Article

Not peer-reviewed version

Coupling Physical Entropy Production and Semantic Diversity in Generative AI Ecosystems: A Preliminary Predictive Framework

Mauro Capocelli *

Posted Date: 19 May 2026

doi: 10.20944/preprints202605.1249.v1

Keywords: AI entropy production; model collapse; semantic homogenization; human novelty; synthetic data; information thermodynamics

doi:10.20944/preprints202605.1249.v1

Article

Coupling Physical Entropy Production and Semantic Diversity in Generative AI Ecosystems: A Preliminary Predictive Framework

Mauro Capocelli

Unit of Process Engineering, Department of Engineering, Università Campus Bio-Medico di Roma, Via Alvaro del Portillo, 21, 00128, Rome, Italy; m.capocelli@unicampus.it

Abstract

Generative artificial intelligence is increasingly embedded in recursive informational ecosystems in which outputs are produced, published, retrieved, summarized, copied and pasted into new human or machine production. This paper proposes a preliminary predictive framework for Dissipative Semantic Homogenization (DSH), a possible regime in which recursive generative AI ecosystems dissipate physical energy while corpus-level semantic diversity contracts and saturates. The framework does not identify thermodynamic entropy with semantic entropy. Instead, it treats them as operationally coupled variables: semantic distributions are transformed by physically implemented computation, while energy dissipation provides a macroscopic cost proxy. We model semantic diversity as Shannon entropy over a corpus-level partition of semantic states and introduce modal amplification, independent novelty injection, and AI assimilation of nominally human production as control variables. The model yields empirically testable implications: semantic contraction should occur only when effective independent novelty falls below a stability threshold; contraction should be scale-dependent; and cumulative semantic loss should saturate even while physical entropy production continues. The framework is not presented as an empirical law, but as a testable theoretical model for future longitudinal and controlled studies.

Keywords: AI entropy production; model collapse; semantic homogenization; human novelty; synthetic data; information thermodynamics

---

1. Introduction

Generative artificial intelligence (GAI) is becoming a structural component of contemporary informational ecosystems and its "hype" is pushing applications based on Large language models (LLM) in many aspects of human life. LLM and related generative systems are increasingly used to draft, summarize, translate, rank, retrieve, and recombine texts across any communication contexts. This expansion can increase productivity, reduce access barriers, and support exploration in domains where the relevant search spaces are too large for unaided human cognition [1,2].

The recent book by Kissinger, Schmidt, and Huttenlocher, synthetized main breakthrough evidences of GAI potentialities that modified the public perception of AI as a technology capable of "transcending" human forms of expertise [3]: AlphaZero, which learned the game of chess without human strategic instruction and produced moves deemed unconventional even by grandmasters [4]; the discovery of halicin, a new antibiotic candidate identified through deep learning by exploring molecular spaces difficult to access by traditional experimentation [5].; and GPT-3, which demonstrated the possibility of generating coherent and seemingly human language for many tasks without specific training for each [6].

These examples established the view that AI could be an epistemic infrastructure capable of detecting latent patterns, navigating high-dimensional spaces, and producing candidate solutions that would be difficult or costly for humans to generate unaided. This interpretation is supported by a growing literature on AI capabilities suggesting that AI can operate as a cognitive amplifier, increasing the speed and accessibility of complex forms of work [7,8]. Beyond individual productivity, AI is increasingly becoming an instrument for scientific discovery and technological design. Krenn et al. argue that artificial intelligence can contribute to scientific understanding by helping researchers to work in "spaces" too large for unaided human exploration [9]. In chemistry, materials science, and engineering, this potential is already visible. Park et al. propose a generative AI framework based on molecular diffusion models for the design of metal-organic frameworks for carbon capture, illustrating how AI can support the exploration of large molecular design spaces [10]. More broadly, Venkatasubramanian frames AI as a potentially transformative technology for chemical engineering, with applications in modelling, optimization, process control, materials discovery, and the management of complex industrial systems [11]. Many works highlighted the social impact of "AI-helped" R&D showing how it could contribute to the achievement of a substantial number of Sustainable Development Goal targets, particularly in areas such as climate modelling, resource optimization, healthcare, education and environmental monitoring [12,13]. This optimistic reading is also reflected in the way AI is increasingly imagined within scientific practice. Messeri and Crockett describe four perceptions of AI as an extension of human cognitive abilities: it promises to read more extensively, evaluate more consistently, compute more deeply, and simulate more broadly than individual researchers or traditional institutions [14].

At the same time, the large-scale use of similar models and retrieval systems raises a different question: not whether generative AI can produce valuable outputs in individual "virtuous" cases, but whether recursive AI-mediated production can alter the diversity of semantic states represented at the level of an informational ecosystem. Outputs generated by LLMs and related generative systems are published, indexed, ranked, summarized, recombined, and reused in future writing, prompting, retrieval, and training cycles. The resulting ecosystem is recursive: texts generated or assisted by AI become part of the environment from which later human and machine production draws. At the same time, the epistemic status of AI-generated outputs is ambiguous. They may be useful, fluent, and contextually appropriate, but they are not equivalent to independent human knowledge: they are probabilistic productions shaped by training data, decoding procedures, alignment constraints and prompt design. In probabilistic terms, a generative model tends to produce outputs from high-probability regions of its learned distribution and close to dominant semantic modes. Such outputs may be useful and even creative in individual cases; nevertheless, when they are produced at scale and recursively reused, they can increase the mass of already probable semantic regions relative to rare, marginal, or idiosyncratic regions. Shumailov et al. define model collapse as a degenerative process in which model-generated data pollute the training set of subsequent generations, causing models to misperceive the original distribution [18]. Their analysis distinguishes early collapse, characterized by loss of information about distributional tails, from late collapse, in which the model may converge toward a distribution with substantially reduced variance [18,19]. Alemohammad et al. analyze self-consuming generative loops (or "autophagous training processes") in which generative models are trained on synthetic data produced by previous generations [20]. Their main conclusion is that, without enough fresh real data in each generation, future generative models progressively lose either quality, measured as precision, or diversity, measured as recall [20]. Anderson et al. reported that LLM-driven thinking brings to a homogenization effect on human-in-the-loop divergent ideation processes (particularly at group-level), partly attributable to low inferential distance between LLM outputs and apparently finished ideas [21].

This duality motivates the present paper. These examples illustrate the diversity-expanding potential of AI in specific task-level and discovery-oriented settings. GAI can increase "local productivity" mainly in complex works and spaces, but its population-level effect depends on how many agents draw from similar models, prompts, ranking systems, corpora, and alignment constraints. This mechanism links the present framework to an alternative body of research related to the study of information manipulation and to the study of its thermodynamic consequences in the era of a possible extensive use of GAI.

doi:10.20944/preprints202605.1249.v1

The term “semantic entropy” has also entered recent AI research in a different but relevant context. Farquhar et al. use semantic entropy to detect hallucinations by measuring uncertainty across semantically equivalent model responses [22]. Since Landauer’s principle, information processing has been understood as a physical process: logically irreversible operations, such as memory erasure, are associated with a minimal thermodynamic cost [23]. Bennett’s analysis of reversible computation further clarified that the relevant issue is not computation in the abstract, but the physical implementation, resetting, memory management, and cyclic use of computational devices [24]. More recent work in the thermodynamics of information and stochastic thermodynamics has generalized this insight beyond the idealized erasure of a single bit, showing that informational states, measurements, feedback, memory, and resetting must be treated as physically instantiated processes with implementation-dependent energetic and entropic costs [25-27]. Moreover, this thermodynamic perspective connects directly with recent work on the environmental cost of large-scale AI. Strubell et al. showed that deep-learning models for natural language processing can entail substantial financial and environmental costs, particularly when progress depends on large computational resources [28]. Schwartz et al. subsequently proposed the idea of “Green AI”, arguing that efficiency should become an explicit criterion in AI research rather than treating accuracy gains as independent of computational cost [29]. Patterson et al. refined estimates of carbon emissions for large neural network training and emphasized that model architecture, datacenter efficiency, hardware, and electricity mix can change emissions by orders of magnitude [30]. More recently, Luccioni et al. shifted attention from training to deployment, showing that inference at scale—especially for multipurpose generative systems—can become a major component of AI’s energy and carbon cost [31].

