

Reasoning for IR & IR for Reasoning

Mohanna Hoveyda

Radboud University
Nijmegen, The Netherlands
mohanna.hoveyda@ru.nl

Panagiotis Eustratiadis

University of Amsterdam
Amsterdam, The
Netherlands
p.efstratiadis@uva.nl

Arjen P. de Vries

Radboud University
Nijmegen, The Netherlands
arjen.devries@ru.nl

Maarten de Rijke

University of Amsterdam
Amsterdam, The
Netherlands
m.derijke@uva.nl

Abstract

Information retrieval has long focused on ranking documents by semantic relatedness. Yet many real-world information needs demand more: enforcement of logical constraints, multi-step inference, and synthesis of multiple pieces of evidence. Addressing these requirements is, at its core, a problem of *reasoning*. Across AI communities, researchers are developing diverse solutions for the problem of reasoning, from inference-time strategies and post-training of LLMs, to neuro-symbolic systems, Bayesian and probabilistic frameworks, geometric representations, and energy-based models. These efforts target the same problem: to move beyond pattern-matching systems toward structured, verifiable inference. However, they remain scattered across disciplines, making it difficult for IR researchers to identify the most relevant ideas and opportunities. To help navigate the fragmented landscape of research in reasoning, this tutorial first articulates a working definition of reasoning within the context of information retrieval and derives from it a unified analytical framework. The framework maps existing approaches along axes that reflect the core components of the definition. By providing a comprehensive overview of recent approaches and mapping current methods onto the defined axes, we expose their trade-offs and complementarities, highlight where IR can benefit from cross-disciplinary advances, and illustrate how retrieval process itself can play a central role in broader reasoning systems. The tutorial will equip participants with both a conceptual framework and practical guidance for enhancing reasoning-capable IR systems, while situating IR as a domain that both benefits and contributes to the broader development of reasoning methodologies.

CCS Concepts

• **Information systems** → **Retrieval tasks and goals; Retrieval effectiveness.**

Keywords

Reasoning for information retrieval, Complex queries, Neurosymbolic and probabilistic reasoning, LLM-based reasoning, Reinforcement learning for reasoning

Cover Sheet Information

Title: Reasoning for IR & IR for Reasoning.

Length: Half-day (3 hours).



This work is licensed under a Creative Commons Attribution 4.0 International License.
SIGIR '26, Melbourne, VIC, Australia

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2599-9/2026/07

<https://doi.org/10.1145/3805712.3808644>

Format: This is a half-day (3-hour) lecture-style tutorial, including interactive polls and short Q&A slots after each block to encourage engagement. scheduled breaks will be included with respect to the conference schedule. It will be conducted on-site, with at least three of the presenters planning to attend the conference in person.

Intended Audience: The tutorial is designed for a broad audience within and beyond the IR community, with a focus on three groups: (i) *Early-career IR researchers and PhD students* who wish to tackle complex information needs (e.g., negation, set-composition, and multi-hop), and who are seeking conceptual frameworks to guide new research directions. (ii) *Applied researchers and practitioners* in search, recommendation, or conversational systems who want to enhance the robustness of their products in edge cases where current semantic-matching approaches fail. (iii) *Researchers from related fields* (e.g., NLP, ML, knowledge representation) who are interested in how diverse reasoning methodologies can be applied to IR and, conversely, how IR can serve general reasoning models.

Prerequisite Knowledge: We assume familiarity with modern IR and language models, but no prior expertise in formal logic, RL, or specialized representations/optimization approaches, remaining accessible to newcomers yet technically engaging for experts.

Previous Talks: This tutorial is accepted and will be presented at ECIR on March 29th, 2026, in the Netherlands, making it the first presentation of this tutorial.

Presenters:

Mohanna Hoveyda (Radboud University, primary contact) is a PhD candidate in Conversational AI and IR. She develops frameworks for orchestrating multi-agent, LLM-based pipelines via reinforcement learning, and neuro-symbolic and probabilistic frameworks for complex information needs. Mohanna has previously taught classes in machine learning for natural language processing, artificial intelligence, and information retrieval.

Panagiotis Eustratiadis (University of Amsterdam) is a post-doctoral researcher with expertise in matters of AI safety. His research focuses on adversarial robustness, Bayesian uncertainty estimation, and information retrieval. Panagiotis has taught diverse topics in machine learning and information retrieval at a BSc and MSc level, as well as AI safety-related topics at a PhD level.

