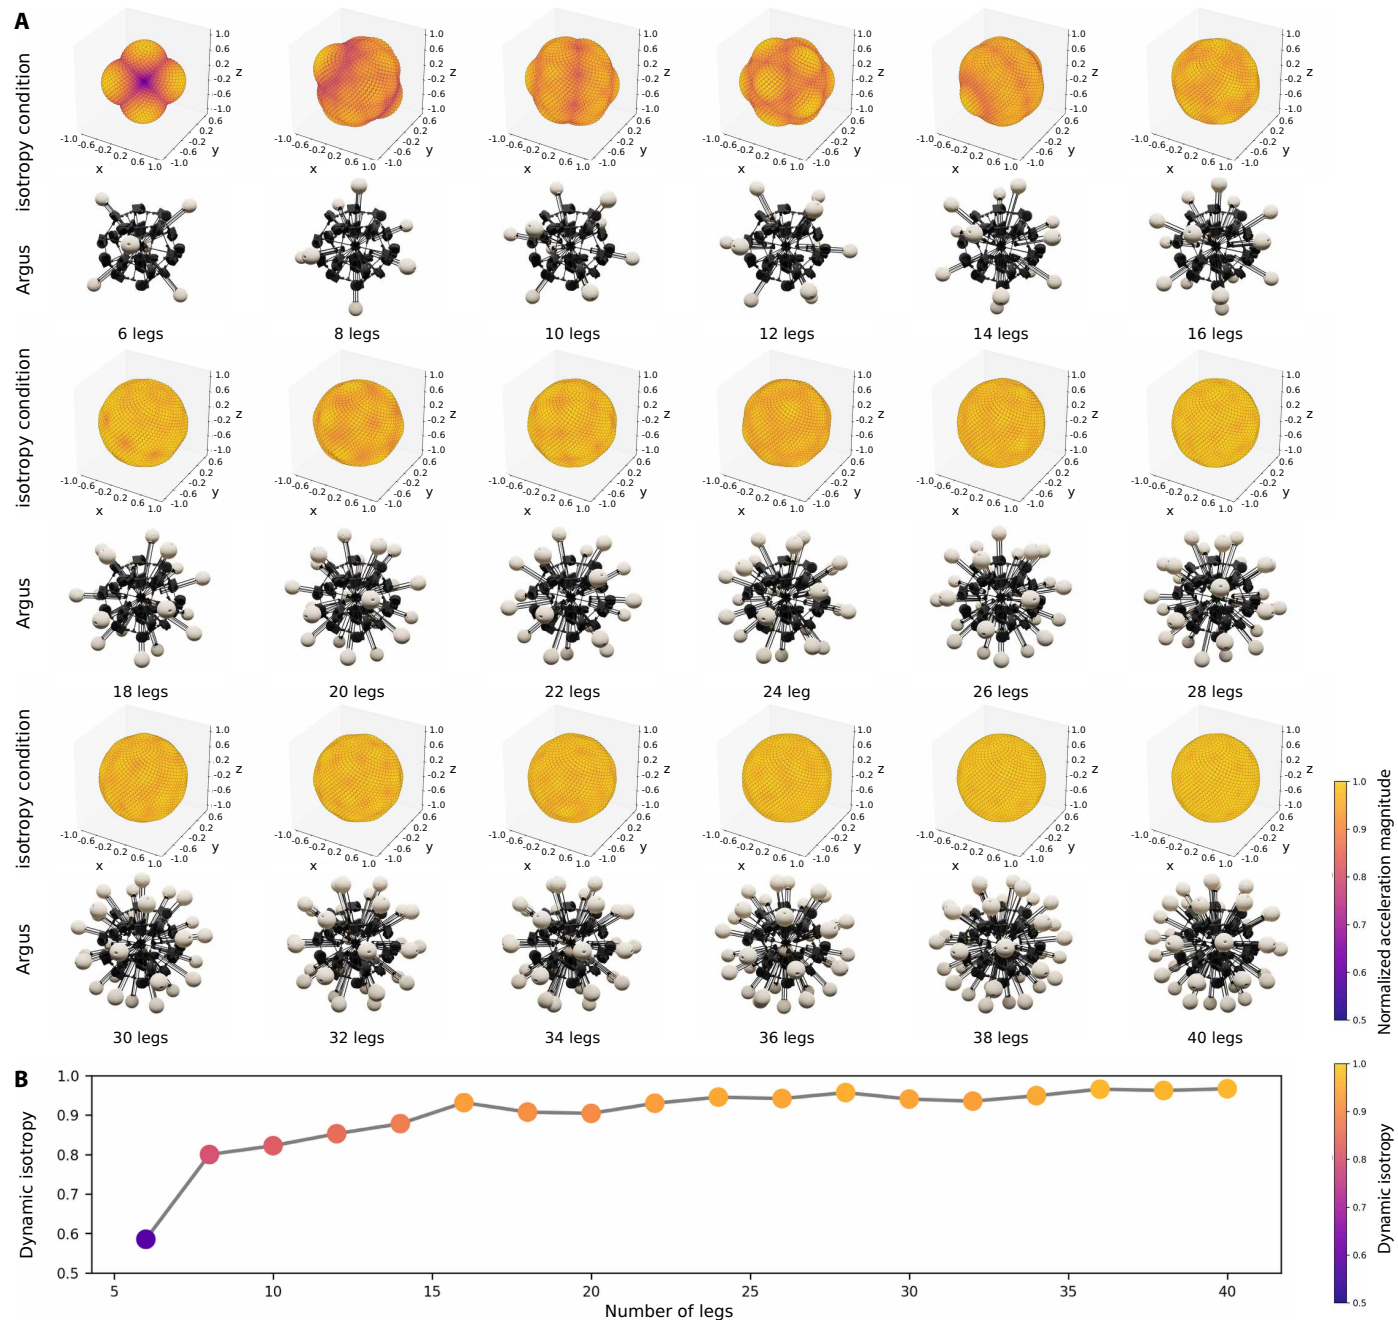


Fig. 3. Dynamic isotropy in symmetric Argus variants with different numbers of legs. (A) Visualization of attainable CoM acceleration magnitudes for symmetric Argus variants with increasing numbers of radially oriented legs. Each cloud was constructed from 2048 uniformly sampled directions; point radius and color indicate normalized maximum acceleration magnitude along that direction. As leg count increased, the attainable acceleration clouds became more uniform and spherical. **(B)** Dynamic isotropy score η for the corresponding variants. Increasing the number of legs monotonically raises η but with diminishing returns: Isotropy began to plateau once the number of legs exceeded roughly 16 to 22, indicating that these designs are already close to the theoretical limit of uniform CoM acceleration for this actuation model.

isotropy ($0.5 \lesssim \eta \lesssim 0.9$), but their marginal advantage diminished as isotropy approached its extreme. Redundancy remained beneficial for leg-disabling scenarios, where extra legs provided alternative contact configurations that allowed locomotion to be preserved despite hardware failures. On discrete terrain, the number of legs played a smaller role, whereas dynamic isotropy continued to be a strong predictor of success, reflecting the importance of uniform acceleration

capability in environments that demand frequent reconfiguration and contact changes.

Agile omnidirectional locomotion over complex terrains in the physical Argus

We trained the locomotion policy using deep reinforcement learning in a large-scale distributed simulation (see the “Policy training,




Fig. 4. Performance scaling with dynamic isotropy in symmetric Argus variants. For 18 symmetric Argus variants with 6 to 40 legs, we plotted trajectory tracking error and COT across all four tasks. As dynamic isotropy rose, both tracking error and COT decreased, whereas task success rates increased. These improvements saturated between roughly 12 and 22 legs, closely aligning with the regime where dynamic isotropy itself begins to plateau (around 16 to 22 legs).

tasks, and evaluations” section in Materials and Methods). In randomized velocity tracking tasks with commanded speeds from -0.8 to 0.8 m/s along both x and y axes, Argus achieved an average tracking error of 0.35 m/s, with a maximum error of 1.13 m/s. When deployed on the physical robot directly, the same policy followed user-commanded velocities and trajectories, including sharp 90° turns (movie S1) and arbitrary path tracking (Fig. 6A and movie S1). This agility arises from real-time redistribution of contact forces across isotropic leg configurations, which counteract momentum and maintain balance without requiring explicit body reorientation for alignment. Despite being trained only on flat terrain, the policy transferred zero-shot to unstructured environments, including concrete bricks, grass, dense foliage with high vegetation, soft sand, and slippery wet surfaces (Fig. 1C and movies S6 and S7). Stability was preserved under orientation disturbances

