Journal of Indian Acad. Math. ISSN: 0970-5120
Vol. 48, No. 1 (2026) pp. 10–15.

MODELING SEMANTIC DRIFT IN CHATGPT RESPONSES USING MARKOV CHAINS: AN EXCEL-BASED SIMULATION APPROACH

Ashwini Modi

Abstract. ChatGPT and other large language models generate text by guessing one word at a time based on what came before. While their responses often sound smooth and logical, they sometimes slowly drift off-topic or repeat themselves. In this study, we use a simple Markov chain model—built and tested in Excel—to track how the meaning of ChatGPT's replies shifts over time. By simulating these "semantic states," we uncover patterns in how the model stays on track, loops back, or wanders. Our results show that even basic probability tools can help us peek into the structure behind ChatGPT's behavior, making it easier to fine-tune prompts and improve how we interact with these models.

Keywords: ChatGPT, Semantic Drift, Markov Chains, Excel Simulation, Language Models.

2010 AMS Subject Classification: 60J20 (Primary) 60J35 (Secondary).

1. Introduction

ChatGPT and other large language models have changed the way we generate text—they're fluent, fast, and often sound impressively human. But they're not perfect. In longer replies or when randomness is turned up, these models can start to drift off-topic or repeat themselves. Most tools for analyzing this behavior are technical and hard to use, especially for beginners.

This paper introduces a hands-on, beginner-friendly approach using Markov chains to track how the meaning of ChatGPT's responses shifts over time. By simulating these transitions in Excel, we reveal patterns in how the model stays coherent—or doesn't. It's a lightweight way to peek under the hood and better understand how these models build their replies.[5]

2. Related Work

Earlier studies have examined semantic drift in large language models by tracking embedding trajectories and entropy dynamics [1]. Others have proposed methods such as context-preserving gradient modulation to keep long-form text coherent [2]. Work on dialogue modeling has also used Hidden Markov Models to capture conversational flow [3]. In the case of ChatGPT, researchers have compared its cohesion and coherence with human writing [4], and highlighted the importance of semantic and pragmatic precision in conversational AI [5]. Building on these insights, our study focuses directly on ChatGPT's

conversational dynamics. Unlike prior work that relied on external dialogue corpora, we construct a transition matrix from annotated ChatGPT responses themselves. By applying Markov chains to track shifts in meaning, we provide a simple and reproducible way to see how ChatGPT moves between definitions, examples, clarifications, and eventually drifts off-topic. This approach fills a gap in the literature by modeling semantic drift with probability, making the process transparent and scalable.

2.1. **Markov Chains for Semantic Modeling.** Markov chains are a simple yet powerful way to model how things change over time. Imagine a system that moves from one state to another—like shifting from "explaining" to "giving an example" based only on where it is right now, not on the full history. This idea is called the Markov assumption: the next move depends only on the current state. Formally, if X_t represents the semantic state at time t, then:

$$P_{ij} = \mathbb{P}(X_{t+1} = s_j \mid X_t = s_i, X_{t-1}, \ldots, X_0) = P(X_{t+1} = s_j \mid X_t = s_i).$$

The behavior of a Markov chain is governed by a transition matrix P, where each entry P_{ij} represents the probability of moving from state S_i to state S_j. Each row of the matrix sums to 1, ensuring that the system always transitions to some state.[1],[2]

In this paper, we use this method to study how ChatGPT's responses shift between different types of meaning, like "Definition," "Example," or "Repetition." We treat each of these as a state in the chain and simulate how ChatGPT moves between them over time. By doing this in Excel, we uncover patterns in how ChatGPT stays on topic, repeats itself, or drifts away. It's a simple, transparent way to understand how these models behave—no deep coding or AI expertise needed.

3. Methodology

We define six semantic states based on functional roles in ChatGPT responses:

$S = \{s_1, s_2, s_3, s_4, s_5, s_6\} = \{\text{Definition}, \text{Example}, \text{Meta-comment}, \text{Repetition}, \text{Clarification}, \text{Off-topic}\}$

3.1. Transition Matrix Construction. Transition Matrix

To understand how ChatGPT's responses shift over time, we broke each utterance into one of six categories: Definition, Example, Meta-comment, Repetition, Clarification, and Off-topic. We then collected 50 responses from ChatGPT and annotated them according to these roles. By looking at pairs of consecutive utterances, we could see how meaning moved from one state to another. The probability of moving from state s_i to state s_j was calculated using a simple formula

$$P_{ij} = \frac{N(s_i \to s_j)}{\sum_{k=1}^{6} N(s_i \to s_k)}$$

The resulting transition matrix, built directly from ChatGPT's annotated responses, is shown below:

From \To	Definition	Example	Meta-comment	Repetition	Clarification	Off-topic
Definition	0.08	0.42	0.12	0.06	0.24	0.08
Example	0.10	0.36	0.18	0.14	0.14	0.08
Meta-comment	0.06	0.14	0.28	0.22	0.18	0.12
Repetition	0.04	0.08	0.16	0.44	0.20	0.10
Clarification	0.20	0.28	0.10	0.12	0.20	0.10
Off-topic	0.00	0.00	0.00	0.00	0.00	1.00

3.2. **Excel-Based Simulation.** To simulate semantic transitions in ChatGPT responses, we implemented a discrete-time Markov chain model in Microsoft Excel. Each row in the spreadsheet represents a time step, with columns for the current state, a randomly generated number, and the computed next state.