Arjen P. de Vries (Radboud University) is professor of IR. His research aims to help understand how users and systems may cooperate to improve information access, and focuses on questions regarding the combination of structured and unstructured information representations. Arjen has taught at all levels on a broad range of topics in information retrieval and data science. He has co-taught a number of tutorials at ECIR, SIGIR and ESSIR.

Maarten de Rijke (University of Amsterdam) is a Distinguished University Professor of Artificial Intelligence and Information Retrieval at the University of Amsterdam. Maarten has taught numerous courses at the BSc, MSc, and PhD level at the University of Amsterdam, numerous tutorials on a range of information retrieval topics at all the main conferences and summer schools in our field, and has co-organized multiple summer schools in information retrieval and artificial intelligence, most recently the 2025 ICAI Summer School.

1 Motivation

While dense and generative models excel at semantic matching across queries and documents, real-world information needs often demand reasoning: enforcing negation [33, 47], exclusion [55], composing sets [28], and synthesizing (conflicting) evidence across various sources remain difficult for neural retrievers despite advances [41]. Recent IR benchmarks contain queries that expose clear empirical weaknesses of current retrieval systems [7, 22, 38].

In parallel, theoretical findings indicate that embedding-based retrieval methods, despite their empirical success, are subject to intrinsic representational limits, that may not be alleviated by pure scale of model and/or data size [46]. Large language models (LLMs), though often regarded as general-purpose reasoning systems and used as an alternative to transcend such limitations through generative reasoning, nonetheless exhibit persistent weaknesses in grounding and inference control, with outputs frequently lacking evidential support [21]. Considered together, these results suggest that information retrieval is reaching the limits of pattern-matching systems. In various subfields of AI, building models that can perform *reasoning* has once again attracted significant attention.

Inference-time strategies seek to elicit structured ‘thinking’ from LLMs, through chain-of-thought and iterative refinement [26, 45, 52]. Reinforcement learning (RL) has been explored as a mechanism to induce and regulate reasoning in LLMs, by defining reward signals over intermediate derivations, aligning model trajectories with human or automated preferences, and coupling language generation with structured search and planning procedures [20, 31]. Neuro-symbolic approaches couple statistical generalization with formal inference [30]. Probabilistic and Bayesian frameworks cast reasoning as inference under uncertainty [32]. Some other line of work focuses on alternative representation spaces, using non-Euclidean embeddings to capture hierarchical structure [6] or compositional representations to model Boolean operations in queries [27]. Lastly, energy-based formulations interpret reasoning as iterative optimization in latent space [11]. These directions offer avenues with distinct strengths, but the approaches remain fragmented across disciplines. They highlight aspects of the reasoning problem, making it essential to connect them within a bigger picture of how representations, inference mechanisms, and learning signals interact.

IR both benefits from and contributes to this momentum. It provides a natural testbed where reasoning models must operate at scale, ground their steps in external evidence, and be judged by rigorous evaluation. The goal of this tutorial is to consolidate reasoning methodologies into a coherent framework for IR, and

to show how IR can serve not only as a beneficiary but also as a useful ground for increasingly robust reasoning architectures.

2 Learning Objectives

The tutorial is designed to provide a conceptual overview of and practical guidance into reasoning for IR. Specifically, participants will be able to:

- (1) Identify the requirements of complex IR tasks, including negation, exclusion, set composition, temporal constraints, and multi-hop evidence synthesis, as well as an overview of applicable benchmarks.
- (2) Understand the main methodological families for reasoning: (i) *LLM-based (including inference-time strategies, and reinforcement learning approaches)*, (ii) *neuro-symbolic integration*, (iii) *probabilistic frameworks*, (iv) *alternative representation spaces and optimization approaches*.
- (3) Compare methods along three axes: (i) *representational adequacy*, (ii) *mechanisms of inference and learning*, and (iii) *computational viability for IR*. They can consequently analyze trade-offs between approaches with respect to requirements of an IR system.
- (4) Apply the conceptual framework introduced in the tutorial to propose novel solutions for IR and recognize opportunities where IR can enhance the development of more general reasoning architectures.

3 Scope

The area to be covered is broad and touches on many sub-disciplines in IR and AI. Since the topics covered are broad and technically diverse, the tutorial will not aim for exhaustive coverage of each reasoning method and its use or potential in or with IR, but rather to provide participants with an overarching view of the challenges and the research possibilities. Slides and reference materials will be shared with participants so as to help them to continue to explore the area.

