eomer

Enterprise AI: predict, explain & classify

The fastest path from data to deployed AI

Why enterprise AI projects fail today

Costly infrastructure, scarce ML talent, and slow deployment block adoption

70%

of AI project budget
goes to infrastructure

$500k+

typical setup cost
for in-house AI projects

12+ months

average time-to-value
for enterprise ML

What we hear from enterprise teams

“We have the data, but not the ML team to use it.“

“We know we're losing customers but can't predict who's next.“

“We can't justify the infrastructure cost for one use case.“

From data to decisions in days, not months

Enterprise AI without infrastructure, hiring, or long deployment cycles

We **remove** the 3 biggest **barriers** to **enterprise AI**

No infrastructure setup

- **No data pipelines** or ML infrastructure required
- Connect via **CSV, database, or API**
- Fully managed **compute and scaling** or **on-premise**

No ML team required

- No model tuning or feature engineering
- Platform selects and **optimises models automatically**
- Built for business users and data teams

Production-ready in days

- First predictions **within hours**
- Deploy via **API** or dashboard instantly
- Continuous monitoring and retraining included

AI that works directly for your data

From raw data to production-grade AI in days, not months - without needing in-house AI expertise

1

Connect your data

- **Upload** time series, tabular, and alternative data **via API or database**
- We handle ingestion at **terabyte scale**

Day 1

2

We train & optimise

- Platform **selects** the **best model architecture** & adjusts to your data
- **Tunes automatically** using foundation model backbones

Days 2-5

3

Deploy & forecast

- **Production-ready** predictions via API & **on-premise deployment**
- Continuous retraining, monitoring, and **fine-tuning for** your **specific use cases**

Day 7+

Click me: eomer platform intro video (3min)

Built for your industry

eomer serves any domain with structured tabular data

Finance & Trading

- Risk scoring
- Alpha signals
- Portfolio optimisation

Energy & Utilities

- Load forecasting
- Price prediction
- Grid optimisation

Supply Chain & Logistics

- Demand planning
- Inventory optimisation
- ETA prediction

Manufacturing & Industrial

- Predictive maintenance
- Quality control
- Yield forecasting

Predict demand, revenue and risk

Used in ERP-driven environments to bring clarity to inventory, pricing, planning

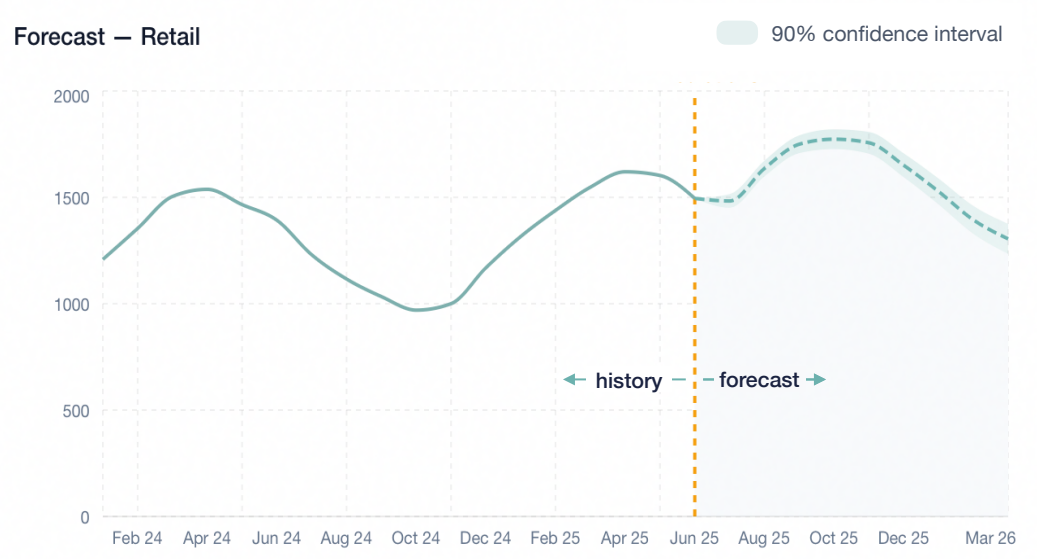

Works out of the box on your data

Pre-trained on billions of time points for superior zero-shot & fine-tuned performance

Confidence interval not just points

Full prediction distributions with calibrated confidence intervals, not just point estimates