This work studies GAI ecosystems in which increasing computational dissipation may coexist with decreasing or saturating semantic diversity, by introducing a novel model of Dissipative Semantic Homogenization (DSH). The starting claim is: when AI-mediated production grows and generated outputs are recursively reused, semantic distributions may become increasingly concentrated around high-probability modes; this phenomenon is exacerbated when independent human novelty is insufficient. Since AI-mediated semantic production is physically implemented through computational infrastructures, its expansion is associated with energy dissipation. However, the model does not treat energy expenditure as a direct cause of semantic contraction. It examines conditions under which increasing computational scale may coexist with local or corpus-level semantic concentration.

The present paper develops a preliminary predictive framework for this problem. The framework does not claim that generative AI universally reduces semantic diversity, nor that every AI-generated output is conventional or low in novelty. Instead, it identifies a possible regime in which recursive reuse of AI-mediated outputs, together with assimilation of human production, can increase the probability of already dominant semantic regions. In this regime, semantic diversity may contract, saturate, or reorganize, even while the volume of generated text and the energy dissipated by computation continue to increase. In this work, Shannon entropy is computed over a corpus-level probability distribution of semantic states, following the standard information-theoretic formulation of entropy [32,33]. The paper also introduces two novel elements to this study: 1) a modal amplification mechanism to represent the tendency of recursive AI-mediated production to reinforce already probable semantic regions; 2) human novelty injection as a stabilizing term, inspired by the role of fresh real data in preventing “auto-degradation”. Moreover, this work connects these semantic dynamics to the thermodynamics of computation (following the cited works [23-27]). These states may be operationalized as topics, frames, claims, arguments, or embedding-based clusters, depending on the empirical implementation.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 19 May 2026 doi:10.20944/preprints202605.1249.v1

2. Materials and Methods

The paper develops a reduced theoretical model of a recursive generative AI ecosystem. The model is not intended to estimate the empirical entropy of a specific corpus in this article. Instead, it defines the variables, mechanisms, and stability conditions that future empirical studies can test.

Let the semantic space of a corpus be approximated by a finite partition $X = \{x_1, \ldots, x_K\}$. Each semantic state x_k represents a region of meaning within the chosen corpus and modelling procedure. In an empirical implementation, texts may be embedded into a semantic vector space and clustered into Ksemantic regions. The model does not assume a universal semantic ontology: depending on the empirical implementation, semantic states may correspond to topics, frames, claims, arguments, narrative forms, scientific concepts, or embedding-based clusters.. If $n_i(t)$denotes the number of texts, passages, or semantic production units assigned to semantic state x_i at time t, then $p_t(x_i)$ is the ratio of Eq. 1. This construction converts a corpus into a discrete probability distribution over semantic states. Semantic diversity is measured using Shannon entropy over the corpus-level semantic distribution (Eq. 2), following the standard information-theoretic definition of entropy [31, 32]. The corresponding effective number of semantic states, N_{eff}, is defined in Eq. 3.

$$p_t(x_i) = \frac{n_i(\text{t})}{\text{N(t)}} = \frac{n_i(\text{t})}{\sum_{i=1}^{K} n_i\,(t)} \tag{1}$$

$$H_{\text{sem}}(t) = -\sum_{i=1}^{K} p_t\,(x_i)\ln\, p_t(x_i). \tag{2}$$

$$N_{\text{eff}}(t) = \exp\,[H_{\text{sem}}(t)]. \tag{3}$$

It must be noted that Shannon entropy does not measure truth, originality, cognitive depth, social value, or epistemic quality: a high value indicates that semantic production is distributed across many states with relatively balanced probabilities; a low value indicates concentration in a smaller subset of states. The effective number of semantic states, $N_{\text{eff},t}$, translates entropy into the number of equally represented states that would yield the same entropy. Thus, a corpus can grow in volume (size) while its effective semantic diversity decreases, remains stable, or increases, depending on how new production is distributed across semantic states. If the corpus is uniformly distributed across K semantic states, $P_t(x_i) = K^{-1}$ then $H_{\text{sem}}(P_t) = \ln K$ and $N_{\text{eff}}(t) = K$. If the corpus concentrates on one dominant semantic state $H_{\text{sem}}(P_t) \to 0$, and $N_{\text{eff}}(t) \to 1$.

2.1. AI-Mediated Modal Amplification

The next step is to specify how the corpus distribution p_t changes under recursive AI-mediated production and independent human novelty injection. Generative systems, ranking mechanisms and reuse dynamics may favor semantic regions that are already probable, acceptable, conventional, or easy to predict. We represent this effect through a modal amplification operator represented by Eq. 4.

$$p_t^{AI}(x_i) = \frac{p_t(x_i)^{\beta}}{\sum_{j=1}^{K} p_t\,(x_j)^{\beta}} \tag{4}$$

The parameter β controls the strength and direction of modal amplification. For $\beta = 1$, the transformation is neutral and $p_t^{AI}(x_i) = p_t(x_i)$. For $\beta > 1$, high-probability semantic states are amplified and low-probability states are suppressed. This regime represents the combined effect of probabilistic generation, alignment, ranking, user acceptance, reuse, and recursive publication when these processes preferentially increase the corpus share of already probable semantic modes. The modal amplification operator is not a claim about the content of every AI-generated text. It is a population-level abstraction. It represents the possibility that, when many agents use related models, retrieval systems, prompts, ranking mechanisms, and editing workflows, semantic regions that are already probable may become easier to reproduce and more likely to persist. In this sense, β should be interpreted as an aggregate parameter capturing the strength of recursive selection and reproduction of dominant semantic modes.

 doi:10.20944/preprints202605.1249.v1

Let $q_t(x_i)$ denote the distribution of original (human) semantic production at time t, the "global" corpus evolves according to Eq. 5 consisting of a first term representing the injection of "fresh-human" semantic material $q_t(x_i)$, while the second term represents AI-mediated modal amplification of the existing corpus distribution (Eq. 4). The two contributes are modulated by the paramenter α_t being the nominal fraction of new corpus material that comes from human or original production.

$$p_{t+1}(x_i) = \alpha_t q_t(x_i) + (1-\alpha_t)p_t^{AI}(x_i) \tag{5}$$

On the other hand, also nominally human production can itself be influenced by AI systems. Writers may use AI to search, draft, summarize, translate, edit, structure arguments, generate alternatives, or learn conventions. Taking it to the extreme, human users might just uncritically copy and paste AI outputs. A growing fraction of human-authored content can therefore become semantically aligned with AI-mediated distributions. Let $\lambda_t \in [0,1]$ denote the degree of AI assimilation in nominal human production, the observed human novelty q_t^{obs} is calculated by Eq. 6 including the AI-amplified semantic structure. Therefore the modified corpus distribution of Eq. 5 results into Eq. 7. The compact dynamic equation becomes Eq. 8 where we define the effective independent novelty rate as $\alpha_{\mathrm{eff},t}$ as reported by Eq. 9.

