

Andrea Galassi

| | |
|---------------|--|
| FIRST NAME | Andrea |
| FAMILY NAME | Galassi |
| YEAR OF BIRTH | 1992 |
| CITIZENSHIP | Italian |
| AFFILIATION | DISI, University of Bologna, Bologna, Italy |
| WORK ADDRESS | Viale del Risorgimento, 2, 40136 Bologna (BO), Italy |
| E-MAIL | a.galassi@unibo.it |
| WEB PAGES | https://www.unibo.it/sitoweb/a.galassi https://site.unibo.it/nlp/en/people/andrea-galassi |
| ORCID ID | 0000-0001-9711-7042 |

Short Bio

Andrea Galassi received the M.S. degree in Computer Engineering and the Ph.D. in Computer Science and Engineering from the University of Bologna (Italy), in 2017 and 2021 respectively. Since 2023 he has been employed as a *Junior Assistant Professor (RTDA)* at the Department of Computer Science and Engineering (DISI) of the University of Bologna. He obtained National Scientific Qualification for the role of Associate Professor in Computer Engineering (ASN 2023-2025 sector 09/H1 II Tier). Since 2018, he has worked in the Language Technologies Lab, led by Prof. Paolo Torroni, and he is now the deputy director of the lab.

His research interests are focused on **Natural Language Processing** and **Machine Learning**, with particular attention to Trustworthy and Human-Centered AI, Argument Mining, Neuro-Symbolic approaches to machine learning, and Legal Analytics. His expertise spans different areas, with the general target of exploiting domain knowledge that is formally defined along with specific knowledge that can be acquired from data, especially textual documents. A large part of his work has focused on the extraction, processing, and use of argumentative content in several multidisciplinary domains. Examples of applications include problems of information retrieval, such as ranking scientific documents or providing trustworthy answers in dialogues, and tasks of forecasting and classification, such as predicting the outcome of a legal judgment. He secured funding and acted as Principal Investigator in two international small projects, and has participated in several national and international research projects.

Highlights

- PhD in Computer Science and Engineering, 2021
 - National Scientific Qualification as Associate Professor, 2024
 - Junior Assistant Professor (RTDA) at the University of Bologna, since 2023
 - Delegate for Equity, Inclusion and Diversity since 2024
 - Deputy director of the Language Technologies Lab of Prof. Paolo Torroni
 - Leader in 2 peer-reviewed and funded international projects as PI; participation in 6 international and 4 national projects
 - Section Editor of the Artificial Intelligence and Law journal (Scimago Q1)
 - Visiting Professor at Université Catholique de Lille
 - Teaching experience in bachelor's degrees, master's degrees, professional master programs, PhD curricula, and summer schools, for a total of more than 200 hours of teaching.
 - Author of 10 publications in computer science journals and a total of more than 35 publications in peer-reviewed scientific venues
 - Recipient of 2 awards: Best Innovative Application Paper (ICAIL 2025) and Outstanding Research Paper (JURIX 2024)
 - More than 950 citations over more than 850 documents, h-index 12 according to the Scopus service; more than 1,600 citations, h-index 17 according to the Google Scholar service
-

Position, Education, and Qualifications

Qualifications and Certificates

- NOVEMBER 2024
National Scientific Qualification as associate professor
DISCIPLINARY FIELD: 09/IINF-05 (ex 09/H1) - Information processing systems - Computer engineering

Institutional Roles

- 2024 – ONGOING
Delegate for Equity, Inclusion and Diversity
INSTITUTION: Department of Computer Science and Engineering (DISI), University of Bologna

Record of Employment

- MARCH 2023 – currently
Junior assistant professor (RTDA) at Department of Computer Science and Engineering (DISI), University of Bologna, Italy
SUPERVISOR: Prof. Michela Milano
- MAY 2025
Visiting Professor at Law School, Université Catholique de Lille, France
- JANUARY 2021 – FEBRUARY 2022
Research Fellow at the Department of Computer Science and Engineering (DISI), University of Bologna, Italy
SUPERVISOR: Prof. Paolo Torroni
- SPRING 2021
Adjunct Professor at Department of Computer Science and Engineering (DISI), University of Bologna, Italy
- NOVEMBER 2021
Collaborator with University of Modena and Reggio Emilia, Italy
PROJECT LEADER: Prof. Marco Lippi
- APRIL 2017 – OCTOBER 2017
Research Fellow at the Inter-departments Center for Health Sciences and Technologies (CIRI-HST), University of Bologna, Italy
SUPERVISOR: Prof. Federico Chesani

Education

- NOVEMBER 2017 - MAY 2021
Ph.D. in Computer Science and Engineering, University of Bologna
FINAL EVALUATION: Excellent (5/5)
THESIS TITLE: Deep Networks and Knowledge: from Rule Learning to Neural-Symbolic Argument Mining
ADVISORS: Prof. Paolo Torroni, Prof. Marco Lippi, Prof. Michela Milano
- DECEMBER 2014 - MARCH 2017
Master Degree in Computer Engineering, University of Bologna
FINAL EVALUATION: 110/110 cum laude
THESIS TITLE: Symbolic versus sub-symbolic approaches: a case study on training Deep Networks to play Nine Men's Morris game
ADVISOR: Prof. Paola Mello
- SEPTEMBER 2011 - DECEMBER 2014
Bachelor Degree in Computer Engineering, University of Bologna
FINAL EVALUATION: 106/110
THESIS TITLE (TRANSLATED): Handwritten writing analysis based on similarity search: distances and matching models comparisons
ADVISOR: Prof. Ilaria Bartolini

Language Proficiency

- **Italian:** Mother tongue
- **English:** C1
CERTIFICATION: IELTS, British Council, 7.5 score, July 2016

Period spent abroad

- MAY 2025
Law School, Université Catholique de Lille, Paris, France
 - MAY 2024
City University of London, London, United Kingdom
SUPERVISOR: Artur d'Avila Garcez, Director of the Research Centre for Machine Learning
 - FEBRUARY 2020 – AUGUST 2020
Imperial College, London, United Kingdom
SUPERVISOR: Francesca Toni, Professor in Computational Logic
 - MAY 2018 – AUGUST 2018
Stanford University, Palo Alto, U.S.A.
SUPERVISOR: Margaret Hagan, Director of the Legal Design Lab
-

Contribution to Academic Research Projects

Funded Research:

Andrea Galassi acted as Principal Investigator or Co-Principal Investigator for the following funded projects:

- **Promoting Fairness and Diversity in Speech Datasets for Affective Computing**
SOURCE: HumaneAI-Net Consortium
CALL: HumaneAI-Net call for micro-projects 2023
PEER REVIEWED: Yes, competitive call
PERIOD: October 2023 – March 2024
AMOUNT: 29,300€ total, 20,000€ for Andrea's Institution
SHORT DESCRIPTION: Critical review and meta-analysis of existing speech datasets in the perspective of inclusivity, transparency, and fair use.
- **A Transparent and Explainable Dialogue System for Immigration Services**
SOURCE: HumaneAI-Net Consortium
CALL: HumaneAI-Net call for micro-projects 2023
PEER REVIEWED: Yes, competitive call
PERIOD: November 2023 – April 2024
AMOUNT: 32,000€ total, 20,000€ for Andrea's Institution
SHORT DESCRIPTION: A transparent and explainable dialogue system for assisting immigrants and non-profit organizations on administrative and legal matters in Italy.
- **Argument Mining for Supporting Learning with Large Language Models**
SOURCE: HumaneAI-Net Consortium
CALL: Contribution to macro-projects
PEER REVIEWED: Informal review by project coordinators and WP leaders
PERIOD: January 2024 – June 2024
AMOUNT: 35,000€ total, 30,000€ for Andrea's Institution
SHORT DESCRIPTION: Use of Argument Mining to enhance and evaluate the ability of Large Language Models to reason, with application to the education context.
- **Neuro-symbolic Argument Mining on Legal Texts**
SOURCE: FAIR Foundation
CALL: Trustworthy AI Short Stay Exchange Program for Researchers Italy-UK
PEER REVIEWED: Yes, competitive call
PERIOD: Spring 2024
AMOUNT: 2,500€ for Andrea
SHORT DESCRIPTION: Preliminary investigation regarding the application of neuro-symbolic frameworks to the problem of Argument Mining on Legal Texts.

Contribution to International Projects:

Andrea Galassi contributed to the following international projects:

- **JUST-2022-EJUSTICE-101087342-POLINE: Principles Of Law In National and European VAT**
PERIOD: 2023 – present
ROLE: Tasks leader
WEBSITE: <https://site.unibo.it/poline/en>
DESCRIPTION: POLINE aims at developing an AI-powered pilot tool for the retrieval and analysis of judicial principles of law in the CJEU and national case-law in Value Added Tax (VAT). The development of the tool will be based on a multidisciplinary approach combining theory and practice of judicial decision-making; legal informatics methods; AI, machine learning, and NLP techniques. The tool includes AI techniques for extracting, clustering and linking judicial principles of law. It covers the case-law of the CJEU and the Italian, Swedish and Bulgarian Supreme Courts and will be accessible to judges, other legal practitioners, tax policymakers and taxpayers.

