Full-Stack Engineer, Muse Team

You are interested in **helping build and run an API powered by large machine learning models answering real-world business needs**. You want to bring state-of-the-art artificial intelligence technology to the world, democratising its use through a simple, reliable, and responsive API.

Our vision. At LightOn, **we are building transformative artificial intelligence.** We believe that scale enables significantly increased capabilities in machine learning models: we seek to democratise extreme-scale models, and we are looking for passionate people to join us in that mission.

Your mission. **You will be joining a tightly knit team, seeking to bring to the world extreme-scale models with human level performance on real-world tasks**. We distribute access to our natural language processing models through our Muse API, serving up to millions of words every month — but we are just getting started. You will contribute to the end-to-end development of the Muse API. In the back-end, you will build and maintain serving, monitoring, and billing services in Python, optimised for speed, reliability, and scale. In the front-end, you will implement and maintain responsive and engaging websites in Javascript, showcasing our API and its capabilities.

| Your daily responsibilities

- 💪 > Draft, build, and maintain high-quality services in Python to power the Muse API;
- 🌐 > Implement, and maintain engaging websites in Javascript to showcase Muse;
- 📈 > Optimise software for speed, scale, engagement, and reliability;
- 🗒️ > Document clearly, internally and in public, tools, techniques, and best practices;
- 🚀 > Help shape the roadmap for our API, collaborating with a cross-functional team.

| Your ideal background

- ☁️ > Extensive knowledge of backend architectures and development;
- 👨‍💻 > Strong knowledge of Python (e.g. Django), Javascript (Vue), and other web technologies;
- 🗄️ > Experience with relational databases (PostgreSQL) and in-memory store (Redis);
- 💪 > Knowledge of containerisation (Docker) and cloud services automation (GCP) is a plus;
- 🧠 > Passionate and thoughtful about artificial intelligence and natural language processing.

We are not credential-driven. We want to understand your unique background and what you can bring to the team. We are excited about people who may not already have a track-record in the field, but who are self-starter and show high potential, enthusiasm, and a willingness to grow.

Benefits. Salary according to your profile, meal vouchers, and flexible hours.

Everyone is welcome. Depending on the global health situation, the position may either take place in our offices in Paris, or remotely. You should be eligible to work in the EU.

To apply, send your resume to:
eng_search@lighton.ai