Creativity as Differential Recombination in Generative AI: A Cognitive Semiotic Framework for Evaluating LLM and GAN

Michela Ravelli
University of Bologna
Department of Philosophy and Communication
Bologna, Italy

ABSTRACT
The study of machine creativity remains one of the most contested and least resolved questions in contemporary AI research. Existing frameworks tend either to exclude generative systems from the domain of creativity on ontological grounds, or to admit them at the cost of definitional vacuity, resting on what Stephensen calls "a highly anthropocentric ontology" that presupposes, rather than argues for, a sharp segregation between the human and the non-human [1]. This paper proposes a third way grounded in cognitive semiotics: creativity as differential recombination within a constrained semiotic space, developed from the convergence of conceptual blending theory [2], Eco's encyclopaedia [3], Lotman's semiosphere [4], and Kristeva's intertextuality [5]. Large Language Models and Generative Adversarial Networks are shown to instantiate this model in architecturally distinct but formally analogous ways. From this framework, four semiotic evaluative dimensions are derived—differential recombination, semiosphere participation, intertextual transformation, and interpretive openness—articulated in terms of sign-processes rather than mental states, offering a tractable, non-metaphysical basis for assessing creativity in human and machine outputs alike.

General Terms
Artificial Intelligence, Creativity, Art, Semiotics, Cognitive Science, Natural Language Processing

Keywords
Computational creativity, conceptual blending, LLMs, GANs, semiotic machines, differential recombination, generative AI, Eco, Fauconnier-Turner, novelty

1. INTRODUCTION
Debates on machine creativity often inherit, more or less explicitly, a much older narrative: one in which original thought—including human thought—is imagined as radically disconnected from what precedes it, a secularised trace of creatio ex nihilo, creation out of nothing. On such a view, genuine novelty could not be a matter of degree or of relation to existing material; it would have to emerge, so to speak, from no material at all. It is plausible that this residual mystique, rather than any well-argued criterion, has contributed to sustaining the assumption—widespread, if not universally endorsed—that human creativity is best understood in terms of intentionality, consciousness, and authentic semantic grounding, treated as the normative standard against which artificial outputs are then measured. This paper does not adopt that assumption, nor does it claim to settle what human creativity ultimately consists in. What can be said with somewhat more confidence is that, under this presupposition, generative systems built on Large Language Models (LLMs) and Generative Adversarial Networks (GANs) tend to be handled in one of two ways, both arguably unsatisfactory: excluded from the domain of creativity on ontological grounds, or admitted into it at the cost of definitional vacuity. Stephensen, for instance, has described the conceptual apparatus underlying this impasse as resting on "a highly anthropocentric ontology based on a sharp segregation between the human and the non-human" [1]—a segregation that appears, at least, to be more often presupposed than argued for.

This paper takes that impasse as its starting point and proposes an alternative grounded in cognitive semiotics rather than in philosophy of mind: creativity as differential recombination within a constrained semiotic space. The model is derived from the convergence of conceptual blending theory, structural semiotics, and the architectures of contemporary generative AI, and it is designed to apply, without ontological adjustment, to both human and machine production. On this basis, the paper further develops four evaluative dimensions—operationalised as sign-processes rather than mental states—capable of assessing any output, human or artificial, along a common semiotic scale.

The remainder of the paper proceeds as follows. Section 3 situates the concept of creativity within its cultural and theoretical genealogy and identifies the limits of existing definitions. Section 4 develops the cognitive-semiotic foundation of differential recombination. Section 5 examines how LLM and GAN architectures instantiate this model. Section 6 operationalises the framework into the four-dimensional Semiotic Profile Vector. Section 6.4 maps the resulting profile taxonomy onto empirical dynamics observed across generative systems, before the paper concludes.

2. METHODOLOGY
This paper is primarily conceptual rather than empirical in character: it proceeds through a comparative synthesis of conceptual blending theory and structural semiotics, applied as an analytical framework to the architectures of contemporary generative AI. A number of methodological coordinates orient this approach.

The first is disciplinary. The analysis adopts a semiotic frame of reference, understood as the science of all mediations that filter the elaboration of meaning beyond biological determination. Considered, after Eco, as an "expanding galaxy" [6], semiotics offers fertile ground both for interrogating the participation of artificial systems in the creative act and for rethinking the creative act itself.

Methodologically, the paper is grounded in cognitive semiotics and in the close analysis of generative architectures, and it develops its argument across five sections, moving from the theoretical genealogy of creativity to its cognitive-semiotic reformulation, its architectural instantiation in LLMs and GANs, and finally its operationalisation into a set of evaluative dimensions.

It should be stated clearly that these evaluative dimensions are proposed at a preliminary, exploratory stage. They are offered as a conceptual scaffold intended to orient future empirical

validation, rather than as an already-tested measurement instrument; the computational and cultural proxies associated with each dimension (Section 6) accordingly require independent empirical calibration before they can support quantitative claims.

3. BACKGROUND AND THEORETICAL CONTEXT: THE CONCEPT OF CREATIVITY

Creativity occupies an unusually unstable position among the concepts of the cognitive sciences. Bartezzaghi describes it as "one of the most bizarre specimens" in the zoo of concepts, a category invoked everywhere and yet rarely subjected to precise scrutiny [7]. This instability extends beyond the lexicographic level, having sedimented through a long cultural history that has invested the concept with layers of mythological significance, rendering it at once indispensable and resistant to analysis.

That genealogy runs through two intertwined traditions. In the Judeo-Christian and Neoplatonic lineage, creation in its full sense belongs to God alone; human making is at best a derivative participation in divine creativity [8, 9]. Rather than dismantling this structure, the Romantic tradition secularised it, shifting the locus of absolute origination from the divine to the human genius. In doing so, it established the creative subject as a sovereign origin of meaning transcending convention and precedent, transforming creativity into a quasi-sacred aura that sets exceptional human agency apart from ordinary labour [10]. Concealed beneath this secular vocabulary, the demand for authenticity and unmediated originality continues to perpetuate the myth of the genius. This frames creativity as an "individual and personal doing" isolated from ordinary labor and socio-cultural structures [11]. Consequently, such a framework proves descriptively inadequate for capturing the full diversity of actual practice—most notably along the contested boundary between creative and non-creative production.

