

# Giacomo Frisoni



CONTACTS      Via dell'Università, 50  
47522 Cesena (FC)  
Italy

✉ giacomo.frisoni@unibo.it  
🏠 Academic Website  
🆔 0000-0003-1104-2014  
📖 DBLP  
👤 Google Scholar  
📄 Scopus  
🌐 LinkedIn

CITIZENSHIP      🇮🇹 Italy

LANGUAGES      🇮🇹 Italian (Mother language)  
🇬🇧 English (C1 level)

CURRENT PLACEMENT      **University of Bologna**  
**Department of Computer Science and Engineering**  
Postdoctoral Researcher in Natural Language Processing

RESEARCH THEMES      Natural Language Processing, Large Language Models, Graph Neural Networks, Clinical and Biomedical Applications

BIO      Giacomo Frisoni, Ph.D., is a postdoctoral researcher in Computer Science and Engineering at the University of Bologna. His research interests lie in Natural Language Processing, Large Language Models, and Graph Neural Networks, particularly for medical applications. He has co-authored over 30 papers in prestigious conferences and journals, including ACL, NAACL, EMNLP, and AAAI, receiving best paper and reviewer awards. He has actively contributed to the research community by serving in various scientific roles, including editorial board member, program committee member, area chair, and workshop organizer for many international venues. He has collaborated with esteemed academic and industrial institutions, including the University of Glasgow, IBM Research, EURECOM, and the University of Edinburgh. He served as an adjunct professor and tutor for multiple AI courses at the University of Bologna, co-supervising over 55 theses.

## MAIN BIBLIOMETRICS

PUBLICATIONS		CITATION METRICS	
	Count	Google Scholar	Scopus
Journal articles	12★	613★	408
Conference articles	15	15★	12

★ **ASN thresholds** (Sector 09/H1, Associate Professor, call 2023–2025): journal articles  $\geq 5$ , citations  $\geq 126$ , H-index  $\geq 7$ . The *Abilitazione Scientifica Nazionale* (ASN) was the Italian national habilitation for university professorships; it was formally abolished by Law n.79 of 5 June 2025, and these thresholds are reported solely as a recognized reference point.

---

## EDUCATION

---

OCT 2020–JAN 2024  
CESENA, ITALY

### **Ph.D., Computer Science and Engineering**

Department of Computer Science and Engineering, University of Bologna

- Ministerial scholarship; ranked first out of 132 candidates in the competitive call
- Supervisor: Gianluca Moro
- Tutors: Antonella Carbonaro and Ilaria Bartolini
- Thesis topic: Knowledge-Enhanced Natural Language Processing
- Graduation score: 5/5 (Excellent)
- Areas of study: Large Language Models, Graph Neural Networks, Retrieval-Augmented Generation, Information Extraction, Semantic Parsing

SEP 2017–MAR 2020  
CESENA, ITALY

### **M.S., Computer Science and Engineering**

University of Bologna

- Cumulative GPA: 4.0 (transcript)
- Graduation score: 110/110 cum Laude
- Graduation class: LM-18, Computer Science
- Thesis topic: A New Unsupervised Methodology of Descriptive Text Mining for Knowledge Graph Learning
- Supervisor: Gianluca Moro
- Co-Supervisor: Antonella Carbonaro
- Area of study: Text Mining

SEP 2014–OCT 2017  
CESENA, ITALY

### **B.S., Computer Science and Engineering**

University of Bologna

- Cumulative GPA: 4.0 (transcript)
- Graduation score: 110/110 cum Laude
- Graduation class: L-8, Information Engineering
- Thesis topic: Design and development of a software system for studying and researching rare diseases
- Supervisor: Dario Maio
- Area of study: Databases

SEP 2009–JUL 2014  
RIMINI, ITALY

### **Scientific High School**

ITIS Leonardo Da Vinci

- Final score: 100/100 cum Laude
- Member of the National Register of Excellence
- Focus on C# and Cryptography (Bletchley Park Visitor)

---

## PUBLICATIONS

---

For conferences and journals, the venue ranking is reported inline with each entry: **ICORE 2026** for conferences (together with the acceptance rate), and **Scimago 2026** quartile for journals. This convention is adopted throughout all sections of this document. Within each category, entries are listed in descending chronological order.

CONFERENCES  
×15

- 
- [1] [Graph-of-Mark: Promote Spatial Reasoning in Multimodal Language Models with Graph-Based Visual Prompting](#)  
[G. Frisoni](#), L. Molfetta, M. Buzzoni, G. Moro (2026). *Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 40, No. 36, pp. 30726-30734)*. **A\***, **18%**.

---

  - [2] [PORTS: Preference-Optimized Retrievers for Tool Selection with Large Language Models](#)  
L. Molfetta, [G. Frisoni](#), N. Monaldini, G. Moro (2025). *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (pp. 10018-10041)*. **A\***, **22%**.

---

  - [3] [ZeroNER: Fueling Zero-Shot Named Entity Recognition via Entity Type Descriptions](#)  
A. Cocchieri, M. M. Galindo, [G. Frisoni](#), G. Moro, C. Sartori, G. Tagliavini (2025). *Findings of the Association for Computational Linguistics: ACL 2025 (pp. 15594-15616)*. **A\***, **20%**.

