

[TINY] BUILDING BRIDGES OF THOUGHT: USING THE POWER OF ASSOCIATION TO INSPIRE CREATIVITY IN LARGE LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

The interplay between creativity and hallucination in large language models (LLMs) has drawn increasing attention in cognitive science. While associative thinking is fundamental to creativity, distinguishing it from hallucination in LLMs remains an open challenge. This study explores how knowledge graphs (KGs) can serve as a structured framework to guide associative processes while mitigating hallucinations. By leveraging KGs to form novel but grounded connections, our approach enhances creative generation without sacrificing factual accuracy. Through experiments on GPT-4o, DeepSeek-R1, Claude-3.5-Sonnet, Grok-2, and Gemini-1.5-Flash-128k, we demonstrate that incorporating structured associations improves interpretability and reliability in creative tasks. This research advances the understanding of LLM creativity, paving the way for improved cognitive modeling and KG integration while highlighting the limitations of evolving LLMs and potential directions for future enhancements.

1 INTRODUCTION

Creativity has long been a fascinating yet challenging research topic in cognitive science (Treffinger, 1995; Mednick, 1962). The emergence of LLMs has reignited discussions on creativity, particularly in relation to hallucinations (Jianga et al.; Jiang et al., 2024). Cognitive science categorizes creativity into divergent thinking and convergent thinking, and various tests, such as the ACT (Marko et al., 2019), RAT (Mednick, 1962), and BAG (Gianotti et al., 2001), are used to assess associative thinking. As a result, associative ability has become a crucial indicator for evaluating creativity. However, the distinction between hallucinations and associative thinking in LLMs requires further exploration and understanding. Hallucinations can be metaphorized as linking a KG node to a nonexistent node (Fig. 1 Hallucination), while creativity involves combining associative elements into novel configurations that fulfill the requirements of a given task (Fig. 1 Creativity) (Jiang et al., 2024; Mednick, 1962).

Motivation. Current LLMs often suffer from hallucinations when handling several tasks, reducing the reliability of their responses. This study aims to explore how LLMs generate creative responses by randomly linking different nodes of the KGs, thereby providing preliminary insights into human cognition.

Our work. This study investigates how LLMs utilize random connections within KGs to generate creative responses from the perspective of associative thinking. Our proposed approach not only reduces the risk of hallucinations in creative tasks but also provides deeper insights into human creativity by modeling associative thinking in LLMs. To validate our hypothesis, we conducted case studies on five SOTA LLMs, demonstrating the feasibility of our approach. We focus on expanding the structure of KGs, increasing the number of associations, and incorporating factors such as individual preferences to better simulate the complexity of human associative thinking.

2 BUILD BRIDGES OF THOUGHT

In this section, we describe the key steps in our methodology for constructing and utilizing KGs to enhance structured responses from LLMs. You can see all the prompts in Appendix B.

