Robots vs Bodies

By Professor Justin Stebbing

Human and robotic movement both involve pushing matter through space under constraint, but they solve this problem with completely different architectures. The body is a wet, self-repairing continuum of muscles, tendons, fascia and fluid, orchestrated by a noisy, predictive low energy consuming brain; the robot is a stack of rigid links and algorithms, precise but brittle, living in a world of coordinates, torque limits and a constrained power supply. That divergence runs all the way down: from how each system senses its own posture, to what "balance" and "falling" mean, to what it is to 'will a movement' rather than merely execute a command. Thus far, movement in the human body and in robots can be made to look superficially similar, but evidence to date shows that they are built on fundamentally different design principles and control logics. The body realises motion through a soft, heterogeneous musculoskeletal system whose passive dynamics and neural control are so deeply intertwined that biomechanics and "software" cannot be cleanly separated, whereas robots usually begin life as explicitly modelled kinematic chains whose motion is prescribed in an abstract configuration space and only later "embodied" in hardware. We've now moved well beyond Isaac Asimov, the plains of Tatooine and the Terminator movies. With integration of predictive AI systems and use of new materials, we are now for the first time, seeing some convergence here.

Passive versus active dynamics

Standing upright is a small miracle of unstable dynamics. The human body is effectively an inverted pendulum with multiple segments stacked on top; the centre of mass is high, the base of support small, and yet most adults can stand for hours with minimal sway. That stability emerges not from stiffening everything to maximum, but from allowing controlled "wobble": muscles around the ankles and hips continuously adjust tone, making tiny corrections that keep the centre of mass over the feet without freezing the joints. When something unexpected happens—a push, a slip—the response is phase-dependent: sometimes an ankle strategy (small corrections), sometimes a hip strategy (bigger, faster torques), sometimes a step to create a new base of support. Robots by contrast, especially early humanoids, tend to stabilise by doing what a biomechanist would never recommend to a human: locking joints and maintaining a precise, precomputed zero moment point within the support polygon. This is very energy-efficient and precise on a flat factory floor, but brittle on irregular terrain or when interacting with humans. Thus, human locomotion exploits compliant mechanics in a way that most classical robots do not. Recent human biomechanic studies of human hopping[1] or running[2] show that joints and soft tissues act like tuneable springs and struts, with the knee, ankle and hip changing their apparent stiffness and function as gait frequency and task demands vary. This "embodied" compliance offloads computation from the nervous system: the body's mechanical design filters perturbations and recycles energy, so that control can be relatively low-bandwidth and reflex-like while still producing stable, efficient movement.

By contrast, many robots, especially industrial arms and early legged platforms, are deliberately built to be rigid and torque-transparent so that their motion can be computed and controlled with high precision from explicit models. Where the human leg behaves like a mass–spring system whose natural dynamics are harnessed, a robot leg must often simulate that behaviour in software, using actuators to impose virtual compliance or adding dedicated elastic elements after the fact. Newer designs are now incorporating compliant joints and more human-like strategies: they allow controlled deviations, use stepping reactions, and sometimes deliberately "fall" into a new configuration. Yet they still confront the central fact that biological tissue can yield, deform and absorb impact in ways that metal and hard plastic struggle to emulate. Recent work, for example, on bioinspired hoppers[3], snakes, the Octobot[4] (see lower left), or wearable exoskeletons[5] (see right below) explicitly borrow from human body insights, using pneumatic artificial muscles and tuneable compliance to push robots closer to the human strategy of letting the mechanics "do some of the thinking":

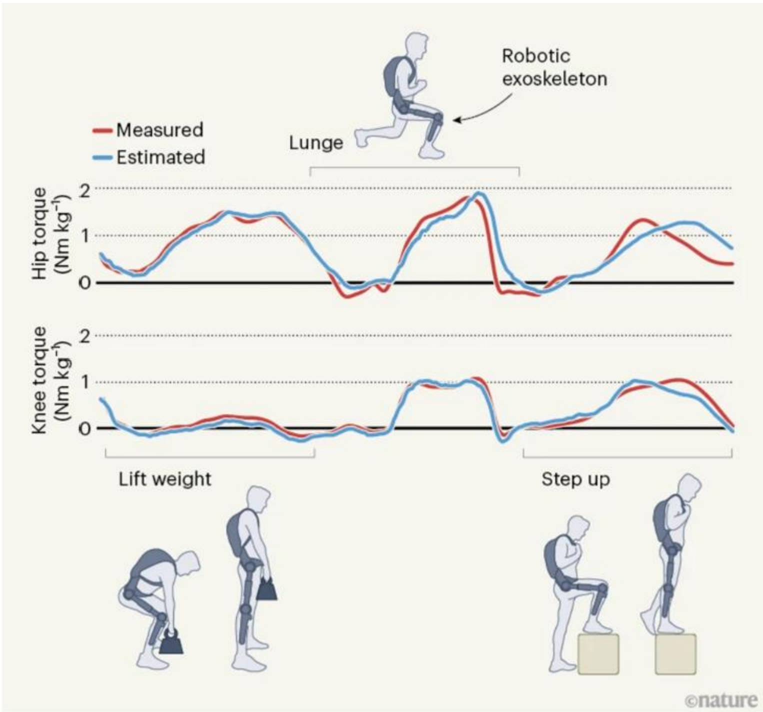

[1] Rashty, Aida Mohammadi Nejad, Maziar A. Sharbafi, Omid Mohseni, and André Seyfarth. "Role of Compliant Mechanics and Motor Control in Hopping – From Human to Robot." *Scientific Reports*, March 21, 2024. https://www.nature.com/articles/s41598-024-57149-0.
[2] Wu, Amy R. "Human Biomechanics Perspective on Robotics for Gait Assistance: Challenges and Potential Solutions." *Proceedings of the Royal Society B: Biological Sciences* 288, no. 1956 (2021). https://royalsocietypublishing.org/rspb/article/288/1956/20211197.
[3] S. Mohamad Hoseinifard and Majid Sadedel, "Standing Balance of Single-Legged Hopping Robot Model Using Reinforcement Learning Approach in the Presence of External Disturbances," *Scientific Reports* 14 (2024), https://www.nature.com/articles/s41598-024-83749-x.
[4] "7 Bio-Inspired Robots that Mimic Nature," *Machine Design*, August 18, 2017, https://www.machinedesign.com/mechanical-motion-systems/article/21835853/7-bio-inspired-robots-that-mimic-nature.
[5] Myunghee Kim and Matthew J. Major, "Robotic Exoskeleton Adapts to Changes in Leg Movements in Real Time," *Nature*, November 13, 2024, https://www.nature.com/articles/d41586-024-03546-4.

