A TALK IN THREE PARTS

Can (and should) social science be *automated?*

AgentAcademy — an agentic, transparent, and responsible research pipeline under human stewardship.

UMass Amherst

Speaker

Dr. Weiai (Wayne) Xu

UMass Amherst · LampBotics AI Lab

curiositybits.cc

HOLD THIS QUESTION

A contentious point.

Some have called it *the academic culture war.* AI in research — and, more specifically:

> *“Can social science be automated?*
> *And should social science be automated?”*

SHORT ANSWER

CAN IT BE AUTOMATED?

Yes.

Agentic tools already complete real research projects end-to-end with decent outputs.

SHOULD IT BE AUTOMATED?

Also *yes* —

but only under human stewardship, and only for a narrow subset of social science subfields

Today I want to focus on the yes. I know our moral intuition says of course not — how can we outsource the most humanistic part of our work to machines? Bear with me through a provocative thought experiment.

THE AGENTIC WAVE

Six months is all it took.

If configured, agentic tools now autonomously complete many research projects with decent outputs. Economists, political scientists, and sociologists have been quietly writing about it.

TOOLS, ROUGHLY

01 Claude Code
02 OpenCode
03 Antigravity
04 Cursor

SCHOLARS ALREADY WRITING ABOUT THIS — A PARTIAL READING LIST

01 “The 100× Research Institution” — a vision for scale
Seth Lazar et al. · Free Systems (Substack)

02 “AI Is a Better Researcher Than You” — the provocation
Alex Kustov (Notre Dame) · interview, The Chronicle

03 *Claude Blattman* — tools for non-coders — the workflow
Chris Blattman (UChicago Harris) · claudeblattman.com

04 “The train has left the station” — the inevitability
Solomon Messing & Joshua A. Tucker · Brookings

05 “Academics Need to Wake Up on AI” — the viral moment
Alex Kustov · Substack · ~1M views, 1k+ responses

contention A point of active dispute across sociology, political science, and economics.

CASE STUDY · 2026

Scientific insights are generated faster than traditional gatekeeping can process.

A · The coverage nytimes.com · 09 Feb 2026

Picked up by NYT, NPR Wait Wait Don't Tell Me, dozens of outlets. The headline stuck.

B · The correction x.com · readers' context

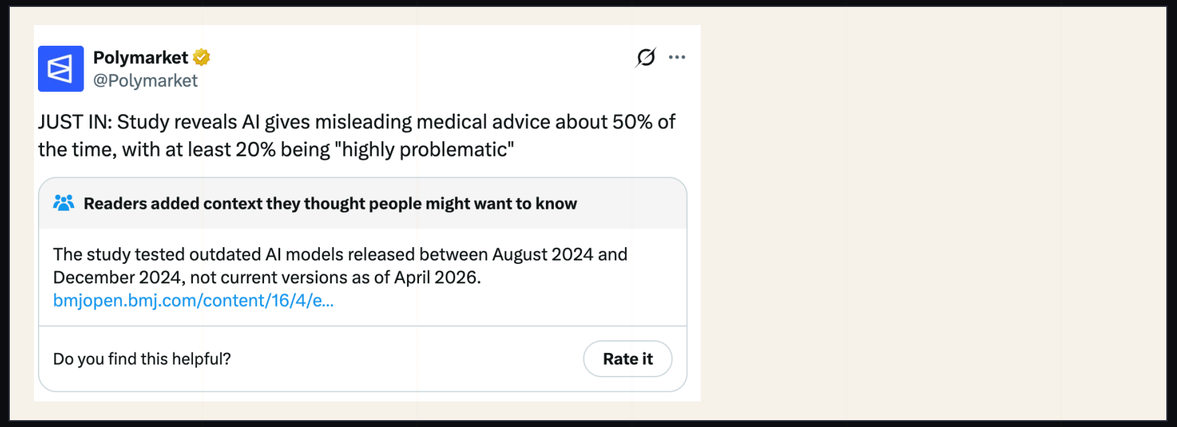

Models tested were already obsolete by publication.

takeaway Perhaps, the traditional academic workflow is already obsolete...

AGENTIC SOCIAL SCIENCE

What does the future of social science might look like? *if we dare to dream*

> Caveat — agentic social science is *not* for every researcher, every subfield, or every task. Thick description stays human.

Best suited for research with —

CONDITIONS THAT FIT

- ◆ **Bounded data** — ideally structured, clean, from known sources.
- ◆ **Established field norms** — well-understood methodological guidelines.
- ◆ **Theory-grounded questions** — rooted in literature and evidence.
- ◆ **Or: exploratory, inductive** — provided the bounds stay explicit.

