The Embodied Octopus: Distributed Intelligence and Active Inference in a Flexible Body

Dr Alexander D. Shaw & Lioba C.S. Berndt
Computational Psychiatry and Neuropharmacological Systems Lab
Department of Psychology, University of Exeter, UK
a.d.shaw@exeter.ac.uk

December 17, 2025

Abstract

Octopuses exhibit a striking form of embodied intelligence: a nervous system in which over two-thirds of neurons reside not in a central brain, but in their highly flexible arms [Hochner, 2012, 2023]. This distributed architecture enables each arm to perform local sensing, prediction, and motor control, while the central brain coordinates global behaviours such as hunting, navigation, and camouflage [Hanlon and Messenger, 1996]. We argue that this organisation is best understood through the lens of hierarchical Active Inference, in which local generative models operate at the periphery of the body while higher-level priors govern goal-directed policy selection [Friston, 2010, Buckley et al., 2017]. Unlike vertebrate organisms, the octopus must control thousands of degrees of freedom in a soft-bodied morphology, making purely centralised control computationally intractable [Flash and Zullo, 2023, Kier and Smith, 1985]. Instead, its nervous system exhibits a multi-scale division of labour between local and global inference, offering a powerful biological blueprint for distributed, sensorimotor AI [Hale, 2025, Carls-Diamante, 2022]. In this perspective, we explore how octopus cognition provides a living model of decentralised, hierarchical predictive processing and discuss implications for artificial systems that seek to embody similar principles [Costa et al., 2022, Mazzaglia et al., 2022].

1 Introduction

Octopuses occupy a unique position in the study of biological intelligence: despite being invertebrates with relatively short lifespans, they display remarkable behavioural flexibility, including problem solving, rapid learning, adaptive camouflage, and exploratory play [Hanlon and Messenger, 1996, Ponte et al., 2022]. What makes these abilities especially intriguing is their neuroanatomy. Unlike vertebrates, where cognition is largely centralised, octopuses distribute most of their neurons throughout their eight arms [Hochner, 2012, Carls-Diamante, 2022]. Each limb exhibits a degree of local autonomy, capable of sensing, acting, and adapting with minimal reference to a central command structure [Hochner, 2023, Olson and Ragsdale, 2023]. This challenges classical models of cognition that assume a singular locus of control and instead highlights the role of the body itself as a site of computation.

Embodied intelligence foregrounds the idea that cognition is not solely the product of a central brain but emerges from the dynamic interplay between neural structures, bodily morphology, and the environment [Pfeifer and Bongard, 2007, Carls-Diamante, 2022]. In soft-bodied organisms like the octopus, this interplay becomes extreme: the animal must coordinate thousands of degrees of freedom, integrate multimodal sensory feedback, and execute precise motor actions using a body that resists simple kinematic description [Kier and Smith, 1985, Flash and Zullo, 2023]. Classical control-theoretic frameworks based on centralised planning or rigid body dynamics struggle to account for such fluid, decentralised coordination.

Active Inference offers a principled framework for understanding these phenomena. By modelling perception and action as two sides of the same process (minimising variational free energy under a generative model) it re-casts control not as purely top-down instruction, but as hierarchical prediction [Friston, 2010, Buckley et al., 2017]. This perspective naturally accommodates systems in which local, peripheral controllers make fast, context-specific inferences while higher levels encode longer-term goals and task-relevant priors [Friston et al., 2018, Mazzaglia et al., 2022]. In this sense, the octopus provides a natural model for thinking about distributed predictive processing: a biological system in which intelligence is literally spread through the body [Hochner, 2012, Hale, 2025].

In this paper, we propose that the octopus exemplifies a form of hierarchical, decentralised Active Inference that allows it to control a massively under-actuated and soft morphology. By examining its nervous system through this computational lens, we aim to clarify how octopus behaviour emerges from the coordination of local and global generative models and how these principles may inspire new approaches in artificial intelligence and soft robotics [Costa et al., 2022, Nakajima et al., 2013, Cianchetti et al., 2015].

2 The Octopus Nervous System: A Distributed Architecture

2.1 Central and Peripheral Control

The octopus nervous system is striking not only for its size relative to body mass, but for its radical decentralisation. Fewer than half of the animals neurons reside in the supra-esophageal brain; the majority are distributed throughout the arms, embedded within axial nerve cords, peripheral ganglia, and dense sensorimotor networks [Hochner, 2012, Kuuspalu et al., 2022, Hale, 2025]. This arrangement stands in sharp contrast to vertebrate architectures, where peripheral nerves primarily relay information to and from a centralised control structure. In the octopus, the arms themselves constitute semi-autonomous computational units [Carls-Diamante, 2022].

