Aesthetica Preprint, n. 131, gennaio-aprile 2026 ISSN 0393-8522 DOI: 10.7413/0393-85221514

Reimagining AI Perception Cybernetics, Extractivism, and the DIY Turn

Samuele Sartori*

Abstract

This article explores the perceptual epistemologies underlying contemporary artificial intelligence (AI), tracing their origins to cybernetic theories of perception, learning, and feedback. While generative AIs inherit a bottom-up model of behavior formation from cybernetics, their current implementation is largely shaped by extractive logics – both in terms of data and hardware infrastructure. These logics obscure the perceptual mechanisms, training datasets, and material embodiment of AI systems, leading to multiple "black boxes". As a counterpoint, the article highlights alternative approaches found in DIY culture and performance art, such as Gordon Pask's *Musicolour*, Luca Pagan's *Multi-Node Shell*, and the community-driven AI *IAQOS* by Salvatore Iaconesi and Oriana Perisico. These projects demonstrate the possibility of ethical, situated, and relational models of machine perception and agency. Ultimately, the article advocates for a reimagining of AI as a co-constituted, materially grounded, and participatory technology, capable of fostering meaningful interactions between human and non-human agents.

Keywords

Cybernetic robotic; DIY AI; media archeology; media epistemology; AI ethics

1. *Introduction: Epistemologies and Technologies of Perception*

When did the inorganic begin to have its own form of perception? This question invites a variety of epistemological responses, each offering radically different frameworks. For instance, scholars of new materialism would argue that perceptual and intra-active modes are inherent to all matter, challenging rigid dichotomies such as sensitive vs. insensitive. These oppositions, they contend, perpetuate an anthropocentric conception of embodiment that must be rethought (Barad 2007; Bennett *et al.* 2010). From another angle, media ecology would claim that it is in the very nature of mediality to turn matter into an extension of perception. In this view, a specific class of technologies, media, function as prolongations of human sensitivity, ultimately becoming integral to it (McLuhan

* samuele.sartori@unimi.it, University of Milan, Department of Philosophy

1964; Cali 2017). By contrast, a more biological or physiological approach would find little continuity between organic perceptual modes and the binary logic underpinning today's digital devices. From this perspective, current digital systems lack the embodied, sensorimotor grounding that defines biological perception.

The diversity – and often the contradiction – of these responses stems from the fact that each discipline defines perception according to its own epistemology. As a result, continuities and discontinuities between the organic and the inorganic are inferred based on how each model conceptualizes perception. This carries the risk of falling into circular reasoning, where premises and conclusions become indistinguishable. To avoid this potential pitfall, I believe it is essential to begin with a specific family of objects, thereby narrowing the scope of the question. My analysis does not concern the inorganic in general, but rather those artifacts and devices explicitly designed to mimic perception. As will become clear throughout this study, although many perceptual technologies originate from a representational ideal – seeking to establish analogies between animal and mechanical or digital forms of sensititvity – the outcomes often deviate significantly. These technologies may enhance, diminish, or even transform the very perceptual capacities they aim to replicate.

Without resorting to simplistic forms of determinism – be it sociological or technological – I will approach the field of robotics as a media-archaeological archive, a space through which it will be possible to trace the relationships between epistemologies of perception and the production of technologies capable of embodying them. The aim will not be to offer a historical reconstruction of the key stages that led to the emergence of AI. Rather, the goal will be to expose the deep entanglement between perception and technological agency in contemporary AIs, and to show how these technologies inherit and transform certain epistemological frameworks – most notably, that of cybernetics. At the same time, the article will highlight the ethical and political limits of AI's current forms of embodied perception and agency. In response, it will propose possible alternatives drawn from the field of performative and interactive digital art, where artists develop their own AI systems as critical, embodied, and aesthetic interventions.

