

Forty AI ideas. Three worth funding.

An industrial services company had a list of AI ideas and no way to choose between them. We put a cost against every item on it. Eleven days later they had three funded projects and thirty-seven closed.



The short version

- 23 died because the data was not there. Not messy data. No data, or data sitting on a client's server we cannot read.
- 9 died in legal and security. One of them took eight minutes.
- 5 lost to something the company already owned. Usually a database query written years ago.
- 3 survived. Together they save Rs 3.19 crore a year against a build cost of Rs 1.1 crore.
- The model is the cheap part. On the biggest project, tokens are 35% of the running cost. The people around it are 57%.

The company had been arguing about this list for seven months. The argument was never really about AI. It was about which of forty things deserved money, and nobody had put a number against any of them.

Everything here is composite. It is built from patterns we see in Indian industrial and services businesses between Rs 250 Cr and Rs 500 Cr of revenue. No client data appears in this document, and the full model behind it is available on request.

How the list was made

Nobody had costed a single one

A hackathon in March produced twenty-two ideas. Three vendors pitched eleven between them over the summer. The board asked about AI in the November pack, which generated seven more in a fortnight. Two of the forty turned out to be the same idea submitted by different teams under different names.

This is normal. Ideas are cheap to produce and nobody is rewarded for killing one. So the list grows and the decision keeps getting deferred, usually to a committee.

Three questions, in this order

We put every idea past three questions. The order is the whole trick. Most companies start at the third question, spend a month building a business case, and only then discover the data does not exist.

	QUESTION	WHY IT COMES FIRST
1	Is the data there today?	Not could it be. Is it, right now, in a system we control, with enough history to test against. This kills the most ideas and takes the least time to answer.
2	Are we allowed?	Client consent, regulation, audit trail, personal liability. A twenty-minute call with legal saves a twelve-week build.
3	Does it beat what we already do?	Against the real current cost, including the people. Some tasks are already cheap. A model that costs more than a clerk is not progress.

What the three questions did to the list

40 ideas on the list

17 have the data to try

8 clear legal and security

3 beat what we already do

Each idea drops out at the first question it fails. An idea that fails on data is never assessed for cost.

Question one

Twenty-three fell at the first hurdle

Over half the list failed before anyone opened a spreadsheet. Not one of them failed because the data was untidy. They failed because it was somewhere we cannot reach, or because nobody had ever recorded the one field that mattered.

IDEA	WHAT STOPPED IT
Predict equipment failure from sensor logs	The sensors belong to the client. We read them during a site visit, not continuously. The idea is sound and we cannot execute it.
Rank inbound sales leads	Leads are marked 'closed' whether they were won or ignored. There is no outcome to learn from.
Predict late invoice payment	Payment dates are recorded. Disputes are not. The signal that explains lateness is the missing one.
Auto-classify procurement spend	Category coding was quietly abandoned in 2023. Seventy-one per cent of lines now read 'General'.
Answer engineer questions from manuals	The manuals are scans of print. Many have no text layer and the OCR on the older ones is poor.
Summarise competitor tenders	They hold nine competitor tenders in total. That is not a corpus, it is a folder.

Three of these become viable within a year for reasons that have nothing to do with AI. Log the dispute reason. Record whether a lead was won. Put the manuals through decent OCR. None of that is a model project and all of it is cheaper than one.

That became the second deliverable. A short list of boring data fixes, each with the AI project it unlocks written next to it. The head of operations found it more useful than the funded projects, because it gave her a reason to do work she had wanted to do for years.

Questions two and three

Permission, and the things we already do cheaply

Nine ideas cleared the data question and failed on permission. These conversations are short. We recommend having them in week one, not week ten.

IDEA	WHY IT STOPPED
Auto-generate compliance certificates	A statutory certificate issued on the strength of model output is not defensible, and the engineer who signs it is personally liable. Legal answered in eight minutes.
Summarise client contracts for renewal	The contracts carry client commercial terms under NDA. Doing this properly needs written consent from about forty clients. Three of them already prohibit it outright.
Draft safety incident reports	Reportable incidents go to the regulator. The cost of the wrong wording is far larger than the cost of writing it by hand.
Screen CVs for technician roles	Hiring decisions need an audit process the company does not have yet. Worth building. Not this year.

