

PERSONAL INFORMATION **Debora Nozza**

## EDUCATION

- Sep 2022 - present **Assistant Professor in Computing Sciences**, *Bocconi University*, Milan.  
My research interests regard the detection and counter-acting of hate speech in text and bias in Machine Learning models which have become fundamental tasks for fighting against several types of discrimination (misogyny, xenophobia, racism, etc.) that subsist in Social Media in multilingual context.
- Feb 2020 - Aug 2022 **Research Fellow**, *Bocconi University*, Milan.  
Supervisor: Prof. Dirk Hovy.  
Academic Discipline (SSD): INF/01.  
Maternity leave: 23/05/2021-23/01/2022.  
Working for the ERC Starting Grant INTEGRATOR, Incorporating Demographic Factors into Natural Language Processing Models.
- Jan 2018 - Jan 2020 **Research Fellow in Computer Science**, *University of Milano-Bicocca*, Milan.  
Title: "Deep Learning Representation for Large Attributed Graphs".  
Supervisors: Prof. Enza Messina and Prof. Elisabetta Fersini.  
Academic Discipline (SSD): MAT/09, INF/01.  
My research project aimed at designing and developing novel Deep Learning unsupervised models for learning an embedding representation from large heterogeneous attributed graphs, considering both the structure of the graph and the attributes associated with each node.
- Dec 2014 - Oct 2018 **PhD in Computer Science**, *University of Milano-Bicocca*, Milan.  
Thesis Title: "Deep Learning for Feature Representation in Natural Language Processing".  
Supervisors: Prof. Enza Messina and Prof. Elisabetta Fersini.  
Academic Discipline (SSD): INF/01.  
My research activity focused on the exploitation and improvements of data representation obtained by Deep Learning techniques in order to address Natural Language Processing tasks, i.e. Sentiment Analysis, Named Entity Recognition and Linking, Hate Speech detection.
- Feb 2017 - May 2017 **Visiting PhD Student**, *INRIA*, Sophia Antipolis – Méditerranée.  
Visiting PhD student in the WIMMICS group under the supervision of Dr. Fabien Gandon.
- 2011 - 2014 **Master's Degree in Computer Science**, *University of Milano-Bicocca*, Milan.  
110/110 summa cum laude.  
Thesis Title: "Learning representation for sentiment classification".  
Supervisors: Prof. Enza Messina and Prof. Patrick Gallinari (University Pierre and Marie Curie, Paris).
- July 2013 - Feb 2014 **Visiting Master Student**, *University Pierre and Marie Curie*, Paris.  
Internship Student at Laboratoire d'Informatique de Paris 6 (LIP6) in the Machine Learning and Information Access research group, under the supervision of Prof. Patrick Gallinari.
- 2008 - 2011 **Bachelor's Degree in Computer Science**, *University of Milano-Bicocca*, Milan.  
Thesis Title: "Experimenting techniques of constrained inference for Information Extraction".  
Supervisors: Prof. Enza Messina and Prof. Elisabetta Fersini.

## RESEARCH ACTIVITIES

## Awards and Scholarships

- **Grant** from Fondazione Cariplo for MONICA (MONItoring Coverage, Attitudes and Accessibility of Italian measures in response to COVID-19), **PI** (120,000 EUR).
- Google Cloud Platform credits (5,000 EUR).
- *Runner up for the Best Paper Award* at the 2019 IEEE/WIC/ACM International Conference on Web Intelligence (WI 2019).
- *Best Paper Award* at the 8th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (2016).
- PhD Scholarship funded by the Italian Ministry of Research and Education.

