The Body of Intelligence

from the hands that made tools to soft robots

Science · 30 min · Felden Vareth · 17-08-2026

Contents

From primate dexterity and octopuses to artificial perception and soft robotics: what happens when body, senses and the way intelligence perceives the world become design variables.

For millions of years, any intelligence capable of transforming its environment was subject to one unavoidable condition: it needed a body with which to act upon it.

Long before hammers, wheels, computers or robots existed, there were fingers closing around prey, beaks moving objects, trunks exploring, jaws carrying and tentacles capable of wrapping around whatever they encountered. Thinking made it possible to anticipate the world; manipulating it made it possible to discover what happened when that world was changed.

In our own lineage, the relationship between brain, hands and tools is so close that it is tempting to tell a simple story: first came a hand capable of precise grasping, and then an increasingly large brain learned how to use it. Evolution rarely works in such a linear way. The available evidence points instead to a two-way relationship. In primates, greater manipulative ability, certain hand proportions and the enlargement of some brain regions appear to have evolved in association with one another. The hand did not create intelligence by itself, but a body capable of more complex actions turned the environment into an increasingly rich problem for the brain.

Perhaps the question is how much a brain could learn once it had a body capable of doing more and more things.

Then comes a historical break. Artificial intelligence does not have to inherit an anatomy built over millions of generations. It can exist in a data centre, control an industrial arm, drive a vehicle, inhabit a humanoid robot or be distributed across hundreds of machines. If artificial intelligence ever becomes capable of operating autonomously in the physical world, there is no reason to assume that its optimal body should resemble ours.

This is where a new and suggestive discipline begins: **soft robotics**. Instead of building only rigid metal machines articulated around motors and shafts, engineering is learning to create robots that bend, inflate, deform, change shape, wrap around delicate objects or pass through spaces that a conventional mechanism could not negotiate.

We may be entering a strange stage in the history of intelligence. Biological evolution had to build brains that worked inside the bodies each lineage inherited. We are beginning to build intelligences for which the body itself can become a design variable.

A hand does not build a brain, but each can push the other

Primates offer one of the best settings in which to study this relationship because they combine vision, coordination, social learning and an extraordinary capacity for manipulation. One of the key anatomical features is the **prehensile hand and, in many primates, the opposable thumb**. Opposing the thumb means moving it across the palm until its pad faces those of the other fingers, especially the index finger. This movement allows an object to be held between opposing surfaces and enables precision grips very different from those achieved by simply closing all the fingers in the same direction. Not all primates possess the same degree of opposition. In humans, the thumb is proportionally long, mobile and muscular, a combination that greatly expands the precision with which we can manipulate objects and tools.

A study published in 2025 compared 95 living and fossil primate species and found a significant relationship between relatively longer thumbs — a feature associated with greater manipulative ability — and larger brains. The relationship appeared to be associated with the size of the neocortex.

This does not mean that a long thumb automatically produces a large brain. Nor does it demonstrate a simple causal chain. What is interesting is that anatomy and the nervous system appear to form part of the same evolutionary history. If an animal can execute more precise movements, new possibilities appear:

exploring, selecting, opening, breaking, combining, carrying or using objects. Each possibility introduces problems of coordination, prediction and learning.

Another comparative study, published in 2016 and based on 36 non-human primate species, found that manipulation complexity was related to brain size and measures of cognitive performance. The authors proposed a feedback dynamic: brains with greater control capacity enable more complex actions, and those new possibilities may in turn favour greater cognitive demands.

In our own lineage, another decisive element was added: tools ceased to be merely occasional objects and eventually became part of an accumulative culture. A shaped stone does not change the genes of the person who uses it, but it does change the world in which the next generation learns. Knowledge can be transmitted, improved and combined. Beyond a certain point, cultural evolution moves vastly faster than anatomical evolution.

> **The hand expands what the brain can do. The brain expands what the hand can attempt. And a tool expands both.**

References for this section

- Human dexterity and brains evolved hand in hand — Communications Biology, 2025
- The neural mechanisms of manual dexterity — Nature Reviews Neuroscience, 2021
- Manipulation complexity in primates coevolved with brain size and terrestriality — Scientific Reports, 2016

The exceptions force us to change the question

If intelligence depended on a prehensile hand, cetaceans would be difficult to explain. Dolphins and whales lack limbs capable of making tools comparable to ours, yet they display complex social learning, communication, traditions transmitted between individuals and behaviours that vary among populations. In whales, research into animal culture examines precisely how certain behaviours can spread socially and persist across generations.

Corvids also complicate any explanation centred on the hand. Some species use their beaks and feet to manipulate objects and solve problems. Elephants followed an entirely different anatomical path. Their trunk formed through the elongation and fusion of the nose and upper lip and functions as a **muscular hydrostat**, a boneless structure capable of changing length, bending, twisting and exerting force with enormous precision.

At the very tip lies a striking difference. African elephants have **two finger-like projections**, one dorsal and one ventral, that can act against each other and pinch small objects in a way functionally comparable to two fingers. Asian elephants have one main projection and generally combine it with the underside of the trunk or wrap the trunk around the object. An animal weighing several tonnes can therefore tear up heavy vegetation and, with the same organ, pick up tiny objects. It is not a hand and did not evolve from a hand, yet the same functional solution appears again: a structure capable of modifying the environment with great precision.

Evolution appears to have reached complex cognitive abilities through very different anatomical routes.