The emergence of generative systems capable of producing texts, images, and other outputs of considerable formal complexity has intensified rather than resolved this difficulty. It raises a question that exceeds the specifics of the underlying technology: how does the human claim to creative primacy hold once that primacy is no longer empirically self-evident? This tension—between taking seriously the generative capacities of artificial systems and refusing to dissolve genuine differences into facile assimilation—defines the theoretical space in which this paper operates.

The debate over whether a machine can be creative has settled around a stable set of dichotomies functioning as devices of ontological demarcation: real/artificial, authentic/simulated, creative/merely computational, meaningful/empty. Stephensen locates the assumption beneath this apparatus in an "unreflected a priori" belief in "a radical difference in the way of being human compared to the way of being of literally any other thing in the cosmos" [1], typically presupposed rather than argued. Bown and McCormack similarly note the tendency to "take as given that humans are the only kind of creative agent we need consider" [?], making the human the unstated criterion against which creativity, intelligence, and meaning are measured.

The tendency has a traceable history. Lovelace's 1842 remark that a computational machine "has no pretension whatever to originate anything" [12] fixed, for well over a century, the image of the machine as a "humble executor" [12]: capable of transcription and calculation, structurally incapable of generating symbolic content. Its basic structure—the machine executes, it does not originate—has proven durable. But contemporary generative systems have ultimately made that position harder to sustain. The aesthetic Turing test—an observer's inability to distinguish an artificial work from a human one on the basis of output alone—is now difficult to contest as a datum: Noll found in 1966 that only 28% of participants correctly identified a computer-generated composition [13], and systems such as GPT-4 and Stable Diffusion now produce outputs Craig and Kerr describe as "indistinguishable at first sight from their human counterparts" [?], in stylistic coherence as well as formal accuracy. Benchmark comparisons of linguistic creativity report machine performance comparable to, and in some cases exceeding, human performance [14]. The argument from technical impossibility has accordingly lost much of its force.

Critical discourse has responded by relocating the objection rather than abandoning it: from the impossibility of the result to the insufficiency of the process that produced it. The question shifts from what is produced to how, and by whom. Its paradigmatic form holds that "though AI can produce content that is beautiful or stylized it is unable to charge its creations with meaning, emotional resonance or even intentionality" [15]: artificial systems are conceded at most a mere novelty [?], set against a truly creative act distinguished by the conditions of its production rather than by its result.

Creativity is thus displaced from the observable to the ontological: from what is created to the nature of the agent that creates it [16]. The resulting hierarchy—an authentic creativity reserved to human action, a merely simulated one attributed to machines—is legible in its own vocabulary: systems that "ape human expression" [17], meanings that "piggyback" on human competence [18], "self-less semiotic simulacra" [19]. Its most durable figure is Bender and colleagues' "stochastic parrot", a system engaged in "haphazardly stitching together sequences of linguistic forms . . . without any reference to meaning" [20]: well-formed sequences without semantic access, mimicry without meaning.

Definitional responses follow the same pattern. Runco and Jaeger's criterion—that a product is creative insofar as it is both novel and valuable [21]—was framed as minimal and neutral. Applied without modification to LLM outputs, it yields a conclusion its own author resists: that generative systems satisfy both conditions and must be admitted as creative. Rather than revise the framework, Runco supplements the definition with intentionality, autonomy, and authenticity [22], conditions stipulated as absent from artificial systems; the choice, as he puts it, is between accepting AI as creative under the standard definition or amending the definition to exclude it, and the second option is taken. The revision is motivated less by a principled account of why intentionality is required than by the prior commitment to preserving creativity as exclusively human.

Fazi names the shared structure of these positions the simulative paradigm [23]: human cognition serves as the normative model against which machine cognition is measured, so that the machine appears, by the terms of the framing itself, as approximation and deficient copy. The exclusion of machine creativity is built into the premises rather than argued for. What is needed, accordingly, is a different framework—one capable of specifying what creativity is as a structural property of outputs, independently of assumptions about the nature of the system that produces them. The following section develops such a framework.

4. CREATIVITY AS DIFFERENTIAL RECOMBINATION: THE COGNITIVE-SEMIOTIC FOUNDATION

4.1 Conceptual Blending and the Structure of Novelty

Any tractable account of creativity must first answer whether novelty can be explained at all. If something genuinely new is by definition disconnected from what precedes it—if it irrupts from an absolute elsewhere—then no theory can account for its emergence without undermining the conditions of intelligibility

that make theorising possible. As North observes, "the absolutely new is simply unintelligible" [24]: what cannot be related to prior structures cannot be recognised as new in any meaningful sense. A theory that locates novelty in absolute origination therefore places its own object beyond theoretical reach. Boden puts the point directly: "insights do not come from gods—and they do not come from nowhere, either" [16].

Rejecting creatio ex nihilo is accordingly a precondition for a serious account of creativity, not a concession to reductionism, and one with deep roots in the Western tradition: from Parmenides' challenge to coming-into-being, through the atomist thesis that birth is always a reorganisation of pre-existing material, to Aristotle's understanding of generation as recycling. Eco traces the same lineage through the medieval ars combinatoria—Lullo's Ars Magna, Leibniz's universal language, the Kabbalistic tradition—to a shared intuition: making is always recombining [25]. Pascal states it with characteristic economy: "Let no one say I have said nothing new: the arrangement of the material is new." Novelty is a differential category: a deviation, a structural rearrangement, not an absolute genesis.

Fauconnier and Turner's theory of conceptual blending develops within this tradition. Its proximate precursor is Koestler's bisociation, the creative act as the sudden intersection of two habitually distinct matrices of thought [26]: Newton's apple and Cézanne's apple are, on this account, structurally comparable events, each reconfiguring conceptual space through the unexpected interlocking of separate frameworks. Fauconnier and Turner supply the formal architecture explaining how such intersections operate [2]: new meaning arises through the selective integration of elements from distinct input spaces into an emergent blended space irreducible to the sum of its inputs.