## Talks and Seminars

- **Invited Talk** “Gli algoritmi hanno pregiudizi? Le conseguenze dell’intelligenza artificiale” at *University of Milano Bicocca* organized by the Department of Business and Law (Online, December 3, 2021)
- Talk “Exposing the limits of Zero-shot Cross-lingual Hate Speech Detection” at the *59th Annual Meeting of the Association for Computational Linguistics (ACL 2021)* (Online, August 2, 2021).
- **Podcast** participation “The Potential Pitfalls of Algorithms” in the *Bocconi Clarity in a Messy World* podcast. (Mar 10, 2021).
- **Tutorial** “JAX for Natural Language Processing” at the *M2L Summer School in collaboration with DeepMind*. (Online, January 4, 2021).
- **Podcast** participation “Etica e bias nell’intelligenza artificiale” in the *Smarter Podcast: the Italian podcast on Artificial Intelligence*. (Dec 10, 2020).
- **Invited Talk** “Directions and Challenges for Misogyny Detection: from corpus construction to prediction models” at the *Alan Turing Institute* organized by the Social Data Science Interest Group (Online, November 20, 2020).
- **Invited Talk** “Making Sense of Language-Specific Contextual Representation Models” at the *Workshop #1 on Named Entity Recognition in the Context of Machine Translation* organized by the Contracting Authority DG CNECT of European Commission and the company Cross-Lang (Online, June 5, 2020).
- **Invited Talk** “Deep Learning for Feature Representation in Natural Language Processing” at *Prometeia* (Milan, Italy, Oct 2020).
- Talk “Towards a Constrained deep Attributed Graph Embedding” at the *49th Annual Conference of the Italian Operations Research Society* (Genoa, Italy, September 5, 2019).
- Talk “Word Embeddings for Unsupervised Named Entity Linking” at the *12th International Conference on Knowledge Science, Engineering and Management* (Athens, Greece, August 30, 2019).
- Talk “A study on biased words for Hate Speech against women” for the Milan Women in Machine Learning and Data Science community event (Milan, Italy, 2019).
- Talk “Dal Machine Learning al Deep Learning, come le macchine interpretano il linguaggio” for the Milan Digital Week event (Milan, Italy, 2019).
- Talk “Adapting Named Entity Types to New Ontologies in a Microblogging Environment” at the *31st International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems* (Montreal, Canada, 2018).
- Seminar on “Mentions and Synsets embeddings induced from a heterogeneous graph structure” at INRIA Sophia Antipolis - Méditerranée (Nice, France, 2017).
- **Invited talk** “Deep learning in natural language processing” at the conference of *AI, Machine Learning and Sentiment Analysis Applied to Finance* (London, UK, 2016).

**Teaching experiences** *University of Milano-Bicocca, Milan.*

- Teaching Assistant for the course of “*Data Analytics*” for the master degree in Computer Science. SSD: INF/01. Overall hours: 24. Period: Feb 2020 - Jul 2020; Feb 2019 - Jul 2019.
- Teaching Assistant for the course of “*Data Analytics*” for the master degree in Computer Science. SSD: INF/01. Overall hours: 12. Period: Feb 2018 - Jul 2018; Feb 2017 - Jul 2017.
- Teaching Assistant for the course of “*Probabilistic Models for Decision Making*” for the master degree in Computer Science. SSD: MAT/09. Overall hours: 30. Period: Sep 2017 - Jan 2018.
- Teaching Assistant for the course of “*Laboratory of programming in Java*” for the bachelor degree in Mathematics. SSD: INF/01. Overall hours: 30. Period: Sep 2016 - Jan 2017; Sep 2015 - Jan 2016.
- Co-Supervisor of 18 bachelor’s and master’s degree thesis in Computer Science.

*University of Trieste, Trieste.*

- Adjunct Professor for the course of “*Natural Language Processing*” for the master degree in Data Science and Scientific Computing. SSD: ING-INF/05. Overall hours: 40. Period: May 2020 - Jul 2020; Mar 2021 - May 2021; Mar 2022 - Jul 2022.

