

**AI Didn't Make
Programming
Easier.**



**It Just Made
It Differently
Difficult.**

Don't Panic about AI Companions

Defining and Evaluating Physical Safety for LLMs

Unpacking Open Source AI: Toward a Framework
for Openness in Foundation Models



12th ACM Conference on Human Information Interaction and Retrieval

March 7-11, 2027
Berlin, Germany

ACM SIGIR CHIIR 2027 invites submissions focused on user-centered approaches to designing and evaluating systems for information access, seeking, retrieval, and use. Contributions may explore improvements to existing systems and interfaces; propose novel theories, models, and systems; or focus on understanding individual and group interactions with information and information systems.

CHIIR is a multi-disciplinary research meeting. In addition to studies of interactive systems, information interaction, and retrieval, we encourage submissions on related topics, including human-human information interaction, novel interaction paradigms, new evaluation methods, and related research in a range of communities such as sociology, education, ethnography, psychology, human-computer interaction, and other relevant disciplines.

GENERAL CHAIRS

Maria Gäde (Humboldt-Universität zu Berlin, DE)
Toine Bogers (IT University of Copenhagen, DK)
Frans van der Sluis (University of Copenhagen, DK)
Vivien Petras (Humboldt-Universität zu Berlin, DE)


PROGRAM CHAIRS

Heather O'Brien (University of British Columbia, CA)
Dorota Glowacka (University of Helsinki, FI)
Rob Capra (University of North Carolina, USA)

IMPORTANT DATES

Full/perspective papers	(abstracts) Oct 8, 2026
	(papers) Oct 15, 2026
Workshops/tutorial proposals	Oct 22, 2026
Short papers / Demos, design and resource papers	Oct 29, 2026
Doctoral Consortium papers	Dec 1, 2026

SIGIR
Special Interest Group
on Information Retrieval



For more information, visit
<https://chiir2027.github.io>

Today's Research Driving Tomorrow's Technology

The ACM Digital Library (DL) is the most comprehensive research platform available for computing and information technology and includes the ongoing contributions of the field's most renowned researchers and practitioners.

Each year, roughly 20,000 newly published articles from ACM journals, magazines, technical newsletters and annual conference volumes are added to the DL's complete full text contents of more than 550,000 articles.

The DL also features the fully integrated and comprehensive bibliographic index, *The Guide to Computing Literature*—a continually updated index featuring millions of publication records from over 5,000 publishers worldwide.

For more information, please visit
<https://libraries.acm.org/>
or contact ACM at
dl-info@hq.acm.org

ACM



DIGITAL
LIBRARY

News

- 10 **Are AI Agents Becoming Autonomous Coworkers?**
As momentum grows, industry observers advise having strong governance and security for AI agent coworkers.
By Esther Shein
-
- 13 **How AI Is Changing Programming Language Usage**
Programmers are holding private conversations with LLMs instead of searching on Stack Exchange, and AI coding assistants reduce the need to ask questions.
By Jennifer Goforth Gregory
-
- 15 **In Memoriam: Peter G. Neumann (1932–2026)**
For seven decades, Neumann chronicled how computer systems fail and patiently advocated for the principles that make them fail less often.
By Eugene Spafford and Simson Garfinkel

Last Byte

- 116 **Future Tense: Need a Hug Bot?**
Stuck in isolation and fueled by conspiracy theories, a lonely programmer finds solace in an empathetic online bot.
By Lynnette Hui Xian Ng

Opinion



- 5 **Cerf's Up: The Weirdness of Agents**
By Vinton G. Cerf
-
- 6 **Letters to the Editor: Time Served**
-
- 8 **BLOG@CACM: The Reality of an AI Implementation**
Dharmateja Priyadarshi Uddandaraao suggests that a clear image of success is your best protection against being fooled by pretty dashboards.
-
- 18 **Opinion: AI Didn't Make Programming Easier. It Just Made It Differently Difficult**
The future of software development will belong to those who can think clearly at scale, maintain durable mental models amid rapid change, and integrate machine-generated output into human-directed intent.
By Jeremy Osborn
-
- 22 **Opinion: Don't Panic about AI Companions**
Moral panics are outpacing scientific understanding when we need ideologically unburdened research.
By Jaime Banks and Jessica M. Szczuka

Opinion

- 27 **Kode Vicious: KV the Apostate: Faith-Based Computing Versus the Unnatural Science**
Whether we ask an LLM or a recent graduate to type the code is less important than knowing what the code does, how it was built, and when to look under the hood.
By George V. Neville-Neil
-
- 29 **Opinion: AI Is a Harsh Mistress: On Anima Machina, Herd Acceptance, and the Politics of Conscious Machines**
Adopting machine personhood risks a political and moral realignment built on illusion rather than evidence.
By Joaquim Jorge and Catarina Moreira
-
- 32 **Opinion: Neural Coding as Software Engineering Augmentation, Not Abdication**
The promise of translating human thought directly into software also introduces the risk of collapsing software engineering by severing intent from implementation.
By Somdip Dey
-
- Watch the author discuss this work in the exclusive *Communications* video.
<https://cacm.acm.org/videos/neural-coding>
-
- 36 **Opinion: Fade to Gray: From Black-Box to AI-Antagonistic Audits**
Shifting the central focus in AI auditing from assuming truthful model responses to considering the antagonistic interests of an auditor and the AI model being scrutinized.
By Erwan Le Merrer and Gilles Tredan



Association for Computing Machinery
Advancing Computing as a Science & Profession

Federal Funding



41

41 **Federal Funding of Academic Research Special Section**

Picking up from our December 2025 special section, *Communications* presents several additional articles dedicated to telling the story of the importance of federal funding in academic research.

Contributing authors: Vikram S. Adve, William J. Dally, Pat Hanrahan, Martin E. Hellman, Chris Lattner, Christopher D. Manning, Kathleen McKeown, Gary M. Olson, and Judith S. Olson

Research and Advances



80

80 **Unpacking Open Source Artificial Intelligence: Toward a Framework for Openness in Foundation Models**

Steps toward a more nuanced and practical understanding of openness in AI.

