

At Crisalix, we are looking for a junior engineer to join the Deep Learning team. You will work across the full model development lifecycle: from generating datasets and training models from scratch, to experimenting with new algorithms and deploying them into production. This is a role where you will grow into an engineer who can bridge research and engineering: from reading a research paper to shipping it to production.

You can apply here: <https://www.linkedin.com/jobs/view/4379677318>

Essential Requirements

- **Deep learning fundamentals:** Solid understanding of how models are trained: loss functions, optimization, backpropagation, overfitting, regularization, etc.
- **Python and PyTorch:** Proficient in Python and hands-on experience training models in PyTorch or another major deep learning framework.
- **Practical engineering skills:** Comfortable with Git, GitHub (pull requests, code reviews), command line, and the general software development workflow. Should be able to pick up an open-source repo, set it up, and get it running.
- **Research literacy:** Ability and willingness to read, understand, and implement ideas from academic papers, journals.

Nice to Have

- Experience with **computer vision** (image classification, segmentation, landmark estimation, etc.).
- Experience with **3D reconstruction**. Parametric models (SMPL), neural fields (NeRF, gaussian splats).
- Experience with **generative models** (training, fine-tuning, inference). Diffusion, GANs, VAEs, etc.
- Exposure to **ML infrastructure and deployment** (MLOps, model serving, containerization, CI/CD for ML).