JOB OFFER

Postdoctoral Researcher

AI LLM Engineer

LiGHT

Laboratory for intelligent Global Health & Humanitarian Technologies

EPFL

School of Computer and Communication Sciences

HARVARD T.H. CHAN SCHOOL OF PUBLIC HEALTH

Position

Postdoctoral Researcher or AI LLM Engineer

Percentage & Contract

100% - Fixed term, full-time, 1 year with possible extension

Location

Lausanne, Switzerland (EPFL) or Boston, USA (Havard/ Ariadne Labs)

Start Date

ASAP (published November 1st, 2025)

Application

Apply via the following link:

http://www.light-laboratory.org/joinus

About LiGHT

The Laboratory for Intelligent Global Health and Humanitarian Response Technologies (LiGHT) is an impact-driven research group creating novel AI tools specifically adapted to limited-resource and volatile global health settings and bringing them to scaled, locally owned implementation.

LiGHT is jointly hosted at:

- Harvard T.H. Chan School of Public Health & Ariadne Labs – Ariadne Labs is a world-leading center for implementation science and innovation in health systems, co-hosted by Harvard T.H. Chan School of Public Health and Brigham and Women's Hospital in Boston.

- Swiss Institute of Technology (EPFL "École Polytechnique Fédérale de Lausanne") – In the top 10 computer science institutions globally, EPFL is a deeply international hub for technical excellence in AI.

Our mission is to develop safe, scalable, and context-aware digital health tools. We work with international NGOs such as WHO, MSF, ICRC, and national ministries of health to design, create, validate and deploy trustworthy AI in real-world healthcare systems—especially in settings where clinical resources are limited and stakes are high.

About the Role

We are looking for a technically exceptional and impact-focused Senior AI Engineer or Postdoctoral Fellow to lead research and development within the Meditron initiative—our suite of evolving open-source medical LLMs and multimodal foundation models.

You will work across large-scale model training, multimodal data integration, medical safety evaluations, and real-world benchmarking to advance usable, trustworthy AI for global health. This is a research-driven role at the intersection of machine learning, medicine, and humanitarian response.

You'll be embedded in a large, fast-paced team of 25 to 50 graduate and undergraduate students, engineers, and fellows working across LLM training, safety, and real-world validation. The role demands both technical leadership and collaborative agility—
supporting junior contributors, exchanging ideas with diverse experts, and maintaining scientific rigor across a distributed team environment.

Your Responsibilities

- Architect and scale LLM training pipelines using FSDP / DeepSpeed / HuggingFace Accelerate across multi-node, multi-GPU environments (A100/H100, 4k+ GPUs)
- Lead modular model development across base models (e.g., LLaMA, Mistral, Phi, Gemma), applying techniques like LoRA, QLoRA, MoE, FlashAttention-2, and sparse transformer variants
- Contribute to safety alignment efforts via techniques including RLHF, DPO, rejection sampling, red-teaming, chain-of-thought guidance, and calibration audits
- Collaborate on data pipeline audits: from tokenizer vocab design to dataset deduplication, differential privacy, and synthetic supervision (e.g., self-instruct)
- Benchmark across general (lm-eval-harness, HELM) and domain-specific tasks (clinical QA, guideline adherence, safety hallucination detection)
- Publish impactful findings and present internally (research rounds) and externally (NeurIPS, ICLR, ML4H, etc.)
- Track, version, and analyze experiments using modern MLops tools (e.g., WANDB, Hydra, MLFlow)
- Build with humility and clarity—collaborating with clinicians, health ministries, and humanitarian orgs to translate model capabilities into real-world safety

Required Skills

- End-to-end experience training foundation models >1B parameters, including distributed pretraining and multi-stage fine-tuning
- Proficiency in frameworks like PyTorch, DeepSpeed, HuggingFace Transformers, Safetensors, and FSDP
- Deep knowledge of LLM internals: rotary embeddings, kv-caching, positional encoding variants, optimizer/learning rate schedules, etc.
- Strong Python software engineering practices (type safety, reproducibility, tests, containerization, CI/CD pipelines)
- Experience managing and optimizing multi-modal or hierarchical input structures (e.g., combining free text with imaging metadata or signal inputs)
- Excellent scientific communication: publication track record, research storytelling, and clear documentation

Preferred Skills

- Familiarity with medical LLM applications or clinical decision support systems
- Experience with safety-critical evaluation protocols (e.g., benchmark leakage detection, hallucination profiling, audit trails)
- Experience contributing to impactful open-source projects
- Deep interest in equity-centered deployment and global health implications of LLMs
- Prior work in low-resource or global health contexts

Our Stack

- **Frameworks:** PyTorch, HuggingFace, DeepSpeed, FSDP, TensorBoardX, Hydra, Safetensors
- **Data Infra:** WebDataset, DVC, Petastorm, HuggingFace Datasets
- **Orchestration:** Slurm, Kubernetes, Docker, WandB, MLFlow
- **Evaluation:** lm-eval-harness, OpenCompass, MedQA
- **Compute:** SwissAlps (CSCS), local HPC, and cloud GPU instances

What We Value

At LiGHT, we cultivate a working culture that is mission-driven, ambitious, and deeply human. Our values guide how we build, how we collaborate, and how we hold ourselves accountable in high-stakes environments.

What We Offer

- The opportunity to shape real-world deployments with international organizations and teams
- Close collaboration with extraordinary researchers
- Travel opportunities, community, and impact at global scale
- Competitive salary and benefits, commensurate with experience

The 11 Guiding LiGHTs

1. **Curious Creativity**. We value bold exploration and context driven innovation, blending imagination with technical precision to solve real-world problems with novel solutions.
2. **Radical Ambition**. We believe that how we spend our time can shape global equity. We pursue work that matters with intensity, excellence, urgency, and endurance.
3. **Autonomous Mission-Driven Collaboration**. In the spirit of ubuntu ("I am because we are"), we value initiative grounded in shared responsibility. Everyone leads, and everyone contributes.
4. **Modular Rigor**. We move quickly within a flexible structure of organization that supports iteration, enabling disciplined work in dynamic settings.
5. **Generous Honesty**. We speak with clarity, evidence, and care. Feedback is direct, constructive, and rooted in shared purpose, respect, and a commitment to integrity.
6. **Embracing Uncertainty**. We see ambiguity as opportunity. Continuous learning and responsive adaptation are part of our core process.
7. **Lived Intelligence**. We curate extraordinary people leading full, interesting, adventurous lives, and believe that this diversity is the foundation of our creativity and collaborative strength.
8. **Constructive Joy**. We value humour, warmth, connection, and a sense of play to build resilient teams capable of sustaining focus, care and a culture of passionate quality.
9. **Grounded Perspective**. We work in close partnership with the communities we serve; originating from, travelling to, and embedding in the places where we co-design inclusive technologies rooted in local realities.
10. **Impact During and With Research**. We integrate implementation into the research process—using academic rigor and resources to build tools that are not only publishable, but provably useful.
11. **Principled Research**. We are guided by the humanitarian principles that have long anchored global health. As an academic institution, we operate with evidence, openness, and integrity—free from profit or political influence. We build knowledge into action, sharing both success and failure with humility and rigor.
    - **Humanity**. No one is too remote for innovation, or too underserved for dignity.
    - **Independence**. We act free from financial or political conflict, using academic freedom to advance knowledge in service of people, not agendas.
    - **Neutrality**. Our tools, partnerships, and research respond to need—not affiliation—and are applied with care across all contexts.
    - **Impartiality**. We focus on urgency and equity, directing our work where it's most needed—not where it's most visible.