Automatic model selection

Platform benchmarks multiple architectures and selects the best for your data

Up to 40%
MAPE improvement

< 1 hour
to first forecast

1B+
pre-training data points

eomer

Explain what drives your business

Predict continuous outcomes from mixed tabular and temporal features

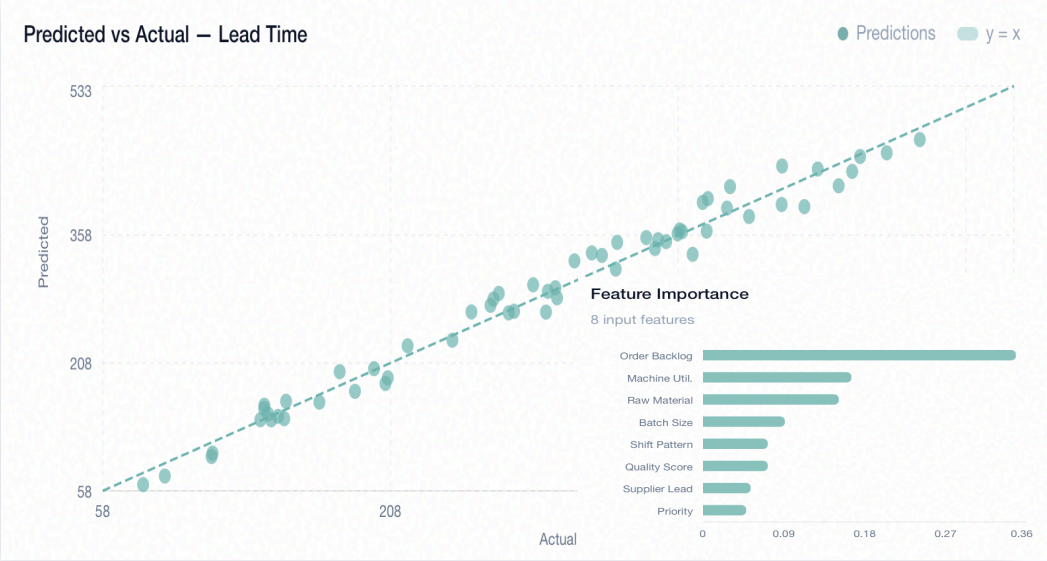

Revenue optimisation

Identify the key drivers behind your business outcomes

Risk scoring

Estimate continuous risk metrics combining static attributes with time-varying indicators

Yield prediction

Forecast manufacturing yields or crop outputs from sensor streams and environmental data

$R^2 > 0.95$
on benchmark datasets

Automatic
feature engineering

Tabular + Temporal
data fusion

eomer

Identify churn, fraud and anomalies

Categorise events, detect anomalies, and trigger alerts from streaming data

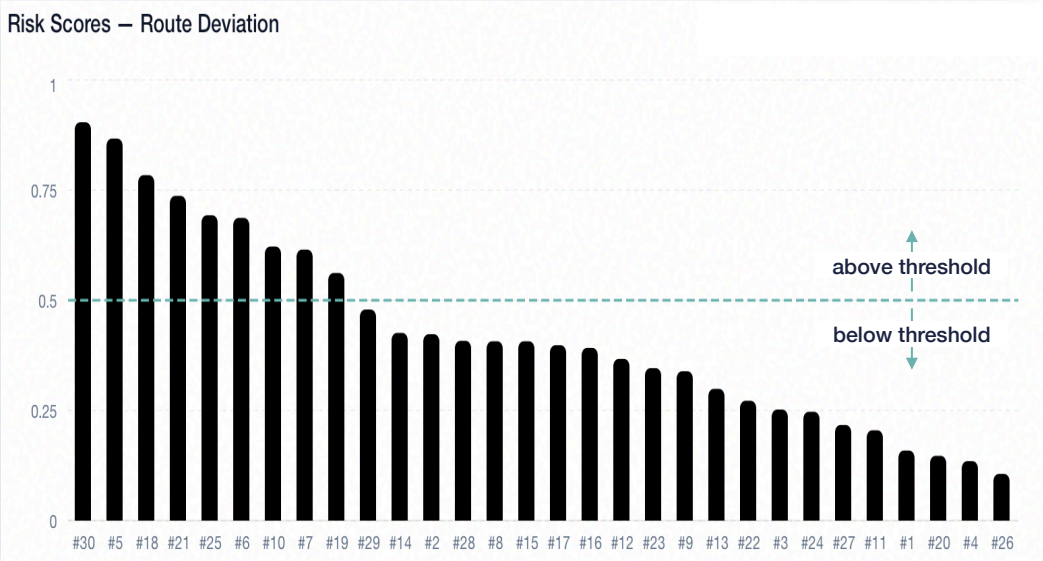

Anomaly & fraud detection

Identify unusual patterns and **flag suspicious transactions** in real time by combining temporal patterns with entity features

Churn prediction

Classify at-risk customers using behavioural sequences and engagement time series

Predictive alerting

Trigger automated alerts when classification confidence exceeds configurable thresholds

> 94%
AUC across benchmarks

Real-time
streaming inference

Configurable
alert thresholds

eomer

Plug into your existing systems in days

Deploy in your environment with the tools you already use

REST API, SDK and MCP

Python and REST APIs for seamless integration, MCP for AI agent readiness

Data connectors

S3, GCS, Snowflake, BigQuery, Kafka, and more

Enterprise security

SOC2 Type II, data encryption at rest and in transit

Auto-scaling compute

Elastic GPU/CPU scaling for training and inference

Typical deployment timeline

Day 1	Connect data sources & define schema
Day 2-3	Automated model training & benchmarking
Day 4-5	Fine-tuning for your specific use cases
Day 6-7	API integration & production deployment
Ongoing	Continuous monitoring & retraining

Why eomer

Industry-proven performance, research-driven technology, enterprise-ready execution

Trusted in production

- **Pilot projects** with **enterprise clients** across **energy & logistics**
- **Benchmarked** on state-of-the-art public & academic datasets[1)]

10x faster deployment

- From data connection **to live AI production in a few days**
- **No ML engineers needed** on your side - **we handle** the **full stack**

The Team

- Built by **AI leaders** and **quant researchers** with **experience** across **QRT, BlackRock & NUS**
- **Deep expertise** across **finance, energy** & **industrial applications**

Enterprise security

- **SOC2 Type II compliant - Data encrypted** at rest and in transit
- VPC peering and **on-premise deployment** options **available**

Note 1) Data sets: Consumer Retail: Corporación Favorita Grocery; Energy: BDG2, ERCOT, ENTSO-E; Academic: M4 & M5

You focus on insights.
We handle all the rest.

- AI as a Service
- Enterprise-grade scale & reliability
- No ML infrastructure or experience required

Let's discuss your use case.

Book a demo

development@eomer.ai

eomer