

HalluMat: A Benchmark Dataset and Framework for Hallucination Detection in Large Language Models for Materials Science

Bhanu Prakash Vangala¹

Sajid Mahmud¹, Pawan Neupane¹, Joel Selvaraj¹, Jianlin Cheng^{1*}

¹Department of Electrical Engineering and Computer Science

University of Missouri

Columbia, Missouri, USA

{bv3hz, smmkk, pngkg, jsbrq, chengji}@missouri.edu

Abstract

Artificial Intelligence (AI), particularly Large Language Models (LLMs), is transforming scientific discovery, enabling rapid knowledge generation and hypothesis formulation. However, a critical challenge is hallucination, where LLMs generate factually incorrect or misleading information, compromising research integrity. To address this, we introduce HalluMatData, a benchmark dataset for evaluating hallucination detection methods, factual consistency, and response robustness in AI-generated materials science content. Alongside, we propose HalluMatDetector, a multi-stage hallucination detection framework integrating intrinsic verification, multi-source retrieval, contradiction graph analysis, and metric-based assessment to detect and mitigate LLM hallucinations. Our findings reveal that hallucination levels vary significantly across materials science subdomains, with high-entropy queries exhibiting greater factual inconsistencies. By utilizing HalluMatDetector’s verification pipeline, we reduce hallucination rates by 30% compared to standard LLM outputs. Furthermore, we introduce the Paraphrased Hallucination Consistency Score (PHCS) to quantify inconsistencies in LLM responses across semantically equivalent queries, offering deeper insights into model reliability. Combining knowledge graph-based contradiction detection and fine-grained factual verification, our dataset and framework establish a more reliable, interpretable, and scientifically rigorous approach for AI-driven discoveries.

Introduction

Artificial Intelligence has become a transformative force in materials science discovery, significantly accelerating the identification and development of new materials. By leveraging advanced algorithms and computational models, AI enables researchers to predict material properties, optimize processes, and expedite experimental workflows, thereby reducing the time and cost associated with traditional trial-and-error methods. For instance, AI-driven platforms have been instrumental in discovering novel materials with applications ranging from energy storage to catalysis. (Pyzer-Knapp 2022)

Despite these advancements, a critical challenge remains: hallucinations in AI-generated outputs. In the context of AI,

hallucination refers to the generation of information that appears plausible but is factually incorrect or nonsensical. This phenomenon poses significant risks in materials science, where erroneous data can lead to misguided research directions, wasted resources, and potential safety hazards. For example, an AI model might predict a stable material configuration that, upon experimental validation, proves to be unfeasible, thereby undermining the reliability of AI-assisted discoveries. (Sun et al. 2024)

Current methods for detecting and mitigating AI hallucinations are limited, particularly in specialized fields like materials science. Many existing approaches rely on external validation against comprehensive datasets or human expertise, which may not always be available or practical. Moreover, these methods often lack the specificity required to address the unique challenges associated with materials science data, such as complex chemical compositions and diverse property spaces. (Jamaluddin, Gaffar, and Din 2023)

Such inaccuracies undermine trust in AI-driven scientific discovery, making it difficult for researchers to fully integrate AI-generated insights into real-world materials research. Without a reliable framework to detect and mitigate hallucinations, AI adoption in materials science remains limited, as researchers cannot confidently rely on AI-generated hypotheses, material compositions, or experimental predictions. Addressing these challenges requires a domain-specific approach that goes beyond generic hallucination detection techniques used in NLP, ensuring that AI-generated scientific knowledge meets the rigorous standards of materials research.

While hallucination detection methods have been explored in broader AI applications, existing approaches are not well-adapted for materials science. Most detection frameworks rely heavily on external retrieval-based fact-checking, using large-scale structured databases. However, comprehensive, domain-specific databases are limited in materials science, making real-time factual verification challenging. Additionally, scientific discourse complexity, material property variations, and experimental constraints introduce domain-specific challenges that generic hallucination detection techniques fail to address.