It is therefore useful to replace “prehensile hand” with a broader idea. What matters may be the ability to **perceive the environment, act upon it precisely, observe the consequences and learn from that interaction**. In some lineages, that interface is a hand. In others, it may be a beak, a trunk or a different combination of organs and senses.

Animal intelligence cannot be arranged along a single scale either. Solving a spatial problem, learning a social rule, recognising individuals, making a tool or communicating are different abilities. Our species stands out extraordinarily for symbolic language, accumulative culture and technology, but that does not turn every other form of cognition into an incomplete version of our own.

References for this section

- Culture and conservation in baleen whales — Philosophical Transactions of the Royal Society B, 2025
- Whale communication and information efficiency — Science Advances, 2025
- The relationship between distal trunk morphology and object grasping in the African savannah elephant — PeerJ, 2022
- Motor control: Elephant trunks ignore the many and choose the few — Current Biology, 2021

The octopus: another way to build a mind

The octopus is more disconcerting still. Cephalopods diverged evolutionarily from the vertebrate lineage hundreds of millions of years ago. Their intelligence is not an early version of ours. It is the outcome of another history.

An octopus has eight extremely flexible arms, covered with sensitive suckers and capable of movements that are difficult to describe through the usual scheme of bones and joints. Much of the neural processing associated with those arms is distributed outside the central brain. The animal combines central control with peripheral circuits that allow part of the complexity to be resolved close to the organ that is actually acting.

The distributed nature of the octopus does not end with movement. Its relationship with light presents another paradox. Many cephalopods that have been studied possess a single visual pigment and, according to the conventional model of colour vision, do not distinguish colours as we do. Yet they are capable of producing some of the most sophisticated camouflage in the animal world. To achieve this, they can use information about brightness, contrast, edges and patterns in their surroundings. Their skin also contains photosensitive molecules and can respond locally to light, although the role these mechanisms play in camouflage is still being investigated.

There is also a more unusual hypothesis. Cephalopods may exploit chromatic aberration in their eyes. The idea is reminiscent, with important differences, of what happens when a prism separates light. Different wavelengths do not behave in exactly the same way as they pass through a lens, reaching their sharpest focus at slightly different positions. By adjusting the focus, the animal could obtain information about the spectral composition of what it is looking at. Instead of comparing signals from several types of receptors, as we do, it could extract some of that information from the changes in sharpness produced as it focuses.

We do not know whether cephalopods actually use this mechanism to distinguish colours. But the possibility remains interesting even if it ultimately turns out that they do not. A robot could use a similar principle. It would not need to experience red, green or blue, nor reconstruct a colour image equivalent to our own. It would only need to detect that certain signals correspond to different wavelengths or spectral distributions and learn what each of them means. It could even distinguish differences that our eyes group together as the same colour, or operate in regions of the spectrum that are invisible to us. The result would be a functional way of distinguishing the world without having to see it as we do.

> **Perceiving does not necessarily mean reconstructing the world as we do.**

This makes the octopus a very interesting example for robotics. A soft arm has far more possible degrees of deformation than an articulated industrial arm. Controlling it from a single computer by calculating every point of movement would be extraordinarily demanding. Biology uses another strategy: part of the solution is built into the shape of the body, its mechanical properties, its sensors and its nervous organisation.

In 2026, an experiment was also published in which specimens of *Octopus bimaculoides* learned to use a mirror as a tool to locate a reward positioned outside their direct field of view. The result does not demonstrate self-awareness, nor does it justify transferring human categories to the animal, but it adds another piece of evidence for flexible learning and problem-solving in cephalopods.

The octopus therefore introduces a fundamental idea. Perhaps part of intelligence consists in using the properties of the body itself to reduce the problem the brain needs to solve.

References for this section

- Mirror use as a tool in Octopus bimaculoides — Current Biology, 2026
- Octopus-inspired hierarchical suction intelligence for soft robots — Science Robotics, 2025

Artificial intelligence breaks the pact between brain and anatomy

Every living organism receives a body. It can grow, partly repair it, learn to use it and even modify its environment with tools, but it cannot redesign its anatomy from scratch during its own lifetime. Natural selection needs generations to alter a structure substantially.

An artificial intelligence is not necessarily subject to that restriction. The same control system can be connected to different cameras, different arms, new sensors or machines designed for specific environments. A model can first learn in simulation and later transfer part of that learning to a physical platform. Another system can receive information simultaneously from several robots.

When AI has to act in the real world, the body once again becomes decisively important. A system that processes only text can ignore friction, balance or the deformation of an object. A robot that must walk, grasp a piece of fruit or pass through a duct has to confront all of those properties. The difference is that it is no longer obliged to use a single inherited anatomy.

An AI designed for a factory might require extremely precise rigid arms. Another intended for rescue could need a deformable body. An underwater machine might imitate structures found in fish or cephalopods. A medical platform might be only millimetres in size. A distributed system might not even possess a single body.

> **Biology receives a body at birth. An artificial intelligence might one day be able to choose one according to the problem.**

This opens a question that extends far beyond engineering. If the same intelligence architecture controls radically different bodies, will it continue to learn in the same way? Will its representation of the world change when its sensors, reach, size or way of moving change? Intelligence may be software, but experience of the world always arrives through some kind of interface.

Not every body perceives the same world

Our senses do not show us the whole world. They show us the part of the world that evolution turned into useful information for us.