---

  - [4] [OpenBioNER: Lightweight Open-Domain Biomedical Named Entity Recognition Through Entity Type Description](#)  
A. Cocchieri, [G. Frisoni](#), M. M. Galindo, G. Moro, G. Tagliavini, F. Candoli (2025). *Findings of the Association for Computational Linguistics: NAACL 2025 (pp. 818-837)*. **A\***, **22%**.

---

  - [5] [FEAST: Retrieval-Augmented Multi-Hierarchical Food Classification for the FoodEx2 System](#)  
L. Molfetta, A. Cocchieri, S. Fantazzini, [G. Frisoni](#), L. Ragazzi, G. Moro (2025). *ECAI 2025 (pp. 4169-4176)*. *IOS Press*. **A**, **23%**.

---

  - [6] [Magic Mirror on the Wall, Which Is the Fairest Prompt of All? A Survey on Automatic Prompt Learning](#)  
S. Fantazzini, [G. Frisoni](#), G. Moro, L. Ragazzi, M. Ciccioni, C. Sartori (2025). *ECAI 2025 (pp. 4444-4451)*. *IOS Press*. **A**, **23%**.

---

  - [7] [Predicting Protein Functions with Ensemble Deep Learning and Protein Language Models](#)  
[G. Frisoni](#), M. Fuschi, G. Moro (2025). *ECAI 2025 (pp. 5139-5142)*. *IOS Press*. **A**, **32%**.

---

  - [8] [Unknown Claims: Generation of Fact-Checking Training Examples from Unstructured and Structured Data](#)  
J. F. Bussotti, L. Ragazzi, [G. Frisoni](#), G. Moro, P. Papotti (2024). *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (pp. 12105-12122)*. **A\***, **21%**.

---

  - [9] [To Generate or to Retrieve? On the Effectiveness of Artificial Contexts for Medical Open-Domain Question Answering](#)  
[G. Frisoni](#), A. Cocchieri, A. Presepi, G. Moro, Z. Meng (2024). *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) (pp. 9878-9919)*. **A\***, **21%**.

---

  - [10] [Cogito Ergo Summ: Abstractive Summarization of Biomedical Papers via Semantic Parsing Graphs and Consistency Rewards](#)  
[G. Frisoni](#), P. Italiani, S. Salvatori, G. Moro (2023). *Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 37, No. 11, pp. 12781-12789)*. **A\***, **20%**.

---

  - [11] [BioReader: a Retrieval-Enhanced Text-to-Text Transformer for Biomedical Literature](#)  
[G. Frisoni](#), M. Mizutani, G. Moro, L. Valgimigli (2022). *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (pp. 5770-5793)*. **A\***, **22%**.
-

- [12] [Enhancing Biomedical Scientific Reviews Summarization with Graph-Based Factual Evidence Extracted from Papers](#)  
G. Frisoni, P. Italiani, F. Boschi, G. Moro (2022). *Proceedings of the 11th International Conference on Data Science, Technology and Applications* (pp. 168-179). **C**, **20%**.
- 
- [13] [Text-to-Text Extraction and Verbalization of Biomedical Event Graphs](#)  
G. Frisoni, G. Moro, L. Balzani (2022). *Proceedings of the 29th International Conference on Computational Linguistics* (pp. 2692-2710). **B**, **28%**.
- 
- [14] [Unsupervised Descriptive Text Mining for Knowledge Graph Learning](#)  
G. Frisoni, G. Moro, A. Carbonaro (2020). *Proceedings of the 12th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (Vol. 1, pp. 316-324)*. **C**, **21%**.
- 
- [15] [Learning Interpretable and Statistically Significant Knowledge from Unlabeled Corpora of Social Text Messages: A Novel Methodology of Descriptive Text Mining](#)  
G. Frisoni, G. Moro, A. Carbonaro (2020). *Proceedings of the 9th International Conference on Data Science, Technology and Applications* (pp. 121-132). **C**, **14%**.
- 
- JOURNALS  
×12
- [16] [COMMA: A Multi-Task and Multi-Lingual Dataset of Constitutional Verdicts](#)  
L. Ragazzi, G. Frisoni, G. Moro, P. Italiani, L. Molfetta, V. Folin (2026). *Artificial Intelligence and Law, 1-55*. **Q1**.
- 
- [17] [OpenBioNER-v2: A Suite of Lightweight Models for Zero-Shot Medical Named Entity Recognition via Type Descriptions](#)  
A. Cocchieri, G. Frisoni, F. Zangrillo, L. Ragazzi, M. M. Galindo, G. Tagliavini, G. Moro (2026). *Expert Systems with Applications, 131725*. **Q1**.
- 
- [18] [LAWSUIT: A Large Expert-Written Summarization Dataset of Italian Constitutional Court Verdicts](#)  
L. Ragazzi, G. Moro, S. Guidi, G. Frisoni (2025). *Artificial Intelligence and Law, 33(4), 1151-1187*. **Q1**.
- 
- [19] [Abstractive Summarization through the Prism of Decoding Strategies](#)  
G. Frisoni, L. Ragazzi, D. Cohen, G. Moro, A. Carbonaro, C. Sartori (2025). *Neural Networks, 108249*. **Q1**.
- 
- [20] [Evidence, my Dear Watson: Abstractive Dialogue Summarization on Learnable Relevant Utterances](#)  
P. Italiani, G. Frisoni, G. Moro, A. Carbonaro, C. Sartori (2024). *Neurocomputing, 572, 127132*. **Q1**.
- 
- [21] [Graph-Enhanced Biomedical Abstractive Summarization via Factual Evidence Extraction](#)  
G. Frisoni, P. Italiani, G. Moro, I. Bartolini, M. A. Boschetti, A. Carbonaro (2023). *SN Computer Science, 4(5), 500; Revised Version of Best Student Paper Award*. **Q2**.
- 
- [22] [Efficient Text-Image Semantic Search: A Multi-Modal Vision-Language Approach for Fashion Retrieval](#)  
G. Moro, S. Salvatori, G. Frisoni (2023). *Neurocomputing, 538, 126196*. **Q1**.
- 
- [23] [Efficient Memory-Enhanced Transformer for Long-Document Summarization in Low-Resource Regimes](#)  
G. Moro, L. Ragazzi, L. Valgimigli, G. Frisoni, C. Sartori, G. Marfia (2023). *Sensors, 23(7), 3542*. **Q1**.
- 
- [24] [Comprehensive Analysis of Knowledge Graph Embedding Techniques Benchmarked on Link Prediction](#)  
I. Ferrari, G. Frisoni, P. Italiani, G. Moro, C. Sartori (2022). *Electronics, 11(23), 3866*. **Q2**.
- 
- [25] [Human Being Detection from UWB NLOS Signals: Accuracy and Generality of Advanced Machine Learning Models](#)  
G. Moro, F. Di Luca, D. Dardari, G. Frisoni (2022). *Sensors, 22(4), 1656*. **Q1**.
- 
- [26] [Unsupervised Event Graph Representation and Similarity Learning on Biomedical Literature](#)  
G. Frisoni, G. Moro, G. Carlassare, A. Carbonaro (2021). *Sensors, 22(1), 3*. **Q1**.