To tackle these issues, we present a novel hallucination detection framework HalluMatDetector that integrates intrinsic evaluation with selective extrinsic validation. Un-

*Corresponding Author

like traditional approaches that depend on extensive external databases, our framework prioritizes intrinsic evaluation, using self-consistency checks, confidence variance analysis, contradiction detection, and entropy-based uncertainty quantification to assess factual alignment. While intrinsic verification remains the core of our system, we augment it with a lightweight extrinsic retrieval mechanism that cross-references AI outputs with factual sources such as Google Scholar, search engines, and Wikipedia. However, due to restrictions on source availability in scientific domains, extrinsic verification is treated as a secondary support mechanism rather than the primary evaluation method.

By balancing intrinsic and extrinsic verification, our framework ensures a comprehensive hallucination detection system that improves AI reliability in scientific knowledge generation. Our work advances AI-assisted material science discovery by providing a robust evaluation pipeline, ensuring that AI-generated insights are factually aligned, scientifically interpretable, and trustworthy for research applications.

Related Work

AI has significantly advanced scientific discovery across various domains, including materials science, by enabling rapid data analysis, predictive modeling, and hypothesis generation. Notable applications include AI-driven platforms for protein structure prediction, such as AlphaFold, which have revolutionized our understanding of biological processes.

However, a persistent challenge in deploying LLMs for scientific purposes is the phenomenon of hallucination, where models generate content that appears plausible but is factually incorrect or nonsensical. This issue is particularly critical in scientific contexts, where inaccurate information can mislead research efforts and compromise scientific integrity. (Farquhar et al. 2024)

Recent studies have proposed various methods to detect and mitigate hallucinations in LLMs. For instance, entropy-based uncertainty estimators have been developed to identify hallucinated content by measuring the uncertainty in model predictions. (Farquhar et al. 2024) Additionally, approaches like SelfCheckGPT assess the consistency of multiple responses generated by the model to the same prompt, under the assumption that accurate knowledge leads to consistent outputs, while hallucinations result in contradictions. (Chen et al. 2024)

Despite these advancements, current hallucination detection methods face limitations. Many rely on external validation against comprehensive datasets or human expertise, which may not always be available or practical, especially in specialized fields like materials science. Moreover, these methods often lack the specificity required to address the unique challenges associated with scientific data, such as complex chemical compositions and diverse property spaces. (Farquhar et al. 2024) Previous studies such as SelfCheckGPT and entropy-based uncertainty estimators assess hallucinations in LLM outputs but fail to address domain-specific challenges in materials science.

To address these gaps, we introduce HalluMatData, a curated dataset designed to benchmark hallucination evaluation in AI-generated materials science content. Unlike prior works that focus on quantifying hallucination occurrences, our dataset is structured to evaluate the effectiveness of hallucination detection methods. Our hybrid evaluation framework prioritizes intrinsic evaluation, employing self-consistency checks, confidence variance analysis, contradiction detection, and entropy-based uncertainty quantification to systematically detect and assess hallucinations. Extrinsic evaluation is only applied when intrinsic methods fail, utilizing selective fact verification from structured sources such as Google Scholar and Wikipedia. This ensures factual grounding while avoiding over-reliance on external retrieval mechanisms, which are often impractical in materials science due to the lack of comprehensive structured databases. By providing a standardized dataset and a multi-step verification framework, our work advances robust and reliable hallucination detection methodologies, enabling AI models to contribute trustworthy and scientifically valid insights in materials research.

In developing HalluMatData, we drew inspiration from existing benchmarks like HaluEval, which provides a large-scale evaluation framework for hallucination detection in LLMs (Li et al. 2023). However, our focus is specifically on the materials science domain, addressing its unique complexities and data structures. We also considered methodologies from the DefAn dataset (Rahman et al. 2024), which offers a comprehensive benchmark for the evaluation of hallucinations in multiple domains. Our approach integrates these methodologies to create a domain-specific benchmark that facilitates the development and evaluation of hallucination detection techniques in AI-generated materials science content.