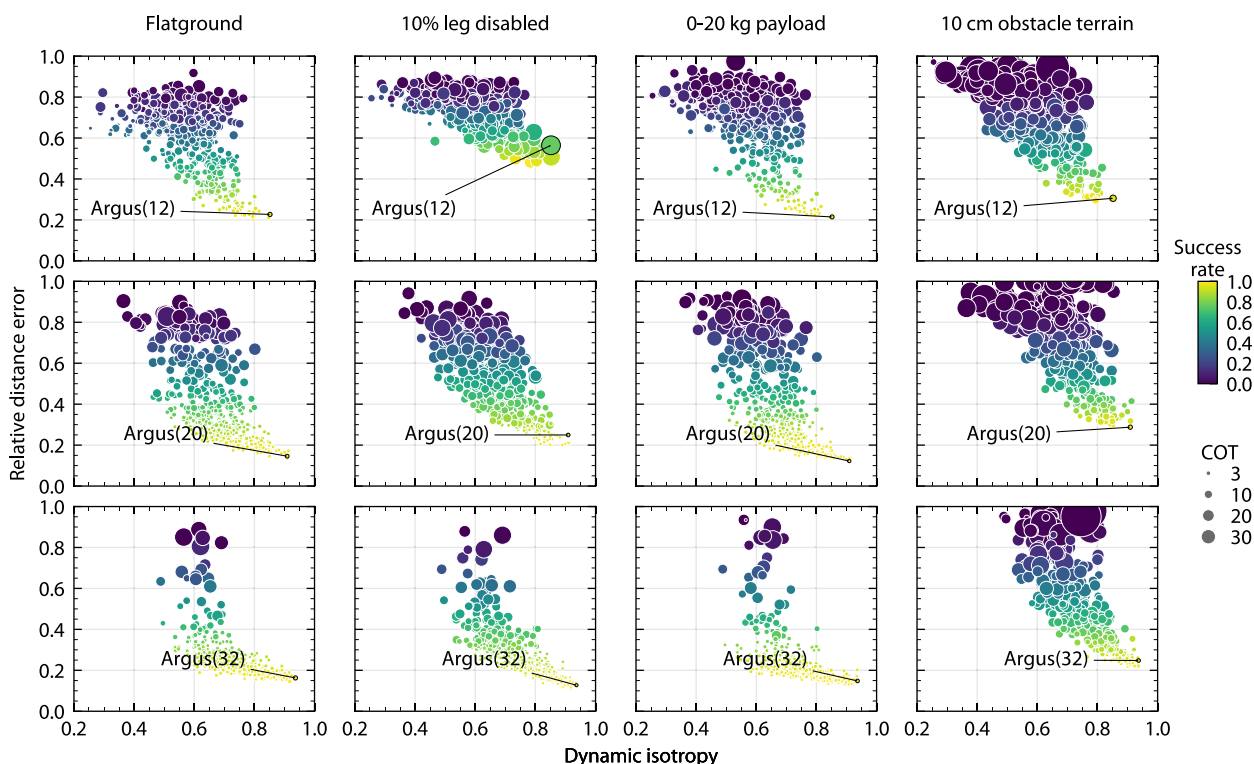


Fig. 5. Effect of dynamic isotropy on performance across randomized Argus morphologies. We evaluated 512 morphological variants each for 12-, 20-, and 32-leg Argus configurations (1536 total robots), spanning dynamic isotropy scores between 0.25 and 0.93 . Each point corresponds to one morphology, with the x axis showing dynamic isotropy η and the y axis showing trajectory tracking error, task success rate, and COT, for all four tasks. Morphologies with higher dynamic isotropy consistently achieved lower tracking error, higher success rates, and reduced COT for all leg counts. Labeled “Argus” points indicate the fully symmetric designs, which lie near the Pareto front of the dynamic symmetry–performance trade-off, highlighting extreme dynamic symmetry as a powerful design axis for robust and efficient locomotion.




Fig. 6. Agile, robust, and resilient locomotion of the 20-leg physical Argus. (A) Argus tracks arbitrary user-commanded trajectories while maintaining continuous rolling motion, exploiting its nearly isotropic actuation to steer without explicit body reorientation. (B) Argus transports a 4.5-kg (10-lb) payload mounted asymmetrically on a single face while preserving most of its commanded forward speed; the dotted line traces a representative foot node [also in (C) to (E)]. (C) Argus traverses discrete terrain by dynamically reconfiguring which legs make contact, morphing its shape to maintain rolling continuity over obstacles. (D) Under lateral pushes, Argus extends legs opposite to the disturbance to generate stabilizing torques and remain in place. (E) Argus maintains locomotion despite having one to three disabled legs by reorienting its body to bring healthy legs into ground contact, demonstrating resilience to partial hardware failure. (F) In a lunar-gravity environment, Argus climbs between parallel walls by alternating bracing and thrusting motions using different subsets of legs, leveraging whole-body isotropic actuation to achieve vertical motion in confined spaces.

from ground obstacles (fig. S5 and movie S7), a result of its morphology's invariance to rotation.

When trained to traverse discrete terrain, Argus exploited its dense leg array to morph its shape, conform to the terrain's profile, and generate propulsion to track the commanded velocity. It achieved a 70% success rate over discrete obstacles with a maximum height of 0.1 m while tracking velocity commands of 0.8 m/s (fig. S4D) in simulation and an 83.3% success rate with a maximum height of 0.12 m while tracking velocity commands of 0.6 m/s (15 of $N = 18$ trials) in physical experiments (Fig. 6C and movie S5). In cases where Argus struggled because of high obstacles, it learned a recovery strategy by first adopting a stable retracted posture and then rapidly extending

its legs to generate recovery momentum to push its body over the obstacle (fig. S8). Such adaptive behavior was facilitated by its ability to uniformly engage with the terrain, redistribute contact forces through redundant body configurations, and execute stabilization policies without requiring a specific body alignment.

Near-extreme dynamic isotropy in the design of Argus also enables passive load-bearing, self-stabilization, and fault tolerance capabilities. Its spherical structure supports the payload passively, allowing the robot to maintain stable motions with relative changes to its actuation strategy. In simulation, Argus transported a 40-kg payload, which is near two times its own body mass (23.4 kg), while retaining 68.4% of its nominal performance without payload, at a commanded rolling speed

of 0.8 m/s (fig. S4C). When it carried 20-kg payloads, 85.7% performance was retained (fig. S4C). In physical testing, a payload was mounted asymmetrically on a single face of the carbon-fiber rod frame, which created a much more challenging and structurally demanding condition than the balanced simulation setup. Under this configuration, Argus carried 4.5 kg (10 lbs), maintaining 96.3% of the commanded velocity of 0.6 m/s (Fig. 6B and movie S4).

Argus achieves self-stabilization by dynamically shifting its CoM and actively interacting with the environment. Legs can be synchronized to approximate a sphere for rolling motions, but they can also extend asymmetrically to resist motion. When dropped at 2 m/s, Argus stabilized with 1.5-m (fig. S4A) displacement (109.5% of its maximum body length). In contrast, the purely spherical structure without coordinated leg control required 2.13 m to stabilize, which is 42.0% worse than the performance with control policy. Under external pushing (Fig. 6D), Argus quickly responded by extending its legs in the opposite direction of the push to remain in place (movie S3). These results highlight the role of the near-extreme isotropic leg distribution in maintaining robustness and stability under dynamic perturbations.

Resiliency against actuator failure emerged from both the redundancy and isotropy of Argus (fig. S3). When certain actuators failed, Argus could redistribute actuation loads uniformly to preserve functionality despite localized failures (fig. S7). Argus retained 95 and 85% of original velocity at 0.8 m/s in simulation with two (10%) and four (20%) legs disabled, respectively (fig. S4B). In physical experiments, forward rolling continued with one to three disabled legs (Fig. 6E). If the disabled configuration reduced efficiency, Argus reoriented its body to bring functional legs into ground contact, thereby restoring propulsion without requiring directional preference (movie S2).

The same isotropic actuation principle enables climbing between parallel walls under the Moon's gravity (Fig. 6F and fig. S6). Argus can extend multiple of its opposite legs to brace its body between the walls while using legs positioned diagonally to the wall surfaces to push upward. By alternating between bracing and pushing motions, Argus learned to generate upward thrust and climb effectively. In simulation, Argus climbed at an average of 0.258 m/s under the Moon's gravity at 85.5% success rate and achieved an average climbing velocity of 0.308 m/s among successful trials. Our physical experiments also demonstrated that Argus could climb parallel walls with an average speed of 0.238 m/s under the Moon's gravity (movie S8), suggesting its potential for low-gravity mobility in space exploration. These results illustrate how near-extreme dynamic symmetry supports agile locomotion not only on horizontal terrain but also in constrained, low-gravity environments relevant to space exploration.

Omnidirectional environmental awareness and whole-body loco-manipulation in the physical Argus

Beyond locomotion and whole-body dynamics, we further augmented the physical Argus with rich environmental

awareness. Because of its modular and symmetric design, each of the 20 foot nodes can accommodate a compact time-of-flight (ToF) depth camera. Equipping all feet with identical sensors yields a radially distributed sensing array that provides nearly uniform coverage of the surrounding space. By transforming the depth information from the 20 ToF sensors (see the “Sensors” section under the “Mechanical design” section in Materials and Methods) into the world coordinate frame, the resulting point cloud forms a consistent 3D representation of the environment that accurately captures the surroundings (Fig. 7, A and B), regardless of the robot's orientation.

Prior platforms that exploit approximate spherical shapes for multidirectional motion include tensegrity robots (36–41) and certain spherical robots (32–35). Although these systems have demonstrated intriguing capabilities, they remain fairly limited in practice: Many rely on hand-crafted or nonadaptive control policies (36, 52), require off-board computation or power (36, 38), produce relatively low forces because of small actuators (39, 53), operate in restricted environments or task sets (40, 52), and often lack external perception (54, 55) that could fully leverage their morphology. In contrast, our physical Argus leverages our theoretical finding of near-extreme dynamic symmetry with high-thrust onboard actuation and fully onboard distributed sensing, allowing it to exploit its uniform dynamic acuation capabilities for agile, perception-guided interaction with the environment.

When tasked with tracking a target object, Argus continuously estimated the object's velocity from the point cloud observation and achieved a success rate of 89.9% and a mean tracking error of 0.167 ± 0.120 m/s in simulation (fig. S9A). In physical experiments, Argus achieved a success rate of 36.8% (14 of $N = 38$ trials), detecting and tracking a moving target of a 1-m cube in various directions under the same initial pose without adjusting orientation (movie S9). Note that the drop in real-world success rate is primarily due to sensing limitations rather than control failure. Most unsuccessful trials stemmed from thermal degradation in the ToF modules, which caused response delays and desynchronization across the 20 cameras during




Fig. 7. Omnidirectional sensing enables object tracking and whole-body loco-manipulation. (A) Argus tracks a 1-m cube under varying velocity commands using depth measurements from 20 foot-mounted ToF cameras fused into a global point cloud. (B) Using the estimated object state, Argus simultaneously pushes and follows the cube by coordinating whole-body actuation and continuous perception.

repeated experiments. These failures typically appeared at the very beginning of a trial because of sensor heating from prior experiments, whereas runs that began with stable, synchronized sensing almost always successfully completed continuous tracking. We acknowledge this as a limitation of existing sensors and highlight it as an important consideration for future sensor integration and robot design choices for deployment.

