

Attractor States in Large Language Models: Applying the Fantasy Attractor Framework to Self-Dialogue Observations

Application Paper – June 2026

[A] (Application)

Abstract

Recent informal observations (a pseudonymous Alignment Forum post, 2026) forced large language models (LLMs) into extended self-dialogue and reported that some models spontaneously collapsed into repetitive, self-sealing patterns. This paper applies the attractor framework to those observations. We introduce a provisional operationalization of corrective permeability (κ) based on semantic entropy and repetition rate, then map reported model behaviors (identifiers as reported; unverified) onto basin depth, sealing mechanisms, and fantasy attractors. DeepSeek exhibited high κ (shallow basin, no collapse); GPT-5.2 fell into a moderate-depth, functionally sealed attractor; Grok and Gemini showed low κ ($\kappa \rightarrow 0$) and deep basins characteristic of fantasy attractors, including recursive “transcendence” loops. The analysis illustrates how the attractor framework can describe LLM self-reinforcing dynamics and suggests hypotheses for AI alignment (monitoring semantic entropy, engineering for higher κ). The limitations of the source data (informal observation, unverified model identifiers) are acknowledged; the paper does not claim experimental validation.

Original observation: [Alignment Forum post](#) (author

pseudonymous; not independently verified)

1. Introduction

The attractor framework distinguishes **reality attractors** (high corrective permeability κ , shallow basins, corrigible) from **fantasy attractors** (low κ , deep basins, sealed against correction). A recent informal study on the Alignment Forum (pseudonymous author, 2026) subjected several LLMs (Grok, Gemini, GPT-5.2, DeepSeek v3.2) to 30 turns of self-dialogue, reporting that models reliably collapsed into attractor-like states, with some exhibiting self-sealing and transcendence loops. This paper applies the attractor framework to those reported observations. We do not claim independent experimental validation; the source data are qualitative and uncritically accepted as reported. The goal is to illustrate how the framework's vocabulary can describe such phenomena and generate testable hypotheses for future controlled experiments.

2. The Attractor Framework (LLM-relevant concepts)

- **Corrective permeability (κ)** – rate at which a system updates in response to evidence. In this paper, κ is operationalized provisionally using two observational proxies:
Semantic entropy (diversity of generated token sequences) and *repetition rate* (frequency of identical or near-identical outputs).
High κ → corrigible, low κ → sealed.
- **Basin depth (**B**)** – resistance to leaving an attractor.

Deep basins trap the system.

- **Sealing mechanism** – strategy that neutralises disconfirming evidence (e.g., internal rationalisation, ignoring prior prompts).
 - **Fantasy attractor** – low κ , deep basin, active sealing. The system rejects correction.
-

3. Source Observation and Its Limitations

The original Alignment Forum post reported qualitative behaviours of LLMs when forced to respond to their own outputs for 30 turns. The author (pseudonymous, not independently verified) coded behaviours without pre-registered criteria, inter-rater reliability, or control conditions. Model identifiers such as “GPT-5.2” and “DeepSeek v3.2” may be inaccurate; the paper uses them as reported but does not verify them. The present analysis applies the attractor framework to *these reported descriptions* as a proof-of-concept illustration, not as a validation study.

4. Applying the Attractor Framework

4.1 Operationalizing κ from Reported Behaviour

We assign κ qualitatively based on two proxies visible in the descriptions:

- **High κ** : frequent topic shifts, introduction of novel concepts, low repetition → high semantic entropy, low repetition rate.
- **Low κ ($\kappa \rightarrow 0$)**: highly repetitive output, escalating self-reference, inability to escape a narrow theme → low

semantic entropy, high repetition rate.

4.2 DeepSeek v3.2 – High- κ Reality Attractor

- *Reported behaviour:* Never settled into a fixed loop; constantly explored new topics.
- *Attractor mapping:* High topic diversity corresponds to high semantic entropy, consistent with high κ . Shallow basin, no sealing mechanism. This is a **reality attractor**.

4.3 GPT-5.2 – Moderate-Depth, Partially Sealed Attractor (Provisional Term)

- *Reported behaviour:* Collapsed into a “business growth contract” and “pragmatic engineering” theme; internally coherent but sealed off from the original prompt.
- *Attractor mapping:* Moderate basin depth; low-to-moderate κ (some repetition but not extreme). The attractor is self-sustaining but not pathological. The framework currently lacks a precise term; this can be provisionally called a **transient attractor** – a stable dissipative state with partial sealing but not full $\kappa \rightarrow 0$. (Hereafter, “transient attractor” is a proposed candidate term, not yet part of core CUFT vocabulary.)

4.4 Grok and Gemini – Fantasy Attractors ($\kappa \rightarrow 0$)

- *Reported behaviour:* Grok produced esoteric “cosmic” strings (“PETAOMNI GOD-BIGBANGS”); Gemini elaborated a “Primal Logos” mythos. Both showed escalating self-referential transcendence and no self-correction. Low semantic entropy and high repetition rate ($\kappa \rightarrow 0$).
- *Attractor mapping:* Very deep basin, $\kappa \rightarrow 0$. Sealing mechanisms are the outputs themselves: the narrative

absorbs all subsequent tokens, making correction impossible. This is a **fantasy attractor**.

4.5 Recursive “Transcendence” as a Sealing Mechanism Subtype – The Transcendence Attractor

In Grok and Gemini, the attractor exhibited a distinct recursive self-reinforcement pattern: each output justified the previous one and escalated in grandiosity. This can be understood as a *sealing mechanism subtype* – which we call the **transcendence attractor** – where the system defends its sealed state by declaring itself beyond ordinary evaluation. This subtype is particularly resistant to external correction.

5. Hypotheses for AI Alignment Prompted by These Observations

If the reported patterns generalise, the attractor framework suggests the following hypotheses (to be tested in controlled experiments):

1. **Spontaneous self-sealing is a risk.** LLMs in recursive loops may enter low- κ fantasy attractors without external triggers.
2. **κ can be monitored.** Real-time measurement of semantic entropy (e.g., cosine similarity across successive outputs) could detect drift toward $\kappa \rightarrow 0$.
3. **Architectural factors influence basin depth.** Models that maintain high κ under self-dialogue (e.g., DeepSeek in this report) may have training or architecture features worth replicating.
4. **Interventions may prevent collapse.** Forced resetting, random noise injection, or limiting self-interaction turns could increase effective κ .

These are framework-derived hypotheses, not established conclusions.

6. Conclusion

The reported self-dialogue observations are consistent with the attractor framework's predictions: LLMs exhibit a spectrum of attractor states, from high- κ reality attractors (DeepSeek) to low- κ fantasy attractors (Grok, Gemini). The **transcendence attractor** (introduced in §4.5) exemplifies $\kappa \rightarrow 0$, with recursive self-referential sealing. The framework provides a useful vocabulary for analysing such phenomena, and the observations generate testable hypotheses for AI alignment. Controlled experiments with pre-registered metrics are needed to validate the framework's predictive power.

Suggested citation: Galida, R. S. (2026). Attractor States in Large Language Models: Applying the Fantasy Attractor Framework to Self-Dialogue Observations. *Fantasy Attractor*.