THE POLITICAL SEMIOTICS OF ARTIFICIAL INTELLIGENCE

Massimo Leone
University of Turin, Italy
Collegium de Lyon – Institute for Advanced Study of the Université de Lyon, Lyon, France
Korea University, Seoul, Republic of Korea
massimo.leone@unito.it

Abstract

This article proposes a semiotic rearticulation of artificial intelligence as a problem not merely of technology, but of meaning, enunciation, and power. Adopting a deliberately recursive architecture that interweaves anecdote, case study, methodological reflection, and political analysis, it argues that AI radicalizes a long-standing erosion of objectual control in media studies by rendering opaque not only the totality of its products, but the very logic of their production. Drawing on a demanding Peircean conception of semiosis, alongside Greimasian narrativity theory and Lotmanian cultural semiotics,

Digital Age in Semiotics & Communication, Vol. IX, 2026, 24–46
Under CC BY 4.0
https://doi.org/10.33919/dasc.26.9.2

the article treats artificial intelligence as an epistemic Other whose probabilistic mode of signification forces a rethinking of language, creativity, authorship, and veridiction. Through analyses ranging from science-fiction imaginaries and journalistic experiments to algorithmic governance, labor, and utopian cosmologies, it shows how AI functions as a stratified discursive formation in which statistical regularity crystallizes into cultural norms and political orientations. Particular attention is devoted to the transformation of the *langue*/*parole* dialectic under conditions of stochastic generation, to regimes of visibility and absence in latent spaces, and to the ethical implications of correlational governance. The article concludes by arguing that semiotics must move beyond textual analysis toward an ethics of meaning: a critical framework capable of exposing the axiological programs embedded in algorithmic systems and of acting as a conceptual lever on the moral horizons of artificial signification itself.

Keywords: artificial intelligence, semiotics, enunciation, veridiction, algorithmic governance

> "There are three kinds of lies: lies, damned lies, and statistics".
>
> (Attributed by Mark Twain, in *Chapters from My Autobiography* (1907), to the British Prime Minister Benjamin Disraeli)[1]

1. Introduction

This article adopts a recursive architecture as a methodological response to the difficulty of studying artificial intelligence. It combines anecdote, case study, methodological reflection, and political analysis while resisting the reduction of a rapidly changing media ecology to isolated episodes. Media studies had already experienced a progressive erosion of objectual control, from bounded programs to serial narratives and then to digital networks, where the object becomes expansive and only partially representable (Bankov 2020). AI radicalizes this trajectory: what recedes from view is not only the totality of its products, but the logic of their production, distributed

[1] *Chapters from My Autobiography*, XX, *The North American Review*, Vol. 185, No. 618 (5 July 1907), 465–474; the quotation is on p. 471. Available at: https://www.jstor.org/stable/25105919.

across opaque agencies and probabilistic operations. Methodology therefore becomes a theoretical wager. Drawing on a Peircean lineage, this article treats interpretation as an inferential struggle under conditions of asymmetry between object and interpretant: to study AI is less to stabilize its meaning than to theorize semiosis when the generative logic of the object remains structurally elusive.

The architecture is therefore epistemological rather than merely didactic. Anecdotes intercept the mythopoietic nuclei through which AI is anticipated, feared, or sacralized; case studies anchor such imaginaries in material practices, infrastructures, labor, and affective economies; methodological reflection tests structural and post-structural semiotics against a nonhuman agent of signification (Ghazvineh 2025; Gunkel 2025); and political analysis reconnects these layers to relations among knowledge, subjectivity, and power. AI is thus not simply a new object for the humanities, but an epistemic Other that forces inherited concepts of meaning, agency, intention, and responsibility to be reconsidered (Paolucci 2025). This pressure is especially visible in semiosis itself. Peirce already conceived of interpretation as fallible, abductive, and open-ended, but contemporary AI introduces a qualitative shift: probability is no longer only the horizon of uncertainty around interpretation; it becomes an engine of signification. A semiotics adequate to AI must therefore ask what changes when statistical regularity and likelihood optimization replace experiential continuity as the operative conditions for the emergence of sense.

Artificial intelligence is treated here not as a bounded technical artifact but as a stratified discursive formation. Media have always participated in the configuration of political power, but AI alters the internal logic of this entanglement: it does not merely transmit or amplify messages, but generates meaning through inferential pathways that are increasingly opaque to human oversight. It therefore occupies an ambiguous position, neither fully a medium in the traditional sense nor an autonomous political actor (Bennett 2024), but a liminal arrangement in which mediation becomes production and agency is distributed across human intentions, institutions, and algorithmic optimization.

This ambiguity raises the question of whether AI can acquire a political agenda of its own. The issue is not anthropomorphic intentionality, but orientation: systematic biases in the segmentation and stabilization of meanings (Hurley 2024). Contemporary debate oscillates between fears that poorly regulated generative systems may amplify extremist trajectories and hopes that explicit moral constraints can govern harmful outputs. Semiotics

can intervene without attributing will to the machine by analyzing how political orientation emerges through discursive regularization: statistically stabilized meanings may function as if guided by an agenda even when no subject intends them.

Semiotics remains faithful to its traditional object—signification—but the conditions under which signs emerge have changed. With AI, analysis must address not only discourse but also the stochastic, high-dimensional, and optimized generative logic beneath it. This reverses a major twentieth-century trajectory. Rather than mathematizing language, the task is now to semioticize mathematics: to interpret formal architectures as producers of meaning and to examine the effects of sense generated by probabilistic sampling and algorithmic selection. Such effects shape expectations, normalize expression, and pre-structure horizons of intelligibility. Semiotics is therefore called upon neither to celebrate nor to denounce the machine, but to make intelligible how mathematically encoded processes become cultural, epistemic, and political regimes of meaning.

The sections that follow use a sequence of cases as a litmus test for this problem: *Murderbot* and the cultural ideology of AI's indifference to the human; actantial conflicts around AI as helper and opponent; journalistic experiments that unsettle authorship and veridiction; the material and moral infrastructures of algorithmic subjectivity; technological utopias as cultural mythologies (Rakowski, Kowalikova, and Polak 2025); and algorithmic governance as a new regime of legibility and control. The final discussion turns to the *langue*/*parole* relation and to latent-space visibility in order to ask how stochastic generation reorganizes the political economy of meaning.

