Generating AI Slop in LambdaPi

An Interactive Demo with Claude Code

LMF, Université Paris-Saclay

March 19, 2026

Ciarán Dunne

From Chatbots to Agents

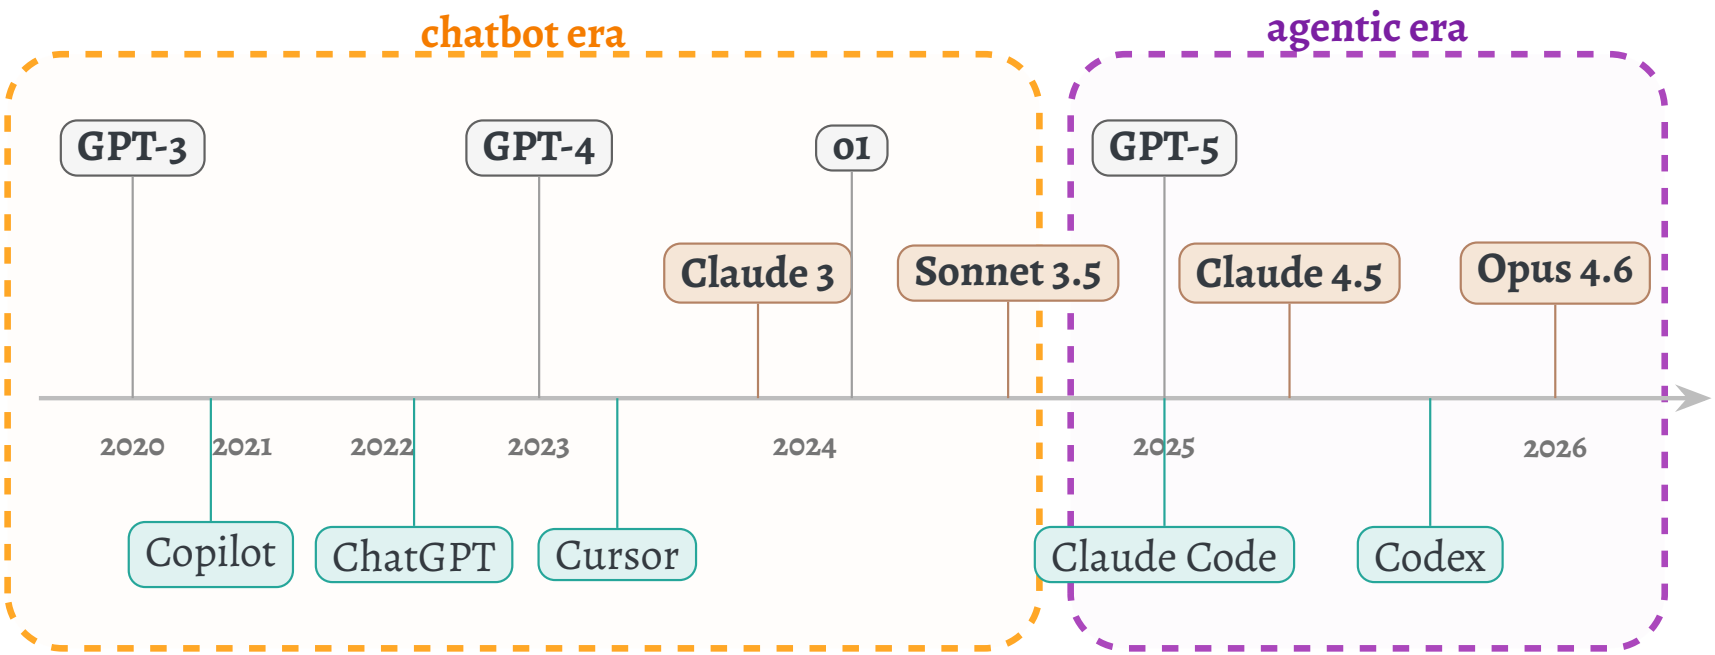

AI and Formal Math

Lean

- LeanCopilot, LeanDojo/ReProver, Leanstral (Mistral)
- HILBERT (Apple) — 99.2% on miniF2F
- AlphaProof (DeepMind) — IMO 2024 silver

Rocq

- CoqPilot (JetBrains), Rocq-MCP

Isabelle

- HybridProver (59.4% miniF2F), IsaMini

All benefit from **extensive training data** — large proof corpora (Mathlib, Isabelle AFP, etc.) in the LLM pretraining mix.

