Appsin.Ai

SERVICE

AI and ML Development Services for Businesses That Run on Data

End-to-end AI and machine learning development, from a feasibility check and proof of concept to models running inside your ERP, CRM or apps.

https://appsin.ai · +91-9025555584 · sales@appsin.ai
AI development for businesses in India, the GCC and the USA

OVERVIEW

AI & ML Development

We build AI and machine learning systems that solve one business problem at a time, such as forecasting demand, reading documents or spotting fraud. Each starts from your own data and ends as a model running inside the software your team already uses, with its accuracy measured.

Our work covers the whole range of applied AI: classic machine learning on tables of business data, deep learning for images, speech and video, natural language processing, and generative AI built on large language models such as GPT, Claude, Gemini and Llama. We choose the simplest approach that meets the target, not the most fashionable one.

Around the models we do the engineering that keeps them useful: data pipelines, integration with ERP, CRM and mobile apps, MLOps for versioning and monitoring, and retraining as your data changes. You can engage us for a fixed-scope project or a dedicated team of ML engineers and data scientists.

BUILT FOR

- Founders adding AI to a product
- CTOs and IT heads in growing companies
- Operations teams buried in manual work
- Enterprises with data but no ML team
- SaaS and software companies

FEATURES

What you get with AI & ML Development

01 **AI consulting and feasibility**
We review your processes and data, rank use cases by value and effort, and tell you plainly which ones are not ready for AI yet.

02 **Proof of concept**
A small working model on a sample of your real data, with an agreed accuracy target, so you can decide before committing a full budget.

03 **Predictive models**
Classification, regression and forecasting for demand, churn, credit risk, fraud and predictive maintenance, trained on your own history.

04 **Natural language processing**
Text classification, entity extraction, summarisation, sentiment analysis, translation and search across contracts, emails, tickets and reports.

05 **Computer vision**
Image classification, object detection, OCR and visual quality inspection, tested on photos and video from your own cameras and production lines.

06 **Generative AI and LLM apps**
Assistants, document automation and knowledge search built on GPT, Claude, Gemini or Llama, grounded in your documents through retrieval-augmented generation (RAG).

07 **AI agents and chatbots**
Multilingual chatbots and agents that answer questions, update records and route work across support, sales and HR, handing over to a person when needed.

08 **Data engineering**
Pipelines, ETL and ELT jobs, warehouses and data lakes that collect and clean the data your models depend on, and keep it flowing after launch.

09 **MLOps and integration**
Docker and Kubernetes deployment, MLflow or Kubeflow experiment tracking, monitoring and retraining, plus APIs into your ERP, CRM and apps.

WHY IT MATTERS

The difference it makes to your business

01 **Less manual work**
Documents, tickets and approvals that staff now read and key in by hand can be classified, extracted and routed for a person to check.

02 **Decisions based on forecasts**
Models trained on your own history give demand, risk and churn estimates that planners can act on, with the uncertainty shown alongside.

03 **One partner from data to deployment**
The same team handles data preparation, modelling, integration and support, so nothing is lost between a consultant's report and a working system.

04 **Spend matched to evidence**
A measured proof of concept comes first, so a full build is funded only once the pilot has shown it works on your data.

HOW WE DELIVER

How we deliver AI & ML Development

01 **Define the problem and the data**
We agree one business problem, how success will be measured and which data exists to learn from, and check that the data can support the target.

02 **Prepare data and prove the model**
We clean and prepare the data, compare approaches from simple baselines to deep learning or an LLM, and test the strongest on held-back data it never saw in training. The results are reviewed with your team against the agreed target.

03 **Build and integrate**
We deploy to your cloud or servers and connect the model to the systems where people work, with access controls, logging and human review where a wrong prediction would be costly.

04

Launch and train your team

Your team learns to use the predictions on real cases, then we roll out in stages and stay close through the first weeks.

05

Monitor and retrain

Accuracy, drift and usage are tracked after launch, and models are retrained as your data and business change.

FAQ

Questions about AI & ML Development

How much does an AI or ML project cost?

It depends on the problem, how much data preparation is needed, whether an existing model can be adapted or a new one must be trained, the number of integrations and where it is hosted. A proof of concept costs a fraction of a full build. After a requirements review you get a fixed quote for the agreed scope.

How long does it take to build an AI solution?

A proof of concept typically takes 3 to 6 weeks. A full custom solution typically takes 2 to 6 months, and large enterprise platforms take longer. Most of the variation comes from data preparation and the number of systems the model has to connect to.

What if the model is not accurate enough?

We agree the accuracy target before any build and measure the proof of concept against it on data the model has never seen. If it falls short, we explain why, whether that is too little data, noisy labels or a problem framed too broadly, and what would close the gap: more or better data, a narrower scope, or a person reviewing the uncertain cases. If none of those is worth it, we recommend stopping, and you have spent only the proof-of-concept budget.

Do we need our own data scientists?

No. We bring the machine learning engineers and data scientists. From your side we need someone who knows the process and can make decisions, and someone in IT who can give access to systems and data. We document everything and train your team, so an in-house team can take over later if you build one.

Which technologies do you use?

Python with scikit-learn, PyTorch and TensorFlow for custom models; GPT, Claude, Gemini, Llama and Mistral for language tasks; MLflow or Kubeflow, Docker and Kubernetes for deployment; and AWS, Azure, Google Cloud or your own servers for hosting. We pick per project by testing candidates on your data, not by habit.

What data do we need to get started?

It depends on the use case. It can be database records, documents, images, video, sensor readings or public datasets. You do not need perfectly clean data on day one: preparing and cleaning it is part of our work, and the feasibility review tells you whether there is enough to learn from.

Should we build a custom model or use an off-the-shelf AI tool?

If a ready-made tool already handles your task well, use it. Custom work makes sense when your data, process or accuracy needs are specific to your business, or when the AI has to sit inside your own systems. We often combine both, for example a hosted LLM with your own retrieval and rules.

Can you integrate AI with our existing systems?

Yes. We connect models to ERP, CRM, accounting, e-commerce, mobile and custom systems through their APIs, webhooks or data exports, and return results into the screens your team already uses.

Who owns the code and the trained models?

You do. On completion and final payment, the code, prompts, trained or fine-tuned models and documentation we build for you are handed over, as set out in the contract. Third-party and open-source models stay under their own licences.

What support do you provide after launch?

Monitoring, performance tuning, model retraining, security updates, bug fixes, infrastructure management and new features, under a support plan agreed at handover.

NEXT STEP

Start with one problem worth solving

Book a free 30-minute call. Tell us the process or decision you want to improve and we will tell you honestly whether AI fits, what data it needs and what a proof of concept would involve.

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

WhatsApp or call: +91-9025555584

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