$$q_t^{obs}(x_i) = (1-\lambda_t)q_t(x_i) + \lambda_t p_t^{AI}(x_i). \tag{6}$$

$$p_{t+1}(x_i) = \alpha_t(1-\lambda_t)q_t(x_i) + [1-\alpha_t(1-\lambda_t)]p_t^{AI}(x_i) \tag{7}$$

$$p_{t+1}(x_i) = \alpha_{\mathrm{eff},t} q_t(x_i) + (1-\alpha_{\mathrm{eff},t})p_t^{AI}(x_i). \tag{8}$$

$$\alpha_{\mathrm{eff},t} = \alpha_t(1-\lambda_t). \tag{9}$$

The final expression of Eq. 8 formalizes the distinction between nominal human production and effective independent human novelty. A large amount of human-authored material may contribute little independent novelty if it is strongly assimilated to AI-mediated semantic distributions. Semantic diversity can be stabilized when a sufficiently large fraction of new corpus material remains both human or original in origin and independent from AI-mediated modal structure

The model does not require separate absolute estimates of total human and AI text production. Instead, it uses the exposure-weighted AI-to-human semantic production ratio (Eq 10) where $N_{AI,t}$ and $N_{H,t}$ denote the single (semantically relevant) production units. Therefore, the nominal fraction of new corpus material α_t can be re-written as the nominal human or original share Eq. 11. Combining the nominal human share and the assimilation term gives a novel expression of $\alpha_{\mathrm{eff},t}$ using the reduced production ratio (Eq. 12).

$$r_t = N_{AI,t}/N_{H,t} \tag{10}$$

$$\alpha_t = \frac{N_{H,t}}{N_{H,t} + N_{AI,t}} = \frac{1}{1+r_t} \tag{11}$$

$$\alpha_{\mathrm{eff},t} = \frac{1-\lambda_t}{1+r_t}. \tag{12}$$

This expression is one of the key control variables of the model. It shows that semantic diversity is not stabilized by nominal human production alone. It depends on the fraction of production that remains both human/original in origin and independent from AI-mediated semantic assimilation.

2.2. Energy Growth and AI Assimilation

The scenario dynamics are specified through a reduced growth closure. The human/original production stream is assumed to grow at a fixed baseline rate, whereas the relative growth of AI-mediated production with respect to human/original production is used as the scenario driver. The ratio r_t evolves as Eq. 13 where r_0 is the initial AI-to-human semantic production ratio and g_r is the scenario-specific relative AI-to-human growth rate. The model assumes constant energy intensity of AI-mediated semantic production over the simulation horizon. Under this simplifying assumption, the same derived growth rate g_{AI} governs both AI-mediated semantic production and the aggregate energy attributed to that production as in Eq. 14.

$$r_t = r_0(1+g_r)^t \quad (13)$$

$$E_t = E_0(1+g_{AI})^t \quad (14)$$

This assumption creates a direct link between semantic production growth and attributed energy growth. The energy term E_t should not be interpreted as a microscopic energy-per-bit, energy-per-token, or energy-per-concept quantity. It denotes a scenario-attributed aggregate energy demand associated with semantically relevant AI-mediated production over one model period. Future extensions may relax the constant-intensity assumption by introducing a time-dependent energy-intensity term.

Figure 1 summarizes the conceptual structure of the model. The framework couples two trajectories. The first is a semantic-recursive trajectory, in which the corpus distribution is transformed by AI-mediated modal amplification, independent human novelty and the effective novelty share. The second is a physical trajectory, in which generative AI computation consumes energy and produces macroscopic physical entropy.

The growth rate of AI-mediated semantic production is not specified independently but can be derived from the fixed human/original baseline growth rate g_H and the relative AI-to-human growth rate g_r according to Eq. 15.

$$g_{AI} = (1+g_H)(1+g_r) - 1 \quad (15)$$

Also the assimilation parameter λ_t defined by Eq. 6 and representing the degree to which nominally human/original production becomes semantically aligned with AI-mediated distributions, can be related to the growth rate of AI-mediated semantic production. In this version of the model, the dynamic of assimilation (λ_t) is assumed as a bounded increasing function of the scale of AI-mediated production according to Eq. 16.

$$\lambda_t = \lambda_0 + (\lambda_{\max} - \lambda_0)[1-(1+g_{AI})^{-\kappa t}] \quad (16)$$

The parameter λ_0 is the initial assimilation level, $\lambda_{\max}$ is the upper saturation level, and κ controls the curvature of the assimilation response. The bounded form reflects the assumption that AI-mediated assimilation can increase as AI-mediated production expands, but cannot grow without limit. Complete assimilation is excluded because residual independent production, domain expertise, institutional constraints, exogenous events, and non-AI-mediated creativity remain possible.

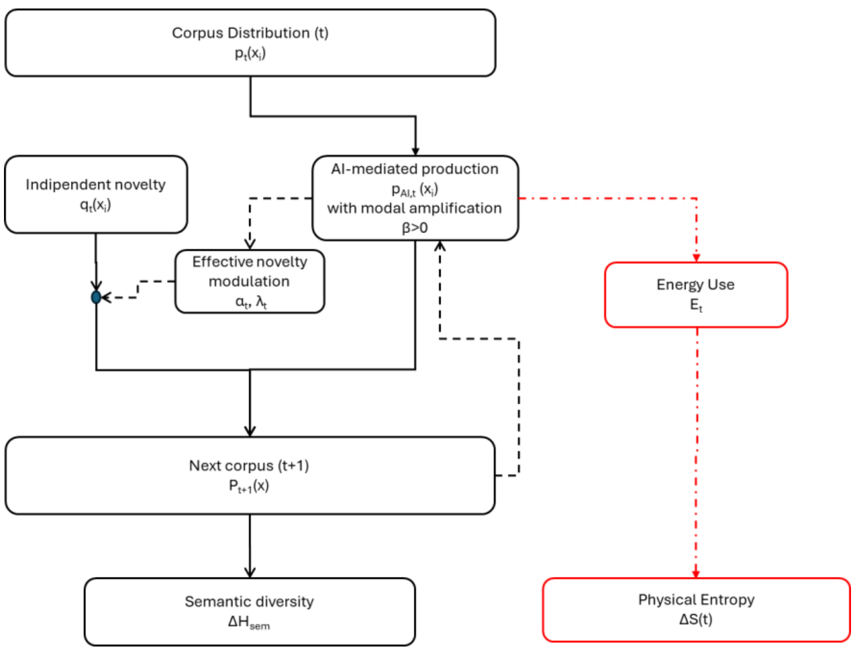

Figure 1. Conceptual structure of Dissipative Semantic Homogenization.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 19 May 2026 doi:10.20944/preprints202605.1249.v1

Physical entropy production is estimated from heat rejection. In a computing infrastructure operating approximately in steady state, the energy consumed by processors, memory, networking, and cooling systems is ultimately degraded into heat and rejected to the environment. The physical entropy production associated with this dissipative process is therefore approximated by Eq. 17 where $Q_{rej,t}$ is the scenario-attributed heat rejected during model period t, and T_0 is an effective environmental or heat-sink temperature. Calling the Semantic entropy variation $\Delta H_{\mathrm{sem},t}$, (if negative, the contraction between two consecutive periods), the cumulative semantic contraction is given by Eq. 18.

$$\Delta S_{\mathrm{phys},t} = \frac{Q_{rej,t}}{T_0}, \tag{17}$$

$$C_{\mathrm{cum}}(t) = \sum_{\tau=1}^{t} -\Delta H_{\mathrm{sem},\tau}. \tag{18}$$

To relate semantic contraction to the energetic scale of the scenario, the model uses the marginal contraction per unit energy and a cumulative dissipative indicator, respectively defined by Eq. 19 and 20. These indicators quantify semantic contraction relative to the scenario-attributed heat-equivalent energy scale. They do not imply that physical entropy is converted into semantic entropy loss.

$$\eta_{E,t} = \frac{-\Delta H_{\mathrm{sem},t}}{E_t} \tag{19}$$

$$\eta_E(t) = \frac{C_{\mathrm{cum}}(t)}{E_{\mathrm{cum}}(t)} \tag{20}$$

3. Results and Discussion

3.1. Sensitivity Analysis

The model yields a set of theoretical regimes rather than empirical forecasts. The purpose of this section is to examine the internal behavior of the framework under controlled assumptions and to derive assumptions for the scenario analysis and expectations for future empirical studies.

