



PDF Download  
3742414.3794961.pdf  
06 April 2026  
Total Citations: 1  
Total Downloads: 100

Latest updates: <https://dl.acm.org/doi/10.1145/3742414.3794961>

EXTENDED-ABSTRACT

## **CURE2026: Communicating Uncertainty to foster Realistic Expectations via Human-Centered Design**

**JASMINA GAJCIN**, IBM Research Europe, Ireland, Dublin, Ireland

**JOVAN JEROMELA**, Trinity College Dublin, Dublin, Leinster, Ireland

**JOEL WESTER**, University of Copenhagen, Copenhagen, Hovedstaden, Denmark

**SARAH SCHÖMBS**, School of Computing and Information Systems, Melbourne, VIC, Australia

**STYLIANI KLEANTHOUS**, Open University of Cyprus, Nicosia, Cyprus

**KARTHIKEYAN NATESAN RAMAMURTHY**, IBM Thomas J. Watson Research Center, Yorktown Heights, NY, United States

[View all](#)

Open Access Support provided by:

[Utrecht University](#)

[IBM Thomas J. Watson Research Center](#)

[Trinity College Dublin](#)

[School of Computing and Information Systems](#)

[Open University of Cyprus](#)

[Ludwig-Maximilians-University Munich](#)

[View all](#)

Published: 22 March 2026

[Citation in BibTeX format](#)

IUI '26: 31st International Conference on Intelligent User Interfaces  
March 23 - 26, 2026  
Paphos, Cyprus

Conference Sponsors:  
[SIGCHI](#)  
[SIGAI](#)

# CURE2026: Communicating Uncertainty to foster Realistic Expectations via Human-Centered Design

Jasmina Gajcin  
IBM Research  
Dublin, Ireland  
jasmina.gajcin2@ibm.com

Jovan Jeromela  
School of Computer Science and  
Statistics  
Trinity College Dublin  
Dublin, Ireland  
jeromelj@tcd.ie

Joel Wester  
University of Copenhagen  
Copenhagen, Denmark  
Aalborg University  
Aalborg, Denmark  
joel.wester@di.ku.dk

Sarah Schömbms  
School of Computing and Information  
Systems  
The University of Melbourne  
Melbourne, VIC, Australia  
sschombs@student.unimelb.edu.au

Styliani Kleanthous  
Cyprus Centre for Trustworthy AI  
Open University of Cyprus  
Nicosia, Cyprus  
styliani.kleanthous@ouc.ac.cy

Karthikeyan Natesan  
Ramamurthy  
IBM Research AI  
IBM T.J. Watson Research Center  
Yorktown Heights, New York, USA  
knatesa@us.ibm.com

Hanna Hauptmann  
Utrecht University  
Utrecht, Netherlands  
h.j.hauptmann@uu.nl

Rifat Mehreen Amin  
LMU Munich  
Munich, Germany  
rifat.amin@ifi.lmu.de

## Abstract

CURE2026 workshop provides a venue for research on how interactive and intelligent interfaces can be leveraged to ensure effective uncertainty communication. With the increased capabilities of black-box AI models, particularly large language models (LLMs), there is a growing urgency to understand exactly what these models can and cannot do. Effective uncertainty communication can help users calibrate their trust and reliance on the AI system. However, effectively conveying uncertainty, be it through conversation or by visualizing their confidence levels, remains a major challenge, as users often get confused or misinterpret uncertainty information due to poorly designed interfaces. CURE2026 workshop brings together researchers from HCI, AI, design, and psychology to investigate how we can design, develop, and evaluate intelligent and interactive interfaces for uncertainty communication. The workshop includes a mini-conference section aimed at inspiring discussion and investigating the role of intelligent and interactive interfaces in uncertainty communication. The workshop also features keynotes by Professors Jessica Hullman and Katrien Verbert. Finally, a part of the workshop is dedicated to an interactive activity to discuss and develop guidelines for communicating uncertainty across different domains and tasks.

## CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; *Interaction design*; *Visualization*; *Accessibility*; • **Computing methodologies** → **Artificial intelligence**; *Machine learning*.

## Keywords

intelligent user interfaces, user explanations, uncertainty, communication, trust calibration

## ACM Reference Format:

Jasmina Gajcin, Jovan Jeromela, Joel Wester, Sarah Schömbms, Styliani Kleanthous, Karthikeyan Natesan Ramamurthy, Hanna Hauptmann, and Rifat Mehreen Amin. 2026. CURE2026: Communicating Uncertainty to foster Realistic Expectations via Human-Centered Design. In *Companion Proceedings of the 31st International Conference on Intelligent User Interfaces (IUI Companion '26)*, March 23–26, 2026, Paphos, Cyprus. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3742414.3794961>

## 1 Introduction

The increased capabilities of black-box AI models and, in particular, large language models (LLMs), have brought forth the need for trustworthiness and explainability for their adoption in real-world tasks [37]. Communicating system uncertainty is essential for achieving transparency [4] and can help users calibrate their trust in, reliance on, and expectations from an AI system [6, 17, 19, 24, 31, 36]. This issue is exacerbated with LLMs which can generate convincing content, leading users to overestimate their certainty [30]. However, uncertainty communication is plagued by challenges such as cognitive biases [4, 25, 35], numeracy skills [11, 26], calibrating risk perception, [29] and increased cognitive load [13], with research finding that lay users can struggle to interpret probabilities and uncertainty visualizations [20].



This work is licensed under a Creative Commons Attribution 4.0 International License.  
*IUI Companion '26, Paphos, Cyprus*  
© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-1985-1/26/03  
<https://doi.org/10.1145/3742414.3794961>

The discrepancy between user expectations and system capabilities can result in poor user experience [33], lower usage willingness, and poor task performance [16, 23, 31]. From a visualisation-driven uncertainty representation on mobile screens through quantile dot-plots [18], to a text-based frequency [5] or first person uncertainty communication via LLMs [19], to an embodied uncertainty communication through an agent’s hesitation gestures [26], interfaces play a central role in enabling personalized, task-specific and interactive uncertainty communication across devices and modalities. This challenge is even more pronounced in agentic systems [34], where users need to navigate uncertainties across multiple agents without direct access to their actions and interactions [27]. Despite a plethora of research on uncertainty quantification [28], tools [12], visualization techniques [15] and taxonomies [3, 9, 14, 32], there is still no consensus on guidelines for uncertainty communication. For IUI, this raises the question of how to design affordances that do not solely display uncertainty but support appropriate interaction with it, across modalities and contexts.

CURE2026 investigates how intelligent user interfaces can help bridge the gap between system uncertainty and users’ expectations, aligning with the IUI theme “Where HCI meets AI”. The workshop discusses and disseminates novel research on designing, developing, and evaluating intelligent user interfaces for uncertainty communication. CURE2026 complements IUI topics of interest, such as explainable AI, user modelling, and innovative interface design, while providing a dedicated venue for investigating, discussing, and disrupting this under-researched topic. CURE2026 is a point of convergence for researchers across several IUI fields, including interface and interaction design, Human-Computer Interaction (HCI), personalization, data visualization, and cognitive science. We expect interest and participation from domain experts where uncertainty communication plays a vital role (e.g. healthcare [5, 7], intent modeling [16, 38], and financial forecasting [1, 21]).

## 2 Workshop History

This is the first edition of the CURE workshop. In the past, a number of workshops dealt with quantifying and leveraging uncertainty in AI (e.g. UNCV@CVPR 2025, Quantify Uncertainty and Hallucination in Foundation Models: The Next Frontier in Reliable AI @ICLR 2025, UncertainNLP@EMNLP 2025, UDL@ICML 2021). These workshops approach uncertainty from an AI perspective, and do not focus on how it is presented to and utilized by the end user. Additionally, a few workshops explore techniques for presenting uncertainty (e.g. Uncertainty Visualization@IEEE VIS 2025). Our workshop is not limited to visualization techniques, but spans the entire process of design, development, and evaluation of interfaces for uncertainty communication, including conversational interfaces.

