Jukka Suomela
Aalto University

AI tools in TCS research

Focus of this talk

Using chatbots to actually ***do*** TCS research, ***solve*** our research questions — not just to assist in research

Goals of this talk

Short-term: share practical advice that might help our community

Long-term: understand what are fundamental characteristics of problems that AI can solve and why

- *what makes a TCS question easy?*

Plan for this talk

1. **Success stories:** flavor of problems
2. **Setup:** tools and best practices
3. **Future:** beyond problem-solving

Success stories

2-coloring cycles

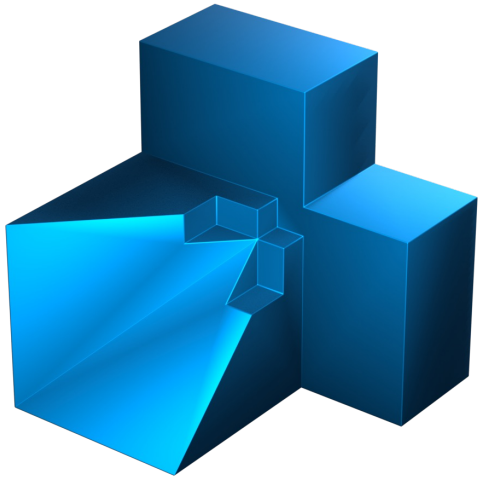

- **Question:** 1-round randomized distributed algorithms in cycles, 2-coloring, fraction of monochromatic edges?
- **Prior:** 0.20 … 0.25 [trivial]
- **New:** 0.23879 … 0.24118
- **Techniques:** ***new lower-bound technique,*** graphs, symmetries, SDPs, solvers

Quantum advantage

- **Question:** does quantum help with distributed symmetry breaking?
- **Prior:** major open question
- **New:** not if we require success probability 1
- **Techniques:** linear algebra, derive a contradiction in a clever new way

Lane changing

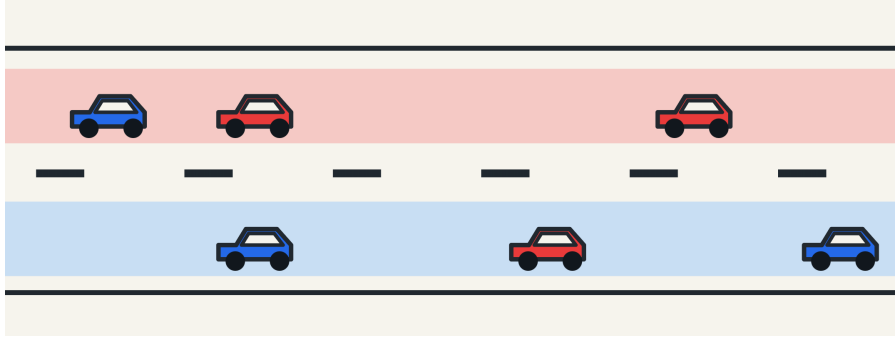

- **Question:** move cars on a 2-lane road to the right lanes with minimum delay
- **Prior:** 1.5-approximation [ATMOS 2018]
- **New:** optimal poly-time algorithm
- **Techniques:** a dynamic programming algorithm for calculating the optimal cost in linear time

Matrix multiplication

- **Question:** how many communication rounds are needed to multiply large d-sparse matrices
- **Prior:** $O(d^{1.832})$ from numerical optimization [SPAA 2022, SIROCCO 2025]
- **New:** $O(d^{1.790})$ with a closed-form expression
- **Techniques:** combinatorics, graph theory

Finitely dependent

- **Question:** smallest ***d*** so that there is a ***d***-dependent stochastic process
- **Prior:** 4-coloring: **1**-dep, 3-coloring: **2**-dep [Holroyd–Liggett 2016, *Forum Math. Pi*]
- **New:** weak 2-coloring: **3**-dep, maximal independent set: **6**-dep …
- **Techniques:** large linear system

Problem flavors

- distributed algorithms, quantum computing, centralized polynomial-time algorithms, parallel algorithms, stochastic processes …
- graph theory, combinatorics, algorithms, probability theory, linear algebra, optimization …

Setup

My tools

- ChatGPT Pro 20x (229 €/month)
- Codex + GPT-5.6
- Docker sandbox: ***github.com/suomela/cursor-sandbox***