---

[27] [A Survey on Event Extraction for Natural Language Understanding: Riding the Biomedical Literature Wave](#)  
G. Frisoni, G. Moro, A. Carbonaro (2021). *IEEE Access*, 9, 160721-160757. **Q1**.

---

BOOKS

[28] [Knowledge-Enhanced Information Retrieval](#)  
 Z. Wang, J. Fang, G. Frisoni, Z. Dai, Z. Meng, G. Moro, E. Yilmaz (2025).

---

BOOK CHAPTERS

[29] [Phenomena Explanation from Text: Unsupervised Learning of Interpretable and Statistically Significant Knowledge](#)  
G. Frisoni, G. Moro (2021). *Data Management Technologies and Applications. DATA 2020. Communications in Computer and Information Science*, vol. 1446, pp. 293-318. Springer, Cham. Revised Version of Best Paper Award.

---

WORKSHOPS AND SYMPOSIUMS

[30] [Enhancing Representation Learning for Content-Based Information Retrieval: A Knowledge-Enhanced Geometric Approach](#)  
 M. Goyo, G. Frisoni, G. Moro, C. Sartori (2025). *International Workshop on Knowledge-Enhanced Information Retrieval* (pp. 36-68). Cham: Springer Nature Switzerland.

---

[31] [Refining Triplet Sampling for Improved Self-Supervised Representation Learning](#)  
 M. Goyo, G. Frisoni, G. Moro, C. Sartori (2024). *SEBD* (pp. 227-246).

---

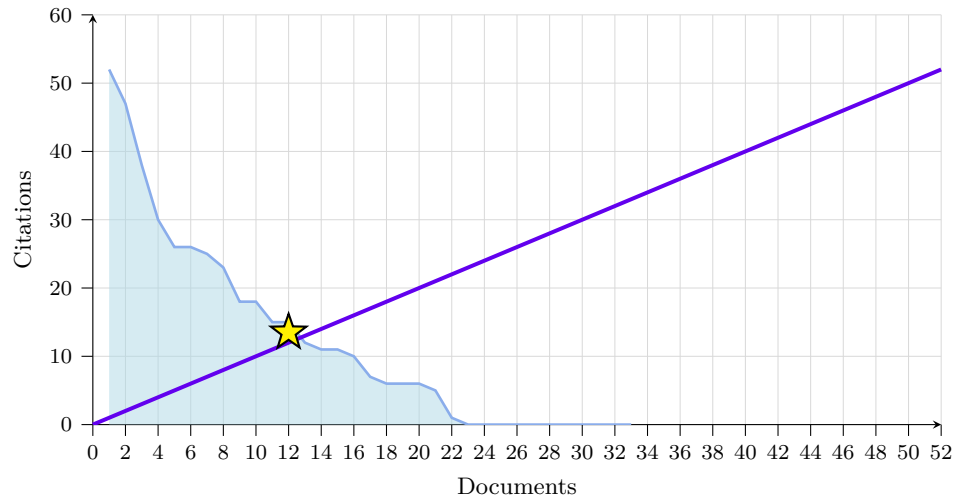
PREPRINTS

[32] [Mixture of Masters: Sparse Chess Language Models with Player Routing](#)  
G. Frisoni, L. Molfetta, D. Freddi, G. Moro (2026). *ArXiv*.

