Debtirtha Saha

Data Scientist, Rakuten Mobile Inc.
Tokyo, Japan
debtirthasaha.github.io

+81 70-8549-1320
debtirthasaha1@gmail.com
LinkedIn
GitHub · Hugging Face

Experience

•Rakuten Mobile Inc. *Jan. 2024 – Present*
Data Scientist Tokyo, Japan

- **Telco-domain SLM:** Building a domain-adapted Qwen3-1.7B for low-latency, on-edge telco deployment via CPT, SFT, and DPO; own RAN-domain data extraction and curation, with the model evaluated across **7 telco domains** (RAN, Core, OSS, IPTX, Security, BSS, Cloud).
- Built the LLM-as-judge evaluation harness (GPT-4o judge + BERTScore over **∼7K Q&A pairs** across all 7 domains, scoring correctness, factuality, grounding, reasoning, safety, and latency) and designed a 75/25 telco/generic CPT mix to curb catastrophic forgetting.
- Benchmarking selected the 1.7B over a 4B candidate at **matched quality and ∼2.5× lower CPU latency**, enabling edge deployment.
- **Autocall optimization:** Modeled the optimal hour to contact churn-risk and payment-default customers from 2-week behavioral history — drove **+23M JPY/month** in additional collections over the payments-team baseline and retained **20%** of reached at-risk customers.
- **Agentic analytics:** Built a multi-agent LangGraph system (MinIO data lake + ClickHouse, served via Viber) automating customer and competitor analytics — natural-language data Q&A, scheduled HTML/PPT reporting (daily churn/onboarding, weekly app usage, network-quality hotspots), and competitor-campaign monitoring (web-scraped daily, summarized to marketing) — cutting **∼40 analyst-hours/month**.

Selected Projects

•Ukiyo-e Haiku VLM *2026*
Bilingual vision-language model (personal) GitHub · HF · Demo

- Bilingual (JP/EN) multimodal generation model producing haiku from ukiyo-e prints; curated paired image–text data and fine-tuned via SFT then preference optimization (ORPO / KTO) with a λ_U sweep.
- Shipped two models to the Hugging Face Hub with a live interactive demo (flagship JP model 4.36/5 avg eval), showing preference gains require a real chosen/rejected quality gap.

•Math SLM (SFT + DPO) *2026*
LLM post-training pipeline (personal) GitHub · HF

- SFT + DPO on DeepSeek-R1-Distill-Qwen-7B; teacher self-distillation data generation via vLLM (TP=8), LoRA training with DeepSpeed ZeRO-3 on 8×H100.
- Built the full pipeline (data gen → SFT → LoRA merge → eval); SFT added **+6.4 pp** average on math benchmarks; model live on the HF Hub.

•Mathematics Distillation Challenge – Equational Theories *2026*
LLM reasoning / evals competition GitHub · Competition

- Designed an 8.7 KB cheatsheet for LLM reasoning over magma equational identities, including a compact algebraic rule for ruling out false implications between equations.
- Ranked **#85 overall**, with #13 on GPT-OSS (across problem sets), #20 on Gemma (extra-hard), and #21 on GPT-OSS (order-5); $0.00066/problem at 100% parse success, in a field of ∼1K participants.

Technical Skills

- **Languages**: Python, C/C++, SQL
- **ML / LLMs**: LLM & VLM fine-tuning + multimodal generation (SFT, DPO, ORPO, KTO), knowledge distillation, CPT, LoRA / PEFT, evals (LLM-as-judge, BERTScore), churn & time-series forecasting
- **Frameworks**: PyTorch, Hugging Face (Transformers, PEFT, TRL), vLLM, DeepSpeed (ZeRO-3), NumPy, Pandas, Matplotlib
- **Infrastructure**: Multi-GPU training (A100 / H100), distributed parallelism (TP / DP / ZeRO), Git

Education & Achievements

•IIT Indore, B.Tech, Mechanical Engineering (Minor: HSS) — CGPA 8.41 *2019 – 2023*

•JEE Advanced — All India Rank 4879, among 245,000 qualified candidates *2019*