Two Ways to Extend an Agent

MCP (Model Context Protocol)

- runtime 'tool' integration
- constrained access to databases, APIs, services

Agent Skills

- literally just Markdown files
- instructions for using a CLI, where to find docs, etc.

The LambdaPi Skill

How it was built — written by Claude, iteratively refined

- Read lambdapi source, docs, and stdlib to draft `SKILL.md`
- Wrote 1,753 lines of test proofs — discovered 47 edge cases
- Refined: 670 → 343 lines + supporting reference files

Final layout

- `SKILL.md` — 343 lines (loaded first)
- `STDLIB_REFERENCE.md` — 485 lines
- `CLI_REFERENCE.md` — 133 lines

Isabelle/ZF → LambdaPi

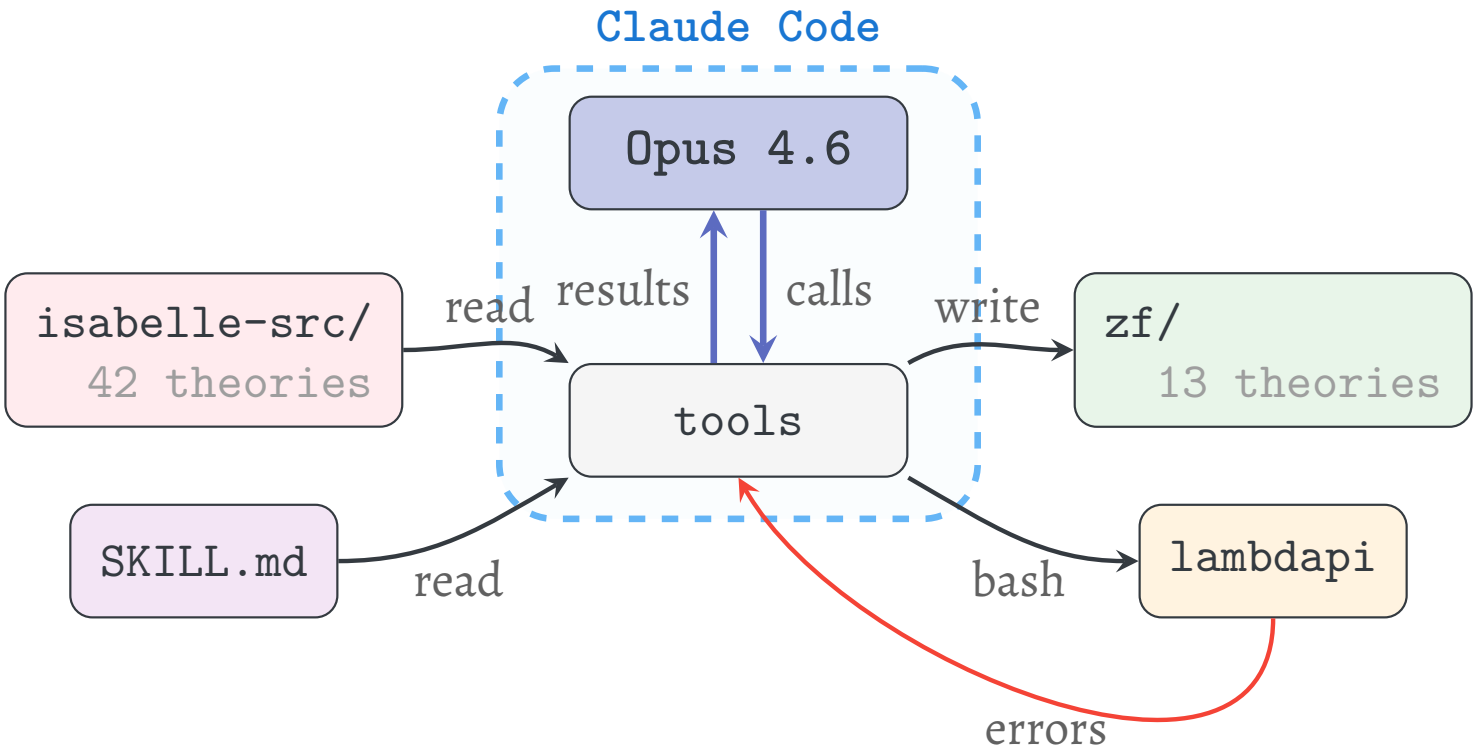

Goal: formalize up to `trans_induct3`. ✓

13/42 theories · 6,468 lines · 306 lemmas · 63 defs · ~3 hours

Demo: Waffle

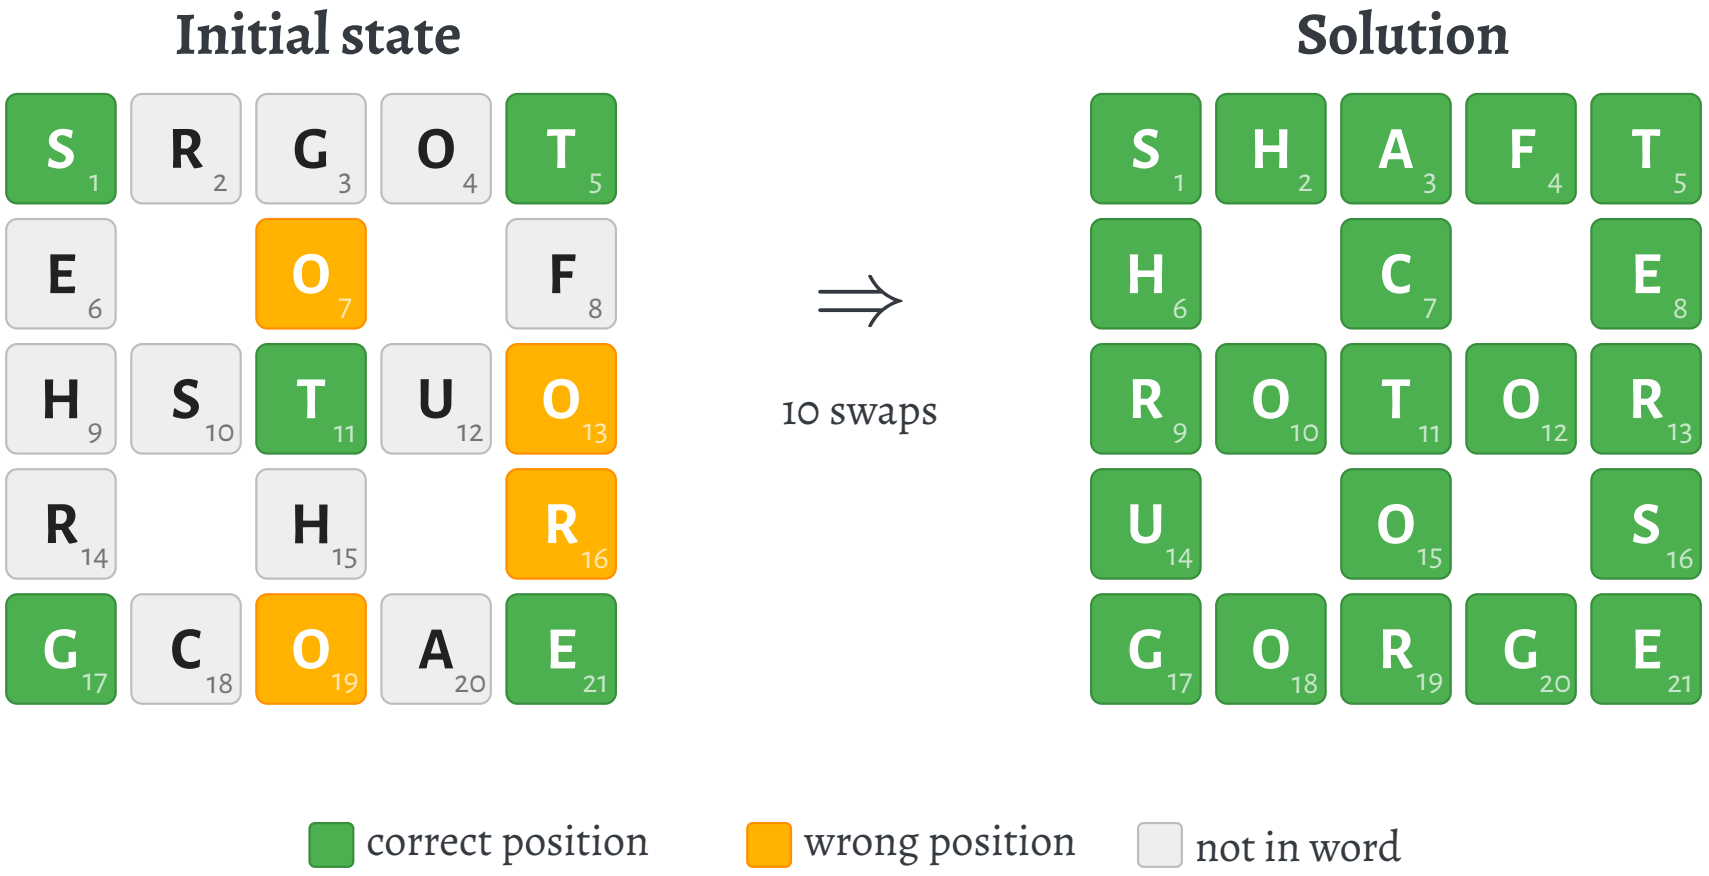