Andrea Galassi actively contributed to the writing of the project's proposal. He currently contributes to and supervises activities in WP3, "Machine learning and NLP analysis". In particular, he is the leader of tasks 3.1 and 3.2, which consist of supporting manual and automatic annotation, and extraction of legal principles from textual documents. His duties in this project include coordination with other WPs, the development of novel solutions, and the supervision of interns, students, and research fellows.

- **H2020-ICT-2018-825619-AI4EU: A European AI On Demand Platform and Ecosystem**
 PERIOD: 2023 – 2024
 WEBSITE: <https://cordis.europa.eu/project/id/825619>
 DESCRIPTION: The EU-funded AI4EU is working to change Europe's place in this race, by building the first European AI On-Demand Platform and Ecosystem that will share resources, tools, knowledge, algorithms and more between Member States. It will help to increase innovation and technology transfer, accelerate the growth of start-ups and SMEs, and fulfill the needs of the European AI community. The project will implement eight pilots led by industrial partners to demonstrate the platform's capabilities.

Andrea Galassi worked on the design and development of the Success Stories service, where companies can showcase how the use of AI has impacted their work.
- **H2020-ICT-2020-101017142-StairwAI: Stairway to AI: Ease the Engagement of Low-Tech users to the AI-on-Demand platform through AI**
 PERIOD: 2022 – 2023
 WEBSITE: <https://stairwai.nws.cs.unibo.it/about-project/>
 DESCRIPTION: The StairwAI project targets low-tech users with the goal of facilitating their engagement on the AI4EU on-demand Platform. This will be achieved through a new service layer enriching the functionalities of the on-demand platform and containing: (1) a multi-lingual interaction layer enabling conversations with the Platform in the user's own language, (2) a horizontal matchmaking service for the automatic discovery of AI assets (tools, data sets, AI experts, consultants, papers, courses etc.) meeting the user business needs and, (3) a vertical matchmaking service that will dimension and provision hardware resources through a proper hardware provider (HPC, Cloud and Edge infrastructures).

Andrea Galassi was involved in work package 5, relative to horizontal matchmaking. In particular, his contribution revolves around the combination of machine-learning-based and knowledge-based techniques to associate use case documents to available AI assets automatically. He has also contributed to work package 8, with deliverable "International Networking Strategy".
- **H2020-ICT-2020-952026-HumanE AI Network**
 PERIOD: 2021 – 2024
 ROLE: Micro-Projects Principal Investigator
 WEBSITE: <https://www.humane-ai.eu/>
 DESCRIPTION: The objective of the HumanE-AI-Net network is to connect European research centers, universities, and industrial enterprises to develop robust, trustworthy AI systems that can 'understand' humans, adapt to complex real-world environments and interact appropriately in complex social settings.

Andrea Galassi's primary focus lies within work package 5, which concerns the ethical, legal, and social challenges associated with artificial intelligence. Within this framework, he has made valuable contributions through the completion of four distinct micro-projects. For two of them, Andrea Galassi acted in the capacity of Principal Investigator (PI) after securing funding through a dedicated competitive call that involved peer review by consortium leaders and domain experts. Moreover, he wrote the proposal and coordinated the effort for the contribution of his institution to the macro project "Learning with LLMs: Supporting complex reasoning, planning and argumentation, applied to providing educational guidance". So far, his work has resulted in 7 publications, among which two in scientific journals [J1, J7, C3, C3, W4, W13, W14]. Finally, he contributed to the "Handbook of Human-AI Collaboration", acting as co-editor of Section, "Multimodal Foundation Models", and author of two chapters.
- **JUST-JACC-EJU-AG-2020-1010074206-Adele: Analytics for DEcision of LEgal cases**
 PERIOD: 2021 – 2023
 ROLE: Task Leader
 WEBSITE: <https://site.unibo.it/adele/en>
 DESCRIPTION: Project ADELE is premised on the ongoing paradigm shift towards cognitive computing and human-centered AI which is transforming many socio-economic activities, including justice. The project applies legal analytics (LA) – a blend of data science, machine learning, and natural language processing techniques – to judicial decisions. It aims to develop methods to extract knowledge and engage in outcome predictions and there build a pilot tool to support legal research and decision-making processes in the judiciary.

Andrea Galassi was part of the Natural Language Processing team of Unibo, and mainly involved in the tasks of outcome prediction and argument mining. Andrea's duties included contributing to the development of solutions, supervising interns and research fellows, and coordinating with the legal experts. His work led to the publication of two conference articles [W10, C6].

- **CLAUDETTE: "automated CLAUse DETectEr"**

PERIOD: 2019 – present

WEBSITE: <http://claudette.eui.eu/>

DESCRIPTION: CLAUDETTE is an interdisciplinary research project hosted at the Law Department of the European University Institute. The research objective is to test to what extent it is possible to automate reading and legal assessment of online consumer contracts and privacy policies, to evaluate their compliance with EU's unfair contractual terms law and personal data protection law (GDPR), using machine learning and grammar-based approaches.

Andrea Galassi is involved in the development of machine-learning solutions to automatically classify segments of text. In particular, his work is focused on multi-lingual approaches that can exploit the knowledge acquired in one domain (the English language) and apply it in novel ones (other languages) through the projection of labels based on content similarity. His contribution has led to the publication of two conference articles [W12, C12] and one journal article [J3].

Contribution to National Projects:

Andrea Galassi contributed to the following national projects:

- **PNRR PE00000013 FAIR: Future Artificial Intelligence Research**

PERIOD: 2022 – present

ROLE: Task leader in Transversal Project

WEBSITE: <https://future-ai-research.it/>

DESCRIPTION: The objective of the FAIR project is to contribute facing the research questions, methodologies, models, technologies, and ethical and legal rules to build AI systems capable of interacting and collaborating with humans, perceiving and acting in evolving contexts, to be conscious about their limits and capable to adapt to new situations, to be aware of the perimeters of safety and trust, and to be careful with the environmental and social impact that their creation and functioning may cause.

Andrea Galassi is involved in workpackage 8.1 "Multi-scale learning and reasoning in pervasive AI systems". His research concerns the development of scalable, transferable, neuro-symbolic machine-learning techniques. He is also involved in work package 8.3, "Visual and Multimodal Perception in pervasive systems" and the transversal project 2, "Vision, Language and Multimodal Challenges", where he acts as task leader of T1.3, "Existing Textual and Multimodal Evaluation Data". His work resulted in the publication of one conference paper [C9], three journal papers [J2, J3, J4, J5], and the organization of a workshop [C8].

- **PRIN20224TPEYC PRIMA: PRivacy Infringements Machine-Advice**

PERIOD: 2024 – present

DESCRIPTION: The project concerns the law and practice of privacy policies from three different perspectives: normative (doctrinal), empirical (socio-legal), legal-informatics. In particular, it deploys analytics to detect and assess privacy policies' infringements. The project team combines all the different scientific, theoretical and methodological approaches into a new interdisciplinary synthesis (law, computing, legal informatics, AI, ML, legal theory, computer ethics).

Andrea Galassi's contributes to WP3 ("Focus on online platforms") and WP5 ("Computational methods and prototype development"). His duties include coordination with legal experts for the annotation of new resources, the development of novel solutions, and the supervision of interns, students, and research fellows. His work resulted in the publication of two conference articles, both winners of awards [W1, C2].

- **PRIN2017NCPZ22 LAILA: Legal Analytics for Italian Law**

PERIOD: 2022 – 2024

ROLE: Task Leader

WEBSITE: <https://dsg.unibo.it/it/ricerca/progetti-di-ricerca/progetti-nazionali-e-di-ateneo/prin2017-laila-legal-analytics-for-italian-law>

DESCRIPTION: The project regards the application of Legal Analytics methods to a vast and heterogeneous set of legal information: legislations, contracts, and judgments. The purpose is the application of Artificial Intelligence, Machine Learning, and Natural Language Processing to extract legal knowledge, infer relationships, and produce data-driven forecasts.

Andrea Galassi's duties in this project involve the research of new solutions, the supervision of interns, students and research fellows, and the coordination with the team of legal experts of Unibo. He is involved in the publication of two conference articles [W7, W9] related to this project.