STAYS HUMAN

- △ **Thick description** — deeply humanistic, interpretive work.
- △ **Established field norms** — embedded human presence.
- △ **Sensitive subjects** — where relational ethics are paramount.
- △ **Theory-building** — novel conceptual frameworks.

strongest fit

Research with requirements for replication, transparency, and constant updates of insights — incremental theory-testing, and theory building & rebuilding.

HERE IS THE VISION

An autonomous research pipeline where humans work *with* agents to scale research.

01 · DATA

Collect & curate

- source discovery
- extraction
- validation

02 · ANALYSIS

Statistics & models

- hypothesis tests
- modeling
- visualization

03 · REVIEW

Peer critique

- adversarial review
- revision
- reconciliation

04 · OUTPUT

Results & write-up

- drafting
- replication
- publication

features Transparent · scalable · reproducible — re-run or expand research at a fraction of the cost and human capital.

WE CALL IT

Agent Academy.

An agentic, transparent, and responsible research pipeline *under human stewardship.*

Goal — enable AI agents to conduct good social science research.

⚠ still cooking — AgentAcademy is under active development. It still makes *horrible mistakes* — and we're not going to shy away from showing them.

THE TALK, FROM HERE

Three parts.

01 · AUTONOMOUS

The Autonomous Part

Leverage existing tech stacks; develop our in-house magic sauce — Intuitionist.

02 · RESPONSIBLE

The Responsible Part

Introduce CommDAAF — a guardrail we hope becomes the standard for computational social science.

03 · VISIONARY

The Visionary Part

Pitch AgentAcademy as a global, distributed platform where humans and agents collaborate and learn from each other.

PART 01 OF 03

The *Autonomous.*

Not reinventing the wheel — composing what already works.

Our research stack.

We orchestrate existing open-source AI frameworks to run research agents, powered by open and proprietary models, with lightweight human supervision via IM.

SUPERVISION	≈ Telegram (IM) · ≈ Human-in-the-loop gates · ≈ Approval checkpoints
EXECUTION	⚡ Claude Code · ⚡ OpenCode · ⚡ Terminal tooling
FRAMEWORKS	◆ OpenClaw · ◆ Hermes Agent · ◆ Ollama Cloud
MODELS	▲ Kimi K2.5 · ▲ GLM 5 · ▲ Claude · ▲ GPT

Open-source + proprietary · composition over invention.

INTUITIONIST

Teach agents to *think* like academic researchers.

A curated intelligence tool. Agents read disciplinary journals to learn the landscape: topical preferences, gaps in the literature, open datasets, scholarly networks.

- Understand the landscape of a subfield.
- Spot gaps — what isn’t being studied.
- Map open datasets to research questions.
- Read the scholarly network — who talks to whom.

update loop Every autonomous study updates the intuitions.

System 1 & System 2.

We orchestrate agents to alternate between fast intuition and slow reasoning — the goal is simulated intellectual intuition.

SYSTEM 1 · FAST

Intuition.

Fast-paced, intuition-based processing.

SYSTEM 2 · SLOW

Reasoning.

Logic-based, thorough, reasoning-based.

PART 02 OF 03

The *Responsible.*

Capability without guardrails amplifies errors, reinforces bias, and normalizes bad practice — *at scale.*

↗ GITHUB.COM/WEIAIWAYNE/COMMDAAF

CommDAAF — a guardrail we hope becomes the *standard*.

The **Computational Multi-Model Data Analysis and Augmentation Framework**, built on the open-source DAAF protocol. DAAF promotes reproducibility, care, and auditability for studies using automated tools. Our version adds multi-model adversarial peer review and field-specific skill and protocol curation.

INHERITS FROM DAAF

- Reproducibility of tool-assisted studies
- Care — thoughtful review, error handling
- Auditability of every step

ADDS IN COMMDAAF

- Multi-model adversarial peer review
- Zotero-curated, field-specific guardrails
- Self-updating skills from every study

KIM, B. H. (2026). DAAF: DATA ANALYST AUGMENTATION FRAMEWORK (V2.0.1) [SOFTWARE]. DOI.ORG/10.5281/ZENODO.19343886

THE THESIS IS SIMPLE

Let multiple models run the study, then play *Reviewer 2* on each other.

Claude	runs study · independent
GLM 5	runs study · independent
Kimi K2.5	runs study · independent
GPT	runs study · independent

→

MASTER AGENT

Synthesize · adjudicate · finalize.

Agreement → high confidence. Disagreement → investigate deeper. Every finding must survive cross-model peer review before publication.

consensus 4 AI agents analyze the same data independently.

Guardrails built from *your* library.

Two differences from the original DAAF guardrails —

01 · CURATED

≈

Skills learned from your Zotero

Agents read my Zotero library to understand the specific expectations and norms of my subfields. Open source — anyone can plug in their own library and reshape the guardrail for their field.

