We are looking for an engineer with robust experience in machine learning and strong mathematical foundations to join our growing ML team and to help drive the direction of our ML platform.

Machine learning is a critical pillar of Epirus's global business. Our ever-evolving trading environment serves as a unique, rapid-feedback platform for ML experimentation, allowing us to incorporate new ideas with relatively little friction. Our ML team is full of people with a shared love for the craft of software engineering, and for designing APIs and systems that are delightful to use.

We'll rely on your in-depth knowledge of the ML ecosystem and understanding of varying approaches — whether it's neural networks, random forests, gradient-boosted trees or sophisticated ensemble methods — to aid decision-making so we apply the *right* tool for the problem at hand. Your work will also focus on enhancing research workflows to tighten our feedback cycles. Successful ML engineers will be able to understand the mechanics behind various modelling techniques, while also being able to break down the mathematics behind them.

If you've never thought about a career in finance, you're in good company. Many of us were in the same position before working here. While there isn't a fixed list of qualifications we're looking for, if you have a curious mind and a passion for solving interesting problems, we have a feeling you'll fit right in.

We're looking for someone with:

- Experience building and maintaining training and inference infrastructure, with an understanding of what it takes to move from concept to production
- A strong mathematical background; Good candidates will be excited about things like optimisation theory, regularisation techniques, linear algebra and the like
- A passion for keeping up with the state of the art, whether that means diving into academic papers, experimenting with the latest hardware or reading the source of a new machine learning package
- A proven ability to create and maintain an organised research codebase that produces robust, reproducible results while maintaining ease of use
- Expertise wrangling an ML framework – we're fans of PyTorch, but we'd also love to learn what you know about Jax, TensorFlow or others
- An inventive approach and the willingness to ask hard questions about whether we're taking the right approaches and using the right tools
- Fluency in English