PROPOSAL FOR ICLR 2026 WORKSHOP ON LOGICAL REASONING OF LARGE LANGUAGE MODELS

Haoxuan Li[1], **Fenrong Liu**[2,3], **Michael Witbrock**[4], **Arman Cohan**[5], **Peter Clark**[6], **Johan van Benthem**[7,2], **Zhouchen Lin**[1]

[1] Peking University, China
[2] Tsinghua University, China
[3] Universiry of Amsterdam, Netherlands
[4] The University of Auckland, New Zealand
[5] Yale University, USA
[6] Allen Institute for AI, USA
[7] Stanford University, USA

> *"The **logic-inspired** approach:*
> *The essence of intelligence is **reasoning**.*
> *This is done by using **symbolic rules** to manipulate **symbolic expressions**.*
> *Learning can wait. Understanding how knowledge is represented in **symbolic expressions must come first**."*
>
> ***— Geoffrey Hinton, 2025.7.26, in WAIC, Shanghai, China***

1 WORKSHOP SUMMARY

Large language models (LLMs) have achieved remarkable breakthroughs in natural language understanding and generation, but their logical reasoning capabilities remain a significant bottleneck. Logical reasoning is crucial for tasks requiring precise deduction, induction, or abduction, such as medical diagnosis, legal reasoning, and scientific hypothesis verification. However, LLMs often fail to handle complex logical problems with multiple premises and constraints, and they frequently produce self-contradictory responses across different questions. These limitations not only restrict the reliability of LLMs in complex problem-solving but also hinder their real-world applications.

In response to these emerging needs, we propose the workshop on Logical Reasoning of LLMs. This workshop will explore the challenges and opportunities for improving deduction, induction, and abduction capabilities of LLMs, implementing symbolic representation and reasoning via LLMs, avoiding logical contradictions across responses to multiple related questions, enhancing LLM reasoning by leveraging external logical solvers, and benchmarking LLM logical reasoning and consistencies. As LLMs continue to expand their role in AI research and applications, this workshop will serve as a platform to discuss and refine the methods for advancing logical reasoning within LLMs.

Overview of Discussion Topics. The workshop will cover topics including but not limited to:

(1) Deduction, Induction, and Abduction Capabilities of LLMs: LLMs' logical reasoning lies in three core modalities—deduction (deriving specific conclusions from general rules), induction (inferring general patterns from specific observations), and abduction (identifying plausible explanations for observed phenomena).

- For **deduction enhancement**: How can we design training frameworks that enable LLMs to strictly adhere to formal logical rules (e.g., syllogisms, first-order logic) when deriving conclusions, reducing errors like "non sequiturs" or inconsistent inferences?
- For **induction optimization**: What techniques (e.g., few-shot meta-learning, synthetic data augmentation of diverse case sets) can help LLMs generalize reliable patterns from limited or noisy observations, while avoiding overfitting to spurious correlations?
- For **abduction refinement**: How can we equip LLMs to generate logically coherent, contextually relevant explanations for observations—especially in ambiguous scenarios with multiple potential causes—and rank these explanations by plausibility?

(2) Symbolic Representation and Reasoning: A core component of enhancing logical reasoning is the ability to map fluid natural language to structured symbolic representations that are amenable to formal manipulation.

- For **identifying relevant information**: What types of natural language information—such as entities, relations, quantifiers, temporal expressions, and causal links—are most critical to extract and represent symbolically to support robust downstream logical inference?
- For **achieving precise translation**: How can we develop and refine neuro-symbolic techniques to achieve precise, context-aware translation from natural language into formal representations, minimizing the loss or distortion of semantic meaning during this conversion process?
- For **selecting the target language**: How do we choose the most appropriate target symbolic language (e.g., first-order logic, description logic, modal logic, knowledge graphs) for a given task, balancing the trade-offs between expressive power, computational complexity, and ease of translation?

(3) Avoiding Logical Contradictions Across Responses to Multiple Related Questions: LLMs are prone to producing responses contradicting themselves *across* different questions. For example, a state-of-the-art question-answering LLM Macaw, answers *Yes* to both questions *Is a magpie a bird?* and *Does a bird have wings?* but answers *No* to *Does a magpie have wings?*.