Each arm is capable of generating and modulating its own motor patterns, integrating local proprioceptive and tactile information, and initiating behaviours that do not require explicit commands from the central brain. Experiments have shown that isolated arms can execute coordinated reaching movements, maintain rhythmic patterns, explore novel objects, and perform grooming-like actions [Hochner, 2012, Levy et al., 2015, Olson and Ragsdale, 2023]. Such findings indicate that substantial elements of decision-making and control are embedded locally within the limbs rather than exclusively governed by a central controller. The central brain instead appears to provide high-level coordination, such as setting global goals, integrating multisensory context, and resolving conflicts between arms, while leaving moment-to-moment sensorimotor inference to distributed peripheral circuits [Hochner, 2023, Hale, 2025].

2.2 Control at the Edge of the Body

Control in the octopus is pushed unusually far toward the periphery of the body. Each arm contains extensive sensory machinery: tactile receptors, stretch receptors, and chemotactile cells allowing it to gather rich local information without relying on central processing [Hochner, 2012, Ponte et al., 2022]. This sensory input is tightly coupled to local motor circuitry, enabling each limb to generate context-appropriate actions such as probing, grasping, or adjusting posture in response to immediate environmental contingencies. Rather than functioning as passive effectors, the arms continuously perform local inference about the causes of their sensations [Carls-Diamante, 2022, Olson and Ragsdale, 2023].

This arrangement invites comparison with vertebrate reflex arcs, yet the octopus goes well beyond simple stimulus-response loops. Whereas traditional reflex pathways are fixed, stereotyped mechanisms, octopus arm control appears to involve flexible generative processes: peripheral circuits form predictions about expected sensory consequences of movement and update motor output accordingly [Hochner, 2012, 2023]. In the language of Active Inference, the arms maintain their own low-level generative models that guide local action-perception cycles [Friston, 2010, Friston et al., 2018].

Perhaps the clearest demonstration of this autonomy comes from classic severed-arm experiments. Even when disconnected from the central brain, isolated arms retain the ability to execute coordinated reaching, grasping, and withdrawal movements in response to tactile stimulation [Hochner, 2012, Levy et al., 2015]. These behaviours reflect an embodied inference process occurring entirely at the limb level. The central brain, rather than micromanaging each degree of freedom, interacts with these peripheral controllers by providing broad priors about global goals, leaving fine grained sensorimotor control to the distributed architecture at the edge of the body [Carls-Diamante, 2022, Hale, 2025].

Understanding how such semi-autonomous limbs coordinate with central goals requires a framework in which perception and action are tightly integrated at multiple scales; a role naturally served by Active Inference [Friston, 2010, Buckley et al., 2017].

3 Active Inference and Hierarchical Generative Models

3.1 Active Inference in a Nutshell

Active Inference provides a unifying mathematical framework in which perception, action, and learning arise from the same underlying imperative: the minimisation of *variational free energy* [Friston, 2010, Buckley et al., 2017]. The starting point is a generative model that encodes beliefs about how sensory data are produced from hidden states of the world (and the body). These beliefs take the form of probability distributions, which are continuously updated as the organism compares predicted sensory inputs with incoming evidence.

Formally, the central quantity is the variational free energy F, which upper-bounds the negative log evidence. That is, it provides a tractable proxy for surprise under the generative model:

$$F[q(s)] = E_{q(s)}[\ln q(s)] - E_{q(s)}[\ln p(s, o)],$$

where $q(s)$ is an approximate posterior over hidden states s, and $p(s, o)$ is the generative model linking states to observations o. Minimising F with respect to $q(s)$ yields Bayesian perception:

$$q^*(s) \approx p(s \mid o),$$

realised via gradient flows on free energy [Friston, 2010, Buckley et al., 2017]. Expressed in differential form, perceptual inference corresponds to:

$$\dot{\mu} \approx -\frac{\partial F}{\partial \mu},$$

where μ denotes the sufficient statistics (e.g., mean) of the posterior beliefs.