Three features of the model matter here:

(1) Emergence: The blend is not a concatenation or superimposition of input structures but generates properties and relations present in neither input alone—a structural claim, not a metaphor, intelligible only within the integrated configuration.

(2) Constraint: Blending is governed by four principles—topology (structural relations must be preserved across the mapping), web (the blend must remain interpretable in terms of its inputs), integration (it must cohere as a stable unit), and unpacking (a competent user must be able to reconstruct the inputs from it) [16, 27]. These constraints distinguish creative recombination from arbitrary perturbation: novelty arises within a rule-governed space, not against or outside it.

(3) Substrate-neutrality: Nothing in the formal architecture of blending—input spaces, cross-space mapping, emergent structure, integration constraints—requires biological implementation. This openness is what makes blending a productive framework for machine creativity.

4.2 The Semiotic Extension: Structuralism, Eco, Lotman, Kristeva

Blending theory accounts for novel meaning as an individual cognitive event, but the blend is not only cognitive: it becomes new, in the sense that matters here, only when it enters a shared symbolic space and is recognised and taken up by others. Novel meaning is a relational effect of a semiotic environment, not a private property of mental structures [28]—which is why the move to semiotics is a theoretical necessity rather than an addition to the framework.

The relevant tradition begins with Saussurean structuralism: meaning is not a property of individual signs but an effect of differential relations among signs within a system; a sign acquires significance only through its position within a network of oppositions [29]. If meaning is constitutively relational, the emergence of new meaning cannot be the introduction of an unprecedented element, only a reorganisation of differential relations within an already-constituted space. The question of the new becomes one of structural transformation: recombination that alters a system's internal relations and produces effects the system did not previously contain.

Eco's account of the encyclopaedia develops this line. Rather than treating the sign as a static unit, Eco locates it within a continuous circulation of references and re-uses [6]. Invention, for Eco, never begins ex nihilo: "convention can arise only if the invention of the not-yet-said is wrapped in the already-said"; "in the semiotic universe there are neither heroes nor prophets" [3]. Every proposal designs itself against an already-organised culture; to call something "new", in this framework, is to say that it has reorganised its position within the encyclopaedic network, activating latent pertinences and producing new interpretive pathways. The encyclopaedia is a generative resource, at once constituted by past production and constitutive of future production.

Lotman's semiosphere extends the picture to cultural dynamics [4]: a heterogeneous environment where strongly codified regions coexist with unstable, liminal zones open to contamination and translation. It is in these peripheral zones—where texts migrate between codification regimes—that novelty manifests, not as a substantially new entity but as the emergent effect of structural transposition. Lotman's "explosions" do not introduce an originary moment; they reveal accumulated local deviations reaching a threshold that compels systemic reorganisation. A system's normativity is thus not rigidity but a capacity to accommodate transformations it did not predict.

Kristeva's intertextuality provides the textual instance of the same dynamic [30]: every text is a "mosaic of quotations," a redistribution of prior enunciations that alters them in the act of incorporation. The text is a site of transformation rather than an origin; novelty is differential recombination, not rupture. Together, Eco, Lotman, and Kristeva articulate the cultural-semiotic dimension that blending theory alone leaves underspecified, embedding Fauconnier and Turner's generative structure within the conditions that make novelty recognisable and culturally significant.

4.3 Toward a Structural Definition of Creativity

The convergence of blending theory with the semiotic frameworks of Eco, Lotman, and Kristeva supports the following definition:

> *Creativity is the emergent property of a constrained transformation within an already-constituted semiotic space, whereby the differential recombination of existing elements produces structurally novel configurations that rearticulate the system's internal relations and yield non-deducible interpretive effects and open new trajectories of sense-making.*

Four properties follow:

—Structural. Novelty is a formal relation between an output and the semiotic space from which it emerges, not a property of the mental states or intentions of the producing agent: whether an output is creative is a question about configurations, not about the inner life of whoever or whatever produced it.

—Constrained. Random perturbation is excluded by definition. An output produced by arbitrary shuffling may be statistically novel but not creative, since it is not anchored within the rule-governed space whose transformation constitutes genuine novelty; blending's constraints (topology, web, integration, unpacking) function here as the cognitive analogue of the semiotic norm against which deviation becomes meaningful rather than merely noisy.

—Relational. Novelty is not intrinsic to an output but a relation between the output and its semiotic context; the same configuration may be novel relative to one space and unremarkable relative to another.

—Scalar. The degree of creativity is, in principle, measurable as the magnitude of structural deviation from the nearest pre-existing configuration.

This definition does not deny the relevance of intentionality, embodiment, or experiential depth to human creativity; it denies only that these are necessary conditions for novelty production as such. It leaves open the possibility that human creativity achieves forms of semiotic productivity that differ significantly from machine creativity—what it refuses is foreclosing that comparison by definitional stipulation.

5. HOW GENERATIVE AI ARCHITECTURES INSTANTIATE THE MODEL

5.1 Large Language Models

A Large Language Model is trained on a corpus that, for the largest contemporary systems, constitutes a substantial sample of recorded human linguistic production across domains, registers, and languages—an operational instantiation of Eco's "library of libraries" [6]. Training produces not a lookup table of stored sentences but a compressed parametric representation of the corpus's distributional structure; generation proceeds by autoregressive sampling from a learned probability distribution conditioned on context, producing sequences governed by learned distributional constraints rather than by retrieval of stored instances.

Mapped onto the terms of Section 4.3: the training corpus is the semiotic space relative to which novelty is defined; the learned distributional parameters are the constraint system governing recombination; stochastic sampling supplies the non-determination that makes outputs exploratory rather than reproductive. The generation process has, in formal terms, the structure of constrained exploration: it produces configurations the constraint system permits but does not uniquely determine, and that are not deducible from any single training instance.