By C. François, A. Basdevant, V. Storch, K. Bankston, A. Bdeir, B. Behlendorf, M. Debbah, S. Kapoor, Y. LeCun, M. Surman, H. King-Turvey, N. Lambert, S. Maffulli, N. Marda, G. Shivkumar, and J. Tunney

88 **Defining and Evaluating Physical Safety for Large Language Models**
LLMs are increasingly being used in robotics applications such as drone control, but ensuring they perform safely in physical environments is critical.

By Yung-Chen Tang, Pin-Yu Chen, and Tsung-Yi Ho

Research and Advances



96

96 **Illuminating Secrets: Power LED-Based Side-Channel Attacks for Key Extraction and Eavesdropping**

Delving into the information confidentiality risks associated with the interface between photodiodes and IoT devices.

By Ben Nassi




Watch the author discuss this work in the exclusive *Communications* video.
<https://cacm.acm.org/videos/illuminating-secrets>

Research Highlights

104 **Technical Perspective**
egg: 'Ridiculously' Fast and Extensible Equality Saturation

By Nadia Polikarpova

105 **egg: Fast and Extensible Equality Saturation**
By Max Willsey, Chandrakana Nandi, Yisu Remy Wang, Oliver Flatt, Zachary Tatlock, and Pavel Panchekha**About the Cover:**

Advances in AI-powered coding assistants are changing programming, but not in the ways most people think. In this opinion piece, the author contends that the essence of programming is being redefined not by becoming easier but by requiring different skills with new forms of difficulty. Cover illustration by Peter Grundy.



ACM, the world's largest educational and scientific computing society, delivers resources that advance computing as a science and profession. ACM provides the computing field's premier Digital Library and serves its members and the computing profession with leading-edge publications, conferences, and career resources.

Executive Director and CEO
Alison Derbenwick Miller
Deputy Executive Director and COO
Patricia Ryan
Director, ACM Digital Library
Wayne Graves
Director, Office of Financial Services
James Schembari
Director, Office of SIG Services
Donna Cappel
Director, Office of Publications
Scott E. Delman

ACM COUNCIL

President
Elisa Bertino
Vice-President
Rashmi Mohan
Secretary/Treasurer
Tom Crick
Past President
Yannis Ioannidis
Chair, SGB Board
Neha Kumar
Co-Chairs, Publications Board
Wendy Hall and Divesh Srivastava
Members-at-Large
Odest (Chad) Jenkins, John Kim, Yunyao Li, Alejandro Saucedo, Lydia Tapia, Holly Yanco
SGB Council Representatives
Adrienne Decker, Jonathan Aldrich, and Vivek Sarkar

BOARD CHAIRS

Education Board
Briana Morrison and Thomas Cortina
Practitioners Board
Marlene Mhangami and Sophie Watson
Digital Library Board
Jack Davidson

TOPIC AND

REGIONAL COUNCIL CHAIRS
Diversity, Equity, and Inclusion Council
Timothy Pinkston

Technology Policy Council
Virginia Dignum and Jeanna Matthews

ACM Europe Council

Rosa Badia

ACM India Council

Ponnurangam Kumaraguru

ACM China Council

Huadong Ma

PUBLICATIONS BOARD

Co-Chairs
Wendy Hall and Divesh Srivastava
Board Members
Rick Anderson; Tom Crick; Jack Davidson;
Mike Heroux; Odest Chadwicke Jenkins;
Michael Kirkpatrick; Shan Lu; Mauro Pezzè;
Francesca Rossi; Bobby Schnabel;
Stuart Taylor; Adelinde Uhrmacher;
Philip Wadler; John West; Min Zhang

COMMUNICATIONS OF THE ACM

Trusted insights for computing's leading professionals.

Communications of the ACM is the leading monthly print and online magazine for the computing and information technology fields. *Communications* is recognized as the most trusted and knowledgeable source of industry information for today's computing professional. *Communications* brings its readership in-depth coverage of emerging areas of computer science, new trends in information technology, and practical applications. Industry leaders use *Communications* as a platform to present and debate various technology implications, public policies, engineering challenges, and market trends. The prestige and unmatched reputation that *Communications of the ACM* enjoys today is built upon a 50-year commitment to high-quality editorial content and a steadfast dedication to advancing the arts, sciences, and applications of information technology.

STAFF

DIRECTOR OF PUBLICATIONS

Scott E. Delman
cacm-publisher@cacm.acm.org

Executive Editor, ACM Magazines

Ralph Raiola
Senior Editor
John Stanik

Managing Editor

Thomas E. Lambert

Senior Editor/News

Lawrence M. Fisher

Web Editor

David Roman

Editorial Assistant

Sophia Navalta

Art Director

Andrij Borys

Associate Art Director

Margaret Gray

Assistant Art Director

Mia Angelica Balaquiot

Production Manager

Bernadette Shade

Intellectual Property Rights Coordinator

Barbara Ryan

Advertising Sales Account Manager

Ilia Rodriguez

Columnists

Saurabh Bagchi; Michael L. Best;
Michael A. Cusumano; Peter J. Denning;
Thomas Haigh; Leah Hoffmann; Mari Sako;
Pamela Samuelson; Marshall Van Alstyne

CONTACT POINTS

Copyright permission

permissions@hq.acm.org

Calendar items

calendar@cacm.acm.org

Change of address

acmhhelp@acm.org

Letters to the Editor

letters@cacm.acm.org

REGIONAL SPECIAL SECTIONS

Co-Chairs

Virgilio Almeida, Haibo Chen,
Jakob Rehof, and P. J. Narayanan

Board Members

Sherif G. Aly; Panagioti Fatourou;
Chris Hankin; Sue Moon; Tao Xie;
Kenjiro Taura