- For **maintaining conversational memory**: How can we develop model architectures or memory-augmented systems that effectively track, store, and recall previously stated facts and inferred conclusions within a single conversational context?
- For **identifying and resolving conflicts**: What self-correction mechanisms can be embedded within LLMs to enable them to detect potential contradictions in their own generated statements and proactively resolve them based on a predefined knowledge hierarchy or by flagging the ambiguity to the user?
- For **ensuring long-range coherence**: Beyond immediate conversational turns, how can we design systems that enforce logical consistency over extended dialogues or across multiple documents, ensuring that the model's reasoning remains sound and non-contradictory over complex, multi-step interactions?

(4) Enhancing LLM Reasoning by Leveraging External Logical Solvers: We explore the potential of combining the natural language understanding capabilities of LLMs with the rigorous, verifiable reasoning power of external symbolic systems like logical solvers and theorem provers.

- For **effective integration**: What are the most robust techniques for teaching LLMs to translate unstructured natural language problems into the structured formal representations (e.g., first-order logic, SMT-LIB) required by external solvers, and to interpret the solver's symbolic output back into a human-readable explanation?
- For **intelligent task delegation**: How can we train LLMs to discern when a problem is best solved using its internal, heuristic-based reasoning versus when it is necessary and more reliable to delegate the task to an external solver?
- For **handling ambiguity and uncertainty**: How can such hybrid systems manage real-world problems that contain inherent ambiguity or incomplete information, which are challenging for the rigid input requirements of most formal solvers?

(5) Benchmarks and Evaluation for Logical Reasoning and Consistencies: Current benchmarks often fall short in testing for deep, multi-step reasoning and multiple logical consistencies (e.g., negative consistency, transitivity consistency, etc.), creating a need for new evaluation tools.

- For **next-generation benchmark design**: How can we create challenging and diverse benchmarks that rigorously test a wide spectrum of logical reasoning skills, are resistant to "shortcut" learning, and can effectively distinguish between genuine reasoning and superficial pattern matching?

- For **metrics on consistency and faithfulness**: What automated methods and metrics can be developed to reliably quantify the logical consistency of an LLM's outputs, both within a single response (e.g., chain of thought) and across multiple related responses, ensuring the reasoning is sound and free of internal contradictions?
- For **interpretable and fine-grained evaluation**: How can we move beyond simple outcome-based accuracy to develop evaluation frameworks that provide deeper insights into the model's reasoning process itself, such as assessing the validity of intermediate reasoning steps or the logical structure of its generated arguments?

2 Recent Experience

Our team has published a survey paper Empowering LLMs with Logical Reasoning: A Comprehensive Survey accepted by **IJCAI 2025 (Survey Track, Accepted Rate: 19.6%)**, also selected for **oral presentation**, as shown in Figure 1(a). This paper has also **been reported by MIT Technology Review** and many mainstream social media outlets. In addition, we presented a **tutorial at IJCAI 25** in the Guangzhou session, as shown in Figure 1(b). Both events attracted a large number of participants.

Until now, we have never organized a workshop on this topic. We believe the time is now ripe and look forward to organizing a workshop on this topic at ICLR 2026.

(a) Oral Talk of our IJCAI 2025 Survey in Montreal

(b) Tutorial at IJCAI 2025 in Guangzhou

Figure 1: Our research topic has attracted widespread interest and strong engagement.

3 Tentative Workshop Schedule

Our hybrid workshop will feature 6 invited talks (30 min), 4 contributed talks (15 min), two 1-hour poster sessions, and a 1-hour panel discussion on the future of LLM logical reasoning. All sessions will be accessible for in-person and virtual attendance. Our initial website is available at https://sites.google.com/view/lr4llm-2026.

- 09:00-09:15: Introduction and Opening Remarks
- 09:15-09:45: Invited Talk 1
- 09:45-10:15: Invited Talk 2
- 10:15-11:15: Poster Session 1
- 11:15-11:30: Contributed Talk 1
- 11:30-12:00: Invited Talk 3
- 12:00-13:30: Lunch break
- 13:30-14:00: Invited Talk 4
- 14:00-14:30: Invited Talk 5

- 14:30-14:45: Contributed Talk 2
- 14:45-15:45: Poster Session 2
- 15:45-16:15: Invited Talk 6
- 16:15-16:30: Contributed Talk 3
- 16:30-16:45: Contributed Talk 4
- 16:45-17:45: Panel Discussion
- 17:45-18:00: Paper Awards and Closing Remarks

4 INVITED SPEAKERS AND PANELISTS

This workshop will feature eleven invited speakers and panelists organized into six invited talks and a panel discussion. This invited group consists of experts from diverse countries in America, Asia, Europe and Oceania. To date, all of the speakers and panelists have confirmed their participation.