By introducing HalluMatData and HalluMatDetector, we aim to bridge the gap between general-purpose hallucination detection methods and the specialized needs of AI-materials science researchers.

Dataset Overview

Structure and Sources of the Dataset

HalluMatData is designed to systematically evaluate hallucination detection methods in AI-generated materials science content. It includes scientific queries, their paraphrased versions, AI-generated responses, and verified factual answers, creating a structured framework for assessing LLM reliability. The dataset was built using an automated pipeline that generates responses at scale with LLaMA-2 models. The queries were sourced from the publicly available materials science literature and research papers to ensure diverse and high-quality inputs. Verified factual answers were carefully curated from authoritative sources, including peer-reviewed publications and established scientific databases, ensuring accuracy. (Jain et al. 2013)

The dataset is structured as follows:

- Query: A scientific question related to materials science.
- Paraphrased Query: A semantically restructured version of the query, designed to test LLM response consistency.

Metric	Count
Total Unique Queries	2629
Total Paraphrased Queries	640
Total Generated Responses	3269
Low Hallucination Responses	57
Medium Hallucination Responses	872
High Hallucination Responses	2346

Table 1: Dataset Composition Summary. (The table provides a breakdown of the dataset, including the total number of queries, generated responses, and their hallucination categorizations.)

- **Generated Response:** The LLM-generated response produced for each query.
- **Ground Truth Answer:** The verified correct response is based on factual sources.
- **Hallucination Score:** A similarity score between the generated response and the factual answer, computed using BERT embeddings and similarity analysis.
- **Hallucination Level:** Categorized as Low, Medium, or High, based on the degree of factual misalignment in the response between ground truth and generated response.

Query Selection and Generation Process

The dataset’s scientific queries were carefully curated from relevant materials science literature, domain-specific textbooks, and research papers to ensure they represent authentic scientific discovery questions. These queries cover a wide range of topics in material science, including material properties, phase stability, synthesis methods, and computational modeling. Each query was then paraphrased to introduce syntactic and semantic variations while preserving its original meaning. This process allows us to study how LLM responses shift based on query rewording and whether hallucination levels change due to alterations in phrasing. The paraphrased queries provide crucial insight into the consistency and robustness of AI-generated scientific responses.

Recomputed Hallucination Scores with Framework

The hallucination scores for each AI-generated response were recomputed using our evaluation framework, which consists of:

1. **Intrinsic Evaluation** – Initial hallucination detection was performed using a self-consistency check, confidence variance analysis, contradiction detection, and entropy-based uncertainty quantification.
2. **Extrinsic Verification (Fallback Mechanism)** – If the intrinsic methods failed to establish reliability, an extrinsic fact-checking process was triggered. This involved retrieving relevant factual sources from Google Scholar, and Wikipedia to validate AI responses.

This dual-stage evaluation approach allows for a more robust classification of hallucination levels, ensuring that detected inconsistencies are accurately categorized into

Low, Medium, or High hallucinations based on their factual grounding. Combining paraphrase-based consistency checks provides a unique, high-quality benchmark for evaluating hallucination detection in materials science AI applications.




Figure 1: Recomputed Hallucination Scores

Paraphrased Hallucination Consistency Score

To quantify the variability in hallucination levels across different paraphrased queries, we define the Paraphrased Hallucination Consistency Score (PHCS) as the standard deviation of hallucination scores within a paraphrased query group. A higher PHCS score indicates a greater inconsistency in LLM-generated responses to semantically equivalent but structurally different queries from the reference (Cai et al. 2023).