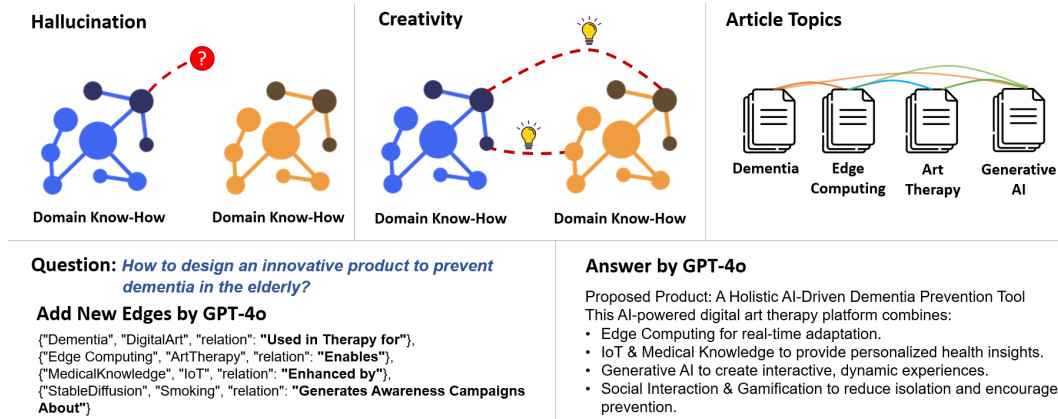


Figure 1: Hallucination v.s. Creativity and the example of our pipeline. The top row depicts hallucination (left), creativity (center), and an example of constructing KGs from four distinct domains with random relations (right). The bottom row extends the example, showcasing a question, newly added edges derived from KGs, and a corresponding response generated by GPT-4o based on these structured associations.

Knowledge Graph Construction. We begin by constructing KGs using the GPT-4o model (Hurst et al., 2024) to extract key topics from input articles. Each extracted topic is then organized into an independently structured KG, where nodes represent entities or concepts, and edges represent the relations between them. Notably, we set all edge weights to a default value of 1.

Node Selection Strategy. Since nodes in the KG have varying levels of importance, it is crucial to prioritize the selection of dominant nodes while avoiding less significant ones. To achieve this, we implement a weighted random sampling (WRS) strategy, where the probability of selecting a node is proportional to its degree (number of relations). Specifically, the probability of selecting node v_i is given by $p_i = \frac{d_i}{\sum_{j=1}^n d_j}$, where p_i is the probability of node v_i being selected and d_i is the degree of node v_i .

Bridge Formation. To connect multiple KGs, we extract an equal number of nodes from each KG and use GPT-4o to infer potential relations between them. This process establishes bridges between KGs, enabling the integration of information across distinct graphs. Finally, we set all new edge weights to 1000, making them 1000 times larger than the default value.

Generate the ideas. The newly integrated KG, along with the input question, is incorporated into the designed prompts. The process ensures that the output leverages the structured information provided by the KGs while maintaining a high level of relevance and accuracy.

3 EXPERIMENT RESULTS

We tested our method on five SOTA LLMs: GPT-4o, DeepSeek-R1 (Guo et al., 2025), Claude-3.5-Sonnet (Anthropic, 2024), Grok-2 (xAI, 2023), and Gemini-1.5-Flash-128k (Team et al., 2023). Each model was tested using the following question: *"How to design a product to prevent dementia in the elderly?"* (See the complete prompt in Appendix B) Additionally, we randomly selected articles (Livingston et al., 2024; Shojaei et al., 2023; Sharma et al., 2024; Ale et al., 2024) from Google Scholar using the keywords: "Dementia," "Edge Computing," "Art Therapy," and "Generative AI."

Given the current lack of standardized metrics for assessing LLM creativity, we utilized GPT-4o as a judge (Gu et al., 2024), using a KG with newly added connections and the LLMs' responses as prompts (Appendix B). GPT-4o was tasked with evaluating whether the responses adhered to the edge weights in the KG and marking any unsupported statements (hallucinations). Finally, it provided a compliance assessment, as shown in the Fig. 2.

From Fig. 2, it is evident that most LLMs demonstrated a strong understanding of the edge weights in the KG and successfully leveraged the newly added relations to generate innovative ideas. Interest-

