odooAI

The Odoo AI Readiness Checklist: 12 Questions Before You Automate

A practical pre-flight assessment for Odoo teams. Score your data quality, process clarity, and EU compliance posture before you spend a euro on AI automation — and find out which workflows are actually worth automating first.

Why Readiness Beats Enthusiasm

Most failed Odoo AI projects do not fail because the model is bad. They fail because the automation is bolted onto messy data and undefined processes. An AI that drafts purchase orders from a vendor list with 30% duplicate contacts will confidently produce 30% wrong orders — faster than a human ever could. Speed amplifies whatever you point it at, including your mistakes.

This checklist exists to stop you from automating chaos. Before you connect an LLM to your `sale.order`, `account.move`, or `stock.picking` models, you need honest answers to 12 questions across three dimensions: is the data clean enough to trust, is the process defined enough to encode, and is the use case legal enough to deploy in the EU.

Work through it as a self-assessment. Score each question 0 (not true), 1 (partially true), or 2 (fully true). A total below 12 out of 24 means you have foundation work to do first — and that foundation work is usually cheaper and higher-leverage than the AI itself.

Part 1 — Data Quality (Questions 1-4)

Q1. Is your master data deduplicated and normalized? Check `res.partner`, `product.template`, and `product.product`. Run a quick count: `SELECT name, COUNT(*) FROM res_partner GROUP BY name HAVING COUNT(*) > 1;`. If you have 'Acme Ltd', 'Acme Ltd.', and 'ACME LIMITED' as three partners, an AI agent matching invoices to vendors will scatter transactions across all three. Score 2 only if duplicates are under 2% and naming follows a convention.

Q2. Are your key fields actually populated? AI reasons over fields, not vibes. If `product.template.standard_price`, `categ_id`, or `default_code` (internal reference) are blank on half your catalog, any margin analysis or reorder suggestion is guesswork. Pull a fill-rate report per critical field before trusting automation on it.

Q3. Is historical data consistent over time? If your team changed how they tag customers or categorize expenses 18 months ago, the AI will learn two contradictory patterns. Decide on a clean lookback window (e.g. 'only data after Jan 2025 is reliable') and document it.

Q4. Can you export a clean sample to inspect? If you cannot produce a 200-row CSV that you would be comfortable showing a customer, the AI cannot work with it either. This is your fastest reality check.

Part 2 — Process Clarity (Questions 5-8)

Q5. Can you write the process as deterministic rules? If you can describe a workflow as 'when X, do Y, unless Z', it is automatable. If your answer is 'it depends, the sales manager just knows', you have tacit knowledge that must be extracted into rules before any AI can encode it. Example that passes: 'Any quotation over EUR 5,000 needs manager approval before confirmation.'

Q6. Is the process stable, or does it change monthly? Automating a workflow that your team is still actively redesigning means you will rebuild the automation every few weeks — negative leverage. Automate the boring, settled processes first; leave the experimental ones manual.

Q7. Are exceptions documented? The 80% happy path is easy. The 20% edge cases (partial deliveries, credit notes, backorders, multi-currency rounding) are where automation silently breaks. List your top 5 exceptions per process. If you cannot, you do not understand the process well enough to automate it yet.

Q8. Is there a clear human checkpoint? Decide upfront where a human reviews AI output. Good pattern: AI drafts, human confirms. Dangerous pattern: AI confirms and posts to the ledger autonomously on day one. Score 2 only if you have defined who approves what.

Part 3 — EU Data Protection (Questions 9-12)

Q9. Does the automation touch personal data? Customer names, emails, phone numbers in `res.partner`, and any employee data in `hr.employee` are personal data under GDPR. If your AI workflow reads or writes these, you need a lawful basis (Art. 6) documented. Internal product catalogs and stock levels are usually low-risk; CRM and HR are high-risk.

Q10. Where does the data physically go? If you send Odoo records to a US-based LLM API, that is a third-country transfer. Confirm your provider offers EU data residency or operates under an adequacy decision / Standard Contractual Clauses. For sensitive workloads, consider EU-hosted or self-hosted models.

Q11. Is the provider use covered by a Data Processing Agreement (DPA)? You need a signed DPA with any AI vendor that processes personal data on your behalf, plus confirmation they do not train their base models on your prompts. Check the vendor's terms explicitly — 'zero data retention' and 'no training' are different promises.

Q12. Can you explain and audit an automated decision? Under GDPR Art. 22, individuals have rights regarding decisions made solely by automated means. Log every AI action with input, output, timestamp, and the model version. If a customer asks 'why was my order rejected', you must be able to answer. Build the audit log before you go live, not after.

