



Tutorial on Reasoning for IR & IR for Reasoning

Mohanna Hoveyda¹(✉), Panagiotis Eustratiadis², Arjen P. de Vries¹,
and Maarten de Rijke²

¹ Radboud University, Nijmegen, Netherlands
{mohanna.hoveyda, arjen.devries}@ru.nl

² University of Amsterdam, Amsterdam, Netherlands
{p.efstratiadis, 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.

1 Motivation

While dense and generative models excel at semantic matching across queries and documents, real-world information needs often demand reasoning: enforcing negation [32, 43], exclusion [51], composing sets [27], and synthesizing (conflicting) evidence across various sources remain difficult for neural retrievers despite advances [12]. Recent IR benchmarks contain queries that expose clear empirical weaknesses of current retrieval systems [6, 21, 37].

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 [42]. 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 [20]. 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 [25, 41, 48]. 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 [19, 30]. Neuro-symbolic approaches couple statistical generalization with formal inference [29]. Probabilistic and Bayesian frameworks cast reasoning as inference under uncertainty [31]. Some other line of work focuses on alternative representation spaces, using non-Euclidean embeddings to capture hierarchical structure [5] or compositional representations to model Boolean operations in queries [26]. Lastly, energy-based formulations interpret reasoning as iterative optimization in latent space [10]. 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

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 Tutorial Format

The format is primarily lecture-style, organized in four main blocks (see the detailed outline below), but will include interactive polls and short Q&A slots after each block to encourage engagement.

6 Detailed Outline

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

6.1 Introduction (15 Minutes)

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

6.2 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 [12, 32, 43]; ExcluIR [51]; QUEST [27]; BRIGHT [37]; BrowseComp-Plus [6]; complex retrieval tasks [21].
 - (b) *Limits of current systems*: theoretical limits of embedding-based retrieval [42]; taxonomy for negation [32]; LLM hallucinations [20].

6.3 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 [41], IRCOT for multi-step questions [38], Self-Refine for iterative self-feedback and revision [25], test-time adaptation/active fine-tuning [18], and implicit chain-of-thought without explicit prompts [40]. Complementary results link these behaviors to procedural knowledge acquired during pretraining of LLMs [35].
- (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 [30], recent systems such as *DeepSeek-R1* [2] and *Search-R1* [19] use structured rewards and interaction with search or planning environments to promote step-wise reasoning. In parallel, complementary analyses question whether such RL-based fine-tuning genuinely improves reasoning ability beyond the pre-trained model [50], and examine the limits of language-based policy optimization [34].
- (3) *Neuro-symbolic approaches*: Early efforts to ground IR in non-classical logic [39] 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 [29]. LINC use the LLM as a semantic parser from natural language to first-order logic, delegating deductive inference to external theorem provers [29]. Other approaches target reasoning under uncertainty by prompting LLMs to produce formal representations (e.g., code, probabilistic logic) that can be executed or queried [28], or by synthesizing task-specific probabilistic formulations on demand to support bayesian inference for open-world cognition [44, 45].
- (4) *Probabilistic and Bayesian Frameworks*: This line of research approaches reasoning by explicitly modeling uncertainty and integrating Bayesian inference to achieve a probabilistic reasoning system [13, 17, 33, 49].
- (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: Box-Embeddings [7], Set-Compositional [22], , and hyperbolic representation

spaces [5,9,46,47]. Complementarily, some alternative approaches model reasoning as iterative inference in latent space via energy-based optimization [10,14].

6.4 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.

7 Target 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. We assume familiarity with modern IR models, but no prior expertise in formal logic or specialized representations/optimization approaches, remaining accessible to newcomers yet technically engaging for experts.

8 Tutorial History and Related Events

This tutorial is new and prepared exclusively for ECIR 2026. 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 [11]
- (2) ECIR’23: Neuro-Symbolic Approaches for Information Retrieval [8]
- (3) ACL’23: Complex Reasoning in Natural Language [52]

- (4) AAAI'23: Advances in Neuro-Symbolic Reasoning and Learning [36]
- (5) KDD'23: Knowledge Graph Reasoning and Its Applications [24]

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, 3, 4, 15, 16, 23]. 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.

9 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>.