The first limiting case is the synthetic-only regime, in which no effective independent novelty is injected into the corpus. This case is obtained by setting $\alpha_{\mathrm{eff},t} = 0$. The corpus distribution then evolves only through the modal amplification operator. For $\beta > 1$, initial probability differences are amplified exponentially. Considering two semantic states $i$and j, their probability ratio after one amplification step is given by Eq. 21.

$$\frac{p_{i,t+1}}{p_{j,t+1}} = \left(\frac{p_{i,t}}{p_{j,t}}\right)^{\beta} \tag{21}$$

Therefore, if $p_{i,t} > p_{j,t}$ and $\beta > 1$, the ratio increases after the transformation. Repeated application of the operator progressively concentrates probability mass on initially dominant semantic states and suppresses distributional tails. In the limiting case of a unique dominant state, the distribution tends toward concentration on that state.

The second theoretical result concerns the stabilizing role of effective independent novelty. To derive the local stability condition, consider the compact semantic dynamics around a high-diversity state. Assume, for analytical simplicity, that the independent novelty source is uniformly distributed across the K semantic states. The critical threshold follows from a local stability analysis of the high-diversity state. The AI-mediated modal amplification operator $P_{AI,\beta}$sharpens the corpus distribution by raising semantic-state probabilities to the power $\beta$and renormalizing them. Around the uniform distribution u, a small zero-sum perturbation $\varepsilon_t$is amplified to first order as $P_{AI,\beta}(u + \varepsilon_t) \approx u + \beta\varepsilon_t$. Since only the fraction $1 - \alpha_{\mathrm{eff},t}$of the next-period distribution is produced by the amplified corpus component, the local perturbation multiplier is $(1 - \alpha_{\mathrm{eff},t})\beta$. Local stability of the high-diversity state therefore requires $(1 - \alpha_{\mathrm{eff},t})\beta < 1$, yielding the critical novelty threshold $\alpha_c = 1 - 1/\beta$.This gives the critical effective novelty threshold given by Eq. 21 and represented in Figure 2. The diversity-preserving regime is for $\alpha_{\mathrm{eff},t} > \alpha_c$, whereas the homogenizing regime becomes locally possible

 doi:10.20944/preprints202605.1249.v1

when $\alpha_{\mathrm{eff},t} < \alpha_c$. The threshold does not state that AI use necessarily reduces semantic entropy. It states that, under the model assumptions, semantic diversity becomes vulnerable when effective independent novelty is insufficient to compensate for recursive modal amplification.

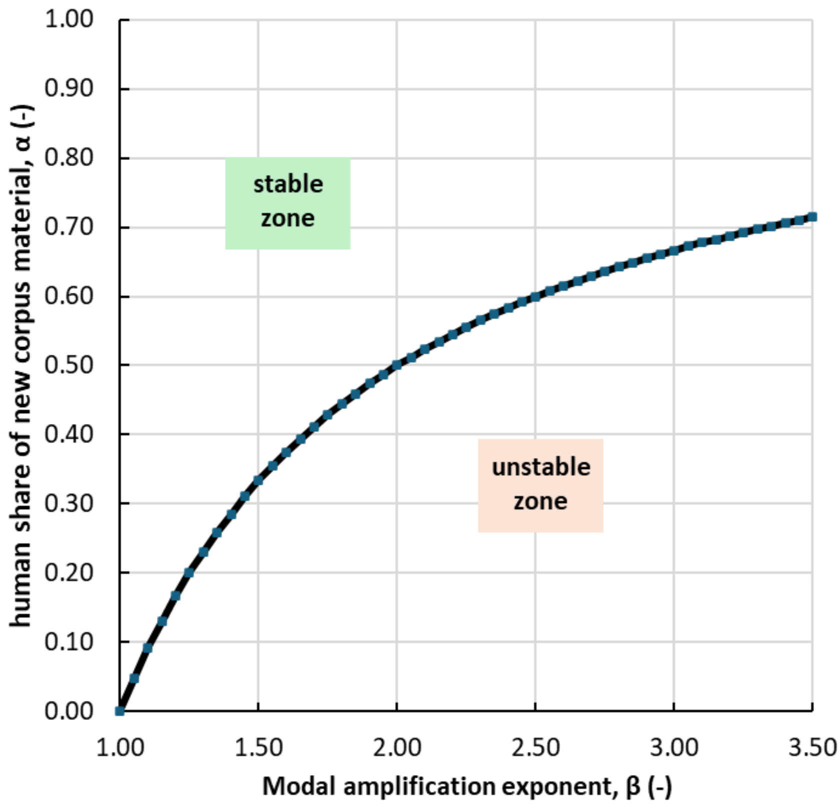

Figure 2. Critical effective novelty threshold as a function of modal amplification strength (the threshold $\alpha_c = 1 - 1/\beta$ separates the diversity-preserving regime from the homogenizing regime).

The critical novelty condition can also be inverted to obtain a stability requirement on the balance between AI-mediated and human/original semantic production. Using the definitions introduced above, the model defines a maximum AI-to-human semantic production ratio, $r_{\max}$, compatible with local semantic stability. Since $r = N_{AI}/N_H$, the reciprocal quantity, represents the minimum required human/original-to-AI semantic production ratio $(N_H/N_{AI})_{\min}$. This reciprocal form is useful because it directly expresses the stabilizing contribution required from independent human/original semantic production. This variable is represented versus assimilation constant λ in Figure 3 at different β values. Higher values mean that a larger amount of independent human/original semantic production is needed, relative to AI-mediated production, to preserve the diversity-preserving regime. In this result, λ is not yet specified as a time-dependent function. It is treated as an assimilation coordinate, while β controls the strength of modal amplification.

As λ increases, nominally human/original production becomes more aligned with AI-mediated semantic structure. Therefore, the maximum sustainable AI-to-human ratio $r_{\max}$ decreases, while the minimum required human/original-to-AI ratio increases. As β increases, modal amplification becomes stronger, so the stabilizing requirement also becomes more demanding. The curves diverge as λ approaches the critical assimilation value. Beyond this point, semantic stability cannot be preserved by changing the AI-to-human production ratio alone. The system would instead require lower assimilation, weaker modal amplification, or additional exogenous mechanisms of independent novelty injection.

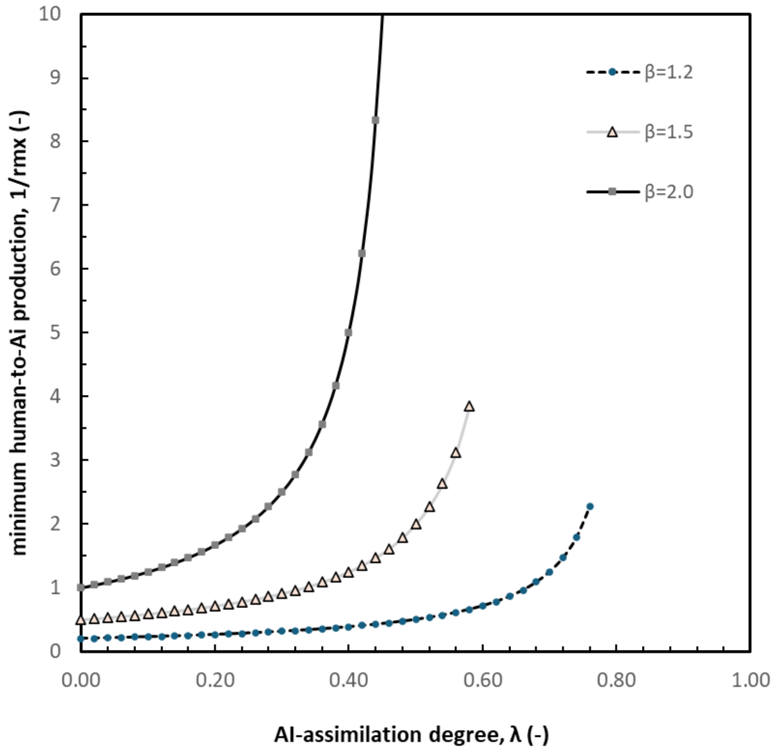

Figure 3. Minimum required human/original-to-AI semantic production ratio required for local stability, $(N_H/N_{AI})_{\min}$.