Co-adaptation

Experiments comparing human balance with humanoid or legged robots highlight another asymmetry: humans achieve upright stance and gait by blending passive stability, reflex loops and predictive control under substantial delays, whereas robots typically rely on fast, centralised optimisation. Numerous papers[6] have dissected[7] human bipedalism, showing that effective posture and gait control depends on integrating sensory information over time to predict how the centre of mass will move, rather than simply reacting to instantaneous errors. Humanoid robots are beginning to adopt similar strategies. Reinforcement-learning[8] based locomotion controllers can now yield gaits with human-like arm swing and coordinated whole-body motion, but they do so by building large, learned policies that map high-dimensional sensor streams to torques, rather than by reproducing the layered spinal–cerebellar–cortical hierarchy of the nervous system. Scientists now have started to discuss this in evolutionary terms: robots are being "fitted" into a tree of morphologies and control schemes, with performance niches that parallel but do not duplicate[9] those occupied by biological bodies. A recurring theme here is that the sharp separation between "plant" and "controller" in classical robotics has no real analogue in biology. Human movement augmentation studies, for example, show that when people wear supernumerary robotic limbs or powered exoskeletons, both the human nervous system and the device adapt over time, arriving at new joint movement patterns that cannot be attributed to either partner alone. Control is not something that sits on top of a fixed body; it is in fact negotiated between changing tissue and changing device.

Modern robots are slowly moving in this direction. Newer reports on humanoid locomotion describe systems in which morphology, actuation and control policies are co-designed, sometimes via automated search, so that the physical body and the learned controller form a mutually adapted pair. Flexible spines, series elastic actuators and variable-stiffness joints—now appearing in high-end humanoids, documented in both the popular[10] and technical publications[11]—are efforts to bring robots closer to the human pattern, where movement quality emerges from the resonance between compliant mechanics and predictive, learning-based control rather than from top-down specification alone.

Movement as behaviour versus trajectory

Underneath these technical differences lies a conceptual one that's now increasingly acknowledged. Human movement is not just a set of joint trajectories. In discussions of human–robot interaction and augmentation, it is framed as behaviour situated in goals, affordances and social context. A reach, a step, or a turn is evaluated not only in terms of kinematic accuracy, but in terms of comfort, effort, expressiveness and the sense of agency it preserves for the person. Until now however, robots, even when they move with strikingly human-like grace—as in demonstrations

[6] Jonathan Eden, Mario Bräcklein, Jaime Ibáñez, Deren Yusuf Barsakcioglu, Giovanni Di Pino, Dario Farina, Etienne Burdet, and Carsten Mehring, "Principles of Human Movement Augmentation and the Challenges in Making It a Reality," *Nature Communications* 13, article no. 1345 (2022), https://www.nature.com/articles/s41467-022-28725-7.

[7] Robert J. Peterka, "Comparison of Human and Humanoid Robot Control of Upright Stance," *Journal of Physiology-Paris* 103, no. 3-5 (2009), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2767299/.

[8] Ilija Radosavovic, Tete Xiao, Bike Zhang, Trevor Darrell, Jitendra Malik, and Koushil Sreenath, "Real-World Humanoid Locomotion with Reinforcement Learning," *Science Robotics* 9, no. 89 (April 17, 2024), https://www.science.org/doi/10.1126/scirobotics.adi9579.

[9] Chuanyu Yang, Kai Yuan, Shuai Heng, Taku Komura, and Zhibin Li, "Learning Natural Locomotion Behaviors for Humanoid Robots Using Human Bias," *IEEE Robotics and Automation Letters* 5, no. 2 (2020): 2610–2617, https://ieeexplore.ieee.org/document/8990011/.

[10] Futurobots Team, "Robot Locomotion Explained: How Humanoid Robots Learn to Walk on Two Legs," *Futurobots*, July 2025, https://www.futurobots.com/robot-locomotion-explained-humanoid-robots-learn-to-walk/.

[11] Alex Li, Samy Bengio, and Hugo Larochelle, "Learning Dexterous Human-like Walking with Latent Motor Primitives," *arXiv*, January 5, 2025, https://arxiv.org/abs/2501.02116.

of learned dancing[12] and cat-walk poses[13]—still typically treat motion as an optimisation problem in a high-dimensional state space, with objectives like stability margins, energy, or similarity to a reference motion-capture trace. The most forward-looking work starts to blur that line, embedding LLMs into control loops[14] so that movement decisions depend on semantic understanding of tasks and environments as well as on dynamics. Yet even there, what for a human is a lived act—moving a body that is felt from within—remains, for the robot, the execution of a policy over a body that is still, fundamentally, a very sophisticated object.

In this context, skeletal muscle is not a single "motor" at a joint; it is a fusiform cloud of fibres, each fibre a bundle of myofibrils, each myofibril a lattice of actin and myosin filaments whose sliding generates force. The same anatomical structure can generate tiny adjustments to keep you upright, or massive bursts of power in a sprint; it can hold an isometric contraction for minutes, or absorb a landing by acting as a brake. Crucially, muscles act across joints through tendons, and often span multiple joints at once (biarticular muscles), so that a single contraction redistributes forces around a whole limb in non-obvious ways. Robots however typically choose the opposite: discrete actuators at or near joints, often rotary electric motors coupled by gears, belts or harmonic drives. Where a muscle's contractile structure is distributed along its length, a motor tends to be concentrated, and the robot must then transmit torque through engineered components. The upshot is that human limbs change shape and stiffness continuously along their length, whereas a classic robotic arm is modelled as rigid segments separated by idealised hinges. Even so-called "artificial muscles" in robotics are designed to approximate the behaviour of a single muscle-tendon unit, not the entire, overlapping orchestra of fibres and fascia in a limb.

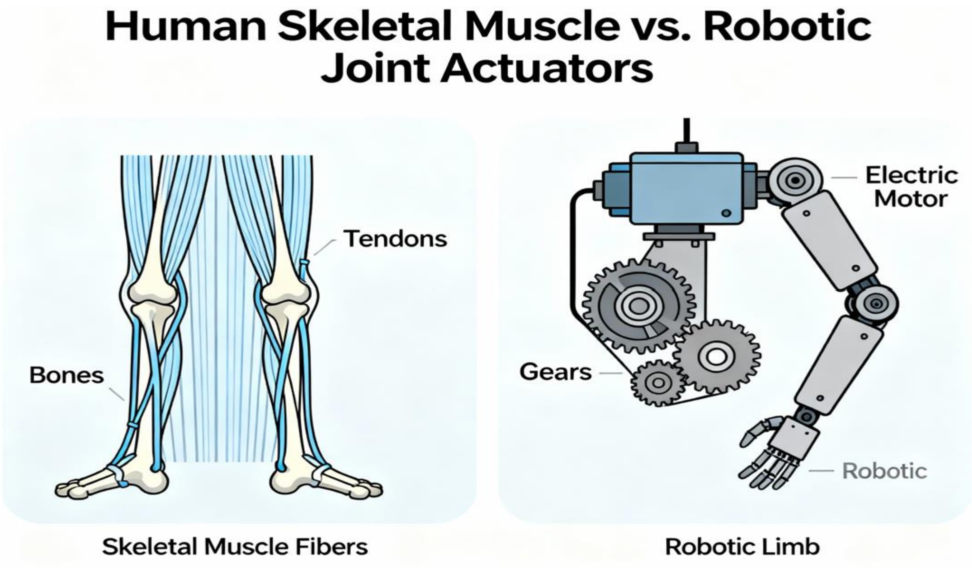

Source: Author's own work.