Action arises from the same principle. Because actions a influence observations $o(a)$, an agent can minimise free energy by selecting actions that fulfil its own sensory predictions:

$$\dot{a} \approx -\frac{\partial F}{\partial a}.$$

In this view, "motor commands" are replaced with the optimisation of proprioceptive prediction errors:

$$\epsilon_{\text{proprio}} = o_{\text{proprio}} - g(\mu),$$

where $g(\mu)$ encodes predicted bodily states. Muscles act to reduce these errors, effectively bringing the world into alignment with the agent's expectations. This resolves the perception-action dichotomy: movement is simply inference expressed through the body [Friston, 2010, Friston et al., 2018].

A crucial implication of Active Inference is that intelligent behaviour fundamentally relies on a closed sensorimotor loop. A generative model without a body cannot complete the cycle of prediction, sensation, and correction, and therefore cannot enact the counterfactual sampling that underwrites adaptive behaviour [Pezzulo et al., 2024]. Embodiment is not an optional add-on but a computational necessity: action provides the very data required for self-evidencing, and perception is inseparable from the motoric means of testing hypotheses about the world [Costa et al., 2022].

This makes Active Inference particularly well-suited for understanding systems like the octopus, where cognition is distributed through sensorimotor structures rather than concentrated in a central brain. In such organisms, intelligence emerges from ongoing cycles of prediction and correction carried out at multiple levels of the nervous system, tightly coupled to the mechanics of the body and the statistics of the environment [Hochner, 2012, Hale, 2025, Carls-Diamante, 2022].

3.2 Hierarchical Control in Flexible Systems

The octopus provides a vivid example of a biological system in which control is distributed across multiple spatial and temporal scales [Hochner, 2012, Hale, 2025]. At the lowest level, each arm engages in rapid sensorimotor inference, continuously updating beliefs about local tactile, proprioceptive, and chemotactile states. These peripheral circuits implement their own prediction-error minimisation loops, allowing each limb to adjust its posture, initiate grasping motions, or explore environmental features with minimal involvement of the central brain [Olson and Ragsdale, 2023]. In effect, each arm maintains a generative model for the expected sensory consequences of its own movement, forming the basis of an autonomous action-perception cycle.

However, this autonomy is not unconstrained. Higher-level structures within the central brain encode broader contextual priors, such as whether the animal is hunting, escaping, camouflaging, or navigating-and these shape the space of permitted behaviours in the arms [Carls-Diamante, 2022, Ponte et al., 2022]. This forms what we describe as a pattern of *hierarchical bounded autonomy*: peripheral controllers optimise their own local free-energy landscapes, but within boundary conditions imposed by global goals. Under Active Inference, this corresponds to a nested hierarchy in which higher levels modulate the priors and precisions that govern lower-level inference [Friston, 2010, Friston et al., 2018].

Within this hierarchical architecture, global inference loops operate over slower timescales and larger spatial contexts. High level beliefs about task demands, environmental affordances, or imminent threats propagate downward as priors on expected sensory states. For example, during prey capture, the central brain sets expectations for coordinated reaching and envelopment, biasing peripheral circuits toward grasping-like trajectories [Hanlon and Messenger, 1996, Levy et al., 2015]. Conversely, during rapid escape, high-level priors favour symmetric propulsion and rapid withdrawal, altering the precision of proprioceptive prediction errors in each arm. In both cases, local and global inference processes interact dynamically: the periphery handles fast, high-bandwidth sensorimotor contingencies, while the central brain maintains coherence with the animal's overarching behavioural goals [Hochner, 2023, Hale, 2025].

This hierarchical separation of timescales and responsibilities enables the octopus to achieve control over an extraordinarily high-dimensional body without relying on centralised micro-management. The result is a flexible, adaptive system in which intelligence emerges from the coordinated interplay of distributed generative models operating at multiple levels of the nervous system [Friston et al., 2018].

4 Degrees of Freedom and the Impossibility of Centralised Control

Octopus arms present one of the most extreme challenges in biological motor control. Unlike vertebrate limbs, which rely on rigid skeletal constraints that drastically limit their degrees of freedom, each octopus arm is a muscular 'hydrostat' capable of bending, elongating, shortening, twisting, and stiffening at virtually any point along its length [Kier and Smith, 1985, Kier, 2016, Flash and Zullo, 2023]. This yields a control space with effectively thousands of continuous degrees of freedom, coupled through non-linear mechanical interactions and high-dimensional sensory feedback. From the perspective of classical control theory, such a system is vastly underdetermined and computationally unwieldy [Flash and Zullo, 2023, Nakajima et al., 2013].