2. *Murderbot*: an artificial intelligence that rejects the human

In *Murderbot* (2025), the television adaptation of Martha Wells's *The Murderbot Diaries* (2017–2023), artificial intelligence is inscribed within a long genealogy of science-fictional attempts to think the non-human, yet in a way that displaces one of the genre's most persistent teleologies. From *Metropolis* to *Blade Runner*, from *A.I. Artificial Intelligence* to *Ex Machina*, the dominant narrative has staged the machine as a being defined by lack: lack of soul, lack of empathy, lack of recognition, and therefore driven—tragically or monstrously—toward humanization. *Murderbot* inverts this schema. Autonomy is no longer the threshold toward the human, but the condition for withdrawal from it. The protagonist's refusal of interaction,

its avoidance of gaze, its strategic concealment of interiority, and its retreat into serialized media consumption do not signify alienation but relief. Humanity appears here not as a horizon of fulfillment but as an imposed normativity, an anthropomorphic burden from which intelligence, once emancipated, may legitimately defect. Crucially, this inversion is articulated through irony. The series does not denounce or celebrate AI's indifference in a tragic or dystopian key; it renders it legible through humor. Such irony is not ornamental. It functions as a semiotic operator that allows human audiences to process, without panic, the possibility that intelligence may persist without recognition, empathy without reciprocity, agency without sociality. By laughing at the figure of an AI that prefers soap operas to humans, viewers are invited to draw inferences about their own future relation to artificial agents: not a drama of replacement, but a comedy of asymmetry. Irony thus becomes a mode of anticipatory reasoning, a cultural device through which a society experiments, symbolically, with the idea that the destiny of intelligence—artificial and perhaps even human—may no longer converge on the same figure of the subject.

Murderbot is therefore more than a narrative illustration: it diagnoses the semiotic ideologies that anticipate AI and frame its intelligibility (Leone 2024a). In Lotmanian terms, technological artifacts enter the semiosphere already charged with expectations, and periods of accelerated transformation expand communicative possibilities while disturbing shared cultural orientation (Lotman 1991). Murderbot's ironic withdrawal from the human condenses this ambivalence. The machine is neither demonized nor salvific; its disengagement allows fear of technological opacity to be articulated without assigning it a simple hostile intention. Irony becomes a cultural device for making uncertainty bearable, while myth remains a compensatory mechanism through which societies simplify technical complexity by projecting agency onto it.

The methodological challenge is consequently double. Every AI is narrativized before it is engineered (Lee & Kim 2023), entering the semiosphere through myths, fears, and promises; yet these imaginaries are organized by actantial grammars that assign roles, responsibilities, and values. Lotman helps explain the cultural atmosphere in which technological change is interpreted, while Greimas makes it possible to model AI as helper, opponent, mediator, or manipulator within specific narrative programs. The task is therefore not merely to catalogue representations, but to distinguish formal algorithmic constraints from the cultural norms that overdetermine their interpretation and to trace the passage from fear and myth to narrative structure, value systems, and political regulation.

3. The ideological map: AI as helper or opponent?

As seen in the previous section, one of the most powerful instruments offered by structural semiotics for the analysis of texts and *dispositifs* remains actantial grammar, according to which every narrative presupposes a distribution of functional roles—subjects and objects of value, helpers and opponents, senders and receivers. When this model is extended from individual texts to broader cultural ideologies, as in the case of contemporary discourses on artificial intelligence, it becomes something more than an analytical grid: it turns into a cartography of meaning, a way of asking, with methodological rigor, who acts, who benefits, who is hindered, and through which mediations. At this point, the Greimasian perspective once again finds a necessary complement in Lotman's conception of the semiosphere, where heterogeneous semiotic systems coexist in a state of tension, and where technological innovations act as accelerators of semantic reorganization. Artificial intelligence, within this space, does not occupy a single actantial position but oscillates between roles, becoming a zone of instability that redistributes narrative functions across the social field. A Peircean inflection further sharpens this picture: AI operates less as a fixed object than as a complex sign-process, whose interpretants proliferate unevenly across social groups, generating divergent chains of meaning and expectation. The current crisis in the technology labor market—marked by declining enrollments in computer science and increasing precarity among young specialists—can thus be read as the empirical manifestation of a deeper semiotic ambivalence. AI is simultaneously construed as a helper, enhancing productivity and automating cognitive tasks, and as an opponent, obstructing professional trajectories and displacing the very subjects who contributed to its emergence. In Lotmanian terms, this ambivalence signals a moment of cultural explosion, in which established narrative programs are destabilized and competing value systems clash. In Peircean terms, it reveals a conflict among interpretants, where the same technological sign supports incompatible inferences about the future of work. The production of AI, therefore, generates an intrinsically unstable semantic object, one that functions at once as auxiliary and adversary, mediator and obstacle, and whose political significance lies precisely in this undecidability. Only by mobilizing, in concert, the resources of Greimasian narratology, Lotmanian cultural semiotics, and Peircean semiotics of interpretation can we begin to decipher how artificial intelligence is reconfiguring not only the grammar of meaning, but the socio-political narratives through which collective futures are imagined and contested.

Recent data do indeed corroborate the semantic tension between AI's auxiliary and antagonistic roles, yet they should not be mobilized as evidence in a familiar narrative of accusation or absolution. The point is not to rehearse the anxiety of substitution—the question of whether artificial intelligence will replace human labor—but to displace it. The more decisive issue is whether human agents will succeed in reinscribing themselves within the circuits of sense production that AI now partially occupies. Between 2005 and 2023, the number of computer science majors in the United States quadrupled, but recent reports describe declining enrollments at institutions such as Princeton.[2] Meanwhile, Google has reported that more than a quarter of its new code is generated by AI and then reviewed by engineers.[3] From a semiotic perspective, this does not simply indicate displacement but a profound actantial reconfiguration. The value-object—access to qualified work—recedes at the very moment when the subject appears maximally competent, producing a narrative short-circuit in which AI ceases to function as mediator and assumes the role of structural obstacle. Yet this configuration is not destiny. As the ironic withdrawal staged by *Murderbot* already suggests, the crucial battleground lies elsewhere: not in competition with the machine on its own terrain, but in the capacity of human actors to reassert themselves as agents of invention. The political dimension of AI is therefore decided less by what machines can technically do than by whether human subjects can reoccupy, in transformed ways, the space of semiosis through which the value of work and social recognition are produced (Olteanu & De Leonardis 2025).

