

Build a predictive maintenance pilot around a decision

Connect sensor behavior to operating context, maintenance events and a false-alarm budget.

In short

- A pilot should first write down the maintenance decision the warning will support, such as scheduling an inspection in the next shift.
- Try a simple robust statistical baseline first and compare any candidate model against it over the same evaluation period.
- Starting in shadow mode records warnings without automatic intervention, and false alerts per asset per day, missed events and lead time are reported together.

01 / Specify the decision

Write down what the maintenance team will do with a warning, such as scheduling an inspection next shift. The warning lead time must allow that action. A model that detects the instant of failure may be too late for planning. Store speed, load and production recipe alongside sensor values.

02 / Link labels to work orders

Align sensor timestamps with work orders, shutdowns and replaced components. Not every vibration increase is a fault: startup and load changes also matter. Unknown labels must not silently become normal labels. Document units, sampling intervals and sensor replacements. Filling missing readings with zero can manufacture a healthy-looking interval.

03 / Establish a simple baseline

Try an operating-regime threshold or a robust statistical baseline first. Kutup Lab demonstrates median and MAD. Compare a candidate such as Isolation Forest over the same evaluation period. An anomaly score is not a failure probability; without verified event labels you cannot calculate fault precision or recall.

04 / Start in shadow mode

Initially record warnings without allowing automatic intervention. Collect the maintenance team's outcome for every alert. Report false alerts per asset per day, missed events and warning lead time together. Success on one pump does not demonstrate generalization to a different pump or operating regime.

Try it

Compare normal, spike and drift scenarios in the lab. Explain why one threshold gives different outcomes.

References and online guide

Open the online guide

scikit-learn: novelty and outlier detection

Proposed pilot steps do not claim field accuracy. Synthetic examples are not measurements from a production site.