

Generic AI gives you filler. These 451 open-source SKILL.md files teach the AI the structure a senior professional actually uses — quality checks and anti-patterns included — so the first draft is one you can ship.

451 skills 65 bundles 25 professions 208 eval-scored · avg 4.8/5 11 chained recipes MIT · client-side · your key

⚡ THE one place to start — no install, no signup

1. Open mohitagw15856.github.io/pm-claude-skills (free Gemini key, in-browser local model, or your own Claude/OpenAI key)
2. Type your task into the `command bar` — “a blameless postmortem for Friday’s outage” — it routes you to the right skill instantly
3. Fill the short form → Run → export as PDF / Word / slides Liked it? `npx pm-claude-skills add --agent claude`

🧑 Six things no other skills library does

- The Boardroom** Five AI executives debate your PRD in rounds — severity-tagged openings, cross-examination, a Chair’s verdict memo with *the changes that would survive the room*. Theatre mode: they argue **out loud**. Verdicts export as hash-verifiable attestations.
- The Firm** Charter a standing AI staff (CFO, CCO, CTO, Strategy...) who file delta-aware memos, **hold board meetings without you**, and log falsifiable predictions you score — each staff member earns a visible calibration record.
- The Gauntlet** Job posting → decoded role → your stories forged → live cross-examination (answer **by voice**) → salary negotiation vs a hidden band → readiness /100 + a frameable certificate.
- Document X-ray** Every sentence heat-mapped by evidence strength — data · sourced · assertion · superlative · hidden assumption — load-bearing lines outlined, weakest claims listed to fix first.
- Living artifacts** RICE renders as a live board — drag the confidence slider, the ranking re-sorts. Journeys draw their emotion curve; health checks animate gauges.
- The Gym** Negotiate a salary / renewal / vendor deal against an AI with a **secret bottom line**; the debrief reveals what was achievable and scores you against it.

Plus: `vision` the Skill Galaxy star map (take the ▶ Sky tour) · `vision` skills that read whiteboard photos & competitor screenshots · `verify.html` for attestations

📁 Reach for these first

You need...	Skill
Product docs that ship	<code>prd-template</code> <code>rice-prioritisation</code> <code>okr-builder</code>
Engineering rituals	<code>incident-postmortem</code> <code>architecture-decision-record</code> <code>ai-code-review</code>
Exec & board comms	<code>executive-update</code> <code>stakeholder-update</code>
Customer success	<code>cs-health-scorecard</code> <code>churn-analysis</code> <code>renewal-playbook</code>
The agent era <small>NEW</small>	<code>mcp-server-spec</code> <code>agent-readiness-audit</code> <code>agent-era-pricing</code>
AI at work <small>NEW</small>	<code>ai-usage-policy</code> <code>ai-roi-audit</code> <code>role-redesign-for-ai</code>
Your career	<code>pm-personal</code> bundle + the Gauntlet

Every Production-Ready skill ships `references/` (the judgment calls) + `templates/` (fill-in worksheets). 🇪🇸 The whole tier is in Spanish too.

🧩 Chained recipes (11)

- `/ship-a-feature` — fuzzy idea → problem → PRD → RICE → roadmap → GTM
- `/land-a-job` — decode → research → apply → interview
- `/ship-an-mcp-server` NEW — spec → readiness audit → pricing → oversight
- `/adopt-ai-properly` NEW — policy → roles → reviews → ROI audit
- ...plus `/launch-a-product`, `/close-the-quarter`, `/rescue-an-account`, `/run-discovery`

🧠 Memory that compounds

- **The Brain** — local plain-markdown memory, provenance-tagged `[data]→[hunch]`; skills read it before answering, write back approval-gated. No vector DB.
- `outcome-tracker` — every decision’s predictions recorded & scored against reality; your frameworks earn calibration records.
- `action-runner` — recommendations become real tickets/messages, risk-rated, approved per action. Nothing acts silently.

🔧 Install anywhere

Claude Code (+4 agents, 21 commands, hooks)	<code>npx pm-claude-skills add --agent claude</code>
Cursor / Codex / Windsurf / 60+	<code>... add --agent cursor</code>
Any MCP client	<code>claude mcp add pm-skills -- npx -y pm-claude-skills-mcp</code>
ChatGPT / Claude.ai connector	<code>pm-skills-mcp.pm-claude-skills.workers.dev</code>
Python (LangChain/CrewAI)	<code>pip install pm-skills</code>
Plugin marketplace	<code>/plugin marketplace add mohitagw15856/pm-claude-skills</code>
GitHub CI (PR descs, changelogs, doc lint)	copy a file from <code>action/examples/</code>
Browser extension (insert skills + lint your writing)	<code>extension/</code>

🏗️ Infrastructure, not just a library

- **SkillsSpec** — the formal SKILL.md standard (conformance L1-L4); SkillCheck + security audit gate every PR
- **SkillBench** — the professional-work benchmark: 12 frozen tasks, models scored bare vs skill-instructed
- **Community registry** — publish `you/skill-name` via one PR; served live over REST + MCP
- **Agents can hire it** — A2A discovery card + `message/send`; MCP pairings make skills act (Notion, Slack — approval-gated)
- **REST API** — `/v1/skills` · `/v1/search` · `/v1/workflows` · `/v1/community` (no auth) → n8n, Lovable, Obsidian
- **Training data** — the library distilled to a dataset, on Hugging Face each release

✅ Quality, out loud

- **Eval-scored, not just counted** — an LLM judge rates structure / completeness / usefulness / grounding → public [leaderboard](#)
- **PR gates** — SkillCheck (structure) + Security Auditor (injection/exfiltration) + a 20-page web smoke suite
- **Honest tiers** — 🟢 Production-Ready (50, with full depth) · 🟡 Stable · 🟠 Experimental