

Toward Enactive Artificial Intelligence

Banafsheh Rafiee¹, Richard S. Sutton^{2,3}

rafiee@ualberta.ca, rsutton@ualberta.ca

¹ Independent Researcher

² Department of Computing Science, University of Alberta, Canada

³ Alberta Machine Intelligence Institute (Amii)

Abstract

In this paper, we advocate for incorporating enactive approaches to perception and cognition into artificial intelligence (AI). Enactive approaches view perception as an active, skillful engagement with the world, where agents perceive by acting and by understanding how their actions shape their experience. This contrasts with classical views that treat perception as a passive internal process in which the brain receives sensory input, processes it, and issues commands for action. Enactive views emphasize the dynamic, embodied, and interactive character of perception, grounded in the lived experience of agents embedded in their environments. We identify and develop four key enactive concepts that we find most relevant to AI: experience, action–perception inseparability, autonomy, and embodiment. Much of mainstream AI, from classical rule-based systems to large language models, has largely neglected these insights, treating cognition as internal processing detached from embodied interaction and intrinsic normativity. Reinforcement learning (RL), however, exhibits structural resonance with enactive principles through its emphasis on action, agent–environment interaction, feedback-driven adaptation, and agent-centered evaluation. However, this resonance should not be taken as theoretical equivalence, as RL approximates some enactive insights, but key elements remain absent or weakly developed. Building on this analysis, we suggest a broader incorporation of enactive ideas into both mainstream AI and RL.

1 Enactive Cognition

Enactive perspectives on cognition and perception view cognition, perception, and action as mutually constitutive (Varela et al., 1991). Perception does not precede action nor does it merely guide it; rather, they unfold together in ongoing interaction with the environment. Actions shape perceptual experience by modulating sensory input, while perception, in turn, is realized through patterns of action. In other words, perception consists of mastering the ways in which our bodily movements produce changes in our sensory stimulations, known as *sensorimotor contingencies* (Noë, 2004). To perceive is thus to skillfully navigate these contingencies, to understand how the world reveals itself through active exploration, rather than to passively register incoming data.

Enactive perspectives contrast with classical representationalist views, in which perception is framed as a passive process where sensory input is transmitted through the body’s sensory apparatus and transformed into internal representations of the world. These internal representations function as stand-ins for the environment, providing an internal map that allows the system to reason about and act upon the environment without requiring the environment’s direct presence (Marr, 1982; Fodor, 1983). Enactive perspectives reject this view, emphasizing continual and real-time interaction with the environment. In the representationalist views, the brain functions as a central processing unit: it encodes sensory information, manipulates symbols, and generates action plans based on its internal representations, assuming a sharp distinction between the perceiving agent and the external world,

with perception reduced to the accurate construction and manipulation of internal representations. By contrast, in enactive perspectives, the adequacy of perception is not measured by representational fidelity but by the agent’s capacity for skillful engagement with the environment.

The term enactive was first introduced within Enactivism (Varela et al., 1991), a subfield of cognitive science, which frames cognition as *enacted* rather than pre-given. To enact is to bring forth or constitute a meaningful world through embodied engagement with the environment. O’Regan and Noë (2001) further argued that perception is not something that happens to an organism but something the organism actively does. These works contributed to a broader research landscape challenging representationalist accounts of cognition, including embodied and situated approaches. Embodied cognition (Clark, 1997) emphasized the role of the body in shaping the mind. Situated cognition (Hutchins, 1995; Lave, 1988) posits that cognitive processes are deeply embedded in the activity, tools, and social interactions rather than arising solely from internal processes.

While we borrow the term enactive from enactivism, the perspective developed in this paper is not restricted to that tradition. We bring together a broader set of ideas that appear across philosophy, psychology, and cognitive science. These approaches are not monolithic and differ in important respects. However, they share a common orientation: they treat cognition as arising through active, embodied, and situated engagement with the environment. Our aim is therefore not to provide a historical survey, but to articulate a coherent perspective that synthesizes these related strands.

