

Machine Learning Operations Engineer

Description

Grammarly is excited to offer a [remote-first hybrid working model](#). Team members work primarily remotely in the United States, Canada, Ukraine, Germany, or Poland. Certain roles have specific location requirements to facilitate collaboration at a particular Grammarly hub.

All roles have an in-person component: Conditions permitting, teams meet 2–4 weeks every quarter at one of Grammarly’s hubs in San Francisco, Kyiv, New York, Vancouver, and Berlin, or in a workspace in Kraków. This flexible approach gives team members the best of both worlds: plenty of focus time along with in-person collaboration that fosters trust and unlocks creativity.

Grammarly team members in this role must be based in Germany or Poland, and they must be able to collaborate in person 2 weeks per quarter, traveling if necessary to the hub(s) where the team is based.

The opportunity

[Grammarly](#) is the world’s leading AI writing assistance company trusted by over 30 million people and 70,000 professional teams every day. From instantly creating a first draft to perfecting every message, Grammarly’s product offerings help people at 96% of the *Fortune* 500 get their point across—and get results. Grammarly has been profitable for over a decade because we’ve stayed true to our values and built an enterprise-grade product that’s secure, reliable, and helps people do their best work—without selling their data. We’re proud to be one of *Inc.*’s best workplaces, a Glassdoor Best Place to Work, one of *TIME*’s 100 Most Influential Companies, and one of *Fast Company*’s Most Innovative Companies in AI.

To achieve our ambitious goals, we’re looking for an ML Operations Engineer to join our Core Product team. Our AI-powered writing assistant leverages the latest NLP and DL technologies for state-of-the-art Grammatical Error Correction and advanced semantic features to help our users communicate clearly and confidently.

This role is for someone excited about elegant engineering solutions that scale to millions of users and more passionate about making deep learning efficient than training deep learning models. This role is at the heart of the Core Product organization, which develops flagship Premium features to help users improve their communication.

Grammarly’s engineers and researchers have the freedom to innovate and uncover breakthroughs—and, in turn, influence our product roadmap. The complexity of our technical challenges is growing rapidly as we scale our interfaces, algorithms, and infrastructure. You can hear more from our team on our [technical blog](#).

Your impact

As an ML Operations Engineer, you will:

- Design, implement, and improve an ML inference infrastructure to serve large language models to millions of active daily users.

Hiring organization

Candidate-1st

Employment Type

Full-time

Beginning of employment

asap

Job Location

Poland; Hybrid

Working Hours

40

Base Salary

euro USD 60K - 112K *

Date posted

May 16, 2024

- Collaborate with researchers to optimize in-house large language models to find the best latency-cost-quality tradeoffs.
- Own integrations of 3rd party LLM providers. Plan required capacity; implement monitoring, alerting, and throttling to guarantee desired latency and high reliability in a high-load environment.
- Bridge the gap between feature and platform teams, bringing the best from platform tools to the feature teams and providing a feedback loop from feature teams to platform teams.
- Research existing open-source tools and [MLOps](#) approaches other companies take to ensure we build best-in-class technology while leveraging available components.
- Guide the organization through buy vs. build trade-offs.
- Provide a common set of infrastructure, orchestration, and monitoring to help make Grammarly's diverse array of machine learning systems more maintainable.

We're looking for someone who

- Embodies our EAGER values—is ethical, adaptable, gritty, empathetic, and remarkable.
- Is inspired by our MOVE principles, which are the blueprint for how things get done at Grammarly: move fast and learn faster, obsess about creating customer value, value impact over activity, and embrace healthy disagreement rooted in trust.
- Is able to collaborate in person 2 weeks per quarter, traveling if necessary to the hub where the team is based.
- Understands traditional machine learning algorithms as well as modern deep learning approaches.
- Is aware of unique challenges with sparse feature sets and the Transformer-based pre-trained deep learning models central to NLP.
- Understands data structures and algorithms at a level sufficient to write performant code when working with large datasets or incoming data streams.
- Has experience with system design and building internal tools.
- Has a consistent record working with internal partners such as data platforms, analytics, data science teams, and strategic stakeholders (Product Managers and Security teams).
- Stays current in the fast-moving field of ML, DL, NLP, and MLOps.
- Is knowledgeable of existing ML cloud platforms and tools.

Support for you, professionally and personally

- *Professional growth:* We believe that autonomy and trust are key to empowering our team members to do their best, most innovative work in a way that aligns with their interests, talents, and well-being. We also support professional development and advancement with training, coaching, and regular feedback.
- *A connected team:* Grammarly builds a product that helps people connect, and we apply this mindset to our own team. Our remote-first hybrid model enables a highly collaborative culture supported by our EAGER (ethical, adaptable, gritty, empathetic, and remarkable) values. We work to foster belonging among team members in a variety of ways. This includes our employee resource groups, Grammarly Circles, which promote connection among those with shared identities including BIPOC and [LGBTQIA+](#) team members, women, and parents. We also celebrate our colleagues and accomplishments with global, local, and team-specific programs.

We encourage you to apply

At Grammarly, we value our differences, and we encourage all—especially those whose identities are traditionally underrepresented in tech organizations—to apply. Grammarly is an equal opportunity company. We do not discriminate on the basis of race or ethnic origin, religion or belief, color, gender expression or identity, sexual orientation, sexual identity, national origin, citizenship, age, marital status, veteran status, disability status, criminal prosecution, judgment in a criminal case, or any other characteristic protected by law.

For more details about the personal data Grammarly collects during the recruitment process, for what purposes, and how you can address your rights, please see the Grammarly Data Privacy Notice for Candidates [here](#).

#LI-AT1

#LI-Hybrid

All team members meeting in person for official Grammarly business or working from a hub location are strongly encouraged to be vaccinated against COVID-19.

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

USD 60K - 112K *