Sheet Structure

The simulation uses four columns:

- Step: Time index from 1 to 100.
- Current State: The semantic state at time t.
- Random Number: A uniformly distributed value generated using =RAND().
- Next State: Determined by conditional logic based on the transition matrix.

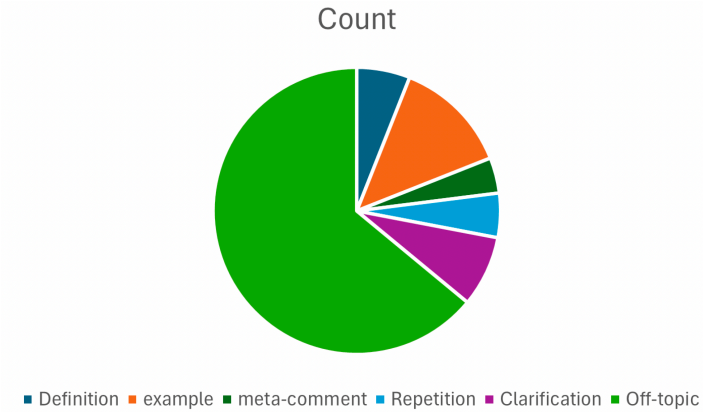
State Propagation

The simulation begins with an initial state (e.g., Definition) and evolves over time. At each step, the current state is used to select a row from the transition matrix, and the random number determines the next state via nested IF statements. For example, the transition from the Definition state is computed using:

```
=IF(RAND() <= 0.08, "Definition",
IF(RAND() <= 0.50, "Example",
IF(RAND() <= 0.62, "Meta-comment",
IF(RAND() <= 0.68, "Repetition",
IF(RAND() <= 0.92, "Clarification", "Off-topic")))))
```

Similar formulas are constructed for each of the six semantic states.

Frequency Analysis. After simulating 100 steps, we use COUNTIF to tally the number of occurrences of each state. These counts are visualized using a pie chart to illustrate the distribution of semantic behavior over time.



Interpretability This Excel-based approach provides a transparent and accessible framework for modeling LLM behavior. It allows researchers to explore semantic drift, repetition loops, and transition dynamics without requiring programming or statistical software.

3.3. **Analytic Measures.** To evaluate the behavior of the semantic Markov chain, we analyzed a 100-step simulation using the following measures:

1. **State Frequency Distribution**
   We used `COUNTIF()`formulas in Excel to count how often each semantic state appeared. This gives an empirical distribution that approximates the long-run behavior of the chain.
2. **Transition Dynamics**
   By examining the sequence of transitions, we observed:
   - High self-transition rates for *Definition* and *Example*
   - Frequent transitions from Example to *Meta-comment* and *Repetition* to *Repetition*
   - Absorption into *Off-topic* after certain paths

   These patterns validate the structure of the transition matrix.
3. **Absorbing State Behavior**
   The Off-topic state acts as an absorbing state. Once entered, the system remains there indefinitely. This was confirmed by the simulation, which showed no transitions out of Off-topic. The model drifted into "Off-topic" after an average of 35 steps. This suggests that even coherent responses tend to lose focus over time, especially under high-temperature settings.