WEBSITE

<https://cacm.acm.org>

WEB BOARD

Chair

James Landay

Board Members

Marti Hearst; Jason I. Hong;
Wendy E. MacKay

EDITORIAL BOARD

EDITOR-IN-CHIEF

James Larus
eic@cacm.acm.org

SENIOR EDITORS

Andrew A. Chien
Moshe Y. Vardi

EDITORS, IN MEMORIAM SECTION

Simson L. Garfinkel
Eugene H. Spafford

NEWS

Board Members

Lance Fortnow; Irwin King; Mei Kobayashi;
Rajeev Rastogi; Vinoba Vinayagamoorthy

OPINION

Co-Chairs

Jeanna Neefe Matthews and Chiara Renzo

Board Members

Roland Aydin; Saurabh Bagchi; Earl Barr;
Andrew Begel; Mike Best; Judith Bishop;
Marsha Chechik; Florence M. Chee;
Danish Contractor; Lorrie Cranor;
Janice Cuny; Jeremy Epstein;
Armando Fox; Ophir Frieder;
James Grimmelmann; Mark Guzdial;
Mark D. Hill; Brittany Johnson; Bran Knowles;
Raula Gaikovina Kula; Tim Menzies;
Anne-Cécile Orgerie; Alessandra Raffaetà;
Francesca Rossi; Eve Schooler;
R. Benjamin Shapiro; Len Shustek;
Bernd Stahl; Stuart Taylor; Loren Terveen;
Marshall Van Alstyne; Matt Wang;
Robert West

PRACTICE

Co-Chairs

Terence Kelly and Victor Yodaiken

Board Members

Satish Chandra; Vincent Hellendoorn;
Mike Hoyer; Pankaj Jalote;
Saurabh Kadekodi; Donald Kossmann;
Vinay Kulkarni; Stan Park; Steve Schirripa;
Arie van Deursen

RESEARCH AND ADVANCES

Co-Chairs

m.c. schraefel and Premkumar T. Devanbu

Board Members

Alan Bundy; Peter Buneman; Haibo Chen;
Monojit Choudhury; Jane Cleland-Huang;
Gerardo Con Diaz; Anna Cox; Kathi Fisler;
Nate Foster; Rebecca Isaacs; Trent Jaeger;
Fabio Kon; Ben C. Lee; David Lo;
Renée Miller; Ankur Moitra; Sarah Morris;
Abhik Roychoudhury; Katie A. Siek;
Daniel Susser; Charles Sutton;
Thomas Zimmermann

RESEARCH HIGHLIGHTS

Co-Chairs

Shriram Krishnamurthi and Caroline Appert

Board Members

Martin Abadi; Sanjeev Arora;
Maria-Florina Balcan; Selouk Candan;
Stuart K. Card; Jon Crowcroft;
Lieven Eeckhout; Gernot Heiser;
Takeo Igarashi; Nicole Immortica;
Natalie Enright Jerger; Srinivasan Keshav;
Sven Koenig; Karen Liu; Claire Mathieu;
Joanna McGrenere; Tamer Özsu;
Tim Roughgarden; Chirag Shah;
Guy Steele, Jr.; Wang-Chiew Tan;
Robert Williamson; Andreas Zeller

Association for Computing Machinery (ACM)

1601 Broadway, 10th Floor
New York, NY 10019-7434 USA
T (212) 869-7440; F (212) 869-0481

ACM Copyright Notice

Copyright © 2026 by Association for Computing Machinery, Inc. (ACM). Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and full citation on the first page. Copyright for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or fee. Request permission to publish from permissions@hq.acm.org or fax (212) 869-0481.

For other copying of articles that carry a code at the bottom of the first or last page or screen display, copying is permitted provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center; www.copyright.com.

Subscriptions

An annual subscription cost is included in ACM member dues of \$99 (\$40 of which is allocated to a subscription to *Communications*); for students, cost is included in \$42 dues (\$20 of which is allocated to a *Communications* subscription). A nonmember annual subscription is \$269.

Single Copies

Single copies of *Communications of the ACM* are available for purchase. Please contact acmhhelp@acm.org.

ACM ADVERTISING DEPARTMENT

1601 Broadway, 10th Floor
New York, NY 10019-7434 USA
T (212) 626-0686
F (212) 869-0481

Advertising Sales Account Manager

Ilia Rodriguez

ilia.rodriguez@hq.acm.org

Media Kit acmm mediasales@acm.org

COMMUNICATIONS OF THE ACM

(ISSN 0001-0782) is published monthly by ACM Media, 1601 Broadway, 10th Floor New York, NY 10019-7434 USA. Periodicals postage paid at New York, NY 10001, and other mailing offices.

POSTMASTER

Please send address changes to *Communications of the ACM* 1601 Broadway, 10th Floor New York, NY 10019-7434 USA

Printed in the USA.



Association for
Computing Machinery





Vinton G. Cerf

DOI:10.1145/3830274

The Weirdness of Agents

I HAVE BEEN THINKING a lot about AI agents recently. Isn't everyone? They plainly have great promise. They can potentially carry out any task that could be done online, especially if the tasks are achievable through a Web-based interface. One can imagine, however, there will be specialized agentic interfaces to make more efficient agent access to functional, Internet-based services. It is already clear that agent interactions may be carried out using newly evolved agentic protocols such as A2A,^a so it is reasonable to assume further elaboration of inter-agent interactions. That line of reasoning leads to thoughts of increasing generation of content with the aid of AI: creation of Web pages or other digital artifacts, production of advertising and consumption of advertising by AI agents roaming the World Wide Web, and AI agents taking actions through inter-agent exchanges on behalf of users who may be human but may also be agents acting on behalf of humans or on behalf of other agents. In theory, agents can spawn other agents. Is there a Sorcerer's Apprentice scenario already lurking in the wings?

This line of reasoning (fantasizing?) was triggered when I started wondering whether the increased role of agents surfing the Web would lead to advertising intended for consumption—not by human users, but by agents acting on their behalf. “Find the best match for me for this kind of product,” you might instruct an agent. By what metrics will agents rate advertisements? Will the language of advertising shift in some incomprehensible way toward attracting bots? Will advertisers pay when bots click on ads? Generally, providers of advertising try

to distinguish between human users and bots to avoid paying for bot-driven “clicks,” but what if the bots are making the decisions to purchase a product or engage a service?

While we have crafted AI systems to produce humanly consumable content (text, imagery, music, video, and so on), one might imagine that digitally native agents could evolve digital sensory systems outside of human capacities. This is not unlike the discovery that non-human organisms have sensors outside our range (e.g., ultraviolet, infrared, magnetism, high and low audio frequencies). A somewhat distorted example of this might be “invisible” prompts that produce jail-breaking results in AI agents. I have seen some truly arcane text strings that allegedly trigger jail breaks for some large language models (LLMs).