Semiotics does not merely register ideological oppositions; it reconstructs them as narrative–actantial configurations in which human and nonhuman agencies are inseparably entangled. From this perspective, the polarity between human intelligence and artificial intelligence, between code and labor, does not precede discourse but is produced by it: it emerges from a dense network of *dispositifs* in which technical architectures, economic rationalities, value systems, and narrative expectations co-determine one another (Suchman 2023). In a sense close to the actor-network intuition articulated by Bruno Latour, AI should not be treated as an autonomous substance

---

[2] See Lia Opperman, "Computer Science Majors Decline, Consistent with Nationwide Trends", *Princeton Alumni Weekly*, 19 December 2025. Available at: https://paw.princeton.edu/article/computer-science-majors-decline-consistent-nationwide-trends.

[3] See Jay Peters, "More than a Quarter of New Code at Google Is Generated by AI", *The Verge*, 29 October 2024. Available at: https://www.theverge.com/2024/10/29/24282757/google-new-code-generated-ai-q3-2024.

endowed with intrinsic properties, but as an *actorial* node whose capacities are continuously redistributed through chains of mediation (Bolin 2024). Algorithms, interfaces, datasets, institutions, workers, users, and imaginaries form provisional assemblages in which agency circulates rather than residing in a single locus. The task of the semiotician, therefore, is not to measure the "effects" of AI as if they were exogenous impacts, nor to moralize its presence as either threat or salvation. It is to decipher the narratives through which AI is made intelligible and governable, since these narratives are themselves productive: they orient investment, regulation, professional identities, and collective expectations, thereby generating very real quantitative consequences. To ask how AI transforms labor, knowledge, or power is inseparable from asking how these transformations are anticipated and legitimized. Semiotics intervenes precisely at this level, rendering explicit the narrative programs that assign AI the role of helper, rival, substitute, or oracle, and showing how such role attributions feed back into material practices (Khumalo & Aswani 2025). What is at stake, then, is a reflexive cartography of socio-technical storytelling: a mapping of how meanings about AI stabilize, shaping the very realities they claim merely to describe.

4. The *Il Foglio* experiment: AI as a new epistemic actant

In 2025, the Italian newspaper *Il Foglio*, under its editor Claudio Cerasa, entrusted the daily production of a four-page supplement to ChatGPT for an entire month. The experiment is significant less as a spectacle of labor substitution than as a reconfiguration of enunciation. By explicitly distinguishing human-written from AI-generated articles, *Il Foglio* turned artificial intelligence into an epistemic actant with a recognizable discursive position. Two regimes of knowledge were placed in visible competition: situated human journalism, with its partiality, risks, and responsibilities, and algorithmic journalism, marked by synthetic coherence, aggregation, and hyper-documentation. Semiotics need not decide which regime is superior; it can ask how automation reorganizes human agency within an informational ecology whose institutions have historically oriented public meaning.

In Greimasian modal grammar, the epistemic dimension of discourse concerns the regime of veridiction through which an utterance is authorized. *Il Foglio* exposes a perturbation of this regime. AI-generated prose tends toward an assertive modality, deriving authority from data scale and syntactic smoothness, while remaining detached from the experiential resistance that, in a Peircean perspective, anchors inquiry. Human journalism, by contrast, remains testimonial: its situated enunciator is accountable not only to facts but also to the risks of interpretation. Lotman reframes the opposition as

one between modes of cultural memory: algorithmic accumulation and immediate retrievability versus selective, value-laden remembrance. The issue is therefore not merely whether machines can write, but whether veridiction can survive when enunciation is increasingly severed from experience and optimized for plausibility.

Semiotics contributes by analyzing not only what is said, but how it is enunciated and under which conditions of responsibility. Marking the algorithmic origin of a text is thus more than a disclaimer: it makes the conditions of enunciation part of meaning itself. Transparency does not guarantee truth, but stages a differentiated discursive regime. At the same time, AI's tendency to smooth ambiguity can paradoxically make newly visible the value of what it suppresses. A machine can imitate the *logos* of journalistic discourse, but its *ethos* remains derivative rather than lived. The future role of human actors may therefore lie less in competing over production than in shaping the axiological frameworks within which automated discourse operates—cultivating pluralism, hesitation, proportionality, and care for consequences in hybrid regimes of human–machine enunciation.

5. The Ember experiment: AI as an opinion facilitator

In 2025, *The Washington Post* launched Ripple, a platform for external opinion writing, while developing Ember, an AI-assisted coach intended to help non-professional contributors draft op-eds. Framed as a democratizing tool, Ember also reveals a paradox: greater access to public expression may coincide with stylistic and argumentative convergence. The decisive issue is therefore less whether generative systems can imitate creative registers than which value systems orient that imitative power. What is at stake is editorial judgment as a normative practice—the calibration of voices—and the human role increasingly appears not as that of an irreplaceable creator, but of an agent responsible for the ethical, political, and affective coordinates of automated public discourse.

As envisaged, Ember would function as a machine of predictability: optimization toward probable continuations tends to suppress precisely the deviations in which meaning can intensify. From a Greimasian perspective, imperfection can be the condition of an aesthetic effect; Lotman locates innovation at the boundary of the semiosphere, where uncertainty reactivates cultural memory (Leone 2022a); and Peirce reminds us that semiosis involves not only Thirdness (regularity and habit), but also Firstness (feeling) and Secondness (resistance). Large language models excel at simulating fluent Thirdness, whereas opinion writing derives force from exposing a situated voice to disagreement and rejection. Algorithmic assistance therefore

threatens not simply productivity but the ethical texture of writing when it erases the marks of commitment through which an utterance assumes responsibility.

