

survey of hallucination in natural language generation

survey of hallucination in natural language generation presents a comprehensive examination of the phenomenon where AI-generated text contains fabricated or inaccurate information. Hallucination in natural language generation (NLG) systems poses significant challenges to the reliability, trustworthiness, and practical deployment of AI models. This survey explores the underlying causes, types, detection methodologies, and mitigation strategies related to hallucinations in NLG. It also delves into the impact of hallucination on various NLG applications, including summarization, dialogue systems, and machine translation. By analyzing recent advancements and research trends, this article provides a detailed understanding of how hallucinations arise and how they can be addressed. The discussion includes both technical approaches and evaluation metrics critical for assessing hallucination severity. The following sections outline the key areas covered in this comprehensive survey.

- Understanding Hallucination in Natural Language Generation
- Types and Causes of Hallucination
- Detection Techniques for Hallucination
- Mitigation Strategies and Model Improvements
- Applications and Implications of Hallucination in NLG

Understanding Hallucination in Natural Language Generation

Hallucination in natural language generation refers to the tendency of language models to produce outputs that are factually incorrect, nonsensical, or not grounded in the input data. This issue undermines the reliability of NLG systems, especially in high-stakes domains such as healthcare, legal documentation, and news generation. Understanding hallucination requires a clear distinction between errors caused by model limitations and those arising from ambiguous or incomplete data.

Modern NLG systems, particularly those based on deep learning and transformer architectures, generate text by predicting the most probable next token given the context. However, this probability-based method can lead to the generation of plausible-sounding yet false statements, a hallmark of hallucination. The challenge lies in balancing creativity and coherence with factual accuracy and relevance.

Definition and Characteristics

Hallucination can be defined as the generation of content by an NLG model that is unfaithful to the

source input or external knowledge. It often manifests as fabricated facts, unsupported claims, or inconsistent details. Characteristics of hallucination include:

- **Fabrication:** False information not supported by any input data.
- **Distortion:** Altered details that misrepresent source content.
- **Irrelevance:** Generation of unrelated or off-topic content.

Importance of Addressing Hallucination

Addressing hallucination is crucial for the adoption of NLG technologies in real-world applications. Hallucinated content can mislead users, reduce trust, and lead to erroneous decisions. Therefore, research focuses on identifying hallucinations early and developing techniques to minimize their occurrence without sacrificing linguistic quality or fluency.

Types and Causes of Hallucination

Hallucination in natural language generation arises from multiple factors, both intrinsic to the model architecture and extrinsic to the training data. Understanding the types and causes helps in designing better detection and mitigation strategies.

Intrinsic vs Extrinsic Hallucination

Two primary categories of hallucination are recognized:

- **Intrinsic Hallucination:** Occurs when the generated content contradicts or diverges from the provided input. This is common in tasks like text summarization, where the summary may introduce details not present in the source document.
- **Extrinsic Hallucination:** Happens when the model generates information that cannot be verified or found in any external knowledge source. This type is prevalent in open-domain dialogue systems and question answering.

Causes of Hallucination

Several factors contribute to hallucination in NLG systems, including:

1. **Training Data Quality:** Noisy, biased, or insufficiently diverse training data can lead to incorrect generalizations.
2. **Model Overconfidence:** Probability-based token prediction can overestimate the likelihood of

certain tokens, causing false information generation.

3. **Exposure Bias:** Discrepancy between training and inference phases, where models rely on ground-truth data in training but on self-generated tokens during inference.
4. **Knowledge Gaps:** Limited or outdated knowledge in pre-trained models can cause erroneous or fabricated outputs.
5. **Task Complexity:** Complex tasks with ambiguous or incomplete input often increase hallucination risks.

Detection Techniques for Hallucination

Detecting hallucination is a critical step toward improving the reliability of NLG systems. Various approaches have been developed to identify hallucinated content, ranging from rule-based methods to advanced neural techniques.

Automatic Metrics and Evaluation

Traditional evaluation metrics such as BLEU, ROUGE, and METEOR focus on lexical overlap and fluency but often fail to capture factual consistency. To address this, newer metrics have been introduced:

- **Fact-based Metrics:** Metrics like FactCC and QAGS evaluate factual correctness by comparing generated text against source data or external knowledge bases.
- **Semantic Similarity:** Embedding-based similarity measures assess whether the meaning of generated text aligns with the input content.
- **Consistency Scores:** Metrics that check for contradictions within generated outputs or between outputs and source data.

