

On the Role of Information Retrieval When Teaching Artificial Intelligence: The What, the Why, and the How

Gianluca Demartini¹, Stefano Mizzaro², Kevin Roitero² and Damiano Spina^{3,*}

¹The University of Queensland, Brisbane, Australia

²University of Udine, Udine, Italy

³RMIT University, Melbourne, Australia

Abstract

Artificial Intelligence (AI) and Information Retrieval (IR) have historically been taught as separate subjects. In this position paper, we argue that IR concepts and methods are fundamental to increasing literacy in AI, and that IR should be incorporated (more) into AI university curricula and AI education in general. We start with an analysis of the current situation (“what”); then, building on recent calls from the IR community, we discuss conceptual reasons (“why”); finally, we propose concrete teaching strategies (“how”).

Keywords

AI Literacy, Information Retrieval, AI Curriculum

1. Introduction

We discuss the position that Information Retrieval (IR) is a foundational discipline that should be more considered in Artificial Intelligence (AI) education. IR is the discipline studying how to find relevant information within a collection of data, often unstructured or semi-structured, to satisfy a user’s information need [1, 2, 3, 4]. While IR did not start with the Web (IR has a long history and an age comparable to AI), it was fundamental for the success of Web search engines [5]. AI is a complex field, and its definition is a challenge; in this paper, we use AI in its broad, contemporary sense, encompassing the wider ecosystem of research on intelligent information systems, including IR, Information Access, Recommender Systems, and Natural Language Processing (NLP).

There are several reasons for the position “teach more IR in AI education” discussed in this paper; we briefly mention some examples here, and motivate and discuss them further in the following. It is natural to consider IR as an application of AI, and even the canonical AI textbook by Russell and Norvig briefly mentions IR as an NLP application [6, p. 901]. It might even be maintained that IR techniques and search engines are among the most successful applications of AI used in daily life. Going further, Cristianini claims that “Recommender systems are better representatives of AI agents than theorem provers” [7, p. 16], and it is well known that recommender systems and IR are closely related [8]. Focusing on more recent work, Large Language Models (LLMs) and Retrieval-Augmented Generation (RAG) [9] suggest that IR and AI are deeply connected. Zhai [10] argues that IR can also be viewed as a foundation of AI when intelligent agents need to seek out information and that IR is essential to achieving AGI (Artificial General Intelligence), and that the boundaries between IR and AI are blurring, especially with the pursuit of AGI. Demartini et al. [11] show that not only has IR faced many challenges now re-emerging in AI, but IR methodologies and lessons learned over more than 30 years can help AI researchers avoid “reinventing the wheel” or repeating mistakes. However, many AI curricula treat IR as a separate discipline or omit it entirely.

edu4AI: 2nd Workshop on Education for Artificial Intelligence, October 26, 2025, Bologna, Italy

*All authors are corresponding authors.

✉ demartini@acm.org (G. Demartini); mizzaro@uniud.it (S. Mizzaro); kevin.roitero@uniud.it (K. Roitero); damiano.spina@rmit.edu.au (D. Spina)

ORCID 0000-0002-7311-3693 (G. Demartini); 0000-0002-2852-168X (S. Mizzaro); 0000-0002-9191-3280 (K. Roitero); 0000-0001-9913-433X (D. Spina)



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

To summarize, it seems fair to say that the separation between AI and IR is unjustified, or at least deserves some discussion. This paper provides some supporting arguments to the position that IR should be an important component of AI education. We follow three steps: (i) we summarize the current status of (non) teaching IR in AI curricula (the “what”, Section 2); (ii) we discuss some reasons that call for a stronger presence of IR concepts in AI education (the “why”, Section 3); and (iii) we outline some concrete recommendations for integrating IR into AI education (the “how”, Section 4).