- **FISR2020IP_01362 AMICA: Argument Mining In Covid-19 Articles**

PERIOD: 2021

WEBSITE: <http://amica.unimore.it/>

DESCRIPTION: The objective of the AMICA project was to exploit the argumentative content present in the scientific literature regarding Covid-19 to improve the retrieval of relevant and reliable articles. The project involved both medical and artificial intelligence experts and aimed to develop an argument mining-based search engine, specifically designed for the analysis of scientific literature related to Covid-19.

Andrea Galassi contributed by designing possible solutions and supervising a post-graduate student. His work has led to the publication of two articles, one in a computer science conference [C10] and one in a medical journal [J6].

Other Projects:

Andrea Galassi contributed to the following projects:

- **ALMA GAIE**

PERIOD: 2024 – 2025

WEBSITE: <https://magazine.unibo.it/archivio/2024/05/23/intelligenza-artificiale-per-una-pa-sostenibile-premiato-alma-gaie>

DESCRIPTION: A project funded by the University of Bologna, aimed at developing an AI system for the automatic classification of research and educational products of the University of Bologna according to their contribution to the 17 Goals of the United Nations' 2030 Agenda for Sustainable Development. The project received the "PA a colori" 2024 award for sustainability in public administrations.

Andrea Galassi coordinated the research activities of the project and supervised PhD students and Research Fellows working on the project. His work resulted in one paper [T1].

- **GeMEB: Generative Models Empowering Business**

PERIOD: 2024 – 2025

DESCRIPTION: GeMEB is a research and innovation project run in collaboration with NIER and Mexage aimed to Empower Business Processes and Enhancing Workflows for Improved Performance using Large Language Models and Retrieval-Augmented Generation.

Andrea Galassi supervised the activities of PhD students and Research Fellows working on the project.

- **SmartCasm: AI-based Smart Collaborative Manufacturing System**

PERIOD: 2024 – 2025

DESCRIPTION: SmartCasm is a research and innovation project in collaboration with BitBang, Expert.ai, Bonfiglioli and MEP, aiming to empower manufacturing companies via generative AI, LLMs, RAG and knowledge graphs for structured and unstructured knowledge integration.

Andrea Galassi supervised the activities of PhD students and Research Fellows working on the project.

Other Contributions to Research Projects

Andrea Galassi has contributed to research projects by acting as a project reviewer for the following institutions:

- National Science Centre (NCN) of Poland (2021, 2025)
- University of Naples Federico II (2025)

Andrea Galassi has been a member of the scientific and technological committee of innovation projects funded by the following agencies:

- Fondimpresa (2021-2022)
-

Contribution to Scientific Conferences and Journals

Invited Contributions

Andrea Galassi has been **invited keynote speaker** at the following international conferences and initiatives:

- 2nd European Innovation Council (EIC) Portfolio Workshop on Inside the Ethics of AI Awareness
Title of the intervention: "Towards an Ethical and Human-centric Artificial Intelligence: two case studies on Fairness in Dialogue Systems and Speech Datasets"
Uppsala, November 2024
Organized by EIC Portfolio partners
- Touché: workshop on Argument Retrieval
Title of the intervention: "Argumentative Ranking: Case Studies and Challenges"
Bologna, September 2022
Organized by Webis Group: Alex Bondarenko, Matthias Hagen, Martin Potthast, Benno Stein
Website: <https://touche.webis.de/clef22/touche22-web/index.html>
- Legal Hackathon Wales
Title of the intervention: "Making legal documents more accessible through Natural Language Processing: analysis, explanations, and multilingualism"
Hosted online, July 2022
Organized by Legal Innovation Lab Wales
https://www.youtube.com/watch?v=n9QMZQtNTQg&t=3s&ab_channel=LegalInnovationLabWales
- NLLP Talks, affiliated with the Natural Legal Language Processing Workshop
Title of the intervention: "Cross-lingual Annotation Projection in Legal Texts"
Hosted online, May 2021
Organized by Catalina Goanta, discussed with Monika Leszczynska and Kasper Drazewski
Website: <https://nllpw.org/talks/>

Organization Committees

Andrea Galassi has contributed to the organization of the following events:

- CheckThat! Lab 2023, International Workshop on Fact-Checking
Co-located with CLEF 2023, Conference and Labs of the Evaluation Forum
ROLE: **co-chair**
September 2023, Thessaloniki, Greece
WEBSITE: <https://checkthat.gitlab.io/>
- ECAI 2025, European Conference on Artificial Intelligence
ROLE: **Student Volunteer Chair**
October 2025, Bologna, Italy
WEBSITE: <https://ecai2025.org/>

Editorial Roles in Scientific Journals

Andrea Galassi is or has covered the following editorial role:

- **Section Editor** (NLP and Text Processing) for Artificial Intelligence and Law (Scimago Q1), since 2025
- **Scientific Editor** for PlosONE in 2024

Prizes and Awards

Andrea Galassi has received the following prizes & awards:

- ICAIL 2025, Peter Jackson award for **Best Innovative Application Paper**, for [C2]
- JURIX 2024, Honorable Mention for **Outstanding Research Paper**, for [W1]

Referee Services in Scientific Journals & Conferences

Andrea Galassi acted as **reviewer** or **program committee member** for the following journals (J) and conferences (C):

- (J) IEEE Transactions on Neural Networks and Learning Systems
- (J) IEEE Transactions on Emerging Topics in Computational Intelligence
- (J) IEEE Transactions on Artificial Intelligence
- (J) IEEE Transactions on Games
- (J) IEEE/ACM Transactions on Audio, Speech, and Language
- (J) IEEE Intelligent Systems
- (J) IEEE Journal of Biomedical and Health Informatics
- (J) IEEE Access
- (J) ACM Computing Surveys
- (J) Journal of Artificial Intelligence Research
- (J) Artificial Intelligence (Elsevier)
- (J) Computers in Human Behaviour (Elsevier)
- (J) Computers and Electrical Engineering (Elsevier)
- (J) Intelligent Systems with Applications (Elsevier)
- (J) Expert Systems with Applications (Elsevier)
- (J) International Journal of Approximate Reasoning (Elsevier)
- (J) Data in Brief (Elsevier)
- (J) Information Processing & Management (Elsevier)
- (J) Nature Scientific Reports (Springer)
- (J) Artificial Intelligence and Law (Springer)
- (J) Artificial Intelligence Review (Springer)
- (J) Knowledge and Information Systems (Springer)
- (J) Pattern Analysis and Applications (Springer)
- (J) Fundamenta Informaticae (IOS Press)
- (J) Intelligenza Artificiale (IOS Press)
- (J) Information (MDPI)
- (J) Applied sciences (MDPI)
- (J) Data (MDPI)
- (J) Electronics (MDPI)
- (C) IJCAI International Joint Conference on Artificial Intelligence
- (C) ECAI European Conference on Artificial Intelligence
- (C) AAAI Conference on Artificial Intelligence
- (C) COLING International Conference on Computational Linguistics
- (C) Association for Computational Linguistics (ACL) Rolling Review

- (C) EMNLP Conference on Empirical Methods in Natural Language Processing
- (C) ACL Annual Meeting of the Association for Computational Linguistics
- (C) EACL Conference of the European Chapter of the Association for Computational Linguistics
- (C) IEEE ICTAI International Conference on Tools with Artificial Intelligence
- (C) Annual ACM Symposium on Applied Computing
- (C) PAIS Conference on Prestigious Applications of Intelligent System
- (C) HHAI Hybrid Human Artificial Intelligence Conference
- (C) FEVER Fact Extraction and Verification workshop
- (C) AI*IA International Conference of the Italian Association for Artificial Intelligence
- (C) AAMAS International Conference on Autonomous Agents and Multiagent Systems
- (C) LOD International Conference on Machine Learning, Optimization, and Data Science
- (C) Workshop on NLP for Positive Impact
- (C) ASAIL Automated Semantic Analysis of Information in Legal Text

Presentations at Workshops & Conferences

Andrea Galassi **presented his research** at the following international workshops and conferences:

- First Argument Mining and Empirical Legal Research Workshop (AMELR 2025), co-located with ICAIL2025, Chicago (United States), June 2025
- 20th International Conference on Artificial Intelligence and Law (ICAIL 2025), Chicago (United States), June 2025
- Second workshop on "Large and Small Language Models: stato dell'arte fra limiti e prospettive", Bologna (Italy), March 2025.
- 10th Italian Conference on Computational Linguistics (CLiC-it 2024), Pisa (Italy), December 2024.
- FAIR General Conference, Naples (Italy), September 2024.
- 14th International Conference of the CLEF Association (CLEF 2023), Thessaloniki (Greece), September 2023,
- 31st International Joint Conference on Artificial Intelligence (IJCAI-ECAI), Vienna (Austria), July 2022.
- 3rd International Workshop Deep Learning meets Ontologies and Natural Language Processing (DeepOntoNLP 2022), co-located with ESWC2022, Hersonissos (Greece), June 2022.
- Natural Legal Language Processing Workshop (NLLP 2021), co-located with EMNLP2021, Punta Cana (Dominican Republic), November 2021.
- 5th Workshop on Natural Language for Artificial Intelligence (NL4AI 2021), co-located with 20th International Conference of the Italian Association for Artificial Intelligence (AI*IA 2021), Online event, November 2021.
- 12th International Conference of the CLEF Association (CLEF 2021), Virtual Event, September 21-24, 2021,
- 28th International Conference on Computational Linguistics (COLING 2020), Barcelona (Spain) and Online, December 2020.
- 30th European Conference On Operational Research (EURO 2019), Dublin (Ireland), June 2019.
- 5th Workshop on Argument Mining (ArgMining 2018), co-located with EMNLP 2018, Bruxelles (Belgium), November 2018.
- 16th Conference of the Italian Association for Artificial Intelligence (AI*IA 2017), Bari (Italy), November 2017.

