LLM Post-Training & Systems Engineer – Alexandra Institute, Danish Foundation Models

Do you want to help build Denmark's open, national language models as part of ***Danish Foundation Models****, and work hands-on with post-training and reinforcement learning at scale? Do you enjoy making large distributed training runs fast and reliable on some of Europe's largest supercomputers? Then you might be our new colleague!*

About the job

Few places in Europe let you train large language models for real – and publish everything you build. Danish Foundation Models (DFM) is one of them. DFM is a government-funded consortium of the Alexandra Institute, Aarhus University, the University of Copenhagen and the University of Southern Denmark, building open-source language models and AI infrastructure for Denmark. At the Alexandra Institute's AI Lab, our team develops and releases DFM's open models. For our recent Munin 2.0 release – a 120B (12B active) mixture-of-experts model post-trained across 128 GPUs on the LUMI and Gefion supercomputers – we ran SFT and reinforcement learning and published both the weights and our RL setup openly.

We are now looking for an LLM Post-Training & Systems Engineer to help design, run and scale the training behind our upcoming models. That means both training models well – iterating on post-training recipes, RL algorithms, hyperparameters and training stability – and the systems engineering needed to run large models reliably and efficiently on finite public supercomputer allocations.

You will join a small, flat team working closely with colleagues who own our data and evaluations, where you own real decisions on both how we train our models and the infrastructure that runs them – seeing your work directly shape our next open releases.

As a GTS institute (a Research and Technology Organisation), we have a special role in Danish society. We act as independent advisors with a strong focus on open standards, open source and independent digital solutions – often with a research foundation and always with a responsible mindset. In the DFM project, you will collaborate with researchers at our partner universities as well as with companies and public organisations putting open Danish language models to use.

(More excited about evaluation, RL environments and data mixtures? Take a look at our companion opening: **LLM Data & Evaluation Engineer**.)

Your future role

You will be in the engine room of our post-training: the distributed training and RL systems that turn our data, environments and compute into better open models.

- **Multi-node post-training at scale (SFT & RL):** Design, optimize, and run distributed post-training for large dense and MoE models across multi-node GPU allocations.
- **Agentic RL infrastructure:** Build and scale the orchestration layer connecting policy training, high-throughput inference/rollout engines, and thousands of isolated execution environments (containers/sandboxes) for coding, tool use, and multi-step reasoning.
- **Systems performance & reliability:** Profile and eliminate bottlenecks across distributed training, rollout generation, data loading, and checkpointing; methodically debug real-world distributed failures (divergence, numerical edge cases, OOMs, straggler nodes, NCCL/RCCL hangs, environment deadlocks).
- **Training & RL recipes:** Iterate on post-training recipes, RL algorithms, and hyperparameters - diagnosing training and RL dynamics (such as loss stability, entropy/KL collapse, and off-policy effects) to improve training runs, and get maximum capability out of finite supercomputer allocations.

•**Open-source releases & publishing**: Ship clean model checkpoints, training code, inference recipes, and technical write-ups.

The competences we're looking for

We imagine that you:

•Have hands-on experience training, post-training, or serving large neural networks across multi-GPU / multi-node setups - in industry, a research lab, a PhD project or serious open-source work.

•Bring strong engineering skills in **Python** and deep hands-on experience with **PyTorch** and modern LLM training/inference stacks (e.g. TorchTitan, Megatron, vLLM, SGLang or prime-rl).

•Have practical experience with Slurm / HPC clusters or large-scale GPU orchestration (shared filesystems, distributed collectives, containers like Apptainer/Docker, and keeping long jobs alive).

•Have a solid grasp of modern LLM post-training methods (e.g. SFT, RL / RLVR / GRPO, MoE architectures, throughput optimization).

•Bring good, pragmatic judgment: you enjoy digging into messy distributed failures, care about reproducibility, and write clean code that a small team can maintain.

•Are interested in building or scaling sandboxed execution environments for agentic RL.

•Are familiar with – or keen to learn – both AMD ROCm and NVIDIA CUDA.

•Ideally have a public track record: open-source ML contributions, released models, technical blog posts or papers (a plus, but not a requirement).

•Communicate fluently in English and can explain complex concepts to both technical and non-technical audiences. Danish is not required for this role.

•Already have the right to work in Denmark (see below).

We are open to different levels of experience and educational backgrounds - from strong recent graduates and PhDs to experienced engineers and researchers from industry or academia. What matters most is real hands-on depth and the drive to learn fast. You don't need to tick every box.

Your personality, passion and approach to work are just as important to us. You are self-driven and take ownership, but you also enjoy being part of a team, collaborating closely, exploring and trying out new technologies – and sharing knowledge with your colleagues.

Your future workplace

You will join the Alexandra Institute's AI Lab, a team of about 25 people split between Aarhus (Katrinebjerg) and Copenhagen (the IT University of Copenhagen). You will be based in Copenhagen, where the DFM team sits, with the option of working from home a day or two a week.

Expect a warm, informal atmosphere and colleagues who are exceptionally good at what they do. We work flexible hours, start every Friday with breakfast together, and offer a lunch scheme and health insurance. You'll have a lot of freedom and be trusted to use it well, with plenty of room to grow, both professionally and personally.

Sounds interesting?

Then we would love to hear from you. Send us your CV and a short note (half a page is plenty) no later than **6 December**. In the note, tell us about a multi-node training run, RL/inference system or hard distributed systems problem you have worked on: what you did, what was hard, and what you learned. We don't need a formal cover letter; we would rather hear about your work. Links to GitHub, Hugging Face, blog posts or papers are very welcome.

First-round interviews will take place December 10th to 14th.

Second-round interviews will take place December 16th to 18th.

Please note: Due to our timeline, we can only consider candidates who already hold the right to work in Denmark (EU/EEA citizens or existing permit holders).

If you have questions about the position, you are welcome to contact Head of AI Lab Mie Hvas at mie.hvas@alexandra.dk / +45 20 63 99 58.
At the Alexandra Institute, we believe that innovation thrives when our employees have many different backgrounds – professionally as well as in terms of gender, religion, nationality and ethnicity. We therefore encourage all qualified candidates to apply, regardless of background.