

Disinformation, Misinformation and Learning in the Age of Generative AI: Joint Proceedings of DISMISS-FAKE'25 and IWILDS'25

Koustav Rudra^{1,*}, Niloy Ganguly¹, Jeanne Mifsud Bonnici², Eric Müller-Budack^{3,*}, Ritumbra Manuvie², Anett Hoppe^{4,*}, Ran Yu⁵, Jiqun Liu⁶ and Nilavra Bhattacharya⁷

¹Indian Institute of Technology Kharagpur, Kharagpur, India

²University of Groningen, Groningen, The Netherlands

³TIB – Leibniz Information Centre for Science and Technology, Hannover, Germany

⁴Philipps University Marburg & Hessian Center for AI (hessian.AI), Marburg, Germany

⁵GESIS – Leibniz Institute for the Social Sciences, Cologne, Germany

⁶University of Oklahoma, Norman, OK, United States

⁷ABB Corporate Research Centre, Mannheim, Germany

1. Introduction

The 1st International Workshop on Disinformation and Misinformation in the Age of Generative AI (DISMISS-FAKE'25) and the 4th International Workshop on Investigating Learning during Web Search were held in conjunction with the 18th International ACM WSDM Conference on Web Search and Data Mining (WSDM 2025) on March 14 in Hannover, Germany. The following proceedings volume contains the accepted contributions of both workshops.

In an era where artificial intelligence and large language models are fundamentally transforming how information is created, distributed, and consumed, the challenge of navigating the digital information landscape has never been more complex—or more critical. These proceedings bring together insights from two complementary workshops that, while addressing different aspects of web-based information processing, share a fundamental concern: how humans interact with, evaluate, and integrate information in our increasingly AI-mediated world. The DISMISS-FAKE workshop addresses the pressing challenge of disinformation and misinformation detection in the era of generative AI. As these technologies make it easier than ever to create convincing yet harmful content, the workshop addresses the critical need for advanced detection methods, trustworthy AI systems, and policy interventions. The IWILDS workshop, focusing on learning during web search, explores how individuals acquire and extend knowledge through online information seeking, examining the cognitive processes involved in navigating, evaluating, and synthesizing information from diverse web sources.

At their convergence lies a shared recognition that the judgment of information quality and the integration of new information with existing knowledge and worldviews represent one of the defining challenges of our time. Both workshops acknowledge that in a world where information abundance meets sophisticated AI-generated content, success depends not merely on accessing information, but on developing the critical faculties to discern, evaluate, and meaningfully incorporate it into our understanding.

Disinformation, Misinformation and Learning in the Age of Generative AI: Joint Proceedings of the 1st International Workshop on Disinformation and Misinformation in the Age of Generative AI (DISMISS-FAKE'25) and the 4th International Workshop on Investigating Learning during Web Search (IWILDS'25) co-located with 18th International ACM WSDM Conference on Web Search and Data Mining (WSDM 2025)

*Corresponding authors.

✉ krudra@ai.iitkgp.ac.in (K. Rudra); niloy@cse.iitkgp.ac.in (N. Ganguly); g.p.mifsud.bonnici@step-rug.nl (J. M. Bonnici); Eric.Mueller@tib.eu (E. Müller-Budack); r.manuvie@rug.nl (R. Manuvie); anett.hoppe@tib.eu (A. Hoppe); ran.yu@uni-bonn.de (R. Yu); jiqunliu@ou.edu (J. Liu); nilavra.bhattacharya@de.abb.com (N. Bhattacharya)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

2. DISMISS-FAKE'25 Workshop

2.1. Workshop Overview & Scope

The International Workshop on Detecting and Mitigating Misinformation with Small and Large Language Models (DISMISS-FAKE) was established as a dedicated venue for examining how AI technologies are reshaping the landscape of misinformation and disinformation. With generative AI capable of both amplifying false narratives and supporting detection, the workshop positioned itself at the crossroads of natural language processing, information retrieval, social computing, and media studies. Its central aim was to foster interdisciplinary dialogue on how to build systems that are not only effective but also transparent, fair, and accountable.

The inaugural edition of DISMISS-FAKE'25 was held at ACM WSDM in Hannover, Germany, on March 14, 2025. Reflecting concerns raised on generative AI and disinformation, the workshop highlighted the dual role of large language models (LLMs) as both potential detectors of misinformation and prolific generators of deceptive content. Research contributions covered themes such as multilingual fact-checking datasets, hallucination detection in model outputs, narrative-driven misinformation, and news values analysis, underscoring the technical, cultural, and regulatory dimensions of the problem.

