

AI Model Training and Fine-Tuning Services on Your Own Data

We prepare your data and fine-tune or train AI models so they follow your terminology, formats and rules, with evaluation before and after every change.

OVERVIEW

Model Training & Fine-Tuning

A general-purpose model knows a lot about everything and little about your business. Fine-tuning closes that gap by training a pre-trained model further on your own examples, so it uses your terminology, follows your formats and handles your edge cases. We run the whole job, from building the dataset to deploying the tuned model.

Fine-tuning is not always the first step. Many problems are solved with better prompts or with retrieval over your documents, at lower cost and with easier updates. We start by testing those options against a baseline, and recommend training only where it clearly beats them on your own test set.

When training is the right choice, the dataset decides the result. Most of the effort goes into collecting, cleaning and labelling examples. We then train with parameter-efficient methods where possible, and check the tuned model for accuracy, safety and any loss of general ability before it goes live.

BUILT FOR

- Teams whose prompts have hit a quality ceiling
- Companies with archives of labelled examples
- Products that need a strict output format
- Businesses wanting a smaller, cheaper model
- Regulated firms that need a self-hosted model

FEATURES

What you get with Model Training & Fine-Tuning

01

Training strategy consultation

We review your goal, data and budget, then choose between prompting, RAG, fine-tuning or training a model from scratch.

02

Dataset preparation

Collecting, cleaning, de-duplicating and labelling prompt-response pairs or classified examples, with a held-out set kept back for testing.

03

Prompt engineering and tuning

Structured instructions, worked examples and iterative prompt refinement to improve consistency and reduce hallucinations before any training.

04

Supervised fine-tuning

Full or LoRA and QLoRA fine-tuning of open-source models, or managed fine-tuning of hosted models where the provider offers it.

05

Preference tuning

Aligning outputs with what your reviewers prefer, using ranked examples and methods such as DPO, so tone and judgement match your standards.

06

Task-specific models

Smaller models trained for one job, such as classification, sentiment analysis or field extraction, that run fast and cheaply at volume.

- 07 RAG integration**
Pairing a tuned model with retrieval, so it writes in your style while drawing facts from your current documents.
- 08 Safety and red-team testing**
Adversarial prompts, bias checks and safety thresholds that catch harmful, biased or off-policy outputs before release.
- 09 Deployment and integration**
Tuned models served in your software, website or chat platform, with versioning so you can return to an earlier model at any time.

WHY IT MATTERS

The difference it makes to your business

- 01 Consistent, on-format output**
A tuned model follows your structure, field names and house style every time, without a page of instructions in each prompt.
- 02 Lower running cost at volume**
A small tuned model can replace long prompts sent to a large one, using fewer tokens per request. We measure the trade-off before you switch.
- 03 Less data and compute than starting over**
Building on a pre-trained model needs far fewer examples and much less GPU time than training a model from nothing.
- 04 Changes you can measure**
Every training run is scored against the same held-out set, so you see exactly what improved and what got worse.

HOW WE DELIVER

How we deliver Model Training & Fine-Tuning

- 01 Consult and set a baseline**
We define the task, agree the metric and measure how a strong base model with good prompts performs today.
- 02 Build the dataset**
Examples are drawn from your records or written with your experts, then cleaned, labelled and split into training and test sets.
- 03 Train**
We pick the base model and method, then train, freezing or adapting layers as the task needs, and track every run.
- 04 Evaluate**
The tuned model is compared with the baseline on the held-out set, checked for lost general skills and red-teamed for safety.

We integrate the model, monitor its output and retrain with new examples when your products, policies or language change.

FAQ

Questions about Model Training & Fine-Tuning

What is the difference between training and fine-tuning?

Training builds a model from scratch on a very large dataset, which takes heavy compute and time. Fine-tuning starts from a model that is already trained and continues training on a smaller, task-specific dataset. It is a form of transfer learning, and it is what most business projects need.

When should we fine-tune instead of using prompts or RAG?

Fine-tune when you need consistent formats, a specific tone, specialised behaviour or lower cost per request, and prompting has stopped improving results. Use RAG when the model needs facts that change often or must cite sources. The two work well together.

How much training data do we need?

Less than most people expect, but it must be good. A focused task can often be improved with a few hundred carefully checked examples, while broader behaviour needs more. A small pilot run on part of your data shows whether adding more examples is likely to help.

What hardware is needed, and how long does training take?

Fine-tuning can technically run on CPUs, but GPUs make it practical. A single run on a smaller model may take minutes to hours, depending on model size, dataset and hardware. Most of a project's calendar time goes on preparing data and evaluating results, not on the training itself.

What is catastrophic forgetting, and how do you avoid it?

It is when a model over-adapts to a narrow dataset and loses general skills it had before. We limit it with parameter-efficient methods that change only a small part of the model, careful learning rates, mixing in general examples, and testing general ability alongside the target task.

What drives the cost of fine-tuning?

How much data has to be collected and labelled, the size of the base model, the number of training and evaluation rounds, GPU time, and how the model will be hosted. Data preparation is usually the largest share.

Who owns the fine-tuned model?

You do. On completion and final payment, the dataset we prepare and the tuned weights are handed over, as set out in the contract. The base model stays under its original licence, so we check that licence permits your intended commercial use before training starts. Models tuned through a provider's managed service stay in your account with that provider.

NEXT STEP

See whether fine-tuning will pay off

Send us a sample of the task and the outputs you expect. We will compare prompting, RAG and fine-tuning on it and recommend the cheapest option that meets your standard.

Book a free consultation: <https://appsin.ai/book-a-consultation/>

WhatsApp or call: [+91-9025555584](https://wa.me/919025555584)

About AppsInAi

AppsInAi Private Limited is an AI development company in Coimbatore, India. We build generative AI assistants, AI agents, chatbots, machine learning models, computer vision and data platforms on our clients' own data, from a first proof of concept to production, with monitoring and support after launch.

Contact

sales@appsin.ai

+91-9025555584

Housing Unit, Sf No. 393, CMS School Backside, 3rd Street, Nehru Nagar, Ganapathy, Coimbatore 641006, India

<https://appsin.ai>