

Elio Musacchio

 m-elio |  Elio Musacchio |  elio.musacchio@phd.unipi.it




SHORT BIOGRAPHY

I was born in Bari (BA) on the 19th of February of the year 2000. I started attending the **University of Bari** in 2018 and hold both a **Bachelor's and a Master's degree in Computer Science**.

I am currently in the third year of the **Italian National Ph.D. in AI**, a program organized by the University of Pisa. Although the Ph.D. program is managed by Pisa, I am based at the University of Bari and work within the **SWAP Research Group**.

While I initially focused on Multimodal Recommender Systems, my current research is centered on **multimodality and multilinguality in Large Vision-Language Models (LVLMs)**. Specifically, I design training strategies and evaluation benchmarks to improve and assess the multilingual and multimodal capabilities of LVLMs.

EDUCATION

-  2023 - present **National Ph.D. in AI, University of Pisa & University of Bari**
Thesis: *Understanding language-based human signals: taking into account multimodality and adaptivity through Large Language Understanding Models*
-  2021 - 2023 **M.S. in Computer Science, University of Bari**
Thesis: *Evaluation of Personalized Recommendation using Multimodal Information*
-  2018 - 2021 **B.S. in Computer Science, University of Bari**
Thesis: *A Python Framework for Content-Based Recommendation: Refactoring and System Optimization*

HIGHLIGHTS

Multimodality and Multilinguality

My current work focuses on **Large Vision-Language Models and their multilingual abilities**. I focus both on designing training strategies for better multilingual understanding and/or generation, and on designing complex tasks to evaluate their capabilities. Relevant examples of my research in this field include **xVLM2Vec**, where a training strategy to improve the multilingual understanding of LVLM-based encoder models is proposed, and **VWSD**, where a multilingual and multimodal benchmark with hard negatives selected using the co-hyponymy relation of semantic networks is proposed.

HPC

I have worked on several projects using **Leonardo**, an Italian supercomputer where each compute node has four NVIDIA A100 64 GB VRAM GPUs. This experience has allowed me to acquire proficiency in handling huge amounts of data and performing large-scale experiments.

RELEVANT EXPERIENCES

 Three months visiting Ph.D. student at the University of Würzburg, supervised by Goran Glavaš

SKILLS AND INFO

 Proficient in Python and Java programming languages

- ✍ Proficient in Deep Learning Python libraries and frameworks
- ✍ Experience in HPC environments
- ✍ Cambridge English C1 Certification

AWARDS

- 👑 Outstanding Reviewer Award at ACM UMAP 2024

PUBLICATIONS

- Pierpaolo Basile, Elio Musacchio, Marco Polignano, Lucia Siciliani, Giuseppe Fiameni, and Giovanni Semeraro. “LLaMAntino: LLaMA 2 Models for Effective Text Generation in Italian Language”. In: *CoRR* abs/2312.09993 (2023). DOI: [10.48550/ARXIV.2312.09993](https://doi.org/10.48550/ARXIV.2312.09993). arXiv: [2312.09993](https://arxiv.org/abs/2312.09993). URL: <https://doi.org/10.48550/arXiv.2312.09993>.
- Pasquale Lops, Elio Musacchio, Cataldo Musto, Marco Polignano, Antonio Silletti, and Giovanni Semeraro. “Reproducibility Analysis of Recommender Systems relying on Visual Features: traps, pitfalls, and countermeasures”. In: *Proceedings of the 17th ACM Conference on Recommender Systems, RecSys 2023, Singapore, Singapore, September 18-22, 2023*. Ed. by Jie Zhang, Li Chen, Shlomo Berkovsky, Min Zhang, Tommaso Di Noia, Justin Basilico, Luiz Pizzato, and Yang Song. ACM, 2023, pp. 554–564. DOI: [10.1145/3604915.3609492](https://doi.org/10.1145/3604915.3609492). URL: <https://doi.org/10.1145/3604915.3609492>.
- Giuseppe Attanasio, Pierpaolo Basile, Federico Borazio, Danilo Croce, Maria Francis, Jacopo Gili, Elio Musacchio, Malvina Nissim, Viviana Patti, Matteo Rinaldi, and Daniel Scalena. “CALAMITA: Challenge the Abilities of LAnGuage Models in ITAlIAn”. In: *Proceedings of the Tenth Italian Conference on Computational Linguistics (CLiC-it 2024), Pisa, Italy, December 4-6, 2024*. Ed. by Felice Dell’Orletta, Alessandro Lenci, Simonetta Montemagni, and Rachele Sprugnoli. Vol. 3878. CEUR Workshop Proceedings. CEUR-WS.org, 2024. URL: https://ceur-ws.org/Vol-3878/116%5C_calamita%5C_preface%5C_long.pdf.
- Pierpaolo Basile, Marco de Gemmis, Elio Musacchio, Marco Polignano, Giovanni Semeraro, Lucia Siciliani, Vincenzo Tamburrano, Vita Santa Barletta, Danilo Caivano, Fabiana Battista, Antonietta Curci, Rosa Scardigno, Gabriella Calvano, and Patrizia Soriano. “Explaining Intimate Partner Violence with LLaMAntino”. In: *Proceedings of the Ital-IA Intelligenza Artificiale - Thematic Workshops co-located with the 4th CINI National Lab AIIS Conference on Artificial Intelligence (Ital-IA 2024), Naples, Italy, May 29-30, 2024*. Ed. by Sergio Di Martino, Carlo Sansone, Elio Masciari, Silvia Rossi, and Michela Gravina. Vol. 3762. CEUR Workshop Proceedings. CEUR-WS.org, 2024, pp. 29–34. URL: <https://ceur-ws.org/Vol-3762/510.pdf>.
- Pierpaolo Basile, Elio Musacchio, and Lucia Siciliani. “ITA-SENSE - Evaluate LLMs’ ability for ITALIAN word SENSE disambiguation: A CALAMITA Challenge”. In: *Proceedings of the Tenth Italian Conference on Computational Linguistics (CLiC-it 2024), Pisa, Italy, December 4-6, 2024*. Ed. by Felice Dell’Orletta, Alessandro Lenci, Simonetta Montemagni, and Rachele Sprugnoli. Vol. 3878. CEUR Workshop Proceedings. CEUR-WS.org, 2024. URL: https://ceur-ws.org/Vol-3878/134%5C_calamita%5C_short.pdf.
- Elio Musacchio, Lucia Siciliani, Pierpaolo Basile, Edoardo Michielon, Marco Pasqualini, Asia Beatrice Uboldi, and Giovanni Semeraro. “A Study on the Soundness of Closed-ended Evaluation of Large Language Models Adapted to the Italian Language”. In: *Proceedings of the Tenth Italian Conference on Computational Linguistics (CLiC-it 2024), Pisa, Italy, December 4-6, 2024*. Ed. by Felice Dell’Orletta, Alessandro Lenci, Simonetta Montemagni, and Rachele Sprugnoli. Vol. 3878. CEUR Workshop Proceedings. CEUR-WS.org, 2024. URL: https://ceur-ws.org/Vol-3878/67%5C_main%5C_long.pdf.