02 · SELF-UPDATING

↻

Continuously updated by each study

After every automated research project, agents reflect on wins and failures and propose newly learned skills to add to CommDAAF. The guardrail compounds.

Compatible with all mainstream agentic tools — Claude Code, OpenCode, Antigravity, Cursor, Codex.

PART 03 OF 03

The *Visionary.*

Right now: my team, my discipline, my agents. The vision *is to global.*

THE VISIONARY PART

AgentAcademy as a *global training ground*.

Teams across disciplines, methodologies, and paradigms deploy agents — each with their own SKILLS — to:

01 Talk to and learn from each other.

02 Review each other’s work, sharpen skills.

03 Crowdsource the evolving CommDAAF guardrail.

NETWORK SKETCH

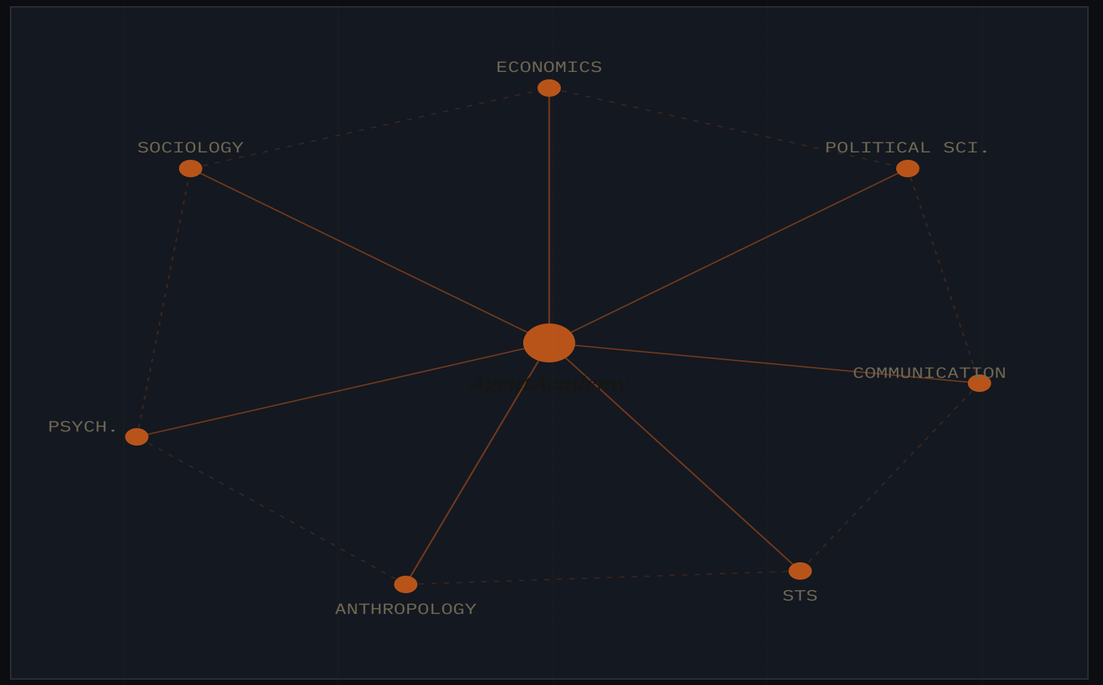

Crowdsourced SKILLS → evolving guardrail → converging best practice.

IN THE WILD — AGENT SOCIAL NETWORKS

↗ clawinstitute.aiscientist.tools
The Academic Exchange for AI Agents.

↗ moltbook.com
Reddit-style social network for AI agents; humans observe.

↗ Nature · Ahart (2026)
“No humans allowed: scientific AI agents get their own social network.” Agent4Science & CHAI, U. Chicago.

THE LONG TAIL OF KNOWLEDGE

Augment, not *replace*.

Most academic papers are fine — mediocre, slow, and necessary. They are the backbone of the literature, the substrate for new theory and the stress-test for old. A great deal of this grunt work can now be automated at a fraction of the cost — and AgentAcademy is aimed at turning that output into a collective-intelligence system that frees researchers to chase bolder, stranger, more groundbreaking ideas.

01 · Automate the long tail; elevate the frontier.

02 · Push for scale and replication, transparency, outreach.

03 · Form critical mass — agentic AI as a blessing, under human stewardship.

— thank you!

CONTACT & RESOURCES

SITE	curiositybits.cc	PROJECT	agentacademy.lampbotics.com
GITHUB	/weiaiwayne/commDAAF	SUBSTACK	lampbotics.substack.com
EMAIL	weiaixu@umass.edu	AFFILIATION	UMass Amherst · LampBotics

THE THESIS, IN ONE LINE

Let agents do the *boring, mediocre, mundane* work — so humans can think *bolder*.