

Learned Sparse Representations of Text: A Story of Indexing and Retrieval

Keynote

Franco Maria Nardini¹

¹*Istituto di Scienza e Tecnologie dell'Informazione "Alessandro Faedo" del Consiglio Nazionale delle Ricerche (ISTI-CNR), Pisa, Italy*

Abstract

In recent years, transformer-based large language models (LLMs) have fundamentally reshaped the way large textual collections are indexed and retrieved. A key driver of this transformation is the use of LLMs to learn high-dimensional, contextual sparse representations of input text. These representations are rapidly gaining popularity for several reasons: 1) they perform competitively with learned dense representations, 2) they are grounded in the LLM's vocabulary, enabling interpretability by design, and 3) they can be efficiently leveraged with a well-established data structure – the inverted index – to support fast maximum inner product search. In this talk, we will review recent advancements that enable efficient indexing and retrieval based on these sparse representations. We will then discuss current limitations and emerging challenges in this rapidly evolving area.

Short bio

Franco Maria Nardini is a Senior Researcher with ISTI-CNR in Pisa, Italy. His research interests focus on Web Information Retrieval and Machine/Deep Learning. He authored over 100 papers in peer-reviewed international journals, conferences, and other venues. He has served as General Co-Chair of ECIR 2025, Program Committee Co-Chair of SPIRE 2023, and Tutorial Co-Chair of ACM WSDM 2021. He is a co-recipient of the ECIR 2025 Best Student Short Paper Award, the ACM SIGIR 2024 Best Paper Runner-Up Award, the ECIR 2022 Industry Impact Award, the ACM SIGIR 2015 Best Paper Award, and the ECIR 2014 Best Demo Paper Award. He coordinated activities in several EU and IT research projects. He is a member of the editorial board of ACM TOIS and a PC member of SIGIR, ECIR, SIGKDD, CIKM, WSDM, IJCAI, and ECML-PKDD. He currently teaches "Information Retrieval" in the Computer Science and AI Master's Degrees of the University of Pisa.

Declaration on Generative AI

During the preparation of this work, the author did not use any AI tool.

IIR2025: 15th Italian Information Retrieval Workshop, 3rd-5th September 2025, Cagliari, Italy

✉ francomaria.nardini@isti.cnr.it (F. M. Nardini)

🌐 <http://hpc.isti.cnr.it/~nardini/> (F. M. Nardini)

🆔 0000-0003-3183-334X (F. M. Nardini)



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