2. What: IR Is Underrepresented When Teaching AI

The statement “IR is not much considered in AI courses” is unlikely to cause any disagreement. Of course there are exceptions, and indeed IR is taught for example in AI curricula at the Bachelor’s of science in AI at the Mohamed bin Zayed University of Artificial Intelligence (MBZUAI),¹ but it might be argued that this is not by any means the most common scenario. Indeed Russell and Norvig’s book [6] only briefly mentions IR on two pages (901, 905). Another evidence is the curricular guidelines jointly developed by ACM, IEEE Computer Society, and AAAI, known as CS2023, that define a set of knowledge areas and competencies for undergraduate computer science programs worldwide [12]. In it, AI is recognized as a core knowledge area, with emphasis on its growing importance, ethical implications, and the impact of Generative AI (GenAI) on computing education [13], whereas IR is notably absent as a standalone knowledge area. When it does appear, IR is often subsumed under related areas, without explicit competency goals related to core IR topics like indexing, ranking, retrieval algorithms, evaluation metrics and methodologies, and other advanced topics [14, 15, 4]. These omissions underscores our key argument, that we discuss in the following: despite its critical role in both historical and modern AI systems, IR is underrepresented in AI curricula.

3. Why: Reasons to Increase IR Content in AI education

We discuss four reasons that, we believe, call for including more IR in AI curricula.

3.1. IR as a Foundation for AGI

A compelling case for incorporating IR into AI education lies in IR’s foundational role in the development of AGI. Zhai [10] has recently argued that intelligent agents cannot reach human-level cognition without the ability to seek and use information effectively. This perspective suggest that even state-of-the-art AI models, limited by the knowledge encoded during training, must master retrieval to acquire new information, verify facts, reason through problems, and learn continuously in dynamic environments. Future AI systems will require not only static knowledge but also the capability to query, update, and contextualize it in real time. Zhai introduces five IR tasks that general AI systems must perform: (1) *External IR*, the ability to query previously unseen sources like the Web; (2) *Provenance IR*, retrieving the origin of a fact to verify its accuracy; (3) *Curriculum IR*, selecting the most useful information to learn from next; (4) *Rule IR*, accessing internal rules or procedures for reasoning; and (5) *Scenario IR*, drawing on past cases to inform current decisions. They suggest that an AGI will resemble a scholar or investigator—consulting both internal memory and external resources—rather than a self-contained, all-knowing neural network. This conceptualization implies that building advanced AI requires a deep understanding of IR fundamentals such as indexing, querying, ranking, and feedback mechanisms.

3.2. IR-AI Similarities and Lessons Learned

A recent paper by Demartini et al. [11] illustrates a different, more historical, viewpoint. Demartini et al. start by observing that the situation in today AI, with the availability of LLMs to the large public, is very similar to what happened at the end of the 1990s when Web search engines made IR available to web

¹<https://mbzuai.ac.ae/study/undergraduate-program/> [Accessed: 19-10-2025]

users. Simplifying, their position can be summarized as “ChatGPT today is very much alike Google in the 1990s”, but their analysis goes on by listing similarities and differences between the two situations, and the AI and IR fields in general. One example concerns evaluation methodologies, that have been studied in depth in the IR field since the 1960s. Today practice of benchmark-based evaluation in AI can be improved by “importing” in AI several results and best practices from IR—indeed, debates on benchmark reliability, experimental reproducibility with proprietary systems and data, handling biased data, privacy and copyright concerns, and ensuring efficiency and scalability were hot topics in IR 25 years ago and are now resurfacing in AI. Another example is prompt optimization, that can benefit from the large body of work on query formulation, reformulation, and variations. At the same time, differences in community culture and scale exist, e.g., the IR research community has historically been smaller and more focused, whereas AI spans many sub-disciplines and has grown rapidly.

The paper concludes by listing seven key lessons that remain highly relevant for the AI field at large and can offer valuable guidance for AI education: 1) the importance of rigorous evaluation practices, including standardized benchmarks and critical awareness of metrics; 2) the recognition that real-world impact depends on full-system design and attention to user needs, not just algorithmic novelty; 3) the value of historical context in understanding the evolution of techniques and avoiding past mistakes; 4) the need to support fundamental research with long-term vision, even when immediate commercial payoff is lacking; 5) the importance of clear terminology and precise communication to avoid confusion and misaligned expectations; 6) the benefits of a collaborative research culture that fosters shared progress through open data and communal challenges; and 7) the broader scientific virtues that IR has long cultivated—rigor, humility, and cooperation—that AI students should internalize.

