Mohamed Rissal Hedna

Hamburg, Germany | +49 176 43402970 | rissalhedna@gmail.com | LinkedIn | Github

SUMMARY

Artificial Intelligence Researcher (M.Sc.) | Generative AI, LLM Safety & Scientific Computing. Junior scientist combining strong research foundations in LLM safety and uncertainty quantification with proven ability to prototype and deploy large-scale AI architectures. Deep expertise in constructing multi-agent RAG workflows and investigating the algorithmic foundations of optimization and high-throughput inference. Technical leader who successfully bootstrapped a tech community of 40,000+ followers. Currently conducting research on hallucination detection and uncertainty estimation in customer-facing products, translating academic rigor into scientific advancement and exploratory research projects.

Technical Skills:

Programming Languages: Python, Java, C, JavaScript/TypeScript.

Model Evaluation & ML: PyTorch, Hugging Face, custom metrics & benchmarking, uncertainty estimation

Distributed Systems & MLOps: Docker, CI/CD, Azure, GCP, high-throughput pipelines

Data & Analysis: NumPy, Pandas, SQL, dashboards & reporting

PUBLICATIONS & PAPERPRINTS

- **M. Rissal Hedna**, Chris Biemann. **"SERL: SElf-Reflecting Language Models"** *Master's Thesis / Manuscript in Preparation* (Targeting ACL/EMNLP 2025).
- **M. Rissal Hedna**, Younes Djemmal, Khaleel Mershad. **"A Model for the Automatic Mixing of Multiple Video and Audio Clips."** *Proceedings of the 2023 International Conference on Cyberworlds (CW)*, Sousse, Tunisia, 2023.
- **Sesugh Nder, M. Rissal Hedna.** *"Simulated Adversarial Patch Attacks on Vision-Based Logistics Systems."* Technical Report, University of Hamburg, 2024. **arXiv:2511.19254 [cs.CV]**.

RESEARCH EXPERIENCE

SERL: SElf-Reflecting Language Models

Hamburg, Germany ***Graduate Researcher (Master's Thesis)*** | *2025 – Present*

- **Unsupervised Self-Calibration in Language Models:** Creating a new paradigm to internally calibrate language models' confidence levels using synthetic data and consistency pseudo-labels to better generalize on unseen datasets.
- Taking advantage of Large Language Models' generation capabilities and internal confidence estimations to force them to self-calibrate, **eliminating the need for external data**.

Simulated Adversarial Patch Attacks

Hamburg, Germany ***Research Assistant (Co-Author)*** | *2024*

- **Differentiable Rendering Pipeline:** Built an end-to-end differentiable rendering pipeline to generate physical adversarial patches, testing the robustness of vision-based logistics systems against security threats.
- **Synthetic Data Generation:** Orchestrated simulation pipelines using Blender and physics engines to create large-scale synthetic datasets for robust model training.

A Model for the Automatic Mixing of Multiple Video and Audio Clips

Beirut, Lebanon ***Lead Researcher*** | *2023*

- **Developed an efficient feature fusion framework** for video/audio processing, successfully reducing computational complexity while maintaining high accuracy for automated video editing tasks.
- **Created a novel blueprint for automatic editing** of social media videos based on Fast Fourier Transform analysis of audio signals.

Advanced Telerobotics Lab

Kent, OH, USA ***Visiting Research Assistant (Computer Vision)*** | *2022*

- **Real-Time GPU Inference:** Engineered a custom C/Python CNN architecture achieving <50ms inference latency on embedded NVIDIA Jetson hardware, optimizing for real-time obstacle detection in GPS-denied environments.

- **Hardware-Constrained Fusion:** Implemented multimodal sensor fusion (LiDAR, Vision) adhering to strict memory and power constraints on edge devices.

PROFESSIONAL EXPERIENCE

PHAROS Labs
AI Engineer | Jan2024 - current

- Engineered a high-throughput data pipeline, performing **memory profiling** and bottleneck analysis to reduce end-to-end latency by 20%.
- Delivered production-grade full-stack features optimized for high concurrency, directly driving a **25% increase in user adoption**.
- Designed **automated evaluation workflows** for multi-agent reasoning systems, including reliability checks and failure case analysis.

Fiindo Hamburg
AI Engineer - contractual until Jan2026 | Oct2025

- **Architected an end-to-end generative content engine** that autonomously ingests raw financial datasets and transforms them into high-retention short-form video content.
- **Led the engineering of a multimodal system** (Text-to-Video/Audio), automating the full production lifecycle from data ingestion to final render to streamline content creation workflows.
- **Designed the retrieval architecture** to handle complex financial queries, ensuring data consistency across the generative pipeline.

AdaLab UG Hamburg
Machine Learning Developer | Jan2023 - Jan2024

- **Scaled a multimodal AI platform** (Text/Image/Video) by implementing parallel execution strategies and memory profiling, successfully resolving critical high-load bottlenecks.
- **Deployed production-ready features** to an active user base, managing the full software development lifecycle from feature implementation to critical hotfixes and system stability.

PROJECTS

Lightweight Git Commit Message Generation (T5-Efficient Optimization)

- **Model Compression:** Engineered a "T5-Efficient" Transformer by optimizing layer distribution, reducing parameter count by **81% (17.7M to 3.3M)** to enable deployment on compute-limited devices.
- **Training Optimization:** Orchestrated end-to-end training in PyTorch using **BF16 mixed-precision**, maximizing GPU throughput to converge in under 12 hours on single-GPU infrastructure.

UniLLM: RAG-based Assistant for International Students

- **End-to-End RAG Pipeline:** Architected a retrieval-augmented generation system using **LlamaIndex** and **Qdrant**, optimizing retrieval chunks to achieve a **13% reduction** in query resolution time.
- **Real-Time Streaming Architecture:** Built the backend to support asynchronous message streaming, ensuring a low-latency user experience for chat interfaces.
- **Automated Data Ingestion:** Engineered a custom web crawler to autonomously fetch, parse, and index living/study information from distributed German data sources, ensuring the vector store remains up-to-date.

EDUCATION

M.Sc. Intelligent Adaptive Systems | University of Hamburg

- **Thesis:** "SERL: SElf-Reflecting Language Models"
- **CGPA:** 1.69/4.0 (German Scale, 1.0=Best)
- **Focus:** Neural Networks, Knowledge Processing, Statistical Methods in Language Technology.

B.Sc. Computer Science | Lebanese American University

- **CGPA:** 3.91 / 4.0 (High Distinction, 4.0=Best)
- **Awards:** MEPI TL Full Scholarship Recipient (highly competitive merit scholarship).

LANGUAGES SPOKEN

German: B2 | **English:** C2 (Fluent) | **French:** C1 (Fluent) | **Arabic:** Native