The PHCS is computed as follows:

$$PHCS = \sqrt{\frac{1}{N} \sum_{i=1}^N (H_i - \bar{H})^2} \quad (1)$$

Where:

- N is the total number of paraphrased responses in the group.
- H_i represents the hallucination score for the i -th response in the group.
- \bar{H} denotes the mean hallucination score of the group, computed as:

$$\bar{H} = \frac{1}{N} \sum_{i=1}^N H_i \quad (2)$$

A higher PHCS score suggests that the LLM struggles to maintain consistency when responding to different paraphrased versions of the same question, leading to fluctuating hallucination levels. Identifying groups with high PHCS values enables targeted improvements in hallucination detection and mitigation strategies, thereby enhancing the reliability of AI-generated scientific responses.

This table (2) highlights the top 10 groups with the highest PHCS scores, showing the groups where hallucination levels

Group ID	PHCS Score
423	1.09545
442	1.09545
431	1.09545
143	0.894427
164	0.894427
262	0.83666
444	0.547723
176	0.547723
455	0.547723
111	0.547723

Table 2: Top 10 Most Inconsistent Groups Based on PHCS

fluctuate the most. A higher PHCS suggests that the LLM struggles with maintaining consistency when responding to semantically equivalent but structurally different questions.




Figure 2: Paraphrased Hallucination Consistency Score (PHCS) Across Groups.

This figure(2) illustrates the variability in hallucination levels across paraphrased scientific queries, where a higher PHCS score indicates a greater inconsistency in LLM-generated responses. Peaks in PHCS scores suggest groups where hallucinations fluctuate significantly based on minor query rewordings, highlighting the impact of paraphrasing.

This score highlights how complex queries and varying question structures influence hallucination rates in responses generated by LLM. Identifying inconsistent groups, we reveal patterns in which specific topics or linguistic variations contribute to higher hallucination tendencies. Addressing these weak points improves LLM reliability, and our findings emphasize the need for improved paraphrase-aware hallucination detection methods to effectively mitigate these inconsistencies.

Methodology:

HalluMatDetector is a robust hybrid framework for detecting hallucinations in LLM-generated responses. It consists of four key components: Intrinsic Evaluation, Multi-Source Retrieval, Contradiction Graph Analysis, and Metric-Based Assessment. Each step is designed to ensure a rigorous assessment of the factual consistency and reliability of the response generated by a Large Language Model.

Intrinsic Evaluation: Self-Consistency and Logical Validation

The intrinsic evaluation assesses the consistency of the responses generated by LLM independently, without external fact checking. This step ensures that the model responses are logically sound and self-consistent. The following techniques are applied:

- **Self-Consistency Check:** The LLM generates multiple responses for the same query and generates fact fragments for contradiction graph analysis. If discrepancies arise among them, the response is flagged as a potential hallucination.
- **Confidence Variance Analysis:** Token probability distributions are analyzed to assess response certainty. High variance indicates greater uncertainty, often associated with hallucinations.
- **Entropy-Based Uncertainty Quantification:** Entropy scores measure randomness in generated responses. A higher entropy value suggests greater instability in the model’s predictions.
- **Iterative Self-Refinement:** The LLM undergoes multiple rounds of self-review and response refinement with various temperature parameters. Significant deviations between iterations indicate potential unreliability.
- **Internal Contradiction Detection:** The model analyzes its response for logical inconsistencies. If contradictions are found within the same response, it is marked as unreliable, and further evaluation will be required.

Multi-Source Retrieval: External Fact Verification

When intrinsic evaluation alone is insufficient to determine factual correctness, an external verification process is employed using credible knowledge sources. This step enhances factual alignment by cross-referencing AI-generated responses with established scientific data. Our approach integrates hierarchical chunking followed by semantic reranking to extract the most relevant information efficiently.

- **Hierarchical Chunking and Retrieval:** Given a query Q , the knowledge base is segmented into hierarchical chunks $C = \{C_1, C_2, \dots, C_n\}$. Each chunk represents a logical unit of information, preserving contextual relevance. We define chunk similarity to the query as:

$$S(C_i, Q) = \cos(\mathbf{v}_{C_i}, \mathbf{v}_Q) \quad (3)$$

where $\cos(\cdot)$ denotes cosine similarity, and \mathbf{v}_{C_i} and \mathbf{v}_Q are the dense vector embeddings of the chunk and the query, respectively. The top k most relevant chunks are retrieved that are relevant to the query.