---

H-INDEX

The figure below illustrates the Scopus h-index plot as of June 8, 2026, where documents are ranked in descending order of citation count along the x-axis and the corresponding citation counts are reported on the y-axis. The diagonal line represents the  $y = x$  reference; the h-index is defined as the largest value  $h$  such that  $h$  documents have each received at least  $h$  citations, identified by the ★ marker at the intersection.



---

## 📄 SCIENTIFIC ACTIVITIES

---

### PARTICIPATION IN RESEARCH GROUPS

#### 🇮🇹 **UniboNLP** ★

2020 – Today

Research group coordinated by [Gianluca Moro](#), comprising faculty, postdoctoral researchers, and PhD students affiliated with the Department of Computer Science and Engineering, University of Bologna, Italy. This is the primary research group I worked with, joined upon completion of my Master’s degree. Within the group, I serve as the main expert on medical NLP and graph neural networks. Beyond scientific cooperation, my responsibilities span co-management of the SLURM cluster infrastructure; co-supervision of PhD, Master’s, and Bachelor’s students; contribution to national and international funding proposals; assistance in teaching material preparation; co-instruction, tutoring, and delivery of specialist seminars within courses.

#### 🇬🇧 **TerrierTeam**

2022 – Today

Research group within the Information, Data & Analysis Section, School of Computing Science, University of Glasgow, Scotland. One of Europe’s longest-standing IR groups, the TerrierTeam develops and evaluates large-scale text retrieval techniques and applications, with regular participation in benchmarking competitions such as TREC. I visited the group for ~four months during my doctoral studies, working with [Zaiqiao Meng](#) on the injection of external knowledge—retrieved or generated—into language models for biomedical question answering, a collaboration that led to the MedGENIE paper [9]. The scientific exchange continued beyond the visit, culminating in the co-organization of the Knowledge-Enhanced Information Retrieval (KEIR) workshop series at ECIR.

#### 🇮🇪 **IBM Research Europe**

2023 – Today

Research lab founded in 2011 and hosted at Trinity College Dublin, Ireland, serving as the only IBM Research facility within the European Union. The lab advances cutting-edge research across artificial intelligence, healthcare and life sciences, and quantum computing, cultivating close relationships with academic and industrial partners across Europe. I collaborate with [Marcos Martínez Galindo](#), [Alessandra Pascale](#), [Javier Carnerero Cano](#), and [James Barry](#) on zero-shot information extraction and multimodal fact checking, co-supervising Master’s and PhD students undertaking research visits and internships at the lab. The collaboration has led to the ZeroNER [3] and OpenBioNER [4] papers.

#### 🇫🇷 **EURECOM, Data Science Department**

2023 – Today

Graduate school and research center in digital sciences located in the Sophia Antipolis technology park on the French Riviera—Europe’s largest science and technology hub, home to over 2,300 companies and 5,000 researchers between Nice and Cannes. The Data Science Department pursues world-renowned research in statistical machine learning, multimodal data processing, and knowledge engineering, with strong ties to both academic and industrial partners. I collaborate with [Paolo Papotti](#) on multimodal fact checking, large language models, and databases, a scientific exchange that also included an invited talk at EURECOM and the Unown paper [8].

#### 🇨🇱 **Universidad Técnica Federico Santa María, Departamento de Informática**

2024 – Today

One of Chile’s leading technical universities, with the Departamento de Informática being among the first computer science departments in the country. The department pursues research across machine learning, artificial intelligence, natural language processing, and data science, with a strong tradition of international academic exchange. I collaborate with [Ricardo Nanculef](#) on representation learning and natural language processing. The collaboration has led to the papers [30] and [31].

#### 🇬🇧 **University of Edinburgh, School of Informatics**

2026 – Today

One of Europe’s largest and most prestigious informatics departments, the School of Informatics at the University of Edinburgh is ranked first in the UK by CSRankings and consistently among the top twenty worldwide. It encompasses the [EdinburghNLP](#) group and the [Institute for Language, Cognition and Computation \(ILCC\)](#), which together cover the full spectrum of NLP research—from statistical and neural methods to reasoning, semantics, and multimodal understanding. I collaborate with [Pasquale Minervini](#) on large language model reasoning, knowledge distillation, and board games.