Then five that lost to the status quo

These are the ones people find hardest to hear, because somebody has usually already demoed them.

IDEA	THE COMPARISON
Detect duplicate invoices	A SQL query written in 2019 catches 94% of duplicates at no marginal cost. The model got to 97% for Rs 40,000 a month. Three points of recall, priced per point, was not worth it.
Route the shared sales inbox	Forty emails a day. One person clears them in twenty-five minutes. The saving is about Rs 11,000 a month against Rs 90,000 to run.
Auto-respond to tender clarifications	Nine a month, each genuinely different. There is no pattern to learn.

A model does not have to be good to fail this test. It has to be better than a clerk who is already quite fast, and cheaper than a query somebody wrote seven years ago and forgot about.

The three that survived

Where the money actually was

USE CASE	VOLUME / MONTH	BY HAND	WITH A MODEL	SAVING / YEAR
Extract fields from supplier invoices 48,000 documents, one format family, six years of history	48,000	Rs 35.00	Rs 13.90	Rs 1.22 Cr
Draft service quotes from scope notes a senior still signs every one	320	Rs 3,120	Rs 432	Rs 1.03 Cr
Triage inbound support tickets 68% route correctly with no human	62,000	Rs 20.00	Rs 7.30	Rs 0.94 Cr

The finding that surprised the board

Break the Rs 13.90 apart and the model is the smallest line on it.

WHAT MAKES UP THE COST OF ONE DOCUMENT	SHARE
Model, embeddings, storage	Rs 4.80 35%
Human review, on the 22% that need a look	Rs 5.50 40%
Exceptions, on the 4% that go wrong	Rs 2.40 17%
QA sampling	Rs 1.20 8%

Fifty-seven per cent of the running cost is people. So the question that decides whether this project makes money is not which model to buy. It is how many documents a human has to look at, and how long each look takes. Nearly every vendor conversation we sat in was about the first thing.

The trade-off nobody offers you

We stopped improving it at 78%

Higher accuracy means fewer documents need a human, which lowers cost. It also takes more engineering to hold, and the errors left over are the hard ones. Those two things move in opposite directions, so somewhere between them is a best answer.

For this use case it is 78% straight-through, with 22% getting a human look.

STRAIGHT THROUGH	HUMAN REVIEW	COST PER DOC	ENGINEERING / MONTH	NET SAVING / MONTH
60%	40%	Rs 18.70	nil	Rs 7.82 L
70%	30%	Rs 16.00	Rs 60,000	Rs 8.52 L
78%	22%	Rs 13.90	Rs 1.50 L	Rs 8.63 L
84%	16%	Rs 12.90	Rs 3.40 L	Rs 7.21 L
89%	11%	Rs 12.97	Rs 7.00 L	Rs 3.57 L
93%	7%	Rs 14.55	Rs 14.00 L	-Rs 4.18 L

Going from 78% to 93% looks like an improvement on every slide a vendor will show you. It turns a project earning Rs 8.6 lakh a month into one losing Rs 4.2 lakh a month.

So the accuracy target was set by the model, in the sense that we ran this table and picked the top row. It was not set by ambition, and it was certainly not set by the number on the vendor's slide.

What we would tell you to do first

Write down every AI idea anyone in the business has had. Then, for each one, answer only the first question: is the data here today, in a system we control, with enough history to test against?

Most companies find that half the list disappears in an afternoon. What is left is short enough to have a proper argument about, and the argument is finally about money.

Troia Value Advisors — product and value advisory. We help founders, product leaders and investors make product and capital decisions that stand up in the P&L and the boardroom.

Our AI Transformation practice is led by Anagha Pavgi. The workbook behind this note, with all forty ideas and every formula, is available on request.

hello@troia.co.in · +91 89209 98019 · troia.co.in

Composite company. All figures constructed to demonstrate method. No client data appears here.