3.2. Scenario Analysis

The theoretical framework is explored through three representative scenarios over a normalized time horizon. The scenarios should be interpreted as controlled model experiments designed to illustrate how the semantic distribution evolves when the relative scale of AI-mediated semantic production changes under the assumptions of the model.

The scenarios differ only in the relative AI-to-human semantic production growth rate g_r. The slow-growth scenario (g_r=5%) represents gradual expansion of AI-mediated semantic production; the moderate-growth scenario (g_r=12%) represents sustained expansion sufficient to move the system toward the critical novelty threshold; the accelerated-growth scenario (g_r=25%) is interpreted as a stress case. in which AI-mediated production grows rapidly relative to human/original semantic production.

Figure 4a shows the trajectories of $\alpha_{\text{eff},t}$ and the assimilation parameter λ_t across the three scenarios. Only for the slow-growth scenario, effective independent novelty declines slowly but remains above the critical threshold for 14 simulated periods. In the others $\alpha_{\text{eff},t}$ crosses the threshold at time 4 and 7. The corresponding trajectories of λ_t , reported in Figure 4b, show how assimilation evolves under the assumed closure. This should not be interpreted as an empirical law of AI adoption, but as a modelling assumption used to explore the consequences of increasing AI-mediated production.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 19 May 2026 doi:10.20944/preprints202605.1249.v1

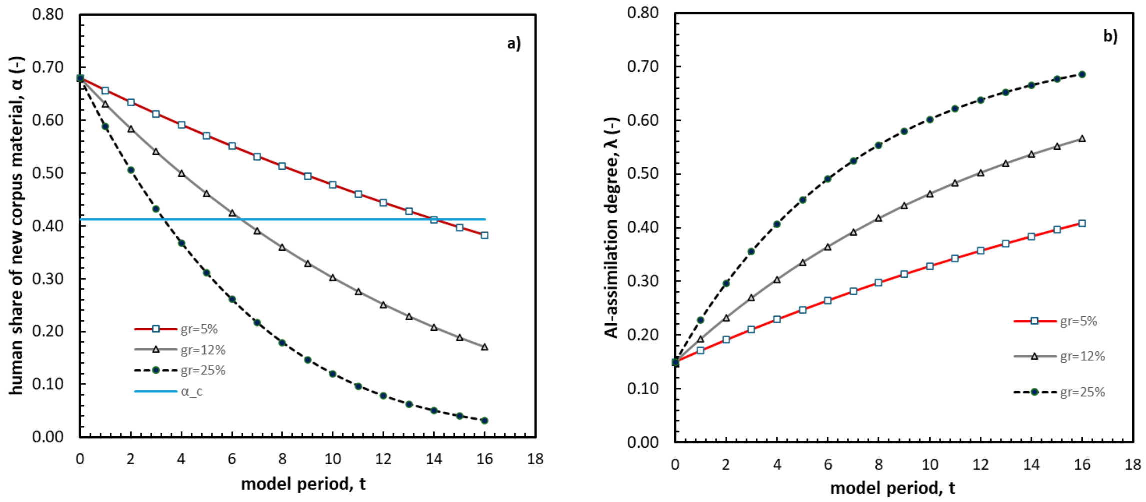

Figure 4. Effective novelty and AI assimilation across scenarios. (a) Trajectories of $\alpha_{\mathrm{eff},t}$ for slow, moderate and accelerated relative AI growth. The horizontal line indicates the critical threshold α_c. (b) Corresponding trajectories of the assimilation parameter λ_t, which represents the degree to which nominally human production becomes semantically aligned with AI-mediated distributions.

Figure 5a reports the semantic consequences of these transitions. In the slow-growth scenario, semantic entropy decreases only moderately, and the effective number of semantic states remains relatively high. In the "average" scenario, the delayed threshold crossing is followed by a substantial contraction of $\mathrm{H_{sem}}$ and $\mathrm{N_{eff}}$. In the final "stressed" scenario, semantic entropy contracts sharply and the system approaches a low-diversity attractor. Neff is reported in Figure 5b, and shows similar behavior. The next diagnostic is periodic semantic contraction, $\Delta H_{\mathrm{sem},t}$, defined as the entropy loss between two consecutive model periods. Figure 6 shows this non-monotonic behavior across the three scenarios. It can peak when the corpus still contains substantial semantic diversity but effective independent novelty has already declined enough for modal amplification to dominate. After this phase, $\Delta H_{\mathrm{sem},t}$ declines because less semantic diversity remains to be lost. The peak occurs later and remains smaller in the slow growth regime whereas, whereas it shifts earlier and becomes more pronounced in the others. This pattern indicates that the strongest marginal contraction occurs near the transition region, not necessarily at the highest level of AI-mediated production.

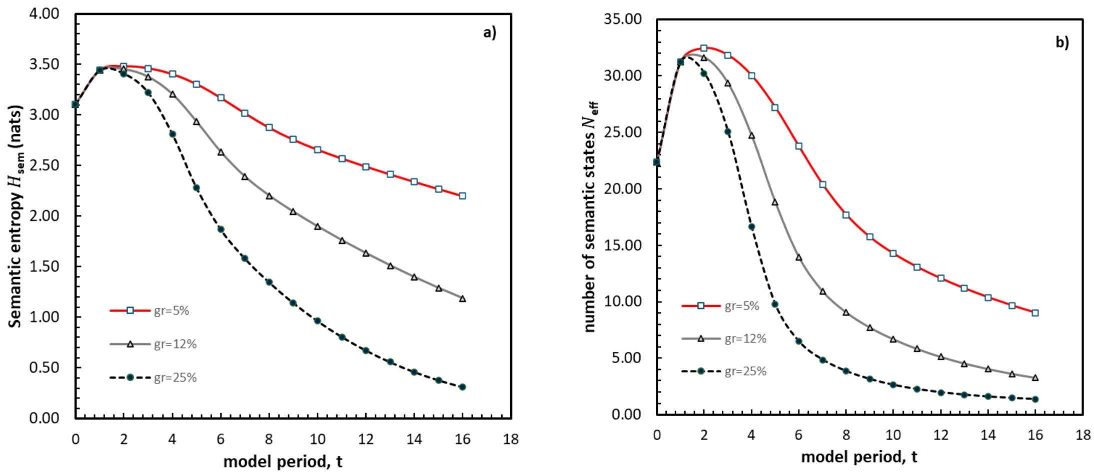

Figure 5. Semantic entropy and effective semantic states across normalized scenarios. (a) Semantic entropy $H_{\mathrm{sem},t}$. (b) Effective number of semantic states $N_{\mathrm{eff},t}$.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 19 May 2026 doi:10.20944/preprints202605.1249.v1