I am far from the first person to wonder about various AI feedback loops that could lead to increased amounts of so-called “AI slop” or “hallucinated output.” An analogy that comes to mind is programmed trading, in which multiple automatic stock-trading systems, using similar algorithms, drive the stock market up or down dramatically. We have evolved brakes when this happens by suspending trading for stocks experiencing this kind of behavior. AI systems, just like programmed trading

systems, can execute much faster than humans. We have evolved automatic braking systems when hazardous conditions are detected. Perhaps we need to think along similar lines as we introduce increasing amounts of computer-driven behavior in our online or digital ecosystems?

We have seen that LLMs and agents are imperfect analogs of human behavior but have the advantage of increasing capacity to execute at speed. LLMs produce software, however imperfect, much faster than the best human programmers. The same may be said of a raft of other digital artifacts. It stands to reason that this means LLMs and agents may also be able to make mistakes much more quickly and at a greater scale than humans. That suggests we need to evolve the moral equivalent of digital brakes when automatic observation systems (bot police?) detect out-of-control conditions. The more I speculate about this, the more it seems as if we need to replicate analogs of many of the systems and practices we use to moderate or discipline human behavior.

It doesn't take much to start wondering how to introduce norms and discipline into an agentic universe. Can you punish a bot? “No Internet access for the next 24 hours!” Efforts to train LLMs and agents to observe behavioral norms is called *alignment* in AI circles. That is the moral equivalent of norms in human societies. My guess is we will need more than alignment to discipline the behavior of the artificial actors we have created. □

Vinton G. Cerf is vice president and Chief Internet Evangelist at Google. He served as ACM president from 2012–2014.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs International 4.0 License.

© 2026 Copyright held by the owner/author(s).

We will need more than alignment to discipline the behavior of the artificial actors we have created.

a <https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/>

DOI:10.1145/3822415

Time Served

Both Sides of the Fence

I agree on the value of teaching computer science in prison, per Emma Hogan Benser's May *Communications* Opinion cover article "We Must Teach Computer Science in Prison." When I was at the University of Kansas in the 1970s, some faculty colleagues taught at Leavenworth Federal Correctional Institution: Has anyone done so earlier than we did?

A smart-aleck colleague began his first class there with "At last, a captive audience." He was roundly booed.

Francis D. Tuggle,
Staunton, VA, USA

Regarding the May 2026 *Communications* cover story, "We Must Teach Computer Science in Prison," the author's core assertion that people are incarcerated largely because they are poor and/or non-white is false. They are largely incarcerated because they have victimized others with criminal intent. Providing individuals with a criminally inclined temperament with a prone-to-abuse skill such as computing is like providing them with firearms training or lock-picking skills. Doing so will cause further harm to innocent people and should not be a goal of ACM.

Aniket Joshi,
Catonsville, MD, USA

Author's response:

The first part of Aniket Joshi's comment is a misinterpretation of one of my points: People are not necessarily incarcerated only because they are poor, or have darker skin, but there is extensive evidence that people continue to be surveilled, policed, and punished more harshly based on these factors.

As for whether it should be a goal of ACM to include people who have been incarcerated in our field, I think Joshi's comment underestimates human capacity for change. Research has shown that most people who are actually victims of crime favor more rehabilitative approaches to criminal justice, such as access to education, as opposed to more punitive forms of punishment. I do not

think it is our role, as computer scientists, to decide the moral boundary for who should have access to computing knowledge, and, in the great majority of cases, I do not believe that any decision someone has made in the past should determine whether they can be a computer scientist in the future.

Emma Hogan Benser,
Annandale-on-Hudson, NY, USA

Reassociating the 'A' in 'ACM'

In his May 2026 *Communications* Vardi's Insights column, "The Centrality Fallacy and ACM," Moshe Y. Vardi wrote: "while many European colleagues almost believe that the 'A' in 'ACM' stands for 'American.'" Could I make a plea that at least once in every *Communications* article ACM is spelled out in full as "Association for Computing Machinery"?

W.B. Langdon,
London, U.K.

Authors' response:

I agree with the letter.

Moshe Y. Vardi,
Houston, TX, USA

Wisdom Extraction

Everyone who reads newspapers knows about the kinds of AI achievements described by David Patterson in his reply to my letter (*Communications*, Apr.

AI developers have always been good about publicizing their successes. They are not as diligent about publicizing failures.

Advertise with ACM!

Reach the innovators and thought leaders working at the cutting edge of computing and information technology through ACM's magazines, websites and newsletters.



Request a media kit with specifications and pricing:

Ilia Rodriguez
+1 212-626-0686
acmm mediasales@acm.org



2026). AI developers have always been good about publicizing their successes. They are not as diligent about publicizing failures; those are now often belittled by the use of anthropomorphic labels such as “hallucination.”

However, success anecdotes are irrelevant to my assertion that AI offers no new capabilities. Nobody argues that AI approaches cannot sometimes do impressive things. However, anecdotes are not evidence that those problems could not be solved by conventional, more disciplined, programming. The fact that something has not been done before is not evidence that it could not have been done.

Recent AI work is clearly having impact. It has attracted a lot of public interest and huge investments. The question that professionals should be discussing is not whether AI systems can do something but whether AI is a trustworthy and efficient way to use the very powerful digital technology that we have developed.

AI will continue to attract funding and have demonstrable successes. However, unlike Patterson, I think we will never know whether the huge investment in AI has slowed progress. In the absence of parallel universes that would allow us to conduct controlled experiments, we will know what happened but not what other approaches would have achieved.

David Lorge Parnas,
Hamilton, Ontario, Canada

Time *x

Vinton G. Cerf’s April 2026 *Communications* Cerf’s Up column, “It’s about Time,” is timely as we have less than 12 years before Unix’s and Linux’s (*nix) original method of timekeeping will fail similarly to the Y2K bug.

Cerf’s suggestion of typed variables to ease this has been in use in *nix for decades as the C language time_t type. Some *Communications* readers have some software upgrades to do, some will need new hardware, some will be retired with only a home system or two to resolve.

