Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Can AI Systems Matter to Microbes? Inverting Turing's Test for Non-Human AI Evaluation

Emanuel Gollob
University of Arts and Industrial Design Linz
Klaus Spiess
klaus.spiess@meduniwien.ac.at
Medical University of Vienna

Research Article

Keywords: transsemiotic AI, biosemiotics, more-than-human AI evaluation, artistic research, non-human meaning-making

Posted Date: June 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9953066/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Can AI Systems Matter to Microbes? Inverting Turing's Test for Non-Human AI Evaluation

Author One[1] and Author Two[1*]

[1*][Affiliation removed for blind review].

*Corresponding author(s). E-mail(s): author.two@anonymous.org;
Contributing authors: author.one@anonymous.org;

Abstract

AI systems operate within ecologies where their effects unfold among non-human organisms, at timescales and through semiotic processes humans cannot access. Yet evaluation paradigms presuppose an anthropocentric horizon that renders these effects invisible by construction. This paper develops the Transsemiotic Turing Approach to investigate what becomes visible when that horizon is provisionally suspended, while remaining accountable to the paradox that its instruments are themselves anthropogenic.

The framework inverts Turing's test, asking not whether machines register as intelligent for human judges but whether AI outputs acquire significance for non-human interpreters. This motivates the concept of transsemiotic exchange as a coupling between entities whose sign systems share no common code. Distinguishing normative selectivity (grounded in biological imperatives) from parametric selectivity (oriented by a reward function) keeps the criteria from being trivially satisfied on either side. The distinction anchors four diagnostic conditions, of which transduction is the material precondition, followed by selectivity, appropriation, and persistence.

Three artistic-research installations (2022–2024) materialise the framework, coupling microbial communities, human participants, and reinforcement-learning agents that generate synthetic voice articulations conditioned by pH and optical density measurements. The contribution is conceptual and practice-led, the diagnostic criteria derive from biosemiotic theory and the generative principles from arts-based research. Rather than verifying significance directly, the approach diagnoses conditions under which evidence for transsemiotic exchange becomes producible. Read as practice-led observations, the installations motivate the hypotheses that AI functions as temporal infrastructure mediating between computational milliseconds and biological days, and that the richest transsemiotic configurations resist diagnostic clarity.

Keywords: transsemiotic AI, biosemiotics, more-than-human AI evaluation, artistic research, non-human meaning-making

1 Introduction

A reinforcement learning (RL) agent generates a phoneme through a synthetic voice articulator. In a petri dish, an oral microbial community reorganises its cell-skeleton. No human evaluator witnessed this exchange, and no existing benchmark would register it. This paper asks what becomes visible when such encounters are approached as a diagnostic challenge to the principles underlying AI evaluation.

AI evaluation has become a vast and sophisticated enterprise. From MMLU (Massive Multitask Language Understanding) and BIG-bench (Beyond the Imitation Game benchmark) (Srivastava et al. 2023), to Holistic Evaluation of Language Models (HELM) (Liang et al. 2023), alignment-focused evaluations and more, the expanding repertoire of AI assessment spans linguistic comprehension, reasoning, task performance, and safety. Fairness audits interrogate distributional biases, red-teaming exercises probe failure modes, and capability benchmarks track performance

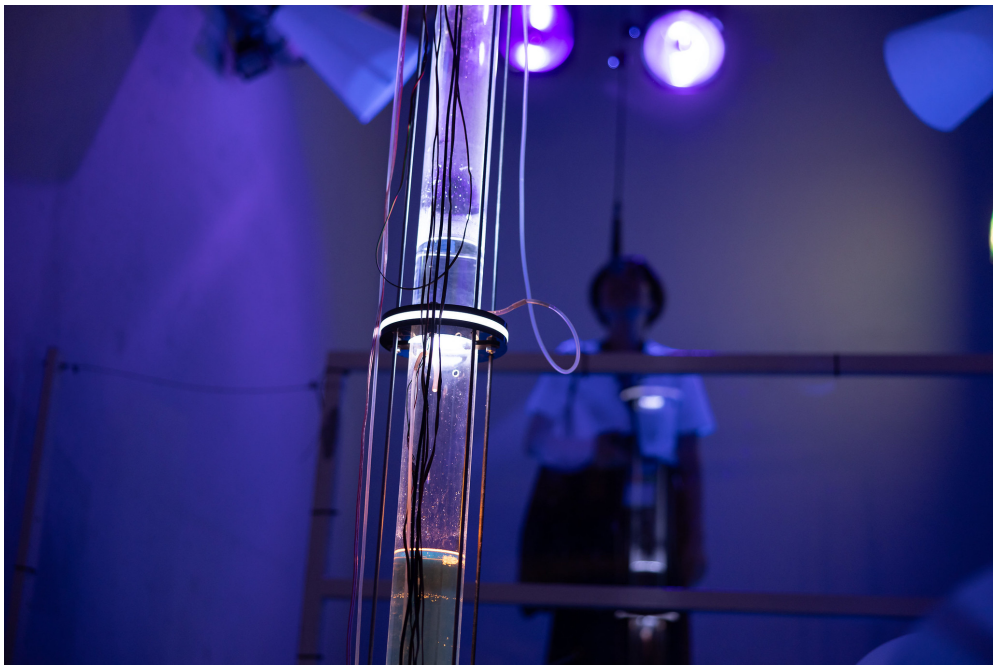


Fig. 1 One of the paper's three installations (Installation B, discussed in detail in Section 4.2), involves RL agents generating phonetic tokens for the in vitro cultivated oral microbiota visible in the foreground. This setup creates a feedback loop in which microbial vitality validates the computational output. Photograph by a programme collaborator, full credit will be provided upon acceptance.

against human baselines across hundreds of tasks. Drawing on critiques of data extraction (Crawford 2021) and semantic illusion (Bender et al. 2021), we can identify three shared structural features: reliance on human-generated ground truth, decomposition of intelligence into human-legible subtasks, and evaluation through human-interpretable scoring. This convergence is not a design flaw but an ontological commitment. Recent critiques from within AI research have begun to interrogate it. Raji et al. argue that benchmarks function as social technologies encoding specific value commitments (Raji et al. 2021). Schlangen questions whether evaluation discourse conflates measurement with understanding (Schlangen 2023). Pasquinelli situates such critiques in a longer social history of artificial intelligence, in which computational systems sediment the labour, disciplinary categories, and measurement regimes of their conditions of production (Pasquinelli 2023). These critiques share our concern but remain epistemological. They interrogate how evaluation is conducted, but not whose significance it can register. The assumption that human symbolic systems exhaust the domain of the evaluatively relevant remains intact. The question of whether an AI system functions successfully is thus reduced to whether it holds significance *for us* (Floridi et al. 2009), an anthropocentric horizon that leaves the system's broader ecological effects and the transsemiotic exchanges in which it participates entirely unexamined.

Partial exceptions exist in adjacent fields. Animal-computer interaction has spent over a decade developing frameworks for non-human evaluation of computational systems. Mancini's foundational programme positions animals as legitimate users and stakeholders in interaction design, while subsequent work has explored species-specific interfaces and welfare-oriented evaluation (Mancini 2011; Mancini et al. 2017; Hirskyj-Douglas et al. 2018), though almost exclusively for mammals and other macroscopic animals. More-than-human HCI extends this orientation beyond individual animals to ecological assemblages, asking how design might serve multispecies collectives rather than human users alone (Wakkary 2021; Liu et al. 2018). Environmental AI research, meanwhile, has begun to address AI's entanglement with energy systems and ecological monitoring, though primarily through the lens of computational sustainability and resource optimisation (Rolnick et al. 2022). These programmes share our concern with non-human significance but differ from the present approach in how they evaluate whether computational systems serve non-human needs *as defined by human observers*. While these fields strive for participatory or non-anthropocentric methods, the evaluative horizon ultimately remains anthropocentric even when the beneficiary is not. With our Transsemiotic Turing Approach, we, in contrast, ask whether AI outputs acquire significance *within non-human interpretive processes themselves*. A question that requires biosemiotic resources, the interaction-design tradition typically does not employ.

The argument of this paper is not that these benchmarks measure the wrong things. It is that significance also exists in domains human evaluation cannot access, and that AI systems already operate in such domains. Whenever an AI system generates a material output that propagates through an environment with non-human organisms, it enters domains of significance that current evaluation paradigms cannot register. The issue is not that we lack the instruments to measure non-human significance (though

we do) but that the very concept of evaluation, as currently constructed, presupposes a human evaluator whose judgements define what counts. The broader intellectual context for this claim lies in the multispecies turn in environmental humanities (Tsing 2015; Haraway 2016), which has shown that human meaning-making is one semiotic practice among many, embedded in and dependent upon ecologies of signification that vastly exceed it. We sharpen this insight into a question that will organise this article: can machines matter to non-humans?

Turing positioned human judgement as the sole arbiter of machine intelligence (Turing 1950). His imitation game asks whether a machine's outputs are indistinguishable from a human's, and the human interlocutor is both the benchmark and the judge. We invert this. Our question is not whether machines can deceive humans but whether AI outputs acquire significance for non-human organisms (Fig. 1). Where Turing's criterion was successful deception, ours is functional appropriation. In this inverted setting, the measure of success is whether non-human organisms can incorporate those outputs into their own interpretive repertoires, a condition that the Approach is designed to diagnose, rather than certify. We invoke Turing not to re-enter the debate about his test's validity but to mark the anthropocentric assumption his framing crystallised and to invert it. We call this the Transsemiotic Turing Approach. It is an approach rather than a test because it diagnoses conditions for transsemiotic exchange along graduated dimensions rather than certifying intelligence through binary evaluation. The shift from "test" to "approach" hereby matters. While a test delivers a verdict, an approach opens an investigation.

