

FEDERICO BELOTTI

belo.fede@outlook.com | (+39) 3496508679 | [GitHub](#) | [LinkedIn](#) | [Google Scholar](#)

PERSONAL INFO

Birth date: July 20th 1993

Spoken languages:

- Italian: C2 (Mother-tongue)
- English: C1 (101/120, taken December 1, 2024)

EDUCATION

Università degli studi di Milano-Bicocca

October 2018 - October 2020

Master in Computer Science, **GPA: 4.0/4.0**

Thesis: "Compass aligned Gaussian Embeddings for unsupervised lexical semantic change detection"

Università degli studi di Milano-Bicocca

September 2015 - July 2018

Bachelor in Computer Science, **GPA: 3.9/4.0**

Thesis: "Recognition of plant varieties in digital images"

I.T.I.S. P. Paleocapa

September 2012 - June 2015

High-school Diploma in Computer Science, **GPA: 3.9/4.0**

WORK EXPERIENCE

University of Milano-Bicocca

January 2024 - Today

Research Fellow

Leveraging Transformer-based models (Encoder-only like BERT and similar, or Decoder-only like GPT and similar) for the Semantic Table Interpretation problem, while applying Meta, Active and/or Reinforcement Learning techniques to adapt them to unseen data and/or human feedback. Addressing interpretability on tabular data through mechanistic interpretability.

Orobix S.r.l

January 2022 - December 2023

Project Leader

Project Leader of a Reinforcement Learning (RL) project, supervising a three-people team with respect to both planning and developing activities. I implemented standard RL algorithms, such as A2C, A3C, PPO, SAC, Dreamer-V1/V2/V3 and curiosity-based algorithms. The project converged to an open-source distributed Reinforcement Learning framework called [sheeprl](#).

Orobix S.r.l

November 2020 - December 2021

Data Scientist

Designed and implemented deep learning models, both Convolutional and Transformer-based, to tackle real world computer vision problems, acquiring knowledge in: unsupervised domain adaptation, object segmentation, supervised and self-supervised learning.

University of Milano-Bicocca

March 2019 - February 2020

Research Fellow

Designed and implemented semantic matching techniques and showed their effectiveness on data relating to descriptions of products managed on eCommerce sites. The semantic matching techniques developed was based on machine learning methods based on deep neural networks. The activity was located inside of the **EW-Shopp** technological innovation project, funded in the context of the H2020 program, which included ten other European research organizations and companies.

University of Milano-Bicocca - Imaging and Vision Laboratory
Internship

February 2018 - July 2018

Developed a state of the art method to recognize fruits and vegetables using Convolutional Neural Networks, that resulted in my bachelor thesis: "Recognition of edible vegetables and fruits for smart-home appliances".

RESEARCH PROJECTS

Machine learning algorithms for annotating and reconciling data through user-interaction mechanisms

January 2024 - Present

University of Milano-Bicocca - Cod. 24A011 and 23A144

Semantic matching of products using algorithms based on deep learning techniques

October 2019 - February 2020

University of Milano-Bicocca - Cod. 19B135

Semantic enrichment techniques to support predictive analytics on big data

March 2019 - June 2019

University of Milano-Bicocca - Cod. 19B024

TECHNICAL STRENGTHS

Programming & Skills
Software & Tools

Python (7yrs), Java (4yrs), C++ (1yr)
PyTorch (6yrs+), PyTorch Distributed (2yr), scikit-learn (2yrs+),
PyTorch-Lightning (3yr), Open-CV (1yr), Hydra (2yr), Docker (2yr),
AzureML (1yr), MongoDB (1yr)

SUMMER SCHOOLS

Mediterranean Machine Learning Summer School

January 2021

The Mediterranean Machine Learning (M2L) summer school was structured around 6 days of keynotes, lectures and practical sessions. The principal topics were: Computer Vision, Natural Language Processing, Generative models, Graph Neural Network, Robotics, Reinforcement Learning and many more. Due to the COVID-19 pandemic, the summer school, which had to be held in august 2020, was postponed to January 2021 in an online format.

SELECTED PUBLICATIONS

Efficient Training of Sparse Autoencoders for Large Language Models via Layer Groups.

Davide Ghilardi*, Federico Belotti*, and Marco Molinari*.

arXiv preprint arXiv:2410.21508 (2024).

Accelerating Sparse Autoencoder Training via Layer-Wise Transfer Learning in Large Language Models.

Davide Ghilardi*, Federico Belotti*, Marco Molinari* and Jaehyuk Lim.

The 7th BlackboxNLP Workshop. 2024.

Evaluating LLMs on Entity Disambiguation in Tables.

Federico Belotti, et al.

arXiv preprint arXiv:2408.06423. (2024).

UNIMIB @ DIACR-Ita: Aligning Distributional Embeddings with a Compass for Semantic Change Detection in the Italian Language.

Federico Belotti, Federico Bianchi, and Matteo Palmonari.

2020, pp. 451-455. ISBN 9791280136329. DOI: 10.4000/books.aaccademia.7688.

SELECTED OPEN-SOURCE CONTRIBUTIONS

alligator

Main contributor of [alligator](#), a Python library to perform entity linking over tabular data.

Lightning-AI

Contributed to the [lightning](#) and [litgpt](#) repositories with multiple PRs, with a complete Reinforcement Learning [example](#).

sheeprl

Main contributor of [sheeprl](#), an open-source PyTorch-based distributed Reinforcement Learning framework.

fwdgrad

A simple functorch implementation of the “[Gradients without Backpropagation](#)” paper.

mdd-domain-adaptation

A simple pytorch-lightning implementation of the “[Maximum Density Divergence for Domain Adaptation](#)” paper.

word2vecpp

A C++ word2vec implementation.

20/03/2024