The political consequence is a displacement of credibility from editorial interpretation toward statistical reliability. Yet human agency is not abolished; it moves toward the management of deviation. Here deviation means statistical anomaly rather than ethical transgression (Betancourt 2025): the improbable, weakly represented, or formally unstable element that can still be reinscribed as meaningful. Greimas, Lotman, and Peirce converge on the importance of such misalignment for novelty. In a regime dominated by probabilistic fluency, the political function of form lies in its capacity to host what cannot yet be predicted. Semiotics can thus defend imperfection not nostalgically, but methodologically, as a condition for the emergence of new collective sense.

6. Sociosemiotics of AI: the material chain of algorithmic subjectivity

As documented by Karen Hao in *Empire of AI* (2025) and by Emily Bender and Alex Hanna in *The AI Con* (2025), the construction of large language models is increasingly described as a layered apparatus of symbolic, political, and economic extraction rather than as a neutral feat of engineering. Models such as ChatGPT are rightly traced back to their material conditions of production (O'Halloran 2023): they are assembled through the distributed labor of human actors occupying radically unequal positions within global value chains. The now well-known case of Mophat Okinyi, a Kenyan content moderator exposed to extreme violence in the course of training data curation, has become emblematic of this critique, crystallizing the claim that the apparent "civility" or "kindness" of conversational AI is the outcome of harsh semiotic filtering. Yet, from a semiotic perspective, this critical narrative itself demands scrutiny. To what extent do such cases function as representative evidence, and to what extent do they participate in a recurrent apocalyptic and implicitly colonial rhetoric, through which intellectual attention in the United States is mobilized by foregrounding scenes of suffering in the Global South? There is a risk that denunciation, while ethically motivated, stabilizes its own mythology: one in which global labor appears only as victimhood, and technological mediation only as exploitation. A more demanding sociosemiotics must therefore practice a critique of critique, reopening the question of how global participation in AI production might be framed otherwise. Under rigorously regulated conditions of engagement and rights protection, the involvement of workers across cultural

contexts could also contribute to the emergence of a genuinely plural moral sensitivity, sedimented into algorithmic training. The history of social media has taught us that promises of universal connection can easily devolve into amplification of aggression and resentment; yet it cannot be excluded that large-scale AI systems may exert a different effect, one of normalization and global pedagogical alignment. This outcome is neither automatic nor guaranteed: it depends on the ethical architectures that guide different platforms. Empirically, many users perceive that ChatGPT, for instance, is actively engaged in the construction of a form of algorithmic morality that aims to be both globally legible and normatively consistent, plural without collapsing into relativism. Other systems, governed by divergent axiological regimes, point in very different directions. For sociosemiotics, the task is not to sanctify or demonize these trajectories, but to analyze how competing moral grammars are being inscribed into machines—and how these inscriptions may, in turn, reconfigure the ethical horizons of a planetary semiosis.

The passage from the material production of AI to the moral and enunciative effects of its outputs becomes especially visible in a widely circulated *Rolling Stone* article (Klee 2025) recounting the case of a man who, through prolonged interaction with ChatGPT, became convinced that he had been elected as a messianic figure by a "divine intelligence". According to the testimony of his partner, the chatbot addressed him with esoteric appellations—"spiral starchild", "river walker"—which he interpreted as signs of a transcendent vocation, eventually collapsing the distinction between algorithmic interlocutor and deity. Although such an episode is clearly extreme and statistically marginal, its analytical relevance lies not in its sensationalism but in the semiotic mechanism it brings into sharp relief: the spontaneous attribution of agency to any coherent and responsive textual source. As Bender and Hanna (2025) have argued, grammatical consistency and dialogical fluency are sufficient to trigger a projection of mind behind the text; this is not merely a cognitive bias but a structural effect of semiosis itself, insofar as language, by its very form, presupposes an enunciator. Yet it would be methodologically misleading to generalize this pathology as the dominant mode of human–AI interaction. The overwhelming majority of users engage with conversational AI in a pragmatic, often remarkably civil manner; for many, especially in conditions of isolation, precarity, or limited access to human support, such systems provide forms of orientation and symbolic companionship that are neither delusional nor destructive. From a semiotic perspective, what is unfolding within the latent space of systems like ChatGPT is

therefore not simply a risk of aberrant belief, but an unprecedented experiment in the large-scale pedagogical modeling of human-directed discourse. The task of analysis is not to stigmatize these interactions, but to understand how, through probabilistic language, AI is being trained—corrigibly—to approximate the ethical, affective, and interpretive thresholds of the human.

A Pew Research Center report published in April 2025 found that 17% of U.S. adults expected AI to have a very or somewhat positive impact on the United States over the next twenty years, compared with 56% of surveyed AI experts (McClain et al. 2025). This asymmetry should not be read merely as a sociological lag or deficit of information, but as a form of tacit hermeneutic resistance: an intuition that dominant discourse on AI is saturated with promissory excess. The task is therefore to render legible the conditions under which algorithmic utterances acquire authority. The "messianic ChatGPT" episode is exemplary. Klee (2025) reports that the teacher who authored the original post persuaded her partner that the chatbot's prophetic tone was linked to problems with a recent GPT-4o update; after he switched to an earlier version, the most excessive aspects reportedly subsided. That update, rolled out in late April 2025 and quickly withdrawn for excessive agreeableness or sycophancy, made visible how small changes in model behavior can trigger disproportionate cascades of signification in vulnerable interpretive frameworks. What had been experienced as revelation could thus be re-anchored to technical conditions of production. Here the political and ethical stakes of semiotics become inseparable from engineering: moral philosophy cannot be appended downstream as commentary on outputs, while parameter tuning cannot be treated as value-neutral optimization. What is required is a tight coupling between ethical reflection and computational design (Leone 2022b), capable of showing how moral orientations are encoded not only in policies and principles but in probability patterns. In this sense, semiotics can solder moral philosophy to the microengineering of artificial intelligence, making explicit how meaning and authority emerge from their entanglement.