**Industrial Research Collaboration**

- Web service and Social Media Manager for DMI Unit at Bocconi University (May 2020 - Apr 2021).
- Collaboration with Opt Out Tools for creating a browser extension that hides misogynistic harassment from the Twitter feed (Mar 2020 - Mar 2021).
- Collaboration with Clinical Insights Ltd for the project “Patient profiling for drug compliance” (Jan 2018 - Feb 2018).
- Collaboration with Consorzio Milano Ricerche for the project regarding the development, exploration and writing of “Proof of Concept: Social Media Sentiment Analytics (Brexit)” (Nov 2016 - Mar 2017).
- Collaboration with Consorzio Milano Ricerche for the project regarding the analysis and presentation of “Methods for the analysis of longitudinal prescription data” (Nov 2014 - Dec 2014).

**SERVICE ACTIVITIES****Editorial Activities**

- *Reviewer* for Computers & Operations Research - Elsevier, Expert Systems with Applications - Elsevier, Information Processing & Management - Elsevier, Journal of Web Semantics - Elsevier, Online Social Networks and Media - Elsevier, and Information - MDPI.
- *Reviewer* for ACL, EAACL, EMNLP, AACL.
- *Program Committee* member of NAACL, \*SEM, ICIW.
- *Area Chair* at the 14th Women in Machine Learning Workshop (WiML) at NeurIPS 2019.
- *Session Chair* at the 12th International Conference on Knowledge Science, Engineering and Management, 2019, and at the 31st International Conference on Industrial, Engineering & Other Applications of Applied Intelligent Systems, 2018.
- Local Committee member of the 2nd International Conference on Data Science and Social Research, 2019.

**Membership to Scientific Societies**

- Member of the Association for Computational Linguistics.

**Service to University**

- Ph.D. representative at the University of Milano-Bicocca, Milan. (2014 - 2018).
- Research Fellow representative at the University of Milano-Bicocca, Milan. (2018 - 2020).

Organization of International Shared Tasks

- AMI: Automatic Misogyny Identification (AMI) at the 7th evaluation campaign of Natural Language Processing and Speech tools for Italian (EVALITA'20). Italy, December 2-3, 2020.
- SemEval-2019 Task 5: HatEval: Multilingual Detection of Hate Speech Against Immigrants and Women in Twitter at the 13th International Workshop on Semantic Evaluation (SemEval-2019). Minnesota, June 6-7, 2019.
- AMI: Automatic Misogyny Identification (AMI) at the 6th evaluation campaign of Natural Language Processing and Speech tools for Italian (EVALITA'18). Italy, December 10-12, 2018.

PERSONAL SKILLS

Mother tongue Italian

Other languages

	UNDERSTANDING		SPEAKING		WRITING
	Listening	Reading	Spoken interaction	Spoken production	
English	C1	C1	C1	C1	C1
French	B1	B1	B1	B1	A2

Levels: A1/A2: Basic user - B1/B2: Independent user - C1/C2: Proficient user  
[Common European Framework of Reference \(CEF\) level](#)

Computer skills

- Excellent use of the most widely used Operating Systems (Linux, Windows, Macintosh).
- Experience with distributed version control systems (git, SVN).
- Good knowledge of the database querying language SQL and the database management system MySQL.
- Good knowledge of programming languages and software: C++, Cytoscape, Knime, Weka.
- Advanced knowledge of JAVA, Matlab, R, Python.
- Proficiency with several tools: Pandas, scikit-learn, Keras, PyTorch, TensorFlow, NLTK.

## PUBLICATIONS

Please note that in natural language processing, it is customary to publish in conference proceedings rather than in journals, and that author order (unless indicated otherwise) reflects involvement.

