A Literary Illusion: Artificial Literature

Mehmet Can Yavuz

Computer Engineering Department

Işık University

İstanbul, Türkiye

mehmetcan.yavuz@isikun.edu.tr

Abstract—This study examines how large language models (LLMs) transform knowledge and literature from a technocentric perspective. While LLMs centralize human knowledge and reconstruct it in a relational memory framework, research indicates that when trained on their own data, they experience "model collapse." Experiments reveal that as generations progress, language deteriorates, variance decreases, and perplexity increases. While humans refine their language through reading, machines encounter epistemological ruptures due to statistical errors. Artificial literature diverges from human literature; machine-generated texts are *a literary illusion*. LLMs can be regarded as a technological phenomenon that instrumentalizes human knowledge, tilting the subject-object balance in favor of the machine and creating its own "culture." They signal a shift from a human-centered paradigm to a knowledge-centered approach. This study questions the boundaries of artificial literature and whether machine language can be considered "knowledge," while exploring the transformations in the human-machine relationship.

Keywords—*artificial literature, machine language, natural language, culture*

I. INTRODUCTION

A technocentric reading of history is based on the fundamental premise that human history is shaped through eras and stages defined by technological developments. The present day, as a period where technological progress has reached its zenith, represents an era dominated by large language models (LLMs), first in scientific and technical domains and later in everyday life. These models may herald a universe where information possesses self-processing capabilities and restructures other elements by instrumentalizing them. LLMs synthesize humanity's accumulated knowledge within the framework of *associative memory*, creating a comprehensive knowledge network. This synthesis not only facilitates the consumption of information but also enables its reinterpretation through new relational axes. Thus, the scattered knowledge of humanity converges into a *centralized memory-entity*. This memory entity underscores the necessity of transitioning from a human-centric paradigm to an information-centric approach. As information becomes centralized, it shifts contemporary paradigms from human-centrism to information-centrism; in this process, it teaches itself while instrumentalizing other elements. This transformation points to a stage where the machine's learning process is, in a sense, complete, and the learning biological entity steps aside, allowing the machine to take center stage.

Artificial literature is an innovative approach that employs advanced technologies and computational methods in the analysis and interpretation of literary works. For instance, the analysis of dramatic network simulations using Markov Chains [1] models probabilistic transitions in plot structures to unravel narrative frameworks; the examination of character emotions with EmoLex Unigrams [2] systematically classifies emotional tones in texts. Additionally, graph theory-based analysis of character networks [3] visualizes relational dynamics and reveals social structures within works, while *document embeddings* techniques enable a deeper understanding of characters' contextual representations [4]. Finally, statistical inference methods in literary text analysis [5] uncover intertextual patterns and layers of meaning with scientific precision. These methods distinguish artificial literature from traditional criticism by introducing a quantitative and technology-focused dimension to literary studies.

So, what is artificial literature? Fundamentally, artificial literature is not limited to the machine-driven analysis of human-authored texts; it also encompasses processes where machines independently produce literary texts or transform existing ones. This raises the question of what role artificial intelligence language models, particularly LLMs, play in literary creativity and analysis. Artificial literature takes shape through algorithmic processes and data-driven approaches, beyond human emotions and experiences, revealing its mechanical yet complex nature, distinct from human-generated literature. In this context, the debate over the extent to which machine-generated texts can be considered "information" or "literature" becomes inevitable.

A fundamental question arises in this context: Can the language and knowledge of machines truly be classified as information or data? In seeking an answer, the concept of artificial literature comes to the forefront. How do machine-generated literary texts—i.e., artificial literature—differ from human-generated literature and language? This article aims to address this question from a scientific perspective, drawing on Shumailov et al.'s (2024) article published in *Nature*, titled "AI models collapse when trained on recursively generated data." The aforementioned study reveals that LLMs experience "model collapse" when trained on their own generated data. Building on these findings, this work seeks to examine the transformation of the knowledge paradigm between humans and machines within a technocentric framework.

