

On the Limits of Self-Improving in LLMs and Why AGI, ASI and the Singularity Are Not Near Without Symbolic Model Synthesis

Hector Zenil^{1,2}

¹ Algorithmic Dynamics Lab, Department of Biomedical Computing, School of Biomedical Engineering and Imaging Sciences, King’s Institute for AI, King’s College London, UK

² Oxford Immune Algorithmics, Oxford University Innovation and London Institute for Healthcare Engineering, UK

Abstract

We formalise recursive self-training in Large Language Models (LLMs) and Generative AI as a discrete-time dynamical system and prove that, as training data become increasingly self-generated ($\alpha_t \rightarrow 0$), the system undergoes inevitably degenerative dynamics. We derive two fundamental failure modes: (1) *Entropy Decay*, where finite sampling effects cause a monotonic loss of distributional diversity (mode collapse), and (2) *Variance Amplification*, where the loss of external grounding causes the model’s representation of truth to drift as a random walk, bounded only by the support diameter. We show these behaviours are not contingent on architecture but are consequences of distributional learning on finite samples. We further argue that Reinforcement Learning with imperfect verifiers suffers similar semantic collapse. To overcome these limits, we propose a path involving symbolic regression and program synthesis guided by Algorithmic Probability. The Coding Theorem Method (CTM) allows for identifying generative mechanisms rather than mere correlations, escaping the data-processing inequality that binds standard statistical learning. We conclude that while purely distributional learning leads to model collapse, hybrid neurosymbolic approaches offer a coherent framework for sustained self-improvement.

Keywords: Large Language Models (LLMs), Model Collapse, Recursive Self-Improvement, Entropy Decay, Coding Theorem Method, Algorithmic Information Dynamics.

1 Introduction

The notion of a technological or AI Singularity, popularised by Vernor Vinge and Ray Kurzweil (12), posits a future inflection point where artificial intelligence surpasses human intellect, leading to an “intelligence explosion” of unforeseeable consequence (23; 13). Central to this hypothesis is the concept of **recursive self-improvement**: an AI system with the capacity to inspect and enhance its own architecture or training processes would initiate a positive feedback loop, with each generation of the AI being more intelligent than the last, leading to exponential growth in its capabilities.

The recent and remarkable successes of Generative Artificial Intelligence (GenAI), particularly Large Language Models (LLMs) like GPT-5 (15) and Diffusion Models for image synthesis (17), have reignited speculation about the proximity of this event. These models demonstrate an unprecedented ability to generate fluent text, create photorealistic images, and synthesise complex data, leading some to believe they are foundational steps towards Artificial General Intelligence (AGI). The assumption is that by scaling these models and enabling them to learn from the vast quantities of data they can generate, we might trigger the prophesied recursive improvement cycle.

This paper challenges this assumption directly. We argue that the very mechanism proposed for self-improvement—training on self-generated data—is, in fact, a pathway to self-destruction. This phenomenon, empirically observed and termed **model collapse** or the ‘curse of recursion’ (19), describes the progressive degradation of a model’s performance as its training data becomes increasingly polluted with its own synthetic outputs. Rather than ascending towards superintelligence, the model’s internal representation of the world contracts and distorts, converging towards a degenerate state of low diversity and high bias.

Our contribution is to move beyond empirical observation and provide a formal mathematical proof of the inevitability of model collapse. We model the self-referential training process as a dynamical system on the space of probability distributions and demonstrate that, under the condition of a diminishing supply of fresh, authentic data, this system is guaranteed to converge to a fixed point that is a distorted and impoverished version of the true data distribution. This conclusion is robust and applies not only to single LLMs but also to complex ecosystems of interacting models and multi-modal systems.

Furthermore, we contend that this mathematical limitation reflects a deeper philosophical boundary. Drawing upon Immanuel Kant’s distinction between analytic and synthetic judgements (10), we argue that current GenAI is fundamentally an **analytic** engine. It excels at analysing, recombining, and interpolating the vast patterns contained within its human-generated training data. It cannot, however, generate **synthetic** knowledge—truly novel concepts, laws, or truths that are not

simply derivative of its input. The Singularity requires a capacity for synthetic knowledge generation, which is absent in the current paradigm.

2 Preliminaries and Notation

This section fixes all notation and technical terms used in the manuscript. We separate (i) distributional objects and learning operators, (ii) information-theoretic quantities and inequalities, and (iii) algorithmic information estimators and perturbation-based mechanism analysis.

2.1 Loss functions, divergences, and empirical estimates

For a discrete distribution P on \mathcal{X} , Shannon entropy is $H(P) = -\sum_{x \in \mathcal{X}} P(x) \log P(x)$. We denote the Kullback–Leibler divergence by $D_{KL}(P \| Q) = \sum_x P(x) \log \frac{P(x)}{Q(x)}$.

In realistic training scenarios, the model does not have access to the full distribution Q_t or P , but only to a finite dataset $\mathcal{D}_t = \{x_1, \dots, x_N\}$ of size N sampled i.i.d. from the source distribution. We denote the *empirical distribution* formed by these samples as \hat{Q}_t .

To analyse dynamics in high-dimensional spaces (like the latent space of an LLM), we assume there exists a feature map $\phi : \mathcal{X} \rightarrow \mathbb{R}^d$ (e.g., the transformer embedding) and define the first moment (mean) of the distribution as $\mu_Q = \mathbb{E}_{x \sim Q}[\phi(x)]$. We denote the model update as an operator \mathcal{T} that maps a target distribution to a model parameterisation, minimising divergence:

$$Q_{t+1} = \arg \min_{Q \in \mathcal{Q}} D_{KL}(\hat{P}'_t \| Q), \quad (1)$$

where \hat{P}'_t is the empirical mixture derived from finite samples.

2.2 Mutual information, DPI and contraction inequalities

Throughout, $I(\cdot; \cdot)$ denotes Shannon mutual information. For discrete random variables U, V it is

$$I(U; V) = \sum_{u,v} p(u, v) \log \frac{p(u, v)}{p(u)p(v)}. \quad (2)$$

We use mutual information only to formalise limitations of *distribution-only* learning pipelines. For a Markov chain $M \rightarrow X \rightarrow Y$, the data-processing inequality (DPI) states

$$I(M; Y) \leq I(M; X). \quad (3)$$

We work with iterative operators that admit one-step bounds of the form

$$D_f^{(t+1)} \leq c D_f^{(t)} + \delta, \quad c \in (0, 1], \delta \geq 0, \quad (4)$$

where c is a contraction factor and δ captures approximation error. Iterating yields

$$D_f^{(n)} \leq c^n D_f^{(0)} + \sum_{i=0}^{n-1} c^i \delta, \quad (5)$$

so smaller c corresponds to stronger correction per iteration, up to an error floor induced by δ .

2.3 Algorithmic complexity, information and probability

Fix a universal prefix Turing machine U . For a computable object o , its Kolmogorov complexity is $K(o) = \min\{|p| : U(p) = o\}$. The algorithmic probability is $m(o) = \sum_{p:U(p)=o} 2^{-|p|}$. The Coding Theorem connects them: $-\log m(o) = K(o) + O(1)$.