Human Evaluation

Human annotators remain essential for detecting nuanced hallucinations that automatic metrics may overlook. Evaluators assess factual accuracy, relevance, and coherence of generated text. Human evaluation protocols often involve:

1. Comparison of generated text against source documents.
2. Verification of facts using external resources.
3. Rating scales for hallucination severity and impact.

Machine Learning-based Detection

Recent progress includes training classifiers to identify hallucinated content directly. These models leverage labeled datasets to learn patterns indicative of hallucination. Techniques include:

- Supervised learning with annotated hallucination examples.
- Adversarial training to distinguish truthful from fabricated text.
- Cross-modal approaches that incorporate knowledge graphs or databases for verification.

Mitigation Strategies and Model Improvements

Reducing hallucination in natural language generation requires targeted interventions during model training, architecture design, and inference. This section reviews key mitigation strategies employed in current research.

Data-Centric Approaches

Improving the quality and diversity of training data is a foundational mitigation method. Techniques include:

- Curating high-quality, factually accurate datasets.
- Augmenting data with external knowledge sources.
- Filtering or correcting noisy annotations.

Model Architecture Enhancements

Architectural modifications help constrain models to generate more faithful content. Examples include:

- Incorporating retrieval-augmented generation to ground outputs in external documents.
- Using fact-aware attention mechanisms to focus on relevant source information.
- Employing constrained decoding strategies to avoid unlikely token sequences.

Training and Inference Techniques

Strategies at the training and inference stages include:

- **Reinforcement Learning:** Optimizing models with reward functions that penalize hallucination.
- **Knowledge Distillation:** Transferring knowledge from more accurate teacher models to student models.
- **Post-Editing:** Employing secondary models or algorithms to detect and correct hallucinated content after generation.
- **Controlled Generation:** Conditioning models on factual constraints or templates to guide output.

Applications and Implications of Hallucination in NLG

Hallucination affects a wide range of natural language generation applications, influencing their effectiveness and user trust. Understanding its implications is essential for deploying NLG systems responsibly.

Text Summarization

In abstractive summarization, hallucination is a prominent challenge as summaries may introduce unsupported details or distort original content. This can mislead readers and reduce the credibility of automated summarization tools.

Dialogue Systems and Chatbots

Conversational agents risk generating hallucinated responses that can confuse or misinform users. Maintaining factual consistency is critical, especially in domains like customer service, healthcare, and education.

Machine Translation

Although hallucination is less common in machine translation, it can occur when models invent or omit information, leading to inaccurate translations. This impacts communication and information exchange across languages.

Content Creation and Journalism

Automated content generation tools may inadvertently produce fabricated facts, raising ethical concerns and highlighting the need for robust hallucination prevention in media and publishing.

Implications for Trust and Ethics

The presence of hallucination affects user trust and the ethical use of AI-generated content. Transparency, explainability, and rigorous validation are necessary to mitigate risks associated with hallucinated outputs.

Frequently Asked Questions

What is hallucination in natural language generation (NLG)?

Hallucination in natural language generation refers to the phenomenon where a model generates content that is factually incorrect, irrelevant, or not supported by the source data, leading to outputs that contain fabricated or misleading information.

Why is hallucination a significant problem in NLG systems?

Hallucination undermines the reliability and trustworthiness of NLG systems, especially in critical applications like summarization, question answering, and dialogue systems, where factual accuracy is crucial for user trust and practical utility.

What are the common types of hallucinations identified in NLG?

Common types include intrinsic hallucination, where generated content contradicts the input data, and extrinsic hallucination, where the content is unsupported but not necessarily contradictory to the input.

What methods are commonly used to detect hallucinations in NLG outputs?

Detection methods include automated fact-checking models, natural language inference (NLI) techniques, human evaluation, and metrics that compare generated outputs against source documents or external knowledge bases.

How do datasets contribute to the study of hallucination in NLG?

Datasets annotated for factual consistency and hallucination help train and evaluate models designed to reduce hallucination, enabling researchers to benchmark performance and develop better hallucination detection and mitigation techniques.