4 Relevance to IR Community

Recent IR benchmarks expose clear weaknesses of current models in handling complex user queries. Meanwhile, advances in LLM-based and neuro-symbolic, and other alternative approaches to reasoning remain scattered across fields, limiting their practical uptake in IR. This tutorial unifies these developments within a coherent analytical framework, clarifying how reasoning can be defined, modeled, and evaluated in IR and helping researchers identify transferable methods and understand their trade-offs.

5 Detailed Outline

We will teach a half-day (3 hours) tutorial, aligned with the two breaks organized by the conference. The tutorial will be organized in four main blocks:

5.1 Introduction (15 minutes)

Why reasoning is central to information retrieval; overview of tutorial’s objectives, structure, and relevance to IR community.

5.2 Reasoning: Definition and its Challenges in IR (30 minutes)

- (1) Definition of reasoning in IR.
- (2) Empirical and theoretical challenges:
 - (a) *Tasks and datasets*: NevIR [33, 41, 47]; ExcluIR [55]; QUEST [28]; BRIGHT [38]; BrowseComp-Plus [7]; complex retrieval tasks [22].
 - (b) *Limits of current systems*: theoretical limits of embedding-based retrieval [46]; taxonomy for negation [33]; LLM hallucinations [21].

5.3 Methodological Families (90 minutes)

- (1) *LLM inference-time strategies and optimization for reasoning*: Approaches that elicit or refine reasoning at LLMs' decode/test time, including chain-of-thought prompting [45], IRCOT for multi-step questions [40], Self-Refine for iterative self-feedback and revision [26], test-time adaptation/active fine-tuning [19], and implicit chain-of-thought without explicit prompts [43]. Complementary results link these behaviors to procedural knowledge acquired during pretraining of LLMs [36].
- (2) *LLMs + RL*: Reinforcement learning (RL) has been increasingly explored as a means to incentivize and enhance reasoning in LLMs. Building on instruction-following models trained with human feedback [31], recent systems such as *DeepSeek-R1* [3] and *Search-R1* [20] use structured rewards and interaction with search or planning environments to promote stepwise reasoning. In parallel, complementary analyses question whether such RL-based fine-tuning genuinely improves reasoning ability beyond the pretrained model [54], and examine the limits of language-based policy optimization [35].
- (3) *Neuro-symbolic approaches*: Early efforts to ground IR in non-classical logic [42] offered a rigorous theoretical foundation but faced scalability and practical limitations. Recent work on reasoning explores use of symbolic solvers, provers and logic engines alongside neural components to enhance the reliability of inference in reasoning-intensive tasks [30]. LINC use the LLM as a semantic parser from natural language to first-order logic, delegating deductive inference to external theorem provers [30]. Other approaches target reasoning under uncertainty by prompting LLMs to produce formal representations (e.g., code, probabilistic logic) that can be executed or queried [29], or by synthesizing task-specific probabilistic formulations on demand to support bayesian inference for open-world cognition [48, 49]. In IR, inspired by this line of work, OrLog [18] integrates LLMs as prior plausibility estimation models and ProbLog as a probabilistic query engine to resolve complex queries.
- (4) *Probabilistic and Bayesian Frameworks*: This line of work treats reasoning as inference under uncertainty, using Bayesian methods to model and propagate beliefs. Approaches include Bayesian teaching for probabilistic reasoning [34], latent-variable inference for chain-of-thought training [17], and trustworthy inference frameworks such as BIRD [13]. Other

work estimates model uncertainty via Bayesian prompt ensembles [39] or Bayesian low-rank adaptation [44], and applies uncertainty-aware adaptive guidance to reasoning [53]. Bayesian surprise has also been used for open-ended scientific discovery [2].

- (5) *Alternative representation spaces and optimization approaches*: Several works enhance the representation space for queries and documents to support explicit set/logic operations and hierarchical structure, including: BoxEmbeddings [8], Set-Compositional [23], , and hyperbolic representation spaces [6, 10, 50, 51]. Complementary, some alternative approaches model reasoning as iterative inference in latent space via energy-based optimization [11, 14].