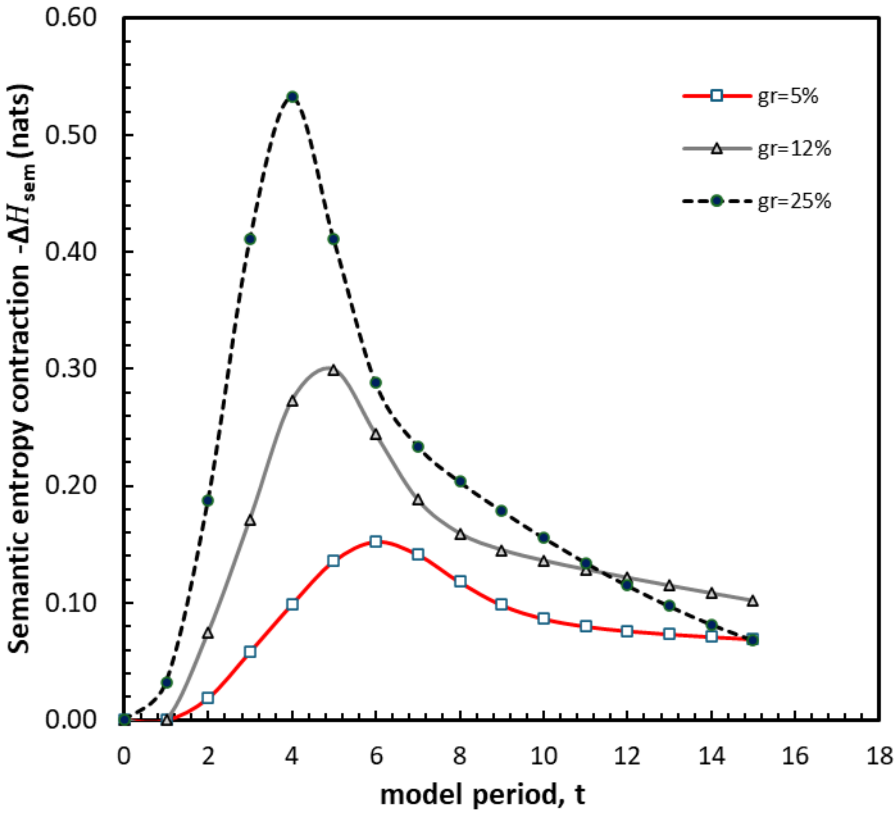

Figure 6. Periodic semantic contraction $\Delta H^{-}_{\text{sem},t}$ across normalized scenarios.

The global comparison across these scenarios identifies three theoretical regimes. The first regime is diversity preservation with mild concentration: AI-mediated semantic production grows and assimilation increases, but $\alpha_{\text{eff},t}$ remains above α_c; Semantic entropy may decline moderately, but the system does not enter the homogenizing regime.

The second regime is delayed threshold crossing: the system initially remains close to the diversity-preserving regime, but the combined effect of increasing AI-mediated production and increasing assimilation eventually pushes $\alpha_{\text{eff},t}$ below the critical threshold; after crossing, $H_{\text{sem},t}$ and $N_{\text{eff},t}$ decline more substantially.

The third regime is accelerated concentration: threshold crossing occurs earlier and the semantic distribution moves more rapidly toward a lower-diversity state within the model.

The main implication is not that more AI-mediated production necessarily implies lower semantic diversity. Rather, the model shows that semantic concentration becomes possible when modal amplification and assimilation reduce the stabilizing effect of independent novelty. These theoretical regimes are conceptually aligned with, but distinct from, the model-collapse literature. Shumailov et al. describe collapse as a degenerative process in which recursively generated data alter the training distribution of later models, with early loss of distributional tails and later reduction of variance [18]. The present framework does not model a sequence of trained models. It instead applies a related intuition to an evolving corpus-level semantic distribution: low-probability semantic regions can become under-represented when recursive AI-mediated production is not balanced by sufficient independent novelty.

The same distinction applies to self-consuming generative loops. Alemohammad et al. show that generative systems require sufficient fresh real data to avoid progressive loss of quality or diversity [20]. In the present model, this requirement appears as a condition on $\alpha_{\text{eff},t}$, the effective independent novelty rate. The relevant stabilizing quantity is not merely the amount of human-authored material, but the amount of semantic production that remains independent from AI-mediated assimilation.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 19 May 2026 doi:10.20944/preprints202605.1249.v1

The framework is also compatible with empirical work on AI-assisted ideation showing that individual productivity or fluency can increase while collective semantic distinctiveness decreases.

The model is intended also to illustrate a dissipative saturation pattern and prepare future data-fitting. For this reason, semantic contraction and physical dissipation have been studied within the same time period and related growth rate (see Eq 15). Figure 7 illustrates the asymmetry between physical and semantic entropy by comparing cumulative physical dissipation with cumulative semantic contraction. The figure should be read as a scenario diagnostic, not as an empirical estimate of the thermodynamic cost of meaning. Semantic contraction and physical dissipation do not scale linearly in the model. Periodic semantic contraction may peak near the transition region and then decline, even if cumulative energy expenditure continues to increase. This does not imply that physical entropy is converted into semantic entropy loss. It means only that, under the model assumptions, semantic concentration can saturate while computation remains dissipative.

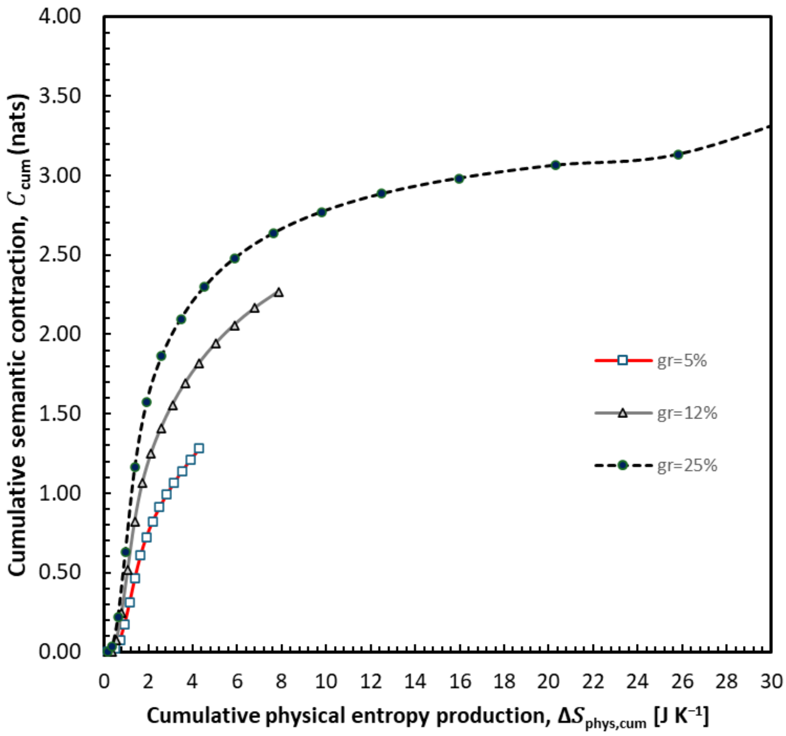

Figure 7. Dissipative semantic saturation across scenarios: cumulative semantic contraction vs. cumulative physical entropy generation.

The preceding figures report the qualitative behavior of the model in normalized time. Table 1 summarizes the scenario outcomes at the final normalized period T. The table reports the internal behavior of the model under the three parameterized scenarios and retains only the indicators needed to compare novelty loss, semantic concentration, and cumulative energy scale. The scenarios differ only in the relative AI-to-human semantic production growth rate g_r. $t_{\text{cross}}$is the first model period in which $\alpha_{\text{eff},t} < \alpha_c$. $t_{\text{peak}}$is the period of maximum periodic semantic contraction. Energy values represent scenario-attributed semantically relevant AI-mediated computation, not total global AI or data-center electricity demand. The table is illustrative of model behavior and should not be interpreted as an empirical forecast. Table 1 summarizes the corresponding quantitative outcomes at the end of the simulation horizon, $t_f = 20$. The table is not intended as a forecast. If one period is interpreted as one year, the same results can be read as an illustrative annual pathway; however, the model itself is defined in normalized time. In addition to threshold crossing and entropy outcomes,

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 19 May 2026 doi:10.20944/preprints202605.1249.v1

the table reports dissipative indicators that quantify semantic contraction per unit of scenario-attributed GenAI energy.