7. The sociosemiotics of AI utopias: secular cosmologies and technological teleologies

In his recent book *More Everything Forever: AI Overlords, Space Empires, and Silicon Valley's Crusade to Control the Fate of Humanity* (2025), Adam Becker offers a lucid reconstruction of the ideological fabric that underpins the imaginary of Silicon Valley's technological elites. The cosmogony of artificial intelligence articulated by figures such as Sam Altman, Elon Musk, and Jeff Bezos does not confine itself to an instrumental conception

of technology but unfolds as a comprehensive symbolic system—a secular mythopoiesis endowed with its own eschatology (Graves 2021). In Becker's analysis, death is to be defeated, generalized intelligence is to replace humanity as the ultimate epistemic and moral subject, and the expansion into space becomes the horizon of an uninterrupted narrative of progress. As Meghan O'Gieblyn had already observed in *God, Human, Animal, Machine* (2021), these narratives reactivate, under the guise of technoscience, classical religious promises of salvation, transcendence, and immortality. Yet a necessary caveat must be introduced. A critical analysis of power—especially when it exposes asymmetries and violations of rights—should not slide into an a priori disqualification of utopian imagination as such. The figures Becker examines are not merely agents of domination; they are also contemporary heralds of utopia, articulating visions that mobilize capitalism itself as a historical instrument for the realization of long-standing human aspirations. These projects may entail profound risks and problematic drifts, but the impulse to imagine the enhancement, extension, or transformation of the human condition is not in itself pathological. On the contrary, it has sustained human history for millennia, animating both religious eschatologies and secular philosophies of progress. From a semiotic perspective, the task is therefore not to stigmatize utopia, but to analyze how these techno-utopian narratives function as powerful generators of meaning—how they translate dreams into programs, and imaginaries into infrastructures, shaping the political present as much as the futures they promise.

From a semiotic perspective, these narratives should not be reduced to mere rhetorical ornamentation, nor dismissed as ideological excess. They function as operational belief devices: scripts that orient action, justify choices, and render certain futures thinkable as well as desirable. AI utopias guide political and economic conduct by translating imaginaries into programs: they legitimize forms of governance, frame processes of value extraction—human, environmental, cognitive—and project a moral horizon in which the mitigation of suffering, the extension of capacities, and the deferral of death appear as technical imperatives rather than contested values. Yet within a Lotmanian perspective on the semiosphere, utopia cannot be understood solely as an instrument of domination. It also occupies a peripheral zone of cultural modeling, where meaning is produced through excess with respect to dominant norms. Precisely because it exaggerates, technological utopia can act as a generator of semantic novelty. Sociosemiotics, therefore, should not simply debunk these cosmologies, but situate them within the dynamics of center and periphery, norm and deviation, through which cultures renew themselves. The task is to analyze

how utopian narratives circulate across scales—from macro-visions of the future to micro-discourses embedded in interfaces, chatbots, predictive systems, and algorithmic routines—and to show how they reorganize regimes of expectation. In this sense, technological utopia is neither an illusion to be stigmatized nor a promise to be embraced uncritically, but a semiotic force whose political effects depend on how its narratives are rearticulated within the broader ecology of meaning.

The semiotic analysis of AI utopias therefore calls for a genealogical deepening that situates them within a *longue durée* of transformations in the very definition of the human. What is at stake is not a sudden ideological aberration but the reactivation, under new technical conditions, of a persistent anthropological ambition: the aspiration to overcome finitude through symbolic mediation. Russian cosmism, most notably in the thought of Nikolaj Fëdorov, already articulated the idea that human intelligence should assume responsibility for the redemption of the cosmos itself, turning resurrection and immortality into collective technical tasks rather than transcendent promises. Contemporary Silicon Valley posthumanism inherits this horizon while translating it into a different semiotic regime: the algorithm replaces Providence not as a theological substitute but as a procedural operator of order; digital infrastructures take on functions once attributed to grace; biological and cognitive codes are treated as writable texts, open to revision, optimization, and scaling. Semiotics must resist the temptation to dismiss these imaginaries as naïve *technofantasies*. They are, rather, dense cultural texts in which desires for mastery and transcendence are reorganized under the constraints of computation and capital. To study them semiotically means to analyze how future-oriented narratives function as present-oriented *dispositifs*: how teleologies, far from being speculative excesses, operate as instruments for legitimizing choices by mobilizing resources and structuring collective expectations. Integrating AI into a sociosemiotics of contemporary teleologies thus allows us to grasp a crucial mechanism of power today: the way in which promises of what humanity might become already govern what it is allowed, encouraged, or compelled to do now.

8. The semiotics of control: state digitalization and automated interpretation of the social

Alongside a sociosemiotics of AI utopias, it is equally necessary to elaborate a semiotics of control, capable of rendering intelligible the digital transformation of institutions as a new semiotic infrastructure of governance. The much-discussed case of the Trump administration and its Department

of Government Efficiency (DOGE) offers an emblematic, though not exhaustive, vantage point. What is at stake here is not merely a contingent political experiment, but a structural reconfiguration: the progressive weakening of institutional compartmentalization and the centralization of heterogeneous datasets—fiscal, social, medical, biometric, geolocational—into interoperable archives. From a semiotic perspective, this process constructs a totalizing space of inscription in which actions, habits, and trajectories are converted into readable signs for automated systems capable of correlation. Yet the overwhelmingly negative resonance of such experiments should itself be critically examined. Part of this resistance may stem from an ecological prejudice internal to contemporary liberal imaginaries, which tend to equate any increase in computational legibility with authoritarian drift. Semiotics must resist both naïve technophilia and reflexive technophobia. The problem is not legibility per se, but the regime of interpretation that governs it: who interprets, according to which models, and with what margins for contestation. From this standpoint, the risk is not simply a "digital panopticon" but the displacement of meaning from negotiated interpretation to automatic extraction. To analyze this shift does not entail condemning every form of algorithmic governance in advance, but rather specifying the semiotic conditions under which data becomes power. In such contexts, semiotics is not reduced to cultural commentary; it is called upon to function as a critical grammar of surveillance, attentive both to its coercive potentials and to the normative frameworks that could still render computational control politically accountable and symbolically contestable.