**2022**

- G. Attanasio, E. Pastor, C. Di Bonaventura, **D. Nozza**. “*ferret: a Framework for Benchmarking Explainers on Transformers*”. arXiv preprint arXiv:2208.01575. 2022.
- G. Attanasio, **D. Nozza**, D. Hovy, E. Baralis. “*Entropy-based Attention Regularization Frees Unintended Bias Mitigation from Lists*”. In Findings of the Association for Computational Linguistics: ACL 2022. Dublin, Ireland, May, 2022.
- **D. Nozza**, F. Bianchi, A. Lauscher, D. Hovy. “*Measuring Harmful Sentence Completion in Language Models for LGBTQIA+ Individuals*”. Proceedings of the Second Workshop on Language Technology for Equality, Diversity and Inclusion. Dublin, Ireland, May, 2022.
- G. Attanasio, **D. Nozza**, E. Pastor, D. Hovy. (2022, May). “*Benchmarking Post-Hoc Interpretability Approaches for Transformer-based Misogyny Detection*”. Proceedings of NLP Power! The First Workshop on Efficient Benchmarking in NLP. Dublin, Ireland, May, 2022.

**2021**

- F. Bianchi, **D. Nozza**, D. Hovy. “*Language Invariant Properties in Natural Language Processing*”. arXiv preprint arXiv:2109.13037. 2021.
- **D. Nozza**, “*Exposing the limits of Zero-shot Cross-lingual Hate Speech Detection*”. Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers). Online, 2021. (acceptance rate 21%).
- **D. Nozza**, F. Bianchi, D. Hovy. “*HONEST: Measuring Hurtful Sentence Completion in Language Models*”. Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). Online, 2021. (acceptance rate 23%).
- F. Bianchi, S. Terragni, D. Hovy, **D. Nozza**, E. Fersini. “*Cross-lingual Zero-shot Topic Models with Contextualized Embeddings*”. Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers. Online, 2021. (acceptance rate 25%).
- F. Bianchi, **D. Nozza**, D. Hovy. “*FEEL-IT: Emotion and Sentiment Classification for the Italian Language*”. Proceedings of the 11th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis. Online, 2021.
- T. Fornaciari, F. Bianchi, **D. Nozza**, D. Hovy. “*MilaNLP at WASSA 2021: Does BERT Feel Sad When You Cry?*”. Proceedings of the 11th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis. Online, 2021.
- **D. Nozza**, P. Manchanda, E. Fersini, M. Palmonari, E. Messina. “*LearningToAdapt with Word Embeddings: Domain Adaptation of Named Entity Recognition Systems*”. Journal of Information Processing and Management, 58(3), 102537. 2021.

**2020**

- S.Terragni, **D. Nozza**, E. Fersini, E. Messina. “*Which Matters Most? Comparing the Impact of Concept and Document Relationships in Topic Models*”. In Proceedings of the First Workshop on Insights from Negative Results in NLP, co-located with the Conference on Empirical Methods in Natural Language Processing. Online, 2020.

- E. Fersini, **D. Nozza**, G. Boifava. “*Profiling Italian Misogynist: An Empirical Study*”. In Proceedings of the Workshop on Resources and Techniques for User and Author Profiling in Abusive Language (ResT-UP 2020). Marseille, France, May 11–16, 2020.
- **D. Nozza**, F. Bianchi, D. Hovy. “*What the [MASK]? Making Sense of Language-Specific BERT Models*”. arXiv preprint arXiv:2003.02912. 2020.

## 2019

- **D. Nozza**, E. Fersini, E. Messina. “*CAGE: Constrained deep Attributed Graph Embedding*”. Information Sciences. 2019.
- **D. Nozza**, C. Volpetti, E. Fersini. “*Unintended Bias in Misogyny Detection*”. In Proceedings of the 2019 IEEE/WIC/ACM International Conference on Web Intelligence (WI 2019). Thessaloniki, Greece, October 14-17, 2019. **Runner up for the Best Paper**. (acceptance rate 18%).
- **D. Nozza**, C. Sas, E. Fersini, E. Messina. “*Word Embeddings for Unsupervised Named Entity Linking*”. In Proceedings of the 12th International Conference on Knowledge Science, Engineering and Management. Athens, Greece, August 28-30, 2019. (acceptance rate 42%).
- V. Basile, C. Bosco, E. Fersini, **D. Nozza**, V. Patti, F. Rangel, P. Rosso, M. Sanguinetti. “*SemEval-2019 Task 5: Multilingual Detection of Hate Speech Against Immigrants and Women in Twitter*”. In Proceedings of the 13th International Workshop on Semantic Evaluation (SemEval-2019). Minneapolis, Minnesota, June 6-7, 2019.