What are some strategies to mitigate hallucination in NLG models?

Strategies include incorporating factual constraints during training, using retrieval-augmented generation to ground outputs in external knowledge, fine-tuning on high-quality datasets, and employing post-generation verification and correction mechanisms.

How do recent transformer-based models address hallucination compared to earlier approaches?

Recent transformer-based models leverage large-scale pretraining and fine-tuning with factuality-focused objectives, often combined with retrieval-based methods, to reduce hallucination, though challenges remain, especially with out-of-domain or ambiguous inputs.

What future directions are suggested for research on hallucination in NLG?

Future research directions include developing more robust hallucination detection metrics, integrating multimodal and external knowledge sources, improving interpretability of hallucination causes, and creating benchmarks that better reflect real-world use cases.

Additional Resources

1. *Hallucinations in Natural Language Generation: A Comprehensive Survey*

This book provides an extensive overview of hallucination phenomena in natural language generation (NLG). It covers the underlying causes, types, and impacts of hallucinated content in generated texts. The authors review state-of-the-art detection and mitigation techniques, offering practical insights for researchers and practitioners to improve model reliability.

2. *Understanding and Controlling Hallucinations in Text Generation*

Focusing on the technical aspects, this book delves into the mechanisms that lead to hallucinations in neural language models. It explains the role of training data, model architecture, and decoding strategies in hallucination occurrences. The text also presents various approaches to control and reduce hallucinated outputs, including reinforcement learning and adversarial training.

3. *Natural Language Generation: Challenges of Hallucinated Outputs*

This work highlights the challenges posed by hallucinations in NLG applications such as summarization, translation, and dialogue systems. It discusses how hallucinations affect user trust and downstream tasks. Case studies and evaluation metrics are provided to help readers assess hallucination severity and model performance.

4. *Mitigating Hallucinations in Neural Text Generation*

Targeted at advanced practitioners, this book explores cutting-edge methodologies for mitigating hallucinated content in neural networks. It covers techniques like data augmentation, confidence calibration, and post-processing filters. Experimental results demonstrate the effectiveness of these methods across different NLG tasks.

5. *Hallucination Detection in Language Models: Methods and Applications*

This book surveys various methods for detecting hallucinations in generated text, including rule-based, statistical, and machine learning approaches. It addresses the challenges in creating reliable detection systems and discusses their practical applications in industry settings. The text also explores the integration of detection tools into production pipelines.

6. Ethical Implications of Hallucinations in AI-Generated Text

Examining the ethical dimensions, this book discusses the consequences of hallucinated information in AI-generated content. It covers misinformation, bias amplification, and user deception risks. The authors propose guidelines and best practices to promote responsible use and development of NLG systems.

7. Survey of Hallucination Phenomena in Multimodal Language Generation

This title extends the discussion of hallucinations to multimodal models that combine text with images, audio, or video. It explores how hallucinations manifest differently across modalities and the unique challenges they present. The book reviews current research efforts aimed at reducing multimodal hallucinations.

8. Evaluating Hallucinations in Natural Language Generation: Metrics and Benchmarks

Providing a deep dive into evaluation, this book catalogs various metrics designed to quantify hallucination in NLG outputs. It compares intrinsic and extrinsic evaluation methods and highlights benchmark datasets used in the community. The book serves as a guide for researchers aiming to standardize hallucination assessment.

9. Future Directions in Hallucination Research for Natural Language Generation

Looking ahead, this book discusses emerging trends and open problems in hallucination research within NLG. It covers advancements in model architectures, interpretability, and user-centered evaluation. The authors outline promising research avenues and encourage interdisciplinary collaboration to tackle hallucination challenges.