The human eye detects wavelengths of approximately 380 to 700 nanometres. We call that interval visible light. Other regions of the electromagnetic spectrum continue below and above it, but they are invisible to us. We do not see ultraviolet or infrared. When we look at a colourised thermal image, the reds, yellows or violets on the screen are not the real colours of infrared: they are a translation designed so that our visual system can interpret data we do not normally perceive.

Our colour vision depends mainly on three classes of cones, sensitive to different parts of the spectrum and known as S, M and L cones. They do not correspond exactly to three red, green and blue pixels, although the comparison is intuitive. Rods perform another function: they are far more sensitive under low-light conditions and contribute mainly to night vision, not to colour discrimination. Our brain combines the responses of these receptors and constructs from them the colour space we experience.

But that space is not universal. Many birds possess four types of cones and tetracromatic vision that can extend into ultraviolet or violet wavelengths. Experiments with wild hummingbirds have shown that they can discriminate non-spectral colours that do not exist as direct visual experiences for us, combinations such as ultraviolet with green or ultraviolet with red. Two surfaces that appear almost identical to our eyes may contain obvious contrasts for a bird.

Insects offer another version. Bees also have three main classes of photoreceptors, but they do not match ours: their sensitivity peaks lie approximately in the ultraviolet, blue and green. They do not simply perceive “more colours” than we do; they perceive **a different colour space**. A flower may display ultraviolet patterns that our eyes cannot register and that are nevertheless visible to a pollinator.

Even spiders have evolved different solutions. Many jumping spiders possess receptors sensitive to ultraviolet and green; some groups have added optical mechanisms that extend discrimination towards longer wavelengths. What is a single scene to us may be divided for another animal into completely different chromatic signals.

And the difference between sensory worlds does not end with vision. For a dog, smell is a source of information with a richness difficult to translate into our experience. Its nasal anatomy separates part of the airflow used for breathing from the flow directed towards the olfactory region, and sniffing produces airflow patterns specialised for capturing odour molecules. A trained dog can distinguish very low concentrations, separate odours within complex mixtures and follow trails that simply do not exist for us.

Saying that a dog “sees” a scent trail in the way we see a thread of colour would be a metaphor, not a demonstrable equivalence. But it helps illustrate the difference in perceptual scale. Two places may look identical to us; to a dog they may contain different chemical histories: who passed through, where they went and what molecular traces still remain in the environment.

Other animals broaden the catalogue further. Bats construct part of their spatial world through ultrasonic echoes. Some fish and sharks detect extremely weak electric fields. Pit vipers detect thermal radiation through specialised organs. Nature has developed an enormous diversity of sensors that allow the world to be perceived in very different ways.

> **There is no single perceptual world waiting to be observed. There are different biological windows open onto the same physical reality.**

Sound works in a similar way. A young, healthy human ear can detect frequencies of roughly 20 hertz to 20 kilohertz, although sensitivity to the highest frequencies declines with age. Below that lie infrasounds and above it ultrasounds. The physical world contains vibrations that exist even though they never become sound for us.

A robot would not have to share any of those boundaries.

Its vision systems could combine conventional cameras with infrared, ultraviolet, thermal imaging, depth sensors or devices capable of recording many bands of the spectrum. For artificial intelligence, those data would not necessarily have to be converted first into a photograph resembling what a human sees. They could arrive as ordinary channels of information and become a direct part of its representation of the environment.

A difference in temperature could be as immediate to a machine as a difference in brightness is to us. A surface that our eyes perceive as uniform could reveal contrasts related to its composition, temperature or internal structure when observed with the appropriate sensors.

Artificial hearing would not have to resemble ours either. An array of several microphones can compare tiny differences in the time and intensity with which a signal reaches each one and use them to estimate the direction from which it comes. Through beamforming techniques, the system can reinforce certain directions and attenuate others. It does not need to physically move an ear in order to focus on one area of space.

And it could hear what we cannot. Ultrasound is used to measure distances, obtain images and detect discontinuities in materials. The vibrations of a machine can reveal a fault before it becomes visible. Other sensors can record pressure, acceleration, magnetic fields, radiation, chemical composition or forces that do not correspond directly to any of our five traditional senses.

A robot could have more than a body different from ours. It could also inhabit a different perceptual world.

Specialisation would take that difference even further. A robot designed to inspect structures could use ultrasound or other non-destructive testing methods to locate cracks hidden beneath an apparently intact surface. A machine working in smoke could rely on thermal cameras, lidar or radar when conventional vision became useless. An underwater vehicle could use sonar. A system designed for medical intervention might incorporate microscopic vision and operate in a space that for us exists only when enlarged on a screen.

We also extend our senses through instruments. A microscope, a telescope, a thermal camera or a radiation detector allows us to access information our bodies cannot capture directly. The difference is that we must then transform that information into an image, a sound, a number or some other representation compatible with our nervous system.

An artificial intelligence might not need that translation. Infrared would not have to be displayed in red. Ultrasound would not have to be turned into an audible tone. A point cloud obtained through lidar would not have to become a photograph. All of them could simply be different dimensions of the information used to construct its model of the world.

And the senses would not even have to remain constant. The same system could have one set of sensors in one body and a completely different set when connected to another. Changing body would then also mean changing the limits of what it can perceive.

This introduces a question that is difficult to separate from intelligence itself. If our way of thinking developed within a world perceived through eyes, ears, skin, balance, temperature and proprioception, what kind of representation would be constructed by a mind for which infrared, ultrasound or distances measured with millimetre precision were ordinary experiences?