The most relevant prior workshop is *Designing for Uncertainty in HCI: When does Uncertainty Help?* [13], hosted at the 2017 Conference on Human Factors in Computing Systems (CHI 2017). The goal of this workshop was to investigate how and when uncertainty should be presented to users and design guidelines for uncertainty communication. Our workshop revisits this challenge after almost a decade of sharp progress in AI, and aims to investigate what effective uncertainty communication looks like for a new generation of AI systems and interactive user interfaces.

## 3 Workshop Programme

The first workshop on Communicating Uncertainty to foster Realistic Expectations via Human-Centered Design (CURE@IUI) took place in person in Paphos, Cyprus on March 23, 2026. The workshop was a half-day event as part of the 31st ACM IUI conference. The workshop was split into three parts: 1) a mini-conference, 2) invited keynote presentations from experts in the field, and 3) an interactive session.

### 3.1 Mini-conference

The first part of the workshop is organized as a mini-workshop, which provides a platform for presenting and discussing novel contributions and works in progress. The workshop received a diverse set of submissions examining uncertainty communication through different lenses, including human centered design, conversational interactions, judgments of generative AI and human and LLM perception for uncertainty communication.

All workshop submissions underwent a rigorous double-blind review process, with each paper evaluated by at least two programme committee members. The selection criteria prioritised technical quality, theoretical soundness, novelty, reproducibility of results, and clarity of presentation to ensure high quality contributions to the field.

The workshop received six submissions, of which four papers were accepted. Although the call for papers invited both full papers and position papers, all accepted submissions were full papers (five page submissions). Below we present the summaries of the accepted papers.

*Conceptual Alignment in Human-Centered Design Teams.* One paper [8] investigates how conceptual alignment can be fostered in interdisciplinary and transdisciplinary Human-Centered Design teams through critical reflection on Large Language Models. Drawing on a case study of an interdisciplinary workshop, the work examines how reflecting on LLM-generated outputs can surface conceptual misalignments among team members. The paper further introduces a graphical user interface based on two-dimensional sentence embedding projections, designed to support shared reflection and discussion, and argues that such visual, collective interactions can enhance mutual understanding and trust when integrating LLMs into design processes.

*Visual Cues in Wrong-Path Conversational Interactions.* This study [2] explores the role of non-verbal visual cues in shaping user experience during wrong-path interactions with conversational AI systems in an e-commerce context. Through a controlled experimental study, the authors examine how elements such as avatars, emojis, and interactive buttons influence task efficiency and emotional comfort when task progress is silently blocked. The findings suggest that while visual cues can improve efficiency, they may also increase user frustration by amplifying expectations, highlighting the need for careful design when communicating system limitations.

*Philosophers’ Longitudinal Judgments of Generative AI.* A third contribution [22] presents a longitudinal mixed-methods study examining how philosophers’ evaluations of generative AI-based Intelligent User Interfaces evolve over time. Combining repeated blind assessments with in-depth interviews over three years, the

study reveals a trajectory from initial resistance to instrumental acceptance and, ultimately, principled skepticism. While quantitative results show increasing appreciation of AI-generated philosophical responses, qualitative findings highlight a persistent tension between perceived performance and normative beliefs about philosophical agency and understanding.

*Human and LLM Perception of Uncertainty in Social Discourse.* The fourth paper [10] investigates how humans and Large Language Models perceive uncertainty in polarized and moderate social media discourse within a multi-agent simulation environment. By comparing human judgments with those of calibrated LLM “mirror personas,” the study finds that both reliably distinguish discourse conditions, but that LLMs systematically amplify differences, particularly for uncertainty. The results reveal fundamental differences in how uncertainty integrates with other perceptual constructs, offering insights into the limits and potential of LLMs as proxies for human judgment in uncertainty-sensitive contexts.