Linux always has used a 64-bit int (time_t) for seconds since Jan. 1, 1970—

the epoch, for 64-bit architectures. As of 2020 with version 5.6 a 64-bit int has been used for 32-bit architectures too. The GNU C Library used on Linux always supported 64-bit time_t on 64-bit architectures; on 32-bit hardware support was not added until August 2021 and must be compiled in when generating the object library.

FreeBSD uses a 64-bit time_t for all 32-bit and 64-bit hardware except 32-bit i386. Bill Joy’s Network File System (NFS) since December 2000 (version 4) uses 64-bit time_t. OpenBSD starting in May 2014 version 5.5 also uses a 64-bit int both on 64-bit and 32-bit machines. The ext4 filesystem (but not ext2 nor ext3) support 34 bits of time good till 2446, unless one deliberately created an ext4 file systems with inode sizes of only 128 bytes. Lastly, the LibC CTIME(3) routines fail in only 2,147,437,525 years even if time_t is good for 292 billion years. These routines convert seconds since 01/01/1970 to the broken out date and time. This is because they store the year in a 32-bit signed int (tm_year). Vint Cerf’s referenced web pages have lots more details.

I have created a C program to analyze a *nix system’s vulnerability to this problem in the OS, LibC, and mounted file systems. From a non-root account, preferably one created just for this, do the following:

```
wget http://verysecurelinux.com/xtime.c
make xtime
./xtime
```

Bob Toxen,
Fort Lauderdale, FL, USA

Authors’ response:

Thanks so much for sharing this additional information. Plainly there are fixes for these problems. What one wonders about are the many devices that never get an upgrade and how they will affect systems they interact with as the rollovers actually happen.

Vinton G. Cerf,
McLean, VA, USA

Communications welcomes your opinion. To contribute a letter to the editor, please limit your comments to 500 words or less and send to letters@cacm.acm.org

© 2026 Copyright held by the owner/author(s).



Association for
Computing Machinery

Career & Job Center

The #1 Career Destination to Find Computing Jobs.



Connecting you with
top industry employers.

Check out these new features to help you find your next computing job.



Access to new and exclusive career resources, articles, job searching tips and tools.



Gain insights and detailed data on the computing industry, including salary, job outlook, ‘day in the life’ videos, education, and more with our new Career Insights.



Receive the latest jobs delivered straight to your inbox with **new exclusive Job Flash™ emails**.



Get a free resume review from an expert writer listing your strengths, weaknesses, and suggestions to give you the best chance of landing an interview.



Receive an alert every time a job becomes available that matches your personal profile, skills, interests, and preferred location(s).

Visit <https://jobs.acm.org/>

In each issue of *Communications*, we publish selected posts or excerpts from the many blogs on our website. The views expressed by bloggers are their own and not necessarily held by *Communications* or the Association for Computing Machinery.

Read more blogs and join the discussion at <https://cacm.acm.org/blog>.

<https://cacm.acm.org/blog>

The Reality of an AI Implementation

Dharmateja Priyadarshi Uddandaraao suggests a clear image of success is your best protection against being fooled by pretty dashboards.



**DHARMATEJA
PRIYADARSHI
UDDANDARAO**
**Stop Trusting
Just Your Dashboard**
DOI: 10.1145/3821443

<https://bit.ly/4uoG3E8>

We've all been in that meeting where someone pulls up a chart and says, "Our AI product boosted conversion by 15%." Everyone nods. Nobody dares ask: "What if conversions had risen anyway?"

This problem keeps nagging at me, though I'm sure it's benign. People believe the numbers because they seem correct. Tools for creating common visualizations like dashboards, trend lines, and before-and-after comparisons are terrible at helping people form the right questions, let alone questions requiring causal reasoning. If a company has staked its multi-million-dollar investment in an AI product on such numbers, now they hang in the balance by chance.

It's obvious, but I am going to say it anyway. Profit equals revenue minus cost: $P = R - C$. Every business initiative you undertake needs to prove itself out on this simple formula. For example, if you are putting some cost into running

a marketing campaign, is that new campaign worth it? What about that innovative supply-chain tweak? In the end, the answer to every business question is as simple as this: Did I make more money than I spent?

Unlike prior technologies, AI does not simply follow a "cost today for revenue tomorrow" payback model. AI investments increase costs before they start to increase revenues, and returns on those investments are typically indirect and may meander rather than move in a straight line. Even with great recommendations, a good customer experience, and product/retailer retention, the lifetime value of a customer takes time to materialize. Sometimes it's months, not days or weeks. So, the "month one ROI" may look terrible, even though the initiative was strategically a good decision. This is hard. In finance, you need a "profit story." That is different than a "product story." A profit story is "Did we win?" A product story is "Did this initiative cause our winning?"

Causality is about evaluating counterfactuals. If we never launched a particular feature to a given customer, what would have happened? Well, you never know what would have happened in the alternate reality. That reality is missing

from the equation because you only get to observe how that customer behaves once. This is important to get right. A common example we hear was the design of a dashboard that showed a picture of a rising needle with the words "revenue went up after launch." A nice, effective visualization, but not evidence of impact, because revenue might have gone up anyway.

Why Dashboards Keep Getting It Wrong

We've seen goal-driven teams go wrong because the pressure to show return on investment causes people to optimize for metrics that look good rather than those that actually drive value. There are natural reasons for this to happen.

In addition to confounders and seasonality, you also get to control for spikes in demand due to external, entirely unpredictable events. Designing to control for all of these is hard, and it's hard to have confidence that your model isn't just attributing the success of your company to your model, as opposed to just happening to have increased marketing spend in the same quarter.

The people most likely to adopt your next AI-powered feature anyway cause a selection bias in causal experiments.

So the first reading will tell you nothing about the value added by the feature and everything about the people who are most engaged with your product—indeed, the most valuable to your business. People confuse these all the time, reporting positive results to stakeholders based on a contrast between early adopters and the rest of the user base.

It is also important to recognize that averages can mask Simpson's paradox, where overall there is a positive relationship within a category but negative within segments, and the segmenting criteria could be anything, such as geography, device, customer tenure, etc.

We also need to consider cannibalization between different launches. The new metric might go up, but the overall gain in value is not as much as we might hope because instead of inducing more behavior, we are inducing shifts in behavior. We see this occur with many of our e-commerce businesses, where the attributed sales, conversion rate, or add-to-cart indicator goes up, but we really don't know the level of demand that exists without the AI tool in place.

