Dario Di Palma

LinkedIn | Online Portfolio | Google Scholar | dario.dipalma@outlook.it

ABOUT: **PhD** in Artificial Intelligence specializing in **Information Retrieval**, **Recommender Systems**, and **Large Language Models**. Published at **ACL**, **SIGIR**, **RecSys**, **UMAP**, and **CIKM**. Experienced in LLM evaluation, Text-to-SQL, sentiment/emotion probing, data generation pipelines, and distributed training on HPC systems. **Collaborated with IBM Research US**, supervised 20+ MSc students, recipient of the **Best Short Paper Award** at SIGIR 2025 and candidate for the **Best LBR Award** at UMAP 2025. **Open to relocation**. Availability: starting from April 2026.

SKILLS: Python (Pandas, NumPy, scikit-learn, PyTorch), SQL, and Hugging Face Transformers. Proficient with Docker, Git, GitHub/GitLab, modern visualization libraries (matplotlib, seaborn, Plotly), and distributed training on HPC systems using SLURM.

EDUCATION

PhD. in Computer Engineering — Politecnico di Bari, Italy — **Jan 2023 – Mar 2026 (expected)**
Topics: LLMs, Recommender Systems, Text-to-SQL, Information Retrieval

M.S in Computer Engineering, Full marks and honors (110/110 cum laude) — Politecnico di Bari — **Sep 2020 – Oct 2022**
Thesis: Point-of-Interest Recommendation via GNNs on Heterogeneous Graphs

WORK EXPERIENCE

Teaching Assistant — Politecnico di Bari — **Sep 2022 – Dec 2025**
- Designed teaching materials adopted by 300+ MSc students in Deep Learning, Databases, and IR courses.
- Mentored 30+ students, leading to research publications in SIGIR, CIKM or opensource projects.
- Developed hands-on labs using PyTorch, TensorFlow, and SQL systems.

Research Scientist Internship (Recommender Systems) — Universidad Autónoma de Madrid — **Apr 2022 – Oct 2022**
- Developed a GNN-based Point-of-Interest Recommender System using Keras, NetworkX, and heterogeneous graphs, improving contextual personalization for tourism recommendations.

Research Scientist (Semantic Knowledge Graph) — Politecnico di Bari — **Dec 2020 – Oct 2021**
- Developed Knowledge Graph–based reasoning modules for a StarCraft II agent and Hyperledger blockchain prototypes, contributing to research on symbolic reasoning and decentralized systems.

SELECTED PUBLICATIONS

ACL 2025 | **Are the Hidden States Hiding Something? Testing the Limits of Factuality-Encoding Capabilities in LLMs**
Factuality in LLMs
- Created realistic datasets combining tables and QA sources to test LLM factuality.
- Demonstrated limits in truthfulness encoding across open-source LLMs.
- Showed failure cases on fabricated content, exposing benchmark weaknesses.

ACL 2025 | **LLaMAs Have Feelings Too: Unveiling Sentiment and Emotion Representations in LLaMA Models Through Probing**
Sentiment Probing in LLMs
- Designed an experimental **layer-editing framework** for LLaMA models to inject downstream classifiers for sentiment tasks.
- Showed that mid-layer activations **outperform prompting by** up to **14%**. **Reduced inference memory** usage by **57%** through a lightweight probing method. **Demonstrated** that **last-token** embeddings are **often suboptimal** for sentiment tasks.

SIGIR 2025 | **Do LLMs Memorize Recommendation Datasets? A Preliminary Study on MovieLens-1M**
🏆 Best Short Paper Award 🏆
First study on LLM memorization in RecSys datasets
- Conducted the first systematic analysis of **data memorization in LLMs** on MovieLens datasets.
- Built a **prompting-based framework** revealing MovieLens-1M memorization in LLMs, showing across GPT and LLaMA models that **memorization impacts** recommendation performance, fairness, and generalization.

UMAP 2025 | **Content-Based or Collaborative? Insights from Inter-List Similarity Analysis of ChatGPT Recommendations**
🎖 Candidate Best LBR 🎖
Analyzing LLM recommendation behavior
- Analyzed whether ChatGPT's recommendations align with Collaborative Filtering or Content-Based Filtering by comparing inter-list similarity recommendation lists.
- Found a stronger alignment with CF, while showing that ChatGPT also recommends more diverse and domain-aware items than CF/CBF baselines, suggesting its potential as a hybrid RecSys component.

RecSys 2023 | **Retrieval-augmented Recommender System: Enhancing Recommender Systems with Large Language Models**
Integrating LLMs to enhance recommendation
- Pioneered the integration of **Large Language Models with Recommender Systems**, introducing retrieval-augmented concepts for hybrid user–item matching.

For a comprehensive list of publications, please refer to my Google Scholar profile.

SIDE PROJECTS

LiteratureRAG
Elasticsearch, Kibana, Milvus, GROBID, Docker Compose, LangChain, HPC, SLURM
- Developed a Retrieval-Augmented Generation (RAG)-based assistant enabling literature retrieval, metadata extraction, and domain-level summarization for scientific research.
- Enabled efficient exploration of large research corpora by combining neural retrieval with automated metadata extraction and domain-aware summarization.

RITA Copilot
Kuksa, Azure, Git, Google Maps API, Python, Java, HuggingFace, Transformers
- Built an AI-powered driver assistant integrating voice recognition, NLP, and text-to-speech to support real-time interaction in vehicle environments.
- Enabled hands-free, context-aware assistance for routing, traffic monitoring, and vehicle support, improving driver safety and usability.