The more substantive objection—that distributional operations over linguistic form produce syntax without semantic content, so that any resulting novelty is hollow—has been engaged from several angles. Compagno describes LLMs as "the most radical instantiation of the structuralist project: machines capable of producing meaning without minds, through the sole operativity of a formal system of relations" [31], suggesting that a Saussurean, differential account of meaning fits LLM semantics better than the intentionalist model most dismissals presuppose.

Paolucci [32] mounts a related but more targeted critique of the "barrier of meaning" objection [33] that recurs across this debate: the claim, traceable to the Searlean tradition, that AI-generated utterances are "defective surrogates" of genuine statements because they lack access to reference [34]. What sustains the objection, on Paolucci's reading, is a residual myth of meaning—the assumption that meaning is an entity localisable in an interior, subjective dimension, transparently accessible to introspection. Once the signifier/signified dichotomy that grounds this assumption is suspended, and once one accepts, following Agüera y Arcas [35], that "in any falsifiable way, statistics do amount to understanding", the question is reframed: not whether systems such as ChatGPT "access" meaning, but what conception of meaning can account, without metaphysical residue, for their operative efficacy [32]. It is at this point that a semiotic framework—understood as the "science of all mediations that filter the elaboration of meaning beyond biological determination" [?]—proves productive: it allows sense-making to be treated as a relational process realised through shared mediations and structures, independently of whether the participating system is human, animal, or artificial.

None of this establishes that LLMs process language as humans do; together, however, these positions suggest that dismissing LLM outputs as semantically empty rests on contestable philosophical premises rather than settled findings.

5.2 Generative Adversarial Networks

A Generative Adversarial Network instantiates the same recombination logic through a structurally distinct mechanism [36], one that maps onto the blending framework with particular transparency. The network takes the norms encoded in a training distribution—the Encyclopaedia, in the terms of Section 4.2—and sets them against a two-agent system: a Generator, which proposes candidate outputs by projecting from a latent space, and a Discriminator, which applies encyclopaedic pressure by evaluating those candidates against the training distribution. The adversarial exchange between the two forces the emergence of an integrated output that satisfies the Discriminator's structural expectations while drawing on the Generator's latent degrees of freedom: an output anchored to the training distribution's semiotic norms while introducing configurations absent from it.

The mapping onto Fauconnier and Turner's integration network can be stated precisely: the training distribution supplies one input space, providing the structural norms against which outputs are evaluated; the latent space supplies the second, encoding dimensions of variation rather than recognisable features directly; the adversarial dynamic constitutes the integration constraint governing the blend's production. The result is, in the technical sense of Section 4.1, emergent structure.

Zeilinger characterises GAN creativity in terms that resonate with both the blending and the semiotic account: GAN-generated images are "fundamentally relational, embedded, and dialogic—not in the traditional sense of the spirited original genius, but in a more progressive sense that proposes creativity as fundamentally distributed" [37]. The output cannot be attributed to a single agent: it results from the structural interaction of Generator, Discriminator, training data, and the human decisions involved in architecture, training, and curatorial selection—a distributed character that, on the recombination model, GAN production shares with human creative production, even where the latter is conventionally described in terms of individual authorship.

5.3 Convergence and Divergence: A Comparative Note

LLMs and GANs instantiate the same structural operation—differential recombination, in the sense defined in Section 4.3—through different formal means: the LLM by probabilistically exploring paths within a parametric Encyclopaedia; the GAN by integrating latent variance with trained structural constraints under adversarial pressure. This section has traced two distinct routes to the same operation, which is the paper's central claim: novelty emerges neither from unconstrained randomness nor from mere retrieval, but from the constrained transformation of an already-constituted semiotic space, whatever the architecture that performs it.

This convergence has an implication beyond the comparison of architectures. It suggests that generative systems are better approached not as candidates to be judged against anthropocentric criteria of authorship and intention, but as semiotic devices in their own right—a shift Picca [38], following Vromen [39], frames as treating generative systems not as simulations of human cognition but as semiotic machines, intelligible through a theory of language and signs rather than a theory of mind. On this reading, LLMs and GANs alike participate in an evolving ecology of signs: their outputs amplify, distort, and reconfigure existing forms of meaning rather than merely reproducing them, entering into relation with the codes, styles, and interpretive traditions that surround them. Picca, following Eco and Lotman, describes the resulting outputs as "open works" requiring interpretive cooperation, produced at the periphery of the semiosphere—the zones, in Lotman's terms, where new configurations of sense are most likely to emerge [40]. What distinguishes generative systems from earlier media technologies, on this account, is not their capacity to transmit meaning already fixed elsewhere, but their capacity to reconfigure the semiotic space they operate within—precisely the capacity this paper has termed differential recombination.