Adopting a highly interactive format, the workshop combined keynote presentations, paper sessions, and open discussions. Central questions animated the sessions: How reliable can LLMs be when they themselves generate convincing falsehoods? What role can smaller, resource-efficient models play in scalable detection? How can cross-lingual and multimodal approaches address misinformation in underrepresented contexts? And what safeguards, policies, and frameworks are needed to ensure responsible deployment? By bringing together perspectives from across disciplines, DISMISS-FAKE'25 provided both a snapshot of current advances and a platform for shaping future directions in countering disinformation in the era of generative AI.

2.2. Key Research Contributions

The event brought together keynote talks, paper presentations, and interactive sessions. *Prof. Krishna Gummadi* opened with a systems-level perspective on how generative AI shapes and amplifies misinformation flows, followed by short papers examining LLMs' role in generating and detecting fake content, categorizing misinformation on Reddit, and analyzing editorial news values. Later in the morning, *Prof. Ritumbra Manuvie* addressed the governance challenges of the EU's Digital Services Act, which frames platform responsibility in moderating online content. This was followed by a long paper on MMTweets, a multilingual dataset for cross-lingual fact-checking, and a short paper on instruction-tuned small models for hallucination detection. Afternoon sessions featured *Prof. Huan Liu* on multimodal social media mining and *Prof. Preslav Nakov* on ensuring factuality in LLMs. The day ended with a high-level panel of experts, encouraging debate across technical, legal, and ethical domains, and closing discussions highlighted the urgent need for interdisciplinary collaboration.

Iknoor Singh, Carolina Scarton, Xingyi Song, and Kalina Bontcheva presented MMTweets, a dataset designed to support cross-lingual fact-checking by linking tweets in multiple languages with verified fact-checks. *Pavan Sanjay Nichani, Ayaan Ahmad Siddiqui, Sakshi Tiwarii, Ark Ikhu, and Marina Ernst* investigated the paradox of whether large language models can recognize disinformation they themselves produce. *Bhavana Ramesh, Durwankur Gursale, Abram Jopaul, and Marina Ernst* explored how LLMs can classify different types of misinformation on Reddit, going beyond binary "true/false" judgments. *Elijah Soba, Harika Abburi, Nirmala Pudota, Jain Aayush, Balaji Veeramani, Edward Bowen, and Sanmitra Bhattacharya* presented research on using instruction-tuned, quantized small models for hallucination detection. *Gullal S. Cheema, Massiollah Azimi, Ralph Ewerth, and Eric Müller-Budack* presented exploratory work on using LLMs to analyze news values, the criteria editors use to decide whether an event is newsworthy.

2.3. Discussion Themes

The interactive format of DISMISS-FAKE'25 encouraged participants to build on ideas raised in the keynotes and paper presentations, leading to lively debates and shared reflections. Four major themes stood out during the day's exchanges: (i). **Multilinguality and Multimodality:** A recurring theme throughout the workshop was the challenge of addressing misinformation that transcends language and media formats, (ii). **Narrative and Cultural Dimensions of Fake Content:** Several exchanges turned to how misinformation is not just factual distortion but often embedded in cultural narratives, humour, and community norms, (iii). **Technical Countermeasures and Guardrails:** The technical challenge of building safeguards against generative misinformation came through strongly, (iv). **Legal and Policy Frameworks:** Participants linked DSA requirements, such as risk assessment, transparency, and accountability, to the practical challenges of building explainable and auditable AI systems.

2.4. Future Directions

DISMISS-FAKE'25 directly addressed the pressing challenges posed by generative AI to the domains of misinformation and disinformation detection. The workshop's discussion-driven format proved especially valuable, allowing participants to move beyond technical presentations toward broader reflections on the interplay between language models, governance, and societal resilience.

A clear consensus emerged: while LLMs and SLMs offer powerful opportunities for advancing detection, they also introduce new vulnerabilities, whether through their capacity to generate convincing disinformation, their biases in classification, or their tendency to hallucinate. This duality underscores the need for stronger guardrails, robust multilingual and multimodal datasets, and interdisciplinary frameworks that integrate technical innovation with legal, ethical, and cultural perspectives.