The Decision Matrix: Value vs Effort

Once you have candidate workflows, do not automate the most exciting one — automate the one with the best value-to-effort ratio. Plot each candidate on a 2x2 grid: vertical axis = business value (time saved per

month x error cost reduced), horizontal axis = implementation effort (data cleanup + integration + review setup).

Top-left (High value, Low effort) — DO NOW. Example: auto-drafting vendor bill lines from incoming PDF invoices using OCR + LLM extraction, with a human confirm step. High volume, repetitive, clean enough data, low legal risk if vendors are companies not individuals.

Top-right (High value, High effort) — PLAN. Example: AI demand forecasting feeding reordering rules. Valuable, but requires clean historical sales data and exception handling. Schedule it after a data cleanup sprint.

Bottom-left (Low value, Low effort) — FILL-IN. Quick wins like auto-summarizing CRM call notes. Nice, not transformative. Do them when you have spare capacity.

Bottom-right (Low value, High effort) — SKIP. Be ruthless here. Automating a rare, complex edge case to save 20 minutes a month is vanity engineering. This quadrant is where most over-engineered AI projects quietly die.

Worked Example — Sales & Purchasing

Sales: AI-assisted quotation drafting. Scenario: a salesperson pastes a customer email request, and the AI drafts a `sale.order` with matched products and quantities. Readiness check: Q1 (deduplicated products) and Q2 (populated prices) are critical — if your `default_code` matching is unreliable, the AI picks the wrong SKU. Process check Q8: AI drafts in 'Quotation' state, the salesperson reviews and confirms. Never let it auto-send. Value: saves ~10 minutes per quote on a team doing 200 quotes/month = ~33 hours/month. Clear top-left candidate if data is clean.

Purchasing: invoice-to-bill matching. Scenario: incoming vendor PDF invoices are extracted and matched against open purchase orders. Readiness check: Q9-Q11 matter because vendor contacts may be sole traders (personal data). Build the match-confidence score and route anything below 95% confidence to a human. Worked metric: if you process 500 bills/month and 70% match cleanly, you have automated 350 bills and focused human attention on the 150 that genuinely need it. The error cost reduction (catching duplicate or mismatched bills) often justifies the project alone.

Worked Example — Finance & Inventory

Finance: expense categorization and anomaly flagging. Scenario: AI suggests the correct analytic account and tax for each `account.move.line`, and flags outliers (e.g. a EUR 12,000 'office supplies' entry). Readiness check: Q3 (consistent historical categorization) is the make-or-break — if your chart of accounts was restructured last year, scope the training window accordingly. Q12 (auditability) is non-negotiable in finance: every AI suggestion must be logged and a human posts the entry. Pattern: AI proposes, accountant approves, ledger stays defensible. Value: reduces month-end close time and catches misclassifications before the auditor does.

Inventory: reorder and stockout prediction. Scenario: AI analyzes `stock.move` history and lead times to suggest reordering rule adjustments. Readiness check: this is a top-right (high value, high effort) item.

It needs clean movement history (Q3), documented exceptions like seasonal spikes (Q7), and stable supplier lead-time data. Do not start here on day one. Start with a read-only 'suggested reorder report' that a planner reviews weekly — get the predictions trusted before you let anything write to `orderpoint` rules automatically.

One-Page Start-Here Summary

Step 1 — Score yourself. Answer all 12 questions, 0-2 each. Under 12/24: fix data and process foundations first. 12-18: you are ready for one carefully scoped pilot. 19-24: you can run multiple automations in parallel.

Step 2 — Pick ONE workflow. Use the value-vs-effort matrix. Choose a single top-left (high value, low effort) candidate. Resist the urge to automate three things at once on the first project.

Step 3 — Define the human checkpoint. Write one sentence: 'The AI does X, then [role] reviews and confirms before Y happens.' If you cannot write that sentence, you are not ready.

Step 4 — Clear the EU checklist. Confirm lawful basis, data residency, a signed DPA with no-training terms, and an audit log. Do this before the first API call, not after go-live.

Step 5 — Pilot for 30 days, measure, then expand. Track time saved, error rate, and human override frequency. If humans override the AI more than 20% of the time, your data or rules need work — go back to Part 1. If overrides are under 5% and value is proven, move to the next workflow on your matrix.

> **Score your team in 15 minutes, then book a free Odoo AI readiness review to turn your top-left workflow into a working pilot.**