2.4 The Coding Theorem and Block Decomposition Methods (CTM and BDM)

The Coding Theorem Method (CTM) approximates algorithmic probability by enumerating a reference class \mathcal{M} of small Turing machines:

$$\hat{m}_{\text{CTM}}(o) = \frac{1}{|\mathcal{M}|} \sum_{M \in \mathcal{M}} \mathbf{1}\{U_M \downarrow = o\}, \quad \text{CTM}(o) = -\log \hat{m}_{\text{CTM}}(o). \quad (6)$$

To scale beyond the small-object regime, BDM decomposes an object o into blocks of size k :

$$\text{BDM}_k(o) = \sum_i (\text{CTM}(b_i) + \log n_i), \quad (7)$$

where n_i is the multiplicity of block b_i .

2.5 Algorithmic Information Dynamics (AID)

AID quantifies the algorithmic causal effect of a perturbation τ by the change in complexity:

$$\Delta_\tau(o) = \text{BDM}_k(\tau(o)) - \text{BDM}_k(o). \quad (8)$$

2.6 Neurosymbolic operators

We define a one-step update as a composition of operators:

$$Q_{t+1} = \mathcal{T}_{\alpha_t} \circ \mathcal{C}_t \circ \Pi_S(Q_t),$$

where Π_S is symbolic projection, \mathcal{C}_t is causal correction, and \mathcal{T}_{α_t} is statistical fitting.

3 Background and Related Work

3.1 Recursive Self-Improvement

The modern conception of the technological Singularity is rooted in the idea of an “intelligence explosion,” first detailed by I. J. Good (7). Good argued that an “ultra-intelligent machine”—defined as a machine that can far surpass all the intellectual activities of any man however clever—would be the last invention that man need ever make, provided that the machine is docile enough to tell us how to keep it under control. The core of the logic is recursive: such a machine could design even better machines, which would in turn design even better ones, creating a positive feedback loop of rapidly accelerating intelligence.

This idea requires two key preconditions: (1) the existence of an AI that is intelligent enough to understand and modify its own source code or training methodology and (2) that these modifications lead to a consistent increase in its own general intelligence. Proponents of the Singularity, such as (13), view the exponential growth in computing power (Moore’s Law) as a direct trajectory towards this point. However, this view often conflates computational capacity with the architectural leaps required for genuine intelligence growth. The core question is not whether a machine can be made more powerful, but whether it can make *itself* more intelligent in a meaningful and unbounded way.

3.2 Generative AI as Distribution Learners

Modern GenAI models, irrespective of their modality, share a common mathematical foundation: they are designed to learn and sample from a complex, high-dimensional probability distribution.

Let \mathcal{X} be a data space (e.g., the space of all possible images, texts, or protein structures). We assume there exists a true, underlying data distribution $P(x)$ for $x \in \mathcal{X}$, from which real-world data is sampled. A generative model, parametrised by $\theta \in \Theta$, aims to learn an approximation of this distribution, denoted $Q_\theta(x)$.

- **Autoregressive Models (e.g., LLMs):** For a sequence of tokens $x = (x_1, \dots, x_L)$, these models learn the conditional probability of the next token given the preceding ones. The joint probability is factorised as $Q_\theta(x) = \prod_{i=1}^L Q_\theta(x_i | x_{<i})$. Training typically involves minimising the negative log-likelihood (cross-entropy) on a large corpus of text, which is equivalent to minimising the Kullback-Leibler (KL) divergence $D_{KL}(P||Q_\theta)$.
- **Diffusion Models:** These models learn to reverse a diffusion process that gradually adds noise to data. They define a sequence of latent variables that corrupt an initial data point $x_0 \sim P(x)$ into pure Gaussian noise x_T . The model then learns the reverse process $Q_\theta(x_{t-1} | x_t)$, allowing it to generate a sample x_0 by starting with random noise and iteratively denoising it. This process implicitly defines a complex distribution $Q_\theta(x)$.
- **Generative Adversarial Networks (GANs):** GANs (8) use a two-player game between a generator G that produces samples and a discriminator D that tries to distinguish them from real samples. The generator G implicitly defines the model distribution $Q_\theta(x)$. The feedback from the discriminator guides the generator to produce samples that are increasingly indistinguishable from the true data distribution $P(x)$.

In all cases, the goal is to make $Q_\theta(x)$ as close as possible to $P(x)$. The Singularity hypothesis, when applied to these models, implies that a model M_t with distribution Q_t could generate data to train a successor model M_{t+1} with distribution Q_{t+1} such that Q_{t+1} is a “better” approximation of P or represents a “more intelligent” distribution. Our paper will show this is not the case.

3.3 Model Collapse: The Curse of Degenerative Recursion

The theoretical foundation for our argument rests on the growing body of literature concerning model collapse. (19) provided a seminal empirical and theoretical study, showing that learning from data generated by other models causes a feedback loop that makes models “forget” the true underlying distribution. They demonstrate this for Gaussian Mixture Models and show empirically for LLMs and diffusion models that diversity is rapidly lost, with the tails of the original distribution being forgotten first.

This phenomenon is not new, though its implications for the Singularity are underexplored. In GAN literature, a similar issue known as “mode collapse” occurs when the generator learns to produce only a few distinct types of samples that can fool the discriminator, failing to capture the full diversity of the data distribution (2). The work in (1) refers to this as “self-consuming” loops, warning that the

proliferation of synthetic data on the internet could “contaminate” the training data for future models, leading to a gradual decay in their quality.

Previous studies on model collapse demonstrate empirically that generative models degrade when trained on synthetic data. By contrast, our contribution is not to show that collapse occurs, but to prove that it must occur under the autonomy condition required for recursive self-improvement. We formalise self-training as a dynamical system and show that, as external grounding vanishes, the system converges to a degenerate fixed point irrespective of architecture, modality, or ensemble structure.

Our work aims to formalise these observations into a rigorous proof of convergence, generalising the argument to any generative model and directly linking this inevitable collapse to the impossibility of a GenAI-driven Singularity-type argument based on current mainstream architectures.

4 Self-Referential Training as a Dynamical System

To analyse the dynamics of recursive self-improvement, we must first establish a formal mathematical framework. We generalise beyond specific architectures to any generative model that learns a probability distribution over a data space.

Definition 1 (Generative Model and Data Distributions). *Let \mathcal{X} be a measurable data space.*

- *Let P be the true data distribution over \mathcal{X} , representing authentic, high-fidelity data (e.g., all human-generated text and images).*
- *A generative model at iteration t , denoted M_t , is characterised by a probability distribution Q_t over \mathcal{X} , which it can sample from. We assume Q_t belongs to some family of distributions \mathcal{Q} representable by the model architecture.*
- *We assume an initial model M_0 with distribution Q_0 has been trained on samples drawn exclusively from P , such that $Q_0 \approx P$. Due to finite data and model capacity, $Q_0 \neq P$.*

Definition 2 (Self-Referential Training Loop). *The transition from model M_t to M_{t+1} occurs via a training process on a new dataset. This dataset is a mixture of authentic and synthetic data.*

- *Let $\alpha_t \in [0, 1]$ be the proportion of new, authentic data drawn from P at iteration t .*

- The remaining proportion, $1 - \alpha_t$, consists of synthetic data sampled from the current model’s distribution, Q_t .
- The training distribution for the next model, P'_t , is a convex combination:

$$P'_t = \alpha_t P + (1 - \alpha_t) Q_t \quad (9)$$

- The new model M_{t+1} is obtained by optimising its parameters to minimise the divergence between its distribution Q_{t+1} and the training distribution P'_t . A common objective is to minimise the KL-divergence:

$$Q_{t+1} = \arg \min_{Q \in \mathcal{Q}} D_{KL}(P'_t \parallel Q) \quad (10)$$

This is equivalent to maximising the log-likelihood of the data sampled from P'_t .