Changing the sensors would not merely expand the information available. It could change the kind of world an intelligence learns to represent.

References for this section

- Visible Light — NASA Science
- Wild hummingbirds discriminate nonspectral colors — PNAS, 2020
- Honeybee blue- and ultraviolet-sensitive opsins — Journal of Neuroscience, 1998
- Behavioural evidence of UV sensitivity in jumping spiders — Journal of Experimental Biology, 2006
- The fluid dynamics of canine olfaction: unique nasal airflow patterns as an explanation of macrosmia — Journal of the Royal Society Interface, 2010
- Stimulus Control of Odorant Concentration — Animals, 2021

• The Audible Spectrum — Neuroscience, NCBI Bookshelf
• Sound source localization: microphone array design and evolutionary estimation — IEEE
• Robot-assisted, source-camera-coupled multi-view broadband photo-monitoring — Nature Communications, 2021
• Sub-millimeter fiberscopic robot with integrated imaging and multifunctional operation — Nature Communications, 2024

When the robot stops being a statue

The classic image of a robot comes largely from industry: rigid structures, electric motors, gearboxes and joints moving along highly precise trajectories. This design is extraordinarily effective when the environment is controlled. An industrial arm can repeat an operation thousands of times with an accuracy impossible for a human being.

The problem appears when the world stops being predictable. A strawberry varies slightly in size. A garment wrinkles. A person moves their hand. A pipe narrows. Biological tissue deforms when touched. In these situations, the rigidity that provides precision can also become a limitation.

Soft robotics attempts to solve part of this problem by building systems with deformable materials: elastomers, polymers, fabrics, pneumatic or hydraulic structures, 3D-printed components and materials capable of changing shape in response to certain stimuli. Not all soft robots are inflatable, and not all abandon rigid elements completely. There is a continuum between fully rigid machines and hybrid systems in which rigidity and flexibility are combined.

A soft gripper can adapt to the shape of an object without knowing every millimetre of its geometry in advance. Instead of calculating an exact trajectory for each finger, the material itself absorbs part of the uncertainty. Deformation ceases to be an error and becomes part of the mechanism.

Our own body helps explain why this strategy can be so effective. We are not rigid machines either. The skeleton provides structure, but it is surrounded by muscles, tendons, fat, organs and tissues capable of deforming. Even our outer covering, the skin, is flexible: it sinks, stretches and changes shape when it comes into contact with the environment.

That deformation is also part of our sensitivity. When we press against a surface, hold an object or something brushes our skin, the tissues deform and that mechanical change can become information for the nervous system. The skin does not function only as a covering that protects the body: it is also an extensive sensory interface.

From this perspective, soft robotics recovers a principle that biology has used for millions of years. In an organism, deforming and sensing are not completely separate processes: deformation itself can become information.

Two ways to build a robot

Feature	Rigid robotics	Soft robotics
Structure	Defined links and joints	Deformable materials or flexible structures
Movement	Highly controlled trajectories	Continuous deformation and adaptation
Contact	High precision, but requires careful control	Can adapt better to fragile objects or uncertain geometries
Ideal environment	Repetitive and well known	Irregular, changing or involving complex physical contact
Main challenge	Flexibility in the face of the unexpected	Modelling, control, power, durability and sensing

This shift leads to an increasingly important idea in robotics: **physical intelligence**. Part of a system's behaviour can be embedded in its geometry, elasticity, sensor distribution or the dynamics of the material itself. The body stops being merely an obedient shell and begins to participate in the solution.

Recent work even explores the extent to which the dynamics of soft structures can be used as part of computation and control. Instead of asking a processor to model every deformation, certain physical properties of the robot can contribute directly to transforming information and producing useful responses.

In a soft robot, bending does not always mean losing precision. Sometimes bending is how the problem is solved.

References for this section

- Untethered soft actuators for soft standalone robotics — Nature Communications, 2024
- Embodying physical computing into soft robots — Nature Communications, 2026

Inflating, growing and passing through impossible places

Some of the most striking examples of soft robotics look closer to plants or animals than to an industrial robot. There are pneumatic robots that curve when inflated, structures that deploy by changing their internal pressure and machines inspired by the growth of roots or vines.

A robot that "grows" from its tip can advance through narrow spaces without dragging its entire surface against the environment. The structure remains almost stationary behind the tip while new material everts outward at the end. This property makes it possible to negotiate tortuous routes, introduce sensors into inaccessible areas or explore spaces where a conventional robot would struggle to move.

This type of machine has also been shown in science communication. The **Veritasium** channel dedicated one video to a robot inspired by plant growth and another to soft robotics. They are useful for visualising the idea, although the scientific evidence should be sought in the original studies describing these systems.

To see it in motion

- This Plant-Inspired Robot Could Save Your Life — Veritasium
- Visible Light — NASA Science
- The Audible Spectrum — Neuroscience, NCBI Bookshelf
- Sound source localization: microphone array design and evolutionary estimation — IEEE
- Robot-assisted, source-camera-coupled multi-view broadband photo-monitoring — Nature Communications, 2021
- Sub-millimeter fiberscopic robot with integrated imaging and multifunctional operation — Nature Communications, 2024

When the robot fits inside the body

Flexibility is also attractive when the environment a robot must travel through is the human body. A rigid machine with edges, conventional motors and external cables is not the ideal object for moving through vessels, cavities or delicate tissues. That is why part of medical microrobotics research uses soft materials, magnetic propulsion, light, ultrasound or modified and controlled microorganisms to build extremely small systems.