RESEARCH  
CONTRACTS

**Postdoctoral Research Fellow @ University of Bologna, Italy** Feb 2024 – Jan 2027

- Supervisor: [Claudio Sartori](#)
- Affiliation: [Department of Computer Science and Engineering](#)
- Research group: [UniboNLP](#)
- Research project: [DARE \[cod. PNC0000002, CUP B53C22006450001\]](#)
  - Duration: Dec 2022–Dec 2026
  - Total budget: €124 million
  - Funder: The Italian Ministry of University and Research under the National Plan for Complementary Investments to the PNRR
  - Partners: Universities, research centers, IRCCS hospitals, local health authorities, environmental agencies, and private companies
  - Goal: Maximizing the potential of data to enhance health promotion and enable lifelong prevention
- Role: Research on representational and generative AI—including language models and graph neural networks—for clinical and biomedical applications

RESEARCH  
VISITS

**Postgraduate Visiting Researcher @ University of Glasgow, Scotland** Sep 2022 – Dec 2022

- Supervisor: [Zaiqiao Meng](#)
- Affiliation: [Information, Data & Analysis Section, School of Computing Science](#)
- Research Group: [Terrier Team](#)
- Object: Combining language models and structured knowledge for natural language processing and understanding in biomedical domains

PARTICIPATION IN  
RESEARCH  
PROJECTS

*National research projects, admitted to funding on the basis of competitive calls that include peer review*

**AI Analysis for Risk Assessment of Empty Lymph Nodes in Endometrial Cancer Surgery** Jul 2024 – Jun 2026

- Total budget: 20,000€
- Funder: The [Fondazione Cassa di Risparmio in Bologna \(CARISBO\)](#)
- Competitive call: “Ricerca Scientifica e Alta Tecnologia 2024”
- Partners: University of Bologna, Department of Computer Science and Engineering; University Hospital of Bologna IRCCS, Sant’ Orsola Polyclinic
- Goal: Develop an AI-based tool capable of estimating the probability that retroperitoneal tissue marked with indocyanine green, collected according to the standardized sentinel lymph node biopsy procedure, will turn out to be an empty lymph node packet at final microscopic analysis
- Role: Co-Principal Investigator (scientific responsibility)

**Artificial Intelligence for Public Administrations Connected** Mar 2023 – May 2026

- Total budget: ~€5.8 million
- Funder: Ministry of Enterprises and Made in Italy
- Competitive call: NextGenerationEU (winner of the Seal of Excellence award)
- Partners: SDA Bocconi, Maggioli Group, University of Bologna, University of Milano-Bicocca, IFEL, Tempo S.r.l.
- Goal: Support public administrations, small medium enterprises, and startups in piloting AI-based solutions
- Role: Assistance in the development of training services and test-before-invest services focused on retrieval-augmented chatbots in low-resource regimes

PARTICIPATION IN  
EDITORIAL BOARDS  
OF JOURNALS

*Special Issues*

**Role and Challenges of Healthcare Cognitive Computing: From Extraction to Data Analysis Techniques** Jun 2025 – Jan 2026

- Journal: [Sensors \(ISSN 1424-8220\)](#), Q1
- Role: Guest Editor
- Objective: Investigate the role and challenges of cognitive computing in the healthcare domain, with particular focus on techniques for extracting and analyzing unstructured data expressed in natural language, which characterize the current clinical information landscape. The collection focuses on the development of hybrid solutions that integrate

neural methodologies (machine learning and deep learning) with symbolic approaches (e.g., semantic web, knowledge graphs, semantic parsing graphs, first-order logic).

ORGANIZATION OF  
WORKSHOPS

**Knowledge-Enhanced Information Retrieval (KEIR-25),  
2nd Edition**

2025

- Role: Co-Chair
- Venue: European Conference on Information Retrieval (ECIR-25), Lucca, Italy, Apr 10, 2025 (in-person); A
- Co-organizers affiliated with University of Amsterdam, University of Glasgow, Google DeepMind, and University College London
- Topics: Agentic retrieval, ontology-based RAG, late chunking, contextual retrieval, learnable sparse embeddings, query rephrasing, multimodal IR
- Submissions: 10 papers submitted, 4 accepted; each manuscript (6–12 pages) peer-reviewed double-blind by three program committee members
- Proceedings: A selection of accepted, extended, and revised papers—together with additional invited contributions—published by Springer in LNCS, vol. 16086
- Responsibilities: Program committee recruitment, call-for-papers dissemination, final acceptance decisions, in-person event coordination and presentation, communication with the editorial office, and proceedings preparation

**Knowledge-Enhanced Information Retrieval (KEIR-24),  
1st Edition**

2024

- Role: Publicity Chair
- Venue: European Conference on Information Retrieval (ECIR-24), Glasgow, Scotland, Mar 28, 2024 (in-person); A
- Co-organizers affiliated with University of Glasgow, MBZUAI, Shandong University, University of Amsterdam, and University College London
- Topics: Knowledge infusion from external memories—structured and unstructured—to improve the effectiveness and interpretability of IR systems
- Submissions: 7 papers submitted, 4 accepted; each manuscript (6–12 pages) peer-reviewed double-blind by three program committee members
- Responsibilities: Program committee recruitment, call-for-papers dissemination, final acceptance decisions, and in-person event co-ordination

SESSION CHAIR IN  
INTERNATIONAL  
CONFERENCES

**The 40th Annual AAAI Conference on Artificial Intelligence (AAAI-26)**

2026

- Venue: Singapore, Jan 25, 2026 (in-person)
- Session: Natural Language Processing 15

**The 28th European Conference on Artificial Intelligence (ECAI-25)**

2025

- Venue: Bologna, Italy, Oct 30, 2025 (in-person)
- Session: Natural Language Processing 7