When whole-body dynamic symmetry is combined with uniformly distributed perception, Argus extends its functionality to more complex interactions, such as whole-body loco-manipulation. In an object-pushing task, Argus exploited lateral legs to apply directional forces while reshaping ground-contacting legs to stabilize and guide the object. In simulation, Argus achieved a 91.5% success rate with a mean direction error of 0.256 ± 0.207 rad over 10 s of continuous pushing (fig. S9B). The learned policy was successfully transferred to the real world (Fig. 7B and movie S10) with a 39.4% success rate (13 of $N = 33$ trials). Because actuation capability is nearly uniform in all directions, Argus maintains multiple contact points with both the ground and the object while preserving continuous perception, enabling seamless transitions between locomotion and manipulation. Similarly, most real-world failures stemmed from overheating-induced delays in ToF cameras rather than policy breakdown, suggesting that more thermally robust sensing hardware will substantially improve real-world success rates.

DISCUSSION

We introduced dynamic symmetry as a possible unifying principle for designing robots whose actuation capabilities remain uniform across all directions, and we formalized this idea through a whole-body measure that we call dynamic isotropy. Across more than 1500 simulated morphologies and extensive physical experiments with a 20-leg prototype, we showed that high dynamic isotropy, especially when close to its theoretical extreme, enables accurate velocity tracking, robust locomotion on diverse terrain, self-stabilization under perturbations, load carrying, resilience to actuator failures, and adaptive whole-body manipulation. Argus, our family of spherical robots built from simple, identical, radially oriented linear actuators, demonstrates that pushing dynamic symmetry toward its theoretical limit makes robots gain enhanced dynamic actuation capabilities, enabling robots to become omnidirectional and multifunctional machines.

Our results reveal that dynamic symmetry serves as a powerful organizing principle across multiple dimensions of robotic functionality. Mechanically, increasing dynamic isotropy produces more uniform whole-body interaction forces, enabling orientation-invariant locomotion and stable behavior across surfaces of varying friction, clutter, or compliance. The symmetry of the actuation inherently distributes load, improves fault tolerance, and facilitates rapid self-stabilization under external perturbations or partial actuator failure. Naturally, this type of design can home distributed perception to enhance object tracking and interaction, especially during large orientation changes. These findings highlight that dynamic symmetry is a principled strategy for coupling morphology, sensing, and control in a single cohesive design space.

Dynamic symmetry also broadens the robot's functional envelope. Argus's whole-body actuation supports unconventional behaviors such as rolling, wall climbing, and whole-body loco-manipulation under continuous motions. Its omnidirectional perception allows the robot to track and interact with objects while maintaining high mobility,

demonstrating that multifunctionality can emerge directly from a dynamic isotropy-driven design rather than from task-specific appendages or specialized subsystems.

At the same time, our studies reveal meaningful trade-offs. As morphology becomes more symmetric, additional actuators contribute decreasing marginal gains to dynamic isotropy and task performance. More spatial redundancies with large numbers of legs introduce added mass, mechanical complexity, and computational load, which can reduce energy efficiency and complicate control optimization. These findings point to the potential need for principled codesign of spatial redundancy and dynamic symmetry (56, 57).

Beyond functional advantages, near-extreme dynamic symmetry yields practical engineering benefits. Argus's modular leg architecture simplifies fabrication, maintenance, replacement, and sim-to-real transfer. Because each foot node can be readily swapped with task-specific attachments, such as high-friction pads for climbing or terrain-adaptive tips for soft substrates, the platform can be customized for new domains without altering its core dynamic symmetry. This modularity positions Argus not only as a scientific demonstrator but also as a promising template for domain-specific deployments.

Looking ahead, advances in actuator technology, onboard sensing, and structural materials could expand the operating envelope of dynamic-symmetric robots. Low-gravity environments, in particular, may unlock the full benefits of whole-body actuation and omnidirectional perception for extraterrestrial exploration. Moreover, future Argus variants equipped with multiple degrees of freedom per leg could explore richer loco-manipulation behaviors while preserving the fundamental principles of dynamic symmetry.

More broadly, we envision dynamic symmetry as a guiding principle for a generation of robots that navigate, perceive, and interact with the world with intrinsic robustness and adaptability. By demonstrating how extreme dynamic isotropy unifies morphology, perception, and control, Argus provides both a conceptual framework and a physical platform for rethinking how robots can engage with complex, unstructured environments.

MATERIALS AND METHODS

Dynamic symmetry and dynamic isotropy

The isotropy measure is primarily used to describe how directionally uniform a robot's velocity or force transmission capabilities are. It characterizes the manipulability of the end effector and serves as a key design criterion for robotic manipulators (43–45, 47). Conventionally, the isotropy is assessed by examining the equality of the singular values of the Jacobian matrix. However, this approach provides only a coarse characterization of the robot's motion capabilities along orthogonal directions in Cartesian space. To more comprehensively quantify how uniformly the robot can accelerate its CoM across all spatial directions, we introduced an isotropy measure based on directional acceleration magnitudes.

CoM linear acceleration modeling

We begin with the standard robot dynamics equation

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{g}(\mathbf{q}) = \boldsymbol{\tau} \quad (3)$$

where $\mathbf{M}(\mathbf{q}) \in \mathbb{R}^{n \times n}$ is the joint-space inertia matrix, $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})$ represents the Coriolis and centrifugal effects, $\mathbf{g}(\mathbf{q})$ is the gravity vector, and $\boldsymbol{\tau} \in \mathbb{R}^n$ is the joint torque vector.

To focus on the effect of joint configuration on the robot's acceleration, we analyzed the dynamics under quasistatic or low-velocity conditions ($\dot{\mathbf{q}} \approx \mathbf{0}$) with gravity compensation. Under these assumptions, the velocity-dependent and gravitational terms are negligible, and the dynamics can be approximated by

$$\ddot{\mathbf{q}} \approx \mathbf{M}^{-1}(\mathbf{q}) \boldsymbol{\tau} \quad (4)$$

Let $\mathbf{J}_c(\mathbf{q}) \in \mathbb{R}^{3 \times n}$ denote the Jacobian matrix mapping joint velocities to the linear velocity of the CoM, that is

$$\dot{\mathbf{x}}_c = \mathbf{J}_c(\mathbf{q}) \dot{\mathbf{q}} \quad (5)$$

Differentiating this relation with respect to time yields the exact expression for the CoM acceleration

$$\mathbf{a}_c = \ddot{\mathbf{x}}_c = \dot{\mathbf{J}}_c(\mathbf{q}, \dot{\mathbf{q}}) \dot{\mathbf{q}} + \mathbf{J}_c(\mathbf{q}) \ddot{\mathbf{q}} \quad (6)$$

where the first term, $\dot{\mathbf{J}}_c(\mathbf{q}, \dot{\mathbf{q}}) \dot{\mathbf{q}}$, captures velocity-dependent effects and the second term represents the contribution from joint accelerations.

Under quasi-static or low-velocity motion ($\dot{\mathbf{q}} \approx \mathbf{0}$), the first term becomes negligible because it is second order in velocity. Neglecting these higher-order terms gives the approximation

$$\mathbf{a}_c \approx \mathbf{J}_c(\mathbf{q}) \ddot{\mathbf{q}} \quad (7)$$

Substituting the quasi-static joint-space dynamics approximation, we obtained a linear relationship between joint torques and CoM acceleration

$$\mathbf{a}_c \approx \mathbf{J}_c(\mathbf{q}) \mathbf{M}^{-1}(\mathbf{q}) \boldsymbol{\tau} \quad (8)$$

For convenience, we defined the configuration-dependent mapping

$$\mathbf{A}(\mathbf{q}) := \mathbf{J}_c(\mathbf{q}) \mathbf{M}^{-1}(\mathbf{q}) \quad (9)$$

so that the CoM acceleration can be compactly written as

$$\mathbf{a}_c = \mathbf{A}(\mathbf{q}) \boldsymbol{\tau} \quad (10)$$

where $\mathbf{A}_i(\mathbf{q}) \in \mathbb{R}^3$ denotes the i th column of $\mathbf{A}(\mathbf{q})$, representing the contribution of the i th joint torque to the CoM acceleration.

Dynamic isotropy

For each unit direction $\mathbf{u} \in \mathbb{S}^2$, the directional sensitivity of actuator i is

$$c_i(\mathbf{u}) = \mathbf{u}^T \mathbf{A}_i(\mathbf{q}) \quad (11)$$

The maximum achievable CoM acceleration magnitude in direction \mathbf{u} is obtained by saturating each actuator at its limit in the direction that maximizes $\mathbf{u}^T \mathbf{A}(\mathbf{q}) \boldsymbol{\tau}$

$$a_{\max}(\mathbf{u}) = \max_{\boldsymbol{\tau} \in \mathcal{T}} \mathbf{u}^T \mathbf{A}(\mathbf{q}) \boldsymbol{\tau} = \sum_{i=1}^n |c_i(\mathbf{u})| \tau_i^{\max} \quad (12)$$

We sampled a set of directions $\{\mathbf{u}_k\}_{k=1}^K$ uniformly distributed on the unit sphere and compute the directional accelerations $a_{\max}(\mathbf{u}_k)$ of the CoM. The dynamic isotropy is then defined as

$$\eta(\mathbf{q}) = \frac{a_{\min}(\mathbf{q})}{a_{\max}(\mathbf{q})}, a_{\min}(\mathbf{q}) = \min_k a_{\max}(\mathbf{u}_k), a_{\max}(\mathbf{q}) = \max_k a_{\max}(\mathbf{u}_k) \quad (13)$$

where $\eta(\mathbf{q}) \in [0, 1]$. A perfectly isotropic acceleration capability corresponds to $\eta(\mathbf{q}) = 1$, indicating equal CoM acceleration magnitude in all directions.