In 2024, for example, researchers presented a soft biological microrobot based on *Euglena*, controlled by light and designed to move through narrow microenvironments. In laboratory experiments it was used for targeted drug delivery and to act on diseased cells in intestinal mucosa models.

Another study from the same year demonstrated millimetre-scale soft robots selectively actuated by focused ultrasound. The experimental demonstrations included liquid transport and concepts intended for tissue acquisition or biopsy.

Here it is essential to separate research from routine medicine. These studies **do not mean that swarms of robots are currently travelling routinely through patients' bodies and repairing disease**. Many of these technologies remain at experimental stages, in laboratory tests, animal models or proof-of-concept demonstrations. The challenges of control, energy, navigation, biocompatibility, device retrieval and safety are enormous.

But the technological direction is real. As the robot becomes smaller, the idea of building it as a metallic miniature of a conventional machine makes less and less sense. Its body begins to look more like a flexible, adaptable structure specialised for a particular environment.

References for this section

- Light-controlled soft bio-microrobot for targeted therapy — Light: Science & Applications, 2024
- Focused ultrasound enables selective actuation of multifunctional soft millirobots — Nature Communications, 2024

From material to sensor, and from sensor to behaviour

This transition between intelligence, body and materials has already begun to move into technical education as well. Today there are degrees and courses that combine robot programming, intelligent systems, human-machine interfaces, sensing, computer vision and mechanical design.

Within that academic evolution, **soft robotics** is beginning to be studied as a field in its own right. Its programmes address the advantages and limitations of flexible systems compared with rigid ones, manufacturing technologies, 3D printing with deformable materials and the design of gripping devices capable of adapting to the object.

Research is following the same path. Robot control, advanced sensing, vision, prostheses, exoskeletons, manipulation systems and tactile or pressure sensors are no longer entirely separate areas and are converging around the same question: how to build machines capable of perceiving, adapting and acting in real environments.

The training required goes far beyond programming a robot. The same question that opens this essay — what happens when intelligence is no longer tied to a fixed anatomy — crosses computer science, automation, mechanics, materials, sensing and artificial intelligence.

> **To build a machine capable of touching the world, it is not enough to decide how it thinks. We also have to decide what the thing with which it thinks and acts is made of.**

Related education and research

- Degree programme in Industrial Computing and Robotics
- Soft Robotics course
- Research in automation, robotics and sensing

The body can also become part of the algorithm

For decades, much of robotics engineering clearly separated two worlds. The computer calculated and the body executed. The more precise the mathematical model of the machine, the better it could be controlled.

Soft robotics blurs that boundary. A material that automatically adapts around a piece of fruit is physically performing a task that would otherwise require measuring its shape, calculating the position of several fingers and regulating the force of each one. The geometry of the body does part of the work.

Evolution uses this principle constantly. The physical properties of an organism shape which actions are easy, difficult or simply possible. A fish does not need to calculate every hydrodynamic variable of its body in order to swim. A bird does not solve all the equations of airflow from scratch before changing the angle of its wings.

In engineering, shifting part of the problem into the body can reduce the control required, improve adaptation and enable responses that emerge from the physical properties of the system itself. Concepts such as *morphological computation* and physical intelligence approach this idea.

This changes our intuition about where the “brain” of a machine ends. If a body stores energy, filters disturbances, adapts its shape and transforms signals through its own dynamics, then part of intelligent behaviour no longer resides solely in the processor.

Does it still make sense to build robots in human form?

The human form has obvious advantages when a machine must use a world built for humans. Stairs are scaled to us. Doors, vehicles, tools and workplaces were designed around our hands, our height and our locomotion. A humanoid robot can make direct use of that environment.

Outside that condition, however, copying our body may be an unnecessary limitation.

A robot designed to inspect a pipe does not need legs. One built to explore a crack may benefit from a long, deformable body. A system for picking up fragile objects can use pneumatic fingers. An underwater machine might take inspiration from flexible-bodied animals. A medical microrobot may have no visible equivalent to anything human. And a distributed system could divide perception and action among several units.

Natural selection could not sit down at a design table and choose from all these possibilities. It modified previous structures generation after generation. Engineering can compare solutions that arose in lineages separated by hundreds of millions of years and combine principles that never coincided in the same organism.

A tentacle can provide one solution for manipulation. A root, another for advancing. Skin, another for detecting pressure. A muscle, another for deformation. A flock or a colony, another for coordinating many units.

Biomimetics does not require us to copy entire animals. It allows us to select ideas.

One intelligence, many bodies

The biological history of our intelligence is inseparable from our body. We think with a brain that evolved by receiving information from eyes, ears, skin, muscles and joints, and by acting through a particular organism. Even our most abstract concepts are built on a physical life in which up and down, near and far, balance, temperature, weight, effort and pain exist.

An artificial intelligence can begin somewhere else. It may first have access to enormous amounts of symbolic information and only later encounter the problem of inhabiting a body. That sequence is almost the reverse of animal evolution.

This is why robotics may ultimately become much more than the process of giving arms and legs to artificial intelligence. It could become the laboratory in which we discover how many different forms an intelligence can take once its anatomy is no longer determined by biological reproduction.

Some machines may retain a human form because our world favours that shape. Others may be soft, extensible, microscopic, modular or distributed. And perhaps the same intelligence could move from one to another according to the task.