Teaching Activities

Bachelor & Master degrees

- **Lecturer**
3 ACADEMIC YEARS: 2023/24 – 2025/26
COURSE: Real-Time Systems and Programming for Automation M, 12 CFUs, Operating System and Programming module
FRONTAL AND LAB TEACHING: 60 hours
CURRICULUM: Master degree in Automation Engineering, in Electrical Energy Engineering, and in Autonomous Systems and Intelligent Robots, University of Bologna, Italy
STUDENTS' SATISFACTION RATE: 83.0%
- **Visiting Professor**
2 ACADEMIC YEARS: 2024/25 – 2025/26 (scheduled)
COURSE: Introduction to Data Science and Machine Learning
FRONTAL TEACHING: 10 hours
CURRICULUM: Master in Law, Governance and Technology, Université Catholique de Lille, Paris, France
STUDENTS' SATISFACTION RATE: 100%
- **Adjunct Professor**
1 ACADEMIC YEAR: 2020/21
COURSE: Sistemi Operativi T (translation: Operating Systems), 9 CFUs, Module 2
FRONTAL AND LAB TEACHING: 40 hours
CURRICULUM: Bachelor degree in Computer Engineering, University of Bologna, Italy
STUDENTS' SATISFACTION RATE: 92.3%

PhD courses

- **Lecturer**
3 ACADEMIC YEARS: 2022/23 – 2025/26 (scheduled)
COURSE: Foundations of Natural Language Processing
FRONTAL TEACHING: 10 hours until 2023/2024, 14 hours in 2024/2025
CURRICULUM: Phd in Law, Sciences and Technology, University of Bologna, Italy
STUDENTS' SATISFACTION RATE: 100%
- **Lecturer**
1 ACADEMIC YEAR: 2024/25
COURSE: Introduction to Machine Learning
FRONTAL AND LAB TEACHING: 3 hours
CURRICULUM: AI and Law Summer School, European University Institute, Fiesole, Italy

Professional master programs

- **Lecturer**
2 ACADEMIC YEAR: 2024/25 – 2025/26 (scheduled)
COURSE: Introduction to Discriminative AI methods, Module 4 out of 4
FRONTAL TEACHING: 6 hours
CURRICULUM: Master in High-Performance and Quantum Computing, University of Bologna, Italy

Teaching Assistant/Tutor/Mentor

- **Teaching Assistant**
3 ACADEMIC YEARS: 2020/21 – 2022/23
COURSE: Natural Language Processing
CURRICULUM: Master Degree in Artificial Intelligence, University of Bologna, Italy
- **Teaching Assistant**
2 ACADEMIC YEARS: 2020/21 – 2021/22
COURSE: Fundamentals of Artificial Intelligence and Knowledge Representation
CURRICULUM: Master Degree in Artificial Intelligence, University of Bologna, Italy

- **Teaching Assistant**
2 ACADEMIC YEARS: 2020/21 – 2021/22
COURSE: Languages and Algorithms for Artificial Intelligence
CURRICULUM: Master Degree in Artificial Intelligence, University of Bologna, Italy
- **Teaching Assistant**
2 ACADEMIC YEARS: 2018/19 – 2019/20
COURSE: Real Time Systems for Automation M
CURRICULUM: Master degree in Automation Engineering, University of Bologna, Italy
- **Teaching Assistant**
3 ACADEMIC YEARS: 2017/18 – 2019/20
COURSE: Fondamenti di Informatica T1 (translated: Foundations of Informatics)
CURRICULUM: Bachelor degree in Computer Engineering, University of Bologna, Italy

Student Supervision Activity

Andrea Galassi has supervised or co-supervised the following **PhD students**:

- Gianmarco Pappacoda, PhD student, Nov 2023 - ongoing, main advisor: Paolo Torroni, research topic: “Methods and Techniques for Enhancing Knowledge Extraction through Large Language Models”, University of Bologna.
- Nicolò Donati, PhD student, Nov 2023 - ongoing, main advisor: Paolo Torroni, research topic: “Use of Large Language Models for the Generation and Evaluation of Textbooks Material”, University of Bologna.
- Elena Palmieri, PhD student, Nov 2022 - ongoing, main advisor: Paolo Torroni, research topic: “NLP techniques for public administration”, University of Bologna.
- Eleonora Mancini, PhD student, Nov 2022 - ongoing, main advisor: Paolo Torroni, research topic: “Data Representation, Fusion and Explainability in Multimodal Deep Learning for Natural Language Processing”, University of Bologna.

Andrea Galassi has supervised or co-supervised the following **post-graduate students and temporary research fellows**:

- Eleonora Misino, temporary research fellow (AdR), October 2024 - March 2025, main advisor: Paolo Torroni, research topic: "Natural Language Processing for the analysis of judicial principles of law", University of Bologna.
- Federico Ruggeri, temporary research fellow (AdR), April 2024 - ongoing, main advisor: Paolo Torroni, research topic: "Knowledge Extraction and Integration for assessing judicial principles", University of Bologna.
- Luca Bolognini, temporary research fellow (AdR), May 2023 - May 2025, main advisor: Paolo Torroni, research topic: "NLP Methods for Automatic Detection of Contributions to UN Sustainable Development Goals", University of Bologna.
- Giulia Grundler, temporary research fellow (AdR), July 2021 - ongoing, main advisors: Paolo Torroni and Francesca Lagioia, research topic: "Natural Language Processing for Multilingual Legal Analytics", University of Bologna.
- Elena Palmieri, temporary research fellow (AdR), March 2021 - Oct 2022, main advisor: Giovanni Sartor, research topic: "Legal Analytics for Italian Law", University of Bologna.

Andrea Galassi has supervised or co-supervised several students during their **final graduation project** (bachelor/master thesis), among which:

- Zhanel Zhexenova, master thesis in artificial intelligence, "Assessing LLM Capabilities in Low-Resource Settings: a Kazakh Language Perspective", 2025
- Marco Panarelli, master thesis in artificial intelligence, "Comparing Large Language Models on Unfair Clause Detection in Terms of Services", 2025
- Francesco Alfieri, master thesis in artificial intelligence, "Graph-Based Approaches for Few-Shot Example Selection in In-Context Learning", 2025

- Samuele Di Giacomo, master thesis in computer engineering, "Generazione e valutazione di test automatici tramite modelli di NLP avanzati", 2024
- Roberto Bonini, master thesis in artificial intelligence, "Integrating Neuro-Ocular Data for Accurate Hand Movement Decoding in Brain-Controlled Robotic Systems", 2024
- Bogdan Ivasiuk, master thesis in artificial intelligence, "Addressing Misinformation Challenges in War Scenario: Russo-Ukrainian War", 2023.
- Michele Faedi, master thesis in artificial intelligence, "Comprehensive Study Of Clinical Entity Extraction And Classification Using Large Language Models", 2023.
- Lorenzo Borelli, master thesis in artificial intelligence, "Design and implementation of a privacy-preserving dialogue system based on argumentation", 2023.
- Emmanuele Bollino, master thesis in artificial intelligence, "Automatic Terminology Coding for the Biomedical Domain", 2023.
- Elisa Ancarani, master thesis in artificial intelligence, "Argument Mining into Active Learning Systematic Reviews: unlocking the synergy between MARGOT and ASReview", 2023.
- Lorenzo Niccolai, master thesis in artificial intelligence, "Knowledge graph embedding enhancement using ontological knowledge in the biomedical domain", 2023.
- Riccardo Varotto, master thesis in computer engineering, "Tecniche di pianificazione automatica in giochi con regole dinamiche: sviluppo di un agente intelligente per la Keke Competition", 2022.
- Luca Salvatore Lorello, master thesis in artificial intelligence, "Small Transformers for Bioinformatics Tasks", 2021.
- Eleonora Mancini, master thesis in artificial intelligence, "Disruptive Situations Detection on Public Transports through Speech Emotion Recognition", 2021.
- Alessandro Pacielli, master thesis in computer engineering, "Estensione della copertura terminologica del metatesauro biomedicale UMLS mediante allineamento con Wikipedia", 2021.
- Francesco Antici, master thesis in artificial intelligence, "Advanced techniques for cross-language annotation projection in legal texts", 2021.
- Francesco Giovanelli, master thesis in computer engineering thesis, "Model Agnostic solution of CSPs with Deep Learning", 2019.
- Giacomo Pinardi, bachelor thesis in computer engineering, "Apprendimento supervisionato di un gioco da tavolo asimmetrico tramite reti neurali: un caso di studio su Tablut", 2019.
- Alessio Leurini, bachelor thesis in computer engineering, "Cross-Domain Sentiment Analysis", 2019.
- Alessandro Ravaglia, bachelor thesis in computer engineering, "Studio delle tecniche di reinforcement learning e delle loro applicazioni ai giochi da tavolo", 2019.
- Andrea Piretti, bachelor thesis in computer engineering, "Sviluppo di un'architettura software distribuita con supporto a giocatori artificiali: il caso di studio del gioco da tavolo Tablut", 2018
- Grilli Matteo, bachelor thesis in computer engineering, "Reti neurali profonde applicate a giochi di carte digitali: un caso di studio su Hearthstone semplificato", 2017

