Industrial AI & ML Glossary

Twenty terms you will meet in a vendor conversation, defined in plant terms — plus a selection tree for choosing a model family. Each entry carries the caveat that usually matters more than the definition.

Model selection tree

Work down the list. The first question you cannot answer is your real project.

Ask	Then
Do you have labelled examples of the event?	No → anomaly detection or clustering. Yes → continue.
Is the answer a category or a number?	Category → classification. Number → continue.
Is the number a time-until-event?	Yes → survival / RUL model. No → continue.
Does the value depend mainly on its own history?	Yes → time-series forecasting. No → regression.
Fewer than ~50 examples of the rare class?	Stop. Instrument and collect first — no model family rescues this.
Is history shorter than one full seasonal cycle?	Prefer a physics-informed or rule-based baseline until you have the cycle.

Glossary

Anomaly detection
Learns what normal looks like and flags departures from it. Needs no failure labels, which is why it is usually the first thing that works in a plant.
Use when you have history but no recorded failures. It tells you something changed, not what will break.

Classification
Assigns a label from a fixed set — pass/fail, which fault, which product.
Needs labelled examples of every class. A class with three examples will not be learned.

Regression
Predicts a number — remaining useful life, expected yield, energy for the next run.
Check the error distribution, not just the average error. Industrial cost is rarely symmetric.

Time-series forecasting
Predicts a value forward in time using its own history and related signals.
Breaks quietly when the sampling interval is irregular. Fix the interval before the model.

Survival / RUL model
Estimates time remaining before an event, handling machines that have not failed yet.
The right shape for maintenance planning; a plain regression on age is not.

Clustering
Groups similar behaviour without labels.
Useful for finding operating modes you did not know you had — often the first honest picture of a process.

Feature
A derived quantity the model actually reads: a rolling mean, a cycle-time variance, a ratio.
Most industrial model quality comes from features, not model choice.

Label
The answer the model is trained against.
In a plant, the label is usually the weakest link: entered late, by hand, under time pressure.

Training / validation / hold-out
Three separated slices of data: fit, tune, and one untouched slice used once.
If a vendor reports one accuracy figure without saying which slice, it is a training figure.

Overfitting
The model memorised the training period rather than learning the process.
Symptom: excellent on history, ordinary in production. Test on a different line, not a different week.

Data leakage
A feature quietly contains the answer — typically a value recorded after the event.
The most common cause of an implausibly good pilot result.

Concept drift
The process changes, so a model trained on the old behaviour degrades.
Assume it. Ask who is told when accuracy falls and what happens next.

Class imbalance
The event you care about is rare, so predicting 'no failure' always scores well.
Judge on precision and recall for the rare class, never on accuracy.

Precision / recall
Precision: of the alarms raised, how many were real. Recall: of the real events, how many were caught.
Pick which one matters before the vendor picks it for you. Nuisance alarms kill adoption faster than misses.

Explainability
Showing which inputs drove a given output.
On a shop floor this is not a nicety — an operator who cannot see why will stop acting on it.

Edge vs cloud inference
Where the model runs: on the line, or in a data centre.
Decided by latency tolerance and network reliability, not by preference.

Digital twin
A model of the asset used to simulate behaviour under conditions you have not run.
Valuable and expensive. Ask what it is calibrated against and how often.

Physics-informed model
A statistical model constrained by known physical relationships.
Needs less data than a pure data-driven model, which matters when history is short.

MLOps
The practice of deploying, monitoring and retraining models in production.
The gap between a pilot and a dependable line is almost entirely this.

Golden batch
A reference run the process is compared against.
Useful shorthand; risky when the golden batch is simply the best one anyone remembers.