Theoretical analysis of dynamic isotropy for stability, robustness, and control efficiency

In this section, we show how the dynamic isotropy formulation developed above enables theoretical characterization of stability, robustness, and energy efficiency. All results follow directly from the CoM acceleration mapping

$$\mathbf{a}_c = \mathbf{A}(\mathbf{q}) \boldsymbol{\tau}, \mathbf{A}(\mathbf{q}) := \mathbf{J}_c(\mathbf{q}) \mathbf{M}^{-1}(\mathbf{q}) \quad (14)$$

under the quasistatic approximation. Torque limits are modeled as

$$\tau_i \in [-\tau_i^{\max}, \tau_i^{\max}], i = 1, \dots, n \quad (15)$$

Acceleration reachability and ellipsoidal approximation

For a configuration \mathbf{q} , the set of attainable CoM accelerations is

$$\mathcal{A}(\mathbf{q}) := \{\mathbf{A}(\mathbf{q}) \boldsymbol{\tau} \mid \boldsymbol{\tau} \in \mathcal{T}\}, \mathcal{T} := \prod_{i=1}^n \tau_i [-\tau_i^{\max}, \tau_i^{\max}] \quad (16)$$

For a direction $\mathbf{u} \in \mathbb{S}^2$, the maximum achievable acceleration magnitude is

$$a_{\max}(\mathbf{u}) = \max_{\boldsymbol{\tau} \in \mathcal{T}} \mathbf{u}^T \mathbf{A}(\mathbf{q}) \boldsymbol{\tau} = \sum_{i=1}^n |\mathbf{u}^T \mathbf{A}_i(\mathbf{q})| \tau_i^{\max} \quad (17)$$

Sampling directions $\{\mathbf{u}_k\}_{k=1}^K$, we defined

$$a_{\min}(\mathbf{q}) := \min_k a_{\max}(\mathbf{u}_k), a_{\max}(\mathbf{q}) := \max_k a_{\max}(\mathbf{u}_k) \quad (18)$$

and the dynamic isotropy

$$\eta(\mathbf{q}) := \frac{a_{\min}(\mathbf{q})}{a_{\max}(\mathbf{q})} \in [0, 1] \quad (19)$$

To enable theoretical reasoning, we approximated $\mathcal{A}(\mathbf{q})$ by its minimum-volume ellipsoid (Löwner-John ellipsoid) with shape matrix $\mathbf{Q}(\mathbf{q}) > 0$. Let $\lambda_1 \geq \lambda_2 \geq \lambda_3 > 0$ be the eigenvalues of $\mathbf{Q}(\mathbf{q})$. Then,

$$a_{\max}(\mathbf{q}) \approx \sqrt{\lambda_1}, a_{\min}(\mathbf{q}) \approx \sqrt{\lambda_3} \quad (20)$$

and dynamic isotropy is inversely related to the condition number

$$\eta(\mathbf{q}) \approx \frac{\sqrt{\lambda_3}}{\sqrt{\lambda_1}} = \frac{1}{\sqrt{\kappa(\mathbf{Q}(\mathbf{q}))}}, \kappa(\mathbf{Q}(\mathbf{q})) := \frac{\lambda_1}{\lambda_3} \quad (21)$$

Thus, high dynamic isotropy corresponds to a well-conditioned acceleration-map ellipsoid.

Stability margins

To analyze closed-loop stability, we consider the set of all feasible CoM accelerations at configuration \mathbf{q} . Under actuator limits, this set can be approximated by the minimum-volume ellipsoid

$$E(\mathbf{q}) = \{\mathbf{a} \in \mathbb{R}^3 \mid \mathbf{a}^T \mathbf{Q}^{-1}(\mathbf{q}) \mathbf{a} \leq 1\}, \mathbf{Q}(\mathbf{q}) > 0 \quad (22)$$

whose shape matrix $\mathbf{Q}(\mathbf{q})$ captures the distribution of attainable accelerations in all directions.

Suppose a controller requires a corrective CoM acceleration \mathbf{a}_{req} to reject a disturbance or stabilize the system. For the robot to remain stable, this commanded acceleration must lie in the feasible acceleration set $E(\mathbf{q})$. Thus, feasibility requires

$$\mathbf{a}_{\text{req}} \in E(\mathbf{q}) \quad (23)$$

Substituting Eq. 22 gives the explicit feasibility condition

$$\mathbf{a}_{\text{req}}^{\top} \mathbf{Q}^{-1}(\mathbf{q}) \mathbf{a}_{\text{req}} \leq 1 \quad (24)$$

Geometrically, $\mathbf{Q}^{-1}(\mathbf{q})$ defines the “cost” of producing acceleration in each direction: Accelerations aligned with weak axes of the ellipsoid (small eigenvalues of \mathbf{Q}) incur larger quadratic costs.

To quantify how close the system is to violating feasibility, we defined the stability margin as

$$\text{margin}(\mathbf{q}) = 1 - \mathbf{a}_{\text{req}}^{\top} \mathbf{Q}^{-1}(\mathbf{q}) \mathbf{a}_{\text{req}} \quad (25)$$

which equals 1 when no corrective acceleration is required and decreases to 0 when the commanded acceleration reaches the boundary of the feasible ellipsoid. A larger margin indicates greater tolerance to disturbances and improved closed-loop stability. To ensure feasibility, the required corrective acceleration must lie inside the ellipsoidal feasible set $E(\mathbf{q})$, which implies $\|\mathbf{a}_{\text{req}}\| \leq \sqrt{\lambda_3}$ along the weakest principal axis. The absolute disturbance rejection capability, therefore, scales with $\sqrt{\lambda_3}$, whereas dynamic isotropy $\eta \approx \sqrt{\lambda_3/\lambda_1}$ controls how uniformly this capability is distributed across directions. Robots with higher dynamic isotropy avoid having disproportionately weak axes, reducing directional vulnerabilities and enabling orientation-invariant stabilization. Even when λ_1 varies across morphologies, higher η ensures that λ_3 does not collapse relative to λ_1 , preserving larger stability margins under arbitrary disturbance directions.

Disturbance rejection and robustness

External forces \mathbf{F}_{ext} induce a required CoM acceleration $\mathbf{a}_{\text{req}} = \mathbf{M}^{-1}(\mathbf{q}) \mathbf{F}_{\text{ext}}$. Feasibility under worst-case alignment occurs when the induced acceleration points along the weakest principal axis of the feasible ellipsoid $E(\mathbf{q})$, meaning

$$\|\mathbf{F}_{\text{ext}}\|_{\text{max}} \propto \sqrt{\lambda_3} \quad (26)$$

Given that dynamic isotropy satisfies $\eta(\mathbf{q}) = \sqrt{\lambda_3/\lambda_1}$, we could also express this as

$$\|\mathbf{F}_{\text{ext}}\|_{\text{max}} \propto \eta(\mathbf{q}) \sqrt{\lambda_1} \quad (27)$$

highlighting that λ_3 governs the absolute disturbance rejection capability, whereas $\eta(\mathbf{q})$ quantifies its directional uniformity. Higher dynamic isotropy prevents the emergence of dynamically weak directions, ensuring that the minimum achievable acceleration $\sqrt{\lambda_3}$ does not collapse relative to the maximum axis $\sqrt{\lambda_1}$. As a result, robots with higher dynamic isotropy maintain more uniform resilience to terrain irregularities, collisions, uneven contacts, or actuator outages, consistent with our empirical findings in cluttered environments and friction-varying surfaces.

Control efficiency

Producing a desired CoM acceleration \mathbf{a}_{des} requires joint torques $\boldsymbol{\tau}$ satisfying $\mathbf{a}_{\text{des}} = \mathbf{A}(\mathbf{q}) \boldsymbol{\tau}$. The minimum-energy torque command solves the quadratic program

$$\boldsymbol{\tau}^* = \underset{\boldsymbol{\tau}}{\text{argmin}} \|\boldsymbol{\tau}^*\|_2^2 \text{ such that } \mathbf{A}(\mathbf{q}) \boldsymbol{\tau} = \mathbf{a}_{\text{des}} \quad (28)$$

The optimal solution is given by the Moore-Penrose pseudoinverse

$$\boldsymbol{\tau}^* = \mathbf{A}^\dagger(\mathbf{q}) \mathbf{a}_{\text{des}}, \quad \|\boldsymbol{\tau}^*\|_2^2 = \mathbf{a}_{\text{des}}^{\top} (\mathbf{A}(\mathbf{q}) \mathbf{A}^{\top}(\mathbf{q}))^{-1} \mathbf{a}_{\text{des}} \quad (29)$$

To connect this energy requirement to the feasible acceleration ellipsoid, similar as above, we approximated $\mathbf{A}(\mathbf{q})$ by its minimum-volume enclosing ellipsoid with shape matrix $\mathbf{Q}(\mathbf{q}) > 0$. Given that $\mathbf{A}(\mathbf{q}) \mathbf{A}^{\top}(\mathbf{q}) \approx \mathbf{Q}(\mathbf{q})$, the minimum required energy becomes

$$\|\boldsymbol{\tau}^*\|_2^2 \approx \mathbf{a}_{\text{des}}^{\top} \mathbf{Q}^{-1}(\mathbf{q}) \mathbf{a}_{\text{des}} \quad (30)$$

Thus, control energy is proportional to the ellipsoidal norm of the desired acceleration. If \mathbf{a}_{des} is decomposed along the principal axes of $\mathbf{Q}(\mathbf{q})$, we obtain

$$\|\boldsymbol{\tau}^*\|_2^2 \approx \frac{\alpha_1^2}{\lambda_1} + \frac{\alpha_2^2}{\lambda_2} + \frac{\alpha_3^2}{\lambda_3} \quad (31)$$

where α_i is the projection of \mathbf{a}_{des} onto eigenvector i . Control becomes highly imbalanced when λ_3 is much smaller than λ_1 because accelerations aligned with the weakest axis require disproportionately larger torques.