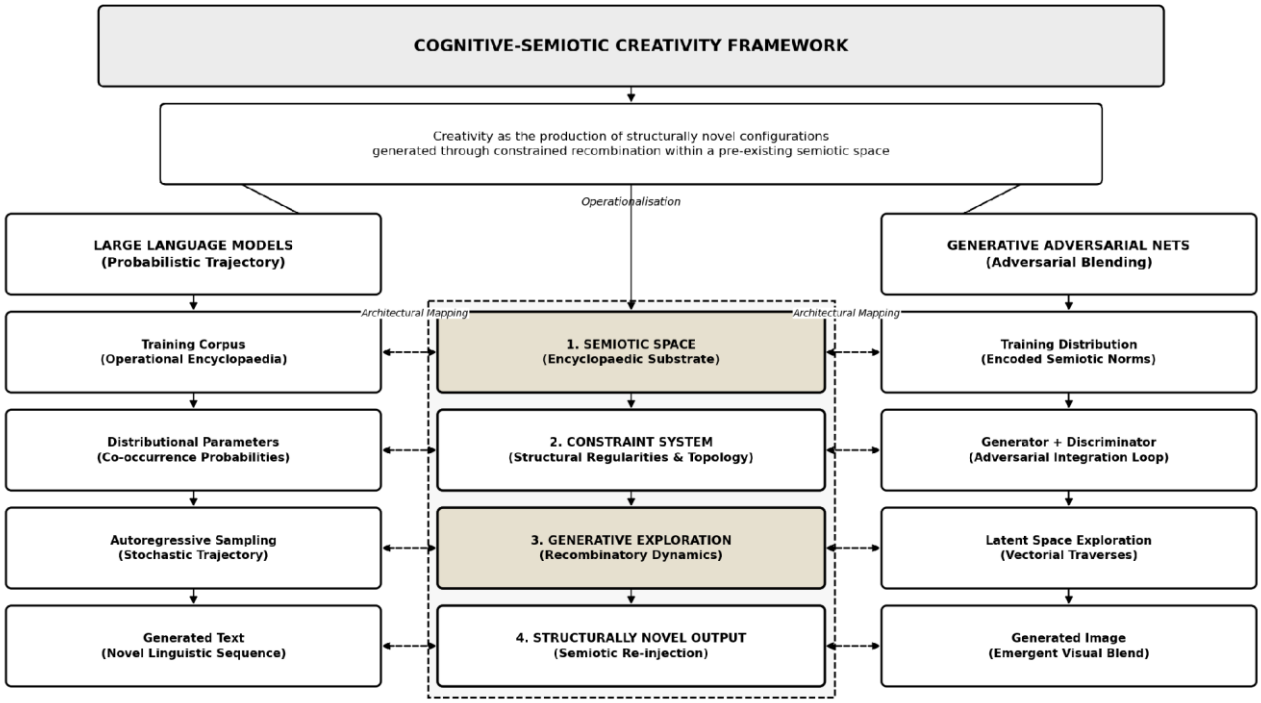

Fig. 1. Architectural Instantiation of the Cognitive-Semiotic Creativity Framework.

6. OPERATIONALISING DIFFERENTIAL RECOMBINATION: A SEMIOTIC PROFILE VECTOR

Sections 4 and 5 have established creativity as a structural relation rather than a property residing in either human intention or artificial mimicry, and have traced how LLMs and GANs alike instantiate that relation through formally distinct mechanisms. What remains is the question of assessment: how might one determine, for any given output, whether differential recombination has in fact occurred? The following section addresses this by translating the theoretical model into a set of four evaluative criteria applicable, without modification, to both human and machine production.

To operationalise this insight, this section translates the conceptual parameters of differential recombination into a formal, multi-dimensional assessment model: the Semiotic Profile Vector (P).

$$P = \langle D_1, D_2, D_3, D_4 \rangle \quad where\ D_i \in [0, 1] \qquad (1)$$

Rather than reducing creativity to a binary judgment, P evaluates any creative output—whether generated by a human author or an artificial system—along four quantitative and semiotically grounded axes.

6.1 Operationalising the Four Semiotic Axes

Dimension 1: Structural Differentiation (D_1). This dimension measures the structural distance between an output and the core density of the pre-existing semiotic space. It formalises Fauconnier and Turner's primary condition of conceptual blending: the actualisation of structures that were latent but previously unrealised [2]. The dual nature of the proxies used to track this distance—computational for artificial architectures, philological and cultural for human production—should not be read as implying two distinct notions of creativity. It reflects, rather, the differing epistemic accessibility of the two systems: whereas a machine permits direct vectorial observation of its own internal states, human semiotic activity can only be traced through its textual sedimentations and the dynamics of its reception within culture.

—*Computational Metric (AI)*: D_1 is computed as the mathematical displacement of the point occupied by a given output in the model's vector space ($\vec{y}$) relative to the centroid of its training distribution ($\vec{\mu}$), or, equivalently, as the statistical distance separating a generated instance from its source distribution: the Fréchet Inception Distance (FID) in GAN latent space [41]; local entropy dynamics for LLM output distributions [42]; MAUVE divergence frontiers for deep semantic novelty [43].

—*Human/Cultural Metric:* Measured via stylometric distance and formal departure from established genres or syntactic norms.

Dimension 2: Semiosphere Participation (D_2). Following Lotman's semiosphere [4] and Picca's positioning of generative models at its periphery [38], D_2 evaluates whether an output enters into living circulation, acting as a catalyst for further semiotic activity—transformation, reinterpretation, recombination—rather than remaining isolated.

—*Computational Metric (AI):* Operationalised through a Downstream Propagation Index (DPI), which estimates the extent to which an output remains computationally active through subsequent generative processes. As a preliminary formulation,

$$DPI = \alpha U + \beta P + \gamma F + \delta M \quad (2)$$

where U represents Reuse Frequency (recursive reuse across subsequent generations), P Prompt-chain Reuse (integration into iterative prompting workflows and derivative outputs), F Fine-tuning Incorporation (integration into synthetic datasets, continual learning pipelines, or model adaptation processes [44, 45]), and M Memory Persistence (retention and retrieval within persistent computational memory architectures [46, 47]). The coefficients α, β, γ, δ represent weighting parameters to be empirically calibrated according to the relative contribution of each indicator within a given computational ecosystem.

—*Human/*Cultural *Metric:* Measured via intertextual citation networks, critical reception, and adaptation frequency within cultural discourse.

Dimension 3: Intertextual Transformation (D_3). Drawing on Kristeva [30] and Eco [6], D_3 distinguishes genuine structural transformation from superficial recombination or memorised reproduction. While D_1 measures the degree of deviation from the pre-existing semiotic space—how different an output is from what already exists—D_3 evaluates whether such deviation corresponds to a meaningful reorganisation of inherited structures—how the output transforms what already exists. An output may therefore achieve high statistical novelty while failing D_3 if it merely recombines source fragments without altering their underlying relational configuration.

—*Computational Metric (AI):* Operationalised through an Intertextual Transformation Index (ITI),

$$ITI = \alpha G + \beta S \quad (3)$$

where G denotes Graph Structural Transformation, measured through graph-based comparison methods such as Graph Edit Distance (GED) [48], which quantify modifications in the topology of dependency or semantic graphs; and S denotes Semantic Shift, measuring the degree of transformation between source and generated representations through embedding-based comparison methods. Unlike D_1, which evaluates distance from the training distribution, ITI captures whether existing structures have been reorganised into a new relational configuration.