Chitta Baral, Arizona State University is a Professor in the School of Computing and Augmented Intelligence at Arizona State University, and the leader of the Cognition and Intelligence Lab. He has published more than 200 scientific papers with 17,000+ citations in top-tier journals and conferences. He has received numerous accolades, including an Outstanding Paper Award at EACL 2023 and an NSF CAREER Award (1995). He served as Co-Chair of the Doctoral Consortium for IJCAI 2019 and Associate Editor for the Artificial Intelligence journal and JAIR.

Thomas Icard, Stanford University is C. I. Lewis Professor of Philosophy, Professor of Computer Science at Stanford University, and the Affiliate Faculty in Stanford AI Lab. He has published more than 70 papers with around 10,000 citations in logic and AI top-tier journals and conferences, spanning automated reasoning, symbolic systems, and logic reasoning for AI.

Yan Liu, University of Southern California is a full professor in the Computer Science Department, Viterbi School of Engineering at USC. Her research includes machine learning for time series, physics-informed machine learning, and interpretable machine learning, with applications to health, sustainability, and social media. She serves as associate Editor-in-Chief of TPAMI, Board Member and Finance Chair of ICLR. In addition, she previously served as general chair (or co-chair) for ICLR 2023 and ACM KDD 2020, program chair (or co-chair) of ICLR 2022, ACM KDD 2022, SDM 2020, WSDM 2018, and associate program co-chair for AAAI 2021.

Yue Zhang, Westlake University is a Tenured Full Professor and Associate Dean for the School of Engineering at Westlake University. His scholarly work has achieved broad recognition, evidenced by over 27,000 citations. He has served as Program Committee Co-chair for EMNLP 2022 and tutorial Co-Chair for ACL 2020. He has consistently acted as a (Senior) Area Chair for top-tier conferences such as ACL, NAACL, and EMNLP. He has garnered significant accolades, including Best Paper Awards at IALP 2017 and COLING 2018, and a Best Paper Nomination at ACL 2023.

Hanna Hajishirzi, University of Washington is a Torode Family Associate Professor in the Allen School of Computer Science and Engineering at the University of Washington and a Senior Director of NLP at AI2. She has published more than 140 scientific papers with 38,000+ citations in top-tier journals and conferences, in which many works have been nominated or received best paper awards. She is also a recipient of **2020 Alfred Sloan Fellowship** and 2021 NSF CAREER award.

Greg Durrett, UT Austin is an Associate Professor of Computer Science at UT Austin. He focuses on critical issues such as reasoning, factuality, and controllability in language models. His research excellence has been recognized with numerous prestigious awards, including a **Sloan Research Fellowship**, an NSF CAREER Award, an Outstanding Paper Award at EMNLP, and a Facebook Fellowship in Natural Language Processing.

Sean Ren, University of Southern California is an Associate Professor, Andrew and Erna Viterbi Early Career Chair in Computer Science, and the Director of INK Research Lab at the University of Southern California. His expertise is recognized by numerous prestigious awards, including honors such as MIT Technology Review Innovator Under 35 (Asia Pacific), Forbes Asia 30 Under 30, Samsung AI Researcher of the Year, NSF CAREER Award, Okawa Research Grant, WSDM Test of Time Award, and the Outstanding Paper Awards in NAACL 2022 and ACL 2023.

Liangming Pan, Peking University is an Assistant Professor at the School of Computer Science, Peking University, and a Researcher at the Institute of Computational Linguistics (ICL) at Peking University. His research primarily focuses on the development of trustworthy LLMs that are logical, truthful, and safe. His work has received the SAC Highlights Award at ACL 2025, Best Paper Runner-up at TRL workshop @ NeurIPS 2024, and the Area Chair Award (Question Answering) at IJCNLP-AACL 2023.

Vaishak Belle, University of Edinburgh is a Reader (Associate Professor), Chancellor's Fellow, Royal Society University Research Fellow, and Alan Turing Faculty Fellow at the University of Edinburgh. He served as Senior Program Committee Member for top conferences like AAAI and IJCAI. He also received a Best Paper Award at IJCLR (2024), the Microsoft Best Paper Award at the Conference on UAI (2015), and Best Student Paper Award at ECML/PKDD.