A purely centralised controller would be required to invert this enormous, dynamically changing mapping from motor commands to sensory consequences in real time. Computing a full inverse model for a soft-bodied limb that must accommodate body-environment interactions, fluid dynamics, internal pressure changes, and continuous deformation, quickly becomes intractable. Even with modern computational resources, roboticists struggle to simulate a single soft arm at interactive timescales, let alone coordinate eight of them concurrently while integrating multimodal sensory input [Cianchetti et al., 2015, 2014, **?**, Tekinalp et al., 2024]. The central nervous system of the octopus, with its limited bandwidth and slow neural conduction speeds relative to vertebrates, could not feasibly issue precise, high-frequency control signals to each muscle fibre across all arms [Hochner, 2012, Hale, 2025].

The octopus circumvents this problem by distributing generative models throughout the body. Peripheral circuits perform fast, local inference about the expected sensory consequences of arm movements, effectively off-loading the computational burden from the central brain [Hochner, 2012, Olson and Ragsdale, 2023]. Under Active Inference, this corresponds to each arm minimising its own local *free energy*, aligning its beliefs with incoming sensory data while higher levels provide global priors that shape the overall behavioural context [Friston, 2010, Buckley et al., 2017]. Instead of explicitly computing motor commands, the system relies on hierarchical prediction errors: the central brain specifies coarse expectations, and the arms resolve fine-grained details through ongoing action-perception cycles [Friston et al., 2018, Hale, 2025].

This approach offers a principled solution to the curse of dimensionality. By decomposing the control problem into nested, distributed inference loops, the octopus transforms an otherwise intractable optimisation problem into one that is solved locally and in parallel [Nakajima et al., 2013, Hochner, 2012]. The success of this strategy has not been lost on the robotics community, where soft robotic arms, continuum manipulators, and morphologically adaptive agents increasingly adopt architectures inspired by octopus-like decentralisation, embedding computation and feedback inside the body itself [Cianchetti et al., 2015, Wang et al., 2017]. The octopus demonstrates that scalable control of high-dimensional, flexible morphologies does not require a monolithic controller, but rather a hierarchy of simple generative models operating at different spatial and temporal scales [Costa et al., 2022, Mazzaglia et al., 2022].

5 The Octopus as a Template for Embodied Intelligence

The octopus offers a compelling biological paradigm for rethinking how intelligent systems can be built. Rather than relying on a centralised controller with complete knowledge of the body and environment, the octopus exemplifies a model in which intelligence arises from the coordination of distributed, embodied generative processes [Hochner, 2012, Carls-Diamante, 2022, Hale, 2025]. This perspective challenges traditional AI architectures that separate perception, planning, and action, instead highlighting the advantages of systems that integrate these functions through ongoing sensorimotor interaction [Costa et al., 2022, Pezzulo et al., 2024].

For AI and robotics, the octopus suggests an alternative organisational principle: multiple generative models arranged hierarchically, with local controllers performing rapid inference on peripheral sensory states and higher-level modules encoding longer-term goals, task contexts, and environmental contingencies [Friston, 2010, Mazzaglia et al., 2022]. In this architecture, local sensory-motor loops operate autonomously but within the constraints of global priors. This mirrors the octopus's pattern of hierarchical bounded autonomy, in which each arm resolves its own prediction errors while the central brain guides overall behaviour through coarse, context-dependent expectations [Hochner, 2023, Carls-Diamante, 2022, Hale, 2025].

Such an arrangement has several practical advantages. First, it enables scalability: by distributing inference across many local modules, the system avoids the exponential explosion of complexity associated with centrally computing actions for all degrees of freedom. Second, it yields robustness: damage or perturbation to one part of the system need not incapacitate the whole, as local controllers retain the ability to adapt their behaviour in situ. Third, it supports flexibility, allowing the agent to exploit the dynamics of a soft body rather than attempting to rigidly control it [Nakajima et al., 2013].

These principles directly inform the design of next-generation robotic agents. Soft robots, which face similar control challenges to biological muscular hydrostats, increasingly employ decentralised sensing, embedded computation, and morphology-aware control [Cianchetti et al., 2015, Wang et al., 2017]. Incorporating Active Inference into these systems provides a principled mechanism for integrating prediction, action, and adaptive belief updating at multiple levels [Costa et al., 2022]. Tactile-driven control, embodied state estimation, and distributed proprioceptive inference become natural components of the architecture rather than add-ons to a central planner [Mazzaglia et al., 2022].