### 3.2 Expert Talks

The second part of the workshop consisted of two invited keynote presentations delivered by Professor Jessica Hullman and Professor Katrien Verbert.

Professor Hullman is a Ginni Rometty Professor of Computer Science and Faculty Fellow at the Institute for Policy Research at Northwestern University. Her research focuses on achieving human-AI complementarity. Central themes in her research are uncertainty quantification and presentation in human-AI interaction. In her keynote talk titled *Aligning Uncertainty Quantification with Human Decision-makers Needs*, Prof. Hullman discusses how to design AI systems that effectively support human decision-making by accounting for the fact that humans need to assess their own reliability alongside the AI’s recommendations and may possess distinct information unavailable to the model. The keynote covers decision-theoretic approaches to designing effective human-AI collaboration and shares insights from observing how people use predictive uncertainty in practice.

The second keynote speaker, Prof. Katrien Verbert, is a Professor at the HCI research group of the Department of Computer Science of KU Leuven working on interactive recommender systems. Prof. Verbert’s research centers around trust, recommendation acceptance, and how to enable end-users to actively steer the recommendation process. In the second keynote titled *Beyond transparency: interactive explanations for trust calibration under uncertainty*, Prof. Verbert presents interactive explanation methods designed for specific end users, including approaches that combine data-centric and model-centric perspectives to help users refine predictive models. The keynote emphasises methods that support trust calibration and enable users to guide AI systems, moving beyond passive transparency to active user engagement.

### 3.3 Workshop Activity: Guidelines for Uncertainty Communication

Finally, the workshop will conclude with an interactive section titled *Guidelines for Uncertainty Communication*. In this section, participants worked in small groups to brainstorm and discuss

guidelines for designing and evaluating interfaces for communicating uncertainty. Such guidelines are essential for the IUI community as AI systems are increasingly deployed across diverse contexts and the range of interface types for interacting with these systems continues to expand. The presented papers, keynotes, and participants’ interdisciplinary backgrounds supported the collaborative effort by synthesising insights from multiple disciplines to inform design considerations and build a shared understanding around effective uncertainty communication in interactive interfaces.

## 4 Conclusion and Future Directions

The CURE2026 workshop at IUI 2026 indicates the growing importance of communicating uncertainty for providing realistic expectations. We expect that the discussions from the workshop will help foster a sense of community, inform future design decisions, and raise awareness in the broader field.

## 5 Program Committee

- Abhraneel Sarma, Northwestern University, USA
- Alex C. Williams, AWS, USA
- Alok Debnath, Trinity College Dublin, Ireland
- Danielle Albers Szafir, University of North Carolina - Chapel Hill, USA
- Eoin Delaney, Trinity College Dublin, Ireland
- Eelco Herder, Utrecht University, Netherlands
- Francesca Naretto, University of Pisa, Italy
- Hariharan Subramonyam, Stanford University, USA
- Hyo Jin Do, IBM Research, USA
- Ivana Dusparic, Trinity College Dublin, Ireland
- Jacy Reese Anthis, University of Chicago, USA
- Jan Leusmann, LMU Munich, Germany
- Jeroen Ooge, Utrecht University, Netherlands Justin Edwards, University of Oulu, Finland
- Leon Fröhling, GESIS Leibniz Institute for Social Sciences, Germany
- Maxwell Szymanski, KU Leuven, Belgium
- Michael Hedderich, LMU Munich, Germany
- Owen Conlan, Trinity College Dublin, Ireland
- Prakash Jamakatel, University of the Bundeswehr Munich, Germany
- Rune Møberg Jacobsen, Aalborg University, Denmark
- Samuel Rhys Cox, Aalborg University, Denmark
- Selene Baez Santamaria, University of Zurich, Switzerland
- Siddharth Swaroop, Harvard University, USA
- Simret Araya Gebreegziabher, University of Notre Dame, USA
- Siya Kunde, IBM Research, USA
- Srishti Yadav, University of Copenhagen, Denmark
- Sunnie S.Y. Kim, Apple, USA
- Timo Speith, University of Bayreuth, Germany
- Yan Zhang, University of Melbourne, Australia
- Yue Fu, University of Washington, USA