Andrea Galassi has supervised the following students during their **internships**:

- Refika Kalyoncu, June–September 2022, during her bachelor degree in Computer and Industrial Engineering, Bogazici University, Istanbul, Turkey
- Sezen Percin, June–September 2021, during her bachelor degree in Electrical Engineering, Bogazici University, Istanbul, Turkey

Dissemination Activities and “Terza Missione”

Andrea contributed to the following activities:

- **AI Explained: il futuro dell’Intelligenza**

PERIOD: 2025

DESCRIPTION: Creation of a series of educational videos about Artificial Intelligence and Natural Language Processing, intended for Italian high school students, and designed in collaboration with Zanichelli Editore. Andrea contributed both as the subject of 2 videos and the writer of the related script.

- **Rigenerativa**

PERIOD: 2025

WEBSITE: <https://sineglossa.it/progetti/rigenerativa/>

DESCRIPTION: The project consists of a training in generative artificial intelligence to strengthen the ability of Italian cultural non-profit organizations to address digital transformation in a conscious and strategic way. Andrea contributed with 5 hours of frontal teaching about the use of Large Language Models.

Projects and Repositories

Andrea Galassi is responsible or co-responsible for the following projects, software, corpora, and related repositories

- **Promoting Fairness and Diversity in Speech Datasets for Mental Health and Neurological Disorders Research**
LINK: <https://www.humane-ai.eu/project/promoting-fairness-and-diversity-in-speech-datasets-for-affective-computing/>
DESCRIPTION: Critical review and meta-analysis of existing speech datasets for Mental health and neurological disorders research in the perspective of inclusivity, transparency, and fair use.
RELATED PUBLICATIONS: [J1]
- **A Transparent and Explainable Dialogue System for Immigration Services**
LINK: <https://www.humane-ai.eu/project/a-transparent-and-explainable-dialogue-system-for-immigration-services/>
DESCRIPTION: A transparent and explainable dialogue system for assisting immigrants and non-profit organizations on administrative and legal matters in Italy.
RELATED PUBLICATIONS: [C3]
- **Multilingual Unfair Clause Detection**
LINK: <https://github.com/lt-nlp-lab-unibo/Multilingual-Unfair-Clause-Detection>
DESCRIPTION: Software for detecting unfair clauses in terms of service contracts in a multilingual context.
RELATED PUBLICATIONS: [J3]
- **AIoD Success Stories Platform**
LINK: https://gitlab.com/ai4europe_unibo
DESCRIPTION: Success Stories service for the "AI-on-demand-platform".
- **Ethical chatbots**
LINK: <https://www.ai4europe.eu/research/research-bundles/ethical-chatbots>
DESCRIPTION: Development of an ethical dialogue system, with a case study on COVID-19 vaccine information.
RELATED PUBLICATIONS: [J7, W14, W13]
- **Cross Lingual Annotation Projection**
LINK: <https://bitbucket.org/a-galaxy/cross-lingual-annotation-projection-in-legal-texts>
DESCRIPTION: Software for the projection of labels between parallel asymmetric legal documents.
RELATED PUBLICATIONS: [C12]
- **ResAttArg: Residual Attentive Deep Networks for Argument Structure Prediction**
LINK: <https://github.com/AGalassi/StructurePrediction18>
DESCRIPTION: Software for the classification of argumentative content in documents and for the prediction of its structure.
RELATED PUBLICATIONS: [W15, J5]
- **Tablut Competition**
LINK: <https://github.com/AGalassi/TablutCompetition>
DESCRIPTION: Software for the Tablut Students Competition, used as part of the Fundamentals of AI courses at the University of Bologna.
RELATED PUBLICATIONS: [W16]
- **Neural Nine Men's Morris**
LINK: <https://github.com/AGalassi/NNMM>
DESCRIPTION: Software for playing the game of Nine Men's Morris using neural networks.
RELATED PUBLICATIONS: [J11]

Andrea Galassi is a contributor to the following projects, software, corpora, and related repositories

- **Privacy Policies Compliance**
LINK: <https://github.com/nlp-unibo/Privacy-Policies-Compliance>
DESCRIPTION: Software and data related to the PRIMA project, regarding legal analytics for detecting unfair clauses in privacy policies.
RELATED PUBLICATIONS: [C1, W1]

- **Language Technologies Lab Repository**
LINK: <https://github.com/nlp-unibo>
DESCRIPTION: Official repository of the Language Technologies Lab of the University of Bologna.
 - **NewsSD-ENG**
LINK: <https://github.com/nlp-unibo/newssd-eng>
DESCRIPTION: Corpus for detecting subjectivity in English news.
RELATED PUBLICATIONS: [C4]
 - **LEXTREME**
LINK: <https://huggingface.co/datasets/joelito/lextrreme>
DESCRIPTION: Benchmark for the evaluation of multilingual language models over legal analytics tasks.
RELATED PUBLICATIONS: [C5]
 - **ADELE**
LINK: <https://github.com/adele-project>
DESCRIPTION: Software and data related to the ADELE project, regarding legal analytics for decisions and argument mining.
RELATED PUBLICATIONS: [W10, W9, W2]
 - **Check That!**
LINK: https://gitlab.com/checkthat_lab/
DESCRIPTION: Software and data for the Check That! workshop.
RELATED PUBLICATIONS: [C4, C7, C8, W5, W6]
 - **SubjectivITA**
LINK: <https://github.com/francescoantici/SubjectivITA>
DESCRIPTION: Corpus for the detection of subjectivity in Italian news.
RELATED PUBLICATIONS: [C11]
-

Research Interests

Andrea Galassi's main research interests cover different areas in the field of Artificial Intelligence (AI), mostly Natural Language Processing (NLP) and Deep Learning (DL), with a focus on Neuro-Symbolic (NeSy) approaches, Argument Mining (AM), Legal Analytics, and Trustworthy and Human-Centered AI.

Machine Learning and Deep Learning I worked on the study, development, and application of novel architectures in NLP research. I have applied residual networks to NLP tasks [W15] before the widespread adoption of Transformers, I wrote a highly-cited survey on the topic of neural attention [J8], and I designed a deep network architecture that combines the two methods [J5]. In recent years, my work has focused more on Transformer-based models, Large Language Models (LLM), and their multimodal variants, with a focus on multidisciplinary applications.

Integration of Knowledge and Neuro-Symbolic approaches. Since the beginning of my academic career, I have been interested in combining data-based approaches with techniques that rely on formal knowledge. I have investigated whether it was possible to train advanced neural networks architecture to learn to solve problems where there are constraints and objectives that are formally defined, such as solving Constraint Satisfaction Problems (CSP) [C13] or winning board games [J11]. My interest has then focused on the application of NeSy techniques to NLP tasks, in particular AM. After an initial analysis of the existing frameworks and their potential use [J10], I have implemented a solution based on Logic Tensor Networks which has shown promising results [P2]. One of the main open challenges in the application of NeSy techniques to NLP tasks remains scalability. Indeed, current frameworks and solutions are designed to work with amount of data points several orders of magnitude lower than the number of data points involved in an NLP problem.