Ultimately, the octopus demonstrates that intelligence does not require a single, centralised cognitive engine, but can emerge from the interaction of many simple inferential systems embedded within a flexible body. By adopting analogous architectures in artificial agents, we may develop robots that are more adaptive, scalable, and capable of thriving in uncertain, unstructured environments [Costa et al., 2022, Pezzulo et al., 2024].

6 Open Questions and Future Directions

The octopus raises profound questions about the nature of intelligence and the architectures required to support adaptive, flexible behaviour. While its nervous system has been extensively characterised anatomically, many aspects of its computational organisation remain poorly understood [Hochner, 2012, Ponte et al., 2022, Hale, 2025]. One outstanding question concerns the precise division of labour between central and peripheral controllers: how do arm-level generative models interface with higher-level priors, and how is conflict resolved when local and global objectives diverge? Understanding this interaction may shed light on how biological systems negotiate autonomy and coordination across multiple scales [Carls-Diamante, 2022, Olson and Ragsdale, 2023].

A second challenge concerns how to translate octopus-style distributed control into artificial systems. While soft robotics has begun to explore decentralised architectures, it remains unclear which computational principles are essential and which are contingent on specific biological constraints [Cianchetti et al., 2015]. Do artificial agents require arm-level generative models, or can similar benefits be achieved through more abstracted forms of local inference? And how should designers balance autonomy and central guidance to achieve the degree of robustness, scalability, and dexterity exhibited by octopus arms? These questions invite systematic exploration using Active Inference as a formal framework for distributed sensorimotor control [Costa et al., 2022, Mazzaglia et al., 2022].

Finally, the octopus offers an intriguing perspective on broader debates about embodied cognition and the future of artificial general intelligence. If intelligence can emerge from the interaction of many simple inferential units distributed through a flexible, sensorimotor body, then highly centralised architectures may not be the most natural pathway to AGI [Pezzulo et al., 2024]. Instead, embodied systems that tightly couple generative models to the physical dynamics of the world may prove more scalable and more capable of adaptive, context-sensitive behaviour. Exploring octopus-inspired architectures therefore has implications not only for robotics, but for the foundations of cognitive science and the design principles underlying future intelligent systems [Pezzulo et al., 2024, Costa et al., 2022].

Taken together, these open questions highlight the need for interdisciplinary research that spans neuroscience, computational modelling, robotics, and theoretical AI. The octopus, with its unique combination of decentralisation, embodiment, and behavioural sophistication, provides an exceptional platform for such inquiry [Hochner, 2012, Hale, 2025, Carls-Diamante, 2022].

7 Conclusion

The octopus exemplifies a form of intelligence that is deeply embodied, distributed, and tightly integrated with the mechanics of a soft, high-dimensional body [Hochner, 2012, Hale, 2025]. Rather than relying on centralised control, the octopus achieves behavioural flexibility through the coordination of peripheral generative models operating within global behavioural contexts. This architecture-combining local autonomy with hierarchical constraint allows the animal to solve control problems that would be computationally intractable for a unitary controller.

Active Inference provides a principled framework for understanding this organisation. By treating perception and action as intertwined processes of prediction-error minimisation, it naturally accommodates systems in which inference is distributed across multiple spatial and temporal scales [Friston, 2010, Buckley et al., 2017, Friston et al., 2018]. The octopus thus serves not only as a biological curiosity, but as a valuable template for rethinking how intelligent agents might be constructed: from soft robots capable of adaptive manipulation to future artificial systems that rely on embodied, decentralised generative models [Cianchetti et al., 2015, Costa et al., 2022].

In embracing the lessons of octopus cognition, we gain a richer understanding of biological intelligence and a roadmap for designing artificial agents that are more robust, scalable, and capable of thriving in uncertain environments [Pezzulo et al., 2024].

References

Christopher L. Buckley, Chang Sub Kim, Simon McGregor, and Anil K. Seth. The free-energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*, 81:55–79, 2017. doi: 10.1016/j.jmp.2017.09.004.

Sidney Carls-Diamante. Where is it like to be an octopus? *Frontiers in Systems Neuroscience*, 16:840022, 2022. doi: 10.3389/fnsys.2022.840022.

Matteo Cianchetti, Tommaso Ranzani, et al. Soft robotics technologies to address shortcomings in today's minimally invasive surgery: The STIFF-FLOP approach. *Soft Robotics*, 1(2):122–131, 2014. doi: 10.1089/soro.2014.0001.

Matteo Cianchetti et al. Bioinspired locomotion and grasping in water: the soft eight-arm OCTOPUS robot. *Bioinspiration & Biomimetics*, 10(3):035003, 2015. doi: 10.1088/1748-3190/10/3/035003.

