

Machine Learning Development Services for Prediction and Automation

Custom machine learning models for forecasting, risk, recommendations and inspection, trained on your own data and evaluated before they reach production.

OVERVIEW

Machine Learning Development

Machine learning finds patterns in your historical data and uses them to predict what comes next: which customers will leave, how much stock to order, which transaction looks wrong, which part on the line is faulty. We design, train and deploy those models, and connect them to the systems where the decisions are made.

Not every problem needs a model built from scratch. Depending on the task and the data you hold, we train a model on your records, fine-tune a pre-trained one such as BERT or Llama, or combine an existing model with your business rules. When data is scarce we use augmentation, transfer learning or synthetic data, and say so when none of these will be enough.

Every model is judged on data it has not seen, against a measure you agree with us before training starts. After launch we watch for drift, the slow loss of accuracy as customers, prices and processes change, and retrain on fresh data before the predictions stop being useful.

BUILT FOR

- Retail and e-commerce businesses
- Banks, lenders and insurers
- Manufacturers and logistics firms
- Healthcare and fitness providers
- Product teams adding predictions to software

FEATURES

What you get with Machine Learning Development

01

ML consulting

An assessment of your infrastructure, data and goals, with a plan covering budget, timeline, security, data confidentiality and compliance.

02

Forecasting and predictive analytics

Demand, sales, cash flow and inventory forecasts that account for seasons, promotions and regional patterns in your history.

03

Fraud and anomaly detection

Models that score transactions, claims or log events for risk and send unusual cases to a reviewer instead of blocking them blindly.

04

Classification and scoring

Churn prediction, lead scoring, credit scoring and ticket routing, with an explanation of which factors drove each score.

05

Recommendation systems

Product, content and next-best-action recommendations from purchase and browsing behaviour, tested against the approach you use today.

06

Vision and quality inspection

Image and video models that spot defects, count items or read labels, for manufacturing, retail and logistics operations.

- 07 NLP for text and voice**
Models that classify, extract and summarise written text and transcribed speech, from support tickets to recorded calls.
- 08 Data preparation**
Collection, cleaning, labelling and feature engineering for structured, semi-structured and unstructured data, including IoT and sensor feeds.
- 09 Deployment and monitoring**
Models served through APIs or scheduled batch jobs, with version tracking, drift alerts and planned retraining.

WHY IT MATTERS

The difference it makes to your business

- 01 Predictions from your own history**
Models learn from your customers, products and seasons rather than generic industry averages, so forecasts reflect how your business behaves.
- 02 Reviewers focus on the hard cases**
Scores and flags let staff spend their time on uncertain cases, while clear-cut ones follow a standard path.
- 03 Results you can explain**
We report which inputs drive each prediction, so managers and auditors can see why a model reached its answer.
- 04 Accuracy that holds up over time**
Drift monitoring and retraining stop performance from fading quietly as your market changes.

HOW WE DELIVER

How we deliver Machine Learning Development

- 01 Assess**
We review your data landscape, systems and goals, and agree the decision the model will support and how its success will be measured.
- 02 Prepare the data**
Data is gathered, cleaned and joined, labels are checked and features are built. This step usually takes the most time on any ML project.
- 03 Train and compare**
We train several approaches, starting from a simple baseline, and keep the one that beats it by a margin worth the extra complexity.
- 04 Validate and deploy**
The chosen model is tested on held-out data, reviewed with your team and put into your workflow through an API or a scheduled job.

05 Monitor and retrain

We track accuracy and drift in production and retrain on new data on a schedule, or sooner when performance drops.

FAQ

Questions about Machine Learning Development

What kind of data does machine learning need?

Structured data such as transactions and customer records, semi-structured data such as logs and JSON, and unstructured data such as text, audio and images can all be used. What matters most is that the data covers the outcome you want to predict, with enough examples of each case, including the rare ones.

What if we do not have much data?

There are options. A pre-trained model can be fine-tuned on a small dataset, existing examples can be augmented, and synthetic data can fill some gaps. We test whether these get close enough to the target, and if they do not, we recommend collecting more data before building.

Do you always train a model from scratch?

No. Training from scratch is the exception. We usually start from a proven algorithm or a pre-trained model such as BERT or Llama and adapt it to your data, which is faster and cheaper and needs fewer examples.

How long does a machine learning project take?

Simple projects typically take 4 to 6 weeks, medium-complexity projects 2 to 3 months and large enterprise projects 4 to 6 months or more. We work iteratively, so you see early results on your data well before the end.

What affects the cost?

The main drivers are the amount of data cleaning and labelling, whether a pre-trained model can be adapted or a custom one is needed, the accuracy required, and the integration and hosting work. You get a written proposal with a fixed price for the agreed scope.

How do you decide whether a model is good enough?

Before training we agree a measure that fits the problem, such as forecast error, precision and recall for fraud, or the share of defects caught, and a baseline to beat, usually your current method or a simple rule. The model goes live only if it beats that baseline on data it has never seen.

Can the model's predictions be explained?

Yes, to a useful degree. For most business models we report the factors that pushed each prediction up or down. Where explanations are a legal or audit requirement, we may choose a simpler, more transparent model over a slightly more accurate one.

NEXT STEP

Find out what your data can predict

Tell us the decision you want to improve and describe the data you hold. We will say whether a model is likely to help and how we would test it before you commit.

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

WhatsApp or call: [+91-9025555584](tel:+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>