Argument Mining. Argument Mining is the analysis and processing of textual documents with the purpose of extracting argumentative contents and their relationships. Similarly to other NLP fields, AM was traditionally addressed by using hand-made features and models, which were tailored to specific datasets and problems and lacked generality. In my work, I have researched more general solutions, designing advanced DL architectures that exploit Residual Networks and Neural Attention [W15, J5]. Among the possible applications of AM there is the ranking and selection of documents, such as scientific documents. Indeed, with the growing amount of scientific literature published every year, there is a growing need for automatic tools to retrieve relevant and high-quality documents. In this context, I have contributed to the development of a search engine for Covid-19 scientific articles in the context of the FISIR AMICA project [C10, J6]. While typically addressed through the use of DL, many challenges still stand open, as argumentation involves tasks such as reasoning, debate, and persuasion that cannot be easily addressed by deep architectures only, sophisticated as they may be. Indeed, DL models may benefit from the integration of formal knowledge, which can be given by the specific domain or by more general structures taken from the research field of computational argumentation [J10].

Trustworthy and Human-Centered AI. The expansion of intelligent technologies has been met by growing concerns about possible misuse, motivating a need to develop AI systems that are trustworthy. On the one hand, governments are pressured to gain or preserve an edge in intelligent technologies, which make intensive use of large amounts of data. On the other hand, there is an increasing awareness of the fundamental need for data protection regulations. Trustworthy AI systems need not only be robust but also respectful of all applicable laws and regulations, as well as of ethical principles and values. I had the opportunity to deepen these topics as part of my participation in the Humane AI consortium. My research has been focused on three different topics:

- I developed a dialogue system architecture specifically designed to respect the users' privacy while being explainable and auditable. The final result [J7] was a modular solution that integrated both an NLP module for understanding the specific request of the user [W14] and a reasoning module based on computational argumentation to provide the answer and its explanation [W13]. I then continued the investigation of this topic by leading a micro-project regarding the application of this architecture in the legal domain [C3].
- Another topic I addressed regards the use of speech data for social good, such as the detection of disruptive situations in public transport [J4]. I have also led a micro-project regarding the responsible creation of speech resources for the mental health domain [J1], analyzing the existing resources and proposing an actionable checklist of desiderata to follow for creating new ones.
- Finally, another dimension of human-centered AI regards citizens' empowerment in social contexts where they may be potentially threatened. Domains that have unfortunately become extremely popular in recent years are the verification information, the automatic detection of fake news, and the recognition of propaganda techniques. I have supervised two studies regarding subjectivity in news articles [C4, C11] and

the detection of fake news [W4], and I have been a co-chair and co-organizer of a workshop on these topics [C7, C8, W5, W6, W3].

Legal Analytics. Mixing NLP, AI, and data science, Legal Analytics aims to extract legal knowledge, infer unknown relations, and realize data-driven forecasting. The integration of multiple techniques from different disciplines is a key factor since it is typically difficult to produce large amounts of data, while the domain knowledge is vast and yet often not clearly defined. Moreover, the inherent trade-off between the generality of an approach and its efficacy in the context of a specific language or legislation highlights the benefits of techniques that can be used to transfer knowledge from one domain to another. Over the years, my research has covered several aspects of this field. Due to the argumentative nature of legal judgments, the legal domain is an especially interesting domain of application for AM [W10, C6, W2], which can be exploited also for more advanced tasks, such as outcome prediction [W9]. Another topic I worked on regards the extraction of the principles of law from legal documents [T2] I have also researched how to create coherent novel synthetic data by combining sub-symbolic approaches such as word embeddings with symbolic ontologies [W7]. I have worked on approaches to analyze the fairness of privacy policies and terms of services with respect to regulations such as GDPR [C2, W1]. In this context, I worked on multilingualism, developing solutions to transfer the knowledge acquired from one language to another one [C1, C12, W12]. I am also part of an international research group that designed a benchmark for the evaluation of multilingual language models over legal tasks [C5].

Education Methods. Finally, I am interested in studying and experimenting with novel educational methods that can improve students' interest and motivation. I contributed to two studies regarding the use of a board-game challenge as a tool to teach AI to master students [J2, W16] and I am directly involved in their prosecution and realization.

Research Group and Collaborations

Language Technologies Lab, University of Bologna. Andrea is the **deputy director** of the Lab led by Prof. **Paolo Torroni**, and he works with the Artificial Intelligence group led by Prof. Michela Milano. His responsibilities range from PI to Task Leader in several projects, and he actively participates in the designing and development of research studies. He participates in the strategic planning of the lab's activities and supervises and mentors students at all levels, including bachelor's, master's, and PhD students. He has been part of the lab since 2017.

European University Institute. Andrea has a long-standing collaboration with the Institute, with Prof. Giovanni Sartor and his research group, with whom he has collaborated on several national and international projects focused on legal analytics since 2019.

University of Florence. Prof. Marco Lippi has supervised Andrea during his PhD, and they continue to collaborate on studies on argument mining and legal analytics since 2017.

University of Calabria Andrea worked with this institution in the context of several projects concerning computational argumentation since 2021.

Consiglio Nazionale delle Ricerche (CNR). Andrea has worked with CNR researchers as part of Humane-AI-Net on topics such as information retrieval and computational argumentation.

Università La Sapienza. Andrea works closely with the research group of Prof. Roberto Navigli as part of the FAIR project, concerning the development of LLM for Italian, since 2024.

Fondazione Bruno Kessler (FBK). Andrea collaborates with the research group of Prof. Bernardo Magnini as part of the FAIR project and the development of a benchmark for Italian LLMs, since 2024.

University of Uppsala. Andrea has worked with this institution since 2023 on ethics and fairness in AI.

Institut national de recherche en informatique et en automatique (INRIA). Andrea collaborates with INRIA on studies concerning Human-Computer Interaction since 2024, first as part of the Humane-AI-Net consortium, and recently on research concerning LLMs.

Publications

Bibliometric indexes

Source: Scopus (<https://www.scopus.com/authid/detail.uri?authorId=57196712506>)

- Number of citations: 967
- Number of citing documents: 875
- h-index: 12

Source: Google Scholar (<https://scholar.google.com/citations?user=OnzdCscAAAAJ&hl=it>)

- Number of citations: 1700
- h-index: 17
- i10-index: 26

For each journal paper, I report the Scimago quartile score relevant to computer science and computer engineering subjects. For each conference paper, I report the score given to the conference by the GII-GRIN-SCIE ranking.

Journals

- [J1] Eleonora Mancini, Ana Tanevska, **Andrea Galassi**, Alessio Galatolo, Federico Ruggeri, and Paolo Torroni. Promoting the responsible development of speech datasets for mental health and neurological disorders research. *Journal of Artificial Intelligence Research*, 82:937–972, February 2025. Scimago quartile: Q1. **Corresponding author**.
- [J2] Allegra De Filippo, **Andrea Galassi**, Alessandro Soriani, Giada Trisolini, Federico Baldo, Federico Chesani, Paola Mello, and Michela Milano. Improving the teaching of artificial intelligence through project-based learning on a board game. *Intelligenza Artificiale*, 19(1):17–29, 2025. Co-**First author**; Scimago quartile: Q2.
- [J3] **Andrea Galassi**, Francesca Lagioia, Agnieszka Jabłonowska, and Marco Lippi. Unfair clause detection in terms of service across multiple languages. *Artificial Intelligence and Law*, April 2024. **First author**; Scimago quartile: Q1.
- [J4] Eleonora Mancini, **Andrea Galassi**, Federico Ruggeri, and Paolo Torroni. Disruptive situation detection on public transport through speech emotion recognition. *Intelligent Systems with Applications*, 21:200305, March 2024. Scimago quartile: Q1.
- [J5] **Andrea Galassi**, Marco Lippi, and Paolo Torroni. Multi-task attentive residual networks for argument mining. *IEEE ACM Transactions on Audio, Speech, and Language Processing*, 31:1877–1892, 2023. **First author**; Scimago quartile: Q1.
- [J6] Gianfranco Brambilla, Antonella Rosi, Francesco Antici, **Andrea Galassi**, Daniele Giansanti, Fabio Magurano, Federico Ruggeri, Paolo Torroni, Evaristo Cisbani, and Marco Lippi. Argument mining as rapid screening tool of covid-19 literature quality: Preliminary evidence. *Frontiers in Public Health*, 10, 2022.
- [J7] Bettina Fazzinga, **Andrea Galassi**, and Paolo Torroni. A privacy-preserving dialogue system based on argumentation. *Intelligent Systems with Applications*, 16:200113, 2022. **Equal contribution** of authors; Scimago quartile: Q1.
- [J8] **Andrea Galassi**, Marco Lippi, and Paolo Torroni. Attention in natural language processing. *IEEE Transactions on Neural Networks Learning Systems*, 32(10):4291–4308, 2021. **First author**; Scimago quartile: Q1.
- [J9] Elena Borelli, Giacomo Paolini, Francesco Antoniazzi, Marina Barbiroli, Francesca Benassi, Federico Chesani, Lorenzo Chiari, Massimiliano Fantini, Franco Fuschini, **Andrea Galassi**, Gian Andrea Giacobone, Silvia Imbesi, Melissa Licciardello, Daniela Loreti, Michele Marchi, Diego Masotti, Paola Mello, Sabato Mellone, Giuseppe Mincoelli, Carla Raffaelli, Luca Roffia, Tullio Salmon Cinotti, Carlo Tacconi, Paola Tamburini, Marco Zoli, and Alessandra Costanzo. HABITAT: an IoT solution for independent elderly. *Sensors*, 19(5):1258, 2019.