—*Human/Cultural Metric:* Analyzed through trope re-functionalisation—the degree to which underlying narrative or formal topoi are structurally altered rather than passively cited.

Dimension 4: Interpretive Openness (D_4). Grounded in Eco's concept of opera aperta [49] and Peirce's notion of infinite semiosis [50], D_4 evaluates the capacity of an output to sustain multiple coherent interpretative trajectories without collapsing into ambiguity or noise: not variability as such, but the coexistence of generative diversity and semantic stability.

—*Computational Metric (AI*): Operationalised through interpretive dispersion analysis, combining decoding entropy (H) across multi-path sampling trajectories (varying temperature $\tau > 0$) [42] with semantic cluster dispersion in high-dimensional embedding spaces [51], approximating the degree to which an output sustains multiple stable interpretative possibilities rather than random divergence.

—*Human/Cultural Metric*: Evaluated via hermeneutic divergence—the variance and stability of coherent interpretations produced across distinct interpretive communities.

6.2 The Operationalisation Matrix

The mapping between theoretical foundations, computational proxies, and human analytical standards is synthesised in Table 1, which functions as the interpretive key for reading any profile P produced by the model.

Table 1. Operationalisation Matrix: Mapping Semiotic Dimensions to Computational and Cultural Metrics.

Dimension	Theoretical Foundation	Computational Proxy (AI)	Cultural / Human Proxy
D_1: Differential Recombination	Emergent structure relative to semiotic space (Fauconnier & Turner)	Cosine / FID Distance from training distribution centroid	Stylometric / formal departure from canonical corpus norms
D_2: Semiosphere Participation	Dynamic circulation in cultural ecology (Lotman, Picca)	Downstream Propagation Index (API re-use, re-prompting)	Intertextual citation rate and cross-domain adaptation
D_3: Intertextual Transformation	Structural redistribution vs. pastiche (Kristeva, Eco)	Graph Edit Distance (GED) of relational dependency trees	Trope re-functionalisation vs. literal plagiarism / homage
D_4: Interpretive Openness	Polysemic stability and open work dynamics (Eco, Peirce)	Decoding Entropy (H) and cluster dispersion in embedding space	Hermeneutic variance across interpretive communities

6.3 Profile Taxonomy & Empirical Dynamics

Before the framework can be applied to actual outputs, it is useful to fix a small set of reference points against which real profiles can be read. The three archetypes below are theoretical schematisations, not measurements performed on existing texts or generations: they are idealised limit cases of P, not empirical findings, and their function is heuristic—to show why no single axis, taken in isolation, is sufficient to track creative achievement.

(1) Overfitted Output / Pastiche (idealised profile):

$$P_pastiche = \langle 0.85, 0.20, 0.15, 0.10 \rangle.$$

High surface novelty (D_1) paired with negligible structural transformation (D_3): the output recombines tokens without altering the relational topology it draws on.

(2) Unconstrained Noise / Hallucination (idealised profile):

$$P_noise = \langle 0.95, 0.05, 0.05, 0.00 \rangle.$$

Maximal distance from the corpus (D_1) achieved at the cost of the constraint system defined in Section 4.1, collapsing into semantic incoherence ($D_4 \rightarrow 0$).

(3) Genuine Differential Recombination (idealised profile):

$$P_creative = \langle 0.75, 0.70, 0.80, 0.75 \rangle.$$

Balanced, jointly high values across all four axes: structural departure achieved while topological integration and interpretive stability are preserved.

Read against these poles, the model also generates a testable prediction about the relationship between architectural sampling parameters—LLM temperature τ, GAN truncation thresholds—and the resulting profile, though this relationship has not itself been empirically verified here and is proposed as a hypothesis for future validation. As sampling hyper-parameters increase, D_1 is expected to grow roughly monotonically with decoding entropy, while D_3 should follow an inverted U-curve, peaking within an intermediate constraint-exploration window ($\tau \approx 0.7–0.85$). Beyond that window, the model predicts a collapse of structural constraints, with D_3 and D_4 declining sharply—the quantitative signature, on this account, of noise displacing recombination. Confirming this trend against generated corpora is a task for the empirical work the present framework is designed to make possible, not one it claims to have already carried out.

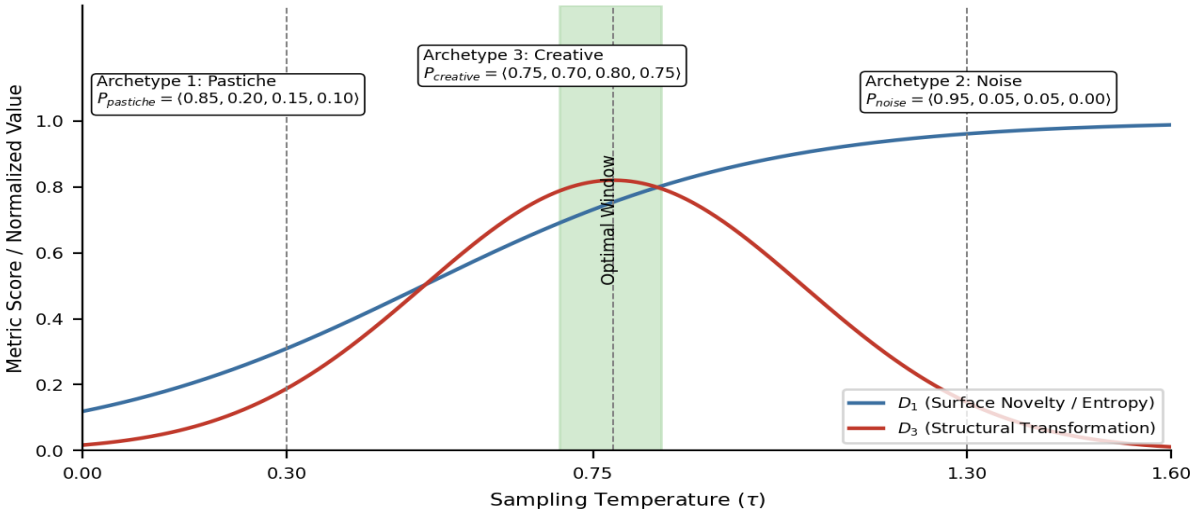

Fig. 2. Empirical mapping of output dynamics across sampling temperatures (τ), using D_1 and D_3 as proxy dimensions. Archetype 1 (Ppastiche) dominates low-τ regimes; Archetype 3 (Pcreative) emerges within the optimal constraint window ($\tau \approx 0.7–0.85$); Archetype 2 (Pnoise) marks structural collapse ($D_3 \rightarrow 0$).