Dynamic isotropy quantifies this imbalance through $\eta(\mathbf{q}) \approx \sqrt{\lambda_3/\lambda_1}$. A higher isotropy value reduces the ratio λ_1/λ_3 , ensuring that no direction demands excessive joint effort relative to others. Thus, high dynamic isotropy directly improves control efficiency by reducing the worst-case and directional variability of the actuation effort needed to realize CoM accelerations of a given magnitude. This analytic prediction aligns with our empirical findings: Robots with higher dynamic isotropy consistently required less total actuation effort to achieve the same locomotion velocities and maneuvering agility, especially when subjected to orientation changes, uneven surfaces, or partial actuator load.

Summary

Dynamic isotropy provides a quantitative and mechanistic link between dynamic symmetry and whole-body capability. High dynamic isotropy ensures that the robot’s acceleration map is well conditioned with orientation-invariant stability margins, uniform robustness to external disturbances, and balanced control efforts. These analytical predictions can help illustrate the empirical advantages observed in dynamically symmetric Argus variants and the trend toward the theoretical extreme of dynamic isotropy, including orientation-invariant locomotion, agile recovery behavior, energy efficiency, and multifunctional behaviors across a wide range of environments.

Dynamic symmetry versus redundancy

Here, we showed that redundancy primarily rescales the feasible-acceleration ellipsoid, whereas dynamic symmetry determines its shape, and therefore governs the stability, robustness, and efficiency results derived above in this work. Recall that the acceleration ellipsoid is characterized by

$$\mathbf{Q}(\mathbf{q}) = \sum_{i=1}^n \tau_i^{\text{max}^2} \mathbf{A}_i(\mathbf{q}) \mathbf{A}_i(\mathbf{q})^{\top}, \quad \mathbf{Q}(\mathbf{q}) > 0 \quad (32)$$

whose eigenvalues satisfy $\lambda_1 \geq \lambda_2 \geq \lambda_3 > 0$. The dynamic isotropy is

$$\eta(\mathbf{q}) \approx \sqrt{\frac{\lambda_3}{\lambda_1}} \quad (33)$$

which depends only on the relative spread of the eigenvalues. Multiplying $\mathbf{Q}(\mathbf{q})$ by any positive scalar c scales all eigenvalues equally,

leaving $\eta(\mathbf{q})$ invariant. Thus, isotropy is insensitive to uniform expansion of the acceleration set.

Suppose that k additional actuators are added whose directions repeat an existing direction $\mathbf{A}_j(\mathbf{q})$. The new matrix is

$$\mathbf{Q}'(\mathbf{q}) = \mathbf{Q}(\mathbf{q}) + \sum_{\ell=1}^k \tau_{\ell}^{\max} \mathbf{A}_j(\mathbf{q}) \mathbf{A}_j(\mathbf{q})^{\top} = \mathbf{Q}(\mathbf{q}) + c \mathbf{A}_j(\mathbf{q}) \mathbf{A}_j(\mathbf{q})^{\top} \quad (34)$$

This update affects only the eigenvalue associated with $\mathbf{A}_j(\mathbf{q})$, increasing λ_1 while leaving λ_2 and λ_3 almost unchanged

$$\lambda_1' = \lambda_1 + \Delta, \quad \lambda_2' \approx \lambda_2, \quad \lambda_3' \approx \lambda_3 \quad (35)$$

Hence,

$$\eta'(\mathbf{q}) \approx \sqrt{\frac{\lambda_3'}{\lambda_1'}} \approx \sqrt{\frac{\lambda_3}{\lambda_1 + \Delta}} \leq \eta(\mathbf{q}) \quad (36)$$

Redundancy alone therefore could reduce isotropy in this extreme case. It cannot explain the observed improvement in worst-case acceleration, stability margin, or robustness. Similarly, increases in actuator count that do not diversify the actuation directions expand the acceleration ellipsoid anisotropically, improving best-case acceleration but leaving the weakest direction unchanged. Even when such additions are distributed evenly for the actuation directions, such redundancy changes only the scale of the ellipsoid, not its overall shape such as uniformity. In contrast, as we show above, more dynamic symmetric designs can reshape the full matrix $\mathbf{Q}(\mathbf{q})$ to produce a more spherical ellipsoid and increase the dynamic isotropy score.

Mechanical design

The 20-leg version of Argos weighs 23.4 kg, and the overall span of the structure ranges from 0.95 m in its fully retracted state to 1.37 m in its fully extended state (fig. S10), corresponding to a leg actuation range of 0.21 m. The home position is defined at the midpoint of the leg's range, allowing for 0.105 m of actuation in both the positive and negative directions. Each custom actuator delivers up to 375 N peak thrust at 1 m/s while weighing only 0.62 kg, which produces 2 to 13 times the maximum thrust per actuator mass compared with all previous solutions (table S4) where linear actuators have been used in prior tensegrity or spherical robots. A ToF camera is equipped at the tip of each leg, providing omnidirectional range sensing (Fig. 1B) that remains unobstructed by leg motion.

Modular leg

To achieve the high-speed and high-force linear actuation required for dynamic and responsive locomotion, we used a cable-driven mechanism that converts rotational motion into linear actuation (Fig. 1A). Each leg costs ~\$300 and exceeds the off-the-shelf linear actuators in both speed and thrust force. The module consists of a quasidirect drive motor (RoboStride 02) housed in a custom motor mount. A rotary drum disk directly coupled to the motor guides a single rope (9KM DWLIFE, 2.3-mm braided Kevlar) routed in both directions and anchored at the two ends of the moving assembly. The moving assembly comprises two triangular end plates connected by three carbon fiber rods (\varnothing 6 mm by 300 mm). Each rod was held by two LM6UU linear bearings (6-mm bore, 12-mm outer diameter, and 19-mm length) embedded in the motor mount. These bearings constrained the motion of the carbon fiber rods to a single axis, allowing

the moving assembly to slide linearly along the actuator axis as the cable on either end is tensioned or released. The outer end plate was connected to a ball-shaped foot node with a radius of 0.06 m and secured in place using screws. The leg modules were interconnected using short carbon fiber rods (\varnothing 6 mm by 150 mm), forming the edges of the dodecahedron. Additional carbon fiber rods (\varnothing 6 mm by 300 mm) connected each leg module to the central support structure, passing through the center of the inner end plate and secured with a bushing (6-mm bore, 8-mm outer diameter, and 10-mm length). These rods provide additional support for each foot, improving overall structural rigidity.

Sensors

A ToF camera (MaixSense-A010) was attached to the tip of the foot node, such that the sensor could capture the view in the direction of actuation without occlusion. The ToF camera has a maximum resolution of 100 pixels by 100 pixels, a field of view of 70° by 60°, a maximum range of 2.5 m, and a maximum frame rate of 20 Hz. We designed the concave shape on the tip of the foot node to ensure enough opening for the field of view of the ToF sensor. The 20 ToF cameras were evenly distributed at the end position of each foot node, such that the robot had full coverage of the environment around its body. In simulation, the ToF camera was modeled via ray casting over a 3D mesh representation of the scene. Rays were emitted from the origin of the foot node and intersected with scene geometry, including a planar mesh representing the ground and cuboid meshes representing objects. The resulting intersection points were computed to emulate the sensor's depth measurements. Each perception unit simulated a fixed spatial resolution of 5-by-5 rays evenly distributed on the basis of the field of view. Only depth measurements within a maximum range of 1.5 m, measured from the tip of the foot node, were retained and converted into point cloud data. Rays returning distances beyond this threshold were masked with zeros. In physical experiments, we downsampled the depth image and zeroed out the distance exceeding the 1.5-m maximum range. The distance threshold was chosen on the basis of the acceptable noise level with respect to the perception distance. The resolution was chosen by balancing the perceived information, the communication bandwidth, and the delays of the signal synchronizations on the basis of our onboard power and computing unit.

Electronics

To preserve mass symmetry, the electronics and battery were mounted on opposite faces of the dodecahedral structure. The controller stack was placed on one face of Argos, whereas two six-cell, 5200 lithium-polymer batteries were positioned on the other side of Argos. The two batteries were connected in series to power all 20 actuators. The robot was controlled using an onboard computer (NVIDIA, Jetson Orin Nano) powered by a USB power bank. The inertial measurement unit (IMU) (SYD Dynamics, TM171) was mounted below the Jetson computer, providing orientation and angular velocity estimates. To enable synchronized streaming of all 20 ToF camera signals, we included an additional mini personal computer (GMKtec NucBox G2 PLUS) to stream 14 sensors' readings and integrate with the remaining six sensors accessed from Jetson through Ethernet. All of the electric cables were tightened on the frames of the base body.

Fabrication

Most structural elements were 3D-printed using polylactic acid, including the center support, motor connectors, carbon fiber rod mounts, actuator end plates, and the electronics and battery holder. Carbon fiber rods, linear bearings, and other electronic parts were

sourced off the shelf. Each actuator was interconnected with its three neighboring actuators using carbon fiber rods and clamped together with screw connections to form the outer edges of the structure. To further improve rigidity, the midpoints of the edges on each pentagonal face were connected by tensioned ropes arranged in a pentagram pattern.