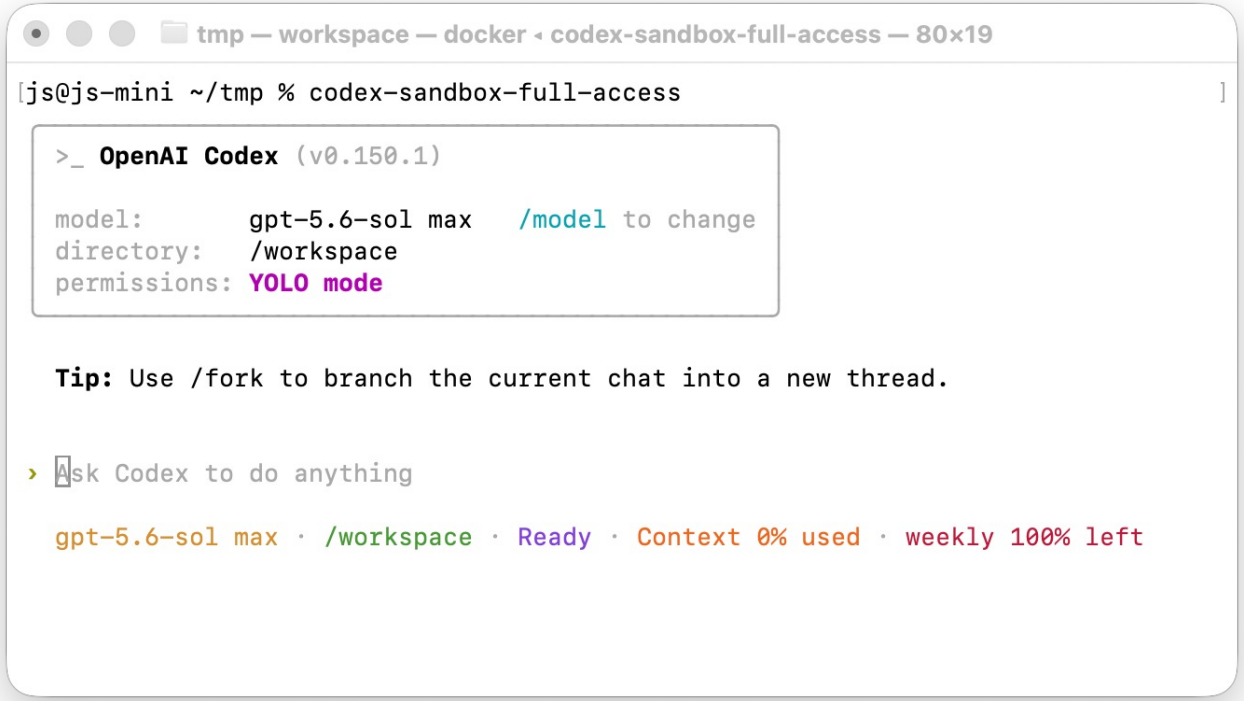

Best practices

- Just prompt in ***Codex***, don't worry too much
- Let it run for hours
- Suggest ***multiagent setup***, adversarial audit, "fresh pair of eyes" (just prompt)
- Ask it to formalize in ***Lean 4*** (in Palomar style)
- Use ***Git***

Pitfalls

- Gets stuck in running large ***numerical experiments*** that will never lead to a proof
- Gets stuck ***reformulating*** the hard problem as another hard problem
- Solves uninteresting special cases
- Cheats in Lean by formalizing conditional versions

What might help

- Explicitly say what you don't want, ***what counts as a solution***, what doesn't, etc.
- Require that all computational experiments must be high-value and ***explicitly justified***
- Refer to e.g. the prompt used for the *Cycle Double Cover* proof
- Prompt for a prompt

Bottlenecks

- Verification
- Understanding

Future

Thinking big

- Tackling problems way beyond what we would dare to touch without AI assistance?
- General reusable ideas, proof techniques, algorithm design paradigms?
- Theory-building?

Our ongoing work

- **Explanations:** ***What really makes maximal matchings hard?*** — systematically explore relaxations, identify a "minimal hard problem"
- **Exploration:** ***A new model of computation that we don't understand yet*** — iteratively prove upper/lower bounds for concrete problems, synthesize, generalize, identify next open questions, repeat…

Our ongoing work

- **Rigor:** ***Solid foundations for distributed complexity theory*** — pedantic definitions of all possible model variants, carefully re-derive all key results, formalize everything in Lean …

Unsolved

- Ethical issues, environmental impact, costs
- Authorship, publishing, reviewing
- Impact on teaching, researcher training, our own ability to do research
- Predicting when it works

Summary

Aalto ALC Slack:
#chatbots

- Yes, chatbots **can solve** many of our research problems, across a wide range of areas of TCS
- Yes, it is often very **easy** — just write a couple of informal prompts in Codex
- No, we don't really **understand** why it works so well, how to predict whether it works, what to do with it, and how it changes our field