The Singularity hypothesis of recursive self-improvement corresponds to the case where this iterative process, predominantly driven by self-generated data (i.e., $\alpha_t \rightarrow 0$), leads to a sequence of distributions $\{Q_t\}$ that become progressively “better” or “more intelligent”. Our thesis is that this process instead leads to convergence towards a degenerate fixed point Q^* , where $D_{KL}(P \parallel Q^*) > 0$ and, critically, the entropy and diversity of Q^* are lower than that of P .

4.1 Generalisation to Different Architectures

This framework is deliberately abstract to encompass various GenAI paradigms:

- For LLMs, \mathcal{X} is the space of token sequences and Q_t is an autoregressive model. Self-referential training means fine-tuning the LLM on a mix of human text (P) and text generated by the LLM itself (Q_t).
- For Diffusion Models, \mathcal{X} is the space of images. The model learns a denoising function. Training on synthetic data means using images generated by the model in a previous iteration as clean examples for a new training run. The update rule (10) corresponds to re-optimising the denoising network on samples from the mixture distribution P'_t .
- For adversarial setups, such as GANs or models where a classifier provides a reward signal, the core logic holds. If the generator (the core generative component) is rewarded based on a classifier’s judgement, and that classifier is itself trained or fine-tuned on the generator’s outputs, a closed loop is formed. The ensemble’s knowledge becomes untethered from the external reality defined by P , leading to a similar collapse.

4.2 Proof of Convergence to a Degenerate State

The update rule in Equation (10) defines a discrete-time dynamical system on the space of probability distributions. Let T_α be the operator that maps Q_t to Q_{t+1} for a fixed $\alpha_t = \alpha$:

$$Q_{t+1} = T_\alpha(Q_t) = \arg \min_{Q \in \mathcal{Q}} D_{KL}(\alpha P + (1 - \alpha)Q_t || Q) \quad (11)$$

If the model family \mathcal{Q} is sufficiently expressive to represent any mixture (e.g., if \mathcal{Q} is the set of all probability distributions), then the solution is simply $Q_{t+1} = \alpha P + (1 - \alpha)Q_t$. This is an exponentially weighted moving average, and its dynamics are straightforward.

Proposition 1 (Convergence of the Idealised Update Rule). *Assuming the model family \mathcal{Q} has **infinite capacity** (i.e., can represent any distribution in the simplex of \mathcal{X}), if $Q_{t+1} = \alpha P + (1 - \alpha)Q_t$ with a constant $\alpha \in (0, 1]$, the sequence of distributions $\{Q_t\}_{t=0}^\infty$ converges to the true distribution P .*

Proof. By recursively expanding the update rule, we have:

$$\begin{aligned} Q_t &= \alpha P + (1 - \alpha)Q_{t-1} \\ &= \alpha P + (1 - \alpha)(\alpha P + (1 - \alpha)Q_{t-2}) \\ &= \alpha P \sum_{k=0}^{t-1} (1 - \alpha)^k + (1 - \alpha)^t Q_0 \\ &= \alpha P \frac{1 - (1 - \alpha)^t}{1 - (1 - \alpha)} + (1 - \alpha)^t Q_0 \\ &= (1 - (1 - \alpha)^t)P + (1 - \alpha)^t Q_0 \end{aligned}$$

As $t \rightarrow \infty$, since $\alpha \in (0, 1]$, we have $(1 - \alpha)^t \rightarrow 0$. Therefore, $\lim_{t \rightarrow \infty} Q_t = P$. *Remark:* This result relies entirely on the assumption that Q_{t+1} can perfectly capture the mixture. In reality, finite capacity introduces an approximation error δ at every step, which accumulates when α is small, as shown in Theorem 2. \square

4.3 The Case of Pure Self-Reference ($\alpha = 0$)

Let us now consider the crucial case for the Singularity hypothesis: a system that improves by learning exclusively from its own output. This corresponds to setting $\alpha = 0$.

Theorem 2 (Entropy Decay in Closed-Loop Training). *Let the training dataset \mathcal{D}_t at iteration t be a finite set of N samples drawn from Q_t . Let Q_{t+1} be the*

empirical risk minimiser over \mathcal{D}_t . In the absence of external ground truth ($\alpha = 0$) and assuming the model family \mathcal{Q} has sufficient capacity to overfit, the differential entropy of the model sequence decreases in expectation:

$$\mathbb{E}[H(Q_{t+1})] \leq H(Q_t) - \Delta(N), \quad (12)$$

where $\Delta(N) > 0$ is a strictly positive term representing information loss due to finite sampling and the discrete approximation of continuous or high-dimensional supports. Consequently, Q_t converges to a minimal-entropy distribution (a point mass or subset of modes) as $t \rightarrow \infty$.

Proof. Let \hat{Q}_t be the empirical distribution formed by N samples drawn from Q_t . By the properties of sampling from high-dimensional distributions, the support of \hat{Q}_t is a sparse subset of the support of Q_t . Specifically, for any distribution with tails (non-compact support), the probability that the finite sample support covers the true support is zero.

The model update Q_{t+1} minimises $D_{KL}(\hat{Q}_t \| Q)$. This is equivalent to maximising likelihood on the finite sample set. While standard maximum likelihood estimation is asymptotically unbiased, it exhibits finite-sample variance that manifests as overfitting to the sampled modes. The “missed” modes in \hat{Q}_t (events with probability $p < 1/N$) are effectively assigned zero probability mass in the empirical target.

Critically, $H(\hat{Q}_t) < H(Q_t)$ for finite N due to the discretisation of the sample space. Since Q_{t+1} is optimized to approximate \hat{Q}_t , the sequence of entropies forms a supermartingale: $\mathbb{E}[H(Q_{t+1}) | Q_t] \leq H(Q_t)$. By the Martingale Convergence Theorem, $H(Q_t)$ converges almost surely to a random variable with minimal entropy consistent with the fixed points of the update operator (i.e., mode collapse). \square

This proves that in a closed loop, no growth in ‘intelligence’ is possible as no new knowledge is generated. The system is information-theoretically closed. This can also be seen from the perspective of the Data Processing Inequality (5).

Corollary 3 (Information-Theoretic Stagnation). *The self-referential training loop cannot increase the mutual information with the true distribution P .*

Proof. Consider the Markov chain $P \rightarrow Q_t \rightarrow Q_{t+1}$. The data processing inequality states that for any Markov chain $X \rightarrow Y \rightarrow Z$, we have $I(X; Z) \leq I(X; Y)$. In our case, this means $I(P; Q_{t+1}) \leq I(P; Q_t)$. The mutual information between the model’s state and the true state of the world can only decrease or stay the same with each iteration of self-training. Any imperfection in Q_t (i.e., information about P that Q_t has lost) cannot be recovered by training on samples from Q_t . \square

4.4 The Realistic Case ($\alpha_t \rightarrow 0$)

The most realistic scenario for a purported Singularity is one where an AI starts with access to human data but gradually becomes more autonomous, causing the proportion of authentic data α_t to approach zero over time. We will now show that this leads to model collapse.