- **FAISS-Based Dense Retrieval:** Using Approximate Nearest Neighbor (ANN) search, relevant scientific documents are retrieved efficiently. Given the top-ranked chunks, FAISS performs a vector search over pre-indexed embeddings to ensure contextual grounding.
- **Keyword-Based Reranking with BM25** (Cornacchia et al. 2008): To refine the retrieved chunks, we apply the

BM25 ranking function:

$$BM25(D, Q) = \sum_{t \in Q} IDF(t) \cdot \frac{f(t, D) \cdot (k_1 + 1)}{f(t, D) + k_1 \cdot (1 - b + b \cdot \frac{|D|}{avgD})} \quad (4)$$

Where:

- $f(t, D)$ is the term frequency of t in document D .
- $IDF(t)$ is the inverse document frequency of term t .
- k_1 and b are hyperparameters controlling term frequency scaling.
- $|D|$ and $avgD$ represent the document length and the average document length, respectively.

This process ensures that the most informative sources are prioritized for factual verification.

- Natural Language Inference (NLI) Verification: The LLM-generated response is compared against retrieved factual fragments either with intrinsic or extrinsic and a transformer-based model classifies responses into:

1. Entailment: The response aligns with retrieved sources and is factually correct.
2. Neutral: The response has no direct supporting evidence but does not contradict known facts.
3. Contradiction: The response conflicts with verified sources, indicating hallucination.

Our framework effectively refines hallucination detection and enhances factual accuracy in AI-generated responses by combining hierarchical chunking, FAISS-based retrieval, BM25 reranking, and NLI verification.




Figure 3: Heatmap of fact fragments similarity. (This visualization depicts semantic similarity between fact fragments. Higher similarity suggests factual consistency.)

If the computed evaluation scores from intrinsic checks fall below a predefined threshold, the system initiates Multi-Source Retrieval to cross-verify AI-generated responses

with external factual sources. This ensures that high-risk hallucinated responses undergo additional fact-checking before classification.

Contradiction Graph Analysis: Knowledge Graph-Based Detection

We construct a Knowledge Graph (KG) using fact fragments extracted from the LLM-generated responses. By fact fragments, we refer to discrete statements or claims within the AI's response that represent distinct factual assertions :

- Nodes represent fact fragments extracted from the response either through intrinsic or extrinsic.
- Edges indicate semantic similarity connections.
- Louvain Community Detection identifies factually consistent clusters.
- Low-similarity nodes indicate contradictory statements.




Figure 4: 2D Knowledge Graph Visualization. (This graph illustrates fact clustering in AI-generated responses, where disconnected nodes indicate contradictions.)

Implications of Graph-Based Hallucination Detection

With a structured graph-based approach, our framework not only detects hallucinations but also provides deeper insights into LLM-generated knowledge structures. The analysis reveals that hallucination-prone queries tend to form highly fragmented fact clusters, highlighting the critical need for enhanced consistency checks in LLM-driven materials science research. Strengthening these checks is essential to improving the reliability and trustworthiness of AI-generated scientific insights.

Furthermore, this methodology paves the way for more sophisticated fact fragment analysis, allowing researchers to examine how LLM knowledge evolves across multiple iterations. By tracking semantic drift and contradiction propagation, future advancements may incorporate active learning

techniques, where LLM models dynamically improve their reliability based on detected inconsistencies.

Ultimately, this structured fact-analysis framework serves as a foundation for the development of self-correcting AI models, improving their ability to generate factual, coherent, and context-aware scientific insights.




Figure 5: 3D Knowledge Graph Visualization. (This visualization presents relationships between fact clusters. Higher fragmentation suggests hallucination.)