Enactive ideas trace back to phenomenology, drawing on Husserl, Heidegger, and Merleau-Ponty, who emphasized the inseparability of organism and world, bound together through continual engagement. Husserl, argued that perception is not a matter of constructing internal representations of the world but of attending to things as they are directly given in lived experience. Heidegger radicalized this insight with his notion of being-in-the-world, which rejects the picture of a detached subject representing the external world. Instead, he emphasized that human existence is fundamentally situated: we do not first perceive the world and then interpret it, but already are immersed in practical, meaningful contexts that shape understanding from the outset. Merleau-Ponty emphasized the body not as an object among others, but as the medium through which the world is experienced, with perception depending on our bodily abilities to move and interact.

In psychology, related enactive ideas appear in early traditions. Gestalt psychology proposed that perceptual experience is shaped by the organism’s own organizing activity, anticipating the emphasis on perception as an active, world-involving process (Koffka, 1935). Goldstein (1939) highlighted the inseparability of perception, action, and context, viewing the organism’s behavior as an integrated, adaptive response rather than a sequence of internal computations. Gibson’s ecological approach treated perception as inherently action-dependent: organisms perceive their environment in terms of what actions they can take, possibilities for action that depend on both the situation and the agent’s own embodied capacities. These possibilities for action, or *affordances*, are not fixed properties of the environment but emerge through interaction with the world (Gibson, 1979).

With this foundation, we turn to enactive ideas that we consider most pertinent for AI research. We identify and develop four key enactive concepts: experience, action–perception inseparability, autonomy, and embodiment, and examine how different AI approaches align with them. Our discussion spans two levels of analysis: one from the enactive perspectives, and one at the computational level of AI systems. These levels must be distinguished. The comparisons we draw are structural aiming to identify shared patterns and organizing principles rather than claims of equivalence.

2 Experience

Enactive cognition is grounded in experience. Central to experience is the continual interaction between agent and environment in which each influences and responds to the other. Perception is not a one-time act of constructing an internal model of the world that functions as a stand-in for the world. No internal model, however elaborate, can capture the inexhaustible richness and open-ended variability of the world. The world is not a fixed set of features or a static object to be represented; it

is a dynamic, evolving field of possibilities that unfolds differently depending on the agent’s activity, context, and engagement history. In other words, the world is an ongoing situation that the agent co-constitutes through action (Varela et al., 1991; Noë, 2004). Because the world exceeds any finite description, the agent must remain in continual interaction with it, drawing on real-time feedback to adjust its actions, recalibrate its expectations, and refine its understanding. In this sense, as Rodney Brooks puts it, “the world is its own best model”: the most reliable, up-to-date, and fine-grained information is always available in the world itself, not in internal surrogates.

In the strong enactive sense, experience is not just continual interaction, but skillful, normative, and embodied. First, experience is skillful: through interaction, agents acquire skills that shape how the world appears to them. The environment shows up as a field of affordances, soliciting certain activities while resisting others. Second, experience is normative rather than neutral. The agent’s actions can succeed or fail, fit well or poorly with the situation, and the agent continually adjusts its behavior in response to this ongoing evaluation. Third, experience is embodied: it is shaped by the ways in which body enables and constrains possible actions. In this section, we focus on the continual interaction aspect of experience; the other aspects are addressed in later sections.

In classic rule-based AI, experience was largely absent. In *What Computers Can’t Do*, Dreyfus (1992) argues that rule-based symbolic systems fall short of replicating human intelligence precisely because they lack the experiential basis. Machine learning introduced hints to the importance of experience by centering learning on data—the byproduct of experience. However, it did not fully embrace experience, as most methods, especially supervised learning, depend on data gathered and labelled by humans, as opposed to data collected by the agent itself through direct interaction with its environment. Moreover, supervised learning treats cognition as a one-time process of learning from fixed datasets rather than as continual engagement and adaptation with the environment.

By contrast, reinforcement learning (RL) placed experience at the core of the learning process, enabling agents to gather their own data by actively interacting with their environment. Although RL does not capture the full, skillful, embodied, and normative sense of experience in the enactive sense, it approximates its emphasis on continual interaction through learning from action and feedback. The importance of allowing AI agents to continually interact with the environment and learn from data gathered by themselves has been well articulated by Silver and Sutton (2025). The central argument is that data must continually improve alongside the agent’s capabilities, a process made possible only through the agent’s own experience.

