

Curriculum Vitae et Studiorum

NAME	Andrea Galassi
YEAR OF BIRTH	1992
CITIZENSHIP	Italian
AFFILIATION	DISI, University of Bologna
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 Engineering from the University of Bologna (Italy), in 2017 and 2021 respectively. Since 2023 he is employed as a *Junior Assistant Professor ("RTDA")* at the Department of Computer Science and Engineering (DISI) of the University of Bologna.

His research interests are focused on **Natural Language Processing** and **Deep Learning architectures**, with particular attention to Neuro-Symbolic approaches to machine learning, Trustworthy and Human-Centered AI, Argument Mining, 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 from textual documents. Most of his research has been focused on the extraction, processing, and use of argumentative content; with applications that range from problems of information retrieval, such as ranking scientific documents or providing trustworthy answers in dialogues, to tasks of forecasting and classification, such as predicting the outcome of a legal judgment.

Contents

Position & Education	1
Professional Activities	4
Teaching Activities	11
Research Interests	14
Publications	16

Position & Education

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
- JANUARY 2021 – FEBRUARY 2023
Post-Doctoral Research Fellow at Department of Computer Science and Engineering (DISI), University of Bologna, Italy
SUPERVISOR: Prof. Paolo Torrioni
- NOVEMBER 2021
Collaborator with University of Modena and Reggio Emilia
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 Torrioni, 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

Period spent abroad

- FEBRUARY 2020 – AUGUST 2020
Imperial College;
SUPERVISOR: Francesca Toni, Professor in Computational Logic

 - MAY 2018 – AUGUST 2018
Stanford University;
SUPERVISOR: Margaret Hagan, Director of the Legal Design Lab
-

Professional Activities

Contribution to National & International Academic Research Projects

INTERNATIONAL PROJECTS:

- **H2020-ICT-2018-825619-AI4EU: A European AI On Demand Platform and Ecosystem**

PERIOD: 2023 – present

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 is among the managers of 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 – present

WEBSITE: <https://cordis.europa.eu/project/id/101017142>

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 is 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 automatically associate use case documents to available AI assets. He has also contributed to work package 8, writing deliverable 8.4 regarding the "International Networking Strategy".

- **H2020-ICT-2020-952026-HumanE AI Network**

PERIOD: 2021 – present

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 is mainly involved in work package 5, relative to the ethical, legal, and social challenges in artificial intelligence. He has contributed to the network with 4 micro-projects. The first was on the development of ethical human-centered chatbots combining natural language

understanding with argumentative-based reasoning. The second one was about the detection of misinformation connected to the war between Russia and Ukraine. The last two are still ongoing and concern the analysis of fairness in speech datasets and the application of the human-centered chatbot to the context of legal documents. His work resulted in 4 publications [C14, C15], among which one in a prestigious journal [J3]

- **JUST-JACC-EJU-AG-2020-1010074206-Adele: Analytics for DEcision of LEgal cases**

PERIOD: 2021 – 2023

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 is part of the Natural Language Processing team of Unibo, which is mainly involved in the tasks of outcome prediction and argument mining. Besides contributing actively to the development of solutions, Andrea’s duties include the supervision of interns and research fellows working on the project and the coordination with the team of legal experts of Unibo. His work led to the publication of one workshop article [C9] and one conference article [C1].

- **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 is it 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 [C12, C16].

NATIONAL PROJECTS:

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

PERIOD: 2022 – present

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's work concern Spoke 8 of the project, which is focused on Pervasive AI, and in particular on 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

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

PERIOD: 2022 – present

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 [C6, C8] 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 prestigious computer science conference [C10] and one in a medical journal [J2].

Reviewer for Grants and Research Proposals

Andrea Galassi has served as a reviewer for national and international grants for the following institutions:

- National Science Centre (NCN) of Poland

Program Chair & Organization Committees

Andrea Galassi has been **co-chair** of:

- CheckThat! Lab 2023, International Workshop on Fact-Checking
Co-located with CLEF 2023, Conference and Labs of the Evaluation Forum
September 2023, Thessaloniki, Greece
WEBSITE: <https://checkthat.gitlab.io/>

Andrea Galassi has been a member of the **program committee** of:

- IEEE International Conference on Tools with Artificial Intelligence (ICTAI), since 2020
- International Workshop on Natural Language Processing for Social Media (SocialNLP), in 2021
- Workshop on Argument Mining (ArgMining), in 2022

Presentations at Workshops & Conferences

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

- Touché: workshop on Argument Retrieval
Title of the intervention: "Argumentative Ranking: Case Studies and Challenges"
Bologna, September 2022
Organized by Alex Bondarenko, Matthias Hagen, Martin Potthast, Benno Stein (Webis Group)
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
Hosted online, May 2021
Organized by Catalina Goanta
Website: <https://nllpw.org/talks/>

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

- 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. Hangzhou (China), October 2021.
- 12th International Conference of the CLEF Association, (CLEF), 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), 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.