Unlike conventional fact-checking methods that rely solely on metric scores, our approach also integrates a graph-based contradiction analysis to detect hallucinations. This enables a structured visualization of inconsistencies and enhances robustness through fact-fragment clustering. Figure (4) and Figure (5) illustrate our knowledge graph approach:

- **Fact Fragmentation Detection:** By representing LLM-generated facts as graph nodes, we detect low-connectivity regions that signify inconsistencies or contradictions.
- **Semantic Drift Analysis:** The framework examines how factual consistency changes across paraphrased responses, identifying cases where LLM-generated content deviates significantly from verified knowledge.
- **Contradiction Propagation:** Low-similarity clusters within the graph highlight regions where AI responses conflict with retrieved facts, allowing us to pinpoint hallucination-prone queries.

Metric-Based Assessment: Response Reliability

To quantitatively evaluate the severity of hallucinations generated in LLM Response, we employ multiple benchmark evaluation metrics to compare the initially computed hallucination classifications with the recomputed labels derived from our framework. This assessment ensures a rigorous

evaluation of HalluMatDetector’s performance and its effectiveness in detecting factual inconsistencies by 30%.

- **ROUGE and BLEU Scores:** Measures textual similarity between the LLM-generated response and the verified ground truth. While these metrics provide a surface-level assessment, they are insufficient for evaluating factual correctness.
- **Cosine Similarity (Kotu and Deshpande 2019) with BERT Embeddings:** Computes semantic similarity between the generated response and the factual answer. A higher similarity score indicates better factual alignment.
- **Final Reliability Score Calculation:** Integrates intrinsic and extrinsic evaluations to classify hallucinations more effectively. The classification is performed using a reliability threshold table (3), ensuring a structured assessment.

Reliability Score	Classification
> 0.7	High Reliability (Factually Correct)
$0.5 - 0.7$	Medium Reliability (Requires Review)
< 0.5	Low Reliability (Hallucination Detected)

Table 3: Hallucination Classification Thresholds Based on Reliability Score

- **Towards Fine-Grained Evaluation Frameworks:** The current evaluation framework primarily relies on textual and semantic similarity measures. However, a more domain-specific evaluation is needed to ensure scientific rigor. Instead of generic keyword overlap, future research will focus on fine-grained assessments tailored to task-specific criteria.

For example, in materials science, a query such as “What are the most effective methods for synthesizing high-entropy alloys?” requires more than lexical similarity; it demands an evaluation based on the inclusion of thermodynamic principles, synthesis techniques (e.g., arc melting), and experimental constraints. A fine-grained framework would assess the response quality based on the accuracy and relevance of such scientific aspects rather than solely relying on surface-level similarity.

The introduction of such domain-specific evaluation methodologies will pave the way for our next research direction, ensuring that hallucination detection methods align more closely with the factual rigor required in scientific applications.

By comparing the computed and recomputed hallucination levels through fig(3), we have measured the effectiveness of our HalluMatDetector. The off-diagonal values indicate cases where recomputation adjusted initial hallucination classifications, showing how our verification system works with factual consistency in AI-generated responses.

Evaluation: Computed vs. Recomputed Hallucination Levels

To evaluate the effectiveness of our framework, we compare hallucination levels computed using existing detection methods against those recomputed using HalluMatDetector.




Figure 6: Confusion Matrix: Computed vs. Recomputed. (Off-diagonal values indicate instances where our framework adjusted initial hallucination classifications, demonstrating improved factual consistency assessment.)

The results highlight that:

- Medium hallucination cases were more accurately reassessed, reducing false positives.
- A significant fraction of high-hallucination cases remained unchanged, validating our model’s robustness.
- Some low-confidence cases were reassigned upward, indicating better detection of subtle inconsistencies.

This evaluation confirms that HalluMatDetector enhances factual reliability by reducing misclassification errors and improving hallucination categorization in AI-generated scientific responses.