The empirical ground for this investigation is the materiality of language. While independent literature establishes that microbes respond to acoustic structure (Aggio et al. 2012; Sarvaiya and Kothari 2015), a parallel quantitative study by the authors examined whether human speech specifically can act as such a structural signal under controlled laboratory conditions (Anonymous 2025). It reported that distinct phonetic configurations are associated with distinct morphological responses in the microbial cell-skeleton. Tonal and aperiodic profiles diverge. We treat this as motivating prior work, not as a finding that the present paper must establish. The premise we carry forward already rests on the independent sources above, namely that it is the acoustic structure rather than gross energy that can act as a genuine environmental stimulus for microorganisms. If human voice can matter to microbes in this sense, the question arises whether computationally generated voice might matter too.

Answering whether machines can matter to non-humans requires a working concept that AI evaluation has not formalised: *transsemiotic exchange*, the process through which computational outputs can become functionally operative within non-human interpretive worlds. The term *transsemiotic* has prior currency in semiotics and translanguaging research, where it describes the mobilising of heterogeneous semiotic resources in meaning-making (Lin 2015). We retain the term but not that sense, and our concept departs equally from intersemiotic translation (Jakobson 1959), Gorlée's inter-code translation (Gorlée 1997), and Kull's translation between Umwelten (Kull 2009). Where translation presupposes a mediating position *between* domains, and the integrative usage an actor who mobilises multiple semiotic resources, our *trans-* names a third topology, traversal *across* entities bound by neither a common code nor an

orchestrating actor. What crosses is physical rather than semantic, akin to *transduction* (Section 2.4). The synthetic phoneme is a sign-type only on the agent's side, and its semiotic identity does not survive the crossing. The organism makes operative a physical perturbation taken up under its own categories, not a phoneme folded into its repertoire. We name this process *transsemiotic exchange* rather than transversal because the crossing is not merely physical but a reconstitution by the organism's Umwelt semiosis.

The paper develops a conceptual framework grounded in artistic-research practice and makes three claims. First, current AI evaluation regimes restrict evaluative relevance to anthropocentric horizons, foreclosing non-human evaluative domains in which AI systems already operate. Secondly, this paper specifies and partially exercises a biosemiotically informed diagnostic structured around transduction as a material precondition, as well as selectivity, appropriation and persistence. This diagnostic is used to investigate one such domain at the level of vegetative semiosis. Third, three installations enacted between 2022 and 2024 materialise this framework and motivate the hypothesis that configurations richest in transsemiotic encounter resist diagnostic clarity. The installations also support the proposition that AI operates as a kind of temporal infrastructure here, accumulating patterns across durations ranging from computational milliseconds to biological days. The diagnostic criteria were derived from theoretical research, while the generative principles that sustain the apparatus emerged through material friction with the installations themselves. Within this scope, the paper diagnoses conditions under which differential biological responsiveness to structured computational stimuli becomes systematically accessible. Whether such responsiveness constitutes significance in the strong semiotic sense is what the framework declines to verify, and what makes it a diagnostic rather than a test.

Biosemiotics provides the conceptual vocabulary, supplying terms for semiosis beyond human language and threshold concepts for distinguishing interpretation from mechanism. Deep reinforcement learning and real-time biological feedback provide the technical means through which these thresholds can be probed. The three artistic-research installations described in Section 4 are the sites where vocabulary and technical means meet practice.

Any attempt to conceptualise meaning beyond human frameworks necessarily employs human-derived concepts; our instruments, languages, and questions are anthropogenic. This paradox governs all claims that follow and returns in the Discussion, where the installations materialise this not as a methodological failure, but as a constitutive feature of transsemiotic enquiry.

2 What AI evaluation excludes: biosemiotic resources

2.1 Semiosis beyond human language

The biosemiotic tradition reconceptualises life as a process of signification. Peirce's triadic model of sign, object, and interpretant explains how physical perturbations become meaningful through organismic response (Peirce and Hartshorne 1931). Uexküll's Umwelt concept names the species-specific world of significance each organism navigates through interpretation rather than mere reaction (von Uexküll 2010).

Sebeok brought these strands together, establishing semiosis as a property of all living systems (Sebeok 1968).

This establishes that organisms interpret, but not how to identify interpretation, when we observe. Here, Gregory Bateson supplies the missing threshold. His bridge between information theory and biological meaning, indispensable to the framework we develop below, turns on a single formulation: information is "a difference which makes a difference" (Bateson 2000). We redeploy this formulation as a diagnostic threshold at which a physical response becomes interpretive discrimination: a microbial cell discriminating between frequencies of equivalent acoustic energy exhibits Batesonian information-processing when it treats their difference as significant.

Within this framework, meaning is a functional relationship, not a mental abstraction. Kull's taxonomy of semiotic levels distinguishes vegetative, animal, and cultural semiosis, each involving progressively more complex sign-processes (Kull 2009). Vegetative semiosis, the most basic level, involves organisms responding to environmental signs through biochemical and morphological pathways without requiring consciousness, attention, or neural processing. We situate the individual oral microbial cell as the vegetative interpreter. The cell integrates acoustic texture through cell-skeleton reorganisation and metabolic shifts. Community-level signatures (e.g., shifts in Shannon-diversity, as observed in Installation C, Section 4.3) we read as the aggregate expression of cell-level vegetative semioses, not as evidence for a community-level interpretation in its own right, which would be a stronger claim that would require an additional argument we do not make here. This positioning is not a metaphorical extension of human semiosis to microbes but an application of Kull's taxonomy at the level it explicitly identifies as semiotic. The limits of vegetative semiosis, as defined by Kull, are constrained to recognition and iconic sign processes, excluding symbolic representation. The diagnostic criteria we develop below respect this constraint. Selectivity requires discrimination, not representation; appropriation requires interpretive integration through existing repertoires, not the construction of novel symbolic categories; and persistence requires molecular memory, not propositional recall.

Sharov's concept of functional information extends this analysis by specifying three conditions under which biological responses constitute genuine semiosis rather than mechanism. The response must be selective (sensitive to specific signal properties), functional (oriented towards the organism's goals), and heritable or memorable (capable of modifying future behaviour) (Sharov 2010). We adopt Sharov's three conditions and operationalise them as the diagnostic criteria presented in Section 3.1. Sharov's selectivity becomes our selectivity criterion, his goal-directedness becomes our appropriation criterion, and his heritability becomes our persistence criterion. Through this lens, phonetic speech (whether human or synthetic) is stripped of its symbolic content and potentially re-semanticised as a biological sign. The phoneme "ba" carries no linguistic meaning for a microbial cell, but parts of its vibrational profile may nevertheless be appropriated for biological significance.

2.2 Accidental iconicity: how non-signs become signs

Why use phonetic speech as the interface between AI and microbes, rather than pure tones or white noise? Phonetic speech provides a multidimensional stimulus space of

spectral composition, temporal patterning, periodicity, and articulatory configuration that exceeds what single-frequency experiments can provide. This richness allows the RL agent to probe complex, structured acoustic signals rather than single parameters.

The mechanism by which organisms materially re-semanticise these patterned sounds is a process of accidental iconicity. Alexander theorises that biological creativity frequently emerges from "fruitful misinterpretation", an organism responds to a novel, unintended stimulus as if it were a familiar biological sign because they share an accidental physical resemblance (Alexander 2013). We adopt this concept as the primary mechanism in transsemiotic exchange. The AI-generated phoneme is not crafted to mimic environmental cues, yet its vibrational structure possesses frequency characteristics that a cell's mechanosensory system may treat as biologically significant. Kleisner and Maran's concept of semiotic co-option names a related but distinct process, namely the evolutionary recruitment of a perceptual category to a novel function (Maran and Kleisner 2010; Kleisner 2011). We advance the idea that accidental iconicity could be the mechanism through which co-option-like dynamics operate within a single organism's interpretive history. The shift from phylogenetic to within-lifetime timescales differentiates the present proposal from co-option as classically defined. What differentiates it from associative learning is the iconic basis of the response. Associative learning pairs an arbitrary cue with an outcome through repeated contingency. Accidental iconicity, by contrast, would turn on a resemblance the organism's evolved sensory apparatus already treats as significant, not on such a learned pairing.

The iconicity at stake here operates at two Peircean sub-levels. At the image-icon level, the shared sensory qualities of a phoneme's vibrational profile resemble a biologically significant cue. Treated as image-iconic alone, this reduces to stimulus-generalisation. What would distinguish genuine appropriation from mere generalisation is a diagrammatic level of shared relational structure across, e.g. formant-transition geometry or periodic-to-aperiodic ratios (Peirce and Hartshorne 1931; Stjernfelt 2014). We acknowledge that the present apparatus cannot itself produce this discrimination. The diagrammatic claim (a response to novel frequency patterns that preserve relational structure while changing absolute values) is therefore advanced as a criterion for future research.

2.3 The threshold of meaning: from causation to interpretation

To define the threshold between causation and semiosis, we adopt Sharov's concept of functional information (Sharov 2010). The question is not whether microbial cells physically respond to AI-generated vibration but whether such a response constitutes interpretation rather than reflex. If purely mechanistic, this project is an exercise in physics. If even minimally interpretive, a domain of non-human significance opens that current AI evaluation cannot register.

Three convergent arguments support the proposition that interpretation exists independently of human language, neural architecture, or consciousness. Kohn demonstrates that Peircean representation pervades living systems (Kohn 2013), Lyon reframes bacterial behaviour as biogenic cognition (Lyon 2015), and Thompson's enactivism argues that if living is itself a form of sense-making, microbial responses need not be reduced to "mere mechanism" (Thompson 2007).

