

# Lucas Pacheco | AI Research Scientist | PhD, University of Bern | LLM Post-Training, Evals & Agentic AI

Belém, Pará – Brazil

📞 +55 (91) 98639-8908 • ✉ lucassidpacheco@gmail.com • 🌐 lsidd  
in lucas-pacheco-l13 • Lattes CV: <http://lattes.cnpq.br/8745332530634533>

## Professional Summary

---

AI Research Scientist (PhD, University of Bern) with **223+** citations, **h-index 9**, and the **IEEE PerCom 2024 Best Paper Award**. Since November 2025 at **CPQD** as Pesquisador Especialista I, contributing to **Project INSPIRE**, Brazil's R\$390M national AI-for-government program, with work spanning the national AI governance methodology, technical training curricula for federal practitioners across maturity levels, and consulting on the Unified Federal AI Platform. Research background in federated learning, edge AI, and distributed ML; applied work covers LLM engineering, RAG, and multi-agent systems deployed in government and industry.

**LLM Engineering:** full-stack expertise across **post-training** (SFT, DPO, RLHF/RLAIF, reward modeling, LoRA/QLoRA), **LLM evaluation** (agentic and multi-turn harnesses, LLM-as-judge, red-teaming, regression suites), and **inference & serving** (vLLM, quantization, latency/throughput optimization). PhD research in federated learning transfers directly to **distributed LLM training** (FSDP, DeepSpeed, mixed precision, scaling behavior). **Agentic AI** systems with **LangGraph** and **CrewAI** production-deployed in federal government services (INSPIRE) and for industry clients including Vale and Nitro Química.

## Core Competencies

---

**LLM Post-Training:** SFT, DPO, RLHF, RLAIF, **LLM Evaluation:** Eval Pipelines, Agentic & Reward Modeling, LoRA/QLoRA Multi-Turn Harnesses, Tool-Use Simulation, LLM-as-Judge, Reward-Hacking Detection, Regression Suites, Custom Benchmarks

**Distributed Training:** FSDP, DeepSpeed, **Inference & Serving:** vLLM, Quantization, Mixed Precision, Scaling Behavior, AdamW Variants, LR Scheduling, Federated Learning Batching, Latency/Throughput Optimization, Self-Hosted High-Volume Inference, Edge Inference

**Agentic AI:** LangGraph, CrewAI, LangChain, **AI Security:** LLM Red-Teaming, Prompt Injection, RAG, Tool/Function Calling, Guardrails, Multi-Agent Orchestration, Jailbreak Defense, Adversarial Robustness, Fortinet Certified

**MLOps & Cloud:** Kubernetes, Docker, Continuous Deployment, Google Vertex AI **AI Policy & Governance:** Federal AI Methodology, Responsible AI, Sovereign AI Platforms, Capacity Building

**Frameworks:** PyTorch (primary), JAX, Tensor- **Programming:** Python, C++, Rust, JavaScript Flow

**Languages:** Portuguese (Native), English (Fluent), German (Intermediate)

## Education

---

**University of Bern & Federal University of Pará** **Switzerland/Brazil**  
*PhD in Computer Science, Co-tutelle Program* *2020–2025*

*Thesis:* “Federated Learning in Mobile Edge Computing”

*Supervisor:* Prof. Torsten Braun (University of Bern)

**Achievement:** Best Paper Award at IEEE PerCom 2024

**Federal University of Pará** **Brazil**  
*Master of Science in Electrical Engineering* *2019–2020*

*Thesis:* “Mobility and Cloud Management in Wireless Heterogeneous 5G Networks”

*Focus:* Edge Computing, Network Optimization

**Federal University of Pará** **Brazil**  
*Bachelor of Engineering in Computer Engineering* *2015–2019*

*Thesis:* “A Handover Algorithm For Video Distribution Over Vehicular Networks”

**Honor:** CNPq Research Scholarship Recipient

## Professional Experience

---

**CPQD — Centro de Pesquisa e Desenvolvimento em Telecomunicações** **Brazil**  
*AI Research Scientist (Pesquisador Especialista I)* *2025–Present*

- Core researcher on **Project INSPIRE** (R\$390M, FINEP/MCTI & Ministry of Management): Brazil’s national AI-for-government program, with 94+ active workstreams across data, AI platforms, intelligent applications, and cybersecurity

- Co-designed the national **AI governance methodology** and project frameworks adopted by federal agencies

- Developed **technical training curricula** on **RAG, LLMs, LLMOps, and knowledge bases** for government practitioners across multiple technical maturity levels

- Technical consultant on the **Unified Federal AI Platform**: sovereign AI infrastructure serving millions of Brazilian citizens through gov.br

- Applied AI research for industry clients: **Vale, Nitro Química, Gaia Greentech, TrinitySys**