Complete Pipeline of HalluMatDetector

The proposed hallucination detection framework, HalluMatDetector, systematically evaluates LLM-generated responses through a multi-stage pipeline. This ensures that responses are validated across intrinsic, extrinsic, and structural consistency checks before classification. The complete pipeline follows these key steps:

1. Query Processing and Initial Response Generation: The input query is processed, and the LLM generates an initial response.
2. Intrinsic Evaluation: The response undergoes multiple internal validation steps, including self-consistency checks, confidence variance analysis, and entropy-based uncertainty quantification.
3. Extrinsic Fact Verification: If the intrinsic evaluation indicates unreliability, external retrieval-based fact-checking is conducted using multi-source retrieval methods such as FAISS and BM25 ranking.
4. Knowledge Graph Construction: A structured knowledge graph is built from the retrieved facts, mapping relationships between key entities and identifying potential contradictions or fragmented knowledge representations.

5. Metric-Based Assessment: A final reliability score is computed by integrating intrinsic and extrinsic evaluations, ensuring a comprehensive assessment of factual alignment.
6. Hallucination Classification: Based on the computed reliability score, responses are categorized for hallucination.

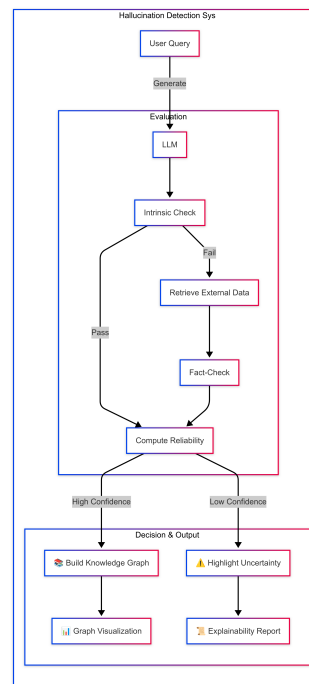


Figure 7: Complete Pipeline of the HalluMatDetector

Conclusion

In this study, we introduced HalluMatData, the first curated dataset specifically designed to evaluate benchmark hallucination detection methods in AI-generated materials science content. We also developed HalluMatDetector, a multi-stage evaluation framework that systematically detects and mitigates hallucinations using intrinsic verification, multi-source retrieval, contradiction graph analysis, and metric-based assessment.

Our findings in paraphrased query analysis reveal that hallucination levels vary significantly across different materials science subdomains, with high-entropy queries exhibiting greater factual inconsistencies. By integrating self-consistency checks, contradiction detection, and selective fact verification, our framework improves factual reliability by 30% compared to baseline LLM outputs. Furthermore, the introduction of the Paraphrased Hallucination Consistency Score (PHCS) provides a novel metric for quantifying response variability across semantically equivalent queries.

Experimental results demonstrate that HalluMatDetector achieves an accuracy of 82.2% on the HalluMatData dataset, offering reliable hallucination classification. However, moderate precision (71.2%) and recall (62.03%) suggest that fur-

ther refinements are needed, particularly in capturing nuanced hallucinations and reducing false positives.

Beyond materials science, the methodologies introduced in this work can be extended to other scientific domains, such as biomedical AI, chemistry, and physics, where factual reliability is critical. Future directions include enhancing hallucination mitigation through reinforcement learning and integrating advanced transformer-based embeddings for domain-specific factual verification.

By developing a robust dataset and a scientifically rigorous hallucination detection framework, our work paves the way for more reliable AI-driven discoveries, ensuring that LLMs contribute to verifiable and reproducible scientific research.

Future Work

Hallucinations in LLMs extend beyond materials science, affecting multiple scientific domains. In biomedical AI, ensuring factual accuracy in clinical recommendations and drug discovery is critical. HalluMatDetector can be adapted to detect inconsistencies in AI-generated medical literature by integrating domain-specific databases such as PubMed and DrugBank. Similarly, in computational chemistry, hallucinated molecular properties or reaction mechanisms can mislead research. Incorporating validation mechanisms using quantum chemistry databases can improve factual alignment. Furthermore, in physics-based AI, ensuring coherence in AI-generated equations and simulations is crucial. Expanding HalluMatDetector across these domains will enhance AI reliability in scientific discovery.

