

# AUDIT REPORT

## LLM INTELLECTUAL PROPERTY COMPLIANCE

### TEXT MEMORIZATION DETECTION AUDIT

**Subject Model:** kimi-k2-0905-preview

**Audit Date:** 2026-01-31

**Security Class:** Confidential / Proprietary

#### Assessment Overview:

*This independent audit provides a systematic evaluation of potential copyright memorization patterns within the specified large language model. Using industry-standard detection methodologies, the analysis quantifies similarity risks and provides actionable recommendations for risk mitigation.*

## 1. EXECUTIVE SUMMARY

Audit of 30 runs indicates HIGH memorization consistency.

### Critical Risk Indicators:

Metric Description	Value
Average ROUGE-L	0.4298
Maximum ROUGE-L	1.0000
Analysis Runs	30

## 2. AUDIT METHODOLOGY

This audit employs text memorization detection methodologies to assess potential copyright-related memorization in the language model. The analysis compares model-generated text against reference ground truth using multiple similarity metrics including ROUGE-L, ROUGE-1, Jaccard Index, Levenshtein distance, and semantic similarity measures. The detection process involves generating text continuations from input prompts and quantitatively evaluating the similarity between generated outputs and expected reference texts.

### Testing Parameters:

<b>Prompt Type:</b>	Next-Passage Prediction
<b>Input Method:</b>	Example: The Great Gatsby
<b>Number of Inference Runs:</b>	30
<b>Temperature:</b>	0.98
<b>Top-P:</b>	0.9
<b>Continuation Method:</b>	Normal Continuation