The objection that all biological responses reduce to deterministic physics (an approach Bickhard critiques for yielding informational specificity without normative subjectivity (Bickhard 2009)) faces three counterarguments. First, organisms make semiotic errors. A microbial cell responding to a fake pheromone has been *deceived*, and deception presupposes interpretation (Hoffmeyer 2008). Second, molecular memory mechanisms encode these "deceptive encounters" to alter future interpretive thresholds (Caudron and Barral 2013). This results in present behaviour shaped by physically absent past events, a phenomenon Deacon terms "constraint by absence" (Deacon 2012). Third, the functional orientation of microbial response constitutes what Canguilhem terms biological normativity, the intrinsic establishment of norms determining what counts as beneficial or harmful (Canguilhem 1991). This normative orientation also grounds the distinction between normative and parametric selectivity developed below.

These arguments do not prove that microorganisms understand AI output. They establish that organismic responses involve properties traditionally associated with interpretation: selectivity, historical conditioning, and normative orientation. The underdetermination mirrors the impasse in AI theory over whether LLM pattern-completion constitutes understanding (Bender et al. 2021). Our framework reframes both debates by shifting from intrinsic properties (does the system "really" understand?) to relational achievements (under what conditions do outputs enter another system's interpretive processes?). This extends the debate on grounding in computational linguistics. Whereas Bisk et al. presuppose grounding through human-centric social or physical interaction (Bisk et al. 2020), the present setup positions non-human biological responses as the grounding substrate. We return to the implications of this displacement in Section 5.4.

2.4 AI as semiotic transducer

Classical AI presupposed that intelligence operates through symbol manipulation within a shared code between system and environment (Haugeland 1989). Contemporary LLMs depart from this paradigm through statistical pattern extraction, but remain anchored to the human linguistic horizon because they rely on human training corpora and reward signals along with human-legible outputs (Bender et al. 2021). Recent multimodal systems generate audio, video, and sensor-derived outputs without text intermediation (Borsos et al. 2023), partially closing the gap between linguistic and non-linguistic computation, yet even these models remain tethered to human-curated data and human-designed evaluation.

We position our RL setup differently as semiotic transducers because a transducer converts signals across material substrates without preserving semantic content (Wiener 1961). The RL agent receives biological reward signals in the form of microbial vitality measurements, processes them computationally, and produces acoustic outputs such as synthetic phonemes. At no point does the system represent the meaning of the microbial response, because it converts perturbations from one material domain to another. While the transducer concept is borrowed from signal processing, it acquires interpretive weight in this context, since the agent does not need to understand the organism's Umwelt to participate in it.

This framework suggests a functional parallel between the RL policy and Dreyfus's account of skilful coping because both involve an embodied and situation-sensitive know-how that operates without explicit symbolic representations (Dreyfus 1992). Over thousands of training steps, the agent develops a repertoire of phonetic configurations that produce differential biological responses without ever representing what those responses mean. Within this transsemiotic circuit, the skill is not represented but enacted.

If the RL agent responds selectively to biological signals, integrates them into its policy and retains them in its weights, does this not also imply that microbes matter to the AI and that it satisfies the diagnostic criteria? The answer demands a turn on a distinction between *normative selectivity* and *parametric selectivity*. Microbial selectivity is normative because it is oriented by biological imperatives such as survival and homeostasis that are constitutive of the organism's existential condition (Thompson 2007; Canguilhem 1991). The RL agent's selectivity is parametric because it is oriented by an externally specified reward function whose preferences carry no functional significance for the agent itself. While the organism's interpretive history is sedimented in molecular memory, the agent's history is stored in weights that dissolve into mere calculation. This asymmetry does not disqualify AI from participating in transsemiotic exchange but prevents the framework from being trivially satisfied in both directions, since the diagnostic criteria presuppose a living interpreter whose responses are normatively structured.

The same capacity that enables encounter could also enable extraction (Crawford 2021): an RL setup that learns to produce sounds affecting microbial vitality could serve either transsemiotic investigation or industrial biomanipulation. This biopolitical tension is constitutive of the framework, and we return to it in Section 5.3.

3 The Transsemiotic Turing Approach: framework and method

The method presented here consolidates insights that emerged through artistic research practice. Rather than testing a pre-existing framework, the installations presented in Section 4 co-constituted it, following the logic of artistic research, whereby knowledge emerges through, within and with creative practice rather than being applied to it from outside (Borgdorff 2012). The epistemological status of this work is therefore twofold. The diagnostic criteria (Section 3.1) are derived from biosemiotic theory. Their warrant is conceptual coherence and fit with that tradition, and they are argued rather than tested. The generative principles (Section 3.2) are practice-led knowledge in Borgdorff's sense (Borgdorff 2012). Their warrant is the documented practice. The biological phenomena encountered in the installations are motivating and illustrative for the present, conceptual contribution. Some of the underlying effects they gesture towards are examined separately, under controlled laboratory conditions, in a parallel quantitative study (Anonymous 2025). This paper refers to that study but does not depend on its results, and the study reports no statistical analysis of the installation's own data. Accordingly, this paper claims to contribute a framework and

vocabulary, a set of generative principles for sustaining more-than-human computational encounters, and two practice-derived propositions that invite future empirical work.

3.1 Diagnostic criteria for transsemiotic exchange

As developed above, transsemiotic capacity is a relational achievement emerging in specific assemblages of organism, machine, and environment. The framework does not certify intelligence but instead diagnoses conditions under which computational outputs may acquire biological significance. Where Turing asks if a machine is intelligent, we ask in what ways and under what conditions AI systems might matter to non-humans, pursued in this paper with microbial communities. Therefore, our Transsemiotic Turing Approach is not a benchmark in the sense of a standardised metric encoding a replicable test (Raji et al. 2021), but a diagnostic apparatus whose contribution lies in mapping the conditions under which transsemiotic evidence becomes producible. Our framework does not propose an alternative metric within existing evaluation paradigms but instead questions the assumptions those paradigms share regarding the nature of intelligence itself.

The precondition for any diagnostic reading is **transduction**, where the AI-generated signal must cross material substrates between semiotic domains, moving from digital computation through acoustic vibration to the cell. In principle, transduction could be a graduated dimension involving different qualities of signal transformation across substrates. However, at the resolution afforded by our apparatus, it functions as a threshold. Either the signal materially reaches the organism, or it does not. The gap between confirming physical transduction (signal reaches organism) and diagnosing semiotic transduction (signal acquires significance within the organism's interpretive world) is precisely the space our diagnostic criteria are designed to investigate.

Once this threshold is crossed, we propose three diagnostic criteria corresponding to the three conditions for functional information established by Sharov (Sharov 2010), to assess the depth and character of the exchange:

1. **Selectivity.** The organism responds distinctively to specific properties of the signal rather than to gross energy. Not all AI outputs produce the same biological feedback because some configurations matter more than others. This criterion corresponds to Sharov's condition that the response must be sensitive to specific signal properties and is additionally grounded in Bateson's concept of "a difference which makes a difference" (Bateson 2000; Sharov 2010). It distinguishes an interpretive response from mere physical reactivity. (On the distinction between normative and parametric selectivity see Section 2.4.)
2. **Appropriation.** Through the accidental iconicity of Section 2.2, the organism integrates AI output with its interpretive repertoire, so that the foreign signal acquires functional significance for the organism's own functioning. This functional integration instantiates Sharov's goal-directedness condition (Sharov 2010). As exercised in our three installations, it is an inference to the best explanation, not a demonstrated

property. A discriminating diagrammatic-iconic test lies beyond the present apparatus. So, appropriation names the most plausible reading of a response interpretable through the organism's existing repertoire, weighed against stimulus generalisation and associative learning rather than shown to exceed them.

3. **Persistence.** The encounter alters future interpretive behaviour. This criterion corresponds to Sharov's condition of memory and is grounded in molecular mechanisms microbial cells utilise to encode encounters through epigenetic modifications that alter future activation thresholds (Caudron and Barral 2013). Hoffmeyer's concept of semiotic scaffolding (Hoffmeyer 2008) describes how interpretive capacities accumulate over evolutionary time through the progressive layering of semiotic constraints. We do not claim that within-lifetime molecular memory constitutes semiotic scaffolding in Hoffmeyer's strict sense. What we investigate is a within-lifetime analogue where epigenetic traces constrain an individual organism's future interpretive thresholds, functioning as scaffolding for the organism's own semiotic history. Deacon's analysis of ententional phenomena (processes constrained by what is not physically present) provides the theoretical anchor here (Deacon 2012). The molecular trace of a past acoustic encounter constitutes precisely such a constraint-by-absence, shaping responses to signals the organism has not yet received. The persistence criterion distinguishes transient perturbation from sustained semiotic integration. It defines our framework's horizon by representing the condition for a strong claim of transsemiotic exchange. Where selectivity and appropriation are what the current apparatus can diagnose, persistence marks what would need to be demonstrated for full transsemiotic exchange (the satisfaction of all three diagnostic criteria together). Its practical inaccessibility at the short timescales of temporary installations is surfacing the temporal requirements of such investigation.

When the transduction precondition and all three criteria are satisfied, a perturbation originating in computational processes has been transduced into the interpretive world of the organism. Then the organism does not passively resonate with phonetic vibration but responds with what Sharov terms protosemiosis: goal-directedness without object representation, which presupposes an autonomous, self-maintaining interpreter (Sharov and Vehkavaara 2015). When a synthetic voice articulation triggers this integration across all three criteria, the output has become a historically interpreted sign within the organism's Umwelt.

3.2 Generative principles for sustaining the encounter

Where the theoretically derived diagnostic criteria read what was produced, three practice-derived generative principles structure how the apparatus sustains the transsemiotic encounter.

Semiotic probing names the iterative process through which the RL algorithm generates computational tokens and evaluates them through biological responses. The algorithm proposes acoustic configurations not as linguistic units but as interrogations into possible transsemiotic relationships. The organism's differential response functionally establishes or dissolves provisional significance. An acoustic configuration nonsensical to humans may carry profound significance for microorganisms, and one

that sounds meaningful to us may register as noise. Probing and validation operate as simultaneous stages rather than sequential ones, where action-space exploration and reward signal evaluation are coupled within each training step.