Survey Of Hallucination In Natural Language Generation

Find other PDF articles:

<http://www.devensbusiness.com/archive-library-607/files?ID=idB93-2201&title=pre-service-education-definition.pdf>

survey of hallucination in natural language generation: Health Information Processing. Evaluation Track Papers Yanchun Zhang, Qingcai Chen, Hongfei Lin, Lei Liu, Xiangwen Liao, Buzhou Tang, Tianyong Hao, Zhengxing Huang, Jianbo Lei, Zuofeng Li, Hui Zong, 2025-05-14 This book constitutes the refereed proceedings of the 10th China Health Information Processing Conference, CHIP 2024, held in Fuzhou, China, November 15-17, 2024. The CHIP 2024 Evaluation Track proceedings include 19 full papers which were carefully reviewed and grouped into these topical sections: syndrome differentiation thought in Traditional Chinese Medicine; lymphoma information extraction and automatic coding; and typical case diagnosis consistency.

survey of hallucination in natural language generation: Natural Language Processing and Information Systems Ryutaro Ichise, 2025-08-01 The two-volume set LNCS 15836 and 15837 constitutes the proceedings of the 30th International Conference on Applications of Natural

Language to Information Systems, NLDB 2025, held in Kanazawa, Japan, during July 4-6, 2025. The 33 full papers, 19 short papers and 2 demo papers presented in this volume were carefully reviewed and selected from 120 submissions. The proceedings contain novel and significant research contributions addressing theoretical aspects, algorithms, applications, architectures, resources, and other aspects of NLP, as well as survey and discussion papers.

survey of hallucination in natural language generation: Advances in Information Retrieval Nazli Goharian, Nicola Tonellotto, Yulan He, Aldo Lipani, Graham McDonald, Craig Macdonald, Iadh Ounis, 2024-03-19 The six-volume set LNCS 14608, 14609, 14609, 14610, 14611, 14612 and 14613 constitutes the refereed proceedings of the 46th European Conference on IR Research, ECIR 2024, held in Glasgow, UK, during March 24-28, 2024. The 57 full papers, 18 finding papers, 36 short papers, 26 IR4Good papers, 18 demonstration papers, 9 reproducibility papers, 8 doctoral consortium papers, and 15 invited CLEF papers were carefully reviewed and selected from 578 submissions. The accepted papers cover the state of the art in information retrieval focusing on user aspects, system and foundational aspects, machine learning, applications, evaluation, new social and technical challenges, and other topics of direct or indirect relevance to search.

survey of hallucination in natural language generation: Humans and Generative AI Tools for Collaborative Intelligence Zhao, Jingyuan, Kumar, V. Vinoth, Katina, Polinapilinho F., Richards, Joseph, 2025-06-11 The relationship between humans and generative AI tools creates a new era of collaborative intelligence, where human creativity, critical thinking, and empathy merge with the speed, scale, and precision of artificial intelligence. Rather than replacing human roles, generative AI systems help with problem-solving, ideation, and decision-making across diverse fields such as education, healthcare, design, and business. This enables individuals and organizations to reveal deeper insights, accelerate innovation, and tackle complex challenges with enhanced efficiency. As these technologies advance, the focus is shifting toward fostering ethical, inclusive, and human-centered collaborations that amplify collective intelligence rather than diminish it. Humans and Generative AI Tools for Collaborative Intelligence explores recent advancements in collaborative intelligence of humans and generative AI tools, with a focus on state-of-the-art approaches, methodologies, and systems for the design, development, deployment, and innovative use of those technologies and applications to advance industries and organizations. It develops guidelines to help industries and organizations achieve the collaborative intelligence of humans and generative AI tools and put the power of collaborative intelligence to work. This book covers topics such as digital technology, robotics, and deep learning models, and is a useful resource for business owners, computer engineers, academicians, researchers, and data scientists.

survey of hallucination in natural language generation: Databases Theory and Applications Tong Chen, Yang Cao, Quoc Viet Hung Nguyen, Thanh Tam Nguyen, 2024-12-12 This LNCS volume constitutes the referred proceedings of 35th Australasian Database Conference, ADC 2024, held in Gold Coast, QLD, Australia and Tokyo, Japan, during December 16-18, 2024. The 38 full papers included in these proceedings were carefully reviewed and selected from 90 submissions. They focus on latest advancements and innovative applications in database systems, data-driven applications, and data analytics.