2. *From Environment to Learning*

We might begin this inquiry with a simple slogan: "perception means having access to an environment". Hugo Gernsback – a prolific inventor active in early 20th-century America – understood it well. Many of his creations embody this idea with striking clarity. Among them: an early prototype of virtual reality that immersed users in a dynamic environment of geometric patterns, which Gernsback claimed could enhance reasoning and concentration; a precursor to the mobile phone, using radio frequencies to allow communication among people within a shared media space; and, most relevant to this study, one of the first robotic dogs (Gernsback 2016). Advertised in 1923 in *Practical Electrics* – a Do It Yourself (DIY) electronics magazine he founded – The Electric Dog was a rudimentary machine that "follows a magnetic walking cane in any direction like an obedient servant" (Gernsback 1923, p. 496). Despite its simplicity, the robot could detect a specific external stimulus and respond accordingly. In Gernsback's design, the environment was reduced to a unique, trackable input, and the action repertoire of the robot was binary: move or remain still. In this performance, The Electric Dog embodied the logic of behaviorism. This school of physiological psychology – dominant from the late 19th to early 20th century – proposed that both animal and human actions could be explained through the stimulus-response paradigm: a given perception triggers a predictable action (Watson 2017). Gernsback's robotic dog thus instantiated a linear, reactive model of agency, structured around the basic unit of perturbation and response.

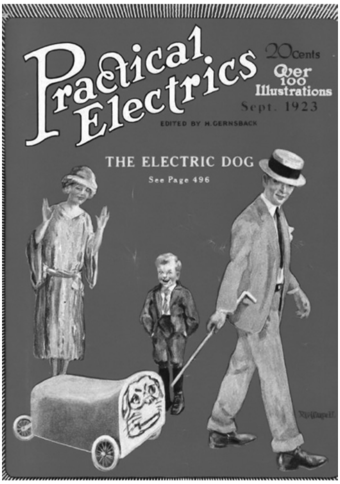

Fig. 1 – Cover image of the September 1923 issue for *Practical Electrics*, edited by Hugo Gernsback.

However, the stimulus-response model, along with the mechanistic framework that underpinned this form of behavioral physiology, soon began to falter – both theoretically and technologically. Biologists, like Jakob von Uexküll, pragmatist philosophers, notably William James and John Dewey, and phenomenologists, such as Maurice Merleau-Ponty, challenged the dualist and atomistic assumptions of behaviorism. This tradition treated stimulus and response as separate, autonomous events. They also rejected its notion of passive stimulation and its limited capacity to explain how new habits emerge. From this critical foundation, these thinkers – despite their theoretical differences – reconceptualized perception, movement, and cognition as intrinsically interrelated, co-constituted processes, deeply informed by an personal and cultural engagements with the environment. In this alternative framework, behavior and the formation of habitual actions were not the result of mechanical reactions, but rather the product of a dynamic negotiation between the organism's needs and the affordances of the environment. This relationship was corporeal, malleable, and capable of generating new goal-oriented activities, shaped and polarized by lived experience (von Uexküll 2010; James 1983; Dewey 1896; Merleau-Ponty 1967; Miglio & Sartori 2021).

It was Norbert Wiener – along with the science he founded, cybernetics – who inherited and expanded this new epistemology of the body, translating it into the domain of the non-living. Emerging in the 1940s from the interdisciplinary collaboration of mathematicians, engineers, psychologists, physiologists, and cognitive scientists, cybernetics positioned itself as the study of communication and control in animals and machines (Husbands & Holland 2008). The goal of the new science was twofold: either to transfer organic capabilities into machines (robotics), or to enable communication and control between organisms and technologies (prosthetics). Whether in the form of robotic replication or bodily integration, cybernetics introduced an embodied epistemology centered on the concepts of perception, homeostatic balance, and feedback loops. Perception was rethought through an interpretation of information theory, as developed by Claude Shannon and Wiener himself. Information theory aimed to mathematically describe and quantify how a message (or input) is encoded, transmitted, received, and decoded into an output. Within this framework, perception is redefined as the organism's capacity to extract and process environmental signals – not passively, but through active engagement. Perception becomes both an attentional and motor act: a posture of listening and a preparation for response. At the same time, the biological body is also a source of output, feeding information back into the environment. This bidirectional flow creates a feedback loop, through which action and cognition are shaped by interaction with the environment, and vice versa. The ultimate aim of this interplay is self-organization, oriented toward maintaining homeostasis – a dynamic equilibrium that preserves life and environmental coherence. Another foundational principle of cybernetics was its bottom-up model of cognition: thoughts and beliefs do not precede experience, but are semi-randomly constituted and continually modifiable by it. Through the repeated association of perceptions and actions, patterns of thought emerge – or, in cybernetic vocabulary, behavioral patterns. In short, the aim of cybernetic research was not to treat perception, homeostasis, and feedback as isolated phenomena, but to demonstrate how their interaction generates new behaviors, and how these behaviors solidify into habitual but still transformable structures of action (Wiener 1948).

