

RL training infrastructure for LLMs. You write the loop. We handle the GPUs.

MinT lets you focus on what matters in LLM post-training — your data, your loss functions, and your RL environments — while handling the heavy lifting of distributed training. Write a simple Python script on your CPU-only machine; MinT figures out how to run it efficiently on a GPU cluster.

You define **what to train, what data, how to optimize**, and **how to evaluate**. MinT handles GPU scheduling, distributed rollout, and weight management. Changing the model requires only a single string in your code.

ARCHITECTURE OVERVIEW

USER LAYER
Model · Data · Loss / LoRA · Metrics

‡ MinT SDK

MINT CORE
Model Router · Task Manager · LoRA Manager

‡ API Gateway

INFRA LAYER
GPU Clusters · Object Storage · Databases

CORE CAPABILITIES



LoRA RL
Low-Rank Adaptation RL made simple, stable, and efficient for mainstream and frontier models.



Distributed
Parallelized rollout and training across GPU clusters, fully managed — no cluster config required.



RLHF / GRPO
Robust pipeline with SFT, DPO, GRPO, and DAPO. Full control over algorithmic details.



Tinker-Compat.
Drop-in upgrade from ThinkingMachines Tinker — migrate with a single import swap.

SUPPORTED MODELS — 6 FAMILIES · 27 VARIANTS

Community

Qwen3-0.6B Qwen3-4B-Instruct-2507 Qwen3-4B-Thinking-2507

Qwen3-30B-A3B-2507 Qwen3-235B-A22B-2507 π0 VLA

Enterprise

DeepSeek-V3 DeepSeek-V3.1 DeepSeek-V3.2 Kimi-K2

GLM-5 GLM-5.1 MiniMax 2.7 Qwen3-VL Kimi-K2.5

DeepSeek-V4

Also technically compatible: Llama 3.x, Gemma 2/3, Qwen2.5. Request via sales@mindlab.ltd

EDITIONS

CAPABILITY	COMMUNITY	ENTERPRISE
Included usage	5M tokens	Custom quota
Model access	Qwen series	All 6 families
Workspace	Shared cloud	Private / VPC
Training loop	Standard	Dedicated
Support	Community	5×8 SLA
Compliance	—	MLPS L3 + audit
Data ownership	—	✓ Tenant-isolated

PRODUCTION CASE STUDIES

−90%

Medical Coding

GPU cost reduction. 27 QC indicators across 8,000+ physicians. Training steps cut by 50%.

340K

Personalized Agent

Mini-apps on shared 1T base. 1M users, only 5B LoRA diff-params per user.

+34%

Smart Customer Service

CSAT lift. 18M-user fintech. Compliance 89%→99.2%. Takeover 38%→15%.

6 hr

Citizen Services

Policy refresh time. 23M citizens, 1,400+ services. Policy errors 23%→5%.

```
$ pip install mindlab-toolkit · export MINT_API_KEY=sk-...
```

Endpoints: mint.macaron.xin · mint-cn.macaron.xin · GitHub: github.com/MindLab-Research/mindlab-toolkit · CLI:

Enterprise inquiries → sales@mindlab.ltd