Because of their actuators, humans and robots literally inhabit different mechanics. The human skeleton is not a Meccano set; it is a dynamically braced framework whose effective degrees of freedom change as muscles tense or relax. Soft tissues wrap the joints, cartilage and menisci distribute loads, ligaments constrain extremes while allowing

[12] "Humanoid Robot Learns to Waltz by Mirroring People's Movements," *New Scientist*, January 16, 2025, https://www.newscientist.com/article/2464178-humanoid-robot-learns-to-waltz-by-mirroring-peoples-movements/.
[13] "XPENG's New IRON Humanoid: Catwalk, Solid-State Battery ...," *YouTube video*, uploaded by Xpeng, 2025, https://www.youtube.com/watch?v=RcNoQICjsLs.
[14] Shilong Sun, Chiyao Li, Zida Zhao, Haodong Huang, and Wenfu Xu, "Leveraging Large Language Models for Comprehensive Locomotion Control in Humanoid Robots Design," *Biomimetic Intelligence and Robotics* 4, no. 4 (2024): 100187, https://www.sciencedirect.com/science/article/pii/S2667379724000457.

tiny micro-motions, and all of this is bathed in fluid and supported by intra-abdominal pressure and fascial tension. When you reach for a cup, you do not rotate a sequence of ideal revolute joints: you shift centre of mass, deform your spine, roll and glide the humeral head in the socket, and engage trunk muscles to counterbalance. The standard industrial robot, by contrast, is designed to be well-posed in a textbook: a kinematic chain of rigid links with explicit joint definitions and a fixed number of degrees of freedom. This makes control tractable—you can write down the forward and inverse kinematics, compute Jacobians[15], optimise trajectories. But it also means the robot's body, as far as its software is concerned, is a collection of angles and link lengths, not a deformable, mass-laden, fluid-filled object. Compliance must be added back in artificially, via spring elements, series elastic actuators or software "impedance control", because the hardware begins conceptually as rigid.

Central control

Movement obviously requires not just actuators and levers, but a sense of self in space. The human body has an unusually rich proprioceptive apparatus: those aforementioned muscle spindles that sense length and rate of stretch, Golgi tendon organs that encode tension, joint capsule receptors for extreme positions, cutaneous receptors in the skin, plus two additional global senses—vision and the vestibular system in the inner ear—that inform the brain about head orientation, acceleration and gravity. All of this is fused in real time into something like an internal model of "me standing here", which is updated continuously, even without conscious awareness. Here, neural control of movement is not like a centralised script issuing discrete commands; it looks more like a partially decentralised, predictive, feedback-laden hierarchy. Spinal circuits can generate rhythmic locomotor patterns (central pattern generators) without cortical input. Cerebellum and basal ganglia provide fine-tuning, error correction, and selection of movement "programmes". Your cortex contributes planning, voluntary initiation, and integration with goals, language, and abstract reasoning. Importantly, the system is shot through with delays: it can take on the order of 100–200 ms from sensory input to muscular response, and yet humans maintain balance and coordinate complex movements in real time. One way the nervous system copes is by running internal forward models that predict the consequences of motor commands before feedback arrives and then correcting on the fly.

Robots, on paper at least, have an easier job. A motor encoder gives a high-resolution joint angle; IMUs and gyroscopes provide linear and angular accelerations; force-torque sensors measure contact loads. The robot controller can, in microseconds, integrate all these signals and estimate configuration and load. But what is missing is not bandwidth or accuracy—it is plastic, context-sensitive integration. A human standing on a moving bus will re-weight vestibular versus proprioceptive input depending on reliability and can maintain balance despite delays and noise. A humanoid robot, even with sophisticated balance controllers, still usually relies on a relatively narrow set of assumptions about contact geometry, timing and friction.

Robotic control, especially in classical frameworks, proceeds from a different starting point. The robot's equations of motion are explicitly known or identified; control laws are designed to stabilise desired trajectories, often assuming negligible delays and exact state information. Model predictive control, whole-body control, and learning-based controllers are bringing robots closer to the biological strategy of prediction plus correction, but the implementation is still relatively modular. There is typically a neat separation between perception, state estimation, planning, and control. In the human, these layers blur: sensory processing is influenced by motor plans; what you intend to do shapes what you perceive as relevant.

[15] Sizhe Lester Li, Annan Zhang, Boyuan Chen, Hanna Matusik, Chao Liu, Daniela Rus, and Vincent Sitzmann, "Controlling Diverse Robots by Inferring Jacobian Fields with Deep Networks," *Nature* 643 (2025): 89–95, https://www.nature.com/articles/s41586-025-09170-0.

Magnetism towards convergence

Despite these profound differences, there is a slow convergence. Soft robotics[16] seeks to replace rigid links with compliant, continuum structures; variable-stiffness actuators blur the line between muscle and motor; learning-based controllers approximate the internal models and predictive control that biological nervous systems evolved. Conversely, our understanding of the body is increasingly computational: we model muscles as actuators, joints as constraints, neural circuits as control policies. Each side uses the other as metaphor and blueprint.

Magnetic actuation, fluidic muscles, shape-memory alloys and other non-traditional mechanisms are part of this trend, but the deeper story is that engineers are trying to build bodies that move more like ours—not by copying anatomy one-to-one, but by adopting the principles of distributed actuation, compliant interaction, prediction under delay, and self-tuning. So, while the human body will likely remain the gold standard for versatile, robust, meaning-saturated movement, robots are now gradually shifting from being rigid manipulators in structured worlds to becoming embodied agents whose motion begins to hint, however faintly, at something like a lived presence in their environment. Here, magnetic actuation in soft robotics works by embedding magnetic particles or small permanent magnets inside compliant materials (silicones, elastomers, gels) and then shaping those materials with external magnetic fields so that they bend, twist, and stretch. External coils or permanent magnet arrays supply the field; because magnetic fields penetrate most materials, they can drive motion without any physical tether, allowing the body of the robot to be entirely soft and unencumbered by cables or gear-trains. Some designs[17] combine magnetism with phase-change polymers, taking advantage of advances in material science, so that the same structure can both move and actively tune its own stiffness, again more like biological muscle that can switch from relaxed to rigid in milliseconds.

Thus while the range and subtlety of motion in a mechanical joint is usually constrained by a few factors (hard end-stops, limited joint geometries, and the fact that force must be transmitted along lines—shafts, belts, tendons—that introduce backlash, slack, and wear), soft, magnetically actuated structures avoid many of these limits because the "joint" is spread throughout the material; large bending angles, coiling, and folding can be induced by simply reorienting the field, with no discrete hinge at all. Several recent systems show why magnetism can plausibly give greater effective range of movement than traditional linear actuators:

- Magnetically actuated soft crawlers and coiling graspers achieve large deformation angles by programming[18] magnetic domains within an elastomer sheet, producing high-strain bending and twisting under relatively modest fields.
- Magnetic "muscles"[19] built from soft composites have been demonstrated that outperform biological muscle in certain metrics, combining high stretch, large stroke, and rapid response, and can drive continuum robotic arms and hands that bend and curve along their entire length rather than about a single axis.

