

Machine Learning Engineer

Machine Learning & EDA · Full-Time · Oxford/London/Cambridge/Remote

Level	Mid / Senior
Reports to	VP Engineering
Location	Oxford, London, Cambridge UK (Hybrid) or Remote

Mach42 is a venture-funded startup at the intersection of machine learning and Electronic Design Automation (EDA). We are developing ground-breaking machine learning tools set to revolutionise the verification of complex integrated circuits – helping chip design teams accelerate simulation coverage, improve product quality and reduce tapeout risk.

THE ROLE

We are looking for a Machine Learning Engineer to build the systems that turn research prototypes into a reliable, production-grade product. You will own the data pipelines, model training infrastructure, inference stack, and export toolchain that get surrogate models into the hands of analog designers. You'll work at the boundary of ML engineering and EDA tooling, shipping software that runs inside some of the world's most demanding chip design environments.

WHAT YOU'LL DO

- Implement and productionise ML training pipelines for surrogate model generation, from raw circuit netlists and simulation data ingestion through to trained, validated models .
- Build and maintain the model export layer — translating PyTorch and Jax surrogate models into code that can run natively in an EDA environment with verified functional equivalence.
- Develop evaluation harnesses and regression suites that compare surrogate model outputs against ground-truth simulation data.
- Profile and optimise model inference speed to meet real-time simulation-in-the-loop latency targets.
- Collaborate with researchers to translate novel algorithms into robust, maintainable production code.

– WHAT WE'RE LOOKING FOR

- Good first degree in a relevant STEM subject.
- 3+ years of industry experience developing production ML systems.
- Strong Python skills and experience with PyTorch or JAX for model training and inference.
- Experience building data pipelines for scientific or engineering data — handling irregular structures and domain-specific formats.
- Solid software engineering fundamentals: testing, CI/CD, code review, clean API design.

- Ability to work across the system end-to-end, from input data processing to model training/inference/export to integration with backend and desktop applications.
- Clear communication — you can write a spec, document an API, and give a live demo to a customer.

NICE TO HAVE

- Familiarity with circuit simulation, behavioural modelling languages and EDA tools.
- Experience with C++/CUDA for performance-critical inference components.
- Prior startup experience or comfort working in a small, fast-moving team.

WHAT WE OFFER

- Competitive salary, discretionary bonus and meaningful equity in a well-funded, early-stage company.
- Flexible remote/hybrid working with London/Cambridge/Oxford home base(s).
- Support for learning and development.
- Close collaboration with researchers and direct exposure to customer workflows.
- A product that ships to real customers — your work ends up running inside chip design tools at semiconductor companies worldwide.

Mach42 is an equal opportunity employer. We welcome applications from all backgrounds and are committed to building a diverse team.