Theorem 4 (Variance Amplification and Mean Shift). *Let the true distribution be P and the model update be $Q_{t+1} = \mathcal{T}(P'_t) + \epsilon_t$, where \mathcal{T} is the ideal update operator and ϵ_t is an approximation error term (due to SGD noise and finite sampling) with variance σ_ϵ^2 . If $\alpha_t \rightarrow 0$, the squared error between the model mean $\mu_t = \mathbb{E}_{x \sim Q_t}[x]$ and the true mean μ_P diverges or follows a random walk bounded only by the support diameter.*

Proof. Consider the mean of the distribution μ_t . The update rule for the mixture $P'_t = \alpha_t P + (1 - \alpha_t)Q_t$ implies the target mean is $\mu'_t = \alpha_t \mu_P + (1 - \alpha_t)\mu_t$. The new model learns this mean with some error: $\mu_{t+1} = \mu'_t + \xi_t$, where ξ_t is a noise term corresponding to ϵ_t . Substituting the target mean:

$$\mu_{t+1} = (1 - \alpha_t)\mu_t + \alpha_t \mu_P + \xi_t. \quad (13)$$

This describes an autoregressive process AR(1). As $\alpha_t \rightarrow 0$, the autoregressive coefficient $(1 - \alpha_t) \rightarrow 1$. The process approaches a random walk: $\mu_{t+1} \approx \mu_t + \xi_t$. Unlike the case where $\alpha > 0$ (which provides a restoring force pulling μ_t back to μ_P), the condition $\alpha_t \rightarrow 0$ removes the restoring force. The variance of the mean $\text{Var}(\mu_t)$ grows linearly with t in the random walk regime. Thus, the model distribution centre drifts away from the true distribution centre purely due to accumulated stochastic errors, confirming model collapse not just as mode-dropping, but as distributional drift. \square

The parameter $\alpha_t \in [0, 1]$ denotes the proportion of *fresh, externally grounded* data drawn from the true distribution P at iteration t . The condition $\alpha_t \rightarrow 0$ means $\lim_{t \rightarrow \infty} \alpha_t = 0$, i.e. for every $\varepsilon > 0$ there exists T such that $\alpha_t < \varepsilon$ for all $t \geq T$. Equivalently, the training mixture

$$P'_t = \alpha_t P + (1 - \alpha_t)Q_t$$

becomes asymptotically self-referential, with the influence of P vanishing in the limit. This regime formalises the *autonomy requirement* implicit in Singularity-style recursive self-improvement arguments: the system must eventually rely predominantly on its own outputs rather than on a persistent external oracle.

Remark 5 (The Limits of Reinforcement Learning and Verifiers). *It is often argued that Reinforcement Learning (RL) with a verifier (e.g., a game engine or*

compiler) allows for self-improvement without external data ($\alpha = 0$), as seen in systems like AlphaZero. However, this relies on the verifier providing a perfect, infinite-precision ground truth signal. In the context of AGI and LLMs operating in open-ended domains (language, reasoning, physical reality), no such perfect verifier exists.

If the verifier is itself a learned model (e.g., a Reward Model in RLHF), it is subject to the same collapse dynamics described in Theorem 2. If the verifier is a static proxy (e.g., string matching or simple heuristics), the model will exploit the metric (Goodhart’s Law), leading to **semantic collapse**, where the model optimises the proxy to the detriment of the underlying complexity. Thus, the impossibility results hold for any system where the ground truth is not strictly encoded in a formal, executable environment.

Lemma 6 (Asymptotic self-reference of the training distribution). *If $\alpha_t \rightarrow 0$, then for any fixed t the mixture distribution satisfies*

$$\|P'_t - Q_t\|_{\text{TV}} \leq \alpha_t \|P - Q_t\|_{\text{TV}} \leq \alpha_t,$$

and hence $\|P'_t - Q_t\|_{\text{TV}} \rightarrow 0$ as $t \rightarrow \infty$, where $\|\cdot\|_{\text{TV}}$ denotes total variation distance.

Proof. By linearity of mixtures, $P'_t - Q_t = \alpha_t(P - Q_t)$, so $\|P'_t - Q_t\|_{\text{TV}} = \alpha_t\|P - Q_t\|_{\text{TV}} \leq \alpha_t$ since $\|P - Q_t\|_{\text{TV}} \leq 1$ for probability measures. Taking $t \rightarrow \infty$ and using $\alpha_t \rightarrow 0$ completes the proof. \square

4.5 Extension to Multi-Modal Ensembles

One might think that an ensemble of multi-modal models could prevent collapse. This is true until the new training data set is exhausted. Let us formalise this scenario and show that it eventually suffers the same fate.

Definition 3 (Multi-Modal Training). *Consider an ensemble of N models including possible different data types $\{M_1, \dots, M_N\}$ with corresponding distributions $\{Q_t^1, \dots, Q_t^N\}$. The synthetic data are drawn from a mixture of these models:*

$$R_t = \sum_{i=1}^N \omega_i Q_t^i, \quad \text{where } \sum \omega_i = 1, \omega_i \geq 0. \quad (14)$$

Each model M_j is then trained on a distribution $P'_{t,j} = \alpha_t P + (1 - \alpha_t) R_t$.

Theorem 7 (Collapse of Multi-Model Ensembles). *As $\alpha_t \rightarrow 0$, the ensemble of models converges to a consensus fixed-point distribution $R^* \neq P$. The individual distributions $\{Q_i^*\}$ may not be identical, but the mixture they form will be a stationary, degenerate distribution.*

Proof. As $\alpha_t \rightarrow 0$, each model M_j is being trained to imitate the ensemble’s average distribution from the previous step, R_{t-1} .

$$Q_t^j \approx \arg \min_{Q \in \mathcal{Q}_j} D_{KL}(R_{t-1} \parallel Q) \quad (15)$$

The entire system’s state is now described by the mixture distribution R_t . The update rule for the mixture is $R_t \approx \sum \omega_i \arg \min_{Q \in \mathcal{Q}_i} D_{KL}(R_{t-1} \parallel Q)$. This is still a closed loop. The ensemble is learning from itself. Any shared biases or errors across the models will be reinforced. Any aspect of the true distribution P that is collectively under-represented by the initial ensemble R_0 will be progressively forgotten as $\alpha_t \rightarrow 0$. The system converges to a fixed-point mixture R^* that represents the consensus reality of the initial models, not the true reality P . Diversity may be lost more slowly than in a single-model case, but the fundamental information-theoretic barrier remains: no new information about P can be created from within the closed system. \square

5 Implications for AGI, ASI and Singularity arguments

One core premise of arguments around Artificial General Intelligence (AGI), Artificial Super Intelligence (ASI) and the so-called AI ‘Singularity’ is a recursive process whereby unbounded growth is possible without external or human intervention. Our analysis shows that the self-referential training process, far from being unbounded, is a convergent process. It leads not to an “explosion” but to an “implosion” of informational diversity and knowledge inference.