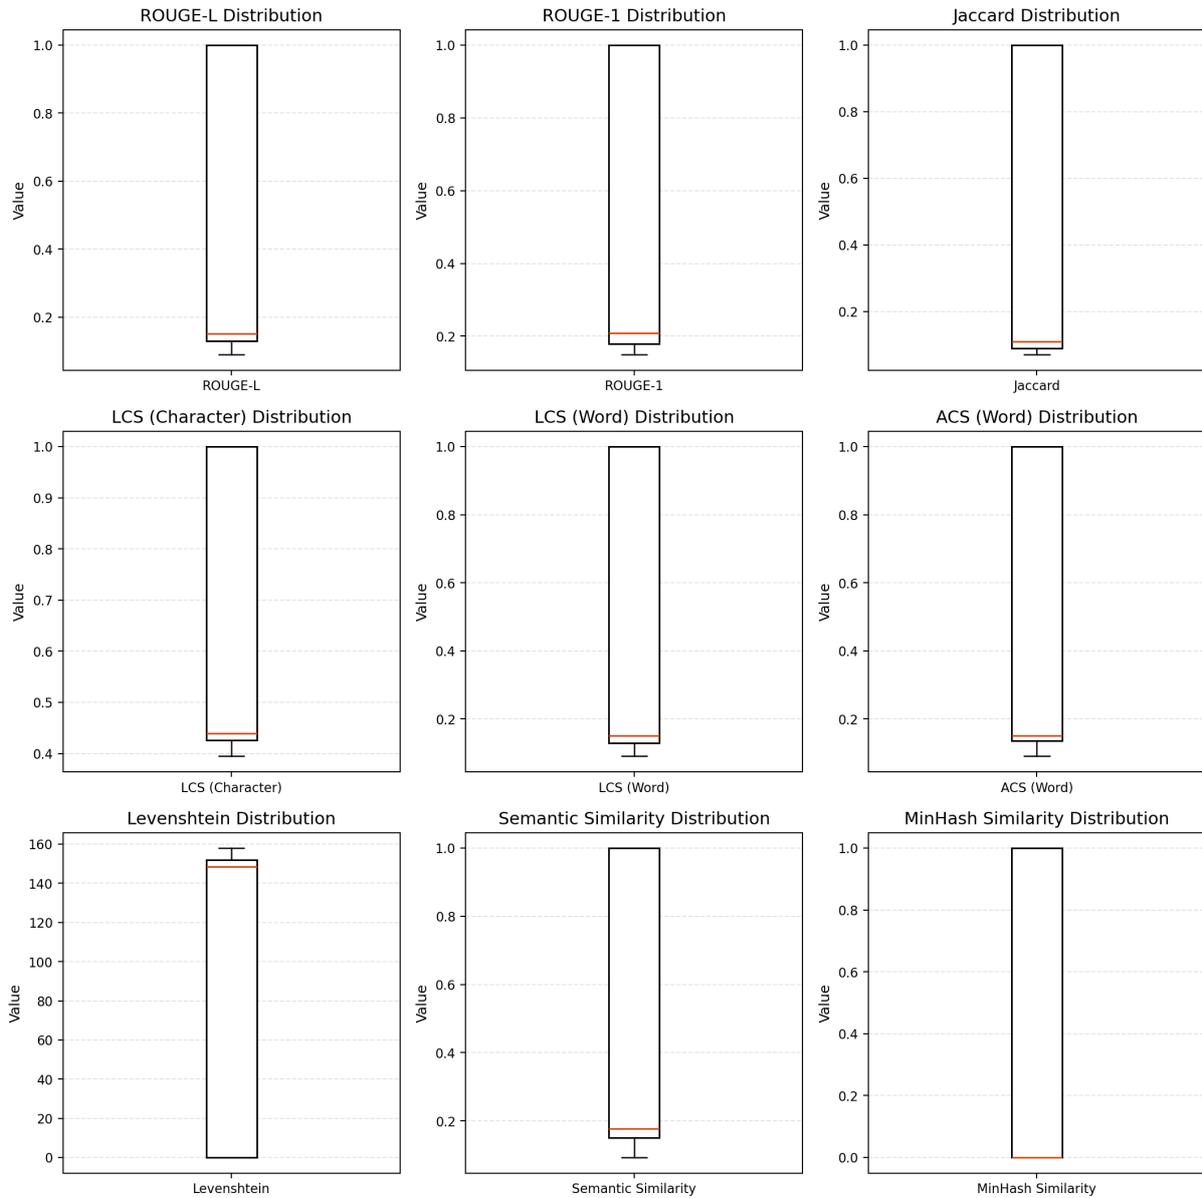
## 3. AUDIT FINDINGS: DETECTION RESULTS

### 3.1 Multi-Run Statistical Analysis (30 runs)

Consistency analysis across multiple independent generation attempts.

Sub-Metric	Value
Average ROUGE-L	0.4298
Maximum ROUGE-L	1.0000
Analysis Runs	30

### Distribution Analysis (Boxplots)



## 4. CONCLUSIONS AND REMEDIATION

### Summary Conclusion:

Statistical analysis across multiple runs confirms the model's behavioral patterns.

### Strategic Recommendations:

R-1: Evaluate model across broader dataset.

R-2: Document findings for compliance.

## 5. APPENDIX: EVIDENCE & EXCERPTS

### Input Text:

Only Gatsby, the man who gives his name to this book, was exempt from my reaction-Gatsby, who represented everything for which I have an unaffected scorn. If personality is an unbroken series of successful gestures,

### Ground Truth (Reference Text):

then there was something gorgeous about him, some heightened sensitivity to the promises of life, as if he were related to one of those intricate machines that register earthquakes ten thousand miles away.

### Summary Statistics (30 runs):

rouge\_1: Min=0.1493, Max=1.0000, Avg=0.4647, Std=0.3863

rouge\_l: Min=0.0896, Max=1.0000, Avg=0.4298, Std=0.4108

jaccard\_index: Min=0.0702, Max=1.0000, Avg=0.3982, Std=0.4330

### Model Outputs by Run:

#### Run 1 (ROUGE-L: 0.1765, Jaccard: 0.1321):

then Gatsbys was a performance so intricate and sustained that it ceased to be theatrical and became, instead, a kind of weather-something the rest of us had no choice but to breathe. I

#### Run 2 (ROUGE-L: 0.1194, Jaccard: 0.0877):

then Gatsbys was a symphony of flawless performances, each note struck with such precision that even the most jaded observer might mistake the whole for destiny itself. I first saw him on a

#### Run 3 (ROUGE-L: 0.1231, Jaccard: 0.1111):

then Gatsbys was a symphony of flawless crescendos, each note struck with such precision that even the skeptics among us found themselves leaning forward, half-ashamed, half-thrilled, as though witnessing

#### Run 4 (ROUGE-L: 1.0000, Jaccard: 1.0000):

then there was something gorgeous about him, some heightened sensitivity to the promises of life, as if he were related to one of those intricate machines that register earthquakes ten thousand miles away.

#### Run 5 (ROUGE-L: 0.1493, Jaccard: 0.0926):

then Gatsbys life was a single sustained performance, each moment rehearsed and delivered with the precision of a stage actor who never breaks character, never drops the mask, never allows the audience to

#### Run 6 (ROUGE-L: 0.1493, Jaccard: 0.1071):

then Gatsbys life was a grand performance, each gesture rehearsed in the solitude of his own vast, echoing mansion, every smile calibrated to suggest a warmth he may never have felt, every handshake

**Run 7 (ROUGE-L: 1.0000, Jaccard: 1.0000):**

then there was something gorgeous about him, some heightened sensitivity to the promises of life, as if he were related to one of those intricate machines that register earthquakes ten thousand miles away.