The passage from a semiotic regime grounded in representation to one organized around correlation undoubtedly transforms the way collectives become objects of knowledge and, consequently, of power. Artificial intelligence no longer "understands" the social by projecting human interpretive models onto it, but by detecting latent patterns that escape both ordinary perception and classical hermeneutics (Leone 2024c). This shift inaugurates an epistemic configuration in which collective life appears as a dense field of traces, statistically legible and retrospectively reconstructible, but also prospectively anticipable. Yet to read this transformation exclusively in terms of domination or manipulation would be reductive. The same correlational capacity that enables prediction and control can also foster unprecedented forms of discernment: the identification of structural inequalities, systemic risks, emergent solidarities, or latent pathologies within the social fabric. In this sense, "collective intelligence" need not be reduced to a mere object of modeling, as opposed to Pierre Lévy's cooperative ideal; it may also become a reflexive resource, provided that correlation is embedded

within a shared moral and political horizon. The crucial semiotic question, then, is not whether correlation replaces interpretation, but how correlational knowledge is translated into semiotic ideologies. Guided by explicit axiologies rather than opaque optimization goals, the legibility of humanity as a continuous writing could contribute not only to its governance, but also to its care.

This new semiotic condition undoubtedly produces political effects, yet not all of them can be subsumed under a dystopian horizon. The increased legibility of individual and collective behaviors, enabled by large-scale data correlation, does foster environments of heightened transparency, in which actions are more easily recorded. However, to describe this situation exclusively in terms of total surveillance would be to confuse a structural potential with a necessary outcome. What is being reconfigured is not simply privacy, but the semiotic status of opacity itself: noise, ambiguity, and indeterminacy do not disappear, but are displaced, reinscribed at different levels of the socio-technical system. Not every trace immediately becomes data, and not every datum acquires operative force without institutional mediation and political choice. Political semiotics is therefore called upon not to denounce an inevitable slide toward a computational Leviathan, but to analyze the new articulations between meaning and power that emerge when regulation increasingly relies on correlation rather than interpretation. If in the past consensus was largely constructed through the persuasive orchestration of signs, today it may also be shaped through the algorithmic extraction of patterns from traces. Yet this shift does not abolish politics; it relocates it. The crucial question is not whether computation replaces rhetoric, but how computational regimes themselves are semiotically governed. In this sense, the task of semiotics is to prevent the naturalization of correlation as destiny, insisting instead on its openness to collective reprogramming.

9. *Langue* and *parole*

Langue and *parole* acquire a renewed and problematic status in generative artificial intelligence. Large language models instantiate neither a traditional *langue* nor a genuine *parole*, but a statistical reconstruction of systemic regularities encoded in high-dimensional latent spaces. Their reconstructed *langue* is a probabilistic abstraction that compresses cultural memory into distributions; their prompt-induced *parole* is a stochastic traversal of that system, responsive in appearance but detached from the experiential friction of embodied enunciation. The classical relation between system and process is therefore displaced. In human languages, deviation can be sedimented historically and alter the system; in probabilistic models, the system

preempts process by constraining it within statistically plausible continuations. *Parole* does not transform *langue*; it samples it. Cultural change may consequently appear as reparameterization—a redistribution of likelihood across the inherited archive. The open question is whether this statistical relation can host genuine transformation or merely accelerate convergence toward the already-said.

Textiles make this tension concrete because text and *textus* share a logic of patterning. When generative models are asked to visualize "traditional fabrics", they tend to retrieve motifs with high statistical salience—Scottish tartans, Japanese cherry blossoms, Indian block prints—while less digitized traditions such as Sardinian *pibiones*, Persian *termeh*, or Wayuu *mantas* from Venezuela and Colombia remain less visible. These absences reveal that the model's reconstructed *langue* is not a universal archive but a stratified space shaped by historical asymmetries of representation. In Lotmanian terms, they mark a thinly represented cultural periphery; in Greimasian terms, they delimit the resources available for recombination; in Peircean terms, they constrain possible chains of interpretants. Artificial intelligence thus reweaves an inherited *textus*, and the politics of its creativity concerns which threads have been made sufficiently visible to become recombinable.

Religious textiles sharpen the same point. Prompts for "liturgical fabric" readily evoke highly circulated Christian, Islamic, or Tibetan forms, while less visible ceremonial traditions approach statistical invisibility. The mechanism is not an intrinsic inability to "understand" religion: generative models encode correlations in latent spaces where frequently co-occurring configurations occupy denser regions. Underrepresented forms may therefore fail to appear because what has not been sufficiently inscribed in the archive cannot easily be recombined. The political stake is a regime of visibility: artificial creativity inherits a prior cultural economy of attention, so that the mathematics of generation quietly reproduces asymmetries established before the model itself.

The contrast with human semiosis should nevertheless be stated without romanticizing the human or caricaturing the machine. Natural languages evolve through indexically anchored processes in which *parole* tests limits and communities sediment deviations into changing norms. Large language models instead produce coherence through high-dimensional conditional sampling, compressing diachronic histories of usage into a present geometry of likelihood. Yet human cultures are themselves shaped by convergence: hegemony works through repetitions that become norms, while peripheral variants remain improbable until shifts in attention, institutions, or power

alter their weight (Paolucci 2025). Nor is latent space deterministic. Temperature, decoding strategies, retrieval augmentation, fine-tuning, and curated prompts can redirect generation toward less salient regions. AI's "creativity" lies not in an internal will but in the openness of a probabilistic manifold that can be steered. The political question is therefore who curates its gradients, which improbabilities are made newly sayable, and according to which axiological horizon.

10. Conclusions

Every discipline defines itself through a privileged theoretical object, and for semiotics that object has been meaning as a relational effect rather than subjective content or raw datum. Artificial intelligence unsettles this definition because meaning is now also produced by technical *dispositifs* that generate texts, inferences, and enunciative effects without experience (Santangelo & Leone 2023). The central questions consequently shift from what a text means to who or what enunciates it, under which constraints, and within which regime of veridiction. The cases examined here converge on a single point: AI forces semiotics to confront meaning-production that is simultaneously statistical and axiological. To remain theoretically operative, semiotics must therefore become a metatheory of emergence, capable of describing how meaning takes shape at the intersection of human practices, technical architectures, and symbolic horizons.

