

Machine learning in production, without a data scientist

How to turn a simple export of your machine data into quantified gains — machine-stoppage prediction and material yield — in a few weeks.

273 h	+9.5 pts	977 → 10
of machine downtime anticipated on an aerospace line	of material yield on a rubber process	scrap parts per batch after optimization

Figures from real client deployments. Detailed at the end of this guide.

What you will learn

- Why most factory AI projects stall — and the shortcut that unblocks them.
- The two use cases that pay back fastest, and how to know which to target first.
- The 5-step method, from CSV export to a gain measured in euros.
- How to run a blind two-week test, with no risk and without tying up your teams.

1. Why factory AI projects fail

Most industrial AI initiatives don't die from a lack of data — they die from too much ambition. A “data program” is launched: hire a data scientist, install a platform, wire up sensors, build a data lake. Twelve months later the budget is gone and nothing has reached production.

The shortcut is counter-intuitive: start with the business question, not the infrastructure. Your machines already produce usable data — in their controllers, historians and quality exports. All you need is a model that answers one precise, quantifiable question: “when will this line stop?”, “where is the material going?”.

The right unit of measure is not model accuracy. It is the euro. A successful project is judged on one thing: downtime hours avoided, yield points gained, scrap reduced.

2. The two use cases that pay back immediately

On the floor, two families of use case stand out because they plug into data you already have and pay back immediately.

Machine-stoppage prediction (anomaly detection). The model learns the line's “normal” behaviour and flags drift before the breakdown. The gain shows up as downtime hours avoided and interventions planned rather than suffered.

Material-yield optimization (process). The model links settings and production conditions to the scrap rate, then recommends the settings that maximise yield. The gain shows up as yield points and material saved.

Which first? Target the use case whose current cost is most visible in your numbers: if unplanned stoppages dominate your losses, start with prediction. If it's material scrap, start with the process.

3. The 5-step method

5. Two real cases

Aerospace — machine-stoppage prediction. On a machining line, the model was trained on the machine history then evaluated blind. Result: **273 hours of downtime anticipated**, i.e. **45%** of the stoppages over the period — detected before they idled the line.

273 h	45%	0
downtime anticipated	of stoppages caught early	new sensor installed

Rubber — material-yield optimization. By linking settings and production conditions to the scrap rate, the model recommended new settings. Yield rose from **68.6% to 78.1% (+9.5 points)**, and scrap per batch fell from **977 to 10**.

+9.5 pts	78.1%	977 → 10
material yield	yield reached	scrap per batch

What now?

Pick one line and one quantifiable problem, gather an export of your data, and run a blind test. That is exactly what Jemba sets up with production teams — no data scientist, from your existing data.

Take action: book a demo at jemba.ai/reserver-demo or test on your own data via jemba.ai/essai-gratuit.