## Acknowledgments

We thank the authors for their valuable contributions to the workshop, the programme committee for their commitment to rigorous

peer review, and the ACM IUI 2026 organisers for providing a platform to discuss and advance the important topic of uncertainty communication in interactive interfaces. This publication has emanated from research conducted in part with the financial support of Research Ireland under grant 18/CRT/6223. This publication is supported in part by the Elitenetzwerk Bayern. This work was funded in part by the EU Horizon project ELIAS (No. 101120237). Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or The European Research Executive Agency.

## References

- [1] Arash Bahrammirzaee. 2010. A comparative survey of artificial intelligence applications in finance: artificial neural networks, expert system and hybrid intelligent systems. *Neural Computing and Applications* 19, 8 (2010), 1165–1195. doi:10.1007/s00521-010-0362-z
- [2] Jie Bai and Alessandra Zarcone. 2026. Visual Cues in Wrong-Path Interactions: An Exploratory Study of User Experience in Conversational AI. In *Companion Proceedings of the 31st International Conference on Intelligent User Interfaces* (Paphos, Cyprus, March 23-26, 2026) (*IUI '26 Companion*). ACM, New York, NY, USA.
- [3] Howard Balslem, Mark Helfand, Holger J Schünemann, Andrew D Oxman, Regina Kunz, Jan Brozek, Gunn E Vist, Yngve Falck-Ytter, Joerg Meerpohl, Susan Norris, et al. 2011. GRADE guidelines: 3. Rating the quality of evidence. *Journal of clinical epidemiology* 64, 4 (2011), 401–406.
- [4] Umang Bhatt, Javier Antorán, Yunfeng Zhang, Q Vera Liao, Prasanna Sattigeri, Riccardo Fogliato, Gabrielle Melançon, Ranganath Krishnan, Jason Stanley, Omesh Tickoo, et al. 2021. Uncertainty as a form of transparency: Measuring, communicating, and using uncertainty. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 401–413.
- [5] Shiye Cao, Anqi Liu, and Chien-Ming Huang. 2024. Designing for appropriate reliance: The roles of AI uncertainty presentation, initial user decision, and user demographics in AI-assisted decision-making. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–32.
- [6] Federico Maria Cau, Hanna Hauptmann, Lucio Davide Spano, and Nava Tintarev. 2023. Supporting high-uncertainty decisions through AI and logic-style explanations. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*. 251–263.
- [7] Yu Chen, Katrien Verbert, Kathrin Gerling, Marie-Elena Vanden Abeele, and Vero Vanden Abeele. 2025. Will Health Experts Adopt a Clinical Decision Support System for Game-Based Digital Biomarkers? Investigating the Impact of Different Explanations on Perceived Ease-of-Use, Perceived Usefulness, and Trust. In *Proceedings of the 30th International Conference on Intelligent User Interfaces*. 1581–1597.
- [8] Simone Ciciliano and Rosella Gennari. 2026. Conceptual Alignment in Human-Centered Design Teams through Critical Reflection on Large Language Models. In *Companion Proceedings of the 31st International Conference on Intelligent User Interfaces* (Paphos, Cyprus, March 23-26, 2026) (*IUI '26 Companion*). ACM, New York, NY, USA.
- [9] Paulo CG Costa, Kathryn B Laskey, Erik Blasch, and Anne-Laure Joussemle. 2012. Towards unbiased evaluation of uncertainty reasoning: The URREF ontology. In *2012 15th International Conference on Information Fusion*. IEEE, 2301–2308.
- [10] Tim Donkers and Jürgen Ziegler. 2026. Uncertainty perception by humans and LLMs in a multi-agent social network simulation. In *Companion Proceedings of the 31st International Conference on Intelligent User Interfaces* (Paphos, Cyprus, March 23-26, 2026) (*IUI '26 Companion*). ACM, New York, NY, USA.