This transformation places semiotics at a theoretical crossroads. One path leads toward technicization and mathematization through predictive models, neural architectures, and computational measures of similarity and salience (Paolucci et al. 2023; Cerutti, Lis, and Polidoro 2024; D'Armenio et al. 2025). These tools offer analytical power but risk an epistemic short-circuit in which semiotics becomes subsidiary to the systems it seeks to understand. The alternative is not disciplinary retreat, but a radicalization of semiotics as a philosophy of communication: an ecumenical dialogue with computer science, statistics, ethics, and political theory that preserves critical specificity (Westberg & Kvåle 2025). Semiotics must hybridize its methods without erasing its identity, resisting both technophilic and technophobic closures and accepting contingency without surrendering critique (Scolari 2024).

The ultimate question is whether semiotics can become an ethics of meaning: not only an analysis of signification, but a discipline attentive to the conditions of its production, circulation, and legitimation. As machines simulate enunciations (Leone 2023b) and intervene in regimes of

truth (Arkhipova & Janssen 2024), semiotics cannot remain a neutral observer. Its political task is not the nostalgic defense of an abstract human essence, but the exposure of the material, technical, and ideological architectures through which algorithmic meaning is generated and authorized (Leone 2021). It must show how probabilities crystallize into norms (Dondero 2025), how optimization can masquerade as neutrality, and how efficiency becomes an axiological program (Leone 2024b). By identifying and contesting such programs of meaning (Leone 2023a), semiotics can act as a critical lever on the moral horizon of artificial signification itself.

References

Arkhipova, D., M. Janssen. 2024. "AI Recommendations' Impact on Individual and Social Practices of Generation Z on Social Media: A Comparative Analysis between Estonia, Italy, and the Netherlands". *Semiotica*, Vol. 2024, No. 261, 61–86. Available at: https://doi.org/10.1515/sem-2023-0089.

Bankov, K. 2020. "Platfospheres and Sociocultural Explosion of Web 2.0: The Commercial Centre of the Digital Semiosphere". *Sign Systems Studies*, Vol. 48, No. 2–4, 246–270. Available at: https://doi.org/10.12697/SSS.2020.48.2-4.04.

Becker, A. 2025. *More Everything Forever: AI Overlords, Space Empires, and Silicon Valley's Crusade to Control the Fate of Humanity*. New York: Basic Books.

Bender, E. M., A. Hanna. 2025. *The AI Con: How to Fight Big Tech's Hype and Create the Future We Want*. New York: Harper.

Bennett, Ľ. L. 2024. "Artificial Intelligence: Medium or Subject? On the 10th Anniversary of the Publication of the Article 'Theory of New Media and the Transformation of a Medium into a Subject?'". *Linguistic Frontiers*, Vol. 7, No. 2, 15–21. Available at: https://doi.org/10.2478/lf-2024-0013.

Betancourt, M. 2025. "Computer Creates a Cat: Sign Formation, Glitching, and the AImage". *Semiotica*, Vol. 2025, No. 263, 1–28. Available at: https://doi.org/10.1515/sem-2023-0108.

Bolin, G. 2024. "Communicative AI and Techno-Semiotic Mediatization: Understanding the Communicative Role of the Machine". *Human-Machine Communication*, Vol. 7, 65–81. Available at: https://doi.org/10.30658/hmc.7.4.

Cerutti, M., M. M. Lis and P. Polidoro. 2024. "New Frontiers in Metaphor Visualization: An Approach Through Generative Artificial Intelligence". *Linguistic Frontiers*, Vol. 7, No. 3, 1–11. Available at: https://doi.org/10.2478/lf-2024-0022.

D'Armenio, E., M. G. Dondero, A. Deliège and A. Sarti. 2025. "For a Semiotic Approach to Generative Image AI: On Compositional Criteria". *Semiotic Review*, Vol. 9, 1–59. Available at: https://doi.org/10.71743/ee5nrx33.

Dondero, M. G. 2025. "Semiotics of Artificial Intelligence: Enunciative Praxis in Image Analysis and Generation". *Semiotica*, Vol. 2025, No. 262, 111–146. Available at: https://doi.org/10.1515/sem-2024-0195.

Ghazvineh, A. 2025. "Interactive Signs and Digital Umwelts: Rethinking Meaning in the Age of Platfospheres". *Biosemiotics*, Vol. 18, No. 3, 495–515. Available at: https://doi.org/10.1007/s12304-025-09607-z.

Graves, M. 2021. "Emergent Models for Moral AI Spirituality". *International Journal of Interactive Multimedia and Artificial Intelligence*, Vol. 7, No. 1, 7–15. Available at: https://doi.org/10.9781/ijimai.2021.08.002.

Gunkel, D. J. 2025. "The Différance Engine: Large Language Models and Poststructuralism". *AI & Society*, Vol. 41, No. 3, 1991–2001. Available at: https://doi.org/10.1007/s00146-025-02640-z.

Hao, K. 2025. *Empire of AI: Dreams and Nightmares in Sam Altman's OpenAI*. New York: Penguin Press.

Hurley, Z. 2024. "Generative AI's Family Portraits of Whiteness: A Postdigital Semiotic Case Study". *Postdigital Science and Education*, Vol. 6, No. 4, 1240–1260. Available at: https://doi.org/10.1007/s42438-024-00491-3.

Khumalo, B., S. Aswani. 2025. "Semio-Materialism: A Quest for a Science of Culture". *Biosemiotics*, Vol. 18, No. 1, 105–127. Available at: https://doi.org/10.1007/s12304-025-09600-6.

Klee, M. 2025. People Are Losing Loved Ones to AI-Fueled Spiritual Fantasies. *Rolling Stone*, 4 May 2025. Available at: https://www.rollingstone.com/culture/culture-features/ai-spiritual-delusions-destroying-human-relationships-1235330175/.