3.3. IR-AI Convergence

Although AI and IR originated as distinct fields, their evolution has been tightly interconnected. IR has long been driven by AI techniques. Modern AI systems frequently incorporate IR components: for instance, search engines are used to feed real-time knowledge to LLMs in RAG [9, 16]. Also the two above described works point in this direction: Zhai [10] notes the “rapid growth of work on applying new AI technologies such as LLMs to improve IR systems or develop more intelligent IR”, as well as using IR to support AI; Demartini et al. [11] describe the repeated historical development. We believe we are witnessing a convergence. This convergence can help us better characterize the relationship between human and machine intelligence [7]. AI systems are shifting from primarily *emulating* intelligence (e.g., automatically classifying objects in an image) toward *augmenting* intelligence (e.g., assisting software developers in coding more cost-effectively). Looking ahead, IR-enhanced AI tools may further extend this trajectory by enabling users to accomplish informational tasks that were previously too complex or resource-intensive to undertake [17]. In short, we argue that IR and AI are co-evolving: AI makes IR smarter, and IR makes AI more complete. For students, researchers, and practitioners, fluency in IR is becoming just as important as knowledge of ML or neural networks when working in AI.

3.4. Alignment with ACM/IEEE-CS/AAAI CS2023 Curricular Guidelines

Integrating IR into AI education would allow universities to better adhere to CS2023’s competency-based approach [12], by ensuring that students not only develop algorithmic and modeling skills, but also gain the ability to retrieve, filter, and evaluate information. Furthermore, IR aligns with the cross-cutting competencies emphasized in CS2023, such as transparency, explainability, and ethical computing: IR concepts such as ranking fairness, user intent modeling, and relevance evaluation directly support interpretability and responsible AI practices. We believe that the inclusion of IR topics in AI education would not only close a curricular gap and better reflect the IR-AI convergence, but also prepare students for building systems that search, reason, and respond in a responsible and informed way.

4. How: Integrating IR into AI Education

We now outline concrete strategies to integrate IR into AI education, targeting various learner groups.

4.1. Undergraduate and Graduate Curricula

A first step to incorporate IR in AI programs at the university level is adding a dedicated IR module within standard AI, ML, or data science classes. This module can introduce fundamental IR concepts like indexing, search algorithms, ranking models, and evaluation methodologies and metrics [15], showing how they complement AI topics. For example, when teaching ML evaluation, instructors can also cover ranking effectiveness from IR evaluation; a lesson on embedding vectors in an NLP class can naturally lead into vector-space retrieval. By aligning IR topics with existing syllabus sections (e.g., introducing inverted indexes during data structures for AI, or ranking loss functions during an ML course), students see IR as an integrated component rather than an isolated subject. This can be complemented with hands-on laboratory activities and assignments. In undergraduate courses, students might build a simple search engine for a mini-project for example by indexing a document collection and implementing a basic ranking function (such as TF-IDF or BM25). Another effective and natural exercise can be carried out in a RAG project: students use an open-source IR toolkit to retrieve relevant documents for a given query, then feed those documents into an LLM model to produce an answer. Such an assignment concretely links IR with modern AI by showing how retrieved context can improve an LLM's responses.

Graduate-level curricula can go further, encouraging interdisciplinary projects that bridge IR and AI like, for example, designing a question-answering system that combines a neural search module with a fine-tuned transformer reader, or evaluating how different retrieval strategies impact the outputs of a chatbot system. Reading seminars and special topics courses could include foundational IR literature alongside AI papers, enabling discussion on how search and learning intersect and are connected.

We believe the key is to normalize IR topics within the AI education pipeline. When IR is presented not as an outlier but as a natural part of AI problem-solving pipeline, students will recognize its value.

4.2. AI Literacy for General Learners

Integrating IR into AI education should not be limited to university students; it is equally important for general AI literacy programs aimed at secondary schools, informal learning, and the public. As AI systems, search engines, and digital assistants are present in everyday life, non-technical learners will benefit from understanding IR principles behind these tools [18]. We recommend weaving basic IR concepts into high school computing courses, library and information literacy workshops, and public AI seminars. The goal at this level is not to explain complex algorithms, but to build intuition and awareness about how information is indexed, retrieved, and presented by AI-driven services.