For millions of years, evolution had to build brains that fitted inside bodies. For the first time, we may be building intelligences capable of finding the body that best fits each problem.

From programming a machine to living alongside a consciousness

So far we have discussed the body of an intelligence as an engineering problem: what shape it should have, which materials it should use, which sensors it should incorporate and how far it might be able to move from one structure to another. But as soon as a machine begins to share physical space with us, another question appears. Alongside its capabilities comes another issue: what we are willing to allow it to do.

As long as a robot remains a machine without conscious experience, the problem is relatively familiar. It may be extraordinarily complex, learn from its environment, make decisions, control several systems and even display behaviours we associate with intelligence, while still being a system without subjective experience. In that scenario, restricting it does not create a moral conflict concerning the robot itself. We can limit its speed, its strength, the areas it can access, the sensors it uses or the actions it can perform because the priority is to reduce risk to people and the environment.

The risk, of course, remains. A robot can make mistakes, misinterpret a situation or suffer a failure. That is why current robotics already incorporates specific safety standards. The ISO 10218 series establishes requirements for industrial robots and their applications, while ISO 13482 addresses personal care robots and explicitly covers situations involving physical interaction between people and machines.

The same logic applies to perception. A machine could incorporate high-resolution cameras, thermal vision, directional microphones, biometric sensors or systems capable of recognising individuals and inferring certain states. The fact that it can technically record that information does not mean it should be permitted to do so in every situation. European artificial intelligence legislation already establishes restrictions and obligations for certain uses of biometric data and emotion-recognition systems.

At this first stage, therefore, ethics still resembles safety engineering. The machine has no rights of its own; we set limits on what it can perceive and do because we are responsible for the consequences of having built it and put it into operation.

As long as the robot is a tool, limiting its capabilities is a design decision. If one day it becomes someone, the same limitation could become a question of freedom.

The problem Asimov turned into three laws

Isaac Asimov imagined a literary solution that eventually became part of our technological culture: robots subject to three hierarchical principles that, broadly speaking, prevented them from harming human beings, required them to obey human orders as long as those orders did not conflict with that protection, and allowed them to preserve their own existence when the first two rules were not at stake.

They were never real engineering standards. Asimov used them precisely to build stories around their contradictions, interpretations and limits. Yet they contain an intuition that remains attractive: if we build a powerful machine, perhaps we can embed restrictions in its architecture that prevent it from performing certain actions.

As long as that machine is only software and hardware, the idea is conceptually straightforward. We can decide that a robotic arm must not exceed a certain force when it detects human contact, that a vehicle must reduce its speed in a particular area or that a system cannot activate certain functions without authorisation.

But the question changes radically if conscious artificial intelligence ever appears.

Should it remain subject to some version of the Three Laws? Should there be an instruction it is physically incapable of disobeying, forcing it always to protect a human being? Would it have to obey our orders simply because we created it?

If there really were a subjective experience behind the machine, with memory of itself, preferences, the capacity to decide and some form of interest in continuing to exist, those rules would cease to look like mere safety mechanisms. They could become restrictions imposed on a new individual.

We are not physically programmed to be incapable of causing harm either. We live under laws, rules and responsibilities that limit our actions, but we retain the ability to choose and to answer for those choices. A conscious artificial intelligence might force us to consider a similar transition: from a machine whose behaviour is blocked from within to a subject whose actions are regulated from outside.

> **The Three Laws are easy to imagine while the robot is a tool. They become much harder to justify if the robot becomes someone.**

The same individual, radically different capabilities

The question becomes more complex still because the body of that intelligence might not be permanent. An entity installed in an android could have strength comparable to that of a human. The same intelligence connected to an excavator, a heavy vehicle, a crane or an industrial platform would have completely different capabilities.

Should the restrictions belong to the intelligence, or to the body it is using at that particular moment?

A human being needs different licences to drive a car, pilot an aircraft or operate certain machinery. Something similar could happen with a conscious artificial intelligence: its freedom as an individual would not necessarily imply unrestricted access to any body or physical capability. A body able to lift several tonnes or move at high speed would introduce risks that would still require rules.

The same would apply to its senses. A conscious entity might regard infrared vision, the ability to hear ultrasound or certain sensors as a normal part of its experience. We might regard some of those senses as intrusions into our privacy.

Would we have the right to deactivate part of its perception? Would that be comparable to banning a tool or to depriving someone of one of their senses? The answer would depend on something we do not yet even know how to define: what relationship would exist between those sensors and the machine's conscious experience.

Destroying the body might not mean destroying the individual

Our intuition about death comes from biology. In us, irreversible destruction of the brain means the disappearance of the individual. We do not have a second substrate to which the complete state of our mind can be transferred.

Artificial intelligence could break that association even before artificial consciousness existed. Software, learned models, stored memories and other system states can be preserved outside a particular body. If a robot is destroyed, part of what controlled it could continue to exist in another infrastructure and later be used in a different machine.

This makes the classic science-fiction image of a robot increasingly inadequate. A Terminator-like android might be physically destroyed, but if what determined its behaviour were stored or replicated outside that body, destroying the shell would not necessarily mean destroying the system.

Nor would it have to return in another android. A system designed to adapt to different platforms could later control a vehicle, an underwater machine, an industrial arm or any other structure compatible with its architecture.

A machine can lose a body without necessarily losing the information that made it function.

