Timur Ganiev

ganiev.tmr@gmail.com
github.com/esceptico
linkedin.com/in/esceptico

Summary

Machine Learning Engineer specializing in post-training and safety alignment for production LLM systems. I build preference learning pipelines, train safety classifiers, and optimize low-latency inference. Previously led alignment work at Replika, shipping models trained on user feedback at scale.

Experience

Replika Lead Machine Learning Engineer

Oct. 2022 – Present

- Owned preference alignment using DPO: built training and evaluation codebase, co-designed feedback instrumentation with product, frontend, and backend teams, implemented data cleaning with safety filtering and roleplay handling, and trained models that delivered statistically significant improvements on business metrics in A/B tests.
- Led safety alignment end-to-end: created adversarial datasets for rejection and recovery behaviors, designed red-teaming evaluations, and trained a safety classifier with entropy-based sampling; improved rejection rate from ~5% to ~60% while reducing false negatives by ~60%.
- Co-led conversation system redesign: contributed to architecture and tool-calling flow, owned context engineering and management, and implemented short-term memory.
- Maintained distributed LLM inference serving 100+ RPS; optimized parallel tool execution, reducing P95 latency by ~40% for multi-tool responses.

Embedika Machine Learning Engineer

Feb. 2022 – Sep. 2022

- Built and deployed active learning service for multi-modal data annotation.
- Developed toxic content classifier achieving 94% F1-score and deployed spell checking service using BERT.

Sber Machine Learning Engineer

Jan. 2020 – Feb. 2022

- Led a team of 4 engineers, managed 3 projects from research to production.
- Built multi-target text classification model achieving 93% F1-score on noisy multitask dataset; improved robustness by 4% F1 using Integrated Gradients and adversarial augmentation.
- Optimized model inference by 80% through distillation, ONNX conversion, and quantization.

Advanced.Careers Data Scientist

Aug. 2018 – Aug. 2019

- Built resume and job posting parsing systems, improving CV upload rate by 10%.

Education

Kazan National Research Technical University B.S. in Computer Science

Sep. 2014 – Aug. 2018

Thesis: "Personality traits prediction based on Twitter account data"

Skills

Post-Training & Alignment	RLHF, DPO, SFT, preference modeling, safety alignment, adversarial evaluation
Inference & Systems	vLLM, DeepSpeed, FSDP, quantization, distillation, distributed serving
LLM Applications	agentic workflows, memory & context engineering, RAG, tool use
Frameworks & Tools	PyTorch, Transformers, TRL, PEFT, Weights & Biases, MLFlow
Infrastructure	Python, FastAPI, Docker, Redis, Linux, CI/CD