Referee Services in Journals & Conferences

Andrea Galassi is a **reviewer** for the following journals (J) and conferences (C):

- (J) Nature Scientific Reports
- (J) IEEE Transactions on Emerging Topics in Computational Intelligence
- (J) IEEE Transactions on Artificial Intelligence
- (J) IEEE Transactions on Games
- (J) IEEE Intelligent Systems
- (J) IEEE Access
- (J) ACM Computing Surveys
- (J) Journal of Artificial Intelligence Research
- (J) Artificial Intelligence (Elsevier)
- (J) Computers in Human Behaviour (Elsevier)
- (J) Intelligent Systems with Applications (Elsevier)
- (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) 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 Systems
- (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

Projects and Repositories

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

- **AIoD Success Stories Platform**
LINK: https://gitlab.com/ai4europe_unibo/success_stories_service_frontend; https://gitlab.com/ai4europe_unibo/success_stories_service_backend
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: [J3, C15, C14]
- **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: [C16]
- **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: [C17, J1]
- **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: [C19]

- **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: [J7]

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

- **SubjectivITA**

LINK: <https://github.com/francescoantici/SubjectivITA>

DESCRIPTION: Corpus for the detection of subjectivity in Italian news.

RELATED PUBLICATIONS: [C13]

- **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: [C9, C8]

- **Check That 2023**

LINK: https://gitlab.com/checkthat_lab/clef2023-checkthat-lab

DESCRIPTION: Software and data for the Check That! 2023 workshop.

RELATED PUBLICATIONS: [C2, C5]

- **LEXTREME**

LINK: <https://huggingface.co/datasets/joelito/lextrema>

DESCRIPTION: Benchmark for the evaluation of multilingual language models over legal analytics tasks.

RELATED PUBLICATIONS: [T1]



Teaching Activities

In the academia (bachelor & master degree)

- **Lecturer**

ACADEMIC YEAR 2023/2024

COURSE: Real-Time Systems for Automation M, 12 CFUs (60 hours)

CURRICULUM: Master degree in Automation Engineering and master degree in Electrical Energy Engineering, University of Bologna, Italy

- **Adjunct Professor**

ACADEMIC YEAR 2020/2021

COURSE: Sistemi Operativi T (translation: Operating Systems), 9 CFUs (40 hours)

CURRICULUM: Bachelor degree in Computer Engineering, University of Bologna, Italy

STUDENTS' SATISFACTION RATE: 92.3%

In the academia (PhD courses)

- **Lecturer**

COURSE: Foundations of Natural Language Processing (10 hours)

CURRICULUM: Phd in Law, Sciences and Technology, University of Bologna, Italy

Course held in Spring 2023

In the academia (Teaching Assistant/Tutor/Mentor)

- **Teaching Assistant**

ACADEMIC YEARS 2020/2021 - 2022/2023

COURSE: Natural Language Processing

CURRICULUM: Master Degree in Artificial Intelligence, University of Bologna, Italy

- **Teaching Assistant**

ACADEMIC YEARS 2020/2021 - 2021/2022

COURSE: Fundamentals of Artificial Intelligence and Knowledge Representation

CURRICULUM: Master Degree in Artificial Intelligence, University of Bologna, Italy

- **Teaching Assistant**

ACADEMIC YEARS 2020/2021 - 2021/2022

COURSE: Languages and Algorithms for Artificial Intelligence

CURRICULUM: Master Degree in Artificial Intelligence, University of Bologna, Italy

- **Teaching Assistant**

ACADEMIC YEARS 2018/2019 - 2019/2020

COURSE: Real Time Systems for Automation M

CURRICULUM: Master degree in Automation Engineering, University of Bologna, Italy

- **Teaching Assistant**

ACADEMIC YEARS 2017/2018 - 2019/2020

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 **post-graduate students and temporary research fellows**:

- 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.
- Giulia Grundler, temporary research fellow (AdR), July 2021 - Dec 2022, main advisor: Paolo Torroni, 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) many students during their **final graduation project** (bachelor/master thesis), among which:

- Luca Salvatore Lorello, master thesis in artificial intelligence, "Small Transformers for Bioinformatics Tasks", 2021.
- Francesco Antici, master thesis in artificial intelligence, "Advanced techniques for cross-language annotation projection in legal texts", 2021.
- Francesco Giovanelli, master 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
Nicola Alessi, project activity, 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**:

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

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.