The sessions highlighted several urgent priorities that will shape future work. These include advancing cross-lingual retrieval to better serve low-resource languages, refining lightweight but accurate SLMs for scalable deployment, and embedding transparency and interpretability into detection pipelines, and pushing the boundaries of adversarial robustness in fact verification systems. Equally critical is aligning technical progress with policy frameworks such as the Digital Services Act and AI Act, ensuring that detection systems remain accountable, privacy-preserving, and inclusive across cultural contexts.

3. IWILDS'25 Workshop

3.1. Workshop Overview & Scope

The International Workshop on Investigating Learning During Web Search (IWILDS) has served as a vital forum for interdisciplinary research on web-based learning since 2019. The fifth iteration, IWILDS'25, took place at WSDM 2025 in Hannover, Germany, bringing together researchers from information retrieval, information management, human-computer interaction and related fields.

This edition placed particular emphasis on understanding how Large Language Models (LLMs) and AI technologies are transforming web-based learning. As users increasingly turn to LLM interfaces and AI-enhanced search engines for knowledge acquisition, the field faces unprecedented questions about how learning occurs in these new environments. The workshop's discussion-focused format enabled seamless transitions between formal presentations and spontaneous discussions, proving valuable for addressing whether traditional "Search as Learning" paradigms adequately capture current and emerging information-seeking behaviors.

3.2. Key Research Contributions

The workshop featured three main presentations examining different aspects of AI's impact on web-based learning:

Video Features for Predicting Knowledge Gain. Wolfgang Bitter and colleagues from TIB Hannover investigated how video interactions during web search affect learning outcomes. Using data from 94 study participants, the research revealed that video interaction features — particularly interaction frequency — are the strongest predictors of learning outcomes, while frequent rewinding showed weak negative correlation, potentially signaling learning difficulties rather than productive engagement.

RAG and Educational Applications. Simon Gottschalk from L3S Research Center explored LLM applications for learning contexts, presenting "EventExplorer," an interactive system using Retrieval-Augmented Generation to help users research historical events through web archives. The presentation highlighted key educational considerations including personalization opportunities, accessibility benefits, teacher support potential, and implementation challenges around data privacy and AI literacy.

AI-Empowered Open Education. Gábor Kismihók from TIB Hannover presented research on leveraging AI for personalized learning through Open Educational Resources, emphasizing that AI should serve as a support tool with pedagogical principles leading technological implementation. The work identified critical risks including "personalization bubbles" that hinder collaboration and "efficiency traps" that prioritize system performance over pedagogical effectiveness.

3.3. Discussion Themes

Four major themes emerged from workshop discussions:

Transformation of "Search as Learning." Participants questioned whether traditional search processes—characterized by query formulation, result evaluation, and iterative refinement—are being replaced by direct AI interactions that provide synthesized answers. This shift challenges existing research methods and theoretical frameworks, with concerns about lost opportunities for serendipitous discovery and critical evaluation. The urgent need for new datasets capturing contemporary behaviors emerged as critical, since current research often relies on data that may not reflect how users integrate AI-generated responses into learning processes.

Trust, Credibility and Information Verification. AI systems introduce significant challenges around credibility assessment. Traditional approaches—examining author credentials, publication venues, and citations—become problematic when AI synthesizes information from multiple sources without transparent attribution. Users need new critical evaluation skills for AI-mediated environments, including understanding how LLMs aggregate sources and maintaining healthy skepticism while benefiting from AI assistance.

Educational Applications and Personalization. While AI's potential to personalize learning could enhance effectiveness by adapting to individual needs, participants identified concerns about "personalization bubbles" that hinder collaboration. Discussion emphasized that pedagogical principles should guide technological implementation, with AI augmenting rather than substituting human judgment. Assessment challenges received particular attention, as traditional methods may be insufficient for AI-mediated environments.

Access, Adoption and Cultural Considerations. Significant concerns emerged about equity and access, with cultural variations in technology adoption creating potential divides. AI tools may exacerbate inequalities if they require high digital literacy or reliable internet access. Representative research requires capturing interaction patterns across diverse demographic, cultural, and socioeconomic groups rather than focusing solely on early adopters.

3.4. Future Directions

Workshop discussions revealed both immediate research priorities and longer-term strategic questions that will shape the field's development:

New Datasets and Continuous Behavioral Observation. The most urgent methodological challenge concerns developing datasets that capture contemporary information-seeking and learning behaviors. Even recent datasets may not adequately capture how users interact with LLMs today, while older datasets fail to represent the technological fluency of younger users. The field requires systematic, ongoing observation of shifting usage patterns across diverse user groups, with deliberate attention to cultural and demographic representation.