**Federal University of Pará** **Brazil**  
*Postdoctoral Research Fellow* *2025*

- Research in **federated learning** optimization and security for vehicular edge networks

- International collaboration with the University of Bern research group

- Supervision of graduate and undergraduate students

**University of Bern** **Switzerland**  
*PhD Researcher* *2020–2025*

- Pioneered **distributed aggregation** approaches for vehicular federated learning

- Developed **layer selection algorithms** for partial federated learning (Best Paper Award)

- Collaborated with international research teams across Europe and South America

- Published **15+ peer-reviewed papers** in top-tier conferences and journals

- **Applied LLMs:** used **LLMs** for technical summarization and hypothesis generation; integrated **RAG** for querying technical repositories

## LLM Engineering Expertise

---

**Post-Training:** supervised fine-tuning (SFT), preference optimization (DPO, RLHF, RLAIF), re-

ward modeling, parameter-efficient methods (LoRA, QLoRA)

**Evaluation:** design of LLM evaluation pipelines and agentic harnesses — multi-turn evaluation, tool-use simulation, LLM-as-judge, reward-hacking detection, red-teaming for robustness, regression suites, and custom domain benchmarks

**Inference & Serving:** vLLM-based serving, quantization, batching strategies, latency/throughput tuning, self-hosted inference for high-volume and edge deployments

**Agentic & RAG Systems:** multi-agent orchestration (**LangGraph, CrewAI**), RAG pipelines and knowledge bases, tool/function calling, guardrails, memory & context management, LLMOps

**Deployment:** production systems in federal government (Project INSPIRE / gov.br sovereign AI platform) and for industry clients; technical training design for practitioners at beginner-to-advanced maturity levels

## Research Impact & Key Publications

---

**Citations:** 223+ citations, h-index 9, i10-index 9 (Google Scholar)

**Top Papers:** [IEEE TNSM 2021, 29 citations] “Predictive UAV Base Station Deployment and Service Offloading with Distributed Edge Learning”

[Computer Networks 2020, 19 citations] “Skipping-based Handover Algorithm for Video Distribution Over Ultra-dense VANET”

[IEEE Access 2024] “A Distributed Aggregation Approach for Vehicular Federated Learning”

**Recent Work:** Additional papers accepted and under review (2025–2026); full list on Google Scholar and Lattes CV

## Awards & Recognition

---

**2024: Best Paper Award** — IEEE PerCom SPT-IoT Workshop: “An Efficient Layer Selection Algorithm for Partial Federated Learning”

**2021: 2nd Place Award** — Bern Data Science Day: “Distributed and Federated Learning Optimization with Federated Clustering”

**2021: Best Poster Award** — NetSys Conference: “Federated User Clustering for non-IID Federated Learning”

**2018: Honorable Mention** — WGRS Workshop for QoE/QoS-aware handover research

## Technical Skills & Certifications

---

**ML Frameworks:** PyTorch, JAX, TensorFlow, **LLM Stack:** vLLM, LangGraph, CrewAI, DeepSpeed, FSDP, FedJAX, LangChain, Hugging Face (Transformers, PEFT/TRL)

**Programming:** Python, C++, Rust, JavaScript, **Infra & MLOps:** Kubernetes, Docker, Continuous Deployment, Linux, Git, MySQL, PostgreSQL

**Cloud & Edge:** Google Vertex AI, Self-Hosted Inference, Edge Computing, 5G Networks, IoT, NS-3

**Professional Certifications**.....

**2025: Fortinet Certified Associate in Cybersecurity** — valid until 2027

**2025: Fortinet Certified Fundamentals in Cybersecurity** — valid until 2027

## Open Source Contributions & Technical Leadership

---

**GitHub:** 92+ repositories focused on federated learning, edge computing, and cybersecurity tools

**Key Projects:**

- **fl-ns3:** NS-3 based federated learning simulation framework (C++)

- **FedNetSim:** TensorFlow-based federated learning simulation tool

- **mysql2postgres:** High-performance database conversion tool (Rust)

- **mycelium\_search:** Full-text search engine built on Tantivy (Rust)

## Research Collaborations & International Network

---

**International Partnerships:** University of Bern (Switzerland), University of Texas at Austin (USA), Northeastern University (USA), Federal University of Pará (Brazil)

**Research Networks:** IEEE, Brazilian Computer Society (SBC), IFIP

**Conference Participation:** IEEE PerCom, NOMS, NetSys, DCOSS, LADC, SBRC

## Additional Information

---

**Security Clearance:** Available for roles requiring security background verification

**Availability:** Open to senior research positions, technical leadership roles, and consulting opportunities

**Specializations:** AI security + governance + engineering: LLM post-training and evaluation at scale, sovereign AI infrastructure, and red-teaming/robustness of production LLM systems