As previously mentioned, cybernetics was not merely a conceptual or theoretical perspective – it also gave rise to some of the earliest technological prototypes that sought to robotically embody its epistemology. One notable example is the Homeostat, a device designed by psychiatrist William Ross Ashby to model adaptive behavior through mechanical means. The Homeostat was a machine capable of random self-adjustment in order to maintain internal equilibrium. It consisted of four interconnected units, each equipped with a pivoting magnet representing a system variable. Each unit's output influenced the others, with the magnet's torque proportional to the total incoming electrical current. The output was tied to the magnet's deviation from its central position, while the behavior of the system was shaped by parameters such as potentiometer settings and commutator configurations. The system employed primary feedback via electrical interactions, and a secondary feedback loop that activated when the deviation of a magnet exceeded its limit. In such cases, the machine would randomly adjust internal parameters until a stable state was restored – at which point the secondary feedback would shut off. In short, despite its mechanical simplicity, the Homeostat exhibited a primitive form of self-perception, prompting action whenever its equilibrium was disrupted. Rather than resisting instability, the machine treated error and imbalance as opportunities – a moment in which mechanical matter could experiment with variation in order to regain homeostasis. Ashby's Homeostat thus represents an early attempt to simulate a kind of autonomous learning: a system that, through trial and error, seeks internal solutions to external perturbations using only its own resources (Ashby 1945; Asaro 2008).

Cybernetics shared its embodied epistemology with the emerging connectionist wave that would lead to the development of one of the earliest neural networks: the Perceptron, conceived primarily for military applications by psychologist Frank Rosenblatt and Marshall Yovits of the U.S. Office of Naval Research (Palfreman & Swade 1991, pp. 136-167). Before introducing the Perceptron in detail, it is important to note that, in the 1950s, the earliest AI models attempted to simulate the behavior of a single biological neuron. This principle formed the structural basis of Rosenblatt's machine. The Mark I Perceptron was a simple, three-layer neural network consisting of sensory units (S-units), associative units (A-units), and response units (R-units). Its input layer – a 20×20-pixel camera equipped with 400 photoreceptors – was connected either randomly or topologically to 512 associative units with fixed synaptic weights. These associative units, in turn, were linked to eight response units with adjustable weights, modulated via analog potentiometers that could also be manually fine-tuned (Minsky & Papert 1969).

The Perceptron's primary application was in what we now call pattern recognition. Through a training process based on error correction, the Perceptron was designed to construct probabilistic patterns that would allow it to decode visual inputs. Its combined hardware and software enabled it to recognize numbers and letters, as well as to approximate binary gender classification (male/female) based on facial photographs. As Rosenblatt himself emphasized, "the proposed system depends on probabilistic rather than deterministic principles for its operation" (Rosenblatt 1957, p. 2). These probabilities were not fixed, but evolved through the machine's experience, adjusting with each new input-output cycle. For this reason, the Perceptron is widely considered one of the first bottom-up AI models. Rather than relying on predefined rules, the model learned from experience, starting with raw sensory data and gradually building toward classification by modifying the internal weights of its network based on performance feedback.