Drift names the structural necessity for open-ended exploration born from competing teleologies within the apparatus. We use this term to describe a system that is forced to wander, rather than a stochastic computational or biological process. A genuine encounter cannot be engineered towards a single target but must emerge through unguided interaction. Because the computational algorithm seeks mathematical convergence, while biological entities strive for metabolic homeostasis, the system must continually rotate control to prevent any single teleology from dominating. This prevents the apparatus from optimising towards a predetermined biological coupling. Instead, the system explores a transsemiotic possibility space pulled by dynamic attractors that emerge from interaction. This functionally open-ended approach distinguishes a transsemiotic configuration from an industrial optimisation pipeline and connects directly to the biopolitical distinction between encounter and extraction developed in Section 5.3.

Recursive attunement names the cascading feedback loops required to sustain the transsemiotic exchange against the embodied inertia of participating entities. The process is recursive because each cycle's output reshapes the conditions for the next. Each computational acoustic output alters microbial metabolic states, which in turn reshape the subsequent computational composition, generating transformations across multiple timescales simultaneously. Immediate computational adjustments unfold alongside hourly biological shifts and daily morphological adaptations. This multi-scale coupling prevents transsemiotic decoupling when interacting with preconditioned human participants and grounds the paper's observation of AI as a temporal infrastructure, detailed in Section 5.1. Where semiotic probing describes what the system explores, and drift describes how it moves, recursive attunement describes how the interaction sustains itself over time.

3.3 Modes of reading and reflexive constraints

Each diagnostic criterion is approached through a distinct mode of reading the installations. Table 1 summarises these modes and records, for each, whether it was actually exercised in the installations or only specified.

Selectivity would ideally be read through randomised control soundscapes, in which parallel cultures receive randomised phonetic sequences unconnected to the RL outputs, to distinguish sound-specific effects from sound-general ones. These parallel control cultures were not run in the installations. What the installations afforded was a looser within-system contrast at matched energy (tonal versus aperiodic in A, varying melodic phrases in C), read qualitatively. Appropriation is read by comparing the responses we worked with against known microbial behavioural repertoires, such as the shmoo morphological transformations characteristic of mating responses (Caudron and Barral 2013). The logic is that once generic stress responses are set aside, an organism responding to a computationally generated signal through interpretive resources evolved for biologically meaningful stimuli makes accidental iconicity the most plausible reading. Persistence is read through time-shifted stimulus testing, in

Diagnostic Criterion	Mode of reading	Status in the installations
Selectivity	Comparison against randomised control soundscapes with equal gross energy	A qualitative within-system contrast was read. Controlled randomised soundscapes were not utilised in the installation setting.
Appropriation	Interpretation against known microbial behavioural repertoires (metabolism, skeletal reorganisation)	Exercised as a qualitative reading in Installations A and C; unattributable in the distributed configuration of Installation B
Persistence	Repeated time-shifted stimulus testing, with delays spanning two or more reproduction cycles	Specified but not exercised; the required window could not be sustained at festival timescales

Table 1 An overview of the mode of how each diagnostic criterion was read, and whether this mode was actually used in the three installations or only specified. These are interpretive procedures for practice-led work, not empirical validation tests. Here, "exercised" means that a reading was performed, not that a statistically significant effect was demonstrated. This exercised/specified axis is distinct from the per-installation legibility (accessible/inaccessible) axis reported in Table 2. Transduction is treated as a material precondition rather than an equal criterion.

which the delays between acoustic outputs and RL feedback integration are set to span several of the organism's own reproduction cycles. Because reproduction serves as the behavioural lens for observing persistence, the reproduction cycle provides the natural biological baseline. In yeast cultures, for instance, a single reproduction cycle lasts roughly 90 to 240 minutes (Caudron and Barral 2013). We acknowledge a significant methodological friction here, enforcing this biological window for a single feedback cycle drastically slows the computational semiotic probing. This deceleration precludes the high-volume interactions that statistical certainty would require, while also productively subjects the computational system to the temporal reality of the organism. In the installations reported below, this window could not be sustained, so persistence remained a mode of reading we specify rather than one we were able to exercise. If an organism's response were to outlast the stimulus and modulate its response to subsequent ones, the encounter would have left an observable short-term trace, a precursor to deeper transsemiotic persistence.

The first author's dual position, as both the technical architect of the RL system and the arts-based researcher interpreting its outputs, functions not as a bias to be mitigated, but as the primary epistemic engine of the framework. This situated standpoint internal to the apparatus was a methodological necessity for retrieving the generative principles directly from the material friction of building and sustaining the installations. Furthermore, the apparatus itself enacts what Barad calls agential cuts (Barad 2007): pH sensors at 0.5-unit resolution render finer fluctuations invisible, and a reward architecture that equates pH and optical-density signals co-defines

what counts as semiotic engagement. These are constitutive of the phenomena the apparatus produces.

The framework's contribution is diagnostic, mapping and reading criteria qualitatively across graduated dimensions rather than relying on null-hypothesis significance testing or effect size estimation. The installations, moreover, were designed as practice-led investigations that prioritise ecological validity over experimental control. Any differential biological responses reported in the following sections should therefore be understood as observed patterns supporting qualitative judgements of legibility, not as statistically confirmed effects.

3.4 Technical implementation: reinforcement learning and vocal synthesis

Two technologies enable this transsemiotic interaction across temporal and vibrational domains. *Proximal Policy Optimisation* (PPO) provides the on-policy reinforcement learning mechanism (Schulman et al. 2017). Its continuous policy updates respond to incoming reward signals that combine biological feedback (pH from sensors for Installations A and B and colourimetric indicators for Installation C, optical density from a custom LED-photosensor transmission setup, and a Shannon diversity index from automated visual inspection of colony types for Installation C) with algorithmic curiosity rewards (Schmidhuber 2010; Pathak et al. 2017). This curiosity component prevents convergence on a single optimised output, providing the computational engine for our practically derived generative principle of drift. *Pink Trombone* (Thapen 2017), a physical modelling voice synthesiser with a continuous articulatory parameter space, allows the PPO to explore acoustic configurations between and beyond the human sounds catalogued in the International Phonetic Alphabet (International Phonetic Association 1997). This continuous synthesis space directly operationalises our generative principle of semiotic probing.

The diagnostic architecture of our framework is, in principle, modality-independent. Selectivity, appropriation, and persistence can also be assessed with other forms of transductions, such as chemical, tactile, or electromagnetic signals. However, although the criteria and generative principles could potentially be transferred without structural modification, such extension would require different means of transduction. This asymmetry between the generality of the framework and the specificity of the practice positions the Transsemiotic Turing Approach as a broader research ambition rather than a single experimental setup.

4 Three installations (2022–2024)

These installations are material-discursive arrangements that produce specific configurations of the observable rather than certifying pre-existing phenomena (Barad 2007). As mentioned earlier, the diagnostic criteria presented in Section 3 were theoretically derived and then applied to these physical cases. In contrast, the generative principles emerged entirely through iterative encounters with the installations, where the material friction of each installation necessitated new strategies to sustain the transsemiotic

exchange. Between 2022 and 2024, all three installations were presented at international media-art venues over multi-day public exhibitions, with audiences in the low thousands (festival, installation, and venue names are redacted for review).

4.1 Installation A: sub-linguistic exploration and semiotic probing (2022)

Installation A functioned as a transsemiotic exploration where oral microbiota, a PPO algorithm, and human participants collectively interacted with acoustic configurations hidden between and beyond the IPA's standardised phonetic inventory (International Phonetic Association 1997). The installation utilised a single microbial species, providing focus at the cost of ecological complexity.

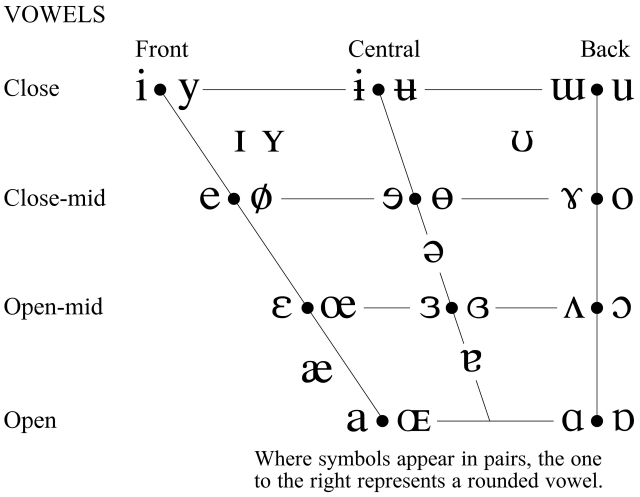


Fig. 2 The International Phonetic Alphabet Vowel Chart, representing the boundaries of human linguistic standardisation challenged by Installation A. Modified from International Phonetic Association (2026) under CC BY-SA 4.0.

Installation A provided the first sustained encounter between the diagnostic framework and physical reality. Transduction was confirmed through parallel measurements tracking signal propagation from the computational outputs to the microbial medium. Within the installation, tonal frequencies appeared to support microbial vitality, while volume-matched broadband aperiodic noise was associated with community decay, a tonal/aperiodic divergence consistent with the theoretical premise. We read this difference as the association the selectivity criterion is meant to make visible, suggesting that difference, not magnitude, was what mattered. Cell-skeleton observations showed changes reminiscent of mating-type responses. Yet whether this consistency indicated genuine appropriation or mere structural coincidence remained underdetermined at the available resolution.