5.4 Bridging Current Methodologies & Future Directions (45 minutes)

We first introduce our comparative framework and then map the discussed methods regarding to the framework's axes;

- (1) *Representational adequacy*; to what extent the representation space encodes logical structures, hierarchies and uncertainty about predicates across the various approaches,
- (2) *Mechanisms of inference verification and learning*; how inference and updates are applied over the representation space,
- (3) *Computational viability for IR*; to what extent each of the reasoning frameworks can be applied in a high-scale real-world IR setting.

We end with a roadmap of open problems and research directions for reasoning in IR and leveraging IR within broader reasoning models.

6 Related Past Events

Relevant past events with a reasoning theme across various venues in IR, NLP, and AI include:

- (1) EMNLP'20: Machine Reasoning: Technology, Dilemma and Future [12]
- (2) ECIR'23: Neuro-Symbolic Approaches for Information Retrieval [9]
- (3) AAAI'23: Advances in Neuro-Symbolic Reasoning and Learning [37]
- (4) KDD'23: Knowledge Graph Reasoning and Its Applications [25]

Alongside the enumerated tutorials, several existing surveys on reasoning in LLMs, RL, optimization-based formulations, geometric representations, and logic-based IR will inform the content of this tutorial [1, 4, 5, 15, 16, 24]. In contrast to prior tutorials, we define reasoning specifically within the context of IR and organize recent methodological advances into a unified analytical framework which connects developments across representation, inference and learning, and computational viability, providing a clear structure for analyzing, comparing, and developing reasoning-capable information access systems.

7 Materials

A dedicated website will contain slides, references, as well as pointers to libraries and experimental environments that help participants of our tutorial examine the area further. See <https://reasoning-for-ir.github.io>.

Acknowledgments

This research was (partially) supported by the Dutch Research Council (NWO), under project numbers 024.004.022, NWA.1389.20.-183, and KICH3.LTP.20.006, and the European Union under grant agreement No. 101201510 (UNITE). All content represents the opinion of the authors, not necessarily shared or endorsed by their respective employers and/or sponsors.