survey of hallucination in natural language generation: Artificial Intelligence in HCI Helmut Degen, Stavroula Ntoa, 2023-07-08 This double volume book set constitutes the refereed proceedings of 4th International Conference, AI-HCI 2023, held as part of the 25th International Conference, HCI International 2023, which was held virtually in Copenhagen, Denmark in July 2023. The total of 1578 papers and 396 posters included in the HCII 2023 proceedings was carefully reviewed and selected from 7472 submissions. The first volume focuses on topics related to Human-Centered Artificial Intelligence, explainability, transparency and trustworthiness, ethics and fairness, as well as AI-supported user experience design. The second volume focuses on topics related to AI for language, text, and speech-related tasks, human-AI collaboration, AI for decision-support and perception analysis, and innovations in AI-enabled systems.

survey of hallucination in natural language generation: Human Brain and Artificial Intelligence Quanying Liu, Youzhi Qu, Haiyan Wu, Yu Qi, An Zeng, Dan Pan, 2025-04-19 This book constitutes the refereed proceedings of the 4th International Workshop on Human Brain and Artificial Intelligence, HBAI 2024, held in Jeju Island, South Korea, on August 3, 2024. The 24 full papers and 5 short papers included in this book were carefully reviewed and selected from 74 submissions. They were organized in topical sections as follows: AI for brain science; AI for brain technology; and brain-inspired AI.

survey of hallucination in natural language generation: The Semantic Web - ISWC 2023 Terry R. Payne, Valentina Presutti, Guilin Qi, María Poveda-Villalón, Giorgos Stoilos, Laura Hollink, Zoi Kaoudi, Gong Cheng, Juanzi Li, 2023-11-01 This book constitutes the proceedings of the 22nd International Semantic Web Conference, ISWC 2023, which took place in October 2023 in Athens, Greece. The 58 full papers presented in this double volume were thoroughly reviewed and selected from 248 submissions. Many submissions focused on the use of reasoning and query answering, with a number addressing engineering, maintenance, and alignment tasks for ontologies. Likewise, there has been a healthy batch of submissions on search, query, integration, and the analysis of knowledge. Finally, following the growing interest in neuro-symbolic approaches, there has been a rise in the number of studies that focus on the use of Large Language Models and Deep Learning techniques such as Graph Neural Networks.

survey of hallucination in natural language generation: Web and Big Data Wenjie Zhang, Anthony Tung, Zhonglong Zheng, Zhengyi Yang, Xiaoyang Wang, Hongjie Guo, 2024-08-27 The five-volume set LNCS 14961, 14962, 14963, 14964 and 14965 constitutes the refereed proceedings of the 8th International Joint Conference on Web and Big Data, APWeb-WAIM 2024, held in Jinhua, China, during August 30–September 1, 2024. The 171 full papers presented in these proceedings were carefully reviewed and selected from 558 submissions. The papers are organized in the following topical sections: Part I: Natural language processing, Generative AI and LLM, Computer Vision and Recommender System. Part II: Recommender System, Knowledge Graph and Spatial and Temporal Data. Part III: Spatial and Temporal Data, Graph Neural Network, Graph Mining and Database System and Query Optimization. Part IV: Database System and Query Optimization, Federated and Privacy-Preserving Learning, Network, Blockchain and Edge computing, Anomaly Detection and Security Part V: Anomaly Detection and Security, Information Retrieval, Machine Learning, Demonstration Paper and Industry Paper.

survey of hallucination in natural language generation: Minds, Machines, and Misinformation Don Donghee Shin, 2025-08-11 Algorithms have become the key organizer through which power is enacted in our society. A huge amount of data regarding our daily routines are monitored and analyzed to make recommendations that manage, control, and lead our behaviors in everyday life. AI, Humans, and Misinformation: How Does AI Alter Human Behavior and How Do Humans Influence Algorithmic Misinformation? is a guide to understanding the dynamics of AI and misinformation in human contexts by addressing meaningful questions—How does AI alter human behavior and how do humans influence algorithmic decision-making? In answering these questions, this book examines the role of misinformation, disinformation, and fake news, and shows readers how to develop AI methods and algorithms that combat misinformation by using AI design choices that provide users and developers alike with meaningful control over AI. This book brings together various perspectives on algorithms into an integrated conceptual framework, and provides a broad socio-technical analysis, addressing critical and ethical issues of misinformation and fake news. The book offers a compelling insight into the misinformation phenomenon and the future of AI-based society. Readers will find an integrated technical analysis of the logic and social implications of algorithmic processes. Reporting from the cutting edge of critical technical methods and research, the result is useful and constructive for developing the relations between algorithms and humans. This is an imperative methodology for understanding what is at stake as industry and government use AI to reshape the world. - Provides a comprehensive examination of the technical foundations and social effects of algorithmic processes and how they are used to purvey and combat