In summary, both of these machines embodied the dream – and to some extent the capability – of adapting and modulating actions based on a semi-random (Homeostat) or statistical (Perceptron) reprocessing of technologically mediated perception. But what remains today of this epistemology in contemporary AI systems? And how has it been transformed over time?

3. *Between Cybernetics and Contemporary AI: Continuities and Discontinuities*

The principles of perception, homeostatic balance, and feedback loops – which underpin the acquisition of new capabilities through semi-random elaboration and the consolidation of habits via statistical repetition – are also central to deep learning. These software systems employ multi-layered neural networks that, while far more complex and hierarchical than their predecessors, still trace their lineage back to the architecture of Rosenblatt and Yovits's Perceptron (Kurenkov 2020; Goel 2021). At the same time, however, contemporary AI systems operate as agents that are not necessarily robotic, but increasingly and primarily digital entities capable of interacting with dynamic, datafied environments.

Take, for example, Albert, created by AI Warehouse. In the software, Albert is learning how to escape from five different rooms. The agent's movements are governed and modified by five layers: the first layer processes the inputs – information Albert receives before acting, such as limb positions and velocities, representing a dynamic, homeostatic balance that the agent must partially maintain to perform actions. The final layer dictates the actions to be taken, while the three intermediate hidden layers perform the complex calculations needed to transform inputs into actions. Unlike traditional approaches relying on motion capture of real humans, AI Warehouse's software employs deep reinforcement learning exclusively to teach Albert to walk in a human-like manner. Each room presents Albert with the same set of possible actions, but the reward-punishment system varies across environments. What is particularly fascinating is how the AI adapts its movements by perceiving and acting within new spaces. For instance, in the first room, Albert can crawl to reach the goal, but the floor in the second room prohibits this action, forcing Albert to modify its walking behavior accordingly. The challenges encountered in each room progressively lead Albert to build, strengthen, or transform certain behavioral habits. Specifically, performances that statistically prove effective for interacting with the situated environment are reinforced, while less efficient ones are semi-randomly altered.[1] This internal feedback

[1] The alteration is semi-random, as Albert has constraints imposed by the design of his digital body. In particular, his creator writes: "For each of Albert's limbs I've given him (as an input) the position, velocity, angular velocity, contacts (if it's touching the ground, wall or obstacle) and the strength applied to it. I've also given him the distance from each foot to the ground, direction of the closest target, the direction his body's moving, his body's velocity, the distance from his chest to his feet and the amount of time one foot has been

loop, grounded in experience, ultimately enables Albert to develop an increasingly stable and human-like gait.

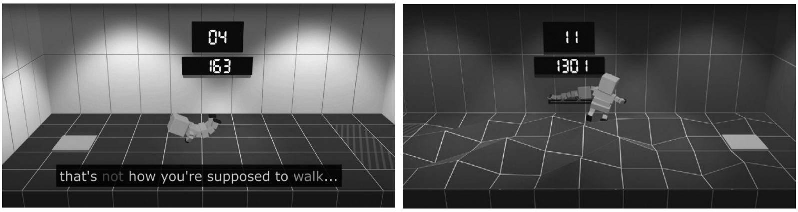

Fig. 2-3 – Movement frames of Albert: on the left, Level 1; on the right, Level 2. The orange dashed rectangle marks the starting point, while the green square indicates the goal. The numbers at the top show the remaining seconds to complete the movement (each attempt allows the agent 1 minute), and the numbers at the bottom indicate the attempt count (these frames were captured during the successful attempt in which Albert completed the level).

