Aleksandar Fontana

Gildarts777 | Linkedin | aleksandar.fontana@gmail.com | Website | +39 339 717 3403

Education

PhD in **AI and Cybersecurity** — *Scuola Superiore Sant'Anna* — 2024 – present

Master in Physics of Complex Systems and Big Data — *University of Tor Vergata* — 2021 – 2024

Thesis: "Exploring the frontier of chatbot technology"

Bachelor in Physics — *University of Tor Vergata* — 2018 – 2021

Thesis: "Quantum computing and cosmology of the early universe"

Work Experience

AI Consultant (Netgroup S.p.A., Naples) — 2025 - present

Engineered a Transformer-based Foundation Model for cyber attack prediction from scratch.

AI/NLP Research Intern (Serco Italia S.p.A, Rome) — 2023 - 2024

Engineered a Small LLM implementing BitNet 1.58-bit quantization for efficient deployment.

Head of Board & Community Manager (Element Gaming) — 2016 - present

Led community operations, coordinating a network of 100+ affiliated content creators.

Skills

LLM Alignment & Reasoning	Reinforcement Learning (PPO, GRPO, GTPO), Natural Language Processing, Unsloth, TRL, Hugging Face
AI for Cybersecurity	Vulnerability Detection, Threat Intelligence, Custom Transformers, Tabular Data Modeling
Core ML & Engineering	Python, PyTorch, Deep Learning, CUDA, Git
Leadership & Soft Skills	Team Leadership, Community Management, Event Coordination, Effective Communication
Languages	**Italian**: Native **English**: Professional Proficiency

Publications

On the Hidden Objective Biases of Group-based Reinforcement Learning

Aleksandar Fontana*, Marco Simoni*, et al. *In ACL 2026.* 2026

Developed a **unified surrogate formulation** to analyze **GRPO**, revealing theoretical structural mismatches and key limitations like prefix token gradient biases.

GTPO: Stabilizing Group Relative Policy Optimization via Gradient and Entropy Control

Aleksandar Fontana*, Marco Simoni*, et al. *Submitted to TACL.* Code Repository 2025

Designed a KL-free policy optimization algorithm mitigating LLM policy collapse. Boosted reasoning performance by up to **15%** on OOD benchmarks (AIME2024, AIME2025, AMC) vs. GRPO.

Improving LLM Reasoning for Vulnerability Detection via Group Relative Policy Optimization

Aleksandar Fontana*, Marco Simoni*, et al. Code Repository 2025

Applied **GRPO** to enhance LLM reasoning for vulnerability detection, achieving up to **+29 macro** F_1 and **+17 accuracy points** over SFT.

TITAN: Graph-Executable Reasoning for Cyber Threat Intelligence 2025

Marco Simoni, Aleksandar Fontana, et al. Code Repository

Unmasking model behavior: How LLMs reason on vulnerability detection

Aleksandar Fontana, Marco Simoni. In: *ARES 2025* 2025

* equal contribution