- **Convergence vs. Divergence:** A Singularity requires a divergent process, where capability grows exponentially. Model collapse is a convergent process, where the model’s distribution approaches a static, limited fixed point.
- **The Necessity of Grounding:** The theorems demonstrate that without a persistent connection to an external, truthful data source (P), an AI system’s model of reality will inevitably drift and degrade. The value $\alpha > 0$ is not a temporary bootstrap but a permanent necessity. AGI (9) cannot be developed in a closed box; it requires continuous, active grounding in the real world.
- **Errors are Amplified, Not Corrected:** In a self-referential loop, any inaccuracies, biases, or “hallucinations” present in model Q_t become codified as “truth” in the training data for model Q_{t+1} . The system lacks an external error correction mechanism and will progressively amplify its own flaws.

5.1 A Fundamental Kantian Distinction: Analytic vs. Synthetic AI

To understand the philosophical limits of GenAI, we can borrow from Immanuel Kant’s epistemology. Kant distinguished between two types of judgements:

- **Analytic Judgements:** The predicate is contained within the concept of the subject (e.g., “All bachelors are unmarried”). These judgements are explicative; they do not add new knowledge but clarify what is already known via logical implication.
- **Synthetic Judgements:** The predicate is not contained within the concept of the subject (e.g., “The cat is on the mat” or “ $7 + 5 = 12$ ”). These require empirical observation or, crucially, **active computation** to verify.

We can apply this framework to AI. Current Generative AI is a fundamentally **analytic** system. It is trained on a massive, but finite, dataset representing a snapshot of human knowledge (P). Its operations consist of identifying patterns, correlations, and structures within this dataset and then interpolating or recombining them to generate outputs. While the outputs may be surprising, they are ultimately derivative of the correlations latent within the training data (the “subject”).

A true AGI, and certainly one capable of initiating a Singularity, must be capable of making **synthetic** judgements. In a computational context, this corresponds to **execution** rather than prediction. An LLM predicts the output of a Python script based on its training data (analytic); a Symbolic engine *runs* the script (synthetic). The latter generates a result that is not statistically entailed by the input but is the result of an irreducible computational process.

Model collapse is the mathematical manifestation of this philosophical limit. When an analytic engine is forced to feed on its own outputs, it has nothing new to analyse. It can only re-process and re-combine its existing knowledge, leading to a caricature of its former self. It cannot synthesise new knowledge because it lacks the mechanism of *computational verification* or external grounding, and therefore it cannot “improve” in any meaningful sense.

5.2 Causality and Neurosymbolic Approaches as a Path Beyond Collapse

The results presented show that recursive self-training with objectives grounded in Kullback–Leibler (KL) divergence leads to model collapse. It is important to emphasise that KL divergence is not an arbitrary choice: it underlies nearly all loss functions used in current deep learning, particularly in the transformer family

of models. Cross-entropy minimisation is equivalent to minimising $D_{KL}(P || Q)$, and maximum likelihood estimation also reduces to this divergence under common assumptions. Thus, our analysis captures the essential training dynamics of statistical deep learning without loss of generality.

Nearly all practical loss functions used in Large Language Models (LLMs) and related deep learning architectures reduce to KL divergence or very close relatives. This establishes that our collapse results apply to the entire class of current statistical deep learning approaches without loss of generality.

- **Cross-Entropy Loss.** The token-level cross-entropy used in transformers is exactly equivalent to the KL divergence between the true data distribution P and the model distribution Q . Specifically,

$$H(P, Q) = H(P) + D_{KL}(P || Q),$$

where $H(P)$ is constant with respect to Q . Minimising cross-entropy is therefore equivalent to minimising KL divergence.

- **Mean Squared Error (MSE).** In regression, MSE minimisation is equivalent to maximum likelihood estimation under Gaussian noise assumptions. Maximum likelihood corresponds to minimising the KL divergence between the empirical data distribution and the Gaussian model distribution, making MSE a special case of KL.
- **Binary Cross-Entropy / Logistic Loss.** Logistic regression minimises the cross-entropy between Bernoulli distributions, which again reduces to KL divergence between P (empirical labels) and Q (model predictions).
- **Categorical Cross-Entropy.** Multiclass classification with softmax uses the categorical cross-entropy, which is the KL divergence between a one-hot true label distribution and the predicted probability distribution.
- **Other Variants.** Many widely used losses (e.g., focal loss, label smoothing) are weighted or regularised versions of cross-entropy, and thus still within the KL framework. Generative Adversarial Networks (GANs) minimise Jensen–Shannon divergence, which belongs to the broader f -divergence family of which KL is a canonical member. Even Wasserstein distances, while distinct, are less commonly used in LLM training and do not alter the essential correlation-based character of the optimisation.

Thus, KL divergence provides a unifying mathematical framework that encompasses the most common objectives in deep learning. Our collapse proofs therefore apply broadly to existing transformer and generative model training regimes.

5.3 Statistical Loss Functions Contribute to Collapse

The reliance on KL means that optimisation is driven by correlation, not causation. Models trained under KL minimise distributional divergence but cannot infer or extrapolate underlying mechanisms. This explains why recursive training contracts distributions, amplifies biases, and converges to degenerate fixed points: KL training lacks a mechanism to preserve diversity beyond what is already present in the data mixture.

An alternative is to replace correlation-based optimisation with causal or algorithmic objectives. Instead of training future models as

$$Q_{t+1} = \arg \min_Q D_{KL}(P'_t \parallel Q),$$

we could envisage

$$Q_{t+1} = \arg \min_Q D_{\text{causal}}(M(P'_t), M(Q)),$$

where $M(\cdot)$ maps raw distributions to causal models, e.g., structural causal graphs. Optimising causal divergence aligns models with invariances under intervention, not just statistical regularities. This allows even synthetic data to yield new knowledge if it encodes counterfactual or interventional information.

Such neurosymbolic architectures, combining statistical learners with causal inference and algorithmic probability (as in CTM and BDM), could in principle break the collapse dynamics. Unlike KL-based systems, they are capable of producing genuinely *synthetic knowledge*, rather than endlessly recombining correlations. However, current LLMs remain confined to KL-like objectives and thus to the collapse trajectory we have proven.

If the incorporation of *causal* and *algorithmic* principles for the update rule is reformulated not as

$$Q_{t+1} = \arg \min_Q D_{KL}(P'_t \parallel Q),$$

but rather as

$$Q_{t+1} = \arg \min_Q D_{\text{causal}}(M(P'_t), M(Q)),$$

where $M(\cdot)$ denotes a causal representation of a distribution, for example, a structural causal model or a symbolic regression (24; 26), then the system optimises not for correlation but for invariance under intervention. In this framework, even synthetic data can yield new information if they encode counterfactual or interventional predictions. Unlike KL-based training, which contracts toward fixed points, causal objectives permit the creation of genuinely *synthetic knowledge*, in line with Kant’s distinction between analytic and synthetic judgments.

5.4 Causal correction and rate of (non-)collapse

We now formalise how the *rate* of convergence to collapse (or escape from it) depends on the quality of (i) the causal component and (ii) the symbolic component in a neurosymbolic pipeline. Throughout, P denotes the target (real-world) distribution, Q_t the model distribution after iteration t , and $\alpha_t \in [0, 1]$ the fraction of fresh (P -drawn) data used at iteration t . We work with an f -divergence $D_f(\cdot\|\cdot)$ that includes KL as a special case; all arguments below therefore cover the standard cross-entropy/negative log-likelihood training used by LLMs and most deep models.