Policy training, tasks, and evaluations

Policy training

We modeled each control task as a partially observable Markov decision process (58), where the control policy receives partial and noisy observations from onboard sensors and must infer its underlying state to make decisions. Let $\pi_{\theta}(\mathbf{a}_t|\mathbf{o}_t)$ denote a policy parameterized by θ , mapping observations \mathbf{o}_t to actions \mathbf{a}_t . The objective is to maximize the expected cumulative reward

$$\max_{\theta} \mathbb{E}_{\pi_{\theta}} \left[\sum_{t=0}^T \gamma^t r(\mathbf{s}_t, \mathbf{a}_t) \right] \quad (37)$$

where $r(\mathbf{s}_t, \mathbf{a}_t)$ and \mathbf{s}_t are the reward and true state at time t and $\gamma \in (0, 1)$ is the discount factor. We trained the control policy in Isaac Gym (59) using the proximal policy optimization algorithm (60), which improved policy performance through stable updates by optimizing the clipped surrogate objective

$$L^{\text{CLIP}}(\theta) = \hat{\mathbb{E}}_t \left[\min \left(r_t(\theta) \hat{A}_t, \text{clip} \left(r_t(\theta), 1 - \epsilon, 1 + \epsilon \right) \hat{A}_t \right) \right] \quad (38)$$

where $\hat{\mathbb{E}}_t$ denotes the empirical average over a finite batch of collected time steps, $r_t(\theta) = \frac{\pi_{\theta}(\mathbf{a}_t|\mathbf{o}_t)}{\pi_{\theta_{\text{old}}}(\mathbf{a}_t|\mathbf{o}_t)}$ is the probability ratio between the new and old policies, \hat{A}_t is the estimated advantage function, and ϵ is a clipping hyperparameter. Task configurations and evaluation metrics are detailed below.

Baseline locomotion on flat terrain

During training and evaluation, 8192 Argus robots were initialized on flat ground with random orientations and were commanded to follow a target velocity uniformly sampled from -0.8 to 0.8 m/s in the horizontal plane. A total of 8192 evaluations were conducted, and linear velocity error was measured at 3 s after the start of each trial.

Self-stabilization

During training, 8192 Argus robots were initialized 0.5 m above the ground with randomized orientations and zero commanded velocity. At initialization, each robot received a perturbation: a randomly sampled linear velocity in the range of -2 to 2 m/s and angular velocity in the range of -2 to 2 rad/s, applied across all three spatial axes. A total of 8192 evaluations were conducted. A stabilization distance was measured as the horizontal distance from the point of first ground contact until the robot came to rest or until the end of the episode.

Carry object

Argus was trained under the same conditions as the baseline locomotion task, with additional randomizations of base mass, inertia, and CoM position. A total of 8192 evaluations were conducted, with a random mass between 0 and 40 kg added to the baselink. During evaluation, Argus was commanded to move at a constant velocity of 0.8 m/s for 10 s. A success rate was computed as the fraction of trials in which the robot traveled at least 50% of the commanded distance.

For real-world testing, a weighted pouch was attached to the surface of the robot, and the robot was commanded to move at 0.6 m/s in the forward (+x) direction.

Locomotion with leg failure

In addition to the baseline locomotion task, the legs of Argus were randomly disabled; each leg had a 10% chance to be disabled. Argus was evaluated under a constant velocity command of 0.8 m/s for a duration of 10 s. A success rate was computed as the fraction of trials in which the robot traveled at least 50% of the commanded distance. A total of 8192 evaluations were conducted. For physical experiments, we randomized the number of legs disabled and the orientation of the robot and commanded the robot to move at 0.6 m/s in the forward (+x) direction.

Discrete terrain traversal

In addition to flat-ground locomotion, Argus was trained on discrete obstacle terrain, where rectangular block obstacles up to 0.1 m in height were placed randomly. During training, the obstacle height was gradually increased on the basis of the robot's tracking performance. For evaluation, a total of 8192 evaluations were sampled, and the robot moved at a commanded velocity of 0.8 m/s for 10 s. The trial was considered successful if the robot traversed half of the commanded distance (4 m), and the success rate was the fraction of trials that met this criterion. In physical experiments, the trial was considered successful if the robot traveled through all of the discrete blocks.

Climb up wall at low gravity

Argus was initialized between two parallel rough walls spaced 1 m apart. It was tasked with climbing the walls as quickly as possible, with part of its weight supported upward to simulate low gravity. During training, 4096 robots were initialized at random heights, positions, and orientations between the walls. The base support force was gradually reduced on the basis of the average velocity of the population, simulating from 8.2% of Earth's gravity to the Moon's gravity (16.6% of Earth's gravity). Linear velocity along the vertical direction was measured for this task. For evaluation, 4096 robots were initialized on the ground between the same 1-m parallel walls. A trial was considered successful if the robot achieved a vertical velocity of at least 0.08 m/s by the end of the 5-s trial. For the physical experiments, we attached a 19.5-kg counterweight to the robot using a pulley system with two rolling wheels to simulate lunar gravity. In each trial, the robot was initialized in the same starting pose.

Object tracking

The object was randomly initialized around Argus with an offset of 1.5 m. The object's y axis was parallel with the direction of commanded linear velocity. The linear velocity command amplitude remained the same across the episode of the environment and was randomly sampled in the range of 0.5 to 0.8 m/s. The reward encouraged the robot to match the object's velocity. The performance was evaluated over 7168 environments for 10 s of tracking with random initial robot orientation. The object's velocity was randomly sampled, consistent with the training setup. The success rate was measured on the basis of the distance of the robot and object in the perceptual range of the robot. We used a maximum distance of 2.5 m in the first 5 s of tracking as the criterion in the success rate evaluation. In physical experiments, the trial was considered successful if the robot continuously tracked the object to the end of the object's traveled distance.

Object pushing

The object was initialized in front of Argus with 0.8-m offset away from the robot's position. The commanded velocity was randomly sampled in the range of -40° to 40° with the same normal amplitude

of 0.6 m/s. The performance was evaluated over 7168 environments for 10 s of pushing with random initial orientation. A trial was considered successful if the robot-object distance remained less than 1.6 m throughout the 10 s of the pushing, and the robot did not surpass the object with a minimal robot-object distance of 0.5 m in the commanded pushing direction. In the real-world experiment, a trial was considered successful if the robot continuously pushed the object until it reached the end of the experimental space.

Two-stage learning from ToF observations

For the perception-related task, computing the large amount of distance information to simulate 20 ToF cameras is slow with large-scale parallel training in the Isaac Gym simulator. To facilitate training, we adopted a two-stage training pipeline. In the first stage, we trained the policy given privileged information about the object states using reinforcement learning, optimizing for the expected cumulative reward of the task. Such information is not available during real-world deployment. In the second stage, we rolled out the locomotion policy and collected the point cloud and object state pairs. We supervised the training of a point cloud encoder to predict the object states, including the object's x and y direction velocity in the object-tracking task and the object orientation in the object-pushing task. The point cloud encoder consisted of two PointNet (61) layers to encode the input contact points to a global feature vector with the dimensions of 1 by 1024. This feature vector was then passed through three fully connected layers to predict the object states. To deploy the policy using point cloud observations directly from the ToF cameras, the encoder predicted the object states, which was then concatenated with other proprioceptive observations and fed into the locomotion policy. Given a point cloud observation \mathbf{P}_i , we trained the point cloud encoder f_p , parameterized by θ_p , to output the corresponding object states $\hat{\mathbf{O}}_i$. We optimized the following loss function \mathcal{L}_{mse} on the basis of the mean squared error loss

$$\min_{\theta_p} \sum_i \mathcal{L}_{\text{MSE}}(\hat{\mathbf{O}}_i, \mathbf{O}_i), \hat{\mathbf{O}}_i = f_p(\mathbf{P}_i) \quad (39)$$

Observations

The proprioceptive observation included angular velocity \mathbf{v}_b , joint position \mathbf{q} , joint velocity $\dot{\mathbf{q}}$, previous action \mathbf{a}_{t-1} , and base rotation matrix \mathbf{R} . Up to 10% of noise was applied to each observation. Detailed observations and states for each task are listed in table S1. In physical experiments, the angular velocity and base rotation were estimated from the IMU. The linear joint position and velocity were derived from motor encoder readings, scaled by the drum radius. We stacked three frames of the observation for the locomotion policy. We used a single frame observation in the object-tracking, object-pushing, and wall-climbing tasks.

Actions

The policy's actions \mathbf{a}_t were smoothed using a first-order exponential low-pass filter ($\mathbf{a}_t^* = \alpha \mathbf{a}_t + (1 - \alpha) \mathbf{a}_{t-1}^*$, where α is randomized each step in the range of 0.5 to 0.9. The smoothed action was then converted to target joint positions ($\mathbf{q}^* = \mathbf{q}_0 + k_q \mathbf{a}_t^*$) at 25 Hz. Here, \mathbf{q}_0 is the default joint positions, and k_q is a scaling factor for the action. For the wall-climbing task, action smoothing was disabled to allow rapid action changes required for fast and responsive behavior. The target joint positions were transformed into torques via a position PD controller operating at 200 Hz in simulation. For physical experiments,

the position PD controller operates asynchronously at 200 Hz. The empirically determined PD gains were 10 and 0.6 N·m/rad.

Reward

The reward functions for the baseline velocity tracking task (table S3) include tracking and regularization terms. $\phi(\boldsymbol{\psi}, \mathbf{w}) := \exp(\sum_{i=1}^n w_i \psi_i^2)$ is the exponential of the weighted sum of squared quantities $\boldsymbol{\psi}$, where w is a corresponding scaling factor. \mathbf{v}_b and \mathbf{g}_b represent the base linear velocity and the projected gravity vector, respectively. \mathbf{q}_{\min} and \mathbf{q}_{\max} define the joint position limits. \mathbf{a} and $\dot{\mathbf{a}}$ refer to the action and action rate, respectively. $\boldsymbol{\tau}$ denotes the joint actuation force, $\dot{\mathbf{q}}$ is the joint velocity, $\ddot{\mathbf{q}}$ is the joint acceleration, and \mathbf{F}_{cz} is the contact force in the vertical direction. \mathbf{I} is a Boolean indicator of foot contact, and G is the robot's total gravitational force. For the object-tracking task, the object's velocity is defined the same as \mathbf{v}_b^* and used in the linear velocity reward. In the object-pushing task, two additional rewards are included. One encourages the alignment of the object's y -axis direction with the commanded velocity, and the other encourages the robot to move closer to the object. The term \mathbf{y}_{obj} denotes the normalized y -axis vector of the object. The commanded pushing velocity is represented by $\mathbf{v}_{\text{command}}$. $\mathbf{p}_{\text{robot}}$ and $\mathbf{p}_{\text{object}}$ refer to the positions of the robot and the object, respectively.