Ehsan Shareghi, Monash University is an Affiliated Lecturer of the University of Cambridge and an Assistant Professor at Monash University, leading a research team focusing on predictive models for language. He serves as (Senior) Area Chair for ACL, EMNLP, and EACL. He co-organized a workshop at NAACL 2022, and received an Outstanding Paper Award at ALTA 2024.

Damien Sileo, Inria Lille is a Tenured Researcher at Inria Lille, France's esteemed national research institute for digital science and technology. His work spans multi-task learning and reasoning tasks creation, and has been published in top-tier journals and conferences such as JMLR, NAACL, and EMNLP.

5 Organizers and biographies

5.1 Program Chairs

Haoxuan Li (Peking University, Male) (hxli@stu.pku.edu.cn) is an assistant researcher at Peking University, also a **Research Fellow of Tsinghua-UvA Joint Research Center for Logic at Tsinghua University**. He has more than 80 publications appeared in top-tier AI conferences and journals, including **more than 40 publications on ICML, NeurIPS and ICLR**, reported by MIT Technology Review and People's Daily. He received the NSFC Young Scientists Fund (2024) and Young Elite Scientists Sponsorship Program by CAST. Moreover, he served as the area chair (AC) of ICML, NeurIPS, ICLR, and has **co-organized the NeurIPS 2024 Workshop** on Causality and Large Models, **ICLR 2025 Workshop** on World Models, and **AAAI 2025 Workshop** on AI with Causal Techniques, and other workshops.

Fenrong Liu (Tsinghua University, Female) (fenrong@tsinghua.edu.cn) is currently a **Distinguished Professor** in the Department of Philosophy, Tsinghua University. She is now the Director of the THU-UvA Joint Research Centre for Logic and Amsterdam-China Logic Chair in ILLC, University of Amsterdam. She is a **Fellow of the Institut International de Philosophie (IIP)** and **Fellow of ASL Code of Ethics committee, International Academy of the Philosophy of Science (AIPS)** (the most important achievements in Philosophy including Logic community). Her research interests are mathematical logic, philosophical logic, modal logic, and LLMs.

Michael Witbrock (University of Auckland, Male) (m.witbrock@auckland.ac.nz) is a Full Professor of computer science at The University of Auckland in New Zealand, executive council of AI Forum NZ, science director of Precision Driven Health, co-founder and chair of the global AI for Good Foundation, the director of the Strong AI Lab (SAIL). His lab focus on **developing intelligence with symbolic AI, reasoning, and natural language understanding,** with an additional focus on maximizing the near-term benefit of AI to NZ entrepreneurs and business, and more generally achieving the best social and civilizational impacts of increasingly powerful AI.

Arman Cohan (Yale University, Male) (arman.cohan@yale.edu) is an **Assistant Professor of Computer Science at Yale University** and holds the position of faculty **Research Scientist at the Allen Institute for AI (Ai2)**. His research covers a range of topics that lie at the crossroads of Machine Learning and Natural Language Processing—these include Language Modeling, Representation Learning, as well as applications in specialized fields like science. Before joining Yale, he worked as a Research Scientist at AI2 and held the role of Affiliate Assistant Professor at the University of Washington. His research achievements have earned him several honors: these include a **best paper award at EMNLP and ACL, an outstanding paper award at EACL, and an**

honorable mention at COLING. Currently, he acts as the **Publications Chair for NAACL 2025**. In the past, he also organized the Workshop on SciNLP - Scientific NLP at AKBC 2020, along with the **Workshop on Scholarly Document Processing at NAACL 2021 and COLING 2022**.