misinformation - Includes deep technical models for AI and ML algorithm developers who seek to alleviate the impacts of misinformation, including credibility assessment models, trust heuristics, and Information Processing Theory - Provides detailed technical understanding of computational approaches to combat misinformation such as automated deception detectors, clickbait detectors, satirical fake detectors, rumor debunkers, and computational fact-checking tools - Focuses on the human aspects of AI and misinformation, and the multidisciplinary study of trends, problems, and algorithmic functions

survey of hallucination in natural language generation: *MultiMedia Modeling* Stevan Rudinac, Alan Hanjalic, Cynthia Liem, Marcel Worring, Björn Þór Jónsson, Bei Liu, Yoko Yamakata, 2024-01-27 This book constitutes the refereed proceedings of the 30th International Conference on MultiMedia Modeling, MMM 2024, held in Amsterdam, The Netherlands, during January 29-February 2, 2024. The 112 full papers included in this volume were carefully reviewed and selected from 297 submissions. The MMM conference were organized in topics related to multimedia modelling, particularly: audio, image, video processing, coding and compression; multimodal analysis for retrieval applications, and multimedia fusion methods.

survey of hallucination in natural language generation: *Proceedings of 17th International Conference on Machine Learning and Computing* Lin Huang, David Greenhalgh, 2025-08-17 This book comprises original and peer reviewed research papers presented at 2025 17th International Conference on Machine Learning and Computing that was held in Guangzhou, China, from February 14 to 17, 2025. The focus of the conference is to establish an effective platform for institutions and industries to share ideas and to present the works of scientists, engineers, educators and students from all over the world. Topics discussed in this volume include Machine Learning Theory and Algorithms, High-performance Computing Models and Data Processing, Large-scale Language Models and Natural Language Processing, Data-oriented Information System Optimization and Intelligent Computing, AI-based Intelligent Control Systems and System Security, etc. The book will become a valuable resource for academics, industry professionals, and engineers working in the related fields of machine learning and computing.

survey of hallucination in natural language generation: *Data Science and Information Security* Hai Jin, Yi Pan, Jianfeng Lu, 2024-04-02 This 3-volume set, CCIS 2058-2060 constitutes the First International Conference, on Artificial Intelligence, IAIC 2023, held in Nanjing, China, in November 2023. The 85 full papers presented were carefully reviewed and selected from 428 submissions. The papers are clustered in parts on: Artificial Intelligence and Machine Learning; Data Security and information Security; Computer Networks and IoT. The papers present recent research and developments in artificial intelligence and its applications in machine learning, natural language processing, computer vision, robotics, and ethical considerations.

survey of hallucination in natural language generation: *Advances in Information Retrieval* Claudia Hauff, Craig Macdonald, Dietmar Jannach, Gabriella Kazai, Franco Maria Nardini, Fabio Pinelli, Fabrizio Silvestri, Nicola Tonellotto, 2025-04-02 The five-volume set LNCS 15572, 15573, 15574, 15575 and 15576 constitutes the refereed conference proceedings of the 47th European Conference on Information Retrieval, ECIR 2025, held in Lucca, Italy, during April 6-10, 2025. The 52 full papers, 11 findings, 42 short papers and 76 papers of other types presented in these proceedings were carefully reviewed and selected from 530 submissions. The accepted papers cover the state-of-the-art in information retrieval and recommender systems: user aspects, system and foundational aspects, artificial intelligence and machine learning, applications, evaluation, new social and technical challenges, and other topics of direct or indirect relevance to search and recommendation.