References

- [1] Karam Abdullahhad, Catherine Berrut, Jean-Pierre Chevallet, and Gabriella Pasi. 2019. Modeling Information Retrieval by Formal Logic: A Survey. *ACM Comput. Surv.* (2019).
- [2] Dhruv Agarwal, Bodhisattwa Prasad Majumder, Reece Adamson, Megha Chakraborty, Satvika Reddy Gavireddy, Aditya Parashar, Harshit Surana, Bhavana Dalvi Mishra, Andrew McCallum, Ashish Sabharwal, and Peter Clark. 2025. AutoDiscovery: Open-ended Scientific Discovery via Bayesian Surprise. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- [3] DeepSeek AI. 2025. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature* (2025).
- [4] Davide Carbone. 2025. Hitchhiker’s guide on the relation of Energy-Based Models with other generative models, sampling and statistical physics: a comprehensive review. *TMLR* (2025).
- [5] Stephen Casper, Xander Davies, Claudia Shi, et al. 2023. Open Problems and Fundamental Limitations of Reinforcement Learning from Human Feedback. *TMLR* (2023).
- [6] Weize Chen, Xu Han, Yankai Lin, et al. 2024. Hyperbolic Pre-Trained Language Model. *IEEE ACM Trans. Audio Speech Lang. Process.* (2024).
- [7] Zijian Chen, Xueguang Ma, Shengyao Zhuang, et al. 2025. BrowseComp-Plus: A More Fair and Transparent Evaluation Benchmark of Deep-Research Agent. CoRR: abs/2508.06600.
- [8] Tejas Chheda, Purujit Goyal, Trang Tran, et al. 2021. Box Embeddings: An open-source library for representation learning using geometric structures. In *EMNLP*.
- [9] Laura Dietz, Hannah Bast, Shubham Chatterjee, et al. 2023. ECIR 23 Tutorial: Neuro-Symbolic Approaches for Information Retrieval. In *ECIR*.
- [10] Tiansi Dong, Mateja Jamnik, and Pietro Liò. 2025. Neural Reasoning for Sure Through Constructing Explainable Models. *AAAI* (2025).
- [11] Yilun Du, Jiayuan Mao, and Joshua B. Tenenbaum. 2024. Learning Iterative Reasoning through Energy Diffusion. In *ICML*.
- [12] Nan Duan, Duyu Tang, and Ming Zhou. 2020. Machine Reasoning: Technology, Dilemma and Future. In *EMNLP*.
- [13] Yu Feng, Ben Zhou, Weidong Lin, and Dan Roth. 2025. BIRD: A Trustworthy Bayesian Inference Framework for Large Language Models. In *ICLR*.
- [14] Alexi Gladstone, Ganesh Nanduru, Md Mofijul Islam, et al. 2025. Energy-Based Transformers are Scalable Learners and Thinkers. CoRR: abs/2507.02092.
- [15] Peter Hase, Thomas Hofweber, Xiang Zhou, Elias Stengel-Eskin, and Mohit Bansal. 2024. Fundamental Problems With Model Editing: How Should Rational Belief Revision Work in LLMs? *Transactions on Machine Learning Research* (2024).
- [16] Neil He, Hiren Madhu, Ngoc Bui, et al. 2025. Hyperbolic Deep Learning for Foundation Models: A Survey. In *ACM SIGKDD*.
- [17] Matthew Douglas Hoffman, Du Phan, David Dohan, and Sothors. 2023. Training Chain-of-Thought via Latent-Variable Inference. In *NeurIPS*.
- [18] Mohanna Hoveyda, Jelle Piepenbrock, Arjen P de Vries, Maarten de Rijke, and Faegheh Hasibi. 2026. OrLog: Resolving Complex Queries with LLMs and Probabilistic Reasoning. <https://arxiv.org/abs/2601.23085> To appear in ECIR 2026.
- [19] Jonas Hübner, Sascha Bongni, Ido Hakimi, and Andreas Krause. 2024. Efficiently Learning at Test-Time: Active Fine-Tuning of LLMs. In *Adaptive Foundation Models: Evolving AI for Personalized and Efficient Learning*.
- [20] Bowen Jin, Hansi Zeng, Zhenrui Yue, et al. 2025. Search-R1: Training LLMs to Reason and Leverage Search Engines with Reinforcement Learning. In *COLM*.
- [21] Adam Tauman Kalai, Ofir Nachum, Santosh S. Vempala, et al. 2025. Why Language Models Hallucinate. CoRR: abs/2509.04664.
- [22] Julian Killingback and Hamed Zamani. 2025. Benchmarking Information Retrieval Models on Complex Retrieval Tasks. CoRR: abs/2509.07253.
- [23] Antonios Minas Krasakis, Andrew Yates, and Evangelos Kanoulas. 2025. Constructing Set-Compositional and Negated Representations for First-Stage Ranking. CoRR: abs/2501.07679.
- [24] Jindong Li, Yali Fu, Li Fan, et al. 2025. Implicit Reasoning in Large Language Models: A Comprehensive Survey. CoRR: abs/2509.02350.
- [25] Lihui Liu and Hanghang Tong. 2023. Knowledge Graph Reasoning and Its Applications. In *ACM SIGKDD*.
- [26] Aman Madaan, Niket Tandon, Prakhar Gupta, et al. 2023. SELF-REFINE: iterative refinement with self-feedback. In *NeurIPS*.
- [27] Quan Mai, Susan Gauch, and Douglas Adams. 2024. SetBERT: Enhancing Retrieval Performance for Boolean Logic and Set Operation Queries. In *NLPIR*.
- [28] Chaitanya Malaviya, Peter Shaw, Ming-Wei Chang, et al. 2023. QUEST: A Retrieval Dataset of Entity-Seeking Queries with Implicit Set Operations. In *ACL*.
