SkillOpt: Microsoft Trains Agent Instructions Instead of Model Weights, Gains +23% Accuracy

Kabui, Charles

2026-06-03

Microsoft Research released SkillOpt, a text-space optimizer that trains natural-language skill documents instead of model weights. A separate optimizer model runs the frozen target agent on scored batches, reflects on failures and successes separately, then proposes bounded

add/delete/replace edits to the skill document. An edit is accepted only if it improves a held-out validation score. The approach borrows deep learning concepts (epochs, learning rates, minibatches) but applies them to text. Across six benchmarks, seven target models, and three execution harnesses (direct chat, OpenAI Codex, and Claude Code), SkillOpt was best or tied-best in all 52 evaluated settings. On GPT-5.5, it lifted average accuracy by +23.5% in direct chat, +24.8% inside Codex, and +19.1% inside Claude Code versus running the same models with no skill. The final artifact is a compact `best_skill.md` file, typically 300 to 2,000 tokens long.

A team can boost agent performance by sharing a text file rather than retraining a model. Skills transfer across model scales and execution environments: a skill optimized on GPT-5.5 in Codex still works when moved to Claude Code or to a nearby math benchmark without further optimization. Since the skill is just context tokens at runtime, there's zero added latency or compute cost. The system is open-source under MIT license and installable via `pip install skillopt`.

This fits a growing pattern alongside Alibaba's SkillClaw, Anthropic's CLAUDE.md files, and VS Code's .instructions.md: structured text instructions are becoming first-class optimization targets. Training the prompt may become as routine as training the model.

Sources:

- SkillOpt Paper (arXiv:2605.23904)
- SkillOpt GitHub Repository
- SkillOpt Project Page
- SkillOpt on HuggingFace Papers