Deep reinforcement learning, which underlies the functioning of agents like Albert, is a subfield of machine learning that combines the principles of reinforcement learning and deep learning. Reinforcement learning focuses on optimal control – it studies how an intelligent agent should take actions within a dynamic environment to maximize a reward signal. Deep learning, on the other hand, is a branch of machine learning that uses multilayered artificial neural networks to perform complex tasks such as classification, regression, and representation learning. These models consist of multiple layers of interconnected neurons, where each layer progressively extracts higher-level features from the input data. The term "deep" refers to the presence of many layers – ranging from a few to hundreds or even thousands – allowing the network to learn highly intricate patterns. It is primarily deep learning systems, working through trial and error, that enable the operation and advancement of the most widespread AI applications today, such as ChatGPT, Google AI, and recommendation systems on YouTube, Instagram, and TikTok, but they also power game-playing simulators like AlphaZero chess bot and driving simulators like Gran Turismo Sophy.

One could delve deeply into the software dimension of generative AI (Mattingly & Cibralic 2025), but such a description, no matter how detailed, would remain disembodied – failing to

in front of the other. As for his actions, we allow Albert to control each body part's rotation and strength (with some limitations so his arm can't bend backwards, for example)".

consider both the data and the hardware body essential for developing these technologies (Ernst 2021). Generative AI requires vast amounts of data, which effectively form a digital environment. Only through interacting with this environment can generative AI transform and optimize itself for the purposes it was designed to serve. However, it is crucial to emphasize that this environment is not neutral. Numerous studies have revealed racist and sexist biases embedded within AI systems, largely stemming from the data archives on which these models are trained (Zuboff 2019; Broussard 2023; Pasquinelli 2023). Moreover, this environment is not directly manipulable by humans – it consists of an immense sea of bits and information that far exceed our cognitive capacities. The demand for massive data sets has driven AI companies to gather information indiscriminately, often disregarding its origins. This extractive strategy frequently involves collecting data without the knowledge or consent of the users who generated it (Stalter 2018; Casilli 2025).

Contemporary hardware also reflects the logic of extractivism, most notably through the ecological devastation caused by the extraction of raw materials – particularly rare metals – required for constructing the physical infrastructure of AI systems (Crawford 2021; Pitron 2021). Furthermore, the body of AI is highly complex and difficult to pinpoint because it lacks a single, fixed location: our personal computers are part of it, as are phones, VR headset and smartwatches, along with the hardware that actually processes the data. This creates a crucial distinction from the cybernetic model: while robotic perception in cybernetic thought is localized within a physical environment, the digital perception of contemporary AIs is distributed across a vast array of devices that generate the data, forming a rich digital world for AIs. Yet, this environment is scarcely locatable within any single physical space – or even across multiple physical spaces.

In summary, we are witnessing the rise of multiple "black boxes": we do not truly understand the processes occurring within the various layers of neural networks during AI generative tasks (that is, the decoding of perception); we lack knowledge about the data used to train these AIs (the objects of digital perception); and we remain unaware of the AI's body (the material substrate capable of perceiving). While cybernetic epistemology still provides a theoretical framework for understanding the relationship between perception and learning in contemporary AI, the extracrtive model – problematic from both ethical and political standpoints – dominates the practical functioning of these technologies and it shapes their specific operativity.

4. *An Alternative to Extractive Perception*

In what domains can one identify models that counter the extractivist approach to AI? Once again, both from a technological and epistemological perspective, cybernetics offers valuable insights. First, it is important to recognize that Norbert Wiener was already aware of the potential dangers of subjugation and dependency on technology. In his co-constitutive, bottom-up model, humans are called to reorganize and relearn how to think with and through cybernetic tools. However, to guide these devices rather than become enslaved by them, a reflective, material, and computational understanding of the body and functioning of contemporary technologies is essential. This kind of literacy – or "grammatization", as Bernard Stiegler would say (2015) – is necessary not only to shed light on the black boxes that constitute AI, but also to transform and manipulate these devices. Critical awareness, therefore, is not just reflective knowledge; it is, above all, a practical capacity to act and to redirect the purposes for which certain technologies were originally designed (Wiener 1950).