Elio Musacchio, Lucia Siciliani, Pierpaolo Basile, and Giovanni Semeraro. “Adapting Large Language Models to Narrative Content”. In: *Proceedings of the 3rd Workshop on Artificial Intelligence and Creativity co-located with 27th European Conference on Artificial Intelligence (ECAI 2024), Santiago de Compostela, Spain, October 20, 2024*. Ed. by Allegra De Filippo, François Pachet, Valentina Presutti, and Luc Steels. Vol. 3810. CEUR Workshop Proceedings. CEUR-WS.org, 2024, pp. 17–29. URL: <https://ceur-ws.org/Vol-3810/paper2.pdf>.

– “LLaVA-NDiNO: Empowering LLMs with Multimodality for the Italian Language”. In: *Proceedings of the Eight Workshop on Natural Language for Artificial Intelligence (NL4AI 2024) co-located with 23th International Conference of the Italian Association for Artificial Intelligence (AI*IA 2024), Bolzano, Italy, November 26th-27th, 2024*. Ed. by Giovanni Bonetta, Claudiu Daniel Hromei, Lucia Siciliani, and Marco Antonio Stranisci. Vol. 3877. CEUR Workshop Proceedings. CEUR-WS.org, 2024. URL: <https://ceur-ws.org/Vol-3877/paper9.pdf>.

Pierpaolo Basile, Lucia Siciliani, Elio Musacchio, and Giovanni Semeraro. “Exploring the Word Sense Disambiguation Capabilities of Large Language Models”. In: *CoRR* abs/2503.08662 (2025). DOI: [10.48550/ARXIV.2503.08662](https://doi.org/10.48550/ARXIV.2503.08662). arXiv: [2503.08662](https://arxiv.org/abs/2503.08662). URL: <https://doi.org/10.48550/arXiv.2503.08662>.

Gaetano Dibenedetto, Elio Musacchio, Marco Polignano, and Pasquale Lops. “Fine-Tuning Large Multimodal Models for Fitness Action Quality Assessment”. In: *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization, UMAP Adjunct 2025, New York City, NY, USA, June 16-19, 2025*. ACM, 2025, pp. 39–44. DOI: [10.1145/3708319.3733684](https://doi.org/10.1145/3708319.3733684). URL: <https://doi.org/10.1145/3708319.3733684>.

Elio Musacchio, Lucia Siciliani, Pierpaolo Basile, and Giovanni Semeraro. “Assessing and Improving the Multilingual Visual Word Sense Disambiguation Ability of Vision-Language Models”. In: *ECAI 2025*. IOS Press, 2025, pp. 4145–4152.

– “Extending Large Language Models to multimodality for non-English languages”. In: *Computer Vision and Image Understanding* (2025), p. 104618.

– “xVLM2Vec: Adapting LVLM-based embedding models to multilinguality using Self-Knowledge Distillation”. In: *CoRR* abs/2503.09313 (2025). DOI: [10.48550/ARXIV.2503.09313](https://doi.org/10.48550/ARXIV.2503.09313). arXiv: [2503.09313](https://arxiv.org/abs/2503.09313). URL: <https://doi.org/10.48550/arXiv.2503.09313>.

Giuseppe Spillo, Elio Musacchio, Cataldo Musto, Marco de Gemmis, Pasquale Lops, and Giovanni Semeraro. “See the Movie, Hear the Song, Read the Book: Extending MovieLens-1M, Last.fm-2K, and DBbook with Multimodal Data”. In: *Proceedings of the Nineteenth ACM Conference on Recommender Systems, RecSys 2025, Prague, Czech Republic, September 22-26, 2025*. Ed. by Mária Bieliková, Pavel Kordík, Markus Schedl, Marco de Gemmis, Sole Pera, Rodrigo Alves, Olivier Jeunen, and Vito Ostuni. ACM, 2025, pp. 847–856. DOI: [10.1145/3705328.3748162](https://doi.org/10.1145/3705328.3748162). URL: <https://doi.org/10.1145/3705328.3748162>.