



# Make your AI a helper you can trust with real work.

AI engineering is the work that turns a clever AI demo into a helper you can trust with real work: hand it your own documents, test it, give it rules and keep checking it.

## Give your AI a proper first week.

**The story, decoded.** In our short film at [algoshred.com/ai-engineering](https://algoshred.com/ai-engineering), Lumi is your AI, the manuals are your documents, the drills are tests, Commander Nadia is the person whose yes the padlocked button needs, and Comet the station cat is anyone who slips the AI a planted note.

### SOUNDS FAMILIAR?

The demo wowed the room. Then it told a customer something nobody had said.

Nobody can tell which document an answer came from.

Every fix to the instructions seems to break something else.

It can send emails or change records, and nobody quite decided that.

### THE IDEA IN PLAIN WORDS

#### Answers from your own documents

It looks things up in your documents before it answers, shows the page it used so a person can check, and hands over when the answer isn't there.

**Experts call this:** retrieval-augmented generation (RAG), grounding

#### Worth knowing

Answers are only as good as the documents behind them. It can still pick the wrong page or make something up, so it shows its source and we test it.

#### Practice before real jobs

Practice questions with agreed good answers or marking rules, run before it gets real work and again whenever anything changes.

**Experts call this:** evaluation (evals), regression testing

#### Worth knowing

Drills check the things you thought to test. Real use finds new ones, so each real slip becomes a new drill, and the drills run again whenever the model, the documents, the instructions or the rules change.

#### Rules that sit outside the AI

Rules the AI can't change: hard locks on the biggest actions, each helper allowed only its own job's tools, and a record of who did what.

**Experts call this:** guardrails, least privilege, human-in-the-loop

#### Worth knowing

Rules reduce risk; they don't remove it. Filters that check what goes in and out can be tricked, so the biggest actions get hard locks, limited tools and a person's yes. Someone on your team stays responsible, and records show who did what. A planted page or note can still fool the AI. It can't be fully prevented, only limited, which is why the lock sits outside the AI and needs a person.

#### Keep checking, and watch the cost

Watch real answers after every update, catch slips early, keep answers short, and send easy questions to a smaller model that still passes the drills.

**Experts call this:** MLOps / LLMOps, monitoring, inference cost

#### Worth knowing

Switching back to an earlier version works only if each version is saved and ready. Model makers also retire old models, so we plan for that. Watching real answers catches some of what the drills miss. Short answers, a smaller model and saved answers are separate levers. Each one helps only if it still passes the drills, and saved answers are refreshed when the documents change.

### HOW WE WORK WITH YOU

#### STEP 01

##### One job

Choose one real job for the AI, agree who owns it, and try it on real and awkward questions.

#### STEP 02

##### Manuals

Prepare the documents it should answer from, show its sources, and agree when it hands over.

#### STEP 03

##### Drills

Turn real and tricky cases into practice questions, and agree which must pass before anything goes live.

#### STEP 04

##### Rules

Lock the biggest actions behind a person's yes, give each helper only its job's tools, and keep records.

#### STEP 05

##### Keep checking

Watch real answers and costs after launch, run the drills again after every change, and hand the routine to your team.

Start small with one of these: “One helper, one shelf of documents”; “Drills for the AI you already have”; or “Rules review for an AI that can act”.

### WHAT WE HELP WITH

Two of eight ways we help (all eight at [algoshred.com/ai-engineering](https://algoshred.com/ai-engineering)):

#### A plain look at the AI you have

A short list of what to fix first, in plain words.

#### Keeping the running cost in view

Running costs you can see and plan for.

## Does your AI need a proper first week?

Start with a plain look at your AI: email [contact@algoshred.com](mailto:contact@algoshred.com) with what your AI does today and what worries you about it, and we'll suggest where to start.

Vignesh T.V., CEO & CTO, Algoshred Technologies Pvt. Ltd.

WATCH THE FILM

[algoshred.com/ai-engineering](https://algoshred.com/ai-engineering)