Appsin.Ai

SERVICE

Recommendation Engine Development for Products, Content and Services

Personalised product and content recommendations built on your own behaviour data, with business rules your team controls and A/B tests that show what works.

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

OVERVIEW

Recommendation Engine Development

We build recommendation engines that show each customer the products, content or services they are most likely to want. Using your own data, such as views, searches, carts, purchases and ratings, models rank items for each person and page, and serve them through a fast API to your website, app, emails and in-store tools.

We use the approach your data supports: content-based filtering for new catalogues, collaborative filtering once there is enough behaviour history, and hybrid models that combine both with context such as time, location and device. Business rules for stock, margin, exclusions and promotions sit on top, under your team's control.

A recommender is only worth what it adds, so we test it. Every engine ships with tracking and A/B testing against your current approach, whether that is best-sellers or manual picks, and we will tell you plainly when your data is too thin to beat simple rules yet.

BUILT FOR

- E-commerce and retail stores
- Media, streaming and publishing platforms
- Online learning platforms
- SaaS products with large catalogues or content libraries
- Travel and hospitality booking sites

FEATURES

What you get with Recommendation Engine Development

01 **Personalised home and category pages**
Each visitor sees products or content ranked for them on home, category and account pages, based on what they and similar users engaged with.

02 **Similar and complementary items**
Similar-item and often-bought-together suggestions for product pages and carts, based on item attributes and purchase patterns.

03 **Content and media recommendations**
Next article, video, course or lesson suggestions based on what each user reads, watches and finishes.

04 **Cold-start handling**
New users and new items get sensible suggestions from attributes, popularity and context until enough behaviour has built up.

05 **Business rules and merchandising controls**
Boost, bury, pin or exclude items, respect stock levels and margins, and run campaigns without retraining the model.

06 **Personalised search ranking**
Search results re-ranked for each user, so two customers typing the same query see what is most relevant to each of them.

07 **Email and push recommendations**
Personalised product blocks for newsletters, abandoned-cart and win-back messages, sent to your email or CRM platform.

08 **Real-time recommendation API**
An API that reacts to what a user does in the current session, with caching and a fallback list if a model is unavailable.

09 **A/B testing and reporting**
Experiment set-up, click and conversion tracking for each placement, and dashboards that compare recommendations with your current approach.

WHY IT MATTERS

The difference it makes to your business

01 **Customers find more of what they want**
Relevant suggestions cut down scrolling and searching, especially in large catalogues and content libraries.

02 **Measured, not assumed**
A/B tests show the effect of each placement on clicks, basket size or watch time before you roll it out to everyone.

03 **Merchandisers stay in charge**
Rules for stock, margin and campaigns let your team steer what is shown without needing a data scientist.

04 **Insight into your customers**
The same models show which items are bought together and which customers behave alike, which helps buying and marketing.

HOW WE DELIVER

How we deliver Recommendation Engine Development

01 **Audit your data**
We check which behaviour and catalogue data you collect, how clean it is and what tracking needs adding before any modelling.

02 **Agree placements and measures**
We choose where recommendations will appear and how success is measured, such as click-through or revenue per session.

03 **Build a baseline and models**
A simple baseline such as best-sellers is set up first, then models are trained and compared against it on historical data.

04 **Integrate and A/B test**
The engine is connected through an API or plugin and tested live against the baseline on a share of your traffic.

05 **Roll out and retrain**
Winning placements go to all users, models retrain on fresh data on a schedule, and performance is monitored.

FAQ

Questions about Recommendation Engine Development

How does a recommendation engine work?
It collects signals such as views, clicks, carts and purchases, turns them and your catalogue attributes into features, and trains models that score how likely each user is to want each item. An API serves the top items for each page in real time, and new behaviour feeds back into the next round of training.

What is the difference between content-based and collaborative filtering?
Content-based filtering recommends items similar to ones a user liked, using attributes such as category, brand or topic. Collaborative filtering recommends what similar users engaged with, without needing attributes. Hybrid systems combine both, which helps with new users and new items.

How much data do we need?
Collaborative filtering needs a good number of users who have each interacted with several items. Small or new catalogues can start with content-based and popularity models and add collaborative signals as data grows. We assess your data at the start and tell you which approach it can support.

Does it work with Shopify, WooCommerce or Magento?
Yes. We connect through each platform's APIs and a plugin or script on the storefront, and the same engine serves your mobile app and emails. Custom websites and apps call our recommendation API directly.

How long does it take?
Most recommendation engines take 3 to 6 months to build after a 3–6 week proof of concept, depending on how ready your tracking data is. One or two placements on clean data can be quicker; real-time systems across many placements and channels take longer.

How much does a recommendation engine cost?
Cost depends on the number of placements and channels, data volume and cleanliness, whether recommendations must update in real time, the integrations and the hosting. You get a fixed quote after a requirements review.

Should we build a custom engine or use an off-the-shelf tool?
Off-the-shelf plugins are a sensible start for a standard store. A custom engine makes sense when you have unusual products or content, several channels, your own business rules or data you do not want to send to a third party. With a custom build, the data stays yours, and once the project is paid for, so are the models we build.

NEXT STEP

See what personalisation could do on your site

Tell us how your catalogue and tracking are set up and we will tell you which recommendation approach your data supports and how we would test it.

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