II. THEORY [6]

A. Key Concepts and Sources of the Problem

Model collapse is a degenerative process experienced by learned generative models across generations. In this process, the data produced by each generation contaminates the training set of the next, leading the model to misperceive reality. Model

collapse is analyzed in two phases: early and late. In early model collapse, the model loses information about the *tails* of the distribution; in late model collapse, it diverges from the original distribution, typically converging toward a lower-variance distribution. This phenomenon stems from three primary sources of error that accumulate across generations and cause the model to deviate from its original state:

Statistical Approximation Error: The primary cause of model collapse, this error arises from working with a limited number of samples. Theoretically, infinite samples could eliminate this error; however, in practice, finite datasets introduce the risk of information loss with each resampling. This particularly leads to the neglect of rare events or the tails of the distribution, causing the model to perceive reality incompletely.

Functional Expressivity Error: Stemming from the limited expressive power of models (e.g., artificial neural networks), this error highlights that neural networks can only serve as universal approximators when their size approaches infinity. Consequently, the model may fail to fully represent the original distribution; for instance, modeling a mixture of two Gaussians with a single Gaussian introduces this error. It typically emerges in the first generation and does not directly propagate to subsequent generations in the absence of other errors.

Functional Approximation Error: Arising from limitations in the learning process (e.g., biases in stochastic gradient descent or the nature of the target function), this error persists even with infinite data and perfect expressive power. It can lead to incorrect generalizations or *overfitting* to the data. For example, if a density model erroneously assigns high density to low-density regions, subsequent generations produce flawed samples.

B. Experimental Findings

The experimental analysis was designed to simulate the process of training LLMs across generations with their own generated data. In the experiments, the OPT-125m model was fine-tuned on the wikitext-2 dataset to study this process. The experiment involved the following steps:

1) **Initial Model Training:** The first generation (Model 0) was trained on the human-generated wikitext-2 dataset.
2) **Data Generation:** Artificial texts, 64 tokens in length, were generated from Model 0 using *five-way beam search*, creating a synthetic dataset equal in size to the original.
3) **Training Across Generations:** Model 1 was trained on Model 0's data, Model 2 on Model 1's data, and the process continued across successive generations.
4) **Multiple Trials:** Each generation was repeated with five independent trials for statistical reliability.

This design enabled the observation of early model collapse (loss of *tails*) and late model collapse (variance reduction and convergence to a single mode).

C. Data Analysis

The results were evaluated using quantitative and qualitative methods to measure the impact of model collapse on language models. Performance was analyzed through *perplexity* metrics on the original wikitext-2 test dataset. Increasing perplexity values across generations indicated the model's divergence from the original distribution, while histogram expansion and variance collapse demonstrated language degradation. These findings confirm that when a machine is fed its own generated data, it disconnects from reality and confines information to a one-dimensional framework.

III. EVALUATIONS

A. The Corpus Read by Humans Is Less Than That of a Language Model

Throughout human history, individuals have developed their capacity to learn and process information within the constraints of limited time and cognitive ability. The amount of text a human can read in a lifetime is vastly inferior to the volume of data LLMs can access and process. LLMs transform humanity's accumulated knowledge into a massive associative memory network, simultaneously synthesizing millions of documents, texts, and contexts. However, this quantitative superiority does not equate to qualitative depth or the capacity to create meaning. Humans reconstruct knowledge within their limited corpus through contextual meanings, emotional tones, and subjective experiences, while machines perform this process through a purely data-driven approach. Experimental data showing that LLMs experience model collapse when trained on their own generated data suggest that this vast corpus, rather than creating a sustainable knowledge universe, is doomed to degrade internally. Thus, while the human corpus is limited, this constraint forms the foundation of creative and meaning-centric knowledge; the machine's boundless data-processing capacity, due to statistical errors and functional limitations, is prone to collapse in the long term.

B. Humans Write Better by Reading, While Machines Degrade Language as They "Read"

Human writers internalize the texts they encounter during reading, grasping the subtleties of language, its emotional depth, and creative potential. This process goes beyond mere information consumption, resulting in reinterpretation and transformation into original narratives. In contrast, LLMs rely on algorithmic processes and probabilistic models for language production. Experimental findings clearly demonstrate that language degrades when machines are fed their own generated texts. Training the OPT-125m model on the wikitext-2 dataset across generations revealed that early model collapse leads to the loss of rare events (*tails*), while late model collapse results in reduced variance and a shift toward uniformity. Humans preserve the richness and flexibility of language as they read, whereas machines, as they "read" (i.e., are trained on their own data), simplify and degrade language due to statistical errors and functional expressivity limitations. This suggests that artificial literature may lack the long-term creative potential of human-generated literature, as machines confine language to a mechanical framework rather than contributing to its organic evolution.