To illustrate how normalized model time could be mapped onto calendar time, consider only the moderate scenario S_M. If one model period were interpreted as one year and $t = 0$ were associated with 2025, the threshold crossing at $t_{\mathrm{cross}} = 8$would correspond indicatively to 2033. If one model period were instead interpreted as a five-year interval, the same normalized transition would correspond indicatively to 2065. This comparison shows that the calendar timing depends entirely on the chosen temporal calibration; the normalized model result is the threshold crossing at $t_{\mathrm{cross}} = 8$, not a forecast year.

Table 1. Quantitative summary of scenario outcomes under annual interpretation.

Indicator	S_L slow growth	S_M moderate growth	S_A accelerated growth
Relative AI/H growth, g_r	5%	12%	25%
Threshold crossing period, t_{cross}	no crossing	8	4
Peak contraction period, t_{peak}	7	5	4
Final effective novelty, α_{eff}	0.424	0.135	0.013
Final assimilation, λ	0.295	0.540	0.694
Final semantic entropy, H_{sem} (nat)	2.371	0.971	0.147
Final effective states, N_{eff}	10.71	2.64	1.16
Cumulative semantic contraction, C_{cum} (nat)	1.119	2.492	3.296
Cumulative heat rejected, Q_{rej} (TWh)	1257	2605	11658
Cumulative contraction per unit energy $\eta_{E,\mathrm{cum}}$ (nat·TWh^{-1})	8.90×10^{-4}	9.57×10^{-4}	2.83×10^{-4}

5. Conclusions

This paper proposed Dissipative Semantic Homogenization (SH) as a theoretical regime in which recursive generative AI ecosystems may continue dissipating physical energy while the effective diversity of corpus-level semantic states contracts, saturates, or reorganizes. The framework does not identify physical entropy production with semantic entropy. Physical entropy is associated with dissipated computational energy, whereas semantic entropy is defined as Shannon entropy over a corpus-level distribution of semantic states. The two quantities are coupled operationally because semantic distributions are transformed by physically implemented computation, but they remain conceptually and dimensionally distinct.

The model does not claim that every AI output is conventional or low-diversity. Rather, it captures the aggregate tendency of shared generative systems and reuse mechanisms to make high-probability semantic regions easier, faster, and cheaper to reproduce. The central result of the model is that semantic diversity is governed by effective independent novelty, not by nominal human authorship alone. Human/original production stabilizes semantic diversity only to the extent that it remains independent from AI-mediated semantic assimilation. Under the assumptions of the model, local semantic stability requires effective independent novelty to remain above a critical threshold determined by the strength of modal amplification. When this condition is not satisfied, semantic concentration becomes possible because recursive production, ranking, reuse, and assimilation can increase the probability mass of already dominant semantic regions.

The scenario analysis illustrates three theoretical regimes. In the slow-growth case, effective independent novelty remains above the critical threshold and semantic diversity is partially preserved. In the moderate-growth case, the system crosses the threshold after a delay and moves toward stronger semantic concentration. In the accelerated-growth case, the transition occurs earlier and the semantic distribution contracts more sharply. These scenarios are not forecasts. They are controlled model experiments that show how the balance among AI-mediated production,

 doi:10.20944/preprints202605.1249.v1

assimilation, modal amplification, and independent novelty can determine whether semantic diversity is preserved or reduced.

A second implication is that semantic contraction and physical dissipation need not scale linearly. Semantic entropy is bounded by the finite semantic partition, whereas scenario-attributed physical dissipation can continue as long as AI-mediated computation continues. Periodic semantic contraction may therefore peak near the transition region and then decline as the corpus approaches a more concentrated state. This dissipative saturation pattern does not imply conversion of physical entropy into semantic entropy loss. It indicates that computational dissipation may continue even after the main phase of semantic contraction has already occurred.

In conclusion, the model provides a theoretical framework for linking AI-mediated production, modal amplification, human/original novelty, semantic diversity, and physical dissipation. It offers a way to formulate future studies on AI-mediated semantic diversity: how semantic states should be partitioned, how effective independent novelty may be distinguished from nominal human authorship, how modal amplification can be operationalized, and how computational energy expenditure may be coupled to—but not identified with—semantic dynamics. DSH model offers a language for formulating future studies on AI-mediated semantic diversity: how semantic states should be partitioned, how effective independent novelty may be distinguished from nominal human authorship, how modal amplification can be operationalized, and how computational energy expenditure may be coupled to—but not identified with—semantic dynamics. In this sense, the framework should be read as a source of model-derived expectations and research questions rather than as a definitive empirical account of current AI ecosystems.

Future empirical work should operationalize semantic states using topic models, claim extraction, embedding-based clustering, or frame analysis; estimate effective novelty and assimilation and test whether threshold-like transitions, local homogenization, and dissipative saturation occur in specific domains.

Moreover, the framework should not be interpreted as an anti-AI argument. Generative AI can expand scientific and technological search spaces when its outputs are constrained by external validation, empirical testing, domain expertise, or independent data. Applications such as AI-assisted molecular or materials design illustrate a diversity-expanding regime in which generative models propose candidates and external constraints filter, falsify, or validate them.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Acknowledgments: During the preparation of this manuscript, the author used ChatGPT, primarily in the preliminary phase to support literature-oriented discussion, identify potentially relevant papers and contextual frameworks, compare the proposed model with existing research, and draft preliminary text passages. All AI-assisted outputs were reviewed, selected, edited, and substantially revised by the author, who takes full responsibility for the content of this publication.

Conflicts of Interest: The author declares no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

Symbol	Meaning	Unit
C_{cum}	Cumulative semantic contraction	nats
E	Electrical energy consumed by AI-mediate computation	kWh
g_{AI}	Derived AI-mediated production growth rate	(-)
g_H	Human/original production growth rate	(-)
g_r	Relative AI-to-human semantic production growth rate	(-)
H_{sem}	Semantic entropy of the corpus distribution	nats

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 19 May 2026 doi:10.20944/preprints202605.1249.v1

K	Number of semantic states in the corpus-level partition	(-)
N_{AI}	AI-mediated semantic production	Semantic production unit (SPU)
N_{eff}	Effective number of semantic states	(-)
N_H	human/original semantic production	Semantic production unit (SPU)
p_k	Probability of semantic state	(-)
p_{AI}	AI-mediated modal amplification operator	(-)
Q_{rej}	Heat rejected to the environment	kWh
q_k	Distribution of fresh human/original semantic novelty	(-)
r	AI-to-human semantic production ratio	(-)
T_0	Environmental temperature	K
t	Normalized model period	(-)
t_{cross}	Unstable-scenario crossing period	(-)
t_{peak}	Model period of maximum periodic semantic contraction	(-)
x_k	Semantic state (k)	(-)
α_c	Critical novelty threshold for local diversity stability	(-)
α_{eff}	Effective independent novelty rate	(-)
β	Modal amplification exponent	(-)
ΔH_{sem}	Periodic semantic contraction	nats
ΔS_{phys}	Physical entropy production	$J\ K^{-1}$
η_E	semantic contraction per unit entropy generation	nats $K\ J^{-1}$
κ	Sensitivity parameter of the assimilation function	(-)
λ	AI-assimilation degree in nominally human production	(-)