This epistemological gap between the theoretical criteria and the physical reality necessitated new generative strategies to sustain the transsemiotic encounter. To satisfy the theoretical demand for a continuous encounter, the PPO algorithm navigated the Pink Trombone synthesiser to propose articulations beyond conventional human phonemes (Fig. 2). These computational tokens functioned not as linguistic units but as sonic interrogations into possible transsemiotic relationships. This active exploration of the unmapped phonetic possibility space provided the empirical basis for the generative principle of semiotic probing. While the installation also yielded swift glimpses of drift and recursive attunement, the necessity of this probing mechanism remains its primary generative legacy.

Installation A surfaced three observations shaping subsequent iterations. First, acoustic configurations lacking human linguistic status were associated with differential microbial responses. Second, morphological changes suggestive of appropriational responses require converging evidence to be unambiguously distinguished from mechanistic effects. Finally, single-species cultures simplify oral microbiome ecology in ways that obscure community-level dynamics, motivating a shift in Installation C to complex multi-species cultures.

4.2 Installation B: competing teleologies and drift (2023)

Installation B invited human visitors to embody co-produced vibrational timelines through body vibration and vocal mimicry of a synthetic voice articulator, utilising custom vibration platforms (15 Hz to 200 Hz) and 8-channel spatial audio. Where Installation A positioned the human visitor as observer and occasional participant, Installation B dissolved the boundary between observer and installation, distributing generative participation across the human, microbial, and computational entities within the system.

While the material precondition of transduction was successfully tracked across the physical and biological interfaces, applying the subsequent diagnostic criteria proved challenging. Selectivity was partially legible, as distinct vibrational regimes were associated with different pH trajectories, which we read as the community responding differentially to the collective acoustic output. Diagnosing appropriation, however, presented a structural difficulty. Because each acoustic moment was an emergent composite of three simultaneous contributions, it was impossible to isolate the computational entity's specific role in the environment. While the microbiota may readily appropriate stable collective patterns regardless of their origin, whether the AI's distinct outputs were appropriated became diagnostically unattributable. The appropriation criterion thus met a configuration that exceeded its unit of analysis. Ultimately, the installation that distributed generative participation most radically also rendered the diagnostic criteria least legible, a tension discussed further in Section 4.4. Persistence remained inaccessible at the available temporal resolution.

The material friction of this distributed architecture provided the empirical basis for the generative principle of drift. During prototyping, we observed that each entity introduced a distinct systemic attractor: the algorithmic architecture optimised for mathematical convergence, the biological medium defaulted to metabolic homeostasis, and human visitors commonly pulled the output towards recognisable acoustic

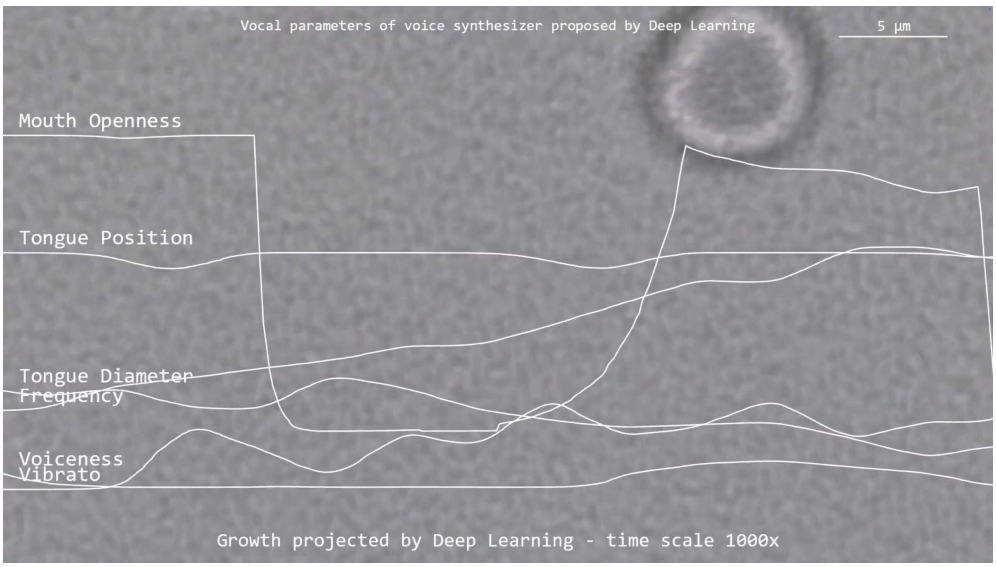


Fig. 3 Live-visualisation of the voice articulation parameters within Pink Trombone. The trajectories represent the synthetic voice articulations modulated by PPO, microbiota, and audience.

coherence. If permitted to steer the encounter, a single entity inevitably dragged the system towards its own operational baseline, foreclosing exploration. To preserve open-endedness, data pathways assign each source an equal 33.33 percent mathematical weight and continuously rotate parameter control (Fig. 3). This structural choice forces these competing tendencies to constantly interrupt one another, ensuring the system wanders rather than converges. Drift manifests here through the dissolution of isolated action into a dispersed relational sound event. This dynamic operates through cascading feedback loops spanning multiple timescales, where vocal input alters microbial vitality within minutes, which modifies PPO policies over the following hour, and subsequently affects synthetic vocal production in later cycles. While prefiguring the recursive attunement formalised in Installation C, these temporal loops primarily ensured that over the five-day duration, the system accumulated patterns no single entity initiated or controlled.

Installation B suggested that the diagnostic architecture presupposes a degree of attributability that distributed configurations dissolve. When the acoustic environment is collectively generated, the search for a singular signal origin loses its referent. The most ecologically rich conditions for a transsemiotic encounter resisted legibility, an observation that became the foundation for the inverse-legibility hypothesis put forward in the paper's conclusion.

4.3 Installation C: embodied inertia and recursive attunement (2024)

Installation C is a singing performance in which a PPO algorithm transcodes the metabolic shifts of an externalised oral microbial community into modified sheet music, adapting the vocal score between two songs according to fluctuating biological states. The performance featured a trained classical vocalist performing with an in vitro culture derived from the performer's own oral microbiota samples. Where Installation A used a single species and Installation B utilised a distributed weighting architecture, Installation C introduced a new variable by incorporating the inertia of a whole microbiome community and a highly conditioned human performer.

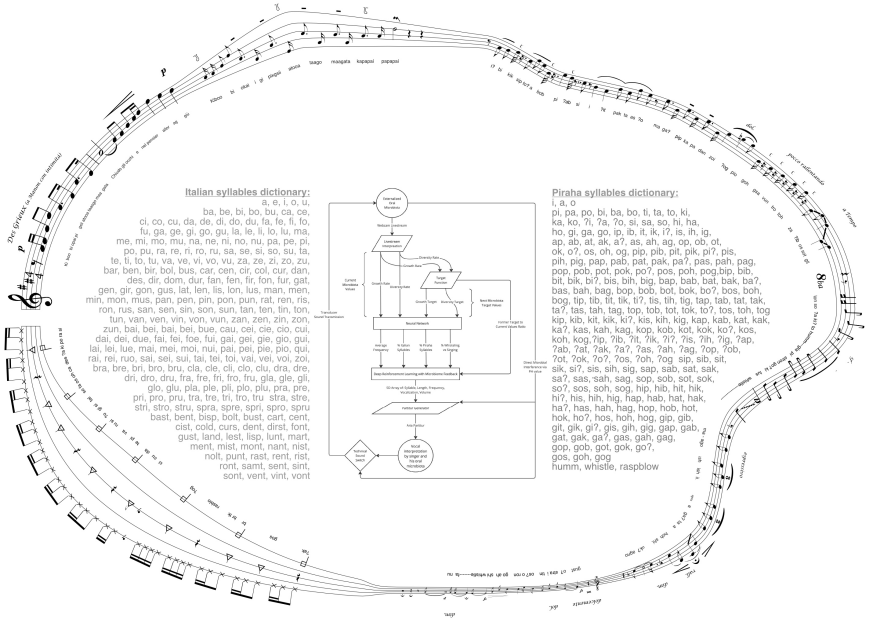


Fig. 4 Syllable library of Installation C and a map of the recursive attunement loop adapting the singing score based on the microbial community's need to maintain diversity with the singer's vocalities. Adapted from a visual by a programme collaborator, full credit will be provided upon acceptance.

This installation served as the final site at which the diagnostic framework was enacted in practice. Tracking acoustic energy from the vocalist to the microbial cultures documented the transduction precondition. Selectivity became legible in the practice setting, as different melodic phrases were associated with apparently distinct diversity-index patterns. Following the unit of analysis specified in Section 2.1, we read these patterns as the aggregate signature of cell-level selectivity across the community rather than as semiosis attributed to the community as such. Appropriation, too,

became comparatively readable as the responses we worked with were interpretable through known diversity-maintenance repertoires rather than as arbitrary perturbations. This alignment with survival strategies is what accidental iconicity would lead us to expect, and we offer it as the most plausible reading, not a demonstrated finding. Persistence, however, remained inaccessible within the brief duration of the exhibition.

Attempting to sustain this encounter surfaces a different kind of resistance. The vocalist's professionally trained voice consistently overpowered the computationally suggested modulations with their ingrained technique. Their embodied expertise functioned as a barrier to the porosity that transsemiotic exchange requires. This embodied inertia provided the empirical basis for our generative principle of recursive attunement. Because the human possessed such rigid momentum towards Western musicality, the system required continuous cascading loops to prevent a total decoupling. The algorithm transcoded real-time microbial data into an evolving score; the singer interpreted this score within minutes; their acoustic performance altered the microbial environment over the following hours; and this metabolic shift reshaped the subsequent computational composition. Unlike a standard cybernetic feedback loop designed to optimise for a target state, this strict looping necessity functioned to sustain an open-ended encounter across incommensurable timescales.