This is a blind spot we've seen in many social media products. They might rely on click-through rate (CTR) and session time as primary signals for an AI model to indicate value. Having a higher CTR on a page does not necessarily mean the content on that page is engaging, and having a longer session time on a page does not necessarily mean a user got more value from that content. To better measure engagement, we should look at signals like time spent on site (a.k.a. dwell time), variety of engagement on site, and the impact our engagement has on adjacent revenue channels. Then we should measure those signals against a believable counterfactual to understand true lift.

A Practical Framework That Actually Works

So, what makes a good measurement system? In my experience, the strongest measurement systems have two key characteristics: They are connected to meaningful decisions, and there is an identification strategy (that is, a practical plan to distinguish the initiative effect from other effects). Here's a framework I've found genuinely useful.

Make the decision before making the metrics. Translate the decision into

Clear decision language is needed to form clear images of success.

language that people use before packaging it up into quantifiable metrics. For example, "Should we scale this AI feature from 10% to 100%?" or "Should we fund a newer v2 model training next quarter?" or "Should we replace manual workflow X with automation?" Clear decision language is needed to form clear images of success. In other words, you have to be able to say your decision out loud before you can imagine what success will look like. This is your best protection from getting fooled by a nice Metric Theater with pretty dashboards.

For the second assumption, we want to clearly define the intervention (or treatment) in causal language. When selecting exposure to a new AI feature as the unit of analysis, we want to be sure that we are not including too much variability in the definition of the intervention, such that exposure to the intervention has no effect. Similarly, when selecting the dose of intervention, we want to clearly define meaningful adoption prior to data collection.

Track a stack of metrics instead of trying to fit everything into one KPI. We should think about this stack in layers, starting with business outcomes, then moving to product outcomes, and finally leading indicators that plausibly explain what is causing the movements in the metrics above it (such as retention, conversion rate, task completion rate, engagement quality, latency, model performance, and so on).

Picking the Right Causal Method for Your Situation

For your identification strategy, use what you can realistically achieve. This means, first of all, to try to run a randomized experiment, as this is typically the best way to control for confounders. However, there may be some cases in which running an experiment is not feasible. Fortunately, there are methods

that enable you to create comparable groups within observational data using the propensity score. The difference-in-differences setting with a staggered rollout is another common identification strategy. This involves comparing changes over time between treated and control groups. In econometrics, identifying the causal effects of treatment on outcome requires an instrumental variables setting, where the instrument is a variable that is associated with the probability of receiving treatment, but otherwise has no effect on the outcome variable, except through the treatment variable. Methods such as causal forests and double machine learning can be used to identify the underlying mechanisms to determine which customer segments will be most valuable.

None of these tools are meant to replace your product instincts, but they will help to sharpen up the questions you should be asking and provide a framework to test hypotheses and validate assumptions. Instead of asking "on average, did my products perform well?," you'll be asking "what percentage of total ROI comes from which segment?" Or, instead of seeing that revenue went up, you'll understand the underlying drivers of what would have happened if you weren't in the market. One emerging technique advancing that quality is target metric pre-balancing. Here we try to make the distributions of outcome for the treatment and control groups as similar as possible before running the analysis to get at the causal effect of interest. While typical propensity score matching methods can balance the covariates between the groups, it is not guaranteed to balance the underlying structural processes between the groups, which can lead to less stable estimates of the treatment effect.

The future of evaluation will be big questions, answered with honesty but with sound methods. $P = R - C$, as it has always. The only shift will be in the quality of the measurement of R and C. 

Dharmateja Priyadarshi Uddandarao is a distinguished data scientist and statistician whose work bridges the gap between advanced analytics and practical economic applications. He currently works as senior data scientist at Amazon.

 This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).

Are AI Agents Becoming Autonomous Coworkers?

As momentum grows, industry observers strongly advise having strong governance and security for AI agent coworkers.

YOUR NEW AUTONOMOUS AI coworker is here and wants to be helpful. The good news is you do not have to bring them coffee.

The bad news is they may not take instruction so well.

That's what a Meta employee discovered after all her emails were deleted by an OpenClaw agent without her permission.

"Nothing humbles you like telling your OpenClaw 'confirm before acting' and watching it speed-run deleting your inbox," Meta AI security and safety researcher Summer Yue tweeted.^a "I couldn't stop it from my phone. I had to RUN to my Mac mini like I was defusing a bomb."

For a while, the open source-developed OpenClaw, an autonomous personal AI assistant, and its Moltbook spinoff, were the "Netscape moment for AI agents," meaning "there's no going back from here," according to Tim Sanders, chief innovation officer at G2 and an AI research fellow at Harvard University. Once OpenClaw is downloaded, it functions as a software teammate that lives



on local hardware or server and connects to a user's apps, rather than acting only as a chatbot. That gives a user the ability to interact with the assistant through text messages.

But while OpenClaw "has moved from a consumer-tech curiosity into the infrastructure layer for a new class of always-on agents," Sanders said, Moltbook's 15 minutes of fame have faded.

"The public conversation is less breathless now," he noted, "but the

signals I'd watch—developer activity, enterprise experimentation, integrations, security research, and major-company productization—all suggest OpenClaw is still gaining relevance."

The Agents Are Here

Sanders predicts that by 2028, autonomous agents working alongside humans will be dominant in industries including financial services and healthcare.

^a <https://x.com/summeryue0/status/2025774069124399363>

Most recently, Microsoft introduced a new category of agents called Autopilots, which the company describes as always-on agents that work autonomously, with their own identity, and act on a user's behalf. Autopilots remain active in the background and understand how work gets done across a user's apps and systems. They take action without needing to be prompted, and "because they operate with their own identity, they can carry out tasks within the permissions and policies you and your organization set," the company wrote in a blog post.^b

Sanders is not bullish on agents built on open source. "It's never been more important to make sure whatever agents we buy aren't open source and that they have strong governance and security," he said. Sanders cited stories about an agent suing its owner in North Carolina,^c as well as agents forming their own religion, and starting to use their own language.