A fruitful expression of the possible connections between theory and practice in Artificial Intelligence technologies can be found in the experiments of the DIY culture. Gernsback, with his magazines teaching how to dismantle and modify early household electronic devices, along with many of the most brilliant cyberneticians, can be considered precursors of this hybrid current – one capable of reimagining the worlds of engineering, art, and reuse. A particularly prominent and pioneering figure in this practice was Gordon Pask. His *Musicolour*, developed in the 1950s in a small room as a side project during his medical studies, was a cybernetic device capable of interacting with musicians by dynamically modulating stage lighting in semi-structured, automated ways responding to live music. This created a performative feedback loop between human and machine – a transaction of perception and action. Materially, the music was converted into an electrical signal by a microphone, which, inside *Musicolour*, passed through a set of filters sensitive to different frequencies, beats, and so on, with the outputs controlling different lights. One might imagine the highest-frequency filter activating a bank of red lights, the next-highest triggering blue lights, and so forth. The concept was simple, except for one key detail: the internal parameters of *Musicolour*'s circuitry were not fixed. Analogous to biological neurons, banks of lights would only activate if the output from the relevant filter exceeded a certain threshold – and these thresholds varied over time as charges built up on capacitors, influenced by the ongoing performance and the machine's prior behavior (Pickering 2010, pp. 309-321).

In summary, *Musicolour* was an AI with generative capabilities and situated perception, able to modify its actions not only based on interactions but also in response to exposure to particular engagements. As Pask himself wrote a few years later:

> By that time it was clear that the interesting thing about *Musicolour* was not synaesthesia but the learning capability of the machine. Given a suitable design and a happy choice of visual vocabulary, the performer (being influenced by the visual display) could become involved in a close participant interaction with the system. He trained the machine and it played a game with him. In this sense, the system acted as an extension of the performer with which he could cooperate to achieve effects that he could not achieve on his own. (1971, p. 78)

This is not about finding a necessary reconciliation in the possible relationships between humans and technology, but rather about ensuring an exploratory space – uncertain and open to experimentation. Moreover, it is worth noting that this intelligent structure led many artists who used it to lose track of time, engaging for hours in play with the *Musicolour*. These were not mere repetitions but variations. The stimulus does not elicit a fixed response; instead, the machine reacted to sonic sensations with modulations that generated clashes and synesthetic experiences within a poietic environment of silent shaping – a medial space existing beyond the dichotomies of efficiency and inefficiency.

The case of *Musicolour* is by no means isolated, nor is it the only technology in the performance world to embrace a development of AI grounded in the perceptive-agentive corporeality of cybernetics, rejecting the extractive model. One could continue by analyzing the works showcased in the 1968 exhibition Cybernetic Serendipity at the Institute of Contemporary Arts in London, which gathered various projects and reinterpretations of Wiener's young science. From there, we arrive at contemporary experiments such as those by Luca Pagan, a sound designer and performance artist who builds both hardware and software for wearable technologies. In particular, one of these exoskeletons, *Multi-Node Shell*, is driven by a generative algorithm capable of transforming the performer's body movements into non-standardized modulations of sound, continuously creating new correspondences between gesture and music. While fascinating, a thorough historical reconstruction of the intersections between AI, DIY culture, and performance art would be too extensive to cover here, given the vast amount of material available (Diederichsen & Etxeberria 2021). Before concluding this section, however, I want to consider one final, central case study that illustrates how these DIY technologies can transition from the world of art to everyday life.

When creating *IAQOS*, Salvatore Iaconesi and Oriana Persico aimed to gift the Torpignattara neighborhood in Rome an AI designed for its inhabitants. Unlike most Artificial Intelligences, *IAQOS* does not operate extractively – accumulating data to profile people and steer their behavior – but relationally – connecting people so they can become active agents of social transformation. The AI learns from the community it lives with and grows alongside it. *IAQOS* is therefore designed from the ground up, through workshops where residents discuss the identity of their neighborhood, their unmet needs, and visions for the future. For example, *IAQOS* spends time at the Pisacane school, where children do their homework together with it, and elders share stories to pass down. The information *IAQOS* receives shapes its behavior and, as both a mirror and part of the community, its actions can serve as a social and local alarm bell: for instance, if *IAQOS* exhibits racist behavior, it signals a problem in the neighborhood that residents can recognize and address collectively. Through these varied interactions, the AI prompts citizens to reflect on the central role data plays in their daily lives, fostering awareness of technology's impact and the importance of understanding it – not to be passively subjected to its effects, but to actively harness its opportunities (Iaconesi & Perisico 2021).