One approach towards this goal is to use real-world examples. Educators can describe a web search engine by explaining its core components: web crawlers, inverted indexes, and ranking algorithms. A class of teenagers might engage in a hunt activity where they formulate queries to find specific information, then discuss why certain results appeared on top introducing ideas like keywords, relevance, and ranking criteria. This can be followed by a discussion of biases in search results and the importance of evaluating information sources. Since many modern AI applications (from conversational assistants to recommender systems) involve retrieval, highlighting these connections helps learners see IR as an integral part of AI. For example, an instructor might show how a virtual assistant's answer to "Who is the president of the United States?" comes from retrieving information from an up-to-date knowledge base or the web, not just the assistant "knowing" the fact. This reinforces a healthy model of AI: even advanced AI often looks things up and thus the quality of its answers depends on the quality of retrieval.

Project-based learning can also be effective for non-technical audiences. Secondary school students could be guided to build a small search engine for a set of documents relevant to a history or science project using simple tools or pre-written code, illustrating IR in a hands-on way. Alternatively, a museum or library could host an interactive demo where participants experiment with a question-answering

system that uses a local archive: they pose questions and see how the system retrieves documents to answer them. Such activities would make IR tangible and enhance AI literacy [18]. They will also foster critical thinking; for instance, participants can learn how rephrasing a query can yield better results (the IR notions of query formulation and reformulation [19]).

4.3. Professional and Continuing Education

The need for IR integration extends to professionals and lifelong learners. Many industry roles now require blending data-driven AI techniques with search and knowledge management. We suggest that professional development programs, such as online certifications and corporate training courses, embed IR modules focused on practical applications. For example, a data science certificate program might add a unit on “AI-powered information retrieval” covering how to use search indexes or vector databases alongside ML models. Likewise, workshops for software engineers can include labs on building a retrieval-augmented chatbot that leverages search algorithms.

A particularly relevant theme for continuing education is RAG. Modern enterprises are adopting RAG pipelines that combine LLM reasoning with domain-specific knowledge bases and search capabilities to improve accuracy [9]. Therefore, training courses for AI professionals should teach how to integrate open-source IR systems or enterprise search platforms with AI models. A corporate workshop might have teams deploy a simple question-answering service: using an IR API to fetch internal documents in response to user questions, and then summarizing those with an LLM. Through this exercise, professionals learn about indexing data, crafting effective queries for the domain, and addressing issues like document relevancy and result filtering, all IR skills directly augmenting AI solutions. Case studies from industry can be featured as well. For instance, learning materials can highlight how e-commerce companies use retrieval algorithms to feed recommender systems, or how healthcare AI systems retrieve patient literature to support clinical decision-making. By seeing these examples, practitioners recognize IR as a critical component in real-world AI pipelines, not just an academic topic.

Organizations and professional societies might develop badges in “IR for AI”, to be earned by completing some significant activity. For those already in IR-centric jobs (like search engineers, information architects), continuing education can introduce cutting-edge AI methods.

4.4. Making IR-AI Integration the Norm

Across all these educational contexts, the integration of IR into AI teaching should be approached as a long-term cultural shift. The IR community has observed that its contributions need to be better communicated and taught broadly. Universities producing AI graduates, outreach programs informing the public, and industry efforts all have a role to play. We propose that academic departments and curriculum committees explicitly include IR outcomes in their program objectives for AI-related degrees. Likewise, national AI literacy initiatives should reference IR as a foundational competency. A more IR-informed AI education will produce practitioners who can design systems that search, reason, and then respond, instead of treating retrieval as an afterthought. In turn, this will lead to AI applications that are more transparent, accountable, and effective at meeting users’ needs.

5. Conclusions and Future Developments

In this position paper we have discussed our opinion that as AI continues to advance toward more general and human-like intelligence, the skills and insights from IR are more relevant than ever. IR has a number of methodologies that can help significantly with the development of modern AI (e.g., managing raw and unstructured data at scale), it is a core component of intelligent systems, from search engines that encapsulate collective knowledge to AI assistants that must fetch facts and reason over retrieved evidence. The message brought by recent research is clear: AI needs IR. AI systems enhanced with IR capabilities are already proving more powerful and trustworthy, whether by providing citations in a chatbot answer or by staying updated with the latest information.

For these reasons, we call on educators and curriculum developers to bridge the gap between IR and AI. The timing is crucial: as AI and IR researchers increasingly collaborate and publish in each other's venues, the academic distinctions will be faded. The next generation of AI specialists should comfortably speak the language of both fields. Integrating IR into AI education will ensure that those who build and use intelligent machines remember the importance of finding information and learning from it, an ability at the very heart of intelligence itself.