Research Priorities and Community Directions. Specific priorities emerged including research on collaborative human-AI learning environments where technology facilitates rather than replaces human interaction. Assessment and evaluation methodologies require significant development to measure learning in AI-mediated environments, necessitating novel approaches that can recognize the interconnected nature of AI and human contributions. Whether the field continues as "Search as Learning" or evolves into something broader will depend on how successfully researchers can adapt to rapidly changing technological and social contexts while maintaining focus on supporting effective human learning processes.

4. Cross-Workshop Insights

While DISMISS-FAKE and IWILDS approached web-based information from different angles—one focusing on detecting false content, the other on learning processes—their discussions revealed striking convergences that illuminate broader challenges in the AI-mediated information landscape.

The Trust Paradox. Both workshops grappled with credibility assessment in AI-generated environments. DISMISS-FAKE highlighted how LLMs can simultaneously generate and detect misinformation, creating fundamental questions about system reliability. IWILDS participants raised parallel concerns about users' inability to evaluate AI-synthesized answers when traditional credibility markers (author credentials, citations, publication venues) disappear. This shared challenge points to an urgent need for new frameworks that help users critically evaluate AI-mediated information, whether they're fact-checking claims or acquiring knowledge.

Cultural and Linguistic Exclusion. Both communities identified how current systems risk marginalizing underrepresented populations. DISMISS-FAKE's work on multilingual fact-checking datasets revealed significant performance gaps for low-resource languages, while IWILDS discussions emphasized how AI tools may exacerbate educational inequalities if they require high digital literacy or reliable internet access. The parallel concerns suggest that equity considerations must be central to system design, not afterthoughts—requiring diverse datasets, culturally-aware models, and deliberate attention to varied user contexts.

The Personalization-Collaboration Tension. A subtle but important theme emerged around individualization versus collective knowledge-building. IWILDS participants worried about "personalization bubbles" that isolate learners, while DISMISS-FAKE discussions touched on how echo chambers and narrative-driven misinformation exploit similar filtering mechanisms. Both point to a broader design challenge: how can AI systems support individual needs while maintaining opportunities for diverse perspectives, serendipitous discovery, and collaborative sense-making?

Methodological Urgency. Perhaps most practically, both workshops identified the critical need for new datasets and evaluation approaches. IWILDS emphasized that existing behavioral data fails to capture how younger users interact with LLMs, while DISMISS-FAKE highlighted gaps in multilingual and multimodal misinformation datasets. Both communities face a common challenge: research methods developed for traditional web search and social media may be inadequate for studying AI-mediated information environments that are evolving faster than our ability to document them.

These convergences suggest that the boundary between "learning" and "misinformation detection" may be less clear than disciplinary divisions suggest. Both ultimately concern how humans construct reliable understanding in information-rich, AI-mediated environments—a challenge that will only intensify as these technologies become more sophisticated and ubiquitous.

5. Closing

This joint proceedings volume spans multiple disciplines – from computer science and information retrieval to educational psychology and law – reflecting the inherently interdisciplinary nature of information processing in the digital age. Whether examining how learners navigate conflicting sources during web search or how detection systems can identify AI-generated misinformation, the work collected here addresses fundamental questions about human-information interaction that transcend traditional academic boundaries.

As we stand at this intersection of human cognition and artificial intelligence, the research presented here illuminates pathways toward more effective, trustworthy, and educationally valuable information systems. The collaboration between these two research communities represents not just academic cooperation, but a recognition that the challenges of information quality, learning, and truth-seeking in the digital age require comprehensive, multifaceted approaches.

We hope this collection serves both as a snapshot of current research frontiers and as inspiration for future work at the critical intersection of technology, cognition, and society's relationship with information.

The Organizing Committees
DISMISS-FAKE Workshop and IWILDS Workshop

Acknowledgement

Koustav Rudra is a recipient of the DST-INSPIRE Faculty Fellowship (DST/ INSPIRE/04/2021/003055 in the year 2021 under Engineering Sciences).

Declaration on Generative AI

During the preparation of this work, the author(s) used Claude.AI in order to: Grammar and spelling check, Paraphrase and reword. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.