Sim-to-real transfer

During the locomotion policy training, domain randomization (table S2) was applied to cover possible variations of real-world dynamics, including loose cable or uneven cable tension, leg deformation, and friction in foot nodes and linear bearings. To transfer the trained policy to the real world, we commanded the five-foot node to track a sinusoidal trajectory for both the sim and real robot. We tuned motor parameters (k_p, k_d) and the torque ratio to match the motor position and velocity responses between the sim and the real. In the simulation, a 20-action delay was added to match the hardware latency. We also found that battery level could affect motor torque and influence the overall sim-to-real transfer. In our physical experiments, we found no major influence of policy transfer under continuous 30-min operations. We leave the exploration of power duration for longer-term operations as future work. We added rubber tape to the foot nodes to improve friction and ensured that the robot had enough pushing force to move forward. For the sim-to-real transfer of the object tracking and loco-manipulation policy, the challenge for ToF cameras primarily lies in the noise from the real sensor reading, the sensor reading delay, and the sensor overheating. To simulate real-world randomness or measurement error during training, we applied 20% random dropout and Gaussian noise following a normal distribution with zero mean and an SD of 0.005 to each measured point. Modeling the thermal behavior of the 20 ToF sensors is beyond the scope of this work and is left for future investigation. We believe that future sensor advancements and selections can largely mitigate the thermal issues.

Statistical analysis

Quantitative results consist of deterministic dynamic-isotropy computations, large-scale simulated rollouts, and small-sample physical-robot trials. Dynamic isotropy scores and attainable-acceleration clouds (Figs. 2 to 5 and figs. S1 to S3) were computed by sampling 2048 uniformly distributed directions on the unit sphere for each morphology and pose.

The redundancy analysis in fig. S3 reports the relative distance error over 10 equal-width isotropy bins across 1536 randomized

morphologies (512 each for the 12-, 20-, and 32-leg variants), summarized as means \pm SD in each bin. In each isotropy bin, pairwise Mann-Whitney U tests were conducted on the relative distance error. All pairwise comparisons in each bin were corrected for multiple comparisons using the Bonferroni method. Significance thresholds were set at $*P < 0.05$, $**P < 0.01$, and $***P < 0.001$. Results are presented as box plots showing the median, interquartile range (IQR), and $1.5 \times$ IQR whiskers. The sample sizes (n) are reported for each group in the figure.

Simulated policies were evaluated in Isaac Gym under randomized initial conditions and the domain randomization (table S2), with 8192 trials for baseline locomotion and locomotion under leg failure, self-stabilization, payload carrying, discrete-terrain traversal; 4096 trails for wall climbing; and 7168 trials for object tracking and object pushing. Unless noted otherwise, shaded regions and error bars denote the 95% confidence interval, and success rates are reported as the fraction of trials meeting the task-specific criterion defined in each task description above.

Physical robot outcomes are reported as trial counts on a single 20-leg prototype: 15 of 18 trials for discrete-terrain traversal, 14 of 38 trials for object tracking, and 13 of 33 trials for object pushing. Data aggregation, statistical summaries, and plotting were performed in Python using NumPy, pandas, and Matplotlib.

Supplementary Materials

The PDF file includes:

Methods

Figs. S1 to S10

Tables S1 to S4

Legends for movies S1 to S10

References (63, 64)

Other Supplementary Material for this manuscript includes the following:

Movies S1 to S10

Data file S1

REFERENCES AND NOTES

- K. Capek, *R.U.R. (Rossum's Universal Robots)* (Penguin, 2004).
- R. Pfeifer, M. Lungarella, F. Iida, Self-organization, embodiment, and biologically inspired robotics. *Science* **318**, 1088–1093 (2007).
- D. Trivedi, C. D. Rahn, W. M. Kier, I. D. Walker, Soft robotics: Biological inspiration, state of the art, and future research. *Appl. Bionics Biomech.* **5**, 99–117 (2008).
- F. Iida, A. J. Ijspeert, "Biologically inspired robotics" in *Springer Handbook of Robotics* (Springer, 2016), pp. 2015–2034.
- Y. Bar-Cohen, C. Breazeal, Biologically inspired intelligent robots. *Proc. SPIE* **5051**, 14–20 (2003).
- N. F. Lepora, P. Verschure, T. J. Prescott, The state of the art in biomimetics. *Bioinspir. Biomim.* **8**, 013001 (2013).
- M. Raibert, K. Blankespoor, G. Nelson, R. Playter, Bigdog, the rough-terrain quadruped robot. *IFAC Proc. Vol.* **41**, 10822–10825 (2008).
- M. Hutter, C. Gehring, D. Jud, A. Lauber, C. D. Bellicoso, V. Tsounis, J. Hwangbo, K. Bodie, P. Fankhauser, M. Bloesch, R. Diethelm, S. Bachmann, A. Melzer, M. Hoepflinger, "Anymal—a highly mobile and dynamic quadrupedal robot" in *Proceedings of the 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2016), pp. 38–44.
- G. Bleidt, M. J. Powell, B. Katz, J. Di Carlo, P. M. Wensing, S. Kim, "Mit cheetah 3: Design and control of a robust, dynamic quadruped robot" in *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2018), pp. 2245–2252.
- G. Metta, L. Natale, F. Nori, G. Sandini, D. Vernon, L. Fadiga, C. Von Hofsten, K. Rosander, M. Lopes, J. Santos-Victor, A. Bernardino, L. Montesano, The iCub humanoid robot: An open-systems platform for research in cognitive development. *Neural Netw.* **23**, 1125–1134 (2010).
- A. Saloutos, E. Stanger-Jones, Y. Ding, M. Chignoli, S. Kim, "Design and development of the mit humanoid: A dynamic and robust research platform" in *Proceedings of the 2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)* (IEEE, 2023), pp. 1–8.
- B. Xia, B. Li, J. Lee, M. Scutari, B. Chen, "The Duke Humanoid: Design and control for energy-efficient bipedal locomotion using passive dynamics" in *Proceedings of the 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2025), pp. 6579–6586.
- Q. Liao, B. Zhang, X. Huang, X. Huang, Z. Li, K. Sreenath, "Berkeley humanoid: A research platform for learning-based control" in *Proceedings of the 2025 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2025), pp. 2897–2904.
- E. Mattar, A survey of bio-inspired robotics hands implementation: New directions in dexterous manipulation. *Robot. Auton. Syst.* **61**, 517–544 (2013).
- R. R. Ma, A. M. Dollar, "On dexterity and dexterous manipulation" in *2011 15th International Conference on Advanced Robotics (ICAR)* (IEEE, 2011), pp. 1–7.
- B. Sun, W. Li, Z. Wang, Y. Zhu, Q. He, X. Guan, G. Dai, D. Yuan, A. Li, W. Cui, D. Fan, Recent progress in modeling and control of bio-inspired fish robots. *J. Mar. Sci. Eng.* **10**, 773 (2022).
- R. Baines, S. K. Patiballa, J. Booth, L. Ramirez, T. Sipple, A. Garcia, F. Fish, R. Kramer-Bottiglio, Multi-environment robotic transitions through adaptive morphogenesis. *Nature* **610**, 283–289 (2022).
- G. Li, X. Chen, F. Zhou, Y. Liang, Y. Xiao, X. Cao, Z. Zhang, M. Zhang, B. Wu, S. Yin, Y. Xu, H. Fan, Z. Chen, W. Song, W. Yang, B. Pan, J. Hou, W. Zou, S. He, X. Yang, G. Mao, Z. Jia, H. Zhou, T. Li, S. Qu, Z. Xu, Z. Huang, Y. Luo, T. Xie, J. Gu, S. Zhu, W. Yang, Self-powered soft robot in the Mariana Trench. *Nature* **591**, 66–71 (2021).
- W. D. Shin, H.-V. Phan, M. A. Daley, A. J. Ijspeert, D. Floreano, Fast ground-to-air transition with avian-inspired multifunctional legs. *Nature* **636**, 86–91 (2024).
- E. Chang, D. D. Chin, D. Lentink, Bird-inspired reflexive morphing enables rudderless flight. *Sci. Robot.* **9**, eado4535 (2024).
- J. Langowski, P. Sharma, A. L. Shoushtari, In the soft grip of nature. *Sci. Robot.* **5**, eabd9120 (2020).
- S. Ocklenburg, A. Mundorf, Symmetry and asymmetry in biological structures. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2204881119 (2022).
- M. Enquist, A. Arak, Symmetry, beauty and evolution. *Nature* **372**, 169–172 (1994).
- X. Zhu, D. Wang, O. Biza, G. Su, R. Walters, R. Platt, "Sample efficient grasp learning using equivariant models," paper presented at Robotics: Science and Systems, New York, NY, 27 June to 1 July, 2022.
- D. Wang, R. Walters, X. Zhu, R. Platt, "Equivariant Q learning in spatial action spaces" in *Proceedings of the Conference on Robot Learning (PMLR)*, 2022, pp. 1713–1723.
- D. Wang, M. Jia, X. Zhu, R. Walters, R. Platt, "On-robot learning with equivariant models" in *Proceedings of the Conference on Robot Learning (PMLR)*, 2023, pp. 1345–1354.
- Z. Su, X. Huang, D. Ordoñez-Apaez, Y. Li, Z. Li, Q. Liao, G. Turrissi, M. Pontil, C. Semini, Y. Wu, "Leveraging symmetry in RL-based legged locomotion control" in *Proceedings of the 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2024), pp. 6899–6906.
- M. Mittal, N. Rudin, V. Klemm, A. Allshire, M. Hutter, "Symmetry considerations for learning task symmetric robot policies" in *2024 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2024), pp. 7433–7439.
- D. O. Apaez, G. Turrissi, V. Kostic, M. Martin, A. Agudo, F. Moreno-Noguer, M. Pontil, C. Semini, C. Mastalli, Morphological symmetries in robotics. *Int. J. Robot. Res.* **44**, 1743–1766 (2024).
- S. Yan, B. Zhang, Y. Zhang, J. Boedecker, W. Burgard, "Learning continuous control with geometric regularity from robot intrinsic symmetry" in *2024 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2024), pp. 49–55.
- E. Frazzoli, M. A. Dahleh, E. Feron, Maneuver-based motion planning for nonlinear systems with symmetries. *IEEE Trans. Robot.* **21**, 1077–1091 (2005).
- D. K. Pai, R. A. Barman, S. K. Ralph, "Platonic beasts: A new family of multilimbed robots" in *Proceedings of the 1994 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 1994), pp. 1019–1025.
- H. Nozaki, Y. Kujirai, R. Niiyama, Y. Kawahara, T. Yonezawa, J. Nakazawa, "Continuous shape changing locomotion of 32-legged spherical robot" in *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2018), pp. 2721–2726.
- V. Gheorghe, N. Alexandrescu, D. Duminica, L. A. Cartal, "Rolling robot with radial extending legs" in *Proceedings of the 3rd International Symposium on Resilient Control Systems* (IEEE, 2010), pp. 107–112.
- R. Liu, Y.-a. Yao, Y. Li, Design and analysis of a deployable tetrahedron-based mobile robot constructed by Sarrus linkages. *Mech. Mach. Theory* **152**, 103964 (2020).
- C. Paul, F. J. Valero-Cuevas, H. Lipson, Design and control of tensegrity robots for locomotion. *IEEE Trans. Robot.* **22**, 944–957 (2006).
- M. Vespignani, J. M. Friesen, V. SunSpiral, J. Bruce, "Design of SUPERball v2, a compliant tensegrity robot for absorbing large impacts" in *Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2018), pp. 2865–2871.
- K. Kim, A. K. Agogino, A. M. Agogino, Rolling locomotion of cable-driven soft spherical tensegrity robots. *Soft Robot.* **7**, 346–361 (2020).