Peter Clark (Allen Institute for AI, Male) (peterc@allenai.org) is a Senior Research Director and founding member of the Allen Institute for AI (AI2), and also served as Interim CEO from 2022-2023. He leads AI2's Aristo Project, aiming at developing AI agents that can systematically reason, explain, and continually improve over time, in particular in the context of scientific discovery. He has published over 250 papers, and has received five Best Paper awards (AAAI, EMNLPx3, AKBC), a Boeing Associate Technical Fellowship (2004), and Senior Membership of AAAI. **He and his team are regarded as the pioneers in the logical reasoning of LLMs**, with their work including **one of the earliest relevant datasets (ProofWriter)**, and **influential approach to logical consistency (BeliefBank).**

Johan van Benthem (Stanford University, Male) (johan@stanford.edu) is a **world-leading logician**, holding professorships at the University of Amsterdam, Stanford University, and Tsinghua University. He is a recipient of the prestigious Spinoza Prize (1996), the highest scientific award in the Netherlands. He was elected member of the Academia Europaea in 1991, the Royal Netherlands Academy of Arts and Sciences (KNAW) in 1992, and a **Foreign Fellow of the American Academy of Arts and Sciences in 2015.**

Zhouchen Lin (Peking University, Male) (zlin@pku.edu.cn) is a **Distinguished Professor in the School of Intelligence Science and Technology at Peking University**. He is a **Fellow of the Institute of Electrical and Electronics Engineers (IEEE)**, Fellow of the International Association of Pattern Recognition (IAPR), Fellow of the Asia-Pacific Artificial Intelligence Association (AAIA), **ICML Board Member, ACM Distinguished Member** and AAAI Senior Member. He earns the Okawa Research Grant and Microsoft SPOT Award. He is a **Senior Area Chair (SAC)** of ICML, NeurIPS, ICLR, AAAI and IJCAI, also an **Associate Editor-in-Chief of the IEEE Transactions on Pattern Analysis and Machine Intelligence** and an associate editor of the International Journal of Computer Vision.

5.2 ORGANIZERS

Chuan Zhou (The University of Melbourne, Male) (chuan.zhou@student.unimelb.edu.au) is a Ph.D student at the University of Melbourne. His research interest are logical and causal reasoning of LLMs. He has published several first-author papers in conferences such as NeurIPS, SIGKDD and ACL. He served as a tutor at the AAAI 26, ICDM 24 and ACML 24 tutorials on LLMs and causal reasoning, and co-organized the AAAI 2026 Bridge LMReasoning.

Fengxiang Cheng (University of Amsterdam, Female) (f.cheng@uva.nl) is currently a Ph.D. candidate in Logic at the University of Amsterdam's Institute for Logic, Language and Computation (ILLC). She has published papers in IJCAI, EMNLP, AAAI and ICLR. She was selected for IJCAI and AAAI Doctoral Consortium. She served as the first organizer of AAAI 2026 Bridge Program.

Ye Mo (**Zhejiang University, Female**) (moye017@zju.edu.cn) is currently a Ph.D. student in the College of Computer Science and Technology, Zhejiang University. Her research interests include logical reasoning, multimodal understanding, and pattern recognition. She has published work at CHI. Her work focuses on the relationship between AI systems and humans.

Hao Wang (**Zhejiang University, Male**) (haohaow@zju.edu.cn) is a Ph.D. candidate in the State Key Laboratory of Industrial Control Technology, Zhejiang University, Hangzhou, China. He is interested in LLM post-training, sequence modeling and causal machine learning. He has authored more than 40 papers in top-tier AI venues, served as the PC member or Area Chair for top conferences such as ICML, NeurIPS, ICLR and SIGKDD, and also the Associate Editor in IEEE SMC.

Licheng Pan (**Zhejiang University, Male**) (licheng.pan@zju.edu.cn) is currently a Ph.D. candidate in the College of Computer Science and Technology, Zhejiang University, Hangzhou, China. He has authored more than 10 papers in top-tier AI venues, such as ICLR, ICML, NeurIPS, and NDSS, etc. His research interests include causality, time-series analysis, LLM post-training and LLM security.

Haocheng Yang (**National University of Singapore, Male**) (haocheng_yang@u.nus.edu) is currently an undergraduate student at the School of Computing, National University of Singapore. He

has published papers at AI conferences such as NeurIPS, AAAI, and ICLR. His research interests include LLM logical reasoning, agentic tool use, LLM post-training, and multi-agent systems.

Chunyuan Zheng (**Peking University, Male**) (cyzheng@stu.pku.edu.cn) is currently a Ph.D. student at the School of Mathematical Sciences, Peking University. He has published over 20 papers in top-tier AI conferences and journals, including ICLR, ICML, NeurIPS, AAAI, SIGKDD, TOIS etc. He has been served as a Workshop Organizer in AAAI 2025 and ICDM 2025, Area Chair for IEEE DSAA, and reviewer for top conferences including lCML, NeurlPS and ICLR.