**The 11th International Conference on Data Science, Technology and Applications (DATA-22)**

2022

- Venue: Lisbon, Portugal, Jul 12, 2022 (virtual)
- Session: Neural Network Applications

AREA CHAIR IN  
INTERNATIONAL  
CONFERENCES

An Area Chair (AC) or Action Editor (AE) at ACL Rolling Review (ARR) oversees the full peer review process for assigned manuscripts. Responsibilities include reviewer selection and management (checking conflicts of interest, expertise fit, and compliance with ethics and anonymity policies), moderation of author-reviewer discussion, solicitation of missing contributions, and adherence to ARR’s monthly editorial schedule. The AC drafts the final meta-review, synthesizing reviewer assessments, integrating author responses, and formulating a motivated editorial recommendation. The role also involves evaluating review quality, flagging negligence or improper use of automated tools, and recognizing good reviewing practices. Appointment is by invitation from the Editors-in-Chief or Program Chairs, based on disciplinary expertise, reputation, and declared availability—a selection reserved for researchers with a mature and recognized academic profile in the NLP community, evidenced by a consistent publication record at ACL venues or equivalent.

**21 meta-reviews**

- ACL ARR May 2026 — 7 meta-reviews 2026
- ACL ARR January 2026 — 7 meta-reviews
- ACL ARR October 2025 — 7 meta-reviews 2025

REVIEWING FOR  
INTERNATIONAL  
CONFERENCES

## 79 reviews

- The 40th Annual Conference on Neural Information Processing Systems (NeurIPS-26); A\* — 4 reviews 2026
- The 43th International Conference on Machine Learning (ICML-26); A\* — 6 reviews
- The 19th Conference of the European Chapter of the Association for Computational Linguistics (EACL-26); Student Research Workshop; A — 2 reviews
- The 14th International Conference on Learning Representations (ICLR-26); A\* — 5 reviews
- The 40th Annual AAAI Conference on Artificial Intelligence (AAAI-26); A\* — 6 reviews
- The 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026); A — 1 review
- ACL ARR May 2025 — 4 reviews 2025
- The 1st International Workshop on Advanced Neuro-Symbolic Applications (ANSyA-25) @ ECAI-25 — 2 reviews
- The 28th European Conference on Artificial Intelligence (ECAI-25); A — 3 reviews
- The 2nd Workshop on Knowledge-Enhanced Information Retrieval (KEIR-25) @ ECIR-25 — 1 review
- ACL ARR February 2025 — 4 reviews
- The 13th International Conference on Learning Representations (ICLR-25); A\* — 3 reviews
- ACL ARR December 2024 — 4 reviews 2024
- ACL ARR October 2024 — 4 reviews
- ACL ARR August 2024 — 4 reviews
- The 38th Annual Conference on Neural Information Processing Systems (NeurIPS-24); Datasets and Benchmarks Track; A\* — 2 reviews
- The 1st Workshop on Knowledge-Enhanced Information Retrieval (KEIR-24) @ ECIR-24 — 2 reviews
- The 46th European Conference on Information Retrieval (ECIR-24); A — 5 reviews
- The 38th Annual AAAI Conference on Artificial Intelligence (AAAI-24); A\* — 1 review
- ACL ARR June 2024 — 4 reviews
- ACL ARR February 2024 — 4 reviews
- ACL ARR December 2023 — 4 reviews 2023
- The 20th Extended Semantic Web Conference (ESWC-23); B — 1 review
- The 22th International Semantic Web Conference (ISWC-23); A — 1 review
- The 37th Annual Conference on Neural Information Processing Systems (NeurIPS-23); A\* — 1 review
- The 37th Annual AAAI Conference on Artificial Intelligence (AAAI-23); A\* — 1 review

REVIEWING FOR  
INTERNATIONAL  
JOURNALS

## 7 reviews

- Engineering Applications of Artificial Intelligence, 2024; Q1 — 1 review 2024
- Engineering Applications of Artificial Intelligence, 2023; Q1 — 1 review 2023
- MDPI Information, 2023; Q2 — 1 review
- MDPI Symmetry, 2023; Q2 — 1 review

- Semantic Web Journal, 2022; Q2 — 1 review 2022
- Expert Systems, 2022; Q2 — 1 review
- MDPI Applied Sciences, 2022; Q2 — 1 review

AWARDS

**Gold Reviewer Award** 2026

- Venue: The 43th International Conference on Machine Learning (ICML-26); A\*
- Granting body: ICML-26 Program Chairs
- Selectivity: Awarded to 4,439 reviewers out of 17,749 (top 25%)
- Criteria: Review quality as assessed by area chairs, based on thoroughness, constructiveness, timeliness, and engagement in author-reviewer discussion

**Outstanding Reviewer Award** 2025

- Venue: The 30th Conference on Empirical Methods in Natural Language Processing (EMNLP-25); A\*
- Granting body: EMNLP-25 Program Chairs
- Selectivity: Awarded to 51 reviewers out of 13,048 (top ~0.4%)
- Criteria: Review quality as assessed by area chairs, based on analytical depth, constructiveness, and timeliness