survey of hallucination in natural language generation: *Computational Science - ICCS 2025 Workshops* Maciej Paszynski, Amanda S. Barnard, Yongjie Jessica Zhang, 2025-07-04 The 6-volume set constitutes the workshop proceedings of the 25th International Conference on Computational Science, ICCS 2025, which took place in Singapore, Singapore, during July 7-9, 2025. The 137 full papers and 32 short papers presented in these proceedings were carefully reviewed and

selected from 322 submissions. The papers are organized in the following topical sections: Volume I: Advances in high-performance computational earth sciences: numerical methods, frameworks & applications; artificial intelligence approaches for network analysis; artificial intelligence and high-performance computing for advanced simulations; and biomedical and bioinformatics challenges for computer science. Volume II: Computational health; computational modeling and artificial intelligence for social systems; and computational optimization, modelling and simulation. Volume III: Computational science and AI for addressing complex and dynamic societal challenges equitably; computer graphics, image processing and artificial intelligence; computing and data science for materials discovery and design; and large language models and intelligent decision-making within the digital economy. Volume IV: Machine learning and data assimilation for dynamical systems; and multi-criteria decision-making: methods, applications, and innovations. Volume V: (Credible) Multiscale modelling and simulation; numerical algorithms and computer arithmetic for computational science; quantum computing; retrieval-augmented generation; and simulations of flow and transport: modeling, algorithms and computation. Volume VI: Smart systems: bringing together computer vision, sensor networks and artificial intelligence; solving problems with uncertainty; and teaching computational science.

survey of hallucination in natural language generation: The Semantic Web: ESWC 2024 Satellite Events Albert Meroño Peñuela, Oscar Corcho, Paul Groth, Elena Simperl, Valentina Tamma, Andrea Giovanni Nuzzolese, Maria Poveda-Villalón, Marta Sabou, Valentina Presutti, Irene Celino, Artem Revenko, Joe Raad, Bruno Sartini, Pasquale Lisena, 2025-01-27 This two volume set constitutes the refereed proceedings of the International Conference, ESWC 2024 Satellite Events, held in Hersonissos, Crete, Greece during May 26-30, 2024. The 67 papers presented were carefully reviewed and selected from 128 submissions. This year conference aimed at acknowledging recent developments in AI with a special tagline, “Fabrics of Knowledge: Knowledge Graphs and Generative AI”. To reflect this year’s special topic, the satellite events of ESWC 2024 featured a Special Track on Large Language Models for Knowledge Engineering, in addition to the poster and demo session, the PhD symposium, the industry track, project networking, and workshops and tutorials.

survey of hallucination in natural language generation: Computer Aided Systems Theory - EUROCAST 2024 Alexis Quesada-Arencibia, Michael Affenzeller, Roberto Moreno-Díaz, 2025-04-23 This three part LNCS volumes constitutes the refereed proceedings of the 19th International Conference on Computer-Aided Systems Theory, EUROCAST 2024, held in Las Palmas de Gran Canaria, Spain, during February 25 to March 1, 2024. The 104 full papers included in this book were carefully reviewed and selected from 150 submissions. They were organized in topical sections as follows : Part I : Systems Theory, Applications, Pioneers, and Landmarks; Theory and Applications of Metaheuristic Algorithms; Mechatronic Product Development; and Model-Based System Design, Verification and Simulation. Part II : Applications of Signal Processing Technology; Applied Data Science and Engineering for Intelligent Transportation Systems and Smart Mobility; Computer and Systems Based Methods and Electronic Tools in Clinical and Academic Medicine ; Systems in Industrial Robotics, Automation and IoT; Systems Thinking: Applications in Technology, Science and Management; and Data Science in Medical and Bio-Informatics. Part III : Modeling, Simulation, and Optimization in Production and Logistics; Green AI and SW-Tools for Sustainable Energy and Materials Consumption; Stochastic Models, Statistical Methods, and Applied Systems Simulations; and Systems Cybersecurity Technologies and Quantum Approaches Potentials.

survey of hallucination in natural language generation: Advances in Artificial Intelligence - IBERAMIA 2024 Luís Correia, Aiala Rosá, Francisco Garijo, 2025-02-24 This volume LNCS 15277 constitutes the refereed proceedings of the 18th Ibero-American Conference on AI, IBERAMIA 2024, held in Montevideo, Uruguay, during November 13-15, 2024. The 36 full papers presented together with 14 short papers were carefully reviewed and selected from 96 submissions. The conference focuses on AI engineering & applications; bio-inspired & soft computing; computer vision & robotics; knowledge representation & reasoning; machine learning; multi agent systems; natural language processing; and social AI.