A related AI research area is continual learning, which studies how agents can learn from a stream of changing data rather than a fixed dataset. While continual learning does not necessarily emphasize agent-gathered data or sustained interaction with the environment, it aligns with enactive views in emphasizing ongoing adaptation under non-stationarity. Continual learning has many definitions (Abel et al., 2023; Khetarpal et al., 2022; Ring, 1997), closely related research areas (Thrun, 1998; Mitchell et al., 2018), and concepts such as loss of plasticity and catastrophic forgetting (Dohare et al., 2024; McCloskey and Cohen, 1989). Despite differences, these approaches share the insight that agents must learn continually to remain adaptive in complex, non-stationary environments. Continual learning becomes particularly important under the *Big World Hypothesis* (Javed and Sutton, 2024), which holds that the world is orders of magnitude larger than the agent and thus is perceived as ever-changing from the agent’s perspective. This assumption echoes the enactive perspectives: because the world exceeds what the agent can internally capture, effective behavior requires continual interaction in which the agent continually updates its understanding and skills.

3 Action–perception inseparability

Enactive approaches emphasize the inseparability of action and perception, viewing perception as understanding sensorimotor contingencies, that is how actions produce changes in sensory input. Visual experience, for example, depends on mastering the patterns of change in visual input as a function of eye, head, and body movements, such as knowing that moving the eyes leftward causes objects in the visual field to shift rightward (Noë, 2004). Similarly, auditory perception relies on

sensorimotor contingencies: head movements systematically alter the acoustic input at each ear, and these changes in timing and intensity, provide reliable cues for locating sound sources (van Dijk and Withagen, 2016). Tactile perception likewise depends on sensorimotor contingencies: surface qualities such as texture or shape are revealed through changes in sensory input as the hand moves across them. Fine textures produce distinct vibration patterns depending on movement speed and direction, which are learned through active touch (Lederman and Klatzky, 1987).

Crucially, the emphasis on sensorimotor contingencies does not imply that perception is merely the passive monitoring of these patterns. Rather, the enactive claim is stronger: perception is itself a form of skillful activity (Noë, 2004)—to perceive is to act. Rather than waiting for input to arrive and then processing it internally, agents engage in purposeful movements to reveal the structure of their environment. Agents often act in order to reveal, stabilize, or make sense of perceptual information, for example, by moving the head to resolve ambiguity in a scene.

Perception, in this sense, is thought of as a feedback loop in which the agent’s understanding of a situation and its response to it become increasingly refined over time. As the agent gets a better grasp of the situation, it becomes able to respond more skillfully and in a more refined manner. The situation comes to afford more refined and differentiated responses, presenting not merely general possibilities for action but more precise and nuanced affordances that were not previously accessible to the agent. The agent’s more refined responses, in turn, further refine its understanding by revealing new aspects of the situation. Thus, refined understanding and response mutually reinforce each other. This feedback loop was referred to as *the intentional arc* by Merleau-Ponty (1945).

In this feedback loop, the agent aims toward a more coherent understanding of the situation, the *maximal grip* (Merleau-Ponty, 1945). In this process, the agent continually adjusts its posture, attention, and movements to bring the situation into sharper focus, both perceptually and motorically. The agent is naturally drawn toward states of greater stability, clarity, and alignment between its body and the environment. This is done often in a fluid and unreflective manner: much like leaning in to see a detail more clearly or tilting the head to resolve an ambiguous sound, the agent acts in ways that reduce uncertainty and restore perceptual equilibrium (Dreyfus, 2002). These refinements are not guided by explicit reasoning but by a bodily sense of tension when the current state deviates from an optimal relation, and by the felt relief that comes with moving closer to that optimal state.

