

Software Engineer, Machine Learning

Description

Join our team at Bumble, where we're revolutionizing online dating through innovative AI-driven experiences. We're seeking machine learning engineers to spearhead the development of next-generation of online dating. In this role, you will be a foundational member of a small, dynamic team dedicated to creating cutting-edge solutions that redefine how people connect and form relationships online. The ideal candidate thrives in a fast-paced environment, is a creative thinker, and has a proven track record in Generative AI and Machine Learning. If you're passionate about leveraging AI to shape the future of online connections, we want to hear from you!

What you'll be doing:

- Work in a small, dynamic team that values agility, creativity, and autonomy
- Contribute to the technical development of new (generative) AI-powered features in our apps
- Collaborate cross-functionally with teams across the business to identify user needs and design AI and Machine Learning solutions that meet these needs
- Collaborate with the [MLOps](#) team to optimize AI and Machine Learning solutions for performance and scalability
- Mentor and coach team members on ML engineering best practices
- Stay up to date with the latest developments in AI to drive innovation

We'd love to meet someone with:

- An advanced degree in Computer Science, Mathematics or a similar quantitative discipline
- Hands-on experience in delivering machine learning models to production at scale
- Demonstrated ability to develop innovative technical solutions to complex problems
- Experience in writing production-quality Python code
- Comfortable working with classic ML frameworks, such as Pytorch or TensorFlow
- Comfortable working with modern LLM deployment frameworks and libraries, such as HuggingFace TGI, VLLM, TensorRT-LLM, or similar
- Strong understanding of machine learning applications development life cycle processes and tools: CI/CD, version control (git), testing frameworks, MLOps, agile methodologies, monitoring and alerting
- Comfortable working with Docker and containerised applications
- Experienced at deploying highload ML applications on container orchestrators (bare-metal k8s, GKE, EKS)
- Understanding of the concepts of GPU-powered workloads, NVIDIA drivers, container runtimes
- Strong communication skills, and the ability to work collaboratively and proactively in a fast-paced environment alongside technical and non-technical stakeholders
- A genuine passion for AI and Machine Learning, and a thoughtful approach to AI fairness, accountability, and transparency.

Hiring organization

Candidate-1st

Employment Type

Full-time

Beginning of employment

asap

Job Location

UK London

Working Hours

40

Base Salary

euro GBP 32K - 59K *

Date posted

May 21, 2024

How the process will look like

Your teammates will gather all requirements within our organization. Then, once priority has been discussed, you will decide as a team on the best solutions and architecture to meet these needs. In continuous increments and continuous communication between the team and stakeholders, you're part of making data play an even more important (and understood) part withing Brand New Day.

Job Benefits

GBP 32K – 59K *