**Best Student Paper Award** 2022

- Venue: The 11th International Conference on Data Science, Technology and Applications (DATA-22); C
- Granting body: Institute for Systems and Technologies of Information, Control and Communication (INSTICC), Portugal
- Selectivity: Awarded to the best full paper with a student as presenting author out of 62 accepted submissions
- Criteria: Double-blind peer review and oral presentation assessment by session chairs, with final deliberation by program, conference, and doctoral consortium chairs based on originality, methodology, and clarity
- Paper: [Enhancing Biomedical Scientific Reviews Summarization with Graph-based Factual Evidence Extracted from Papers Certificate](#)


**Best Paper Award** 2020

- Venue: The 9th International Conference on Data Science, Technology and Applications (DATA-20); C
- Granting body: Institute for Systems and Technologies of Information, Control and Communication (INSTICC), Portugal
- Selectivity: Awarded to the best full paper in the conference program out of 35 accepted submissions
- Criteria: Double-blind peer review and oral presentation assessment by session chairs, with final deliberation by program, conference, and doctoral consortium chairs based on originality, methodology, and clarity
- Paper: [Learning Interpretable and Statistically Significant Knowledge from Unlabeled Corpora of Social Text Messages: A Novel Methodology of Descriptive Text Mining \(first authored publication\)](#)
- Certificate

**Best Master's Thesis Award** 2020

- Prize: Con.Scienze Prize for Scientific Master's Theses, 2020
- Granting body: [National Conference of Presidents and Directors of University Structures in Sciences and Technologies \(Con.Scienze\)](#), Italy
- Selectivity: National prize awarded to the 10 best Master's research theses in scientific disciplines; one nomination per department
- Criteria: Originality, methodological rigor, and scientific relevance

CONFERENCE PRESENTATIONS

-  AACL-26. Singapore. "Graph-of-Mark: Promote Spatial Reasoning in Multimodal Language Models with Graph-Based Visual Prompting" [1]. In-person, Oral. 2026

- 🇨🇳 EMNLP-25. Suzhou, China. “PORTS: Preference-Optimized Retrievers for Tool Selection with Large Language Models” [2]. In-person, Oral. 2025
- 🇮🇹 ECAI-25. Bologna, Italy. “Magic Mirror on the Wall, Which is the Fairest Prompt of All? A Survey on Automatic Prompt Learning” [6]. In-person, Poster.
- 🇮🇹 ECAI-25. Bologna, Italy. “Predicting Protein Functions with Ensemble Deep Learning and Protein Language Models” [7]. In-person, Demo.
- 🇦🇪 EMNLP-22. Abu Dhabi, United Arab Emirates. “BioReader: a Retrieval-Enhanced Text-to-Text Transformer for Biomedical Literature” [11]. In-person, Poster. 2022
- 🇰🇷 COLING-22. Gyeongju, Republic of Korea. “Text-to-Text Extraction and Verbalization of Biomedical Event Graphs” [13]. Virtual, Oral.
- 🇵🇹 DATA-22. Lisbon, Portugal. “Enhancing Biomedical Scientific Reviews Summarization with GraphBased Factual Evidence Extracted from Papers” [12]. Virtual, Oral.
- 🇫🇷 DATA-20. Paris, France. “Learning Interpretable and Statistically Significant Knowledge from Unlabeled Corpora of Social Text Messages: A Novel Methodology of Descriptive Text Mining” [15]. Virtual, Oral. 2020

INVITED TALKS

*Research Centres*

**LLM Evaluation: From Best Practices to Legal Domain Benchmarks** 2026

Alma Mater Research Institute for Human-Centered Artificial Intelligence (ALMA-AI) Workshop Series; Contact: [Monica Palmirani](#)

**NLP Methods Based on Small and Large Language Models in Low-Resource Regimes** 2025

Alma Mater Research Institute for Human-Centered Artificial Intelligence (ALMA-AI) Workshop Series; Contact: [Monica Palmirani](#)

---

*Doctoral Training Programmes*

**LLM Training and Evaluation** 2026

[FutureData4EU](#), international programme funded under the European Marie Skłodowska-Curie COFUND; Contact: [Stefano Lodi](#)

---

*Foreign Universities*

**Knowledge-Enhanced NLP: From Retrieval Augmentation to Generative Augmentation in Medical Domains** 2024

[EURECOM](#), France; Contact: [Paolo Papotti](#)

**Retrieval-Enhanced Language Models and Semantic-Driven Summarization for Biomedical Domains** 2022

[TerrierTeam](#), University of Glasgow, Scotland; Contact: [Zaiqiao Meng](#)

RELEVANT PROGRAM ATTENDANCE

**Advanced Course on Data Science and Machine Learning, ACDL 2024** 2024

Tuscany, Italy; [Certificate](#)

**Oxford Machine Learning School, OxML 2023** 2023

Andrew Wiles Building, Mathematical Institute, University of Oxford; [Certificate](#)

**The Cornell, Maryland, Max-Planck Pre-doctoral Research School, CMMRS 2020** 2020

Max Planck Institute for Software Systems; One of 100 students internationally selected to participate; [Certificate](#)

VOLUNTEERING

*Patient Advocacy*

**AMAE Onlus, Italian Nonprofit Association for Esophageal Achalasia Patients (Rare Disease)** 2019—Today

Member of the Board of Directors. As both a researcher and a patient, I contribute to the association by applying NLP and IT expertise to patient-centered initiatives, supporting research advancement and the development of AI-based tools for rare disease communities.