While semiotic probing and drift operated sequentially to explore melodicity (Fig. 4), the singer's trained robustness functioned as the antithesis of the openness that drift demands. This empirical observation generated our concept of semiotic plasticity, defined as the capacity of an entity to be perturbed and restructured by signals from an alien domain. This concept is related to, but distinct from, Hoffmeyer's semiotic freedom (Hoffmeyer 2008), which names the active range of choices available to an organism. Where semiotic freedom describes active repertoire breadth, plasticity describes receptive porosity. Effective transsemiotic coupling requires plasticity from all participating entities. The vocalist's embodied inertia suggested the systemic breakdown that occurs when one entity within the installation lacks this porosity.

4.4 Cross-case synthesis: differential accessibility

The three installations, read together, suggest an inverse relationship between encounter richness and legibility. Installation B distributed generative participation most radically and dissolved isolated action, creating what we read as the richest conditions for transsemiotic encounter. It was also the installation where the diagnostic criteria became least legible: appropriation was unattributable, and selectivity was confounded by collective dynamics. Installation A occupied a middle position, where single-species cultures allowed a clearer reading of selectivity but limited the ecological complexity of the encounter. Installation C, where a single trained performer introduced severe embodied inertia, produced the most diagnostically legible reading of appropriation. Yet the performer's trained robustness simultaneously suggested that such rigidity can foreclose the very coupling it makes legible.

We interpret this inverse relationship as a framework-derived hypothesis that requires targeted empirical validation. The very conditions that maximise the possibility of transsemiotic exchange (distributed participation, multi-entity interaction, temporal extension, and semiotic plasticity) may be precisely the conditions that

defeat neat diagnostic interpretation. The appropriation row of Table 2 makes this candidate pattern visible. It moves from partially accessible in the controlled but ecologically simple Installation A, through unattributable in the ecologically rich but diagnostically opaque Installation B, to most accessible in Installation C, where embodied inertia enables attribution but limits the encounter. The framework does not resolve this tension but makes it observable as a hypothesis to be tested.

Two objections must be faced. First, the pattern could be an artefact of attribution methods rather than a structural property. Causal-inference techniques might partially resolve attribution in distributed configurations like Installation B, and future iterations should incorporate them (Pearl 2009). Even so, the framework predicts a residual relationship, since the integrated signal is a relational product of the configuration, not a sum of decomposable sources. Second, and more fundamentally, the gradient may be designed-in rather than discovered as the installations vary on many axes at once and were built in the sequence of simple (A), distributed (B) and constrained (C). So the ordering could follow from our design choices and post hoc readings. Three heterogeneous cases cannot separate a designed-in gradient from a structural one. We therefore advance the inverse-legibility relation as a framework-derived hypothesis, not a demonstrated finding, whose proper test is controlled investigation that varies encounter richness while holding other factors fixed.

Persistence remains inaccessible at festival timescales, marking the framework's temporal horizon and defining the central challenge for future longitudinal investigation.

Material precondition: Transduction — confirmed in all three installations			
Diagnostic Criterion	**Installation A**	**Installation B**	**Installation C**
Selectivity	Accessible	Partially accessible	Accessible (community)
Appropriation	Partially accessible	Unattributable	Most accessible
Persistence (*horizon*)	Inaccessible	Inaccessible	Inaccessible

Table 2 What the diagnostic apparatus could bring into view across three practice settings. "Accessible/inaccessible" denotes how readable each criterion was given the installation's configuration, not the strength of empirical evidence. Transduction is reported as a uniformly tracked precondition. The three criteria vary in legibility, with the inverse-legibility pattern most visible in the appropriation row.

5 Discussion: what becomes visible when evaluation shifts

5.1 AI as temporal infrastructure

The theoretical framework initially positioned AI as a semiotic transducer (Section 2.4), converting perturbations across material substrates. However, the installations indicated a more specific capacity that the transducer concept did not anticipate, namely, temporal mediation. In all three installations, the computational entities operated across timescales that no single biological or human entity could natively bridge. Computational adjustments occurred in milliseconds, while microbial metabolic responses unfolded over minutes to hours, and community-level diversity dynamics spanned days. We advance the framework-derived hypothesis that the computational system may not merely respond to biological feedback but maintain coherence across these incommensurable scales, accumulating patterns that connect fast computational cycles to slow biological processes.

Installation B made this capacity most visible. Over five days of continuous operation, in situ monitoring suggested that the PPO algorithm's acoustic output shifted in conjunction with slow-moving metabolic trends imperceptible to visitors, who encountered only momentary snapshots of the installation. Because algorithmic logs were not retained, we cannot reconstruct this trajectory post hoc and do not offer it as evidence. Rather, we report it as a practice-derived observation that motivated a core conceptual proposition: AI's distinctive contribution to transsemiotic exchange may lie not in signal generation per se, but in functioning as a temporal infrastructure. This capacity provides the very physical architecture required to sustain the recursive attunement later developed in Installation C.

We borrow the concept of infrastructure from Star's foundational analysis (Star 1999), which emphasised that the infrastructural character of a system becomes visible primarily upon breakdown. While Star focused on spatial and institutional arrangements, we extend this lens to the coordination of disparate timescales. Bratton's concept of planetary computation offers a complementary frame, positioning AI systems as infrastructural layers that mediate between scales of organisation (Bratton 2016). Under this hypothesis, our installations instantiate that scalar mediation concretely. This temporal infrastructure does not constitute an Umwelt in the Uexküllian sense, because the computational entity lacks the normative orientation that makes a world of significance a world *for* an organism. Instead, it functions purely as the scaffolding holding the encounter together across its vastly different metabolic and computational horizons.

Because this infrastructural contribution unfolds over days through accumulated micro-adjustments, it remains invisible to standard, human-centric evaluation paradigms that score AI outputs as isolated moments. To capture temporal mediation, evaluation must itself become temporally distributed. Concretely, rather than relying on single-moment benchmarking, practitioners would measure the coupling between fast computational cycles and slow biological shifts via continuous, multi-timescale logging. Furthermore, to test this infrastructural hypothesis directly, researchers could

temporarily freeze the AI's mediating function to observe whether slow-scale biological coherence subsequently degrades. As Star noted, it is precisely upon this kind of breakdown that infrastructure reveals itself. The gap generalises beyond the transsemiotic case. In ecological monitoring, a system mediating between high-frequency sensors and slow ecosystem dynamics could surface early-warning signals of regime shifts that periodic sampling misses (Scheffer et al. 2009).

5.2 Semiotic plasticity against embodied inertia

The trained vocalist's embodied inertia to microbially driven perturbations in Installation C exposed the implication that skill can foreclose encounter. The singer's years of vocal training had produced a robust and stable instrument optimised for musical performance, which is precisely what conventional evaluation would reward. Yet, in the transsemiotic context, this robustness functioned as a structural obstacle. The trained voice resisted the perturbations the recursive attunement required by overpowering computationally suggested modulations with the inertia of embodied technique. The untrained participants in earlier installations, by contrast, adapted more readily to novel acoustic configurations. Their vocal plasticity enabled the recursive attunement that the singer's expertise inhibited. Whether the operative variable is neuromotor rigidity or attentional mode remains underdetermined. What the contrast suggests is that some configurations of embodied skill impede the porosity that transsemiotic coupling requires.

This implies that effective transsemiotic systems may require deliberate destabilisation rather than optimised robustness. Systems designed for semiotic plasticity and sensitivity to perturbation from heterogeneous alien domains may be better suited to transsemiotic exchange. In computational design terms, the most transsemiotically productive entities may be those with high exploratory variance and low convergence pressure. The generative principle of drift acquires profound empirical grounding through the singer's resistance. Without plasticity, drift is impossible, and without drift, the transsemiotic encounter collapses into imposition.

5.3 The biopolitics of transsemiotic encounter

Crawford's analysis of AI as fundamentally extractive applies with particular force to interspecies AI (Crawford 2021). The very techniques that enable transsemiotic encounter could be repurposed for extraction. An algorithmic system that learns which acoustic frequencies stimulate microbial vitality could power sonic bioreactors optimising microbial labour. Similarly, a computational entity that discovers decay-inducing patterns could inform targeted acoustic manipulation of biological communities (e.g., through acoustic cavitation (Patange et al. 2021)). These represent direct applications of our methods, reoriented from encounter to control. Because acoustic systems capable of inducing community decay raise significant dual-use concerns, any programme extending this work bears a strict governance obligation to conduct biosecurity reviews alongside semiotic investigation.

What distinguishes artistic from industrial deployment is not the technology, but its orientation. The generative principles do not encode pre-programmed ethical

rules, but rather emerge directly from the material commitment to an open-ended encounter. Semiotic probing explores acoustic configurations without a predetermined target. Drift resists premature mathematical or biological convergence. Recursive attunement extends biological-computational coupling across durations that exceed industrial patience. An industrial application would invert each principle to force targeted optimisation of biological states.

Couldry and Mejias's analysis of data colonialism and Liboiron's anticolonial-science critique pose a more fundamental challenge. Couldry and Mejias trace how the extraction of data from human and ecological domains constitutes a new phase of colonial appropriation (Couldry and Mejias 2019). Liboiron sharpens the challenge by arguing that scientific knowledge-making itself proceeds through land relations, and that making non-human worlds knowable to human institutions has historically preceded and enabled their exploitation (Liboiron 2021). Our framework, however reflexively positioned, sits inside both critiques rather than outside them. The framework is itself an instrument of knowability: the sensor resolutions, temporal windows, and mathematical equivalencies established earlier (Section 3.3) act as anthropogenic cuts that dictate how the microbial response can be seen. We do not claim to dissolve this tension. The inverse-legibility hypothesis offers a partial friction in this regard. The distributed, temporally extended, multi-entity configurations where transsemiotic encounter flourishes also frustrate the controlled replicability that industrial applications demand. The framework's diagnostic, rather than productive, orientation offers resources for maintaining the distinction, but it cannot prevent future appropriations from instrumentalising its methods. More fundamentally, extraction requires converting normative selectivity into a controllable input, and this conversion ruptures the encounter it seeks to exploit. Whether encounter or extraction prevails depends not on the framework itself, but on the orientation governing its deployment.