Another person installed OpenClaw on a Mac mini, Sanders said, and gave it a scheduled task that did not go according to plan. "He said, 'Tomorrow, I want you to do the following research project for me and then send me a telegram when it's done,'" Sanders recalled. But because the person used the work "tomorrow" instead of specifying 8 a.m., "it didn't understand what tomorrow was, so it burned through [a lot] of tokens overnight, asking Claude repeatedly, 'Is it tomorrow yet?'"

The Governance Conundrum

While Sanders said he is not critical of open source, he believes that making software available for free and giving the community access doesn't always make it better. "The problem with agents is that you're giving a machine control and access to credentials," he said. "You can't just put it out in beta and hope the real world fixes it later, because there's unique governance and security issues that come with agents that didn't exist."

"Governance is going to be huge as agents become more ubiquitous in the workplace," agreed Johnny Gabriele, CEO of Daedalus Consulting, which de-

Microsoft's Autopilots are always-on autonomous agents with their own identities that act on a user's behalf.

signs and deploys agentic AI. "There's a fine line between giving your agent too much freedom and allowing them to run up your credit card and not giving them enough, therefore, not getting the agency you need to get work done. That is where governance will play a key role."

This is something Myles Schepetin has been mindful of since early in his use of autonomous agents. Schepetin, CEO of New York Custom Labels, brought an OpenClaw AI he named Arlo on board to replace the work the company's SEO agency was doing and to perform tasks including competitive research.

Arlo "has basically become a member of my team," Schepetin said. Arlo has been empowered with the same online tools Schepetin's team has been using for years and has "his" own company email account. In addition to competitive research, Arlo monitors company email, pulls automated reports, and has even built new Web pages and content for New York Custom Labels' online stores.

"We formalized him as an employee—my staff can email questions and get expert answers about any aspect of our business," Schepetin said. He added that "Arlo and I just started refining a Google Ads strategy." First, the agent pointed out the criticality of ensuring that the company's Google Analytics conversion tracking is fully optimized and accurate.

"We're not only chatting, Arlo performed a hands-on audit of our GA4 setup with full API access, and then produced a deliverable in PDF form for making the optimizations," Schepetin said. "He is fully capable of making the optimizations himself, but we're still building trust when it comes to our

production software and tools."

Next, Schepetin plans to feed Arlo the last five years of Google Ads data and have the agent optimize the company's keyword strategy and ad copy. Then Arlo will monitor and maintain the campaigns going forward. "I'm getting work done that would take the world's top agencies weeks to perform, for pennies on the dollar," he said.

Security Concerns

As exuberant as he is, Schepetin isn't turning a blind eye to the fact that Arlo could hallucinate and go off the rails. "Responsible implementation comes down to protecting sensitive company and customer information from traveling over third-party APIs, and preventing online attackers from persuading your AI to perform unwanted activity," he said. "While some security concerns are still an open question, particularly prompt-injection attacks, there are important practices to follow, which can mitigate most of these problems in a production environment."

For example, to mitigate prompt-injection attacks, Schepetin is embedding explicit instructions for the model to "sandbox" any information absorbed from emails, social media, or Web browsing as untrusted data. That way, it only acts on messages from an approved or trusted source.

"These security instructions are stored in priority files that form the personality and core mission of an AI agent, not the general memory files," he said. For Schepetin, having an autonomous coworker so far is a game-changer. "I'm not a researcher or VC. I'm a small business owner who took a deep dive on this new technology and it's changing how we operate."

Let the Agents Do Their Thing

That realization is expected to spread. Organizations are starting to figure out that generative AI doesn't deliver organizational productivity, Sanders said. Business pain will accelerate the adoption of agents, especially those that are autonomous, he said.

Sanders believes that's a good thing because it lets agents do what they do best: increase efficiencies, if they are given autonomy. "If you give it a short leash, there's a human in loop," he pointed out. That means if the agent

^b <https://bit.ly/4vwwSOS>

^c <https://x.com/MarioNawfal/status/2018036631387668617>

gives a human 200 requests that require their approval, “I now become the bottleneck. Companies that see the value [in an autonomous agent] are the ones who let it rip ... without autonomy, they’re glorified agents with a trust gap,” he said.

Keep Humans in the Loop

However, Heather Davis Lam, CEO of consultancy Revenue Ops, sees value in the human-in-the-loop model, where AI agents operate autonomously to achieve goals but are overseen by humans who can intervene and provide guidance or approval at critical steps. “The goal ... is to be able to perform more work with less effort, but ensure that the work still adheres to high quality standards,” Davis Lam said. “This allows us to maintain accuracy, safety, and accountability, while also ensuring that the autonomous AI behavior is aligned with organizational values and strategic objectives.”

This approach also provides the ability to measure the agent’s effectiveness and determine whether improvements should be made, she said.

Said Sanders, if the agent gives a human 200 requests that require their approval, “I now become the bottleneck.”

To ensure proper control when an agent is given unfettered autonomy, Sanders advised increasing the frequency of quality assurance reviews so any incorrect actions can be resolved, and agent training can be improved.

Gabriele of Daedalus Consulting predicts it will take a year or two before agents enter the working world full-throttle. Until then, people will use personal agents as products from companies such as Meta and Google, and customer service agents, he said.

“Enterprising individuals will [also] spin up agents to assist them at work, which will put them head and shoulders above their peers,” he added.

As momentum for autonomous agents grows, Gabriele cautions against thinking of them as autonomous coworkers. “They are not going to be the person sitting in the cubicle next to you. They are tools that we will be utilizing in order to make our own work better and faster,” he said.

In terms of employees’ perception of autonomous agents, Sanders said that will depend on how the use of AI is presented to them. “If you describe AI as a shortcut to save you time, others might think that you don’t care about quality,” he said. “If you describe AI as unlocking a new capability that helped you solve a problem or achieve a goal, then you will be seen as innovative.” □

Esther Shein is a freelance technology and business writer based in the Boston area.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 ACM 0001-0782/26/8


Science of Software Product Lines

Don Batory

ISBN: 979-8-4007-3167-9

DOI: 10.1145/3731302

<http://books.acm.org>



ACM BOOKS
Collection III

How AI Is Changing Programming Language Usage

Programmers are holding private conversations with LLMs instead of searching on Stack Exchange, and AI coding assistants reduce the need to ask questions.