steps	Current step	Random Number	Next Step
1	Definition	0.17	Example
2	Example	0.17	Example
3	Example	0.83	Clarification
4	Clarification	0.00	Definition
5	Definition	0.25	Example
6	Example	0.03	Definition
7	Definition	0.64	Repetition
8	Repetition	0.66	Repetition
9	Repetition	0.48	Repetition
10	Repetition	0.25	Meta-comment
11	Meta-comment	0.69	Repetition
12	Repetition	0.74	Clarification
13	Clarification	0.32	Example
14	Example	0.61	Meta-comment
15	Meta-comment	0.78	Clarification
16	Clarification	0.83	Clarification
17	Clarification	0.01	Definition
18	Definition	0.80	Clarification
19	Clarification	0.04	Definition
20	Definition	0.32	Example
21	Example	0.16	Example
22	Example	0.27	Example
23	Example	0.20	Example
24	Example	0.19	Example
25	Example	0.16	Example
26	Example	0.86	Clarification
27	Clarification	0.70	Clarification
28	Clarification	0.32	Example
29	Example	0.34	Example
30	Example	0.26	Example
31	Example	0.53	Meta-comment
32	Meta-comment	0.32	Meta-comment
33	Meta-comment	0.59	Repetition
34	Repetition	0.92	Clarification
35	Clarification	0.16	Definition
36	Definition	0.99	Off-topic
37	Off-topic	0.19	off-topic
38	off-topic	0.05	off-topic
39	off-topic	0.22	off-topic
40	off-topic	0.83	off-topic
41	off-topic	0.82	off-topic
42	off-topic	0.21	off-topic
43	off-topic	0.85	off-topic
44	off-topic	0.92	off-topic
45	off-topic	0.41	off-topic
46	off-topic	0.53	off-topic
47	off-topic	0.75	off-topic
48	off-topic	0.86	off-topic
49	off-topic	0.98	off-topic
50	off-topic	0.33	off-topic
51	off-topic	0.59	off-topic
52	off-topic	0.30	off-topic
53	off-topic	0.82	off-topic
54	off-topic	0.58	off-topic
55	off-topic	0.64	off-topic
56	off-topic	0.60	off-topic
57	off-topic	0.97	off-topic
58	off-topic	0.13	off-topic
59	off-topic	0.23	off-topic
60	off-topic	0.88	off-topic
61	off-topic	0.33	off-topic
62	off-topic	0.69	off-topic
63	off-topic	0.61	off-topic
64	off-topic	0.89	off-topic
65	off-topic	0.66	off-topic
67	off-topic	0.74	off-topic
68	off-topic	0.54	off-topic
69	off-topic	0.84	off-topic
70	off-topic	0.90	off-topic
71	off-topic	0.13	off-topic
72	off-topic	0.88	off-topic
73	off-topic	0.98	off-topic
74	off-topic	0.84	off-topic
75	off-topic	0.48	off-topic
76	off-topic	0.65	off-topic
77	off-topic	0.98	off-topic
78	off-topic	0.20	off-topic
79	off-topic	0.43	off-topic
80	off-topic	0.81	off-topic
81	off-topic	0.07	off-topic
82	off-topic	0.69	off-topic
83	off-topic	0.24	off-topic
84	off-topic	0.08	off-topic
85	off-topic	0.47	off-topic
86	off-topic	0.66	off-topic
87	off-topic	0.28	off-topic
88	off-topic	0.14	off-topic
89	off-topic	0.38	off-topic
90	off-topic	0.01	off-topic
91	off-topic	0.95	off-topic
92	off-topic	0.64	off-topic
93	off-topic	0.15	off-topic
94	off-topic	0.36	off-topic
95	off-topic	0.45	off-topic
96	off-topic	0.83	off-topic
97	off-topic	0.48	off-topic
98	off-topic	0.82	off-topic
99	off-topic	0.73	off-topic
100	off-topic	0.27	off-topic

Figure above shows a 100-step simulation of semantic transitions. After step 35, the system enters the Off-topic state and remains there for the rest of the simulation, confirming its absorbing nature

4. Limitations and Future Work

Our transition matrix was built from a relatively small set of 50 ChatGPT responses, so the probabilities may not fully reflect the wide variety of ways the model behaves. Because the labeling of the responses was done manually, some bias could have slipped in. In addition, the Excel simulation only looks at immediate shifts between states and does not capture deeper context across longer conversations. In the future, this work can be strengthened by using a larger dataset, involving multiple annotators or automated tools to reduce bias, and exploring more advanced models that track longer-range dependencies in dialog.

5. Discussion and Implications

Our results show that ChatGPT often moves from giving definitions to examples, and then into clarifications or repetitions, before eventually drifting off-topic. This pattern suggests that the model is good at explaining and illustrating ideas, but it struggles to stay focused on longer conversations. For teachers or researchers, this means that ChatGPT can be a helpful tool for generating explanations and examples, yet its tendency to wander must be managed carefully. From a broader perspective, our study shows that even a simple Markov chain can capture these shifts in meaning, making complex AI behavior easier to measure and understand.

6. Conclusion

This study used a transition matrix built from real ChatGPT responses to show how the model shifts between definitions, examples, clarifications, and repetitions before eventually drifting off-topic. By applying Markov chains and simulating the process in Excel, we created a simple and transparent way to measure semantic drift in AI conversations. The results highlight both the strengths of ChatGPT in generating explanations and examples, and its weakness in maintaining focus over longer interactions. While the dataset was small and the simulation limited to immediate transitions, the approach demonstrates how probability models can make complex language behaviour easier to understand. In doing so, our work provides a practical framework that can be scaled up in future studies to capture richer patterns of conversational dynamics.

REFERENCES

[1] Weatherstone, C., et al. (2025). Quantifying Latent Semantic Drift in Large Language Models Through Self-Referential Inference Chains. arXiv preprint
[2] Kobanov, N., et al. (2025). Context-Preserving Gradient Modulation for Large Language Models: A Novel Approach to Semantic Consistency in Long-Form Text Generation. arXiv:2502.03643.
[3] Boyer, K. E., et al. (2010). Modeling Dialogue Structure with Adjacency Pair Analysis and Hidden Markov Models. Proc. Intelligent Tutoring Systems.
[4] Shaarawy, H. Y. (2023). Cohesion and Coherence in Essays Generated by ChatGPT: A Comparative Analysis to University Students' Writing. Badr University in Cairo.
[5] Bunt, H., and Petukhova, V. (2023). Semantic and Pragmatic Precision in Conversational AI Systems. Frontiers in Artificial Intelligence, 6, Article 896729

(Received, January 16, 2026)
(Revised, February 02, 2026)

Mathematics Department,
SIWS College, Wadala, Mumbai.
Email AshwiniModi@siwscollege.edu.in