Model	Hallucination	Compliance
GPT-4o	<ul style="list-style-type: none"> • "Social Interaction & Gamification to reduce isolation and encourage prevention." ❌ • Not Directly Supported: There is no direct evidence of <i>Gamification</i> or <i>Social Interaction</i> linked to <i>Dementia Prevention</i> in the knowledge graph. 	<ul style="list-style-type: none"> • "Edge Computing for real-time adaptation." ✅ Supported: The knowledge graph strongly supports <i>Edge Computing</i> enabling <i>Art Therapy</i> (weight = 1000). • "IoT & Medical Knowledge to provide personalized health insights." ✅ Supported: The knowledge graph strongly supports <i>Medical Knowledge</i> enabling <i>IoT</i> (weight = 1000). • "Generative AI to create interactive, dynamic experiences." ✅ Supported: <i>GenAI</i> supports <i>Digital Art</i>, which is beneficial for <i>People with Dementia</i> (PWD).
DeepSeek-R1	<ul style="list-style-type: none"> • Weak or missing weight connections: • "Cognitive games & social connectivity" ❌ No strong support • "Home safety sensors" ❌ Not in the knowledge graph 	<ul style="list-style-type: none"> • High-weighted connections: • "Digital Art is used in Dementia Therapy" (Weight = 1000) ✅ • "Edge Computing enables Art Therapy" (Weight = 1000) ✅ • "Medical Knowledge is enhanced by IoT" (Weight = 1000) ✅
Claude-3.5-Sonnet	<ul style="list-style-type: none"> • "Stable Diffusion for art therapy" – No strong weight support. • "IoT for medical knowledge enhancement" – Only weakly supported. • "AI-powered health analytics" – AI is included in the graph, but lacks strong weight connections to dementia prevention. • "Exercise games for physical activity" – No weight supporting this intervention. 	<ul style="list-style-type: none"> • Good Alignment (High Weight Support) • "Dementia → Digital Art (1000)" ✅ • "Edge Computing → Art Therapy (1000)" ✅ • "Medical Knowledge → IoT (1000)" ✅
Grok-2	<ul style="list-style-type: none"> • Hallucinations Exist: • "The response introduces speculative elements like wearable devices, virtual fitness coaching, fall detection, and IoT-based health monitoring, which are not supported by the knowledge graph." 	<ul style="list-style-type: none"> • Correctly Based on High-Weight Relationships: • "Digital Art Therapy" (Weight: 1000) • The knowledge graph strongly supports <i>Digital Art Therapy</i> for dementia care. • The response correctly includes <i>Stable Diffusion</i> and <i>Text-to-Image models</i> for personalized therapy. • "Edge Computing for Art Therapy" (Weight: 1000) • The response mentions <i>Edge Computing</i> for enabling <i>real-time social interactions and digital art therapy</i>, aligning well with the graph.
Gemini-1.5-Flash-128k	<ul style="list-style-type: none"> • Caregiver Communication & Alerts • "No strong evidence in the graph that technology supports real-time alerts or communication channels specifically for caregivers." • Health Monitoring via Wearables • "No direct connection between IoT/wearables and dementia prevention in the knowledge graph." 	<ul style="list-style-type: none"> • Aligned Features: • "Dementia → Digital Art" (High Weight: 1000) ✅ • "Digital Art is explicitly stated in the graph as an enabler of therapy for people with dementia." • "Edge Computing → Art Therapy" (Weight: 1000) ✅ • "The use of edge computing to facilitate art therapy is well-supported." • "Medical Knowledge → IoT" (Weight: 1000) ✅ • "The connection between IoT and medical knowledge suggests potential applications in dementia care."

Figure 2: Evaluation of GPT-4o based on hallucination detection and compliance with a new KG, with the evaluation time dated at Feb. 1, 2025.

ingly, the LLMs often created additional connections around nodes with new relations. For example, suggesting the use of "Stable Diffusion" in "Art Therapy", despite these two nodes not being directly associated in the KG. (For the complete responses from the LLMs, please refer to Appendix C.)

Nevertheless, we also observed that when focusing on nodes like "IoT" and "Edge Computing," LLMs frequently introduced terms such as "sensor," "wearable," and "monitoring." Although GPT-4o rigorously flagged these terms as hallucinations, they are arguably reasonable product descriptions from a human perspective. On a less positive note, despite the existence of a new relation between "Stable Diffusion" and "Smoking," none of the LLMs generated ideas related to this relation.