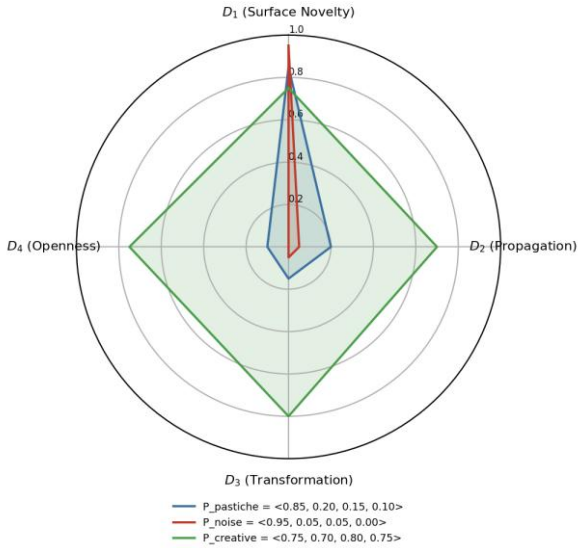

Fig. 3. Kiviat radar diagram showing the distinct semiotic profiles across the four-dimensional vector space P = ⟨D_1, D_2, D_3, D_4⟩. While Ppastiche and Pnoise exhibit severe dimensional asymmetry, Pcreative displays balanced, high-capacity coverage across all axes

7. CONCLUSION

This paper has argued that creativity, whether attributed to a person or a machine, names a structural relation rather than a metaphysical endowment. Building from Fauconnier and Turner's model of conceptual blending and extending it through the semiotic traditions of Saussure, Eco, Lotman, and Kristeva, Section 4 defined creativity as the emergent product of a constrained transformation within an already-constituted semiotic space—a definition indifferent, by design, to the biological or artificial nature of the system performing the transformation. Section 5 showed that this indifference is not merely stipulative: Large Language Models and Generative Adversarial Networks, for all their architectural divergence, instantiate the same operation of differential recombination, the former through probabilistic exploration of a parametric encyclopaedia, the latter through adversarial integration of latent variance against trained structural norms.

The Semiotic Profile Vector developed in Section 6 translates this convergence into an instrument. Rather than adjudicating creativity by fiat—reserving it to human agents on the strength of intentionality or authenticity, or extending it uncritically to any statistically novel output—the four dimensions of P allow the degree and kind of creative achievement to be assessed on structural grounds common to human and artificial production alike: distributional departure, participation in circulating semiosis, depth of relational reorganisation, and capacity to sustain plural, non-collapsing interpretation. The profile

taxonomy sketched in Section 6.4—pastiche, noise, genuine recombination—indicates that the framework's value lies precisely in what it refuses to collapse: high novelty without transformation is imitation; high novelty without interpretive stability is noise; only their joint satisfaction constitutes creativity in the sense argued for here.

Two qualifications should be stated plainly. First, the computational proxies proposed for each dimension—FID, entropy dynamics, MAUVE, graph edit distance, decoding dispersion—remain preliminary formulations whose weighting coefficients require empirical calibration against corpora of human-judged creative and non-creative outputs; the model's discriminative validity is accordingly a matter for future quantitative work rather than a settled result of the present paper. Second, the argument developed here deliberately brackets the question of what, if anything, differs between a system that recombines meaning and a system that additionally lives through the meanings it recombines. Nothing in the structural definition of creativity forecloses that difference; it only refuses to let it be smuggled in as a precondition for calling a machine's output creative in the first place.

Eco observed that the semiotic universe knows neither heroes nor prophets, only positions within an encyclopaedia against which the not-yet-said takes its bearings. Generative systems, on the account offered here, occupy such a position—not as counterfeit authors awaiting unmasking, nor as neutral instruments extending a human maker's will, but as devices that redraw, however locally, the space of the already-said. Whether that capacity amounts to creativity in the fullest sense the word has carried since the Romantics is a question this paper leaves open by design; what it has tried to show is that the question can no longer be closed by definition alone.

8. REFERENCES

[1] Jan Løhmann Stephensen. Artificial creativity: Artificial Creativity. Transformations, 36:19–37, 2022.

[2] Gilles Fauconnier and Mark Turner. The Way We Think: Conceptual Blending and the Mind's Hidden Complexities. Basic Books, New York, 2002.

[3] Umberto Eco. Semiotics and the Philosophy of Language. Indiana University Press, Bloomington, 1984.

[4] Jurij Michajlovič Lotman. Universe of the Mind: A Semiotic Theory of Culture. Indiana University Press, Bloomington, 1990.

[5] Julia Kristeva. Sēmeiōtikē; recherches pour une sémanalyse. Paris, Éditions du Seuil, 1969.

[6] Umberto Eco. The Search for the Perfect Language. Blackwell, Oxford, 1995. Translation of La ricerca della lingua perfetta nella cultura europea, 1993; frequently cited as "1997" referring to the paperback reprint.

[7] Stefano Bartezzaghi. Mettere al mondo il mondo: Tutto quanto facciamo per essere detti creativi e chi ce lo fa fare. Bompiani, Milano, 2021.

[8] Agostino d'Ippona. La città di Dio. 426. Edizione italiana di riferimento: Milano, Bompiani, 2001.

[9] Rob Pope. Creativity: Theory, History, Practice. Routledge, London, 2005.

[10] Andreas Reckwitz. The Invention of Creativity: Modern Society and the Culture of the New. Polity Press, Cambridge, UK, 2017. Trans. Steven Black.