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) [C18] or winning board games [J7]. A possible way of injecting or extracting knowledge in DL architecture is through Neural Attention, a popular module that is now part of most state-of-the-art solutions. I have deeply studied such a method, especially focusing on its application to NLP, resulting in the writing of a survey on the topic [J4]. 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 [J6], 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 that are 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 [C17, J1]. 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, J2]. 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 [J6].

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 for gaining or preserving 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, and in this context, I developed a dialogue system architecture specifically designed to respect the users' privacy while being explainable and auditable. The final result [J3] was a modular solution that integrated both

a NLP module for understanding the specific request of the user [C15] and a reasoning module based on computational argumentation to provide the answer and its explanation [C14]. Another dimension of human-centered AI regards the use of AI tools to empower citizens 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. After supervising a study regarding subjectivity in news articles [C13], I have been involved as co-chair in the organization of a workshop on these topics [C2, C3, C5].

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. Due to the argumentative nature of legal judgments, the legal domain is an especially interesting domain of application for AM [C9, C1], which can be exploited also for more advanced tasks, such as outcome prediction [C8]. I have also researched how to combine sub-symbolic approaches such as word embeddings with symbolic ontologies to create new synthetic data that are specific of the domain [C6]. For what concern multilingualism, I have worked on the development of solutions to transfer the knowledge acquired from one language to another one, by exploiting the projection of labels through similarity-based matching between textual content [C16, C12]. I am also part of an international research group that designed a benchmark for the evaluation of multilingual language models over legal tasks [T1].

Education Methods. Finally, unrelated to the rest of my research, I am interested in studying and experimenting with novel methods of education, which can improve the interest and motivation of students. I contributed to a study regarding the use of a board-game challenge as a tool to teach AI to master student [C19] and I am directly involved in its prosecution and realization.

Publications

Journals

- [J1] Andrea Galassi, Marco Lippi, and Paolo Torrioni. 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.
- [J2] Gianfranco Brambilla, Antonella Rosi, Francesco Antici, Andrea Galassi, Daniele Giansanti, Fabio Magurano, Federico Ruggeri, Paolo Torrioni, Evaristo Cisbani, and Marco Lippi. Argument mining as rapid screening tool of covid-19 literature quality: Preliminary evidence. *Frontiers in Public Health*, 10, 2022.
- [J3] Bettina Fazzinga, Andrea Galassi, and Paolo Torrioni. A privacy-preserving dialogue system based on argumentation. *Intelligent Systems with Applications*, 16:200113, 2022. Equal contribution of authors. Scimago Quartile: Q1.
- [J4] Andrea Galassi, Marco Lippi, and Paolo Torrioni. Attention in natural language processing. *IEEE Transactions on Neural Networks Learning Systems*, 32(10):4291–4308, 2021. First Author. Scimago Quartile: Q1.
- [J5] Elena Borelli, Giacomo Paolini, Francesco Antoniazzi, Marina Barbiroli, Francesca Benassi, Federico Chesani, Lorenzo Chiari, Massimiliano Fantini, Franco Fuschini, Andrea Galassi, Gian Andrea Jacobone, 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. Scimago Quartile: Q1.
- [J6] Andrea Galassi, Kristian Kersting, Marco Lippi, Xiaoting Shao, and Paolo Torrioni. Neural-symbolic argumentation mining: An argument in favor of deep learning and reasoning. *Frontiers Big Data*, 2:52, 2019. First Author. Scimago Quartile: Q1.
- [J7] 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.

Conferences & Workshops

- [C1] Piera Santin, Giulia Grundler, Andrea Galassi, Federico Galli, Francesca Lagioia, Elena Palmieri, Federico Ruggeri, Giovanni Sartor, and Paolo Torrioni. Argumentation structure prediction in cjeu decisions on fiscal state aid. In *ICAAIL*, pages 247—256, Braga, Portugal, June 2023. Corresponding Author. Conference GGS Rating: B-.
- [C2] 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, 2023. Conference GGS Rating: A-.
- [C3] 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, Mücahid 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 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 251–275. Springer, 2023. Conference GGS Rating: B.
- [C4] 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.
- [C5] 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.
- [C6] 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.
- [C7] Pavlo Seroyzhko, 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.
- [C8] 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.

- [C9] 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.
- [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] 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.
- [C12] 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.
- [C13] 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.
- [C14] 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.
- [C15] 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.
- [C16] 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.
- [C17] 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.

- [C18] 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-.
- [C19] 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] 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. *CoRR*, abs/2301.13126, 2023.
- [T2] Andrea Galassi. An upper bound on the complexity of tablut. *CoRR*, abs/2101.11934, 2021.