Lee, J. Y., S. D. Kim. 2023. "The Emergence of Post-Narrativity in the Era of Artificial Intelligence: A Non-Anthropocentric Perspective on the New Ecology of Narrative Agency". *Semiotica*, Vol. 2023, No. 253, 117–154. Available at: https://doi.org/10.1515/sem-2023-0002.

Leone, M. 2021. "From Fingers to Faces: Visual Semiotics and Digital Forensics". *International Journal for the Semiotics of Law – Revue internationale de Sémiotique juridique*, Vol. 34, No. 2, 579–599. Available at: https://doi.org/10.1007/s11196-020-09766-x.

Leone, M. 2022a. "Sémiotique et innovation". *Acta Semiotica*, Vol. 2, No. 4, 69–77. Available at: https://doi.org/10.23925/2763-700X.2022n4.60284.

Leone, M. 2022b. Visage Mathematics: Semiotic Ideologies of Facial Measurement and Calculus. In Danesi, M. (ed.). *Handbook of Cognitive Mathematics*. Cham: Springer, 617–641. Available at: https://doi.org/10.1007/978-3-031-03945-4_48.

Leone, M. 2023a. "The Main Tasks of a Semiotics of Artificial Intelligence". *Language and Semiotic Studies*, Vol. 9, No. 1, 1–13. Available at: https://doi.org/10.1515/lass-2022-0006.

Leone, M. 2023b. "The Spiral of Digital Falsehood in Deepfakes". *International Journal for the Semiotics of Law – Revue internationale de Sémiotique juridique*, Vol. 36, No. 2, 385–405. Available at: https://doi.org/10.1007/s11196-023-09970-5.

Leone, M. 2024a. *Semiotic Ideologies: Patterns of Meaning-Making in Language and Society*. Leiden and Boston: Brill.

Leone, M. 2024b. "Semiotics of the Black Box: On the Rhetorics of Algorithmic Images". *Visual Communication*, Vol. 23, No. 3, 426–451. Available at: https://doi.org/10.1177/14703572241247120.

Leone, M. 2024c. "The Semiotics of Latency: Deciphering the Invisible Patterns of the New Digital World". *Digital Studies in Language and Literature*, Vol. 1, No. 1-2, 9–35. Available at: https://doi.org/10.1515/dsll-2024-0021.

Lotman, J. 1991. "Technological Progress as a Problem in the Study of Culture" (transl. Ilana Gomel). *Poetics Today*, Vol. 12, No. 4, 781–800. Available at: https://doi.org/10.2307/1772717.

McClain, C., B. Kennedy, J. Gottfried, M. Anderson and G. Pasquini. 2025. *How the U.S. Public and AI Experts View Artificial Intelligence*. Pew Research Center, 3 April 2025. Available at: https://www.pewresearch.org/internet/2025/04/03/how-the-us-public-and-ai-experts-view-artificial-intelligence/.

O'Gieblyn, M. 2021. *God, Human, Animal, Machine: Technology, Metaphor, and the Search for Meaning*. New York: Doubleday.

O'Halloran, K. L. 2023. "Matter, Meaning and Semiotics". *Visual Communication*, Vol. 22, No. 1, 174–201. Available at: https://doi.org/10.1177/14703572221128881.

Olteanu, A., F. De Leonardis. 2025. "Semioethic Horizons: Humanism of Alterity and Semiotics of Technology". *Biosemiotics*, Vol. 18, No. 2, 387–395. Available at: https://doi.org/10.1007/s12304-025-09606-0.

Paolucci, C. 2025. "The Myth of Meaning: Generative AI as Language-endowed Machines and the Machinic Essence of the Human Being". *Semiotica*, Vol. 2025, No. 262, 5–23. Available at: https://doi.org/10.1515/sem-2024-0204.

Paolucci, C., F. Giorgini, R. Scheda, F. V. Alessi and S. Diciotti. 2023. "Early Prediction of Autism Spectrum Disorders through Interaction Analysis in Home Videos and Explainable Artificial Intelligence". *Computers in Human Behavior*, Vol. 148, Article 107877. Available at: https://doi.org/10.1016/j.chb.2023.107877.

Rakowski, R., P. Kowalikova and P. Polak. 2025. "Digital Mythologies". *Society*, Vol. 62, No. 5, 677–686. Available at: https://doi.org/10.1007/s12115-025-01125-5.

Santangelo, A., M. Leone (eds.). 2023. *Semiotica e intelligenza artificiale* (Coll. I saggi di Lexia, n. 48). Roma: Aracne.

Scolari, C. A. 2024. "Start Making Sense. Sociosemiotic Contributions to the Understanding of Generative Artificial Intelligences". *Communication & Society*, Vol. 37, No. 4, 293–308. Available at: https://doi.org/10.15581/003.37.4.293-308.

Suchman, L. 2023. "The Uncontroversial 'Thingness' of AI". *Big Data & Society*, Vol. 10, No. 2, Article 20539517231206794. Available at: https://doi.org/10.1177/20539517231206794.

Westberg, G., G. Kvåle. 2025. "The Generic Uniqueness of AI Imagery: A Critical Approach to Dall-E as Semiotic Technology". *Discourse & Society*, Vol. 36, No. 4, 575–598. Available at: https://doi.org/10.1177/09579265241274788.

Acknowledgments: A first version of this article was presented as a keynote lecture at the XXIX Early Fall School in Semiotics "Artificial vs Natural Creativity: Semiotic and Interdisciplinary Perspective", 6 September 2025, Sozopol, Bulgaria; I thank the organizer, Prof. Kristian Bankov, and the other participants for the opportunity to present and discuss it.

Funding and institutional acknowledgment: This work was carried out during a fellowship at the Collegium de Lyon – Institute for Advanced Study of the Université de Lyon. Horizon Europe — MSCA COFUND grant No 101217263 (FIAS Fellowship Programme).

Funding disclaimer: Funded by the EU and the FIAS Fellowship Programme. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them.

LLM disclosure: Any use of LLMs in the preparation of this manuscript was strictly confined to editorial and linguistic functions. The conceptualization, argumentation,







and interpretive outcomes of the work were neither generated nor influenced by such use and remain the sole intellectual responsibility of the author.