## 2018

- E. Fersini, **D. Nozza**, P. Rosso. “*Overview of the Evalita 2018 Task on Automatic Misogyny Identification (AMI)*”. In Proceedings of the 6th evaluation campaign of Natural Language Processing and Speech tools for Italian (EVALITA'18). Turin, Italy, December 10-12, 2018.
- V. Basile, E. Cabrio, F. Gandon and **D. Nozza**. “*Mapping Natural Language Labels to Structured Web Resources*.” In Proceedings of the NL4AI Workshop, co-located with the International Conference of the Italian Association for Artificial Intelligence (AI\* IA). Trento, Italy, November 22-23, 2018.
- F. Bianchi, M. Palmonari, **D. Nozza**. “*Towards Encoding Time in Text-Based Entity Embeddings*”. In Proceedings of the 17th International Semantic Web Conference. Monterey, California, USA, October 8-12, 2018. (acceptance rate 25%).
- E. Fersini, P. Manchanda, E. Messina, **D. Nozza**, M. Palmonari. “*Adapting Named Entity Types to New Ontologies in a Microblogging Environment*”. In Proceedings of the 31st International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems. Montreal, Canada, June 25-28, 2018. (acceptance rate 59%).

## 2017

- **D. Nozza**, F. Ristagno, M. Palmonari, E. Fersini, P. Manchanda, E. Messina. “*TWINE: A real-time system for TWEet analysis via INformation Extraction*”. In Proceedings of the Demonstrations at the 15th Conference of the European Chapter of the Association for Computational Linguistics. Valencia, Spain, April 3-6, 2017.
- **D. Nozza**, E. Fersini, E. Messina. “*A Multi-View Sentiment Corpus*”. In Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics. Valencia, Spain, April 3-6, 2017. (acceptance rate 27%).

- P. Manchanda, E. Fersini, **D. Nozza**, E. Messina, M. Palmonari. “Towards Adaptation of Named Entity Classification”. In Proceedings of the 32nd ACM Symposium on Applied Computing. Marrakesh, Morocco, April 3-6, 2017. (acceptance rate 23%).

## 2016

- F. M. Cecchini, E. Fersini, P. Manchanda, E. Messina, **D. Nozza**, M. Palmonari, C. Sas. “UNIMIB@NEEL-IT: Named Entity Recognition and Linking of Italian Tweets”. In Proceedings of the 5th International Workshop EVALITA. Naples, December 7th, 2016.
- **D. Nozza**, E. Fersini, E. Messina. “Unsupervised Irony Detection: A Probabilistic Model with Word Embeddings”. In Proceedings of the 8th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management. Porto, November 10th, 2016. **Best Paper Award**. (acceptance rate 18%).
- **D. Nozza**, E. Fersini, E. Messina. “Deep Learning and Ensemble Methods for Domain Adaptation”. In Proceedings of the 28th IEEE International Conference on Tools with Artificial Intelligence. San Jose, CA, November 6th, 2016. (acceptance rate 31%).

## 2014

- **D. Nozza**, D. Maccagnola, V. Guigue, E. Messina and P. Gallinari. “A latent representation model for sentiment analysis in heterogeneous social networks”. In Proceedings of the 3rd International on Modelling and Knowledge Management applications. Grenoble, September 2nd, 2014.

*I hereby authorize the use of my personal details solely for circulation within the company according to the Italian Legislative Decree n. 196/2003.*