[16] Yeongju Jung, Kangkyu Kwon, Jinwoo Lee, and Seung Hwan Ko, "Untethered Soft Actuators for Soft Standalone Robotics," *Nature Communications* 15, no. 3510 (2024), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11045848/.
[17] Hongman Wang, Zhihong Zhu, Huan Jin, Rui Wei, Lili Bi, and Wenling Zhang, "Magnetic Soft Robots: Design, Actuation, and Function," *Journal of Alloys and Compounds* 922 (2022), https://www.sciencedirect.com/science/article/pii/S092583882202610X.
[18] Dharmi Chand and Sivakumar M. Srinivasan, "Pattern Architected Soft Magnetic Actuation," *Soft Matter* 21, no. 6 (2025): 1072–1084, https://pubs.rsc.org/en/content/articlelanding/2025/sm/d4sm01450b.
[19] Minho Seong, Kahyun Sun, Somi Kim, Hyukjoo Kwon, Sang-Woo Lee, Sarath Chandra Veerla, Dong Kwan Kang, Jaeil Kim, Stalin Kondaveeti, Salah M. Tawfik, Hyung Wook Park, and Hoon Eui Jeong, "Multifunctional Magnetic Muscles for Soft Robotics," *Nature Communications* 15 (2024): 7929, https://www.nature.com/articles/s41467.

- At the micro-scale, magnetically steered robots can snake through tortuous vascular or anatomical pathways because forces and torques are applied directly to embedded magnetic elements rather than transmitted through stiff catheters, alleviating[20] many of the geometric constraints of cable-driven systems.

This is philosophically interesting: instead of a robot moving like a set of rigid symbols in mechanical "syntax," the whole body becomes a deformable field of potential motion, with the magnetic field playing a role analogous to neural activation patterns recruiting different muscle fibres. The robot's shape is no longer just a set of discrete joint angles; it is a continuous configuration in a high-dimensional space defined by how each region of the soft material responds to the magnetic field.

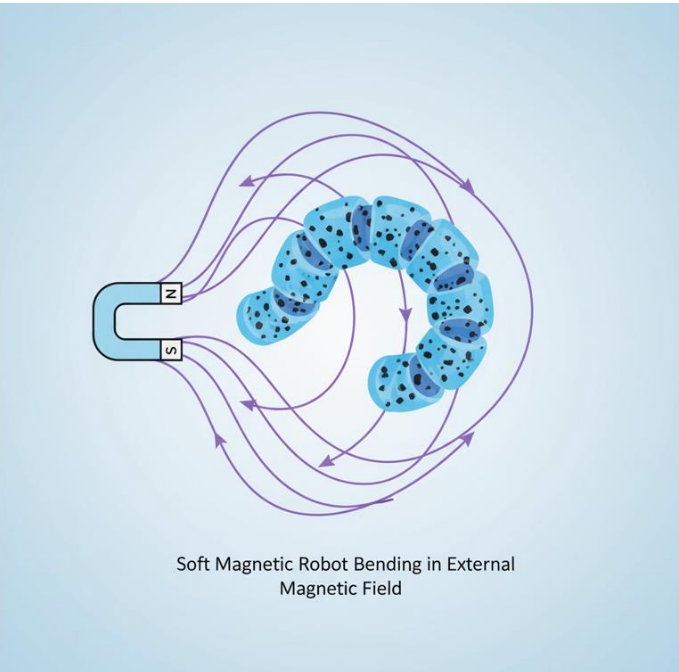

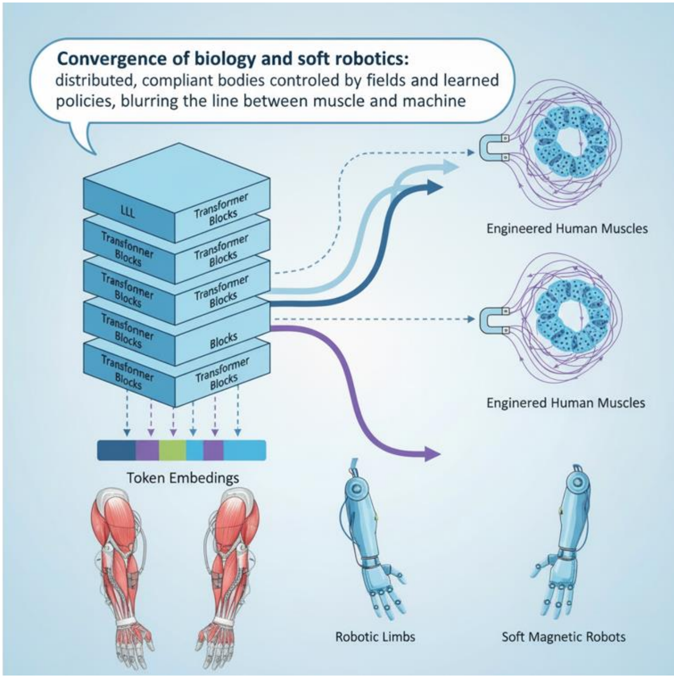

Source: Author's own work.

Despite these advantages, magnetism does not automatically surpass traditional engineering in every dimension. External magnetic field generators—Helmholtz or Maxwell coil systems, rotating permanent magnets—can be bulky, energy-intensive, and impose spatial constraints, especially at larger scales where strong fields over large volumes are required. Scaling laws also matter: magnetic forces drop with distance, and designing field gradients that produce strong, controllable forces deep inside the body of a large robot is non-trivial. There are also material limitations: embedding many magnetic particles can increase density and damping, potentially slowing motion or reducing efficiency, and repeated large deformations demand elastomers that can withstand fatigue without losing magnetization patterns. For high-load, high-speed tasks—industrial arms welding car frames, for example—traditional motors and geared transmissions still provide unmatched power density and robustness. The frontier is therefore hybrid. Research is moving toward robots where:

- Magnetic soft actuators provide large, gentle, tissue-like movements with high compliance.

[20] Tiantian Kong, Qitong Zheng, Jiarong Sun, Chunxiao Wang, Huibin Liu, Zhizheng Gao, and Wenguang Yang, "Advances in Magnetically Controlled Medical Robotics: A Review of Actuation Systems, Continuum Designs, and Clinical Prospects for Minimally Invasive Therapies," *Micromachines* 16, no. 5 (2025): 561, https://www.ncbi.nlm.nih.gov/pmc/articles/PMC12114355/.

- Traditional motors, pneumatic systems, or electrohydraulic actuators supply raw power or long-range motion where needed.
- Embedded sensing and control loops use the same magnetic fields for both actuation and perception, reminiscent of the brain's integration of signalling and sensing in neural tissue.

Philosophically, magnetically actuated robots push robotics closer to an embodied, field-based view of intelligence and movement. The human brain does not simply send binary commands down wires; it modulates fields of electrochemical activity across distributed, soft tissue, and muscles respond by altering their continuous mechanical state—length, stiffness, and tension. Magnetic soft robots mimic this to a surprising degree: a global field shapes a distributed material, and the resulting form and motion emerge from the interaction of field, structure, and environment rather than from rigid, pre-specified joint trajectories.

Different, worse or better?