Acknowledgements. 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 agreements No. 101070212 (FINDHR) and No. 101201510 (UNITE). All content represents the opinion of the authors, not necessarily shared or endorsed by their respective employers and/or sponsors.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Abdulahhad, K., Berrut, C., Chevallet, J., Pasi, G.: Modeling information retrieval by formal logic: a survey. *ACM Comput. Surv.* (2019)
2. AI, D.: Deepseek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature* (2025)
3. Carbone, D.: Hitchhiker’s guide on the relation of energy-based models with other generative models, sampling and statistical physics: a comprehensive review. *TMLR* (2025)
4. Casper, S., Davies, X., Shi, C., et al.: Open problems and fundamental limitations of reinforcement learning from human feedback. *TMLR* (2023)
5. Chen, W., Han, X., Lin, Y., et al.: Hyperbolic pre-trained language model. *IEEE ACM Trans. Audio Speech Lang. Process* (2024)
6. Chen, Z., Ma, X., Zhuang, S., et al.: Browsecomp-plus: a more fair and transparent evaluation benchmark of deep-research agent (2025). *CoRR* [arxiv:2508.06600](https://arxiv.org/abs/2508.06600)
7. Chheda, T., Goyal, P., Tran, T., et al.: Box embeddings: an open-source library for representation learning using geometric structures. In: *EMNLP* (2021)
8. Dietz, L., Bast, H., Chatterjee, S., et al.: Ecir 23 tutorial: neuro-symbolic approaches for information retrieval. In: *ECIR* (2023)

9. Dong, T., Jamnik, M., Liò, P.: Neural reasoning for sure through constructing explainable models. In: AAI (2025)
10. Du, Y., Mao, J., Tenenbaum, J.B.: Learning iterative reasoning through energy diffusion. In: ICML (2024)
11. Duan, N., Tang, D., Zhou, M.: Machine reasoning: technology, dilemma and future. In: EMNLP (2020)
12. van den Elsen, C., Barkhof, F., Nijdam, T., et al.: Reproducing NevIR: negation in neural information retrieval. In: SIGIR 2025 (2025)
13. Feng, Y., Zhou, B., Lin, W., Roth, D.: BIRD: a trustworthy bayesian inference framework for large language models. In: ICLR (2025)
14. Gladstone, A., Nanduru, G., Islam, M.M., et al.: Energy-based transformers are scalable learners and thinkers (2025). CoRR [arxiv:2507.02092](https://arxiv.org/abs/2507.02092)
15. Hase, P., Hofweber, T., Zhou, X., Stengel-Eskin, E., Bansal, M.: Fundamental problems with model editing: how should rational belief revision work in LLMs? *Trans. Mach. Learn. Res.* (2024)
16. He, N., Madhu, H., Bui, N., et al.: Hyperbolic deep learning for foundation models: a survey. In: ACM SIGKDD (2025)
17. Hoffman, M.D., Phan, D., Dohan, D., Sothers: training chain-of-thought via latent-variable inference. In: NeurIPS (2023)
18. Hübotter, J., Bongni, J., Hakimi, I., Krause, A.: Efficiently learning at test-time: active fine-tuning of LLMs. In: Adaptive Foundation Models: Evolving AI for Personalized and Efficient Learning (2024)
19. Jin, B., Zeng, H., Yue, Z., et al.: Search-R1: training LLMs to reason and leverage search engines with reinforcement learning. In: COLM (2025)
20. Kalai, A.T., Nachum, O., Vempala, S.S., et al.: Why language models hallucinate (2025), CoRR [arxiv:2509.04664](https://arxiv.org/abs/2509.04664)
21. Killingback, J., Zamani, H.: Benchmarking information retrieval models on complex retrieval tasks (2025). CoRR [arxiv:2509.07253](https://arxiv.org/abs/2509.07253)
22. Krasakis, A.M., Yates, A., Kanoulas, E.: Constructing set-compositional and negated representations for first-stage ranking (2025). CoRR [arxiv:2501.07679](https://arxiv.org/abs/2501.07679)
23. Li, J., Fu, Y., Fan, L., et al.: Implicit reasoning in large language models: a comprehensive survey (2025). CoRR [arxiv:2509.02350](https://arxiv.org/abs/2509.02350)
24. Liu, L., Tong, H.: Knowledge graph reasoning and its applications. In: ACM SIGKDD (2023)
25. Madaan, A., Tandon, N., Gupta, P., et al.: Self-refine: iterative refinement with self-feedback. In: NeurIPS (2023)
26. Mai, Q., Gauch, S., Adams, D.: Setbert: enhancing retrieval performance for Boolean logic and set operation queries. In: NLPPIR (2024)
27. Malaviya, C., Shaw, P., Chang, M.W., et al.: QUEST: a retrieval dataset of entity-seeking queries with implicit set operations. In: ACL (2023)
28. Nafar, A., Venable, K.B., Kordjamshidi, P.: Reasoning over uncertain text by generative large language models. In: AAI (2025)
29. Olausson, T., Gu, A., Lipkin, B., et al.: Linc: a neurosymbolic approach for logical reasoning by combining language models with first-order logic provers. In: EMNLP (2023)
30. Ouyang, L., Wu, J., Jiang, X., et al.: Training language models to follow instructions with human feedback. In: NeurIPS (2022)
31. Paruchuri, A., Garrison, J., Liao, S., et al.: What are the odds? language models are capable of probabilistic reasoning. In: EMNLP (2024)
32. Petcu, R., Bhargav, S., de Rijke, M., et al.: A comprehensive taxonomy of negation for NLP and neural retrievers. In: EMNLP (2025)