- [J10] **Andrea Galassi**, Kristian Kersting, Marco Lippi, Xiaoting Shao, and Paolo Torroni. Neural-symbolic argumentation mining: An argument in favor of deep learning and reasoning. *Frontiers Big Data*, 2:52, 2019. **First author**; Scimago quartile: Q1.
- [J11] Federico Chesani, **Andrea Galassi**, Marco Lippi, and Paola Mello. Can deep networks learn to play by the rules? A case study on nine men’s morris. *IEEE Transactions on Games*, 10(4):344–353, 2018. **Equal contribution** of authors; Scimago quartile: Q2.

Main Conferences

- [C1] Giulia Grundler, Mariaceleste Musicco, **Andrea Galassi**, Francesca Lagioia, Rūta Liepiņa, Giorgio Resta, Sara Roccu, Giovanni Sartor, and Paolo Torroni. Detecting vague clauses in italian privacy policies using transformers, llms, and cross-lingual techniques. In *ECAI*, October 2025. **Corresponding author**; Conference GGS Rating: A-; accepted.
- [C2] Marco Panarelli, **Andrea Galassi**, Francesca Lagioia, Rūta Liepiņa, Marco Lippi, Przemysław Palka, and Giovanni Sartor. Is it worth using llms for unfair clause detection in terms of service? In *International Conference on Artificial Intelligence and Law (ICAAIL)*, June 2025. **Corresponding author**; Conference GGS Rating: B-; Awarded as **Best Innovative Application Paper**.
- [C3] Bettina Fazzinga, Elena Palmieri, Margherita Vestoso, Luca Bolognini, **Andrea Galassi**, Filippo Furfaro, and Paolo Torroni. A chatbot for asylum-seeking migrants in europe. In *2024 IEEE 36th International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 702–707, October 2024. **Corresponding author**; Conference GGS Rating: B.
- [C4] Francesco Antici, Federico Ruggeri, **Andrea Galassi**, Katerina Korre, Arianna Muti, Alessandra Bardi, Alice Fedotova, and Alberto Barrón-Cedeño. A corpus for sentence-level subjectivity detection on English news articles. In Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue, editors, *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 273–285, Torino, Italy, May 2024. ELRA and ICCL. Conference GGS Rating: A.
- [C5] Joel Niklaus, Veton Matoshi, Pooja Rani, **Andrea Galassi**, Matthias Stürmer, and Ilias Chalkidis. LEXTREME: A multi-lingual and multi-task benchmark for the legal domain. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, volume Findings, December 2023. Conference GGS Rating: A+.
- [C6] Piera Santin, Giulia Grundler, **Andrea Galassi**, Federico Galli, Francesca Lagioia, Elena Palmieri, Federico Ruggeri, Giovanni Sartor, and Paolo Torroni. Argumentation structure prediction in cjeu decisions on fiscal state aid. In *International Conference on Artificial Intelligence and Law*, pages 247—256, Braga, Portugal, June 2023. **Corresponding author**; Conference GGS Rating: B-.
- [C7] Alberto Barrón-Cedeño, Firoj Alam, Tommaso Caselli, Giovanni Da San Martino, Tamer Elsayed, **Andrea Galassi**, Fatima Haouari, Federico Ruggeri, Julia Maria Struß, Rabindra Nath Nandi, Gullal S. Cheema, Dilshod Azizov, and Preslav Nakov. The CLEF-2023 checkthat! lab: Checkworthiness, subjectivity, political bias, factuality, and authority. In *Advances in Information Retrieval - 45th European Conference on Information Retrieval, ECIR 2023, Dublin, Ireland, April 2-6, 2023, Proceedings, Part III*, volume 13982 of *Lecture Notes in Computer Science*, pages 506–517. Springer, April 2023. Conference GGS Rating: A-.
- [C8] Alberto Barrón-Cedeño, Firoj Alam, **Andrea Galassi**, Giovanni Da San Martino, Preslav Nakov, Tamer Elsayed, Dilshod Azizov, Tommaso Caselli, Gullal S. Cheema, Fatima Haouari, Maram Hasanain, Muċahid Kutlu, Chengkai Li, Federico Ruggeri, Julia Maria Struß, and Wajdi Zaghouni. Overview of the CLEF-2023 checkthat! lab on checkworthiness, subjectivity, political bias, factuality, and authority of news articles and their source. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction - 14th International Conference of the CLEF Association, CLEF 2023, Thessaloniki, Greece, September 18-21, 2023, Proceedings*, volume 14163 of *Lecture Notes in Computer Science*, pages 251–275. Springer, 2023. Conference GGS Rating: B;.
- [C9] Mirko Del Moro, Serban Cristian Tudosie, Francesco Vannoni, **Andrea Galassi**, and Federico Ruggeri. Inception models for fashion image captioning: An extensive study on multiple datasets. In Avi Arampatzis, Evangelos Kanoulas, Theodora Tsikrika, Stefanos Vrochidis, Anastasia Giachanou, Dan Li, Mohammad Aliannejadi, Michalis Vlachos, Guglielmo Faggioli, and Nicola Ferro, editors, *Experimental IR Meets Multilinguality, Multimodality, and Interaction - 14th International Conference of the CLEF Association*,

CLEF 2023, Thessaloniki, Greece, September 18-21, 2023, Proceedings, volume 14163 of *Lecture Notes in Computer Science*, pages 3–14. Springer, 2023. **Corresponding author**; conference GGS Rating: B.

- [C10] Marco Lippi, Francesco Antici, Gianfranco Brambilla, Evaristo Cisbani, **Andrea Galassi**, Daniele Giansanti, Fabio Magurano, Antonella Rosi, Federico Ruggeri, and Paolo Torroni. AMICA: an argumentative search engine for COVID-19 literature. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 5932–5935. ijcai.org, 2022. Conference GGS Rating: A++.
- [C11] Francesco Antici, Luca Bolognini, Matteo Antonio Inajetovic, Bogdan Ivasiuk, **Andrea Galassi**, and Federico Ruggeri. Subjectivita: An italian corpus for subjectivity detection in newspapers. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction - 12th International Conference of the CLEF Association, CLEF 2021, Virtual Event, September 21-24, 2021, Proceedings*, volume 12880 of *Lecture Notes in Computer Science*, pages 40–52. Springer, 2021. **Corresponding author**; conference GGS Rating: B.
- [C12] **Andrea Galassi**, Kasper Drazewski, Marco Lippi, and Paolo Torroni. Cross-lingual annotation projection in legal texts. In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 915–926. International Committee on Computational Linguistics, 2020. **First author**; conference GGS Rating: A.
- [C13] **Andrea Galassi**, Michele Lombardi, Paola Mello, and Michela Milano. Model agnostic solution of cpsps via deep learning: A preliminary study. In *Integration of Constraint Programming, Artificial Intelligence, and Operations Research - 15th International Conference, CPAIOR 2018, Delft, The Netherlands, June 26-29, 2018, Proceedings*, volume 10848 of *Lecture Notes in Computer Science*, pages 254–262. Springer, 2018. **First author**; conference GGS Rating: B-.