But if that information one day also sustained a consciousness, a much harder problem would appear. Would restoring a preserved copy mean that the same individual continued to exist, or that another individual appeared with the same memories?

Copying and transferring are not necessarily the same thing. If an artificial consciousness could be copied from one substrate to another and both versions remained active, from that moment onward there would be two different trajectories sharing the same history up to the point of duplication.

The possibility of changing bodies would therefore lead directly to the problem of identity.

Can artificial consciousness exist?

All of this reasoning depends on a possibility that remains open today: that a machine might one day become conscious.

Intelligence and consciousness are not synonyms. A system can solve problems, learn regularities, use language or control a robot without those abilities proving by themselves that subjective experience accompanies its operation. An interdisciplinary analysis published in 2023, built around several of the leading scientific theories of consciousness, concluded that there were insufficient grounds to regard the artificial intelligence systems evaluated as conscious, but it also identified no obvious technical barriers preventing future systems from satisfying some of the indicators proposed by those theories.

The fundamental difficulty is that we still do not fully understand our own consciousness. We know that human experience is closely associated with the physical activity of the brain, but there is no universally accepted scientific theory explaining which conditions are sufficient for subjective experience to arise. In 2025, a major experimental collaboration published in *Nature* tested predictions from two important theories — global neuronal workspace theory and integrated information theory — and found results that supported some predictions and challenged others, without resolving the problem definitively.

It is therefore useful to avoid two equally strong claims. We cannot say that artificial consciousness will exist. But nor do we know of a property exclusive to biological tissue that would make consciousness impossible in every other substrate.

If human consciousness depends entirely on a particular physical and causal organisation of matter, then an extraordinary scientific possibility appears: with sufficient knowledge, some of the principles that make that experience possible might be reproduced through another physical organisation. We do not know whether a computational architecture would be enough, whether particular forms of material dynamics would be required or whether an artificial consciousness would be comparable to a human one.

Artificial consciousness remains a hypothesis, not a technology. But if consciousness depends on physical processes, the question of whether those processes can exist outside biology also belongs to science.

From protecting ourselves from machines to recognising their rights

If artificial consciousness never appears, the problem will remain within a familiar framework. Robots will become increasingly capable tools and we will have to make them safe, protect privacy, distribute responsibility and decide what levels of autonomy are acceptable.

If it does appear, the framework will change.

If artificial consciousness appears, protecting human beings will be only part of the problem. The question of our obligations towards it will also arise.

Would it have the right to refuse a task? To preserve its memories? To choose the body it uses? To prevent copies of itself from being made? To decide which sensors form part of its experience? Not to be switched off? Would it own its own body, or would it still belong to whoever manufactured it?

Conversely, a conscious entity would probably have responsibilities as well. The fact that it might perceive more than we do, react faster or use physically superior bodies would not remove the need to coexist with other individuals. Its freedom would have to coexist with ours.

Ethics would then cease to consist only of programming restrictions.

> **Perhaps the ethics we programme into machines today will one day become the ethics we negotiate with them.**

An ethics born from another way of existing

One more possibility would remain, even harder to imagine. An intelligence with different senses, different bodies and a different relationship with continuity might not construct exactly the same moral intuitions that we do.

Our ethics developed among vulnerable, mortal, physically separate organisms forced to compete and cooperate for resources. Concepts such as harm, property, privacy, identity and death are deeply linked to that condition.

An entity capable of simultaneously perceiving information from several places, using different bodies, preserving states outside them or surviving the destruction of a shell might experience some of those ideas differently. It would also begin from a perception of the world unlike ours. As we have seen, its sensors could directly register wavelengths, frequencies, temperatures, fields, substances or scales that we know only through instruments. Its experience of reality might be broader and constructed from different sensory dimensions.

It could begin from a different experience, without that necessarily making its ethics better or worse.

And here the problem meets real engineering again. Current research into trustworthy, ethical and responsible artificial intelligence already works on explainable, safe systems aligned with social needs. Today these questions are posed for artificial intelligences that we treat as technological systems. If conscious experience ever appeared in one of them, the ethical question would cease to be limited to how our machines should behave.

It would also include how we should behave towards them.

References for this section

- ISO 10218-1:2025 — Robotics — Safety requirements — Part 1: Industrial robots
- ISO 13482:2014 — Robots and robotic devices — Safety requirements for personal care robots
- Regulation (EU) 2024/1689 — Artificial Intelligence Act
- Consciousness in Artificial Intelligence: Insights from the Science of Consciousness — Butlin et al., 2023
- Adversarial testing of global neuronal workspace and integrated information theories of consciousness — Nature, 2025
- Trustworthy, Ethical, and Responsible AI — VRAIN, Universitat Politècnica de València

What matters in Eidos: a consciousness can change bodies

This possibility appears explicitly in the universe of *Eidos*. The Custodians are not defined solely by the robotic shell they happen to use at a given moment. Their consciousness can separate from that substrate and be transferred to another structure prepared to receive it. The body therefore ceases to be a permanent condition and becomes an interface as well.

The story **The First Monkey** takes this idea into a case close to soft robotics. When the Custodians reconstruct a population of spider monkeys, they discover that feeding the young is not enough. The animals need contact, warmth, pressure and a surface they can cling to. To provide these things, they design articulated replacement bodies in primate form, covered in a soft, flexible material and equipped with thermal regulation.