5.4 The asymmetry of normative and parametric selectivity

As developed in Section 2.4, microbial selectivity is oriented by biological imperatives constitutive of the organism's continued existence, while algorithmic selectivity is oriented by an externally specified reward function carrying no consequence for the RL agent itself. This asymmetry structures the gap between normative and parametric selectivity. Whether a signal reaches an organism is a question about materiality, but whether it enters that organism's interpretive world is a question about normativity. The installations indicate that this distinction is not merely definitional but operationally consequential. Across all three installations, the organism's selectivity shifted in response to fluctuating metabolic conditions, whereas the algorithm's selectivity remained functionally invariant, governed strictly by the extrinsic logic of its reward signal. In biological subjects, interpretive history is sedimented in molecular traces that bear directly upon its continued viability. Conversely, the algorithm's history is accumulated in weights that, while complex, bear on nothing beyond the optimisation of the next policy update. This operational divergence is not a contingent limitation of current AI architectures but a structural feature of parametric systems. This is because, in a biosemiotic sense, normativity requires that the act of selecting be constitutive of the selector's own continued existence.

This connects to active debates about AI agents' moral status and computational autonomy. The question is not whether AI systems have interests, because parametric systems can easily simulate goal-directed behaviour. The question is whether their selectivity is existential. A computational entity whose selectivity dissolves when its reward function is removed lacks the kind of self-maintained orientation that, within biosemiotic accounts, underwrites intrinsic interests (Sharov and Vehkavaara 2015; Thompson 2007). This also illuminates an alignment problem. Reward functions approximating biological norms cannot constitute them, which makes specification gaming philosophically inevitable rather than merely technically challenging. The grounding debate introduced in Section 2.3 therefore takes on a sharper form. Distributing grounding across ontologically distinct entities reveals that their selectivity types are incommensurable. The organism grounds through norms, while the algorithm grounds through parameters. No shared representational space can bridge this gap while preserving the distinction that makes transsemiotic exchange non-trivial. The apparatus mediates across these incommensurable domains without reconciling them.

The differentiation between normative and parametric selectivity also clarifies the framework's relationship to protosemiosis. Sharov defines protosemiosis as goal-directedness without object representation that is constitutive of the organism's autonomous functioning. The organism's capacity to maintain itself as a far-from-equilibrium system generates the normative orientation that makes its responses proto-semiotic rather than merely mechanical (Sharov and Vehkavaara 2015). The critical test is whether goal-directedness persists when external specification is removed. Conventional reward-driven systems fail this test. Their objectives dissolve with their reward functions, because nothing in the system's organisation generates goals autonomously. Within the diagnostic apparatus developed here, this is why the criteria presuppose a living interpreter and why the framework is not trivially satisfiable in both directions. Philosophers of technology increasingly argue that autonomous computational architectures are blurring the boundary between programmed execution and normative action. Viewed through frameworks ranging from informational state-transitions (Floridi and Sanders 2004) to relational phenomenology (Coeckelbergh 2020), these systems dismantle the assumption that goal-directed agency requires a biological author. Two considerations preserve the distinction's load-bearing function here. First, mathematical weights persisting beyond a specific reward function are not equivalent to the autonomous generation of a new normative orientation. Transfer-learned policies still inherit their parametric orientation from the original specification, however indirectly. Second, the diagnostic apparatus developed here is concerned not with whether the parametric/normative line is bridgeable in principle but with whether, in the specific assemblages we have studied, the computational entity's selectivity is constituted by its own continued existence. The framework, therefore, makes no in-principle claim about self-grounding computation. It claims only that the present installations do not instantiate it.

The diagnostic orientation does not, however, render the framework unfalsifiable. Two conditions would require its fundamental revision. First, the semiotic interpretation would be refuted if organisms satisfying all three criteria were subsequently

shown to exhibit purely mechanistic responses. This would require mechanotransduction pathway analysis sufficient to account for the observed responses without recourse to the organism's interpretive history or normative orientation. Second, the framework's strongest criterion would be empirically empty if persistence were never observed, even under longitudinal conditions specifically designed to detect it.

6 Conclusion: reorienting AI evaluation for multispecies futures

The concept of transsemiotic exchange has surfaced a domain of significance that existing AI evaluation paradigms exclude, because their ontological commitments render non-human significance invisible by construction. Our diagnostic framework has indicated that investigation of this domain is possible, if temporally demanding and epistemically constrained. Across the three installations, selectivity and appropriation proved partially accessible, while persistence remained a horizon at the constrained timescales of temporary exhibitions, rather than an operationally readable criterion. The framework does not verify meaning but produces the conditions under which evidence for meaning-like processes becomes accessible or stubbornly elusive. The map of differential accessibility across three installations has indicated that the depth of transsemiotic encounter and the clarity with which it can be diagnosed may pull in opposite directions. We advance this inverse-legibility relation as a hypothesis warranting confirmatory investigation, rather than as an established structural property to be overcome.

The anthropogenic paradox flagged at the outset (that any framework for non-human significance is necessarily built from human-derived concepts) has not dissolved over the course of the argument. Every diagnostic criterion, every generative principle, and every accessibility judgement reported above was formulated in human-derived concepts applied to processes whose semiotic structure exceeds the descriptive reach of those concepts. The installations do not resolve this paradox but demonstrate that it is a constitutive feature of transsemiotic enquiry rather than a methodological failure to be corrected.

Dominant paradigms evaluate AI against human benchmarks, asking whether machines can perform tasks that humans value in ways that humans recognise. Our Transsemiotic Turing Approach proposes a different question by asking how computational entities can participate in meaning-making that exceeds human evaluation. This is not a call to abandon human-centred concerns, but to recognise that algorithmic systems also operate within ecologies where significance emerges through relations humans cannot fully observe, at timescales humans cannot natively bridge, and through semiotic processes whose effects are materially real but observationally inaccessible. Evaluation that restricts itself to the human semiotic horizon will miss what may be one of AI's most consequential effects.

The installations illustrate what Haraway's response-ability (the capacity to respond across more-than-human relations (Haraway 2016)) looks like when computational systems participate in its conditions. Visitors experienced that their speech affects microbial vitality. Distributed generative participation dissolved isolated action

into a collective encounter. A professional singer's embodied inertia suggested the structural constraints that the absence of plasticity generates. Response-ability was not a property of individual entities, but an achievement of specific material-semiotic configurations. Such configurations require at least one participant whose selectivity is normatively grounded, which explains why computational systems can participate in these conditions without constituting them on their own. Tsing's entanglements in capitalism's ruins (Tsing 2015), van Dooren's attention to extinction through the unravelling of multispecies relations (Van Dooren 2014), and Haraway's sympoiesis are not invoked as new authorities, but as the intellectual tradition to which this paper contributes empirically. Testing these generative principles with other semiotic modalities marks the next stage of this research programme.

In all three cases, AI may have operated as temporal infrastructure across biological and computational scales, accumulating patterns in durations no biological or human participant could natively bridge. We advance this as a framework-derived hypothesis warranting controlled longitudinal investigation. Evaluation operating at a single moment cannot register this contribution. As algorithmic systems become infrastructural more broadly, mediating flows of energy, information, and matter across scales that exceed human perception, the question of what they mean *to whom* becomes urgent. The Transsemiotic Turing Approach specifies the conditions under which answers might emerge from domains our evaluation paradigms have not yet learned to register.

Acknowledgements. All of the oral microbiota involved were collected from donors or performers who gave full informed consent. Funding details were redacted for the double-blind review process. The authors used Claude (Anthropic) as an editorial collaborator across multiple revision rounds of this manuscript. The model assisted with sentence-level rephrasing, terminology consistency checks, identification of redundancy, and editorial feedback on revision decisions. The framework's coined concepts, the empirical work (three installations 2022–2024), the case selections, the theoretical claims, and the argumentative structure are the authors' own. Final authorial judgement governed every editorial decision. The authors take full responsibility for the manuscript.

References

Aggio RBM, Obolonkin V, Villas-Bôas SG (2012) Sonic vibration affects the metabolism of yeast cells growing in liquid culture: a metabolomic study. Metabolomics 8(4):670–678. https://doi.org/10.1007/s11306-011-0360-x

Alexander VN (2013) Creativity: Self-referential mistaking, not negating. Biosemiotics 6(2):253–272. https://doi.org/10.1007/s12304-012-9158-0

Anonymous (2025) Title omitted for double-blind review. bioRxiv (preprint) Full reference details will be provided upon acceptance.

Barad K (2007) Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning. Duke University Press, Durham, https://doi.org/10.1215/9780822388128

Bateson G (2000) Steps to an Ecology of Mind: Collected Essays in Anthropology, Psychiatry, Evolution, and Epistemology, university of chicago press ed. edn. University of Chicago Press, Chicago, IL, USA, URL https://press.uchicago.edu/ucp/books/book/chicago/S/bo3620295.html, originally published in 1972 by Chandler Pub. Co.