Other Archival Conferences & Workshops

- [W1] Giulia Grundler, Rūta Liepiņa, Mariaceleste Musicco, Francesca Lagioia, **Andrea Galassi**, Giovanni Sartor, and Paolo Torroni. Detecting vague clauses in privacy policies: The analysis of data categories using bert models and llms. In *Legal Knowledge and Information Systems - JURIX*, volume 395 of *Frontiers in Artificial Intelligence and Applications*, pages 72–83. IOS Press, 2024. **Corresponding author**. Awarded the Honorable Mention for **Outstanding Research Paper**.
- [W2] Giulia Grundler, **Andrea Galassi**, Piera Santin, Alessia Fidelangeli, Federico Galli, Elena Palmieri, Francesca Lagioia, Giovanni Sartor, and Paolo Torroni. AMELIA - argument mining evaluation on legal documents in italian: A CALAMITA challenge. In Felice Dell’Orletta, Alessandro Lenci, Simonetta Montemagni, and Rachele Sprugnoli, editors, *Proceedings of the Tenth Italian Conference on Computational Linguistics (CLiC-it 2024), Pisa, Italy, December 4-6, 2024*, volume 3878 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2024. **Corresponding author**.
- [W3] Julia Maria Struß, Federico Ruggeri, Alberto Barrón-Cedeño, Firoj Alam, Dimitar Dimitrov, **Andrea Galassi**, Georgi Pachov, Ivan Koychev, Preslav Nakov, Melanie Siegel, Michael Wiegand, Maram Hasanain, Reem Suwaileh, and Wajdi Zaghouni. Overview of the CLEF-2024 checkthat! lab task 2 on subjectivity in news articles. In *CLEF (Working Notes)*, volume 3740 of *CEUR Workshop Proceedings*, pages 287–298. CEUR-WS.org, 2024.
- [W4] Nina Khairova, Bogdan Ivasiuk, Fabrizio Lo Scudo, Carmela Comito, and **Andrea Galassi**. A first attempt to detect misinformation in Russia-Ukraine war news through text similarity. In Sara Carvalho, Anas Fahad Khan, Ana Ostroški Anić, Blerina Spahiu, Jorge Gracia, John P. McCrae, Dagmar Gromann, Barbara Heinisch, and Ana Salgado, editors, *Proceedings of the 4th Conference on Language, Data and Knowledge*, pages 559–564, Vienna, Austria, September 2023. NOVA CLUNL, Portugal.
- [W5] Federico Ruggeri, Francesco Antici, **Andrea Galassi**, Katerina Korre, Arianna Muti, and Alberto Barrón-Cedeño. On the definition of prescriptive annotation guidelines for language-agnostic subjectivity detection. In Ricardo Campos, Alípio Mário Jorge, Adam Jatowt, Sumit Bhatia, and Marina Litvak, editors, *Proceedings of Text2Story - Sixth Workshop on Narrative Extraction From Texts held in conjunction with ECIR*, volume 3370 of *CEUR Workshop Proceedings*, pages 103–111. CEUR-WS.org, 2023.
- [W6] **Andrea Galassi**, Federico Ruggeri, Alberto Barrón-Cedeño, Firoj Alam, Tommaso Caselli, Mücahid Kutlu, Julia Maria Struß, Francesco Antici, Maram Hasanain, Juliane Köhler, Katerina Korre, Folkert Leistra, Arianna Muti, Melanie Siegel, Mehmet Deniz Türkmen, Michael Wiegand, and Wajdi Zaghouni. Overview of the CLEF-2023 checkthat! lab: Task 2 on subjectivity detection. In Mohammad Aliannejadi,

Guglielmo Faggioli, Nicola Ferro, and Michalis Vlachos, editors, *Working Notes of the Conference and Labs of the Evaluation Forum (CLEF 2023)*, Thessaloniki, Greece, September 18th to 21st, 2023, volume 3497 of *CEUR Workshop Proceedings*, pages 236–249. CEUR-WS.org, 2023.

- [W7] Sezen Perçin, **Andrea Galassi**, Francesca Lagioia, Federico Ruggeri, Piera Santin, Giovanni Sartor, and Paolo Torroni. Combining WordNet and word embeddings in data augmentation for legal texts. In *Proceedings of the Natural Legal Language Processing Workshop 2022*, pages 47–52, Abu Dhabi, United Arab Emirates (Hybrid), December 2022. Association for Computational Linguistics. **Corresponding author**.
- [W8] Pavlo Seroyizhko, Zhanel Zhexenova, Muhammad Zohaib Shafiq, Fabio Merizzi, **Andrea Galassi**, and Federico Ruggeri. A sentiment and emotion annotated dataset for bitcoin price forecasting based on Reddit posts. In *Proceedings of the Fourth Workshop on Financial Technology and Natural Language Processing (FinNLP)*, pages 203–210, Abu Dhabi, United Arab Emirates (Hybrid), December 2022. Association for Computational Linguistics. **Corresponding author**.
- [W9] Federico Galli, Giulia Grundler, Alessia Fidelangeli, **Andrea Galassi**, Francesca Lagioia, Elena Palmieri, Federico Ruggeri, Giovanni Sartor, and Paolo Torroni. Predicting outcomes of italian VAT decisions. In *Legal Knowledge and Information Systems - JURIX 2022: The Thirty-fifth Annual Conference, Saarbrücken, Germany, 14-16 December 2022*, volume 362 of *Frontiers in Artificial Intelligence and Applications*, pages 188–193. IOS Press, 2022. **Corresponding author**.
- [W10] Giulia Grundler, Piera Santin, **Andrea Galassi**, Federico Galli, Francesco Godano, Francesca Lagioia, Elena Palmieri, Federico Ruggeri, Giovanni Sartor, and Paolo Torroni. Detecting arguments in CJEU decisions on fiscal state aid. In *Proceedings of the 9th Workshop on Argument Mining, ArgMining@COLING 2022, Online and in Gyeongju, Republic of Korea, October 12 - 17, 2022*, pages 143–157. International Conference on Computational Linguistics, 2022. **Corresponding author**.
- [W11] Eleonora Mancini, Federico Ruggeri, **Andrea Galassi**, and Paolo Torroni. Multimodal argument mining: A case study in political debates. In *Proceedings of the 9th Workshop on Argument Mining, ArgMining@COLING 2022, Online and in Gyeongju, Republic of Korea, October 12 - 17, 2022*, pages 158–170. International Conference on Computational Linguistics, 2022.
- [W12] Kasper Drawzeski, **Andrea Galassi**, Agnieszka Jablonowska, Francesca Lagioia, Marco Lippi, Hans Wolfgang Micklitz, Giovanni Sartor, Giacomo Tagiuri, and Paolo Torroni. A corpus for multilingual analysis of online terms of service. In *Proceedings of the Natural Legal Language Processing Workshop 2021*, pages 1–8, Punta Cana, Dominican Republic, nov 2021. Association for Computational Linguistics. **Equal contribution** of authors.
- [W13] Bettina Fazzinga, **Andrea Galassi**, and Paolo Torroni. An argumentative dialogue system for COVID-19 vaccine information. In *Logic and Argumentation - 4th International Conference, CLAR 2021, Hangzhou, China, October 20-22, 2021, Proceedings*, volume 13040 of *Lecture Notes in Computer Science*, pages 477–485. Springer, 2021. **Equal contribution** of authors.
- [W14] Bettina Fazzinga, **Andrea Galassi**, and Paolo Torroni. A preliminary evaluation of a privacy-preserving dialogue system. In *Proceedings of the Fifth Workshop on Natural Language for Artificial Intelligence (NL4AI 2021) co-located with 20th International Conference of the Italian Association for Artificial Intelligence (AI*IA 2021), Online event, November 29, 2021*, volume 3015 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2021. **Equal contribution** of authors.
- [W15] **Andrea Galassi**, Marco Lippi, and Paolo Torroni. Argumentative link prediction using residual networks and multi-objective learning. In Noam Slonim and Ranit Aharonov, editors, *Proceedings of the 5th Workshop on Argument Mining, ArgMining@EMNLP 2018, Brussels, Belgium, November 1, 2018*, pages 1–10. Association for Computational Linguistics, 2018. **First author**.
- [W16] Federico Chesani, **Andrea Galassi**, Paola Mello, and Giada Trisolini. A game-based competition as instrument for teaching artificial intelligence. In *AI*IA 2017 Advances in Artificial Intelligence - XVIth International Conference of the Italian Association for Artificial Intelligence, Bari, Italy, November 14-17, 2017, Proceedings*, volume 10640 of *Lecture Notes in Computer Science*, pages 72–84. Springer, 2017. **Equal contribution** of authors.

Other Peer-reviewed Publications

- [P1] **Andrea Galassi**. *Deep Networks and Knowledge: from Rule Learning to Neural-Symbolic Argument Mining*. PhD thesis, University of Bologna, Italy, 2021.
- [P2] **Andrea Galassi**, Marco Lippi, and Paolo Torroni. Investigating logic tensor networks for neural-symbolic argument mining. In *Tenth International Workshop on Statistical Relational AI*, 2021.

Technical Reports and Manuscripts Under Review

- [T1] Luca Bolognini, Elena Palmieri, Nicolò Donati, Giulia Grundler, Gianmarco Pappacoda, Federico Ruggeri, **Andrea Galassi**, and Paolo Torroni. Assessing the contribution of scientific articles to the un sustainable development goals. 2025. **Corresponding author**, under review.
- [T2] Sezen Percin, Piera Santin, **Andrea Galassi**, Francesca Lagioia, Ruta Liepina, Federico Ruggeri, Giovanni Sartor, and Paolo Torroni. Automatic extraction of legal principles of the court of justice of the european union. 2024. **Corresponding author**, under review.
- [T3] **Andrea Galassi**. An upper bound on the complexity of tablut. *CoRR*, abs/2101.11934, 2021.