



SERVICE GUIDE · 2026

Predictive AI


See what is coming early enough
to do something about it.

INSIDE THIS GUIDE

- 01 How it works
- 02 What it can actually do
- 03 Where it applies
- 04 What we build for you
- 05 Accuracy & governance
- 06 Hosting & deployment
- 07 Technical companion guide
- 08 Training & support
- 09 Common questions
- 10 Our delivery process

ENTRA AI

www.entra-ai.com
info@entra-ai.com
+44 2034 323178



WHAT IS PREDICTIVE AI

Turning history into foresight

Your organisation already knows a great deal about what has happened. What it usually cannot do is say what is *about* to happen — which customers are drifting away, which machine is heading for a failure, which orders will not arrive on time.

Predictive AI finds the patterns in your own history and uses them to put a number on what comes next. Not a crystal ball: a well-tested, measurable estimate, delivered early enough that someone can still do something about it.

THINK OF IT LIKE THIS

Think of your most experienced operator — the one who can look at a job and say 'that one will be trouble' before anyone else notices. Predictive AI is that instinct, learned from every record you hold rather than one career's worth, and applied to every case rather than the ones that happen to catch someone's eye.

01

How It Works

A plain-English overview

■ We Learn From Your History

We use records you already hold — sales, jobs, sensors, tickets — and find the patterns that reliably preceded the outcome you care about.

■ We Prove It on the Past

Before anything goes live, we test the model on history it has never seen and show you exactly how often it would have been right.

■ It Runs on What's Next

The tested model scores live cases, so your team sees the risk or the forecast while there is still time to act on it.

02

What It Can Actually Do

The questions it answers

■ Forecasting

Demand, revenue, headcount, stock, or call volume — projected with a stated confidence range rather than a single optimistic number.

■ Churn & Retention

Which customers are drifting, how early the signal appears, and which of them are worth intervening on.

■ Failure Prediction

Equipment and process failures flagged in advance, so maintenance happens on a plan instead of an emergency.

■ Risk Scoring

Which applications, claims, orders, or accounts carry most risk, so scrutiny goes where it is actually needed.

■ Anomaly Detection

Spotting the transactions, readings, or behaviours that do not fit — including the kinds no one thought to write a rule for.

■ Prioritisation & Ranking

Leads, cases, and interventions ordered by likely value, so limited capacity goes to the highest-return work.

03 · WHERE IT APPLIES

Real scenarios, real outcomes

A

Revenue & Customer

Acting before the customer decides for you

■ Churn Prediction

Accounts at risk surfaced early enough to act on, with the specific signals driving the score so your team knows what to address.

■ Lead & Opportunity Scoring

Ranking pipeline by genuine likelihood of closing rather than by who shouted most recently.

■ Demand Forecasting

Sharper forecasts by product, region, and season — the basis for stock, staffing, and cash planning that holds up.

■ Pricing & Elasticity

Understanding how demand actually responds to price, so changes are modelled before they are made.

B

Operations & Assets

Fixing things before they break the week

■ Predictive Maintenance

Sensor and service history turned into early warning, so equipment is serviced before it fails rather than after.

■ Delivery & SLA Risk

Flagging the orders and cases likely to miss their commitment, early enough to reprioritise.

■ Quality Prediction

Identifying which batches or jobs are trending toward a defect while there is still time to correct the process.

■ Capacity Planning

Projected workload by team, shift, and skill, so rotas are built on evidence rather than on last year's shape.

C

Risk, Finance & Compliance

Putting scrutiny where it earns its keep

■ Fraud & Anomaly Detection

Unusual patterns surfaced across high volumes of transactions, including novel behaviour that fixed rules would miss.

■ Cash Flow Forecasting

Forward visibility on receipts and payments, with confidence ranges your finance team can plan against.

■ Credit & Application Risk

Consistent, explainable risk scoring, with the reasoning available for every individual decision.

■ Compliance Monitoring

Continuous checking across all records rather than periodic sampling, with exceptions raised as they appear.

04

What We Build For You

The parts of a working system

■ Data Pipelines

Reliable, repeatable preparation of the data the model needs, so predictions do not quietly degrade when a source changes.

■ Scoring Service

The live component that produces predictions, whether in batches overnight or on demand through an API.

■ Explanation Layer

For every prediction, the factors that drove it. Essential for trust, and often for regulation.