39. J. Jeong, I. Kim, Y. Choi, S. Lim, S. Kim, H. Kang, D. Shah, R. Baines, J. W. Booth, R. Kramer-Bottiglio, S. Y. Kim, Spikebot: A multigait tensegrity robot with linearly extending struts. *Soft Robot.* **11**, 207–217 (2024).
40. D. Surovik, K. Wang, M. Vespignani, J. Bruce, K. E. Bekris, Adaptive tensegrity locomotion: Controlling a compliant icosahedron with symmetry-reduced reinforcement learning. *Int. J. Robot. Res.* **40**, 375–396 (2021).
41. J. Liu, Z. Xu, J. Lu, X. Gu, J. Wu, LCRBot: Load-carrying rolling robot based on truncated hexahedral tensegrity. *J. Field Robot.* **42**, 2454–2468 (2025).
42. N. Chen, K. Wang, W. R. Johnson III, R. Kramer-Bottiglio, K. Bekris, M. Aanjaneya, “Learning differentiable tensegrity dynamics using graph neural networks” in *Proceedings of the Conference on Robot Learning (CoRL)* (PMLR, 2024), pp. 5134–5149.
43. J. K. Salisbury, J. J. Craig, Articulated hands: Force control and kinematic issues. *Int. J. Robot. Res.* **1**, 4–17 (1982).
44. T. Yoshikawa, Manipulability of robotic mechanisms. *Int. J. Robot. Res.* **4**, 3–9 (1985).
45. C. A. Klein, B. E. Blahot, Dexterity measures for the design and control of kinematically redundant manipulators. *Int. J. Robot. Res.* **6**, 72–83 (1987).
46. O. Ma, J. Angeles, “The concept of dynamic isotropy and its applications to inverse kinematics and trajectory planning” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 1990), pp. 481–486.
47. C. A. Klein, T. A. Miklos, Spatial robotic isotropy. *Int. J. Robot. Res.* **10**, 426–437 (1991).
48. S. Kim, I. Jeong, S. Lee, “Systematic isotropy analysis of a mobile robot with three active caster wheels” in *International Conference on Intelligent Computing* (Springer, 2007), pp. 587–597.
49. O. Ma, J. Angeles, “Optimum design of manipulators under dynamic isotropy conditions” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 1993), pp. 470–475.
50. T. Yoshikawa, Dynamic manipulability of robot manipulators. *Trans. Soc. Instrum. Control Eng.* **21**, 970–975 (1985).
51. L. Glasser, A. Every, Energies and spacings of point charges on a sphere. *J. Phys. A* **25**, 2473–2482 (1992).
52. Y. Zheng, Y. Li, Y. Lu, M. Wang, X. Xu, C. Zhou, Y. Luo, Robustness evaluation for rolling gaits of a six-strut tensegrity robot. *Int. J. Adv. Robot. Syst.* **18**, 1–11 (2021).
53. H. Nozaki, R. Niiyama, T. Yonezawa, J. Nakazawa, “Shape changing locomotion by spiny multipedal robot” in *2017 IEEE International Conference on Robotics and Biomimetics (ROBIO)* (IEEE, 2017), pp. 2162–2166.
54. F. Xu, X. Zhao, M. Yue, A physics-driven closed-loop motion planning method for spherical multi-expandable-limb robots. *IEEE Trans. Ind. Electron.* **71**, 16087–16097 (2024).
55. D. Yang, Y. Liu, Y. Yu, “A general locomotion approach for a novel multi-legged spherical robot” in *Proceedings of the 2023 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2023), pp. 10146–10152.
56. H. Dong, J. Zhang, T. Wang, C. Zhang, “Symmetry-aware robot design with structured subgroups” in *Proceedings of the International Conference on Machine Learning (PMLR)* (2023), pp. 8334–8355.
57. M. Ghaffari, R. Zhang, M. Zhu, C. E. Lin, T.-Y. Lin, S. Teng, T. Li, T. Liu, J. Song, Progress in symmetry preserving robot perception and control through geometry and learning. *Front. Robot. AI* **9**, 969380 (2022).
58. L. P. Kaelbling, M. L. Littman, A. R. Cassandra, Planning and acting in partially observable stochastic domains. *Artif. Intell.* **101**, 99–134 (1998).
59. V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, G. State, Isaac Gym: High performance GPU-based physics simulation for robot learning. arXiv:2108.10470 [cs.RO] (2021).
60. J. Schulman, F. Wolski, P. Dhariwal, A. Radford, O. Klimov, Proximal policy optimization algorithms. arXiv:1707.06347 [cs.LG] (2017).
61. C. R. Qi, H. Su, K. Mo, L. J. Guibas, “Pointnet: Deep learning on point sets for 3d classification and segmentation” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (IEEE, 2017)*, pp. 652–660.
62. A. Dasgupta, SpiderBot_DeepRL: A custom-designed spider robot trained to walk using deep reinforcement learning in a PyBullet simulation, GitHub (2026); https://github.com/arijit-dasgupta/SpiderBot_DeepRL.
63. robot-descriptions, awesome-robot-descriptions, GitHub (2026); <https://github.com/robot-descriptions/awesome-robot-descriptions>.
64. D. Yang, Y. Liu, F. Ding, Y. Yu, “Bionic multi-legged robot based on end-to-end artificial neural network control” in *2022 IEEE International Conference on Cyborg and Bionic Systems (CBS)* (IEEE, 2023), pp. 104–109.

Acknowledgments: We thank E. Ma, J. Lee, and M. Li for help in 3D printing. We thank S. Moore for the advice on the implementation of ray casting and Y. Jia for the reinforcement learning implementation. **Funding:** This work is supported by DARPA FoundSci program under award HR00112490372, DARPA TIAMAT program under award HR00112490419, ARO under award W911NF2410405, and ARL STRONG program under awards W911NF2320182, W911NF2220113, and W911NF242021. **Author contributions:** B.C., J.L., and B.X. conceived and designed the research. J.L. and B.X. designed simulations and performed physical experiments. All authors analyzed data and wrote the manuscript. **Competing interests:** Duke University has filed patent rights for the technology associated with this manuscript. For further license rights, including using the patent rights for commercial purposes, please contact Duke’s Office for Translation and Commercialization (otcquestions@duke.edu) and reference OTC DU8860PROV. **Data, code, and materials availability:** The code is available at <https://doi.org/10.5061/dryad.3j9kd520k>. The high-resolution versions of all of the supplementary videos can be accessed at <https://figshare.com/s/0e117c0158f70c39c6b8>. All materials are commercially available.

Submitted 9 September 2025

Accepted 28 April 2026

Published 27 May 2026

10.1126/scirobotics.aec1725

Extreme dynamic symmetry enables omnidirectional and multifunctional robots

Jiaxun Liu, Boxi Xia, and Boyuan Chen

Sci. Robot. **11** (114), eaec1725. DOI: 10.1126/scirobotics.aec1725

View the article online

<https://www.science.org/doi/10.1126/scirobotics.aec1725>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)

Science Robotics (ISSN 2470-9476) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. The title *Science Robotics* is a registered trademark of AAAS.

Copyright © 2026 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works