$$\sigma = (2\ 9)(3\ 20)(4\ 15\ 6\ 8)(7\ 18)(10\ 16\ 13)(12\ 14\ 19) \qquad \text{6 cycles, 10 swaps}$$

Chantal's Conjecture

Conjecture (Keller, 2025):

> *Call a tile **free** if it is not green. If a letter appears free **exactly once**, then we can move it to its correct position and **remain optimal**.*

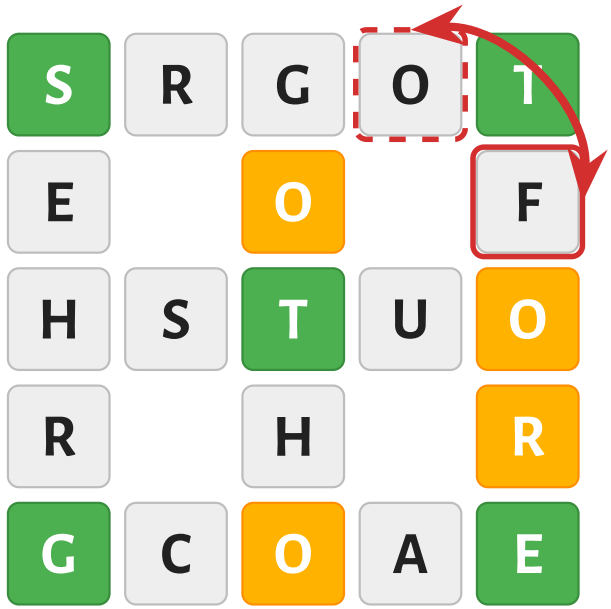

Let's formalize this!

Claude Code + LambdaPi Skill