Action–perception inseparability is largely absent in mainstream AI, where perception is dominantly understood as preceding action. This mirrors the representationalist view, where perception is treated as passive extraction of information from sensory input, later used to guide behavior. This view is reinforced by large-scale video-generation systems, where purely observational learning is sometimes taken to yield understanding of phenomena like intuitive physics without active interaction. Yet the competence of such systems extends only as far as the world continues to behave like their data. As noted by Goddu et al. (2024), a system that has learned traffic light regularities may accurately predict the sequence of green, yellow, and red, but this amounts only to tracking regularities, not to understanding how the sequence could be changed. When the light malfunctions, is interrupted, or action is required to alter the situation, such as stopping traffic, triggering a crossing signal, or diagnosing a fault, the system has nothing to fall back on. By contrast, an enactive system does not merely anticipate what typically occurs, but can intervene, correct, and explore when expectations fail. The difference is not one of accuracy but of kind: a generative video model can continue a pattern, whereas an enactive system can determine what to do next when the pattern breaks.

In contrast to the mainstream view, several strands of AI have explored action–perception coupling. Early work, such as the Pengi system (Agre and Chapman, 1987), showed that competent behavior can arise from tightly coupled perception–action loops without explicit internal world models, where perception is directly implicated in triggering action rather than serving as a detached input stage. Ballard (1991) and Whitehead and Ballard (1990) further demonstrated that perception is an active process of information acquisition: by moving the eyes or body to gather task-relevant information, the agent learns control policies that determine what is sensed, making vision dependent on action strategies rather than passive encoding. Situated planning (Chapman, 1989) highlighted that even

planning is not separable from perception in realistic settings: plans must be continually revised in response to perceptual feedback, since viable actions depend on the evolving interaction with the environment rather than a pre-given model of it. Behavior-based robotics (Brooks, 1991) emphasized this coupling at the architectural level, showing that coherent behavior can emerge without centralized representations precisely because perception and action are distributed across interacting sensorimotor routines, rather than separated into sequential stages. In parallel, formal accounts such as predictive coding, the perception–action cycle, and active inference (Rao and Ballard, 1999; Tishby and Polani, 2010; Friston et al., 2017) framed behavior as a unified loop in which perception and action are jointly organized, either through hierarchical prediction error minimization, information-theoretic optimization, or variational inference under a shared objective. Complementing these approaches, several formal frameworks capture aspects of action–perception coupling by enabling agents to anticipate how sensory inputs evolve under their actions, and by treating action, observation, and internal state as mutually dependent. This includes General Value Functions (Sutton et al., 2011), Predictive State Representations (Littman et al., 2002), Partially Observable Markov Decision Processes (Kaelbling et al., 1998), and world models (Ha and Schmidhuber, 2018).

Similar ideas appear in recent work. Sutton et al. (2022) introduced the *STOMP* framework in which the agent learns subtasks that maximize distinct aspects of perception. As the agent improves at each subtask, it also refines its model of how actions affect perception, creating a feedback cycle in which refined actions lead to refined perception, and vice versa, echoing the *intentional arc*. Machado et al. (2023) proposed a loop where an agent’s representation of its environment and behavior are refined through constant interaction. As the agent develops better ways of representing the world, it discovers behaviors, which in turn enrich representations. Another line develops affordance theory in RL, where the environment constrains and suggests possible actions (Graves et al., 2022; Khetarpal et al., 2020). Rather than treating all actions as always available, agents learn which are afforded in each situation, mirroring selective, skillful responses in enactive perception.

4 Autonomy

Autonomy plays a central role in enactive cognition. Agents are not passive responders to external stimuli, but self-organizing systems whose perception is shaped by their goals and needs (Di Paolo, 2005). This autonomy is often grounded in the notion of *autopoiesis*: agents are self-producing and self-maintaining systems that actively sustain their own organization. As such, perception reflects what matters to the organism from the standpoint of its continued viability (Jonas, 1966; Di Paolo et al., 2017). What is perceived is not simply what is present, but what is relevant to the agent’s self-maintenance. The world does not show up as a neutral array of features, but as meaningful relative to the agent’s goals and needs (Gallagher, 2017).