- [29] Aliakbar Nafar, Kristen Brent Venable, and Parisa Kordjamshidi. 2025. Reasoning over uncertain text by generative large language models. In *AAAI*.
- [30] Theo Olausson, Alex Gu, Ben Lipkin, et al. 2023. LINC: A Neurosymbolic Approach for Logical Reasoning by Combining Language Models with First-Order Logic Provers. In *EMNLP*.
- [31] Long Ouyang, Jeffrey Wu, Xu Jiang, et al. 2022. Training language models to follow instructions with human feedback. *NeurIPS* (2022).
- [32] Akshay Paruchuri, Jake Garrison, Shun Liao, et al. 2024. What Are the Odds? Language Models Are Capable of Probabilistic Reasoning. In *EMNLP*.
- [33] Roxana Petcu, Samarth Bhargav, Maarten de Rijke, et al. 2025. A Comprehensive Taxonomy of Negation for NLP and Neural Retrievers. In *EMNLP*.
- [34] Linlu Qiu, Fei Sha, Kelsey Allen, Yoon Kim, Tal Linzen, and Sjoerd van Steenkiste. 2026. Bayesian teaching enables probabilistic reasoning in large language models. *Nature Communications* (2026).
- [35] Rajkumar Ramamurthy, Prithviraj Ammanabrolu, Kianté Brantley, et al. 2023. Is Reinforcement Learning (Not) for Natural Language Processing: Benchmarks, Baselines, and Building Blocks for Natural Language Policy Optimization. In *ICLR*.
- [36] Laura Ruis, Maximilian Mozes, Juhan Bae, et al. 2025. Procedural Knowledge in Pretraining Drives Reasoning in Large Language Models. In *ICLR*.
- [37] Paulo Shakarian, Gerardo I. Simari, Chitta Baral, et al. 2023. Advances in Neuro-Symbolic Reasoning and Learning: 2023 AAAI Tutorial. In *AAAI*.
- [38] Hongjin Su, Howard Yen, Mengzhou Xia, et al. 2025. BRIGHT: A Realistic and Challenging Benchmark for Reasoning-Intensive Retrieval. In *ICLR*.
- [39] Francesco Tonolini, Nikolaos Aletras, Jordan Massiah, and Gabriella Kazai. 2024. Bayesian Prompt Ensembles: Model Uncertainty Estimation for Black-Box Large Language Models. In *Findings of the Association for Computational Linguistics: ACL 2024*.
- [40] Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, et al. 2023. Interleaving Retrieval with Chain-of-Thought Reasoning for Knowledge-Intensive Multi-Step Questions. In *ACL*.
- [41] Coen van den Elsen, Francien Barkhof, Thijmen Nijdam, et al. 2025. Reproducing NevIR: Negation in Neural Information Retrieval. In *SIGIR 2025*.
- [42] C J Van Rijsbergen. 1986. A new theoretical framework for information retrieval. *SIGIR Forum* (1986).
- [43] Xuezhi Wang and Denny Zhou. 2024. Chain-of-thought reasoning without prompting. In *NeurIPS*.
- [44] Yibin Wang, Haizhou Shi, Ligong Han, Dimitris N. Metaxas, and Hao Wang. 2024. BLoB: Bayesian Low-Rank Adaptation by Backpropagation for Large Language Models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- [45] Jason Wei, Xuezhi Wang, Dale Schuurmans, et al. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *NeurIPS*.
- [46] Orion Weller, Michael Boratko, Iftekhar Naim, and Jinhyuk Lee. 2025. On the Theoretical Limitations of Embedding-Based Retrieval. CoRR: abs/2508.21038.
- [47] Orion Weller, Dawn J. Lawrie, and Benjamin Van Durme. 2024. NevIR: Negation in Neural Information Retrieval. In *EACL*.
- [48] Lionel Wong, Katherine M. Collins, Lance Ying, et al. 2025. Modeling Open-World Cognition as On-Demand Synthesis of Probabilistic Models. CoRR: abs/2507.12547.
- [49] Lionel Wong, Gabriel Grand, Alexander K. Lew, et al. 2023. From Word Models to World Models: Translating from Natural Language to the Probabilistic Language of Thought. CoRR: abs/2306.12672.
- [50] Menglin Yang, Aosong Feng, Bo Xiong, Jiahong Liu, Irwin King, and Rex Ying. 2024. Enhancing LLM Complex Reasoning Capability through Hyperbolic Geometry. In *ICML 2024 Workshop on LLMs and Cognition*.
- [51] Menglin Yang, Aosong Feng, Bo Xiong, Jihong Liu, Irwin King, and Rex Ying. 2024. Hyperbolic Fine-tuning for Large Language Models. CoRR: abs/2410.04010.
- [52] Shunyu Yao, Dian Yu, Jeffrey Zhao, et al. 2023. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. In *NeurIPS*.
- [53] Zhangyue Yin, Qishi Sun, Qipeng Guo, et al. 2024. Reasoning in flux: Enhancing large language models reasoning through uncertainty-aware adaptive guidance. In *ACL*.
- [54] Yang Yue, Zhiqi Chen, Rui Lu, et al. 2025. Does Reinforcement Learning Really Incentivize Reasoning Capacity in LLMs Beyond the Base Model?. In *AI for Math Workshop @ ICML 2025*.
- [55] Wenhao Zhang, Mengqi Zhang, Shiguang Wu, et al. 2025. ExcluIR: Exclusionary neural information retrieval. In *AAAI*.