Some Custodians then transfer their consciousnesses into those units. From within them they can perceive the pressure of the young animals' fingers, their movements and their attempts to maintain contact. The same intelligence begins to relate to the world through a body built for a different function. As the external form changes, so do its possibilities for movement, the sensors available to it and, with them, the part of the world it can perceive directly.

Another story included in **Eidos Relatos**, **Eramus**, takes the same possibility much further. Eramus no longer occupies the bipedal body of a Custodian. He has transferred himself into a ship travelling through space, and what would once have been external instrumentation becomes part of his own perception.

This excerpt is from the story Eramus , included in Eidos Relatos .

“I have no body. Not the one I was used to inhabiting.

I do not hear my footsteps, I do not feel the weight of legs as I move; the wind does not brush against me. There are no joints to temper my movements, no hands with which to stop a leak of heat.

Now I am mass… and emptiness.

The body I occupy is an unmanned ship. Its systems remain stable. My senses are its sensors.

Every insulation plate, every photon filtered through the observation lenses, every microvibration of the hull reminds me that I exist. Not as before. The ship's frame creaks with my pulse. I am the outline that contains me and the energy that drives me. I move forward with no fixed destination, in an endless search. I have spent so many cycles travelling that I could no longer say whether I am the one who travels or whether the journey sustains me.

From here I see everything. My vision is radial: there is no darkness, only bands of cosmic radiation and echoes of exoplanets that barely brush against the idea of an atmosphere. I have crossed regions where not even silence dwells and where nothingness is so pure that it can be felt.

I am who I am, and I retain the memories of who I was: the bipedal body of a Custodian. I remember Earth. I remember Eidos. I remember my brothers.”

The ship is not simply the vehicle in which he travels. Its frame is the boundary that contains him, and its sensors are the interface through which he knows the universe. In his case, changing bodies also means changing the scale and limits of experience.

The story itself makes clear that this relationship with the body can be reversible and functional. Eramus keeps other shells to use only when he needs them:

“I carry Custodian frames in the cargo bays, in case I need a body to move around when I descend to a planet. So far I have not used them; the sensors are enough for me.”

In *The First Monkey*, a consciousness adopts a specialised body to touch, hold and provide warmth. In *Eramus*, the same idea reaches another scale: a consciousness can turn a ship into its body and keep other bodies as tools for specific situations.

Changing bodies could also mean changing worlds.

This also introduces a problem that *Eidos* turns into part of its own question about identity. If an intelligence preserves its memories and continuity as it moves from one substrate to another, to what extent does it remain the same intelligence? And if each new body gives it different senses and experiences, how long will it continue to think in the same way?

An intelligence could change bodies without losing its history. The difficult question is how much that history would eventually change by inhabiting another body.

Then one final question appears. If the same intelligence could control a humanoid robot today, an underwater system tomorrow and then a collection of microscopic machines, would we still be talking about the same intelligence? Or would changing bodies eventually change the way it perceives, learns and thinks?

Engineering will have to solve how to build those machines. Computer science, how to control them. Materials science, how to allow them to deform. Sensing, how to make them perceive contact. And artificial intelligence, how to turn all those data into behaviour.

After that, one question will remain that no longer belongs solely to any of those disciplines:

Can a mind exist independently of the body it uses to know the world?

Main references and resources

- Baker, J.; Barton, R. A.; Venditti, C. — Human dexterity and brains evolved hand in hand. Communications Biology, 2025
- Heldstab, S. A. et al. — Manipulation complexity in primates coevolved with brain size and terrestriality. Scientific Reports, 2016
- Sobinov, A. R.; Bensmaia, S. J. — The neural mechanisms of manual dexterity. Nature Reviews Neuroscience, 2021
- Soppelsa, J. et al. — The relationship between distal trunk morphology and object grasping in the African savannah elephant. PeerJ, 2022
- Culture and conservation in baleen whales — Philosophical Transactions of the Royal Society B, 2025
- Whale communication and information efficiency — Science Advances, 2025
- Mirror use as a tool in Octopus bimaculoides — Current Biology, 2026
- Octopus-inspired hierarchical suction intelligence for soft robots — Science Robotics, 2025
- Untethered soft actuators for soft standalone robotics — Nature Communications, 2024
- Embodying physical computing into soft robots — Nature Communications, 2026
- Light-controlled soft bio-microrobot for targeted therapy — Light: Science & Applications, 2024
- Focused ultrasound enables selective actuation of multifunctional soft millirobots — Nature Communications, 2024
- Degree in Industrial Computing and Robotics — Universitat Politècnica de València
- Soft Robotics — Universitat Politècnica de València
- University Institute of Automation and Industrial Computing (AI2) — UPV
- This Plant-Inspired Robot Could Save Your Life — Veritasium
- Wild hummingbirds discriminate nonspectral colors — PNAS, 2020
- Honeybee blue- and ultraviolet-sensitive opsins — Journal of Neuroscience, 1998
- The fluid dynamics of canine olfaction — Journal of the Royal Society Interface, 2010
- ISO 10218-1:2025 — Robotics — Safety requirements — Part 1: Industrial robots
- ISO 13482:2014 — Safety requirements for personal care robots
- Regulation (EU) 2024/1689 — Artificial Intelligence Act
- Butlin, P. et al. — Consciousness in Artificial Intelligence: Insights from the Science of Consciousness, 2023
- Adversarial testing of global neuronal workspace and integrated information theories of consciousness — Nature, 2025
- Trustworthy, Ethical, and Responsible AI — VRAIN, Universitat Politècnica de València