

Fine-Tune, RL, or Fix the Prompt — Decision Checklist

Four diagnostic checks in order, green lights, red flags, and the budget heuristic.

AI Agents · Fine-tuning · LLM · Companion to <https://aabee.tech/blog/fine-tune-vs-rl-vs-prompting-decision-tree/> · October 2026

Run the checks in order — stop at the first that moves the metric

1. **Context check** (hours): add the expert document; pass@k jumps → build retrieval, not training data
2. **Model ceiling** (afternoon): strongest model, well-specified prompt; pass@k ≈ 0 → decompose the task differently
3. **Format check** (a day): right answer, broken JSON → SFT with 500–2 000 examples
4. **Preference check** (weeks): correct but unpreferred, gradeable → DPO first, RLHF with a real reward model

Green lights / red flags

- ☐ SFT: format instability, small data, clear target — green
- ☐ RL: measurable quality gradient + clean reward — green
- ☐ "Make it smarter" without a metric — red flag
- ☐ No baseline from steps 1–2 — red flag
- ☐ Reward you cannot red-team — red flag

Budget heuristic

- A fix not worth 2 engineer-weeks to validate is not worth a training run
- A training run that can't beat a 2-day prompt fix was not worth running

aabee.tech — AI agent engineering and software performance engineering, delivered as measurable outcomes. IT Park Uzbekistan Resident.

Questions or a project where this applies? b2b@aabee.tech · aabee.tech