Lancelot Da Costa, Pablo Lanillos, Noor Sajid, Karl J. Friston, et al. How active inference could help revolutionise robotics. *Entropy*, 24(3):361, 2022. doi: 10.3390/e24030361.

Tamar Flash and Leonardo Zullo. Biomechanics, motor control and dynamic models of the soft limbs of octopuses and their robotic inspirations. *Journal of Experimental Biology*, 226 (Suppl_1):jeb245295, 2023. doi: 10.1242/jeb.245295.

Karl J. Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010. doi: 10.1038/nrn2787.

Karl J. Friston, Thomas Parr, and Bert de Vries. The graphical brain: Belief propagation and active inference. *Network Neuroscience*, 2(4):381–414, 2018. doi: 10.1162/netn_a_00021.

Melina E. Hale. The octopus as a comparative model for neural control of limb movement. *Current Opinion in Neurobiology*, 83:102759, 2025. doi: 10.1016/j.conb.2023.102759.

Roger T. Hanlon and John B. Messenger. *Cephalopod Behaviour*. Cambridge University Press, Cambridge, UK, 1996.

Binyamin Hochner. An embodied view of octopus neurobiology. *Current Biology*, 22(20):R887–R892, 2012. doi: 10.1016/j.cub.2012.09.001.

Binyamin Hochner. Embodied mechanisms of motor control in the octopus. *Current Biology*, 33 (17):R933–R947, 2023. doi: 10.1016/j.cub.2023.04.072.

William M. Kier. The musculature of coleoid cephalopod arms and tentacles. *Frontiers in Cell and Developmental Biology*, 4:10, 2016. doi: 10.3389/fcell.2016.00010.

William M. Kier and Kathleen K. Smith. Tongues, tentacles and trunks: the biomechanics of movement in muscular-hydrostats. *Zoological Journal of the Linnean Society*, 83(4):307–324, 1985. doi: 10.1111/j.1096-3642.1985.tb00878.x.

Mikk Kuuspalu et al. Multiple nerve cords connect the arms of octopuses. *Current Biology*, 32 (20):R1145–R1146, 2022. doi: 10.1016/j.cub.2022.08.024.

Ganit Levy, Guy Sumbre, Tamar Flash, and Binyamin Hochner. Arm coordination in octopus crawling involves unique motor control strategies. *Current Biology*, 25(9):1195–1200, 2015. doi: 10.1016/j.cub.2015.02.064.

Pietro Mazzaglia, Tim Verbelen, Ozan Çatal, and Bart Dhoedt. The free energy principle for perception and action: A deep learning perspective. *Entropy*, 24(2):301, 2022. doi: 10.3390/e24020301.

Kohei Nakajima, Hiroki Hauser, Tao Li, and Rolf Pfeifer. A soft body as a reservoir: case studies in a dynamic model of octopus-inspired soft robotic arm. *Frontiers in Computational Neuroscience*, 7:91, 2013. doi: 10.3389/fncom.2013.00091.

Ethan E. Olson and Chad W. Ragsdale. Toward an understanding of octopus arm motor control. *Integrative and Comparative Biology*, 63(1):168–181, 2023. doi: 10.1093/icb/icad008.

Giovanni Pezzulo, Thomas Parr, Paul Cisek, Andy Clark, and Karl J. Friston. Generating meaning: Active inference and the scope and limits of passive AI. *Trends in Cognitive Sciences*, 28(2):97–112, 2024. doi: 10.1016/j.tics.2023.10.002.

Rolf Pfeifer and Josh Bongard. *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT Press, Cambridge, MA, 2007.

Giovanna Ponte, Cinzia Chiandetti, David B. Edelman, Pamela Imperadore, Eleonora Maria Pieroni, and Graziano Fiorito. Cephalopod behavior: From neural plasticity to consciousness. *Frontiers in Systems Neuroscience*, 15:787139, 2022. doi: 10.3389/fnsys.2021.787139.

Arman Tekinalp et al. Topology, dynamics, and control of a muscle-architected octopus-analog muscular hydrostat. *Proceedings of the National Academy of Sciences*, 121(20):e2318769121, 2024. doi: 10.1073/pnas.2318769121.

Liyu Wang, Surya Girinatha Nurzaman, and Fumiya Iida. Soft-material robotics. *Foundations and Trends in Robotics*, 5(3–4):191–259, 2017. doi: 10.1561/2300000055.