5.3 Program Committee Members

We have successfully recruited **60 distinguished members** to form our reviewing committee, including 43 leading experts and scholars from the fields of LLMs, along with 17 renowned experts and scholars with deep expertise in logical reasoning. They come from a wide range of leading universities, academic institutions, and industry labs. We are committed to leveraging a global and diverse perspective to ensure a review process that is professional and fair. Recruitment for the committee is currently ongoing, as we continue to seek additional experts from relevant fields to further enhance the committee's professional expertise and expand its breadth of coverage.

5.4 Previous Organizing Experience

Our team members have extensive experience in organizing workshops and symposiums, making us well-prepared to handle workshops organization at large conferences, with important lists as below:

- Logical and Symbolic Reasoning in Language Models: Bridge Program at AAAI 2026 (88 accepted papers);
- World Models: Understanding, Modelling and Scaling Workshop at ICLR 2025 (63/84);
- Reasoning and Planning for Large Language Models Workshop at ICLR 2025 (63/84);
- Causality and Large Model Workshop at NeurIPS 2024 (31/43 accepted papers);
- Causal Representation Learning Workshop at NeurIPS 2024 (37/50 accepted papers);
- Artificial Intelligence with Causal Techniques Workshop at AAAI 2025 (26/40 accepted papers);
- Knowledge-Intensive Multimodal Reasoning Workshop at ICCV 2025;
- Scholarly Document Processing Workshop at NAACL 2021;
- Program co-chair of IEEE International Conference on Data Mining (ICDM) 2024;
- Program co-chair of Uncertainty in Artificial Intelligence (UAI) 2022;
- Co-founder of Conference on Causal Learning and Reasoning (CLeaR), from 2022;
- Co-founder of Asian Workshop on Philosophical Logic (AWPL), from 2012;
- Co-founder of Workshop Series on the History of Logic in China (HOLIC), from 2010;
- Co-founder of Workshop on Logic, Rationality and Interaction (LORI), from 2007.

6 Organization Details

6.1 Important Dates for Review Process

We will follow the suggested submission date of ICLR 2026.

- Workshop paper submission deadline: 30 January 2026
- Workshop paper notification date: 1 March 2026
- Final workshop program, camera-ready, videos uploaded: 11 March 2026

All deadlines are specified in Anywhere on Earth (AoE).

6.2 SUPPORT FOR TINY OR SHORT PAPERS

Following the ICLR initiative, our workshop will feature a **Tiny/Short Paper (2-4 pages)** track to encourage the submission of late-breaking developments and novel ideas. We welcome less-than-full-conference papers, welcoming submissions that, for example, present an implementation and evaluation of an unpublished but simple idea, a modest but self-contained theoretical result, a follow-up experiment or re-analysis of a previously published paper, or a fresh perspective on an existing publication. All submissions must be primarily authored by humans. AI tools may be used for assistance, but AI-generated papers are not permitted for tiny or short papers.

6.3 LLM USAGE POLICY

This workshop follows the ICLR 2026 Policies on LLM Usage. We permit limited use of LLMs by authors (e.g., for proofreading, language refinement, or code generation), provided that the LLM does not make a substantial authorial contribution. LLMs cannot be listed as authors. All submissions must include a clear statement detailing whether and how LLMs were used in the manuscript.

To ensure the integrity and confidentiality of the review process, the use of any LLMs by reviewers is strictly prohibited.

6.4 ANTICIPATED AUDIENCE SIZE

The workshop is open to researchers, practitioners, and industry professionals interested in AI and logical reasoning. There are no specific criteria or maximum number of attendees, and we expect there are **around 350-400 attendees.**

6.5 PLAN TO GET AN AUDIENCE FOR A WORKSHOP

Our plan to maximize attendance and foster a vibrant intellectual environment is centered on clear communication and strategic promotion within the ICLR community. The following are our specific promotional plan measures: (1) Promote the confirmed schedule and speaker highlights on our website and social media platforms such as X (Twitter), Facebook, LinkedIn, and Instagram to showcase compelling content. (2)Announce the final program to relevant mailing lists, emphasizing key topics and learning opportunities for attendees. (3) Highlight the workshop's unique interactive sessions and networking events to stress the value of in-person attendance. (4) Leverage the app used in the ICLR conference for on-site promotion to attract attendees already at the venue.