Operators. Let the one-step update be a composition of three operators acting on distributions:

$$Q_{t+1} = \underbrace{\mathcal{T}_{\alpha_t}}_{\text{statistical learner}} \circ \underbrace{\mathcal{C}_t}_{\text{causal correction}} \circ \underbrace{\Pi_{\mathcal{S}}}_{\text{symbolic projection}} (Q_t),$$

where:

1. \mathcal{T}_{α_t} is the statistical (likelihood) update to fit the mixture $P'_t = \alpha_t P + (1 - \alpha_t)Q_t$ (Sec. 3). In the well-specified idealisation one has $\mathcal{T}_{\alpha_t}(Q) = \arg \min_{Q'} D_{\text{KL}}(P'_t \| Q')$.
2. $\Pi_{\mathcal{S}}$ is a projection onto a feasible set \mathcal{S} of symbolic/axiomatic constraints (e.g., conservation, monotonicity, type rules, physical/biological invariants). Formally, $\Pi_{\mathcal{S}}(Q) = \arg \min_{R \in \mathcal{S}} D_f(R \| Q)$.
3. \mathcal{C}_t is a causal-correction operator built from interventional queries; e.g., it replaces or reweights some conditionals $Q(X | \text{do}(Z))$ using estimates derived from a causal model \mathcal{G} and interventional data. We quantify its *strength* by how much it contracts the divergence to P on the parts it corrects.

Quantifying symbolic and causal “power”. We define two *per-iteration* contraction factors:

$$\text{(Symbolic)} \quad D_f(P \| \Pi_{\mathcal{S}}(Q)) \leq \underbrace{\sigma}_{\in(0,1]} D_f(P \| Q) + \delta_s, \quad (16)$$

$$\text{(Causal)} \quad D_f(P \| \mathcal{C}_t(R)) \leq \underbrace{\kappa_t}_{\in(0,1]} D_f(P \| R) + \delta_{c,t}, \quad (17)$$

where $0 < \sigma \leq 1$ captures the *strength of symbolic constraints* (smaller is better), and $0 < \kappa_t \leq 1$ captures the *effective causal correction* at iteration t (again, smaller

is better). The additive terms $\delta_s, \delta_{c,t} \geq 0$ capture model/estimation imperfections (finite data, solver tolerance, partial coverage of the graph, etc.). In practice,

$$\kappa_t = 1 - \eta_c \phi_t, \quad \phi_t \in [0, 1]$$

where $\eta_c \in (0, 1]$ encodes identification strength of the causal queries (e.g., instrument strength, overlap, do-operator availability), and ϕ_t is the *coverage* of interventional updates at step t (fraction of conditionals actually corrected).

The contraction factor $\sigma < 1$ is not arbitrary but is a direct consequence of the Coding Theorem. Let \mathcal{H} be the total hypothesis space of size $|\mathcal{H}|$. A purely statistical learner effectively searches this entire space for correlations. A symbolic learner projects onto a subspace $\mathcal{S} \subset \mathcal{H}$ defined by programs p with Kolmogorov complexity $K(p) \leq L$. By the Coding Theorem, the probability mass is concentrated on simple programs: $m(x) \approx 2^{-K(x)}$. If the true mechanism M has low complexity ($M \in \mathcal{S}$), the projection $\Pi_{\mathcal{S}}$ eliminates the vast majority of the search space (high-complexity noise and overfitting candidates). The reduction in the volume of the search space corresponds to an information gain of approximately $|\mathcal{H}| - |\mathcal{S}|$ bits. In terms of divergence, this forces the model distribution Q to align with the algorithmic structure of P , ensuring that

$$D_f(P \parallel \Pi_{\mathcal{S}}(Q)) \ll D_f(P \parallel Q),$$

which implies an effective contraction $\sigma \ll 1$ whenever the data is generated by a computable process.

The quantities σ and κ_t in Eqs. (16) and (17) are *contraction factors*. Their magnitude determines the rate at which iterative updates reduce discrepancy with respect to the target distribution or mechanism. Throughout, smaller values correspond to stronger correction.

Formally, consider a generic inequality of the form

$$D_f^{(t+1)} \leq c D_f^{(t)} + \delta, \quad c \in (0, 1]. \quad (18)$$

Iterating (18) yields

$$D_f^{(n)} \leq c^n D_f^{(0)} + \sum_{i=0}^{n-1} c^i \delta, \quad (19)$$

which shows that the discrepancy decays exponentially at rate c up to an error floor determined by δ . When $c \ll 1$, convergence is fast; when $c \approx 1$, improvement per iteration is negligible; and when $c = 1$, no contraction occurs.

In Eq. (16), the factor σ quantifies the *strength of symbolic constraints*. The projection $\Pi_{\mathcal{S}}$ enforces syntactic, grammatical, or invariant structure, thereby eliminating large regions of the hypothesis space. A small value of σ indicates that

symbolic structure is highly informative, leading to a rapid collapse toward low-complexity representations. From an algorithmic information perspective, this corresponds to strong compression of admissible descriptions, consistent with CTM-based estimation.

In Eq. (17), the factor κ_t measures the *effective causal correction* at iteration t . Writing

$$\kappa_t = 1 - \eta_c \phi_t,$$

makes explicit that causal contraction improves with both the identification strength η_c of interventions or perturbations and their coverage ϕ_t . Smaller values of κ_t indicate that causal updates eliminate large classes of incompatible mechanisms per iteration, reflecting strong perturbation coherence and high causal informativeness.

By contrast, the statistical update in Eq. (20) induces contraction only in distribution space. Although the factor $(1 - \alpha_t)$ also yields contraction in the sense of Eq. (19), this improvement is limited by the data-processing inequality and cannot increase information about the underlying generative mechanism. Statistical contraction therefore represents a weak, degenerate regime in which learning reduces to smoothing or reweighting existing distributional information.

Taken together, these observations establish a hierarchy. Statistical updates contract discrepancies between distributions but cannot recover mechanisms. Symbolic updates contract hypothesis space more aggressively by exploiting structural constraints. Causal updates operating in program space provide the strongest contraction by directly eliminating incompatible generative mechanisms. Smaller contraction factors thus correspond to stronger explanatory and corrective power per iteration.

Statistical learner under synthetic drift. Unlike the symbolic and causal updates above, this step does not perform causal correction. It characterises the behaviour of a learner restricted to distributional information only, as arises under synthetic drift or when algorithmic structure cannot be accessed.

For the likelihood update on $P'_t = \alpha_t P + (1 - \alpha_t)Q_t$, standard information-geometry arguments yield, for common choices of D_f including KL:

$$D_f(P \parallel \mathcal{T}_{\alpha_t}(S)) \leq (1 - \alpha_t) D_f(P \parallel S) + \delta_{\text{stat},t}, \quad (20)$$

where $\delta_{\text{stat},t} \rightarrow 0$ in the well-specified, infinite-data limit; otherwise it scales with capacity/optimisation error. In other words, even when statistical learning converges optimally, it may converge to a causally meaningless solution.