There are some clear areas where robots excel. Modern underwater robots routinely operate at depths of around 1,000 metres with dexterous manipulation, and specialized soft or tethered vehicles have gone past 10,000 metres into the deepest ocean trenches[21], with designs often mimicking nature[22], including actually being edible[23]. Similarly, mobile robots have been sent into reactor buildings at Fukushima and similar facilities to map radiation, take samples, and perform simple manipulations in places where human workers would exceed safe lifetime radiation doses in minutes. The same principle applies in space: robotic probes have landed on Venus, flown through Jupiter's intense radiation belts, and operated for years in vacuum and extreme temperature swings, where unprotected humans could not survive and even heavily shielded crewed missions would be prohibitively risky or expensive. Similarly, neuromorphic cameras[24] — which respond to changes in brightness and are inspired by the retina — can be used in drones. The image on the right illustrates the visual input the drone receives from the camera: red indicates pixels that are getting darker and green indicates those that are getting brighter. On the left, the miniature curved compound eye, called CurvACE[25], was inspired by the eyes of insects:

[21] Noah Baker, "How a Soft Robot Survived the Deepest Ocean on Earth," *Nature*, March 5, 2021, https://www.nature.com/articles/d41586-021-00605-y.

[22] Guorui Li, Tuck-Whye Wong, Benjamin Shih, Chunyu Guo, Luwen Wang, Jiaqi Liu, Tao Wang, Xiaobo Liu, Jiayao Yan, Baosheng Wu, Fajun Yu, Yunsai Chen, Yiming Liang, Yaoting Xue, Chengjun Wang, Shunping He, Li Wen, Michael T. Tolley, A-Man Zhang, and Cecilia Laschi, "Bioinspired Soft Robots for Deep-Sea Exploration," *Nature Communications* 14, no. 7097 (2023), https://www.nature.com/articles/s41467-023-42882-3.

[23] Shuhang Zhang, Bokeon Kwak, Ruihao Zhu, Markéta Pankhurst, Lu Zhang, Remko M. Boom, and Dario Floreano, "Edible Aquatic Robots with Marangoni Propulsion," *Nature Communications* 16, article no. 4238 (2025), https://www.nature.com/articles/s41467-025-59559-8.

[24] "How the Natural World Is Inspiring the Robot Eyes of the Future," *Nature*, May 29, 2025, https://www.nature.com/articles/d41586-025-01660-5.

[25] Min Su Kim, Min Seok Kim, Mincheol Lee, Hyuk Jae Jang, Do Hyeon Kim, Sehui Chang, Minsung Kim, Hyojin Cho, Jiwon Kang, Changsoon Choi, Jung Pyo Hong, Do Kyung Hwang, Gil Ju Lee, Dae-Hyeong Kim, and Young Min Song, "Feline Eye-Inspired Artificial Vision for Enhanced Camouflage Breaking under Diverse Light Conditions," *Science Advances* 10, no. 38 (2024): eadp2809, https://www.science.org/doi/10.1126/sciadv.adp2809.

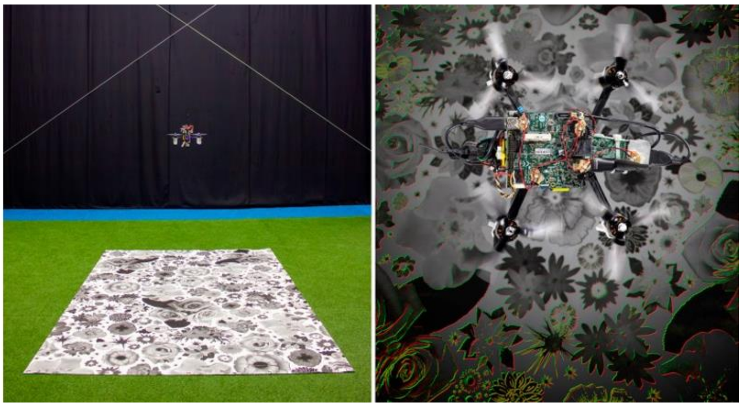

Source: Author's own work.

In terms of everyday tasks, robots are now matching or exceeding[26] human mobility on a few narrow metrics like top running speed and precise parkour, but they still lag badly on agility, robustness and efficiency in unstructured environments. Normal human walking speed is about 3–4 mph (1.3–1.8 m/s), and recent full-size humanoids report typical walking around 3.4 mph, so on level ground they now look broadly human-like. Fast humans jog at roughly 4–6 mph, whereas state-of-the-art humanoids[27] have reached about 5.5–7.3 mph in controlled tests, with some newer prototypes in the 6–6.2 mph range, putting peak straight-line speed into the recreational-runner band rather than superhuman territory.

Quadruped robots illustrate[28] the gap in agility[29]. In dog-style obstacle courses, small dogs clear a calibrated course in about 10 seconds, while a similar-sized quadruped robot currently takes roughly twice as long, suggesting that even in a compact, scripted task, animal-level agility is still at least a factor-of-two away. Humans still dominate in versatility, efficiency and even robustness. Even with strong learning-based controllers, leading quadruped robots complete validated agility benchmarks at about half the speed of untrained dogs and remain markedly more brittle to slips, odd footholds and sensor noise. A human can walk tens of kilometres on roughly 20 W of metabolic power, while comparable humanoids rely on batteries with far lower energy density and require frequent recharging after relatively short, high-power runs. Today's best robotic locomotion results are almost always achieved under tightly constrained conditions—flat floors, known obstacle layouts, or carefully tuned policies—whereas humans retain near-baseline gait on ice, mud, clutter, stairs and in crowds without any redesign of hardware or control.

The trend line is steep, though. Over roughly a decade, top bipedal speed has moved from slow walking to stable running at 5–7+ mph and parkour-style jumps, with new speed records announced year-on-year. On the quadruped side, dog-level agility has become an explicit benchmark: within just a few years, performance has advanced from "cannot complete the course" to "finishes in about twice the dog time," representing order-of-magnitude improvements at the controller level even if the animal gap is not yet closed. Thus, robots have now reached recreational-human levels in straight-line speed and in specific choreographed stunts, but in agility, robustness and

[26] Xiangji Wang, Wei Guo, Siyu Yin, Sen Zhang, Fusheng Zha, Mantian Li, Pengfei Wang, Xiaolin Li, and Lining Sun, "Walking Control of Humanoid Robots Based on Improved Footstep Planner and Whole-Body Coordination Controller," *Frontiers in Neurorobotics* (2025), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11885507/.
[27] Liz Hughes, "Humanoid Robot Breaks Walking Speed Record," *IoT World Today*, March 6, 2024, https://www.iotworldtoday.com/robotics/humanoid-robot-breaks-walking-speed-record.
[28] Peter Grad, "Barkour Benchmark Measures Quadruped Robot Agility," *TechXplore*, June 1, 2023, https://www.techxplore.com/news/2023-05-barkour-benchmark-quadruped-robot-agility.html.
[29] Matthew T. Mason and Josh T. Huang, "A Benchmark for Agile Quadruped Locomotion: The Barkour Dataset," *arXiv*, May 22, 2023, https://arxiv.org/abs/2305.14654.

energy economy they remain one to two qualitative steps behind ordinary humans and even small animals. I wouldn't like to challenge however, Boston Dynamics' Atlas robot[30] in any sport:

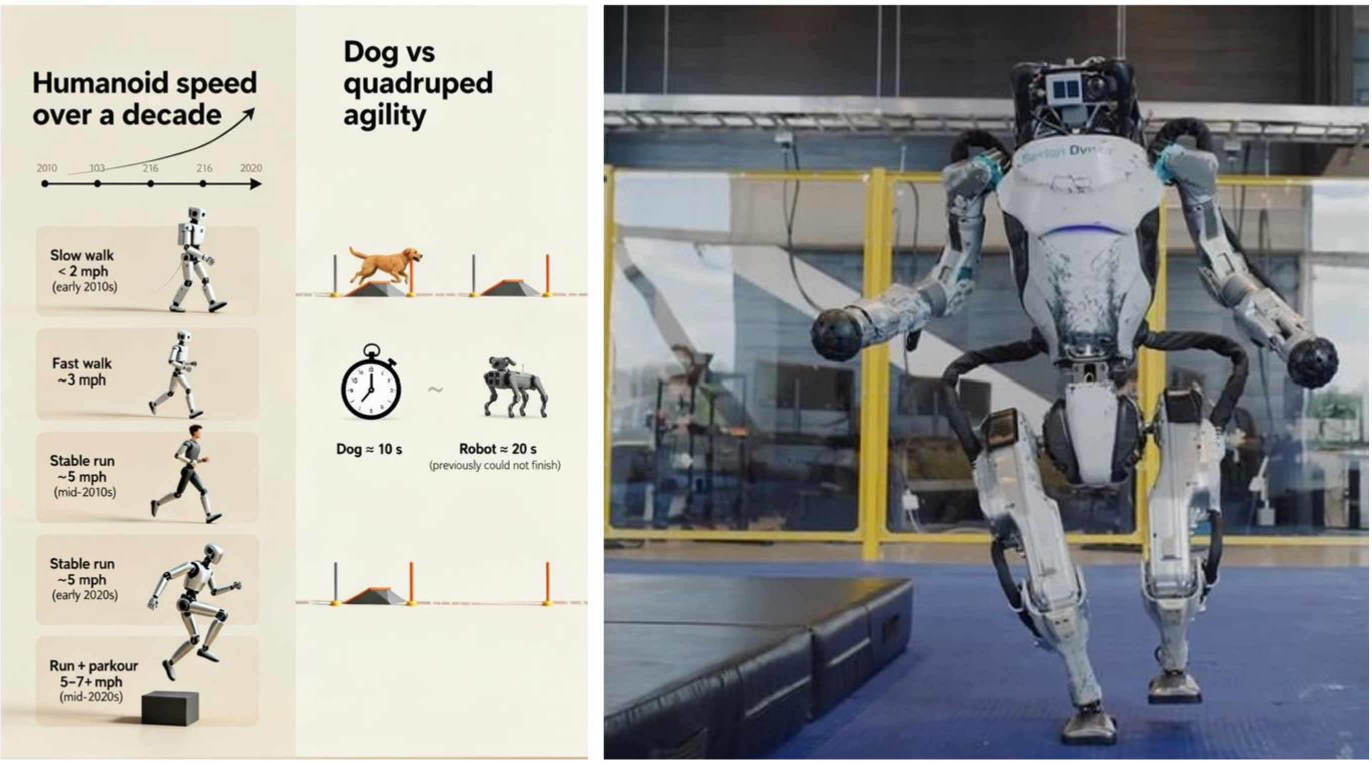

Source: Author's own work.

LLM philosophy

Underneath the engineering, again rests a philosophical split. The robot's body, as used by its control system, is essentially symbolic. It is a set of coordinates in a configuration space, a graph of possible contacts and constraints, a series of parameters in a dynamics model. The "movement" is then a path through that abstract space, computed to satisfy objectives and constraints. The physical chassis merely realises that path; if the model is wrong, the controller adapts or fails, but the body remains an object in equations.

For a human, the body is not just a thing that moves—it is the locus of experience. Movement is not the execution of a script, but the unfolding of a lived intention in the world. To raise a hand is not to set shoulder flexion to 60 degrees; it is to greet, to shield, to point, and those meanings are inseparable from the sensation of the muscles contracting and the weight of the arm in gravity. Philosophers like Merleau-Ponty called this the body-subject: a body that is not owned by a mind, but is itself the means by which the world is disclosed. That is why the same physical motion—a flinch, a reach, a touch—can feel utterly different depending on context; the geometry is similar, but the lived movement is not.

Roboticists are beginning to grapple with this gap as AI systems gain richer internal models of the world. A robot guided by an LLM that predicts human reactions, plans long-term strategies, and reasons about tasks starts to edge

[30] "Boston Dynamics Humanoid Atlas Robot Performs Parkour in Impressive New Video," *Dornob*, August 2021, https://dornob.com/boston-dynamics-humanoid-atlas-robot-performs-parkour-in-impressive-new-video/.

closer to something like a "movement with meaning". Yet without organic embodiment—without tissue that feels pain, fatigue, effort, and pleasure—the robot's motion remains, for now, entirely on the side of execution rather than intention. Its kinematic chains may eventually match or surpass ours in flexibility and precision, but the phenomenological content of that movement is absent. Recently, a team established a framework[31] for creating small-scale soft robots with enhanced environmental intelligence through tightly integrated sensing, actuation, and decision-making architectures, two key elements being mimicry of nature and new materials:

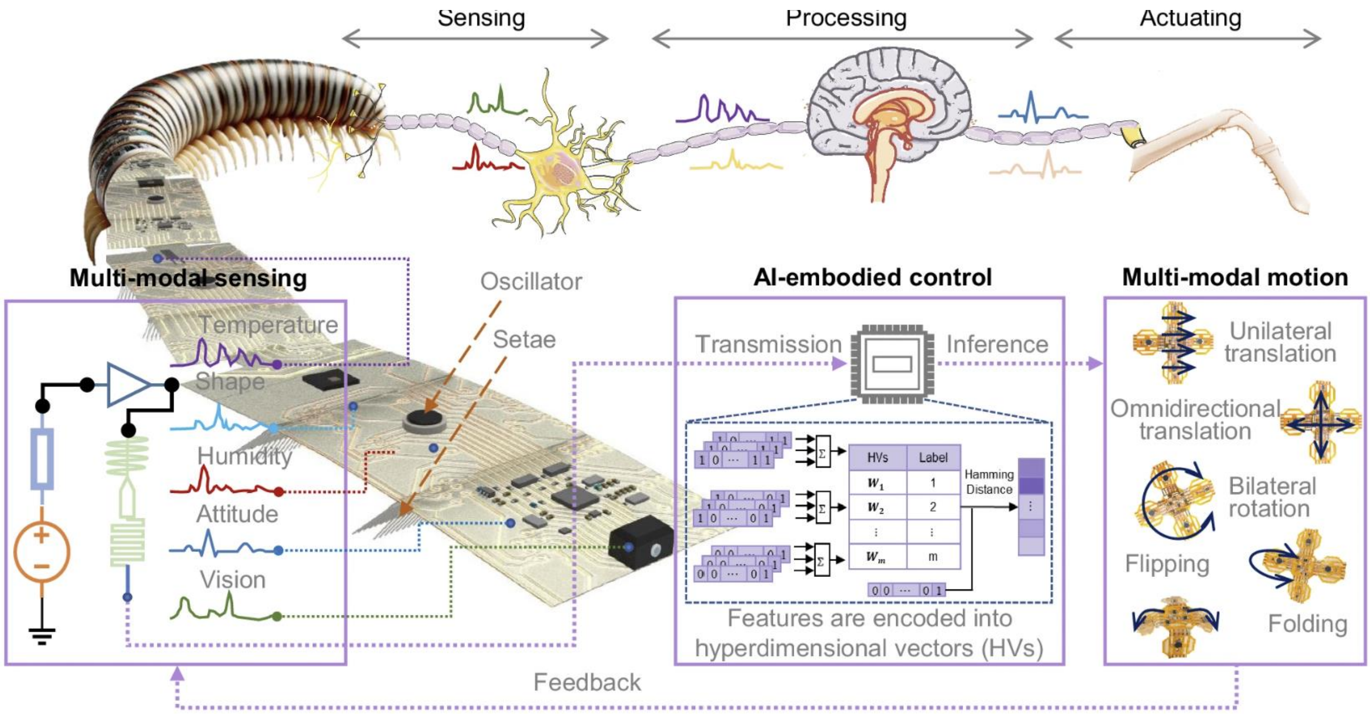

Source: Author's own work.