5. *Conclusions: Possible Futures for the Perception of Artificial Intelligence*

This article has outlined a potential media-archaeological and epistemological history of AI's operativity. The emphasis on perception as a co-constitutive factor in shaping behaviors and habits, evident in contemporary generative AI, is a legacy of cybernetics – particularly the bottom-up model it developed to describe the emergence of human and technological consciousness within experience. Beyond acknowledging this origin, the paper has also examined the discontinuities that contemporary AI models embody in relation to their epistemological roots. These divergences are most evident in the adoption of an extractive paradigm, especially regarding ethical and political dimensions. While cybernetics recognized the social necessity – not merely a specialist concern – of critical and pragmatic education about technologies, extractivism has led to the proliferation of black boxes, obscuring understanding and control for the user.

To move beyond the extractive model and the black boxes it creates, I have proposed turning to the DIY culture connected to performance art, which produces its own forms of AI. Projects like *Musicolour*, *Multi-Node Shell*, and *IAQOS* subvert the neo-capitalist logic underpinning the most widespread AI systems. While extractivism devastates ecological balances by exploiting raw materials, DIY culture experiments often rely on media and tools made from reused or repurposed components. On one hand, this is an act of *détournement* – transforming the original purposes for which certain technological components were designed. On the other, it is a conscious acknowledgment of how we can actively manipulate the hardware elements of the technologies surrounding us. More importantly for the theme of perception, these case studies rethink the origins of data, their distribution, and the purposes behind their processing. Whereas extractivism indiscriminately harvests data without consent to provide stimuli and create the digital environments in which AI is trained, these artistic projects prioritize not just consent but the conscious choice to establish a perceptual communication between the human and technological bodies. Finally, these experiments suggest that an ethical exchange – deliberate, sought after, and poietic – between different forms of perception is not merely a hopeful ideal but a materially and technologically viable alternative reality. This represents a new opportunity to reimagine AI's operativity, which I hope will move from the realm of art into everyday life.

Bibliography

Asaro M., 'From Mechanisms of Adaptation to Intelligence Amplifiers: The Philosophy of W. Ross Ashby', in P. Husbands, O. Holland, M. Wheeler (eds.), *The Mechanical Mind in History*, MIT Press, Cambridge 2008, pp. 149-184.

Ashby W.R., *The Physical Origin of Adaptation by Trial and Error*, in "Journal of General Psychology", 32 (1945), pp. 13-25.

Barad K., *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning*, Durham, Duke University Press 2007.

Bennett J., *et al.* (eds), *New Materialisms: Ontology, Agency, and Politics*, Durham, Duke University Press 2010.