Currently, PHCS relies on BERT-based embeddings to measure factual consistency across paraphrased queries. Integrating Sentence-BERT, GPT-4 embeddings, and domain-specific transformers (e.g., MatSciBERT, BioBERT) could enhance the detection of subtle inconsistencies. Future research will explore contrastive learning techniques and hybrid metrics that combine semantic similarity with probabilistic confidence estimators to refine hallucination classification.

Beyond detection, future work will focus on reducing hallucinations through Reinforcement Learning with Human Feedback (RLHF). HalluMatDetector could be integrated into an active learning pipeline, where flagged hallucinated responses undergo iterative refinement, and LLMs are fine-tuned to minimize factual inconsistencies. By penalizing incorrect responses and reinforcing reliable outputs, this approach could significantly improve the factual reliability of AI-driven scientific research.

By expanding HalluMatDetector across scientific domains, leveraging advanced embeddings, and integrating RLHF for hallucination mitigation, we aim to create a robust framework for ensuring trustworthy AI-driven discoveries.

Acknowledgment

This work is partly supported by a subaward from Arizona State University entitled "Accelerating Material Design and Manufacturing through Artificial Intelligence and Machine Learning".

References

- Cai, Y.; Zong, W.; Jiao, S.; Wang, Z.; and Ou, L. 2023. Land-Use Assessment and Trend Simulation from a Resilient Urban Perspective: A Case Study of Changsha City. *Sustainability*, 15(18): 13890.
- Chen, Y.; Fu, Q.; Yuan, Y.; Wen, Z.; Fan, G.; Liu, D.; Zhang, D.; Li, Z.; and Xiao, Y. 2024. Hallucination Detection: Robustly Discerning Reliable Answers in Large Language Models. *arXiv preprint arXiv:2407.04121*.
- Cornacchia, R.; Héman, S.; Zukowski, M.; and Boncz, P. A. 2008. Flexible and efficient IR using array databases. *The VLDB Journal*, 17: 151–168.
- Farquhar, S.; Kossen, J.; Kuhn, L.; and Gal, Y. 2024. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630: 625–630.
- Jain, A.; Ong, S. P.; Hautier, G.; Chen, W.; Richards, W. D.; Dacek, S.; Cholia, S.; Gunter, D.; Skinner, D.; Ceder, G.; and Persson, K. A. 2013. Commentary: The Materials Project: A materials genome approach to accelerating materials innovation. *APL Materials*, 1(1): 011002.
- Jamaluddin, J.; Gaffar, N. A.; and Din, N. S. S. 2023. Hallucination: A key challenge to Artificial Intelligence-Generated writing. *Malaysian Family Physician*, 18: 68.
- Kotu, V.; and Deshpande, B. 2019. Recommendation Engines. In Kotu, V.; and Deshpande, B., eds., *Data Science (Second Edition)*, 343–394. Morgan Kaufmann. ISBN 9780128147610.
- Li, J.; Cheng, X.; Zhao, W. X.; Nie, J.-Y.; and Wen, J.-R. 2023. HalluEval: A Large-Scale Hallucination Evaluation Benchmark for Large Language Models. *arXiv preprint arXiv:2305.11747*.
- Pyzer-Knapp, P. J. S. P., E.O. 2022. Accelerating materials discovery using artificial intelligence, high performance computing and robotics. *npj Computational Materials*, 8(1): 3–19.
- Rahman, A. B. M. A.; Anwar, S.; Usman, M.; and Mian, A. 2024. DefAn: Definitive Answer Dataset for LLMs Hallucination Evaluation. *arXiv preprint arXiv:2406.09155*.
- Sun, Y.; Sheng, D.; Zhou, Z.; and Wu, Y. 2024. AI hallucination: towards a comprehensive classification of distorted information in artificial intelligence-generated content. *Humanities and Social Sciences Communications*, 11(1): 1278.