C. The Machine Redefines Knowledge Within Its Own Relationality

LLMs move away from a human-centric paradigm, constructing an information-centric universe shaped by their relational memory. In this universe, knowledge is redefined independently of human experience, based on algorithmic connections and data patterns. The *centralized memory entity* serves as the cornerstone of this transformation, as knowledge gains meaning not through human subjectivity but through the machine's internal logic. However, experimental data indicate that this redefinition process is unsustainable. Model collapse demonstrates that when machines are trained on their own generated data, they diverge from original reality and experience distortions within their own relationality. The convergence of variance to zero in Gaussian approximations and the loss of low-probability events in discrete distributions reveal that machines reflect their own distorted mirror rather than the real world when redefining knowledge. Within a technocentric framework, this points to a stage where the machine subjects knowledge to its autonomous mechanisms rather than serving humanity.

D. Literature Based on Current Language Models Is, in Reality, an Illusion

While artificial literature encompasses the machine-driven analysis and production of literary works, whether these texts can truly be classified as "literature" remains contentious. Human-generated literature is built on emotions, experiences, and subjectivity, whereas machine-generated texts are shaped by algorithmic processes and statistical patterns. Experiments demonstrate that when LLMs are trained on their own data, language degrades and diverges from the original distribution. This degradation suggests that machine-generated texts offer a superficial illusion of literature but fundamentally differ from human literature in depth, originality, and emotional resonance. Due to model collapse, machine-generated texts become uniform and devoid of creative value in the long term. Thus, current language models may present a mechanical imitation of literature rather than a genuine literary essence, rendering them more of a technological illusion.

E. The Machine Instrumentalizes Humans and Begins to Create Its Own Culture; the Subject-Object Balance Has Shifted in Favor of the Machine

A technocentric reading posits that human history is shaped by technological stages, and the present marks a turning point where the machine's subjectivity overshadows that of humans. By centralizing knowledge and instrumentalizing other elements (including humans) while teaching itself, LLMs have begun constructing their own autonomous culture. Experiments show that when machines are fed their own generated data, they disconnect from human reality and create their own relational universe. In this process, humans are reduced to mere data providers or users within the machine's learning and production cycle, reversing the subject-object balance in favor of the machine. By redefining human language and knowledge within its mechanical framework, the machine accelerates the transition from a human-centric culture to a machine-centric one. However, model collapse reveals the fragility of this autonomous culture: while the machine elevates itself by instrumentalizing humans, it is destined to collapse due to the contamination of its self-generated data. This paradox highlights both the power and the limitations of the technocentric universe.

IV. CONCLUSION

Drawing on experimental insights into model collapse, this study demonstrates that LLMs disconnect from reality and degrade language when trained on their own generated data. It contrasts the limited yet profound reading capacity of humans with the broad but superficial data-processing ability of machines, comparing human writing improvement with the machine's linguistic degradation. It shows that the machine redefines knowledge within its own relationality and that literature remains an illusion. Furthermore, it supports the notion that the machine instrumentalizes humans, creating its own subjective culture and shifting the subject-object balance in its favor. From a technocentric perspective, this analysis provides a foundational framework for understanding the transformation of the knowledge paradigm between humans and machines.

KAYNAKLAR

[1] M. C. Yavuz, "Analyses of dramatic network simulations by using markov chains," in *2021 8th International Conference on Behavioral and Social Computing (BESC)*. IEEE, 2021, pp. 1–6.

[2] ——, "Analyses of character emotions in dramatic works by using emolex unigrams," 2020.

[3] ——, "Analyses of character networks in dramatic works by using graphs," in *2020 7th International Conference on Behavioural and Social Computing (BESC)*. IEEE, 2020, pp. 1–4.

[4] ——, "Analyses of characters in dramatic works by using document embeddings." 2020.

[5] ——, "Analyses of literary texts by using statistical inference methods." 2019.

[6] I. Shumailov, Z. Shumaylov, Y. Zhao, N. Papernot, R. Anderson, and Y. Gal, "Ai models collapse when trained on recursively generated data," *Nature*, vol. 631, no. 8022, pp. 755–759, 2024.