Scan AI Agent Skills Before You Install Them

NVIDIA SkillSpector · free, Apache-2.0, 23K+ GitHub stars · tested 30 Oct 2026 (v2.12.0)

Why

Skills and MCP add-ons for Claude Code, Codex and Cursor are just text and scripts from strangers. A bad one can hide instructions that hijack your agent.

Install

```
python3 -m venv sv && . sv/bin/activate
pip install git+https://github.com/NVIDIA/SkillSpector
```

Scan a skill folder (free, no API key)

```
skillspector scan --no-llm ./path-to-skill
```

Without `--no-llm` it also runs AI analysis, which needs your own LLM API key or a local Ollama model.

Our test

- A harmless "summarise notes" skill: **0/100, no security issues**.
- A planted bad skill (ignore instructions, reveal prompt, pipe a remote script to bash): **100/100 CRITICAL, 5 HIGH issues**, incl. instruction override and external script fetching.

Habit checklist

- ☐ Scan every new skill or MCP server before installing.
- ☐ Read anything it flags yourself. A tool helps, it doesn't replace judgement.
- ☐ Re-scan after updates.