33. Qiu, L., Sha, F., Allen, K.R., et al.: Bayesian teaching enables probabilistic reasoning in large language models (2025). CoRR [arxiv:2503.17523](https://arxiv.org/abs/2503.17523)
34. Ramamurthy, R., Ammanabrolu, P., Brantley, K., et al.: Is reinforcement learning (not) for natural language processing: benchmarks, baselines, and building blocks for natural language policy optimization. In: ICLR (2023)
35. Ruis, L., Mozes, M., Bae, J., et al.: Procedural knowledge in pretraining drives reasoning in large language models. In: ICLR (2025)
36. Shakarian, P., Simari, G.I., Baral, C., et al.: Advances in neuro-symbolic reasoning and learning: 2023 AAAI tutorial. In: AAAI (2023)
37. Su, H., Yen, H., Xia, M., et al.: BRIGHT: a realistic and challenging benchmark for reasoning-intensive retrieval. In: ICLR (2025)
38. Trivedi, H., Balasubramanian, N., Khot, T., et al.: Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions. In: ACL (2023)
39. Van Rijsbergen, C.J.: A new theoretical framework for information retrieval. In: SIGIR Forum (1986)
40. Wang, X., Zhou, D.: Chain-of-thought reasoning without prompting. In: NeurIPS (2024)
41. Wei, J., Wang, X., Schuurmans, D., et al.: Chain-of-thought prompting elicits reasoning in large language models. In: NeurIPS (2022)
42. Weller, O., Boratko, M., Naim, I., Lee, J.: On the theoretical limitations of embedding-based retrieval (2025). CoRR [arxiv:2508.21038](https://arxiv.org/abs/2508.21038)
43. Weller, O., Lawrie, D.J., Durme, B.V.: Nevir: negation in neural information retrieval. In: EACL (2024)
44. Wong, L., Collins, K.M., Ying, L., et al.: Modeling open-world cognition as on-demand synthesis of probabilistic models (2025). CoRR [arxiv:2507.12547](https://arxiv.org/abs/2507.12547)
45. Wong, L., Grand, G., Lew, A.K., et al.: From word models to world models: translating from natural language to the probabilistic language of thought (2023). CoRR [arxiv:2306.12672](https://arxiv.org/abs/2306.12672)
46. Yang, M., Feng, A., Xiong, B., Liu, J., King, I., Ying, R.: Enhancing LLM complex reasoning capability through hyperbolic geometry. In: ICML 2024 Workshop on LLMs and Cognition (2024)
47. Yang, M., Feng, A., Xiong, B., Liu, J., King, I., Ying, R.: Hyperbolic fine-tuning for large language models (2024). CoRR [arxiv:2410.04010](https://arxiv.org/abs/2410.04010)
48. Yao, S., Yu, D., Zhao, J., et al.: Tree of thoughts: deliberate problem solving with large language models. In: NeurIPS (2023)
49. Yin, Z., Sun, Q., Guo, Q., et al.: Reasoning in flux: enhancing large language models reasoning through uncertainty-aware adaptive guidance. In: ACL (2024)
50. Yue, Y., Chen, Z., Lu, R., et al.: Does reinforcement learning really incentivize reasoning capacity in LLMs beyond the base model? In: AI for Math Workshop @ ICML 2025 (2025)
51. Zhang, W., Zhang, M., Wu, S., et al.: ExcluIR: exclusionary neural information retrieval. In: AAAI (2025)
52. Zhao, W., Geva, M., Lin, B.Y., et al.: Complex reasoning in natural language. In: ACL (2023)