Contact Information

Name Hong Chen

Position PhD Candidate, School of Information, University of Michigan

Visiting Research Fellow, Kellogg School of Management, Northwestern University

Email hongcc@umich.edu

Website hongcchen.com

Address 2200 Hayward Street, Ann Arbor, MI 48109

Graduate Study

I am a PhD candidate at the University of Michigan School of Information and a Visiting Research Fellow at the Kellogg Center for Science of Science and Innovation at Northwestern University. I study the economics of science and innovation.

I entered the PhD program in Information at the University of Michigan in Fall 2022. My coursework covered causal inference, experimental methods, and machine learning. I have advanced to candidacy and expect to graduate in 2028. My advisors are Prof. David Jurgens and Prof. Misha Teplitskiy. My work has appeared in *PNAS Nexus* and at ACL, NAACL, and ICLR. My research is funded by J-PAL and a Rackham Graduate Student Research Grant.

Research Interests

I study how scientific knowledge is produced, selected, and used, and how artificial intelligence changes each of these steps. Science is a central input to long-run growth, yet the institutions that govern it (funding, peer review, and publication) decide which ideas are pursued, which results become public, and which findings reach the people who pay for research.

My work asks three questions. How do these institutions select what we know, and what is lost in the process? What does society get back from public investment in research, beyond patents and productivity? And how does AI change the economics of science: as a tool that makes the hidden parts of science measurable, as a technology that lowers the cost of doing and reporting research, and as a new participant that reads and produces knowledge on its own? I address these questions with large-scale text data built with language models, field experiments, and causal inference.

Dissertation

Working title: *Essays on Artificial Intelligence and the Economics of Science*

Committee: Prof. David Jurgens (co-chair) and Prof. Misha Teplitskiy (co-chair), School of Information, University of Michigan; Prof. Dashun Wang, Kellogg School of Management, Northwestern University; Prof. Chris Rider, Ross School of Business, University of Michigan.

Overview. Standard data on science record only what survives: published papers, posted results, granted patents. The dissertation uses AI to measure what is left out and what is inflated, tests whether AI can bring hidden results back, and asks what AI itself learns from a selected literature. Four projects are in progress.

1. Public returns to R&D: measuring the reach of research with LLMs. Returns to R&D are usually measured through firms: productivity, patents, and spillovers. A better-informed public is a main justification for public research funding, yet existing data track only formal mentions, such as news stories and posts that cite a paper, and miss most informal discussion. Podcasts are now a major news source in the United States, and research is discussed on them often, almost always without a formal reference: only 1 in 68 episodes in our data contains a DOI or full citation. I built an LLM pipeline that reads transcripts, detects these informal mentions, and links each one to the paper discussed; it finds the correct paper about 75 percent of the time against human annotations. A first pass links about 90,000 papers to 210,000 episodes from 30,000 shows. Matching papers to NIH and NSF grants gives public reach per research dollar, and the 2025 federal grant terminations give a first setting to study how funding cuts change the public circulation of research.

2. A counterfactual approach to measuring hype in science. How much does the way a finding is written matter, apart from the evidence behind it? If assertive language draws attention on its own, the signals that readers and reviewers rely on are distorted, and LLM writing tools now make such language nearly free to produce. Observational data cannot answer this, because stronger claims and stronger evidence travel together. I use an LLM to rewrite the same abstract at controlled levels of rhetorical intensity while holding the evidence fixed, which isolates the effect of framing on attention and citations. The same framework tracks rhetorical intensity over time and shows that it has risen in published science and accelerated since LLM writing tools entered research workflows. The next step asks whether readers discount inflated language or whether the market for scientific claims tips toward hype. The framework was published at ICLR 2026.

3. AI research agents and publication bias. AI agents now read literatures and run analyses on their own, but the literature they learn from over-represents significant findings. I am testing whether agents inherit this bias and whether the tests used to evaluate them reward agreement with published results rather than with the truth. I build tasks in which the true answer is known by construction and score the agent's conclusion against it rather than against the published estimate. A companion experiment gives an agent data with no true effect and an open instruction to find something, and records how often it reports a significant result. I am also testing a fix: extracting every finding a paper reports, including minor and null results, and giving them equal weight in what the agent reads.

4. Illuminating science's file drawer: a field experiment on results disclosure. Many completed clinical trials never report their results, which biases the evidence available to clinicians and

reviewers. How much of this gap comes from effort costs that AI can remove, and how much still requires money? With funding from J-PAL, I am running a randomized experiment with the responsible parties of roughly 500 completed trials that have no results on ClinicalTrials.gov. The core treatment is AI-based logistical support. Posting results means assembling numbers and filling registry fields, work that investigators put off, so for every invited trial an AI system drafts a results summary from the public trial record. The investigator's task shrinks from writing to verifying and posting. A pilot compares the AI draft with a short written guide.