- [11] Mirta Galesic and Rocio Garcia-Retamero. 2010. Statistical numeracy for health: a cross-cultural comparison with probabilistic national samples. *Archives of internal medicine* 170, 5 (2010), 462–468.
- [12] Soumya Ghosh, Q Vera Liao, Karthikeyan Natesan Ramamurthy, Jiri Navratil, Prasanna Sattigeri, Kush Varshney, and Yunfeng Zhang. 2022. Uncertainty quantification 360: a hands-on tutorial. In *Proceedings of the 5th Joint International Conference on Data Science & Management of Data (9th ACM IKDD CODS and 27th COMAD)*. 333–335.
- [13] Miriam Greis, Jessica Hullman, Michael Correll, Matthew Kay, and Orit Shaer. 2017. Designing for uncertainty in hci: When does uncertainty help?. In *Proceedings of the 2017 CHI conference extended abstracts on human factors in computing systems*. 593–600.
- [14] Jon C Helton, Jay D Johnson, and William L Oberkampf. 2004. An exploration of alternative approaches to the representation of uncertainty in model predictions. *Reliability Engineering & System Safety* 85, 1-3 (2004), 39–71.
- [15] Jessica Hullman, Xiaoli Qiao, Michael Correll, Alex Kale, and Matthew Kay. 2018. In pursuit of error: A survey of uncertainty visualization evaluation. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 903–913.
- [16] Jovan Jeromela and Owen Conlan. 2023. Voicing Suggestions and Enabling Reflection: Results of an Expert Discussion on Proactive Assistants for Time Management. Article 48 pages. doi:10.1145/3571884.3604317
- [17] Susan L Joslyn and Jared E LeClerc. 2012. Uncertainty forecasts improve weather-related decisions and attenuate the effects of forecast error. *Journal of experimental psychology: applied* 18, 1 (2012), 126.
- [18] Matthew Kay, Tara Kola, Jessica R. Hullman, and Sean A. Munson. 2016. When (ish) is My Bus? User-centered Visualizations of Uncertainty in Everyday, Mobile Predictive Systems. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (*CHI '16*). Association for Computing Machinery, New York, NY, USA, 5092–5103. doi:10.1145/2858036.2858558
- [19] Sunnie SY Kim, Q Vera Liao, Mihaela Vorvoreanu, Stephanie Ballard, and Jennifer Wortman Vaughan. 2024. "I'm Not Sure, But...": Examining the Impact of Large Language Models' Uncertainty Expression on User Reliance and Trust. In *Proceedings of the 2024 ACM conference on fairness, accountability, and transparency*. 822–835.
- [20] Stephanie Lem, Patrick Onghena, Lieven Verschaffel, and Wim Van Dooren. 2014. Interpreting histograms. As easy as it seems? *European journal of psychology of education* 29, 4 (2014), 557–575.
- [21] Spyros Makridakis and Nikolas Bakas. 2016. Forecasting and uncertainty: A survey. *Risk and Decision Analysis* 6, 1 (2016), 37–64. doi:10.3233/rda-150114
- [22] Yibo Meng, Lyumanshan Ye, Bingyi Liu, and Yan Guan. 2026. Can Intelligent User Interfaces Engage in Philosophical Discussions?: A Longitudinal Mixed-Methods Study of Philosophers' Judgments. In *Companion Proceedings of the 31st International Conference on Intelligent User Interfaces* (Paphos, Cyprus, March 23-26, 2026) (*IUI '26 Companion*). ACM, New York, NY, USA.
- [23] Swati Mishra and Jeffrey M Rzeszotarski. 2023. Human Expectations and Perceptions of Learning in Machine Teaching. 13–24 pages. doi:10.1145/3565472.3595612
- [24] Samir Passi and Mihaela Vorvoreanu. 2022. Overreliance on AI literature review. *Microsoft Research* 339 (2022), 340.