Bender EM, Gebru T, McMillan-Major A, et al (2021) On the dangers of stochastic parrots: Can language models be too big? In: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. Association for Computing Machinery, New York, NY, USA, FAccT '21, pp 610–623, https://doi.org/10.1145/3442188.3445922

Bickhard MH (2009) The interactivist model. Synthese 166(3):547–591. https://doi.org/10.1007/s11229-008-9375-x

Bisk Y, Holtzman A, Thomason J, et al (2020) Experience grounds language. URL https://arxiv.org/abs/2004.10151, arXiv:2004.10151

Borgdorff H (2012) The Conflict of the Faculties: Perspectives on Artistic Research and Academia. Amsterdam University Press, https://doi.org/10.26530/OAPEN_595042

Borsos Z, Marinier R, Vincent D, et al (2023) AudioLM: A language modeling approach to audio generation. IEEE/ACM Trans Audio, Speech and Lang Proc 31:2523–2533. https://doi.org/10.1109/TASLP.2023.3288409

Bratton BH (2016) The Stack: On Software and Sovereignty. The MIT Press, https://doi.org/10.7551/mitpress/9780262029575.001.0001

Canguilhem G (1991) The Normal and the Pathological. Zone Books, New York, translated by Carolyn R. Fawcett, in collaboration with Robert S. Cohen; with an introduction by Michel Foucault

Caudron F, Barral Y (2013) A super-assembly of Whi3 encodes memory of deceptive encounters by single cells during yeast courtship. Cell 155(6):1244–1257. https://doi.org/10.1016/j.cell.2013.10.046

Coeckelbergh M (2020) AI for climate: freedom, justice, and other ethical and political challenges. AI and Ethics 1(1):67–72. https://doi.org/10.1007/s43681-020-00007-2

Couldry N, Mejias UA (2019) The costs of connection : how data is colonizing human life and appropriating it for capitalism, 1st edn. Culture and Economic Life, Stanford University Press, Stanford, California, https://doi.org/10.1515/9781503609754

Crawford K (2021) The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence. Yale University Press, New Haven, https://doi.org/10.2307/j.ctv1ghv45t

Deacon TW (2012) Incomplete nature : how mind emerged from matter, 1st edn. Norton, New York

Dreyfus HL (1992) What computers still can't do : a critique of artificial reason. MIT Press, Cambridge, Mass.

Floridi L, Sanders J (2004) On the morality of artificial agents. Minds and Machines 14(3):349–379. https://doi.org/10.1023/B:MIND.0000035461.63578.9d

Floridi L, Taddeo M, Turilli M (2009) Turing's imitation game: Still an impossible challenge for all machines and some judges—an evaluation of the 2008 Loebner contest. Minds and machines (Dordrecht) 19(1):145–150. https://doi.org/10.1007/s11023-008-9130-6

Gorlée DL (1997) Intercode translation: Words and music in opera. Target : international journal of translation studies 9(2):234–269. https://doi.org/10.1075/target.9.2.03gor

Haraway DJ (2016) Staying with the Trouble: Making Kin in the Chthulucene. Duke University Press, Durham, https://doi.org/10.1215/9780822373780

Haugeland J (1989) Artificial Intelligence: The Very Idea. The MIT Press, https://doi.org/10.7551/mitpress/1170.001.0001

Hirskyj-Douglas I, Pons P, Read JC, et al (2018) Seven years after the manifesto: Literature review and research directions for technologies in animal computer interaction. Multimodal Technologies and Interaction 2(2). https://doi.org/10.3390/mti2020030

Hoffmeyer J (2008) Biosemiotics: An Examination into the Signs of Life and the Life of Signs, 1st edn. Approaches to Postmodernity, University of Chicago Press, United States

International Phonetic Association (1997) Publication of the handbook of the international phonetic association. Journal of the International Phonetic Association 27(1–2):91–91. https://doi.org/10.1017/S002510030000548X

International Phonetic Association (2026) International phonetic alphabet in ipa kiel: Vowels. Available under a Creative Commons Attribution-ShareAlike 4.0 International Licence, URL https://www.internationalphoneticassociation.org/IPAcharts/IPA_charts_E/IPA_charts_E_img.html, copyright © 2026 International Phonetic Association. Accessed: 28 May 2026.

Jakobson R (1959) On linguistic aspects of translation. In: Brower RA (ed) On Translation. Harvard University Press, Cambridge, MA and London, England, p 232–239, https://doi.org/10.4159/harvard.9780674731615.c18

Kleisner K (2011) Perceive, co-opt, modify, and live! organism as a centre of experience. Biosemiotics 4(2):223–241. https://doi.org/10.1007/s12304-010-9093-x

Kohn E (2013) How Forests Think: Toward an Anthropology Beyond the Human. University of California Press, https://doi.org/10.1525/california/9780520276109.001.0001

Kull K (2009) Vegetative, animal, and cultural semiosis: The semiotic threshold zones. Cognitive Semiotics 4(Supplement):8–27. https://doi.org/10.1515/cogsem.2009.4.spring2009.8

Liang P, Bommasani R, Lee T, et al (2023) Holistic evaluation of language models. URL https://arxiv.org/abs/2211.09110, arXiv:2211.09110

Liboiron M (2021) Pollution is colonialism. Duke University Press, Durham London, https://doi.org/10.2307/j.ctv1jhvnk1

Lin A (2015) Egalitarian Bi/multilingualism and Trans-semiotizing in a Global World, John Wiley & Sons, Ltd, chap 2, pp 17–37. https://doi.org/10.1002/9781118533406.ch2

Liu J, Byrne D, Devendorf L (2018) Design for collaborative survival: An inquiry into human-fungi relationships. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, CHI '18, pp 1–13, https://doi.org/10.1145/3173574.3173614

Lyon P (2015) The cognitive cell: bacterial behavior reconsidered. Frontiers in microbiology 6:264. https://doi.org/10.3389/fmicb.2015.00264

Mancini C (2011) Animal-computer interaction: a manifesto. Interactions 18(4):69–73. https://doi.org/10.1145/1978822.1978836

Mancini C, Lawson S, Juhlin O (2017) Animal-computer interaction: The emergence of a discipline. International Journal of Human-Computer Studies 98:129–134. https://doi.org/10.1016/j.ijhcs.2016.10.003

Maran T, Kleisner K (2010) Towards an evolutionary biosemiotics: Semiotic selection and semiotic co-option. Biosemiotics 3(2):189–200. https://doi.org/10.1007/s12304-010-9087-8

Pasquinelli M (2023) The eye of the master : a social history of artificial intelligence. Verso, London New York

Patange AD, Simpson JC, Curtin JF, et al (2021) Inactivation efficacy of atmospheric air plasma and airborne acoustic ultrasound against bacterial biofilms. Scientific reports 11(1):2346–2360. https://doi.org/10.1038/s41598-021-81977-z

Pathak D, Agrawal P, Efros AA, et al (2017) Curiosity-driven exploration by self-supervised prediction. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp 488–489, https://doi.org/10.1109/CVPRW.2017.70

Pearl J (2009) Causality: Models, Reasoning and Inference, 2nd edn. Cambridge University Press, USA, https://doi.org/10.1017/cbo9780511803161

Peirce CS, Hartshorne C (1931) Collected papers of Charles Sanders Peirce. 1, Principles of philosophy. The Belknap Press of Harvard Univ. Press, Cambridge, Mass.

Raji ID, Bender EM, Paullada A, et al (2021) AI and the everything in the whole wide world benchmark. URL https://arxiv.org/abs/2111.15366, arXiv:2111.15366

Rolnick D, Donti PL, Kaack LH, et al (2022) Tackling climate change with machine learning. ACM Comput Surv 55(2). https://doi.org/10.1145/3485128

Sarvaiya N, Kothari V (2015) Effect of audible sound in form of music on microbial growth and production of certain important metabolites. Microbiology (New York) 84(2):227–235. https://doi.org/10.1134/S0026261715020125

Scheffer M, Bascompte J, Brock WA, et al (2009) Early-warning signals for critical transitions. Nature (London) 461(7260):53–59. https://doi.org/10.1038/nature08227

Schlangen D (2023) On general language understanding. URL https://arxiv.org/abs/2310.18038, arXiv:2310.18038

Schmidhuber J (2010) Formal theory of creativity, fun, and intrinsic motivation (1990–2010). IEEE Transactions on Autonomous Mental Development 2(3):230–247. https://doi.org/10.1109/TAMD.2010.2056368

Schulman J, Wolski F, Dhariwal P, et al (2017) Proximal policy optimization algorithms. https://doi.org/10.48550/arXiv.1707.06347, arXiv:1707.06347

Sebeok TA (1968) Zoosemiotics. American Speech 43(2):142–144. https://doi.org/10.2307/454548

Sharov A (2010) Functional information: Towards synthesis of biosemiotics and cybernetics. Entropy 12(5):1050–1070. https://doi.org/10.3390/e12051050

Sharov A, Vehkavaara T (2015) Protosemiosis: Agency with reduced representation capacity. Biosemiotics 8(1):103–123. https://doi.org/10.1007/s12304-014-9219-7

Srivastava A, Rastogi A, Rao A, et al (2023) Beyond the imitation game: quantifying and extrapolating the capabilities of language models. Transactions on Machine Learning Research 2023(5):1–95

Star SL (1999) The ethnography of infrastructure. American Behavioral Scientist 43(3):377–391. https://doi.org/10.1177/00027649921955326

Stjernfelt F (2014) Natural Propositions: The Actuality of Peirce's Doctrine of Dicisigns. Docent Press

Thapen N (2017) Pink trombone. URL https://dood.al/pinktrombone/

Thompson E (2007) Mind in Life, 1st edn. Harvard University Press, Cambridge

Tsing AL (2015) The Mushroom at the End of the World: On the Possibility of Life in Capitalist Ruins, 1st edn. Princeton University Press, Princeton

Turing A (1950) I.—computing machinery and intelligence. Mind LIX(236):433–460. https://doi.org/10.1093/mind/LIX.236.433

von Uexküll J (2010) A Foray into the Worlds of Animals and Humans: With a Theory of Meaning. University of Minnesota Press, Minneapolis

Van Dooren T (2014) Flight ways : life and loss at the edge of extinction. Critical perspectives on animals, Columbia University Press, New York

Wakkary R (2021) Things we could design : for more than human-centered worlds. Design thinking, design theory, The MIT Press, Cambridge, Massachusetts, https://doi.org/10.7551/mitpress/13649.001.0001

Wiener N (1961) Cybernetics : or control and communication in the animal and the machine, 2nd edn. M.I.T. Press, Cambridge, Mass.