Another group reports an embodied large-language-model-enabled robot (ELLMER)[32] framework, utilizing GPT-4 and a retrieval-augmented generation infrastructure, to enable robots to complete long-horizon tasks in unpredictable settings. However, no matter how complex its behaviour, any robot has only a limited number of sensors that pick up information about the environment (cameras, radar, lidar, microphones and carbon monoxide detectors, to name a few examples). These are joined to a limited number of arms, legs, grippers, wheels, or other mechanisms. Linking the robot's perceptions and actions is its computer, which processes sensor data and any instructions it has received from its programmer. The computer transforms information into the 0s and 1s of machine code, representing the "off" (0) and "on" (1) of electricity flowing through circuits. Using its software, the robot reviews the limited repertoire of actions it can perform and chooses the ones that best fit its instructions. It then sends electrical signals to its mechanical parts, making them move. Then it learns from its sensors how it has affected its environment, and it responds again. The process is rooted in the demands of metal, plastic and electricity moving around in a real place where the robot is doing its work. When you give an LLM a prompt, a question, request or instruction, the model

[31] Junfeng Li, Zhangyu Xu, Nanpei Li, Kaijun Zhang, Guangyong Xiong, Minjie Sun, Chao Hou, Jingjing Ji, Fan Zhang, Junwen Zhong, and YongAn Huang, "AI-Embodied Multi-Modal Flexible Electronic Robots with Programmable Sensing, Actuating and Self-Learning," *Nature Communications* 16, no. 8818 (2025), https://www.nature.com/articles/s41467-025-63881-6.

[32] Ruaridh Mon-Williams, Gen Li, Ran Long, Wenqian Du, and Christopher G. Lucas, "Embodied Large Language Models Enable Robots to Complete Complex Tasks in Unpredictable Environments," *Nature Machine Intelligence* 7, no. 4 (2025): 592–601, https://www.nature.com/articles/s42256-025-01005-x.

converts your words into numbers, the mathematical representations of their relations to one another. This maths is then used to make a prediction: given all the data, if a response to this prompt already existed, what would it probably be? The resulting numbers are converted back into text. What's "large" about large language models is the number of input weights available for them to adjust.

Thus, a robot can't delicately fillet a salmon if it has only a two-fingered gripper with which to handle objects. If asked how to make dinner, the LLM, which draws its answers from billions of words about how people do things, will suggest actions the robot can't perform. Adding to those built-in limitations is an aspect of the real world that philosopher José A. Benardete called "the sheer cussedness of things." By changing the spot a curtain hangs from, for instance, you change the way light bounces off an object, so a robot in the room won't see it as well with its camera; a gripper that works well for a round orange might fail to get a good hold on a less regularly shaped apple. As the robotics researchers Singh, Thomason and their colleagues put it, "the real world introduces randomness." Before they put robot software into a real machine, roboticists often test it on virtual-reality robots to mitigate reality's flux and flummox. "The way things are now, the language understanding is amazing, and the robots suck," says Stefanie Tellex, half-jokingly. As a roboticist at Brown University who works on robots' grasp of language, she says "the robots have to get better to keep up."[33]

When an LLM comes up with a realistic plan for cooking a meal, "it seems like there's some kind of reasoning there," says Anirudha Majumdar, a professor of engineering at Princeton. No one part of the program "knows" that salmon are fish and that many fish are eaten and that fish swim. But all that knowledge is implied by the words it produces. "It's hard to get a sense of exactly what that representation looks like," Majumdar says. "I'm not sure we have a very clear answer at this point." In one experiment, Majumdar made use of an LLM's implicit map of the world to address what they call one of the "grand challenges" of robotics: enabling a robot to handle a tool it hasn't already encountered or been programmed to use. Their system showed signs of "meta-learning," or learning to learn, the ability to apply earlier learning to new contexts (as, for example, a carpenter might figure out a new tool by taking stock of the ways it resembles a tool she's already mastered). AI researchers have developed algorithms for meta-learning, but in the Princeton research, the strategy wasn't programmed in advance. No individual part of the program knows how to do it, Majumdar says. Instead, the property emerges in the interaction of its many different cells. "As you scale up the size of the model, you get the ability to learn to learn." [34]

Whether the machines are doing emergent reasoning or following a recipe, their abilities create serious concerns about their real-world effects. Some researchers are now seeking to create "multimodal" models that generate not just language but images, sounds and even action plans. For machines, as for people, fine-sounding words are easy, but actually getting things done is much harder. The bottleneck is at the level of simple things like opening drawers and moving objects. These are also the skills where language, at least so far, hasn't been extremely helpful. For now, the biggest challenges posed by LLMs arise not from robot bodies housing them but rather from the way they copy, in mysterious ways, much that human beings do well. An LLM, is perhaps a kind of gestalt of the Internet. So, all the good parts of the Internet are in there somewhere. And all the worst parts of the Internet are in there somewhere, too. Maybe then compared with LLM-made phishing e-mails and spam or with LLM-rendered fake news, putting one of these models in a robot is probably one of the safest things you can do with it.

[33] David Berreby, "Scientists Are Putting ChatGPT Brains Inside Robot Bodies. What Could Possibly Go Wrong?" Scientific American, February 29, 2024, https://www.scientificamerican.com/article/scientists-are-putting-chatgpt-brains-inside-robot-bodies-what-could-possibly-go-wrong/
[34] Ibid.

Moravec and Conclusions

There are a wide variety of other differences. Two immediate examples spring to mind. First, a human body moves on about 20 watts of continuous power—less than a dim light bulb—yet can walk tens of kilometres, climb, lift, and maintain posture. The metabolic pathways that supply this energy, from ATP stores to oxidative phosphorylation, are intimately integrated with the muscles and vascular system, and scale dynamically with demand. Elastic structures like tendons store and release energy with each step, acting as biological springs; the arch of the foot, the Achilles tendon, and passive muscle elasticity all contribute to the extraordinary efficiency of running or hopping. Moreover, the system repairs itself: micro-tears in muscle stimulate hypertrophy, bones remodel in response to load, and neural circuits rewire after injury. Robots, even the most advanced, usually rely on external power supplies or batteries whose energy density and dynamic response are poor in comparison. Actuators often waste energy as heat, particularly when holding loads statically. Some designs incorporate mechanical springs and energy-recycling schemes, but these are add-ons rather than intrinsic properties of the material. And failure modes differ radically: a torn muscle hurts but may heal; a stripped gear or cracked linkage requires human intervention and replacement. From the viewpoint of movement as a life-long process, the body is both actuator and mechanic, constantly tuning and repairing itself in the background.