*Developer Advocacy*

**Streamlit Student Ambassador** 2022—2023

Selected among fewer than 50 students worldwide for the inaugural cohort of Streamlit’s Student Ambassador Education Program. Role involved building and sharing data applications, and promoting interactive ML development tools within academic communities.

**Hugging Face Student Ambassador** 2022

Selected among fewer than 100 students worldwide for the inaugural cohort of Hugging Face’s Student Ambassador Program, a global initiative by one of the leading open-source ML platforms aimed at democratizing machine learning education across diverse academic backgrounds. Role involved promoting open-source ML tools and fostering community adoption of the Hugging Face ecosystem.

CERTIFICATIONS

🌟 **AI Agents Fundamentals**

Hugging Face; 2025; [Certificate](#)

🌟 **Generative AI with Large Language Models**

Coursera; 2023; [Certificate](#)

🌟 **Natural Language Processing Specialization**

Coursera; 2020; [Certificate](#)

---

## TEACHING ACTIVITIES

---

ADJUNCT  
PROFESSOR

### Master in Data Science and Business Analytics @ Bologna Business School, Italy

- Course: Text Mining and Natural Language Processing  
Machine learning and deep learning for textual data. Text representation techniques: bag-of-words, TF-IDF, word2vec, contextualized embeddings. Processing methods: lexicon-based predictors, multilayer perceptrons, recurrent neural networks, transformers. Large language model prompting and fine-tuning. Focus on text classification tasks.
- Scientific disciplinary sector: ING-INF/05
- Role: Co-instructor (with Gianluca Moro); responsible for the laboratory module (12 hours of in-person teaching), including the design and evaluation of coding assignments
- Academic years (satisfaction out of 5): [iii] 2025–2026, [ii] 2024–2025 (4.5), [i] 2023–2024 (4.3)

TUTORING

### M.S., Artificial Intelligence @ University of Bologna, Italy

- Course: Big Data Analytics and Text Mining (Module 1)  
The course introduces students to the fundamentals and state-of-the-art methodologies in the field of Natural Language Processing, actively incorporating recent milestones and newly published research contributions. The content focuses on Large Language Models, covering advanced topics such as prompt optimization techniques, model compression and quantization methods, parameter-efficient fine-tuning, and multimodal language models. It also examines approaches for integrating external knowledge, such as Retrieval-Augmented Generation, Graph Neural Networks, and agents. Assessment consists of a multiple-choice examination and the discussion of a project.
- Scientific disciplinary sector: ING-INF/05
- Role: Assistance in preparing teaching material, seminars, project assignment and assessment (20 hours)
- Course coordinator: Gianluca Moro
- Academic years: [iii] 2025–2026 (cod. B2133), [ii] 2024–2025 (cod. B2133), [i] 2023–2024 (cod. 91262)

THESIS  
CO-SUPERVISION

Co-supervisor of 57 Bachelor's and Master's theses on Natural Language Processing and Deep Learning at the University of Bologna. A complete list is available on AMS Laurea, the University of Bologna's official open-access institutional thesis repository.

#### *Selected theses*

- Buzzoni, M. (2026). *Graph-of-Mark: Graph-Based Visual Prompting for Enhanced Spatial Reasoning in Multimodal Language Models*. M.S. thesis, Artificial Intelligence [LM-DM270]
- Raponi, M. (2025). *Multimodal Generative Information Retrieval of Chest X-Rays Grounded on ICD-9 Taxonomy*. B.S. thesis, Computer Science and Engineering [L-DM270]
- Zecca, A. (2024). *End-to-End Extraction and Injection of Graphs: A Self-Supervised Neuro-Symbolic Method for Explainable Large Language Models*. M.S. thesis, Artificial Intelligence [LM-DM270]
- Monaldini, N. (2024). *Large Action Models: End-to-End Retrieval-Enhanced Learning for Generating Function Calls from Instruction Manuals*. B.S. thesis, Computer Science and Engineering [L-DM270]
- Cocchieri, A. (2023). *Leveraging Large Language Model Distillation to Enhance Zero-Shot Named Entity Recognition and Classification*. M.S. thesis, Artificial Intelligence [LM-DM270]
- Presepi, A. (2023). *To Generate or to Retrieve: On the Effectiveness of Artificial Contexts for Biomedical Question Answering*. B.S. thesis, Computer Science and Engineering [L-DM270]

PARTICIPATION IN  
INTERNATIONAL  
THESIS EVALUATION  
COMMISSIONS

### Universidad Técnica Federico Santa María (UTFSM), Valparaíso, Chile

2026

- Degree: Magíster en Ciencias de la Ingeniería Informática

- Role: External Member of the Thesis Committee; responsibilities include written evaluation of the dissertation and participation in the oral defense
- Topics: counterfactual explanations, explainable AI, natural language processing