■ Trained & Tested Models

Models built for your data and validated properly against held-out history — with the numbers shown to you, not summarised away.

■ Dashboards & Alerts

Predictions delivered where the decision is actually made — your existing tools, a dashboard, or a direct alert.

■ Monitoring & Retraining

Automatic tracking of accuracy over time, with retraining when the world moves on from the data we learned from.

05

Accuracy & Governance

Real numbers, openly reported

■ Measured, Not Claimed

We report accuracy on data the model has never seen, using measures that suit your problem — and we show the failures too.

■ Bias Testing

We test performance across the groups that matter and report differences rather than hoping nobody looks.

■ Human in the Loop

Predictions inform decisions; they do not silently make them. Where stakes are high, a person stays in the chain.

■ Explainability

Every prediction can be traced to the factors behind it, so decisions can be defended to a customer, an auditor, or a regulator.

■ Drift Monitoring

Model accuracy is tracked continuously. When reality shifts, you find out from us, not from a bad quarter.

■ Data Privacy & Security

Your data stays yours, within your security policies and your residency requirements.

06

Hosting & Deployment

Wherever your policies require

■ Your Cloud

Deployed into your own AWS, Azure, or Google Cloud environment under your existing controls.

■ On-Premise

Fully self-hosted where regulation or policy demands it. Nothing leaves your network.

■ Managed by Us

We host, monitor, and retrain it, and you receive a working service with agreed support.

GO DEEPER · TECHNICAL COMPANION GUIDE**07 · Machine learning & deep learning — the detail underneath**

Predictive AI is the outcome; machine learning and deep learning are how it is built. Our technical companion guide covers model selection, feature engineering, validation strategy, neural architectures, and the MLOps practice that keeps models healthy in production. It is written for data and engineering teams who want to interrogate the approach rather than take it on trust.

www.entra-ai.com/assets/docs/Entra_AI_Machine_Learning_Service_Guide.pdf

08**Training & Support**

Making sure the numbers get used

A prediction nobody acts on is an expensive statistic. Most of the value comes from the change in how decisions get made, so we spend real time with the people who will use the output.

■ Team Onboarding

Practical sessions on how to read a score, what the confidence range means, and — importantly — when to overrule it.

■ Documentation

Clear write-ups of the data used, the assumptions made, and the known limitations of every model we hand over.

■ Leadership Briefings

What the model does and does not cover, how accuracy is measured, and what good looks like in the numbers.

■ Ongoing Support

Reviews as your data and business change, plus retraining and extension when you are ready.

09

Common Questions

The things clients ask us first

■ **How much data do we need?**

Less than most people expect, but it must cover enough real examples of the outcome. We tell you at assessment, before you commit.

■ **Can we see why it decided that?**

Yes. Every prediction comes with the contributing factors — we do not hand over systems nobody can explain.

■ **Does it degrade over time?**

All models do as the world changes. That is why monitoring and retraining are part of the service, not an afterthought.

■ **What if it is not accurate enough?**

We find that out in validation, before go-live and before the invoice. Some problems are not predictable, and we will say so.

■ **Will it replace our experts?**

No. It handles the volume and the consistency so your experts spend their judgement on the cases that need it.

■ **How long until it is live?**

It depends on how ready your data is. We assess that before anything is committed, so the timeline you get is a realistic one.

10

Our Delivery Process

From first conversation to a live model

■ **01 — Discovery Call**

We learn what decision you are trying to improve and what it is worth getting right.

■ **04 — Shadow Run**

It predicts alongside your current process so you can compare before you rely on it.

■ **02 — Data Assessment**

We check whether your data can actually support the question before anyone commits budget.

■ **05 — Deploy & Integrate**


We go live and put predictions into the tools where decisions are actually made.

■ **03 — Build & Validate**

We build the model and test it against held-out history, reporting the real numbers.

■ **06 — Monitor & Retrain**

We track accuracy, watch for drift, and retrain as your business moves.



NEXT STEP

Ready to stop reacting and start anticipating?

Tell us about a decision your team makes on instinct or on yesterday's numbers. We will tell you whether your data can do better, and say so plainly if it cannot.

GET IN TOUCH

WEB

www.entra-ai.com

EMAIL

info@entra-ai.com

PHONE

+44 2034 323178