- [25] Valerie F Reyna and Charles J Brainerd. 2008. Numeracy, ratio bias, and denominator neglect in judgments of risk and probability. *Learning and individual differences* 18, 1 (2008), 89–107.
- [26] Sarah Schömbms, Saumya Pareek, Jorge Goncalves, and Wafa Johal. 2024. Robot-Assisted Decision-Making: Unveiling the Role of Uncertainty Visualisation and Embodiment. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (*CHI '24*). Association for Computing Machinery, New York, NY, USA, Article 589, 16 pages. doi:10.1145/3613904.3642911
- [27] Sarah Schömbms, Yan Zhang, Jorge Goncalves, and Wafa Johal. 2026. From Conversation to Orchestration: HCI Challenges and Opportunities in Interactive Multi-Agent Systems. In *Proceedings of the 13th International Conference on Human-Agent Interaction (HAI '25)*. Association for Computing Machinery, New York, NY, USA, 158–168. doi:10.1145/3765766.3765795
- [28] Ola Shorinwa, Zhiting Mei, Justin Lidard, Allen Z Ren, and Anurudha Majumdar. 2025. A survey on uncertainty quantification of large language models: Taxonomy, open research challenges, and future directions. *Comput. Surveys* (2025).
- [29] David Spiegelhalter. 2017. Risk and Uncertainty Communication. *Annual Review of Statistics and Its Application* 4, Volume 4, 2017 (March 2017), 31–60. doi:10.1146/annurev-statistics-010814-020148 Publisher: Annual Reviews.
- [30] Mark Steyvers, Heliodoro Tejada, Aakriti Kumar, Catarina Belem, Sheer Karny, Xinyue Hu, Lukas W Mayer, and Padhraic Smyth. 2025. What large language models know and what people think they know. *Nature Machine Intelligence* 7, 2 (2025), 221–231.
- [31] Madiha Tabassum, Tomasz Kosiński, Alisa Frik, Nathan Malkin, Primal Wijesekera, Serge Egelman, and Heather Richter Lipford. 2019. Investigating Users' Preferences and Expectations for Always-Listening Voice Assistants. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 4 (2019), Article 153. doi:10.1145/3369807
- [32] Anne Marthe Van Der Bles, Sander Van Der Linden, Alexandra LJ Freeman, James Mitchell, Ana B Galvao, Lisa Zaval, and David J Spiegelhalter. 2019. Communicating uncertainty about facts, numbers and science. *Royal Society open science* 6, 5 (2019), 181870.
- [33] Joel Wester, Tim Schrills, Henning Pohl, and Niels van Berkel. 2024. "As an AI language model, I cannot": Investigating LLM Denials of User Requests. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (*CHI '24*). Article 979, 14 pages. doi:10.1145/3613904.3642135
- [34] Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryan W White, Doug Burger, and Chi Wang. 2023. AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation. arXiv:2308.08155 [cs.AI] https://arxiv.org/abs/2308.08155
- [35] Hang Zhang and Laurence T Maloney. 2012. Ubiquitous log odds: a common representation of probability and frequency distortion in perception, action, and cognition. *Frontiers in neuroscience* 6 (2012), 1.
- [36] Yunfeng Zhang, Q Vera Liao, and Rachel KE Bellamy. 2020. Effect of confidence and explanation on accuracy and trust calibration in AI-assisted decision making.

- In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 295–305.
- [37] Haiyan Zhao, Hanjie Chen, Fan Yang, Ninghao Liu, Huiqi Deng, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, and Mengnan Du. 2024. Explainability for large language models: A survey. *ACM Transactions on Intelligent Systems and Technology* 15, 2 (2024), 1–38.
- [38] Çağla Çiğ Karaman and Tevfik Metin Sezgin. 2018. Gaze-based predictive user interfaces: Visualizing user intentions in the presence of uncertainty. *International Journal of Human-Computer Studies* 111 (2018), 78–91. doi:10.1016/j.ijhcs.2017.11.005