6.6 VIRTUAL ACCESS TO WORKSHOP MATERIALS AND OUTCOME

Online Participation. To maximize the workshop's impact and ensure that valuable insights are accessible to a global audience, we will implement a comprehensive digital strategy. We will ensure that all invited talks and oral presentations will be conducted in person on-site, with all speakers attending the venue. Meanwhile, all oral presentations and keynotes will be professionally recorded and made publicly available on our workshop website shortly after the event. This initiative will cater to participants in different time zones and those unable to attend in person, while also creating a valuable resource for the wider research community.

Online Resources. We will launch and maintain a dedicated workshop website that serves as a central, dynamic hub for all materials. Prior to the workshop, the site will feature the complete program, speaker bios, and presentation abstracts to help attendees plan their schedule. Following the event, this platform will evolve into a comprehensive digital archive, hosting all accepted short papers, posters, presentation slides, and industry demos. All resources will be open-access, ensuring long-term availability and extending the reach of our contributors' work.

Fostering a Lasting Research Community. To facilitate sustained dialogue and collaboration, we will establish a moderated online discussion forum or a persistent Q&A channel (e.g., using a platform like Slack or Discord). This space will allow both workshop attendees and the broader public to interact directly with authors, pose follow-up questions, and debate the presented topics. By creating this channel for asynchronous communication, we aim to cultivate a durable and active community around the workshop's themes long after the conference concludes.

6.7 FUNDING AND SPONSOR

We have received sufficient funding from an industrial company, which enables us to fund the registration fee and partially fund the travel fee for the speakers and panelists. **We also plan to host a social dinner to facilitate communication between organizers, speakers, and attendees.**

In addition, we are engaging in discussions with Google, IBM, ByteDance, and UvA, and they have expressed interest in providing funding support for this workshop.

7 DIVERSITY COMMITMENT

Diversity of the Organizing Committee. The foundation of our workshop is an organizing team that embodies a rich spectrum of experiences and backgrounds. Our program chair and organizing committee group consisting of **3 female and 11 male** organizers, spans ten distinct institutions across **6 different countries** (United States, China, Netherlands, New Zealand, Australia and Singapore). This geographic breadth is matched by a remarkable diversity in career stages, uniting individuals **from Assistant to Full, and Chair Professors** on the academic side, alongside **Senior Research Director and CEOs** from the industrial sector. A pivotal strength of our organizing body is the **intentional fusion of perspectives from the logic and artificial intelligence communities**, with 2 members hailing from the former and others from the latter, ensuring a deeply integrated and cross-disciplinary approach.

Diversity of Speakers. Our commitment to a multitude of viewpoints extends directly to our lineup of invited speakers and panelists. Our lineup of speakers is consciously composed of **both females and males**, bringing together a valuable mix of **varying seniority**, from rising talents to seasoned pioneers in the field. **All speakers are confirmed to present in-person**, with affiliations in **5 different countries** (United States, United Kingdom, China, Australia, and France), promising a truly international exchange of ideas.

Diversity of Participants. We are committed to fostering a diverse and inclusive participant base. Our outreach strategy will encompass a broad range of disciplines, professional backgrounds, and geographic regions, with a specific focus on encouraging and supporting the attendance of individuals from underrepresented communities. Through these initiatives, we will cultivate an inclusive workshop environment where a plurality of voices can engage in and meaningfully shape the scientific conversation.

8 PREVIOUS RELATED WORKSHOPS

Several recent workshops have explored themes relevant to our focus, including:

- ICML 2023 Workshop on Knowledge and Logical Reasoning in the Era of Data-driven Learning
- ICLR 2024 Workshop on LLM Agents
- IJCAI 2024 Workshop on Interactions between Analogical Reasoning and Machine Learning
- ACL 2024 Workshop on Natural Language Reasoning and Structured Explanations
- ICLR 2025 Workshop on Workshop on Reasoning and Planning for LLMs
- NeurIPS 2025 Workshop on System 2 Reasoning At Scale
- AAAI 2025 Workshop on Neural Reasoning and Mathematical Discovery

We found that the previous workshop mainly focused on general reasoning and planning of LLMs, or specific reasoning capabilities such as code reasoning or math reasoning, and to the best of our knowledge, we are the first workshop focus on logical reasoning of LLMs.