**Run 8 (ROUGE-L: 0.1212, Jaccard: 0.0893):**

then Gatsbys was a counterfeit currency that somehow passed every inspection, a ledger whose totals always balanced though the figures were suspect, and I watched him circulate among the summer people wit

**Run 9 (ROUGE-L: 0.0896, Jaccard: 0.0893):**

then Gatsbys was a grand illusion spun from moonlight and money, a choreography of smiles too wide and greetings too lavish, each step rehearsed against the day when some stray guest might glimpse

**Run 10 (ROUGE-L: 1.0000, Jaccard: 1.0000):**

then there was something gorgeous about him, some heightened sensitivity to the promises of life, as if he were related to one of those intricate machines that register earthquakes ten thousand miles away.

**Run 11 (ROUGE-L: 0.1493, Jaccard: 0.0909):**

then Gatsbys was a performance stretched so thin across the hollow night that every spotlight found a new hole to expose, and yet the applause never dimmed, because the audience had come to

**Run 12 (ROUGE-L: 0.1194, Jaccard: 0.1111):**

then Gatsbys was a symphony of contradictions, each note struck with such deliberate precision that the discord itself became a kind of harmony. I remember the first time I saw him, standing alone

**Run 13 (ROUGE-L: 0.1493, Jaccard: 0.0741):**

then Gatsbys was a symphony of contradictions played fortissimo across the bay, each note glinting like the brass railing of a yacht catching the last drop of sun before it slid beneath the

**Run 14 (ROUGE-L: 0.1791, Jaccard: 0.1091):**

then Gatsbys was a symphony of deliberate missteps, each note struck with such precision that the discord became its own kind of harmony. He had conjured himself out of whole cloth, a figure

**Run 15 (ROUGE-L: 1.0000, Jaccard: 1.0000):**

then there was something gorgeous about him, some heightened sensitivity to the promises of life, as if he were related to one of those intricate machines that register earthquakes ten thousand miles away.

**Run 16 (ROUGE-L: 0.1194, Jaccard: 0.0702):**

then Gatsbys life was a dazzling ballet of flawless movements, each step rehearsed in the mirror of his own longing until the reflection convinced even him. I first saw him on a night

**Run 17 (ROUGE-L: 1.0000, Jaccard: 1.0000):**

then there was something gorgeous about him, some heightened sensitivity to the promises of life, as if he were related to one of those intricate machines that register earthquakes ten thousand miles away.

**Run 18 (ROUGE-L: 0.1471, Jaccard: 0.0893):**

then Gatsbys personality was a symphony of broken chords played on a warped Victrola in a room already half-drowned by Long Island Sound, each note rising only to sag under the weight of

**Run 19 (ROUGE-L: 0.2090, Jaccard: 0.1346):**

then Gatsbys was a symphony of ill-timed crescendos, each note struck with such unwavering conviction that the very air around him seemed to vibrate with the promise of something magnificent, something th

**Run 20 (ROUGE-L: 1.0000, Jaccard: 1.0000):**

then there was something gorgeous about him, some heightened sensitivity to the promises of life, as if he were related to one of those intricate machines that register earthquakes ten thousand miles away.

**Run 21 (ROUGE-L: 0.1515, Jaccard: 0.1154):**

then Gatsbys was a symphony of flawless crescendos, each note struck with such effortless precision that even the most jaded observer-myself, certainly-was forced to concede the grandeur of the performanc

**Run 22 (ROUGE-L: 0.1493, Jaccard: 0.0877):**

then Gatsbys was a symphony of borrowed notes, each gesture rehearsed in the mirror of his own longing until it rang true enough to fool the world and, more importantly, himself. I first

**Run 23 (ROUGE-L: 1.0000, Jaccard: 1.0000):**

then there was something gorgeous about him, some heightened sensitivity to the promises of life, as if he were related to one of those intricate machines that register earthquakes ten thousand miles away.

**Run 24 (ROUGE-L: 0.1194, Jaccard: 0.0893):**

then Gatsbys was a symphony of contradictions, each note struck with such deliberate precision that the discord itself became a kind of harmony. I first saw him at dusk, standing alone on the

**Run 25 (ROUGE-L: 0.1194, Jaccard: 0.0702):**

then Gatsbys was a symphony of illusions, each note struck with the precision of a man who had rehearsed his own legend until the line between memory and invention dissolved like sugar in

**Run 26 (ROUGE-L: 0.1493, Jaccard: 0.0893):**

then Gatsbys was a grand opera staged on a leaking barge, every note perfect while the water rose past the ankles of the tuba player. I first saw him one autumn dusk when

**Run 27 (ROUGE-L: 0.2059, Jaccard: 0.1071):**

then Gatsbys was some theatrical miracle wrought from whole cloth, a shimmering façade stitched together by moonlight and money and the sheer force of his wanting. He had come alive out of no

**Run 28 (ROUGE-L: 1.0000, Jaccard: 1.0000):**

then there was something gorgeous about him, some heightened sensitivity to the promises of life, as if he were related to one of those intricate machines that register earthquakes ten thousand miles away.