Proposition 8 (DPI bound for distribution-only learning and the resulting causal hierarchy). *Let \mathcal{A} be any distribution-only learner/estimator that maps an observed sample (or an empirical distribution) to a model distribution,*

$$\mathcal{A} : \hat{P} \mapsto Q_{\theta(\hat{P})},$$

and suppose the learning pipeline factors through a Markov chain of the form

$$M \longrightarrow X \longrightarrow \widehat{P}(X) \longrightarrow \theta(\widehat{P}) \longrightarrow Q_\theta,$$

where M denotes the (unknown) data-generating mechanism and X the observed data. Then the mutual information between the mechanism and the learned model is bounded by the data-processing inequality:

$$I(M; Q_\theta) \leq I(M; X). \quad (21)$$

In particular, no distribution-only learning update—including one driven by minimising KL or any f -divergence—can increase the information the learner has about the generating mechanism beyond what is already contained in the observations.

Moreover, if the learner’s update further compresses X into a finite-dimensional summary (e.g. empirical moments, sufficient statistics under a misspecified family, or a finite-capacity parametric model), then typically

$$I(M; Q_\theta) < I(M; X), \quad (22)$$

with strict inequality whenever the map $X \mapsto Q_\theta$ is non-invertible on the support induced by M .

Proof. The assumed pipeline defines the Markov chain

$$M \rightarrow X \rightarrow Q_\theta,$$

because Q_θ is a (possibly randomised) function of X through the intermediate computations $\widehat{P}(X)$ and $\theta(\widehat{P})$. By the data-processing inequality for mutual information (5),

$$I(M; Q_\theta) \leq I(M; X),$$

which proves (21). If $X \mapsto Q_\theta$ is not one-to-one (equivalently, if there exist distinct $x \neq x'$ with $Q_{\theta(x)} = Q_{\theta(x')}$ on a set of non-zero probability under the joint distribution induced by M), then information is lost and the inequality is strict, yielding (22). \square

Proposition 8 makes the hierarchy in this paper explicit: purely statistical updates contract distances between *distributions* (e.g. via KL) but cannot increase information about *mechanisms*. Symbolic and causal updates that operate in program/mechanism space are not constrained in the same way, because they do not merely post-process X ; they introduce additional mechanistic hypotheses and perturbation-coherence constraints that are not representable as a distribution-only map.

5.5 From Statistical to Algorithmic Integration

We now formalise how CTM and BDM allow us to transcend the limitations of statistical learning derived in Section 5.6.

Proposition 9 (Escaping DPI via Universal Priors). *While purely statistical learning is bound by the Data Processing Inequality (DPI) such that $I(M; Q_{t+1}) \leq I(M; Q_t)$, a generative programme synthesis approach \mathcal{G} can effectively increase the mutual information with the true mechanism M by conditioning on the **Universal Distribution** m .*

Proof. The DPI constraint $I(X; Z) \leq I(X; Y)$ for a Markov chain $X \rightarrow Y \rightarrow Z$ assumes that Z is computed solely from Y . In Algorithmic Information Dynamics, the update is $(Q_t, \mathcal{U}) \rightarrow Q_{t+1}$, where \mathcal{U} is an enumeration of Turing Machines ordered by length $2^{-|p|}$. The operator \mathcal{G} selects programmes $p \in \mathcal{U}$ that reproduce the data X within error ϵ . This introduces **exogenous information** (Occam’s bias). If the true mechanism M is simple, the prior $m(p)$ assigns high probability to M . The statistical learner assigns mass based on frequency (correlation); the algorithmic learner assigns mass based on descriptive complexity (causation). When data X is sparse (the collapse regime), the intersection of data constraints and the algorithmic prior $m(p)$ can uniquely identify M , recovering information hidden to statistical projections. \square

We have established that purely statistical updates lead to entropy decay (Theorem 2) and variance drift (Theorem 3). We now detail how the tools defined above counteract these specific failure modes.

Theorem 2 proves that models trained on finite samples \hat{Q}_t lose diversity because the “tails” are statistically invisible. Algorithmic Probability solves this by **generative implication**. Let the observed data be x . An algorithmic learner searches for the minimal program p^* such that $U(p^*) = x$. If x is generated by a lawful mechanism, p^* implicitly defines the *entire* distribution, including the tails missing from the sample. Formally, the algorithmic update expands support via the program’s domain:

$$\text{supp}(Q_{alg}) = \{y \mid \exists p, U(p) = y, K(p) \approx K(p^*)\}. \quad (23)$$

Thus, CTM restores entropy by **re-deriving the underlying law** necessitating the unseen data.

5.5.1 Countering Variance Drift: The Symbolic Anchor

Theorem 3 showed that without external grounding ($\alpha \rightarrow 0$), the model mean follows a random walk. Symbolic constraints act as a **discretisation anchor**.

While a continuous parameter vector can drift infinitesimally, a program cannot; it must change to a distinct valid program. The update rule becomes:

$$p_{t+1} = \arg \min_p [-\log P(D|p) + \lambda K(p)]. \quad (24)$$

The term $K(p)$ acts as a potential barrier. Small statistical noise is insufficient to jump to the next simplest program, locking the model to the simplest explanation p^* .

The algorithmic update ensures that $I(M; Q_{t+1}^{alg}) \geq I(M; Q_t^{stat})$, explicitly escaping the data-processing inequality (as per Proposition 9) by injecting the universal prior m . This allows the system to maintain alignment with the ground truth mechanism M even when $\alpha_t \rightarrow 0$.

5.6 The Algorithmic Solution to Entropy Decay and Drift

We have established that purely statistical updates lead to entropy decay (Theorem 2) and variance drift (Theorem 3) because they rely exclusively on the empirical properties of finite samples. To halt and reverse these dynamics, we must move from statistical inference to **algorithmic inference**. We now detail exactly how the tools defined in the previous section (CTM and BDM) mechanistically counteract the specific failure modes derived in Section 4.

5.6.1 Countering Entropy Decay: Generative Implication

Theorem 2 proves that models trained on finite samples \widehat{Q}_t lose diversity because the “tails” (rare events) are statistically invisible. Standard regularisation (e.g., weight decay) smooths the distribution but does not recover the specific structure of the lost tails.

Algorithmic Probability, approximated by CTM (Eq. 6), solves this by **generative implication**. Let the observed data be a sequence x . A statistical learner sees x as a collection of frequencies. An algorithmic learner searches for the minimal program p^* such that $U(p^*) = x$. Crucially, if the data x is generated by a lawful mechanism (e.g., the sequence 2, 4, 6, 8), the minimal program p^* (“print even numbers”) implicitly defines the *entire* distribution, including the tails that were missing from the sample (e.g., 10, 12, ...).

Formally, while the statistical update shrinks support to the observed samples, $\text{supp}(Q_{stat}) \subseteq \text{supp}(\widehat{Q}_t)$, the algorithmic update expands support via the program’s domain:

$$\text{supp}(Q_{alg}) = \{y \mid \exists p, U(p) = y, K(p) \approx K(p^*)\}. \quad (25)$$

Thus, CTM restores the lost entropy not by adding random noise, but by **re-deriving the underlying law** that necessitates the existence of the unseen data.

5.6.2 Countering Variance Drift: The Symbolic Anchor

Theorem 3 showed that without external grounding ($\alpha \rightarrow 0$), the model mean μ_t follows a random walk $\mu_{t+1} \approx \mu_t + \xi_t$ driven by optimisation noise.