Autonomy, in this sense, gives rise to normativity. Because the agent must continually maintain itself as a coherent system, its interactions with the environment are not neutral; they can succeed or fail relative to the agent’s continued viability. Normativity, in this sense, arises from the agent’s need to maintain its own organization, rather than being imposed from the outside. This normative structure, in turn, shapes what is relevant to the agent. The environment is therefore not encountered merely in terms of what is, but in terms of what matters: what supports or threatens the agent’s ongoing self-maintenance. In other words, it serves to constrain the vast space of possible generalizations: in a world where everything could resemble everything else in countless ways, the agent’s goals and needs filter which similarities are relevant (Dreyfus, 2002). What counts as similar is shaped not by abstract resemblance, but by the agent’s ongoing movement toward success in a particular context.

To analyze the notion of autonomy and normativity within AI, it is useful to distinguish between two related questions. First, whether the agent can evaluate its own behavior, that is, whether it has a sense of success and failure grounded in its own activity. Second, whether the criteria for success and failure arise from the agent itself or are imposed externally.

In much of AI, particularly supervised learning, both aspects are absent. The system does not evaluate its own performance; correctness is determined only by comparison to human-provided labels.

For any given input, the model has no access to whether its output is successful unless an external signal is supplied. Evaluation is therefore not part of the agent’s ongoing activity, but occurs entirely externally. Moreover, success criteria are fully externally specified by the dataset and labeling process. Thus, the system neither self-evaluates nor has its own standards of success. This holds for large language models as well: although trained with self-supervised objectives such as next-token prediction, they effectively learn by imitating patterns in human-generated data and are not able to evaluate their own outputs without outsourcing evaluation to external signals.

A similar pattern appears in classical symbolic planning and automated reasoning (Newell and Simon, 1956; Ghallab et al., 2004) where success is defined by satisfaction of explicitly specified conditions. This introduces a notion of success and failure, but not as grounded in the agent’s activity. The system does not assess its performance as it acts; instead, evaluation is limited to checking whether predefined conditions are met, typically as a binary outcome. There is no ongoing, graded assessment of performance during interaction. Furthermore, success criteria are externally imposed: goals are specified in advance and function as fixed conditions to be achieved.

In planning and control-based systems, evaluation becomes continuous and embedded in the agent’s interaction with the environment. These systems track their state relative to a desired setpoint, trajectory, or cost function, and adjust their behavior through feedback. This enables the system to register deviations and to distinguish, at each moment, whether its behavior is improving relative to the target. In this sense, the system possesses an ongoing, graded assessment of success. However, this evaluation remains tied to instantaneous deviation from a predefined objective, that is, how well the current state matches the target. Moreover, success criteria, such as the target state or cost function, are still specified externally.

RL marks a significant shift in that the agent evaluates its behavior through experience. By interacting with the environment and receiving rewards, the agent can assess how well it is doing across entire trajectories, rather than only in terms of its current state. This allows actions to be evaluated in terms of their consequences, including delayed effects, introducing a temporally extended notion of success. While control systems answer "how close am I to the target now?", RL answers "was this behavior good, given what it led to over time?". In this sense, RL captures an important aspect of normativity: evaluation is embedded in the agent’s ongoing interaction and learning process. However, evaluation criteria remain externally defined through the reward function.

Several lines of work have sought to relax dependence on externally specified criteria by developing more agent-centered notions of evaluation. The perception–action cycle (Tishby and Polani, 2010) and active inference (Friston et al., 2017) both make evaluation emergent from the agent’s ongoing interaction with the environment, arising either through optimization of information flow or through minimization of expected surprise. Intrinsic motivation develops internal reward signals that drive exploration and skill acquisition without external supervision, often based on learning progress, rewarding improved ability to predict or control aspects of the environment (Singh et al., 2004; Oudeyer et al., 2007; Şimşek and Barto, 2006). Goal discovery approaches (Andrychowicz et al., 2017) construct or reinterpret objectives from experience, expanding what counts as achievement beyond external rewards. These methods move toward more agent-centered forms of evaluation, in which the standards for success are less rigidly fixed in advance. Nevertheless, full autonomy in the enactive sense, where normativity arises from the agent’s own organization, remains unrealized.