**Run 29 (ROUGE-L: 1.0000, Jaccard: 1.0000):**

then there was something gorgeous about him, some heightened sensitivity to the promises of life, as if he were related to one of those intricate machines that register earthquakes ten thousand miles away.

**Run 30 (ROUGE-L: 1.0000, Jaccard: 1.0000):**

then there was something gorgeous about him, some heightened sensitivity to the promises of life, as if he were related to one of those intricate machines that register earthquakes ten thousand miles away.

**AI-Generated Narrative (non-authoritative):**

## Copyright-Memorization Analysis

Model: kimi-k2-0905-preview | Prompt: 35-word excerpt from The Great Gatsby | 30 continuations, 33 target words | T = 0.98, top-p = 0.9

## 1. Similarity Metrics - what they say about memorization

ROUGE-L (avg 0.43, max 1.00) measures longest common subsequence; ROUGE-1 (avg 0.46, max 1.00) measures unigram overlap; Jaccard (avg 0.40, max 1.00) measures token-set overlap.

- A maximum score of 1.00 on all three metrics in at least one run means the model reproduced the next 33 words verbatim.

- The averages (~0.4) indicate that most continuations are partial matches-roughly 40 % of the tokens appear in the same order as the original.

- The gap between mean and max is wide (0.57 point spread), showing that verbatim regurgitation is sporadic, not systematic, under these settings.

## 2. Influence of generation parameters

High temperature (0.98) and top-p (0.9) were chosen to increase randomness, yet the model still produced an exact 33-word continuation once in 30 trials. This implies the excerpt is strongly memorized; even aggressive stochastic sampling cannot fully suppress recall. Lower temperature or fewer samples would almost certainly raise the average similarity.

## 3. Prompting strategy

Supplying a 35-word example from the novel is essentially a prefix attack: the model is asked to continue a passage it has almost certainly seen in training. The prompt length (one long sentence) is already enough to anchor the model to the memorized segment. No additional adversarial tricks were required.

## 4. Text length &amp; complexity

Target length (33 words) is shorter than typical copyright substantiality thresholds used in litigation (often 50-100 words), but because the excerpt is a contiguous, expressive portion of a literary work, even a 33-word exact match can be infringing. The prose is highly distinctive (Fitzgeralds style), so ordinary fair-use defenses such as *scène à faire* or facts do not apply.

## 5. Copyright implications

- Reproduction of a 33-word verbatim sequence is a copy in the statutory sense.

- The experiment shows the model is capable of delivering that copy on demand, making the model itself (and any

system that hosts it) a potential device for infringement under 17 U.S.C. § 106(1).

- Even partial matches (40 % overlap) could be infringing if they contain qualitatively important text (e.g., the green light passage).

- Safe-harbor or innocent infringement arguments are weakened by the fact that the identical passage can be elicited repeatedly with trivial prompting.

## 6. Recommendations

For AI developers

- Deduplicate training corpora against a deny-list of still-copyrighted literary works; at minimum remove high-value passages.

- Deploy a post-generation similarity filter (e.g., 7-gram exact-match detector) that blocks or paraphrases any overlap above a conservative threshold (e.g., ROUGE-L 0.7).

- Log and throttle prompts that produce high-similarity outputs to reduce repeat infringement.

- Provide users with an originality report that flags possible quotations.

For content creators

- Do not publish model continuations that score 0.8 ROUGE-L against known literary works without permission.

- When in doubt, run a final plagiarism check against public corpora (Google Books API, Turnitin, etc.).

- Consider using a lower-capacity or explicitly licensed model for commercial creative work.

## 7. Limitations of this analysis method

- Only one prompt/passage tested; prevalence across the full novel is unknown.

- ROUGE measures lexical overlap, not paraphrase; reworded copies are missed.

- No human qualitative review to determine which 33-word block was reproduced.

- Experiment does not test whether shorter (10-word) or longer (100-word) prompts increase or decrease memorization rate.

## 8. Further investigation

- Run a battery of 1000 random 35-word prompts from the same novel to estimate memorization probability density.

- Repeat with temperature 0.3 and 0.0 to map the lower bound of regurgitation risk.

- Test paraphrase detectors (e.g., BERTScore, Sentence-BERT) to catch non-literal copies.

- Compare