Second, humans excel at self-repair through innate biological processes like inflammation, stem cell activation, and tissue remodeling, where damaged cells are cleared via apoptosis and replaced seamlessly without external parts, enabling recovery from wounds, fractures, or even organ stress in days to weeks. Robots, by contrast, depend on human intervention for fixes—ordering bespoke spare components like servos or circuits, which introduces delays, costs, and incompatibility risks, leaving them inert until repaired. Even with large language models powering robotic intelligence, a hardware failure demands physical swaps, whereas LLMs themselves offer redundancy: simply boot the same software on another computer, underscoring how biological embodiment trumps mechanical fragility in resilience.

The table below contrasts the meaning and resilience of movement: human actions are lived, expressive and socially interpreted, emerging from a controller–body system that can self-repair and reorganise after injury, whereas robot motion is typically an optimised trajectory in an abstract state space, with mechanical failures needing external repair and no intrinsic experiential content attached to the movement itself. Around that philosophical core sit the more mechanical differences. The human body realises motion through a soft, as aforementioned heterogeneous musculoskeletal architecture whose passive dynamics and neural control are deeply intertwined, while most robots are assembled as rigid kinematic chains whose motions are first prescribed in configuration space and only then realised in hardware. Human limbs behave like tuneable springs and struts, recycling energy through tendons and arches so that walking and running remain stable and efficient on minimal continuous power, whereas robots must emulate such behaviour in software, enforcing stability with active torque control and paying an energy penalty for every correction. Humans further enjoy dense proprioception from muscles, tendons, joints and skin, fused with vestibular and visual input into a predictive internal model that can tolerate noise, delays and even partial damage while still supporting fine balance and dexterity. Robotic systems instead work with more modular sensor suites—encoders, inertial units, force–torque cells, cameras—whose outputs are combined by explicit estimation and control algorithms, typically under a clean separation between "plant" and "controller". Finally, adaptation and failure play out differently: human motor skills are shaped by lifelong co-adaptation of tissue and nervous system, with injuries triggering biological repair and neural reorganisation, whereas a robot's morphology is usually fixed post-manufacture, most learning happens only in software, and structural failures demand external intervention rather than any intrinsic process of self-healing.

Dimension	Human body movement	Robot movement
Primary actuators	Skeletal muscles: distributed, fibre-based actuators with variable stiffness and non-linear force length and velocity properties.	Electric motors, hydraulics, pneumatics or artificial muscles, typically discrete at joints with more linear and predictable characteristics.
Body structure	Compliant, heterogeneous musculoskeletal system with bones, cartilage, tendons, fascia and fluid; effectively change with muscle activation.	Usually modelled as rigid links with fixed joints; compliance added via springs, series elastic actuators or software impedance control.
Passive dynamics	Gait and hopping exploit natural spring behaviour of limbs and tendons, recycling energy and stabilising motion.	Many robots must emulate passive dynamics in control; stability often enforced by active torque control rather than intrinsic mechanics.
Proprioception and sensing	Rich proprioception (muscle spindles, Golgi organs, joint and skin receptors) plus vestibular and visual input, fused into an internal body model.	Encoders, IMUs, force/torque and vision sensors provide precise but modular signals; integration depends on designed estimation algorithms.
Balance control	Uses multi-level strategies (ankle, hip, stepping) with predictive and reflex components, tolerating delays and noise.	Often relies on fast optimisation (e.g., ZMP, whole-body control, RL policies) with stricter assumptions about contact and terrain.
Energy and efficiency	Locomotion powered by biochemical pathways; tendons and arches release energy, enabling long-range movement on ~20 W.	Batteries or external power; energy density and losses in actuators make comparable agility and endurance challenging.
Adaptation and learning	Lifelong motor learning; nervous system and body co-adapt (plasticity, hypertrophy, bone remodelling).	Controllers can be retrained or updated; morphology usually fixed post-manufacture, though co-design is emerging.
Controller-body coupling	No clean separation: neural control and mechanics evolved together; control policies depend on body physics.	Classical split between plant and controller though increasingly explored now in research.
Movement representation	Movements experienced as intentions and actions (reaching, grasping, walking) with embedded meaning and affect.	Typically represented as trajectories and policies in state space, optimised for stability, efficiency or task success.
Social and expressive aspects	Gait, gestures and posture convey identity, mood and intent; small kinematic differences alter social perception.	Human-like motion improves acceptance but not if timing and expressiveness are slightly off.
Failure and repair	Injuries can self-heal; motor patterns reorganise after damage (neuroplasticity, compensation).	Mechanical failures require external repair; software can re-route control but structural self-repair is rare.

Source: Author's own work.

Delving into philosophy, Moravec's paradox is the unsettling observation that it is relatively easy to build machines that do things we regard as intellectually demanding—playing chess, proving theorems, passing exams—while it is extraordinarily hard to give them the effortless sensorimotor skills of a toddler, like walking across a cluttered room or catching a fly. The paradox arises because evolution has been optimising those "low-level" bodily skills for hundreds of millions of years, distilling perception, balance and movement into dense, unconscious circuitry, while abstract reasoning is a very recent bolt-on in evolutionary time and so can be approximated with relatively simple symbolic or statistical tricks. Frogs are a vivid example: with a nervous system far smaller than a modern AI model, a frog can seamlessly integrate vision, posture, and tongue strikes to catch moving prey in real time, all while maintaining balance on a swaying reed. What looks cognitively trivial—just "being a body" in the world—turns out to be computationally profound. When comparing humans and robots, Moravec's paradox explains why we now have systems that can write code and summarise Nature or Science papers, yet still struggle to achieve the fluid, frog-like integration of seeing, orienting, and acting that underpins something as mundane as walking down some stairs while talking to a friend.

Source: Author's own work.

On the surface, the movements of a body and a robot can be made to look similar, but they arise from radically different relations between matter, meaning, and control. The human body moves as a lived subject: muscles, tendons, joints and fascia express intentions that are already shot through with significance—reaching, recoiling, caressing are not neutral trajectories but modes of being-in-the-world, as Merleau-Ponty would put it. A robot, by contrast, typically moves as an object executing a plan defined in an abstract configuration space: a set of coordinates is optimised under constraints, and the hardware instantiates that path without any intrinsic "aboutness" beyond its task specification. This is why even very smooth humanoid robots can feel uncanny: their kinematics may approximate ours, but they lack the dense web of bodily memory, vulnerability, and anticipation that makes a human movement recognisably human. Philosophically, the contrast is not just between soft tissue and metal, but between movement as expression and movement as implementation. In the biological case, control and embodiment are entangled—the nervous system is shaped by, and continually reshapes, the body it commands—whereas in most current robots, body and controller remain separable modules, and motion is something imposed on matter rather than emerging from it. This difference underlies our ethical and emotional responses: we see in another person's gait or gesture a fellow subject whose movements can reveal joy, pain or intent, while in the robot we see, at best, a mirror engineered to reflect these patterns back to us without sharing the inner life they ordinarily presuppose.