5 Embodiment

Enactive perception is fundamentally embodied: the shape, structure, and capacities of the body influence how an organism perceives (Clark, 1997; Thompson, 2007). More specifically, perception is the mastery of sensorimotor contingencies that depend on the body. Sensorimotor contingencies are not abstract mappings between inputs and outputs that could be specified independently of embodiment; they are grounded in the specific capacities of a body: what movements it can perform, how it can explore, and how its sensory systems are organized. Morphological factors such as joint structure, muscle distribution, and sensory placement constrain and enable the forms of engagement

available to the agent, and determine the space of possible sensorimotor contingencies. Embodied robotics supports this claim: altering a robot’s physical form can significantly change its perceptual abilities (Pfeifer and Bongard, 2006). The body is therefore not an optional component added after the fact, but the condition for perception to be possible in the first place (Merleau-Ponty, 1945)

Embodiment structures what counts as perceptually relevant. The body constrains the forms of engagement available to an agent, and therefore, shapes which similarities in the environment stand out as meaningful. Gibson’s notion of affordances captures this idea: features of the environment appear as “graspable,” “climbable,” or “passable” only relative to the agent’s bodily capacities. Without a body capable of such actions, these distinctions do not exist as meaningful features. In this way, the body does not simply constrain what can be perceived, but provides a structure for understanding the world in terms of action-relevant similarities (Dreyfus, 2002).

Embodiment also connects to autonomy through autopoiesis: an agent is autonomous insofar as it is a self-maintaining system that continually produces and preserves its own organization. The body is not peripheral in this process; it is the site where this self-production is realized. The boundaries, processes, and interactions constituting the agent are grounded in its embodied organization. In this sense, embodiment makes autonomy possible: without a body that sustains and regulates itself through interaction with the environment, there is no autonomous system in the enactive sense.

Mainstream AI largely treats perception in a disembodied manner. Even when trained on large-scale multimodal data, many systems learn mappings from inputs to internal representations without any dependence on sensorimotor engagement or bodily structure. In this setting, perception is decoupled from the kinds of structured sensorimotor contingencies that characterize embodied agents, and instead reduced to pattern recognition over static datasets (Bender et al., 2021; Bommasani et al., 2021). Therefore, these systems remain limited in their ability to develop genuinely situated perception or to adapt through active exploration of novel environments (Harnad, 1990; Chemero, 2009).

Embodied RL and robotics often treat embodiment as an external constraint rather than a constitutive principle of cognition. Many robotic systems still rely on modular architectures that separate perception, planning, and control, preserving a classical decomposition of cognitive functions (Brooks, 1991; Clark, 1997). In such systems, the body functions primarily as an interface for executing pre-computed policies, rather than as a source of structure that shapes perception. Similarly, extensive reliance on simulation and offline training further distances learning from the full variability and constraint structure of real sensorimotor interaction (Jakobi et al., 1995; Cully et al., 2015).

By contrast, work in soft robotics and morphological computation demonstrates that bodily structure can play an active computational role, simplifying control and shaping behavior through physical dynamics (Rus and Tolley, 2015; Pfeifer and Bongard, 2006; Zahedi and Ay, 2013). Despite these insights, such approaches remain peripheral within mainstream robotics and AI, where disembodied data-driven learning and modular system design continue to dominate. Consequently, the enactive emphasis on embodied sensorimotor engagement has yet to be fully integrated into AI.

6 Conclusion

Mainstream AI has largely failed to appreciate the enactive insights. By contrast, RL exhibits several structural resonances: It enables agents to generate their own experience through trial-and-error, places action at the center of learning, and introduces a temporally extended notion of evaluation through reward. However, this alignment is partial. Evaluation remains externally specified through reward functions, and is not grounded in the agent’s own organization; action–perception inseparability is not fully realized, with perception still typically treated as preceding action; and embodiment is treated as an implementation detail rather than a constitutive condition for cognition.

We suggest a deeper incorporation of enactive ideas into RL and AI. This paper takes a step toward articulating key enactive concepts and relating them to existing AI frameworks, but it does not operationalize them. A key future direction is to make these ideas more precise and testable. This includes questions such as: what constitutes a higher degree of action–perception inseparability?

What benchmarks capture skillful engagement rather than pattern reproduction? What does self-maintenance mean for artificial agents: battery state, hardware integrity, or learned competence? What counts as embodiment in AI: a robot body or a software agent with tools and APIs? Addressing such questions is essential for moving from conceptual alignment toward enactive AI in practice.