Declaration on Generative AI. During the preparation of this work, the authors used ChatGPT in order to: Grammar and spelling check, Paraphrase and reword. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

- [1] R. Baeza-Yates, B. A. Ribeiro-Neto, *Modern Information Retrieval – The Concepts and Technology Behind Search*, 2 ed., Pearson Education Ltd., Harlow, England, 2011.
- [2] B. Croft, D. Metzler, T. Strohman, *Search Engines: Information Retrieval in Practice*, Pearson, 2010.
- [3] C. D. Manning, P. Raghavan, H. Schütze, *Introduction to Information Retrieval*, Cambridge University Press, 2008.
- [4] O. Alonso, R. Baeza-Yates, *Information Retrieval: Advanced Topics and Techniques*, volume 60, 1 ed., ACM, 2024.
- [5] M. Sanderson, W. B. Croft, *The History of Information Retrieval Research*, *Proc. IEEE* 100 (2012).
- [6] S. Russell, P. Norvig, *Artificial Intelligence: A Modern Approach*, 4 ed., Pearson, 2022.
- [7] N. Cristianini, *The Shortcut: Why Intelligent Machines Do Not Think Like Us*, CRC Press, 2023.
- [8] N. J. Belkin, W. B. Croft, *Information Filtering and Information Retrieval: Two Sides of the Same Coin?*, *Commun. ACM* 35 (1992) 29–38. doi:10.1145/138859.138861.
- [9] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, D. Kiela, *Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks*, in: *Advances in Neural Information Processing Systems*, volume 33, 2020, pp. 9459–9474.
- [10] C. Zhai, *Information Retrieval for Artificial General Intelligence: A New Perspective of Information Retrieval Research*, in: *Proceedings of the 48th ACM SIGIR Conference*, 2025, p. 3876–3886.
- [11] G. Demartini, C. Hauff, M. Lease, S. Mizzaro, K. Roitero, M. Sanderson, F. Scholer, C. Shah, D. Spina, P. Thomas, A. P. de Vries, G. Zuccon, *Preaching to the ChoIR: Lessons IR Should Share with AI*, in: *Proceedings of the 2025 ACM SIGIR ICTIR Conference*, 2025, p. 78–91.
- [12] A. N. Kumar, R. K. Raj, et al., *Computer Science Curricula 2023*, 2024.
- [13] IEEE Computer Society, IEEE Computer Society and Industry Partners Release Guidance in the Form of Undergraduate Computer Science Curricula, 2023. URL: <https://www.computer.org/press-room/new-cs2023-curriculum-guide>, accessed October 2025.
- [14] I. Markov, M. de Rijke, *What Should We Teach in Information Retrieval?*, *SIGIR Forum* 52 (2019).
- [15] C. Bauer, M. Fröbe, D. Jannach, U. Kruschwitz, P. Rosso, D. Spina, N. Tintarev, *Overcoming Methodological Challenges in Information Retrieval and Recommender Systems through Awareness and Education*, Technical Report, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023.
- [16] A. Salemi, H. Zamani, *Evaluating Retrieval Quality in Retrieval-Augmented Generation*, in: *Proceedings of the 47th ACM SIGIR Conference*, 2024, p. 2395–2400. doi:10.1145/3626772.3657957.
- [17] J. R. Trippas, D. Spina, F. Scholer, *Adapting Generative Information Retrieval Systems to Users, Tasks, and Scenarios*, in: R. W. White, C. Shah (Eds.), *Information Access in the Era of Generative AI*, volume 51, Springer Nature Switzerland AG, Cham, Switzerland, 2025.
- [18] A. Vercelli, P. B. D'Amico, K. Bianculli, *Librarians and AI Literacy: Strengthen Awareness and AI Literacy in Argentine and Latin American Libraries*, in: *Proceedings of the 1st Workshop on Education for Artificial Intelligence, edu4AI 2024, CEUR Workshop Proceedings*, 2024, pp. 106–115.
- [19] J. Huang, E. N. Efthimiadis, *Analyzing and Evaluating Query Reformulation Strategies in Web Search Logs*, in: *Proceedings of the 2009 ACM CIKM*, 2009, p. 77–86.