survey of hallucination in natural language generation: *Proceedings of the 15th International Conference on Ubiquitous Computing & Ambient Intelligence (UCAmI 2023)* José Bravo, Gabriel Urzáiz, 2023-11-25 This book serves as a comprehensive compilation of groundbreaking research endeavors within the realms of ambient intelligence and ubiquitous computing. These initiatives are pivotal in enabling both researchers and practitioners to discern recent breakthroughs and emerging frontiers in these fields. Encompassing a wide array of domains, including Ambient Active and Assisted Living (A3L), the Internet of Things (IoT), Smart Environments, Data Science, and Human-Ambient Interaction, acts as a valuable resource for scholars, professionals, and graduate students alike. The primary aim of this book is to empower individuals within the academic and professional community to harness this wealth of knowledge. It equips them to tackle innovative challenges and engineer smart and ubiquitous solutions that will shape the landscape of the next decade. By amalgamating insights from various facets of ambient intelligence and ubiquitous computing, this book encourages cross-disciplinary collaboration and fosters a holistic understanding of the field. Thus, it not only highlights the recent strides in these areas but also serves as a roadmap for future exploration and innovation, paving the way for a smarter and more interconnected world.

survey of hallucination in natural language generation: *Information and Communication Technology* Wray Buntine, Morten Fjeld, Truyen Tran, Minh-Triet Tran, Binh Huynh Thi Thanh, Takumi Miyoshi, 2025-05-29 This four-volume set, CCIS 2350-2353, constitutes the referred proceedings of the 13th International Symposium on Information and Communication Technology, SOICT 2024, held in Danang, Vietnam in December 2024. The 88 full papers and 68 poster papers presented here were carefully reviewed and selected from 229 submissions. The papers presented in these volumes are organized in the following topical sections: Part I: Multimedia Processing; Operations Research. Part II: AI Applications; Cyber Security. Part III: AI Foundations and Big Data; Human-Computer Interaction. Part IV: Lifelog and Multimedia Retrieval; Generative AI; Software Engineering.

Related to survey of hallucination in natural language generation

Create a survey - Google Surveys Help Where will my survey questions appear? Questions appear throughout sites in our publisher network in order to get a representative sample of respondents. Users complete survey

Create your first form in Google Forms When someone takes your survey, they will be required to enter their email address before they submit the form. Collect verified emails Important: Respondents must confirm their Google

Create a survey - Google Surveys Help Where will my survey questions appear? Questions appear throughout sites in our publisher network in order to get a representative sample of respondents. Users complete survey

Quick Start Guide - Google Surveys Help How to set up screening questions Select the checkbox for each answer that qualifies a respondent for this audience. Having three or more answers helps eliminate

Crear una encuesta - Ayuda de Surveys Cuando Google Surveys recoge respuestas de la "audiencia general de Internet", utiliza conjuntos de datos de población de Internet publicados para realizar la distribución de la

Google Surveys Sunset - Google Surveys Help Historical survey results downloads are no longer available. We began Surveys over 10 years ago to enable businesses of all sizes to run custom market research with an

Device Usage Study Help - Google Help Official Device Usage Study Help Help Center where you can find tips and tutorials on using Device Usage Study Help and other answers to frequently asked questions

Google Surveys Help Official Google Surveys Help Center where you can find tips and tutorials on using Google Surveys and other answers to frequently asked questions

View and export results - Google Surveys Help To view your survey results: Sign in to Google Surveys. Click the survey you want to view on the survey dashboard. Click the text of any question to see individual question results. Keep in

Earn rewards - Opinion Rewards Help - Google Help With Google Opinion Rewards, you'll take surveys that are run by market researchers. Survey frequency may vary, and you don't have to answer every survey you receive. In exchange,

Related to survey of hallucination in natural language generation

Addressing AI hallucinations with retrieval-augmented generation (InfoWorld1y) The hallucinations of large language models are mainly a result of deficiencies in the dataset and training. These can be mitigated with retrieval-augmented generation and real-time data. Artificial

Addressing AI hallucinations with retrieval-augmented generation (InfoWorld1y) The hallucinations of large language models are mainly a result of deficiencies in the dataset and training. These can be mitigated with retrieval-augmented generation and real-time data. Artificial

Back to Home: <http://www.devensbusiness.com>