Symbolic constraints (Section 5.4) and algorithmic complexity act as a **discretisation anchor**. The space of low-complexity programs is sparse. While a continuous parameter vector can drift infinitesimally ($\mu_t \rightarrow \mu_t + \epsilon$), a program cannot “drift” slightly; it must change to a distinct valid program. The update rule under algorithmic constraint becomes:

$$p_{t+1} = \arg \min_p [-\log P(D|p) + \lambda K(p)]. \quad (26)$$

The term $K(p)$ acts as a significant potential barrier. Small statistical noise ξ_t is insufficient to jump the gap to the next simplest program. Consequently, the model state is “locked” to the simplest explanation p^* until significantly contradictory evidence accumulates. This quantisation of the hypothesis space prevents the continuous degradation (random walk) characteristic of purely neural updates.

5.6.3 Refining the Contraction

We can now refine the symbolic contraction inequality (Eq. 16). The factor σ is not arbitrary but is derived from the density of computable objects. By the Coding Theorem, the probability mass of the hypothesis space concentrates on simple programs: $m(x) \approx 2^{-K(x)}$. When we project the learned distribution Q_t onto the set of distributions generated by programs of length L , we effectively discard the high-complexity “noise” that constitutes the drift ξ_t .

The algorithmic update ensures that $I(M; Q_{t+1}^{alg}) \geq I(M; Q_t^{stat})$, explicitly escaping the data-processing inequality (as per Proposition 9) by injecting the universal prior m . This allows the system to maintain alignment with the ground truth mechanism M even when the external data fraction $\alpha_t \rightarrow 0$, provided the underlying reality M remains algorithmically simple.

6 Conclusion

We have demonstrated that the “bootstrap” hypothesis of AGI—that an AI system can recursively improve itself indefinitely using its own outputs—is mathematically unsound under current distributional learning paradigms. Specifically, we proved that the self-referential loop constitutes a supermartingale with respect to entropy, leading inevitably to information loss and mode collapse. Furthermore, in the absence of a persistent external grounding signal ($\alpha > 0$), the model’s centroid drifts from the true distribution via a random walk driven by optimisation noise.

Critically, this limitation extends to Reinforcement Learning systems relying on learned or proxy verifiers, which are themselves subject to semantic collapse. The system is information-theoretically closed: it can re-represent existing errors but cannot generate new information about the environment.

The only viable path to escape these attractor dynamics is to move from *analytic* prediction (correlation-based) to *synthetic* computation (mechanism-based). By integrating statistical learning with algorithmic information theory—specifically the Coding Theorem Method (CTM) and Block Decomposition Method (BDM)—we can construct objectives that penalise causal incoherence rather than just distributional divergence. Future work must focus on these neurosymbolic architectures, where “improvement” is defined not by lower perplexity, but by the discovery of more concise, algorithmically probable programs that explain the observed data.

References

- [1] Alemohammad, S., Casco-Rodriguez, J., Luzi, L., Humayun, A. I., Babaei, H., LeJeune, D., Siahkoochi, A., and Baraniuk, R. G. (2023). Self-Consuming Generative Models Go MAD. *arXiv preprint arXiv:2307.01850*.
- [2] Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein Generative Adversarial Networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pp. 214–223. PMLR.
- [3] Calude, C. S. *Information and Randomness: An Algorithmic Perspective*, 2nd ed., Springer, Berlin, Heidelberg, 2002.
- [4] Chaitin, G. J. (1969). On the Length of Programs for Computing Finite Binary Sequences. *Journal of the ACM*, vol. 13, no. 4, pp. 547–569.
- [5] Cover, T. M. and Thomas, J. A. *Elements of Information Theory*, 2nd ed., Wiley-Interscience, Hoboken, NJ, 2006.
- [6] Delahaye, J.-P. and Zenil, H. (2012). Numerical Evaluation of Algorithmic Complexity of Short Strings: A Glance Into the Innermost Structure of Algorithmic Randomness. *Applied Mathematics and Computation*, vol. 219, pp. 63–77.
- [7] Good, I. J. (1965). Speculations Concerning the First Ultraintelligent Machine. In *Advances in Computers*, vol. 6, pp. 31–88. Elsevier.
- [8] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, vol. 27.

- [9] Goertzel, B. and Pennachin, C. (eds.), *Artificial General Intelligence*, Springer, Berlin, Heidelberg, 2007.
- [10] Kant, I. (1781). *Critique of Pure Reason*. Johann Friedrich Hartknoch, Riga.
- [11] Kolmogorov, A. N. (1965). Three Approaches to the Quantitative Definition of Information. *Problems of Information Transmission*, vol. 1, no. 1, pp. 1–7.
- [12] Kurzweil, R. *The Singularity Is Near: When Humans Transcend Biology*, Viking Press, New York, 2005.
- [13] Kurzweil, R. (2005). *The Singularity Is Near: When Humans Transcend Biology*. Viking Press.
- [14] Levin, L. A. (1974). Laws of Information Conservation (Non-growth) and Aspects of the Foundations of Probability Theory. *Problems of Information Transmission*, vol. 10, no. 3, pp. 206–210.
- [15] OpenAI (2023). GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774*.
- [16] Pearl, J. *Causality: Models, Reasoning, and Inference*, 2nd ed., Cambridge University Press, Cambridge, 2009.
- [17] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. (2022). High-Resolution Image Synthesis with Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695.
- [18] Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, vol. 27, pp. 379–423, 623–656.
- [19] Shumailov, I., Shumaylov, Z., Zhao, Y., Gal, Y., Papernot, N., and Anderson, R. (2023). The Curse of Recursion: Training on Generated Data Makes Models Forget. *arXiv preprint arXiv:2305.17493*.
- [20] Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., and Gal, Y. (2024). AI models collapse when trained on recursively generated data. *Nature*, vol. 631, pp. 755–759.
- [21] Solomonoff, R. J. (1964). A Formal Theory of Inductive Inference. Parts I and II. *Information and Control*, vol. 7, pp. 1–22 and 224–254.
- [22] Hernández-Espinosa, A., Ozelim, L., Abrahão, F. S., and Zenil, H. (2025). SuperARC: An Agnostic Test for Narrow, General, and Super Intelligence Based On the Principles of Recursive Compression and Algorithmic Probability. *arXiv preprint arXiv:2307.01850*.

- [23] Vinge, V. (1993). The Coming Technological Singularity: How to Survive in the Post-Human Era. *Whole Earth Review*, no. 81, pp. 88–95.
- [24] Zenil, H., Kiani, N. A., and Tegnér, J. “An Algorithmic Information Calculus for Causal Discovery and Reprogramming Systems,” *iScience*, vol. 23, no. 3, p. 100911, 2020.
- [25] Zenil, H., Hernández-Orozco, S., Kiani, N. A., Soler-Toscano, F., and Rueda-Toicen, A. “A Decomposition Method for Global Evaluation of Shannon Entropy and Local Estimations of Algorithmic Complexity,” *Entropy*, vol. 20, no. 8, p. 605, 2018.
- [26] Zenil, H., Kiani, N. A., and Tegnér, J. *Algorithmic Information Dynamics: A Computational Approach to Causality with Applications to Living Systems*, Cambridge University Press, 2023.
- [27] Zenil, H. Do-Calculus as an Incomplete Theory of Causation: From Correlation-Bound Inference to Algorithmic Mechanistic Causality.