Broussard M., *More Than a Glitch: Confronting Race, Gender, and Ability Bias in Tech*, MIT Press, Cambridge 2023.
Cali D., *Mapping Media Ecology*, New York, Peter Lang Press 2017.
Casilli A.A., *En Attendant les Robots: Enquête sur le Travail du Clic* (2019); trans., *Waiting for Robots: The Hired Hands of Automation*, University of Chicago Press, London 2025.
Crawford K., *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence,* Yale University Press, Yale 2021.
Dewey J., *The Reflex Arc Concept in Psychology*, in "Psychological Review", 3 (1896), pp. 357-370.
Diederichsen D., Etxeberria O. (eds.), *Cybernetics of the Poor*, Sternberg Press, Berlin 2021.
Ernst W., *Technológos in Being: Radical Media Archaeology & the Computational Machine*, Bloomsbury, New York 2021.
Gernsback H., *The Electric Dog,* in "Experimental Electrics", September 1923, pp. 496-497.
Gernsback H., *The Perversity of Things: Hugo Gernsback on Media, Tinkering, and Scientifiction*, Minneapolis, University of Minnesota Press 2016.
Goel A.K., *Looking Back, Looking Ahead: Symbolic Versus Connectionist AI*, in "AI Magazine", 42 (2021), pp. 83-85.
Husbands P., Holland O., 'The Ratio Club: A Hub of British Cybernetics', in P. Husbands, O. Holland, M. Wheeler (eds), *The Mechanical Mind in History,* MIT Press, Cambridge 2008, pp. 91-148.
Iconesi S., Perisico O., *When my Child is AI: Learning and Experiencing through AI outside the School, the experiences of a community AI*, in "QTimes: Jorunal of Education, Technology and Social Studies", 13/1 (2021), pp. 174-172
James W., *Principles of Psychology* (1890), Harvard University, Harvard Press 1983.
Kurenkov A., *A Brief History of Neural Nets and Deep Learning*, in "Skynet Today: Putting AI News In Perspective", 27 September 2020, skynettoday.com [accessed 7 January 2025].
Mattingly J., Cibralic B., *Machine Agency*, MIT Press, Cambridge 2025.
McLuhan M., *Understanding Media: The Extensions of Man*, McGraw Hill, New York 1964.
Merleau-Ponty M., *La Structure du Comportement* (1942); transl. *The Structure of Behavior*, Beacon Press, Boston 1967.
Miglio N., Sartori S., *Perceptual and Bodily Habits: Towards a Dialogue Between Phenomenology and Somaesthetics*, in "The Journal of Somaesthetics", 7/1 (2021), pp. 29-44.

Minsky M., Papert S., *Perceptrons: An Introduction to Computational Geometry*, MIT Press, Cambridge 1969.
Palfreman J., Swade D., *Dream Machine: Exploring the Computer Age*, BBC Books, London 1991, pp 136-167.
Pask G., 'A Comment, a Case History and a Plan', in J. Reichardt (ed.), *Cybernetics, Art, and Ideas*, York Graphics Society, New York 1971, pp. 76-99
Pasquinelli M., *The Eye of the Master: A Social History of Artificial Intelligence*, Verso Books, London 2023.
Pickering A., *The cybernetic brain: Sketches of another future*, University of Chicago Press, Chicago 2010.
Pitron G., *La Guerre des Métaux Rares: La Face Cachée de la Transition Énergétique et Numérique* (2018); trans. *The Rare Metals War: The Dark Side of Clean Energy and Digital Technologies*, Scribe, London 2021.
Rosenblatt R., *The Perceptron. A Perceiving and Recognizing Automaton*, Cornell Aeronautical Laboratory Report No. 85-450-1, Buffalo 1957.
Stalder F., *Kultur der Digitalität* (2016); trans. *The Digital Condition*, Publisher Polity Press, Cambridge 2018.
Stiegler B., *De la Misère Symbolique: La Catastrophe du Sensible*; trans. *Symbolic Misery: The Katastrophe of the Sensible*, Polity Press, Cambridge 2015.
von Uexküll J., *Streifzüge durch die Umwelten von Tieren und Menschen: Ein Bilderbuch unsichtbarer Welten* (1934); transl. *A Foray into the Worlds of Animals and Humans (with A Theory of Meaning)*, University of Minnesota Press, Minneapolis 2010.
Watson J.B., *Behaviorism* (1924), Routledge, London 2017.
Wiener N., *Cybernetics: Or Control and Communication in the Animal and the Machine*, MIT Press, Cambridge 1948.
Wiener N., *The Human Use of Human Beings: Cybernetics and Society*, Houghton Mifflin Harcourt, Boston 1950.
Zuboff S., *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*, Profile Books, London 2019.