References

- Abel, D., Barreto, A., Van Roy, B., Precup, D., van Hasselt, H. P., and Singh, S. (2023). A definition of continual reinforcement learning. *Advances in Neural Information Processing Systems*, 36:50377–50407.
- Agre, P. E. and Chapman, D. (1987). Pengi: An implementation of a theory of activity. In *Proceedings of the sixth National conference on Artificial intelligence-Volume 1*, pages 268–272.
- Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Pieter Abbeel, O., and Zaremba, W. (2017). Hindsight experience replay. *Advances in neural information processing systems*, 30.
- Ballard, D. H. (1991). Animate vision. *Artificial intelligence*, 48(1):57–86.
- Bender, E. M., Gebru, T., McMillan-Major, A., and Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623.
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, M., et al. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47(1-3):139–159.
- Chapman, D. (1989). *Instruction Use in Situated Activity*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA. Ph.D. thesis.
- Chemero, A. (2009). *Radical embodied cognitive science*. MIT Press.
- Clark, A. (1997). *Being There: Putting Brain, Body, and World Together Again*. MIT Press, Cambridge, MA.
- Cully, A., Clune, J., Tarapore, D., and Mouret, J.-B. (2015). Robots that can adapt like animals. *Nature*, 521(7553):503–507.
- Di Paolo, E. A. (2005). Autopoiesis, adaptivity, and an enactive approach to agency. *Phenomenology and the Cognitive Sciences*, 4(4):429–452.
- Di Paolo, E. A., Buhrmann, T., and Barandiaran, X. E. (2017). *Sensorimotor Life: An Enactive Proposal*. Oxford University Press, Oxford.
- Dohare, S., Hernandez-Garcia, J. F., Lan, Q., Rahman, P., Mahmood, A. R., and Sutton, R. S. (2024). Loss of plasticity in deep continual learning. *Nature*, 632(8026):768–774.
- Dreyfus, H. L. (1992). *What Computers Still Can’t Do: A Critique of Artificial Reason*. MIT Press, Cambridge, MA, revised edition. Originally published in 1972 as *What Computers Can’t Do*.
- Dreyfus, H. L. (2002). Intelligence without representation—merleau-ponty’s critique of mental representation: The relevance of phenomenology to scientific explanation. *Phenomenology and the Cognitive Sciences*, 1(4):367–383.
- Fodor, J. A. (1983). *The Modularity of Mind*. MIT Press, Cambridge, MA.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., and Pezzulo, G. (2017). Active inference: a process theory. *Neural computation*, 29(1):1–49.

-
- Gallagher, S. (2017). *Enactivist Interventions: Rethinking the Mind*. Oxford University Press, Oxford.
- Ghallab, M., Nau, D., and Traverso, P. (2004). *Automated Planning: theory and practice*. Elsevier.
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston.
- Goddu, M. K., Noë, A., and Thompson, E. (2024). Llms don’t know anything: reply to yildirim and paul. *Trends in Cognitive Sciences*.
- Goldstein, K. (1939). *The Organism: A Holistic Approach to Biology Derived from Pathological Data in Man*. American Book Company.
- Graves, D., Günther, J., and Luo, J. (2022). Affordance as general value function: a computational model. *Adaptive Behavior*, 30(4):307–327.
- Ha, D. and Schmidhuber, J. (2018). World models. *arXiv preprint arXiv:1803.10122*.
- Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1-3):335–346.
- Hutchins, E. (1995). *Cognition in the Wild*. MIT Press, Cambridge, MA.
- Jakobi, N., Husbands, P., and Harvey, I. (1995). Noise and the reality gap: The use of simulation in evolutionary robotics. In *European Conference on Artificial Life*, pages 704–720. Springer.
- Javed, K. and Sutton, R. S. (2024). The big world hypothesis and its ramifications for artificial intelligence. In *Finding the Frame: An RLC Workshop for Examining Conceptual Frameworks*.
- Jonas, H. (1966). *The Phenomenon of Life: Toward a Philosophical Biology*. Harper & Row, New York.
- Kaelbling, L. P., Littman, M. L., and Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134.
- Khetarpal, K., Ahmed, Z., Comanici, G., Abel, D., and Precup, D. (2020). What can i do here? a theory of affordances in reinforcement learning. In *International Conference on Machine Learning*, pages 5243–5253. PMLR.
- Khetarpal, K., Riemer, M., Rish, I., and Precup, D. (2022). Towards continual reinforcement learning: A review and perspectives. *Journal of Artificial Intelligence Research*, 75:1401–1476.
- Koffka, K. (1935). *Principles of Gestalt Psychology*. Harcourt, Brace.
- Lave, J. (1988). *Cognition in Practice: Mind, Mathematics and Culture in Everyday Life*. Cambridge University Press, Cambridge.
- Lederman, S. J. and Klatzky, R. L. (1987). Hand movements: A window into haptic object recognition. *Cognitive Psychology*, 19(3):342–368.
- Littman, M. L., Sutton, R. S., and Singh, S. (2002). Predictive representations of state. In *Advances in Neural Information Processing Systems*, volume 14, pages 1555–1561.
- Machado, M. C., Barreto, A., Precup, D., and Bowling, M. (2023). Temporal abstraction in reinforcement learning with the successor representation. *Journal of machine learning research*, 24(80):1–69.
- Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W. H. Freeman and Company, San Francisco.

