



Custom AI models & synthetic data

A company briefing with evidence from Signal Lab

artificialmodelling.in

Company problems we can scope

Workflow	Available starting point	Project output
Operations planning	Historical orders or operational logs	A forecast or ranked exception list
Document processing	Documents with reviewed examples	Extraction or classification with error analysis
Equipment monitoring	Telemetry and event labels	Failure scores and measured alert counts
Software and model testing	A schema and permitted source examples	Synthetic records for specified test cases

Example project scopes. Feasibility depends on the data and the decision the output supports.

Our services

Custom AI models

Baseline and model comparison

Training and error analysis

Batch inference or an API

Model artifacts and documentation

Synthetic data

Schema and generation recipe

Coverage of specified rare cases

Checks for structure and task utility

Dataset and source documentation

Each engagement has agreed deliverables and an evaluation plan.

Signal Lab

Synthetic telemetry and a trained failure classifier

A complete workflow
your team can inspect


Generation controls

Shared-test evaluation


Inference and file exports

Public synthetic benchmark. Engineering demonstrator.

Artificial Modelling



Signal Lab architecture



Validation and test labels enter evaluation. They never fit the generator or classifier.

Benchmark design


Partition	Rows	Failure labels	Purpose
Training	6,000	203	Fit generation and the source model
Validation	2,000	68	Select the starting thresholds
Test	2,000	68	Compare both models on shared rows

Default generation recipe

6,000 rows / 15% failure labels / noise 0.15 / seed 42

AI4I 2020 is synthetic. A random split measures benchmark classification, not future failures at a plant.

Generated training retains similar AP




Same model architecture
Same 2,000 test records

AP is similar in the default run.
ROC AUC is lower with
generated training.

The small AP difference does not
establish superiority.

68 failures in the test set. The no-skill AP reference is 0.034. Public synthetic benchmark.

The threshold changes the review workload



35 to 95

records flagged for review

60 additional reviews
15 additional failure labels

A pilot must agree which errors and review costs matter.

Descriptive threshold exploration on one benchmark test set. Classification scores are not calibrated probabilities.

A manufacturing pilot

Compare current rules, a model and the effect of synthetic data.

Company inputs

Timestamped telemetry and machine IDs

Failure and maintenance event logs

A prediction horizon and operator action

Costs of reviews and missed events

Acceptance measures

Recall at an agreed alert budget

False alarms per machine and time period

Useful lead time before an event

Errors by machine and operating regime

Proposed pilot. Company data must establish operational performance and any financial benefit.

The handover

Deliverable	What your team receives
Model or dataset	Versioned artifacts in agreed formats
Evaluation	Baseline comparison, operating points and error analysis
Reproduction	Training or generation code and run instructions, where scoped
Integration and ownership	Inference contract, source licences and defined responsibilities

The proposal defines the formats, access, ownership and support responsibilities.

A scoped feasibility study

01	Scope	The decision, intended users and success criteria
02	Data review	A permitted dataset and a representative evaluation plan
03	Comparison	A baseline and the agreed model or generation experiment
04	Decision	Continue, refine the data, or stop with a documented reason

Fees and timeline follow the agreed scope. A build or support engagement can follow the study.

An initial project discussion

A first brief can answer four questions

Which decision or workflow needs to improve?

What data and reviewed examples are available?

Who will use the output?

How will you judge a useful result?

aditya@artificialmodelling.in

artificialmodelling.in

[Signal Lab case study](#)

Generation and model specification

Component	Implementation
Generator	Six Gaussian copulas, conditional on failure label and product type
Numeric representation	Air temperature, temperature gap, speed, torque and wear
Classifier	Class-balanced logistic regression with quadratic interactions
Features and parameters	22 features and 23 parameters including the intercept
Training	Full-batch Adam, 160 epochs, learning rate 0.035, L2 0.005

The generator approximates relationships. Structural checks do not establish physical label correctness or privacy.


Quality checks and sensitivity

0

Invalid structural records
Exact source input matches
Repeated generated inputs

Zero exact copies does not
establish privacy or physical label
correctness.

AP across five generation seeds



Fixed recipe and test partition. AP ranges from 0.489 to 0.520.

Evidence and validation boundaries

Public synthetic benchmark with 68 positive test examples

Random partitions rather than a future-time or unseen-plant test

Exploratory recipe changes reuse the same test partition

Company pilots must establish operational value and privacy requirements

[Case study and downloadable evidence](#)

[AI4I 2020 dataset, UCI DOI 10.24432/C5HS5C](#)

CC BY 4.0. Associated introductory paper: S. Matzka (2020). Snapshot: 11 October 2026.