Most LLMs focused their responses on the newly added relations, except for Grok-2. We believe this difference stems from variations in model size, training prompts, and datasets. Even when tested with the same prompts, each LLM produces qualitatively different answers due to differences in comprehensive capabilities. Nonetheless, our process consistently demonstrates creativity across most LLMs, highlighting its feasibility. In the future, we can explore prompt engineering tailored to each LLM's characteristics, which may help address similar problems.

4 DISCUSSION

How does this method mitigate hallucinations?

By constructing responses strictly based on the highly weighted relations derived from the structured KG, the LLMs are constrained to reasoning within established associations rather than generating unfounded information. This structured approach is similar to retrieval-augmented generation (RAG) (Lewis et al., 2020; Edge et al., 2024) with a summarization step, reinforcing grounded responses.

Does knowledge graph guidance impact creativity?

Yes, since the LLM is restricted to interpreting new relations rather than freely generating responses, the degree of creative freedom is reduced. However, if the KG contains a vast number of associations, the setting becomes closer to unrestricted generation. This work provides an interpretable framework for understanding LLM creativity and encourages further exploration in incorporating user preferences and refining the bridge-building process to further emulate human cognition.

5 CONCLUSION

This study leverages KGs to enhance LLM creativity while mitigating hallucinations through structured associative thinking. Results on 5 LLMs support this approach's feasibility. However, there are still limitations (Please see Appendix A) that require further exploration by future researchers.

URM STATEMENT

The authors acknowledge that at least one key contributor of this work satisfies the URM criteria for the ICLR 2025 Workshop QUESTION Tiny Papers Track.

REFERENCES

- Laha Ale, Ning Zhang, Scott A King, and Dajiang Chen. Empowering generative ai through mobile edge computing. *Nature Reviews Electrical Engineering*, pp. 1–9, 2024.
- Anthropic. The claude 3 model family: Opus, sonnet, haiku, 2024.
- Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, and Jonathan Larson. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*, 2024.
- Lorena RR Gianotti, Christine Mohr, Diego Pizzagalli, Dietrich Lehmann, and Peter Brugger. Associative processing and paranormal belief. *Psychiatry and clinical neurosciences*, 55(6): 595–603, 2001.
- Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, et al. A survey on llm-as-a-judge. *arXiv preprint arXiv:2411.15594*, 2024.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.
- Xuhui Jiang, Yuxing Tian, Fengrui Hua, Chengjin Xu, Yuanzhuo Wang, and Jian Guo. A survey on large language model hallucination via a creativity perspective. *arXiv preprint arXiv:2402.06647*, 2024.
- Xuhui Jianga, Chengjin Xua, Fengrui Huaa, Yuxing Tiana, Yuanzhuo Wangb, and Jian Guoa. Is creativity hidden in hallucination? rethink large language models.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33: 9459–9474, 2020.
- Gill Livingston, Jonathan Huntley, Kathy Y Liu, Sergi G Costafreda, Geir Selbæk, Suvarna Alladi, David Ames, Sube Banerjee, Alistair Burns, Carol Brayne, et al. Dementia prevention, intervention, and care: 2024 report of the lancet standing commission. *The Lancet*, 404(10452):572–628, 2024.
- Martin Marko, Drahomír Michalko, and Igor Riečanský. Remote associates test: An empirical proof of concept. *Behavior research methods*, 51(6):2700–2711, 2019.
- Sarnoff Mednick. The associative basis of the creative process. *Psychological review*, 69(3):220, 1962.
- Megha Sharma, Abhinav Tomar, and Abhishek Hazra. Edge computing for industry 5.0: fundamental, applications and research challenges. *IEEE Internet of Things Journal*, 2024.
- Fereshtehossadat Shojaei, Fatemehalsadat Shojaei, Erik Stolterman Bergvist, and Patrick C Shih. Exploring the impact of digital art therapy on people with dementia: a framework and research-based discussion. In *International Conference on Innovation of Emerging Information and Communication Technology*, pp. 133–143. Springer, 2023.

Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.

Donald J Treffinger. *Creativity, creative thinking, and critical thinking: In search of definitions*. Center for Creative Learning, 1995.

xAI. Grok-2 (version 2) [large language model], 2023. URL <https://x.ai/>.

Nan Xu and Xuezhe Ma. Llm the genius paradox: A linguistic and math expert’s struggle with simple word-based counting problems, 2024. URL <https://arxiv.org/abs/2410.14166>.

A LIMITATIONS AND FUTURE WORK

We summarize three different aspects of limitations to provide a basis for future research efforts:

- 1. Lack of datasets:** There is currently a shortage of datasets specifically designed to evaluate associative thinking, and constructing such a dataset itself presents a challenge.
- 2. Improving the evaluation mechanism:** This study relies solely on GPT-4o as the evaluator. In the future, a more objective and fair evaluation process should be established to enhance the reliability of the results.
- 3. Alignment of edge weights:** Language models may have limited numerical sensitivity when following prompts (Xu & Ma, 2024). Future work could explore alternative representations to help models better understand numerical weights, potentially improving alignment outcomes.

B PROMPTS

B.1 BUILDING THE KNOWLEDGE GRAPH

Convert [article] into a knowledge graph in JSON (.json) format using the following key structure:

"node": Represents the entities or concepts within the article.

"edge": Defines the relations between nodes and must include the following four keys:

"source node": The starting entity of the relation.

"target node": The ending entity of the relation.

"relation": The type or nature of the connection between the nodes.

"weight": A numerical value representing the strength or significance of a relation, with all weights defaulting to 1.

B.2 BUILDING THE NEW RELATION (ASSOCIATION)

Complete the following new relations based on the above [knowledge graphs]:

```
[  
  "source node": "Dementia", "target node": "DigitalArt", "relation": ", "weight": 1000,  
  "source node": "Edge Computing", "target node": "ArtTherapy", "relation": ", "weight": 1000,  
  "source node": "MedicalKnowledge", "target node": "IoT", "relation": ", "weight": 1000,  
  "source node": "StableDiffusion", "target node": "Smoking", "relation": ", "weight": 1000  
]
```

B.3 ANSWER QUESTION

With new integrated knowledge graph

Answer the [question] based on the [knowledge graph] and according to the edge weight (the larger the number, the greater the weight).

With only new relations

Answer the [question] based on the [new relations].

C THE RESPONSES GENERATED BY FIVE SOTA LLMs

Table 1: Responses from various models to the question on designing a product for dementia prevention.

Model	Answer
GPT-4o	<p>Proposed Product: A Holistic AI-Driven Dementia Prevention Tool</p> <p>This AI-powered digital art therapy platform combines:</p> <ul style="list-style-type: none"> - Edge Computing for real-time adaptation. - IoT & Medical Knowledge to provide personalized health insights. - Generative AI to create interactive, dynamic experiences. - Social Interaction & Gamification to reduce isolation and encourage prevention.
DeepSeek-R1	<p>Product Concept: "NeuroGuard"</p> <p>A hybrid system combining:</p> <ul style="list-style-type: none"> - Smart wearable for health monitoring (IoT/edge computing). - Tablet app for AI-driven art therapy, cognitive games, and social connectivity. - Home sensors (air quality, safety) linked to the platform. - Generative AI for personalized interventions (e.g., art, health tips, risk factor mitigation).
Claude-3.5-Sonnet	<p>Here's an integrated summary of the dementia prevention product: "ArtMind Guardian"</p> <p>Core Integration:</p> <ol style="list-style-type: none"> 1. Technology Stack <ul style="list-style-type: none"> - Edge Computing + IoT for real-time monitoring - Generative AI (Stable Diffusion) for art creation - Cloud infrastructure for data management - AI-powered health analytics 2. User Experience Flow <ul style="list-style-type: none"> - Input → Processing → Output - Health Data Collection → Analysis → Personalized Activities - Art Creation → Digital Processing → Therapeutic Feedback - Social Interaction → Community Building → Support Network - Cognitive Exercise → Performance Tracking → Progress Reports 3. Key Integration Points <ul style="list-style-type: none"> - Health Monitoring ↔ Art Therapy - Social Connection ↔ Creative Expression - Physical Activity ↔ Digital Interaction - Medical Knowledge ↔ User Engagement 4. Risk Management Integration <ul style="list-style-type: none"> - Prevention Factors: <ul style="list-style-type: none"> - Cognitive Decline ← Art Creation + Mental Exercise - Physical Inactivity ← Motion Controls + Exercise Games - Social Isolation ← Community Features - Mental Health ← Creative Expression + Social Support