-
- McCloskey, M. and Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier.
- Merleau-Ponty, M. (1945). *Phénoménologie de la perception*.
- Mitchell, T., Cohen, W., Hruschka, E., Talukdar, P., Yang, B., Betteridge, J., Carlson, A., Dalvi, B., Gardner, M., Kisiel, B., et al. (2018). Never-ending learning. *Communications of the ACM*, 61(5):103–115.
- Newell, A. and Simon, H. (1956). The logic theory machine—a complex information processing system. *IRE Transactions on information theory*, 2(3):61–79.
- Noë, A. (2004). *Action in Perception*. MIT Press, Cambridge, MA.
- O’Regan, J. K. and Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24(5):939–973.
- Oudeyer, P.-Y., Kaplan, F., and Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2):265–286.
- Pfeifer, R. and Bongard, J. C. (2006). *Understanding intelligence*. MIT press.
- Rao, R. P. N. and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87.
- Ring, M. B. (1997). Child: A first step towards continual learning. *Machine Learning*, 28(1):77–104.
- Rus, D. and Tolley, M. T. (2015). Design, fabrication and control of soft robots. *Nature*, 521(7553):467–475.
- Silver, D. and Sutton, R. S. (2025). Welcome to the era of experience. *Google AI*, 1.
- Şimşek, Ö. and Barto, A. G. (2006). An intrinsic reward mechanism for efficient exploration. In *Proceedings of the 23rd international conference on Machine learning*, pages 833–840.
- Singh, S., Barto, A., and Chentanez, N. (2004). Intrinsically motivated reinforcement learning. *Advances in neural information processing systems*, 17.
- Sutton, R. S., Modayil, J., Delp, M., Degris, T., Pilarski, P. M., White, A., and Precup, D. (2011). Horde: A scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 761–768.
- Sutton, R. S., Rafols, E., Lehnert, L., Modayil, J., and Barreto, A. (2022). Reward-respecting subtasks for model-based reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Thompson, E. (2007). *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press, Cambridge, MA.
- Thrun, S. (1998). Lifelong learning algorithms. In *Learning to learn*, pages 181–209. Springer.
- Tishby, N. and Polani, D. (2010). Information theory of decisions and actions. In *Perception-action cycle: Models, architectures, and hardware*, pages 601–636. Springer.
- van Dijk, L. and Withagen, R. (2016). On the notion of affordance and the importance of environment–organism mutuality. *Ecological Psychology*, 28(3):174–194.
- Varela, F. J., Thompson, E., and Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press, Cambridge, MA.

-
- Whitehead, S. D. and Ballard, D. H. (1990). Active perception and reinforcement learning. In *Machine learning proceedings 1990*, pages 179–188. Elsevier.
- Zahedi, K. and Ay, N. (2013). Quantifying morphological computation. *Entropy*, 15(5):1887–1915.