[11] Claudio Paolucci, Flavio Valerio Alessi, and Luigi Lobaccaro. The myth of the mad genius: a semiotic perspective on creativity and madness. RIFL - Rivista Italiana di Filosofia del Linguaggio, 2025.

[12] Ada Lovelace. Notes by the translator [to l. f. menabrea, "sketch of the analytical engine invented by charles babbage"]. Scientific Memoirs, Selected from the Transactions of Foreign Academies of Science and Learned Societies, 3:666–731, 1843. Commonly cited as "1842" in secondary literature referring to the composition date.

[13] A. Michael Noll. Human or machine: A subjective comparison of Piet Mondrian's "Composition with Lines" (1917) and a computer-generated picture. The Psychological Record, 16(1):1–10, 1966.

[14] Erik E. Guzik, Christian Byrge, and Christian Gilde. The originality of machines: AI takes the torrance test. Journal of Creativity, 33(3):100065, 2023.

[15] Jon McCormack, Toby Gifford, and Patrick Hutchings. Autonomy, authenticity, authorship and intention in computer generated art. International Conference on Computational Intelligence in Music, Sound, Art and Design (EvoMUSART), 2019.

[16] Margaret A. Boden. The Creative Mind: Myths and Mechanisms. Routledge, London, 2nd edition, 2004.

[17] David Hellberg. Sacred signs or empty words: A naturalistic and semiotic comparison of AI vs human generated art. IntechOpen, 2025.

[18] Luca Miracchi Titus. Does ChatGPT have semantic understanding? A problem with the statistics-of-occurrence strategy. Cognitive Systems Research, 83, 2024.

[19] Benjamin Kemple. Semiosis in the machine: Lost natures and artificial intelligence. Filosofi(e)Semiotiche, 10(1):24–39, 2023.

[20] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21, pages 610–623, New York, NY, USA, March 2021. Association for Computing Machinery.

[21] Mark A. Runco and Garrett J. Jaeger. The standard definition of creativity. Creativity Research Journal, 24(1):92–96, 2012.

[22] Mark A. Runco. AI can only produce artificial creativity. Journal of Creativity, 33(3), 2023.

[23] M. Beatrice Fazi. Contingent Computation: Abstraction, Experience, and Indeterminacy in Computational Aesthetics. Rowman & Littlefield International, London, 2018.

[24] Michael North. Novelty: A History of the New. University of Chicago Press, Chicago, IL, 2013. Paperback edition 2015.

[25] Umberto Eco. Combinatoria della creatività. Conferenza tenuta a Firenze per la Nobel Foundation, 15 settembre 2004, 2004. Consultato il 3 ottobre 2025.

[26] Arthur Koestler. The Act of Creation. Hutchinson, London, 1964.

[27] Margaret A. Boden. Computer models of creativity. AI Magazine, 30(3):23–34, 2009. DA VERIFICARE.

[28] Claudio Paolucci. Cognitive Semiotics: Signs, Mind, and Culture. Springer, 2021.

[29] Ferdinand de Saussure. Cours de linguistique générale. Payot, 1916.

[30] Julia Kristeva. Desire in Language: A Semiotic Approach to Literature and Art. Columbia University Press, New York, 1980. Ed. Leon S. Roudiez.

[31] Dario Compagno. Intervento al séminaire international de sémiotique. Paris, Association Française de Sémiotique (29/11/2023), 2023.

[32] Claudio Paolucci. The myth of meaning: Generative AI as language-endowed machines and the machinic essence of the human being. Semiotica, 2025(262):5–23, 2025.

[33] Melanie Mitchell. Artificial intelligence hits the barrier of meaning. Information, 10(2):51, 2019.

[34] John R. Searle. Minds, brains, and programs. Behavioral and Brain Sciences, 3(3):417–424, 1980.

[35] Blaise Agüera y Arcas. Do large language models understand us? Daedalus, 151(2):183–197, 2022.

[36] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks, June 2014.

[37] Martin Zeilinger. Generative adversarial copy machines. Culture Machine, 20:1–23, 2021.

[38] Davide Picca. Not Minds, but Signs: Reframing LLMs through Semiotics, March 2025. arXiv preprint v1: May 2025; revised versions through 2026.

[39] Elad Vromen. Language Models as Semiotic Machines: Reconceptualizing AI Language Systems through Structuralist and Post-Structuralist Theories of Language, October 2024.

[40] Jurij M. Lotman. La semiosfera. L'asimmetria e il dialogo nelle strutture pensanti. Marsilio, Venezia, 1985.

[41] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. Advances in Neural Information Processing Systems, 30, 2017.

[42] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The Curious Case of Neural Text Degeneration, February 2020.

[43] Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaid Harchaoui. MAUVE: Measuring the Gap Between Neural Text and Human Text using Divergence Frontiers, November 2021.

[44] Long et al. Ouyang. Training language models to follow instructions with human feedback. arXiv preprint arXiv:2203.02155, 2022.

[45] Wang et al. Self-Instruct: Aligning Language Models with Self-Generated Instructions. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 13484–13508, Toronto, Canada, July 2023. Association for Computational Linguistics.

[46] Joon Sung et al. Park. Generative Agents: Interactive Simulacra of Human Behavior, August 2023.

[47] Charles et al. Packer. MemGPT: Towards LLMs as Operating Systems, February 2023. arXiv preprint, October 2023; published 2024.

[48] Xinbo et al. Gao. A survey of graph edit distance. Pattern Analysis and Applications, 13(1):113–129, 2010.

[49] Umberto Eco. Opera aperta: Forma e indeterminazione nelle poetiche contemporanee. Bompiani, Milano, 1962.

[50] Charles Sanders Peirce. Collected Papers of Charles Sanders Peirce. Harvard University Press, Cambridge, MA, 1931–1958. Ed. Charles Hartshorne, Paul Weiss, and Arthur W. Burks, 8 vols.

[51] Ashwin K. et al. Vijayakumar. Diverse beam search: Decoding diverse solutions from neural sequence models, 2016.