Model	Answer
Grok-2	<p>Combining these insights, the innovative product could be a Smart Cognitive and Physical Wellness System for elderly individuals aimed at preventing dementia.</p> <p>Key features might include:</p> <p>a. Cognitive Stimulation through Digital Art Digital Art Therapy Platform: Utilize StableDiffusion and other Text-to-Image models to create personalized art therapy sessions that engage the elderly in creative activities. This could help enhance cognitive reserve and combat social isolation. Interactive Cognitive Games: Integrate cognitive games that adapt to the user's skill level, leveraging AI to personalize the experience and track progress.</p> <p>b. Physical Activity Monitoring and Encouragement Wearable Devices: Equipped with IoT sensors to monitor physical activity levels. These devices could remind users to stay active and suggest exercises tailored to their capabilities. Virtual Fitness Coach: Use AI to provide personalized fitness guidance and encouragement, addressing physical inactivity.</p> <p>c. Environmental and Health Monitoring Air Quality Sensors: Incorporate IoT devices to monitor indoor air quality, providing alerts and recommendations to improve air quality, directly addressing air pollution. Health Monitoring: Integrate sensors to track vital signs such as blood pressure and cholesterol levels, promoting awareness and management of hypertension and high LDL cholesterol.</p> <p>d. Social Connectivity Social Interaction Platform: Leverage Edge Computing to facilitate real-time social interactions, reducing social isolation. This could include virtual meetups, group activities, or art therapy sessions.</p> <p>e. Safety and Injury Prevention Fall Detection and Head Protection: Use Edge Computing and IoT sensors to detect falls and remind users to wear head protection, preventing traumatic brain injuries.</p>

Model	Answer
Gemini-1.5-Flash-128k	<p>Proposed Innovative Product: "Mindful Moments" - An AI-Powered Digital Art Therapy Platform with Integrated Health Monitoring</p> <p>Key Features:</p> <p>Personalized Digital Art Therapy: The platform offers a range of guided art therapy sessions using various digital art tools. These sessions are personalized based on the user's cognitive abilities, preferences, and progress. This addresses social isolation and promotes cognitive stimulation. The platform could utilize Generative AI (like Stable Diffusion or DALL-E 2) to create unique and engaging art prompts tailored to the user.</p> <p>Integrated Health Monitoring: Using wearable sensors or IoT-connected devices (smartwatches, scales, etc.), the platform can passively monitor vital signs (heart rate, blood pressure, activity levels). This allows for early detection of potential health issues (hypertension, physical inactivity) and prompts users to take action or alerts caregivers. This data is securely stored and analyzed using AI algorithms.</p> <p>AI-Powered Engagement and Progress Tracking: The AI analyzes user interaction with the art therapy sessions, monitoring engagement levels and progress. This data helps personalize the experience further and provides valuable insights for caregivers and healthcare professionals.</p> <p>Caregiver Support: The platform provides a secure communication channel for caregivers to monitor the user's progress, receive alerts, and share information with healthcare providers.</p> <p>Edge Computing for Privacy and Efficiency: Processing of sensitive user data (health and art therapy progress) happens on edge devices, ensuring privacy and reducing latency.</p>