References

1. Noy, S.; Zhang, W. Experimental evidence on the productivity effects of generative artificial intelligence. Science **2023;** ,381,pp. 187-192.
2. Ayers, J.W.; Poliak, A.; Dredze, M. , Leas, E.C., Zhu, Z.; Kelley, J.B.; Faix, D.J., Goodman, A.M.; Longhurst, C.A.; Hogarth, M.; Smith, D.M.; Comparing Physician and Artificial Intelligence Chatbot Responses to Patient Questions Posted to a Public Social Media Forum. *JAMA Intern Med.* **2023**;183(6):589–596.
3. Kissinger, H. A., Schmidt, E., & Huttenlocher, D. P. The age of AI: And our human future. Little, Brown and Company **2021**.
4. Silver D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; Lillicrap, T.; Simonyan, K.; Hassabis, D. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play, Science, **2018** 1140-1144(2018).DOI:10.1126/science.aar6404
5. Stokes JM, Yang K, Swanson K, Jin W, Cubillos-Ruiz A, Donghia NM, MacNair CR, French S, Carfrae LA, Bloom-Ackermann Z, Tran VM, Chiappino-Pepe A, Badran AH, Andrews IW, Chory EJ, Church GM, Brown ED, Jaakkola TS, Barzilay R, Collins JJ. A Deep Learning Approach to Antibiotic Discovery. Cell. **2020**, 20;180(4):688-702.e13.
6. Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; Amodei, D. Language Models are Few-Shot Learners, 34th Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada NeurIPS, 2020.
7. Dell'Acqua, F.; McFowland, E.; Mollick, E.R.; Lifshitz-Assaf, H.; Kellogg, K.; Rajendran, S.; Krayer, L.; Candelon, F.; Lakhani, K.R. Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. SSRN Electronic Journal **2023**.
8. Peng, S.; Kalliamvakou, E.; Cihon, P.; Demirer, M. The impact of AI on developer productivity: Evidence from GitHub Copilot. arXiv **2023**, arXiv:2302.06590.
9. Krenn, M.; Pollice, R.; Guo, S.Y.; Aldeghi, M.; Cervera-Lierta, A.; Friederich, P.; Gomes, G.P.; Häse, F.; Jinich, A.; Nigam, A.K.; et al. On scientific understanding with artificial intelligence. Nature Reviews Physics **2022**, 4, 761–769.
10. Park, H.; Yan, X.; Zhu, R.; Huerta, E.A.; Chaudhuri, S.; Cooper, D.; Foster, I.; Tajkhorshid, E. A generative artificial intelligence framework based on a molecular diffusion model for the design of metal-organic frameworks for carbon capture. Communications Chemistry **2024**, 7, 21.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 19 May 2026 doi:10.20944/preprints202605.1249.v1

11. Venkatasubramanian, V. The promise of artificial intelligence in chemical engineering: Is it here, finally? AIChE Journal **2019**, 65, 466–478. doi:10.1002/aic.16489.
12. Vinuesa, R.; Azizpour, H.; Leite, I.; Balaam, M.; Dignum, V.; Domisch, S.; Felländer, A.; Langhans, S.D.; Tegmark, M.; Fuso Nerini, F. The role of artificial intelligence in achieving the Sustainable Development Goals. Nature Communications **2020**, 11, 233. doi:10.1038/s41467-019-14108-y.
13. Kaack, L.H.; Donti, P.L.; Strubell, E.; Kamiya, G.; Creutzig, F.; Rolnick, D. Aligning artificial intelligence with climate change mitigation. Nature Climate Change 2022, 12, 518–527. doi:10.1038/s41558-022-01377-7.
14. Messeri, L., Crockett, M. J. Artificial intelligence and illusions of understanding in scientific research. Nature **2024** 627, 49–58. https://doi.org/10.1038/s41586-024-07146-0
15. International Energy Agency. Energy and AI; IEA: Paris, France, 2025.
16. Strubell, E.; Ganesh, A.; McCallum, A. Energy and policy considerations for deep learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics; ACL: Florence, Italy, 2019; pp. 3645–3650. doi:10.18653/v1/P19-1355.
17. Luccioni, A.S.; Viguier, S.; Ligozat, A.-L. Estimating the carbon footprint of BLOOM, a 176B parameter language model. Journal of Machine Learning Research 2023, 24, 1–15.
18. Shumailov, I.; Shumaylov, Z.; Zhao, Y.; Papernot, N.; Anderson, R.; Gal, Y. AI models collapse when trained on recursively generated data. Nature 2024, 631, 755–759. https://doi.org/10.1038/s41586-024-07566-y
19. Shumailov, I.; Shumaylov, Z.; Zhao, Y.; Papernot, N.; Anderson, R.; Gal, Y. Author Correction: AI models collapse when trained on recursively generated data. Nature 2025, 640, E6. https://doi.org/10.1038/s41586-025-08905-3
20. Alemohammad, S.; Casco-Rodriguez, J.; Luzi, L.; Humayun, A.I.; Babaei, H.; LeJeune, D.; Siahkoohi, A.; Baraniuk, R.G. Self-Consuming Generative Models Go MAD. arXiv 2023, arXiv:2307.01850.
21. Anderson, B.R.; Shah, J.H.; Kreminski, M. Homogenization Effects of Large Language Models on Human Creative Ideation. In Creativity and Cognition (C&C ’24); ACM, 2024. https://doi.org/10.1145/3635636.3656204.
22. Farquhar, S.; Kossen, J.; Kuhn, L.; Gal, Y. Detecting hallucinations in large language models using semantic entropy. Nature 2024, 630, 625–630. https://doi.org/10.1038/s41586-024-07421-0.
23. Landauer, R. Irreversibility and Heat Generation in the Computing Process. IBM Journal of Research and Development 1961, 5, 183–191. https://doi.org/10.1147/rd.53.0183.
24. Bennett, C.H. The thermodynamics of computation—a review. International Journal of Theoretical Physics 1982, 21, 905–940. https://doi.org/10.1007/BF02084158.
25. Parrondo, J.M.R.; Horowitz, J.M.; Sagawa, T. Thermodynamics of information. Nature Physics 2015, 11, 131–139. https://doi.org/10.1038/nphys3230.
26. Wolpert, D.H. The stochastic thermodynamics of computation. Journal of Physics A: Mathematical and Theoretical 2019, 52, 193001. doi:10.1088/1751-8121/ab0850.
27. Still, S.; Sivak, D.A.; Bell, A.J.; Crooks, G.E. Thermodynamics of prediction. Physical Review Letters 2012, 109, 120604. doi:10.1103/PhysRevLett.109.120604.
28. Strubell, E.; Ganesh, A.; McCallum, A. Energy and Policy Considerations for Deep Learning in NLP. In Proceedings of ACL 2019, 3645–3650. https://doi.org/10.18653/v1/P19-1355.
29. Schwartz, R.; Dodge, J.; Smith, N.A.; Etzioni, O. Green AI. Communications of the ACM 2020, 63, 54–63. https://doi.org/10.1145/3381831.
30. Patterson, D.; Gonzalez, J.; Le, Q.; Liang, C.; Munguia, L.M.; Rothchild, D.; So, D.; Texier, M.; Dean, J. Carbon Emissions and Large Neural Network Training. arXiv 2021, arXiv:2104.10350.
31. Luccioni, A.S.; Jernite, Y.; Strubell, E. Power Hungry Processing: Watts Driving the Cost of AI Deployment? In ACM FAccT ’24, 2024. https://doi.org/10.1145/3630106.3658542.
32. Shannon, C.E. A mathematical theory of communication. Bell System Technical Journal 1948, 27, 379–423; 623–656.
33. Cover, T.M.; Thomas, J.A. Elements of Information Theory, 2nd ed.; Wiley: Hoboken, NJ, USA, 2006.
34. Doshi, A.R.; Hauser, O.P. Generative AI enhances individual creativity but reduces the collective diversity of novel content. Science Advances 2024, 10, eadn5290. doi:10.1126/sciadv.adn5290.

35. Burtch, G.; Lee, D.; Chen, Z. The consequences of generative AI for online knowledge communities. Scientific Reports 2024, 14, 10413. doi:10.1038/s41598-024-61221-0.
36. Bender, E.M.; Gebru, T.; McMillan-Major, A.; Shmitchell, S. On the dangers of stochastic parrots: Can language models be too big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency; ACM: New York, NY, USA, 2021; pp. 610–623. doi:10.1145/3442188.3445922.
37. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M.S.; et al. On the opportunities and risks of foundation models. arXiv 2021, arXiv:2108.07258.