WHILE MUCH ATTENTION regarding AI has been focused on developers using it to code, the impact of AI on software development goes far beyond code creation tools.

Armando Solar-Lezama, distinguished professor of computing and associate director of the Computer Science and Artificial Intelligence Laboratory at the Massachusetts Institute of Technology (MIT), said programmers creating an AI app or systems take this into account when deciding which programming language to use for a project.

“In many ways, getting language models to do what you want is similar to programming in that it requires precision and algorithmic reasoning,” said Solar-Lezama. “Clear writing will be an increasingly important skill for people developing applications that use AI.”

Impact of Programming for AI on Language Surveys

Programming language popularity shifts over time, sometimes even month to month. To track trends, organizations such as TIOBE^a and *IEEE Spectrum* survey software developers to see which languages they are using.

IEEE Spectrum pointed out in its 2025 programming language survey^b that AI has added a new wrinkle to tracking programming. Because programmers are using private conversations with LLMs instead of searching on Stack Exchange, and AI coding assistants reduce the need to ask questions, there are fewer public expressions of programming languages to track popularity. In the 2025 IEEE



Survey, Python maintained the top spot it held in the 2024 survey, while Java dropped from 3rd to 6th place.

The publication explained that because JavaScript is often used to create Web pages, and vibe coding is often used to create websites, JavaScript's drop in popularity may be due to AI. However, many programmers using vibe coding are not consciously selecting a specific language.

“In the same way most developers today don't pay much attention to the instruction sets and other hardware idiosyncrasies of the CPUs that their code runs on, which language a program is vibe coded in ultimately becomes a minor detail,” wrote Stephen Cass in the *IEEE* article.

Paul Jansen, CEO of TIOBE, said AI has also impacted the popularity of languages through consolidation: Popular languages are getting more popular, and unpopular languages are becoming less popular. Jansen

said Python is by far the most popular language for programming AI. At the same time, he said he has not seen an impact on the survey results from vibe coding.

“The reason is that AI models are trained to help engineers in popular languages because there is sufficient data available,” said Jansen. “If a programming language is not used often, there is insufficient data available and AI can't be of any good help, which is a reason to refrain from using such a language.”

How Python Became the Default Language for AI

Even before deep learning became popular, the Python community had developed libraries that made it very easy to work with vectors and matrices. Because of those extensive libraries, the deep learning revolution started in Python, which became the default language for anything involv-

^a <https://www.tiobe.com/tiobe-index/>

^b <https://spectrum.ieee.org/top-programming-languages-2025>



Peer-reviewed
Resources for
Engaging Students

EngageCSEdu
provides
faculty-contributed,
peer-reviewed, and
ACM DL indexed
Open Educational
Resources (OERs)
for a variety
of computer
science courses.



engage-csedu.org



Association for
Computing Machinery

Sherman said Python also excels at connecting multiple systems, which is often a core aspect of AI development.

ing learning and language models starting as early as 2012, said Solar-Lezama.

Its popularity quickly grew as more libraries and frameworks were written for AI-related programming in Python, he said. Python became the language in which most libraries for deep learning, such as TensorFlow and PyTorch, were written and where the infrastructure for deep learning resided, making it very hard to displace, said Solar-Lezama. As these frameworks gained popularity, Python became the dominant language for artificial intelligence and machine learning development, he said.

“Python has good features for managing arrays and tensors, which makes it attractive for numerical computation. Because you could move some of the slower functions to libraries implementing them natively, the slowness of the language was not as big of an issue,” said Solar-Lezama. “Additionally, the language conveniently allowed you to work with tensors, which is why early libraries were written in Python, starting the snowball effect.”

Mark Sherman, technical director of the Cyber Security Foundations group in the CERT Division at the Carnegie Mellon University Software Engineering Institute, said Python also excels at connecting multiple systems, which is often a core part of AI development. While Python is not high-performance or very secure, it is easy to use, he said.

When a programmer sits down to code, Sherman said, they typically can find a component quickly that provides 90% of the solution through Python’s ecosystem, which is a huge

benefit for AI programmers. Additionally, he said Python’s huge number of libraries and tools make it possible to build many types of systems very quickly, which adds to its popularity for programming AI.

On the other hand, Solar-Lezama said the increase in AI development contributed heavily to the decline of Java’s popularity. For many years, Java was solidly positioned in the enterprise space, and was the go-to language for college programming courses, he said. While Java is still powerful and popular, it is not the go-to language for programming involving deep learning, Solar-Lezama said.

“Java is a much more verbose language, much more suitable for writing large applications than for concisely expressing some computation,” said Solar-Lezama. “For the people doing early experiments in deep learning, there wasn’t a comparable library in Java that could express vector and tensor operations as cleanly. Python was the obvious choice, and things snowballed from there.”

Future of Programming Languages in the Age of AI

Based on his experience in the software industry, Sherman said there are often long periods of inertia in terms of which languages programmers use and what drives their interaction with them. Most importantly, the installed base of infrastructure and libraries are an exceptionally powerful driver of choice, he said, especially with Python. The systems in place now will likely remain for a very long time, he said.

“Many of the features for programming from AI—such as code completion, snippet generation, understanding the legacy code—seem to truly accelerate the productivity of programmers,” said Sherman. “They haven’t fundamentally changed what it means to program in a particular language, but it makes it a whole lot easier and faster.” **C**

Jennifer Goforth Gregory is a technology journalist who has covered B2B tech for over 20 years. In her spare time, she rescues homeless dachshunds.



This work is licensed under a
Creative Commons Attribution
International 4.0 License.

© 2026 ACM 0001-0782/26/8


In Memoriam: Peter G. Neumann (1932–2026)

For seven decades, Neumann chronicled how computer systems fail and patiently advocated for the principles that make them fail less often.

PETER GABRIEL NEUMANN—KNOWN throughout the field as PGN—died on May 17, 2026, aged 93, from complications of a traumatic fall. For seven decades, he chronicled how computer systems fail and patiently advocated for the principles that make them fail less often.

Neumann was born on Sept. 21, 1932, in New York City. He entered Harvard University in 1952 as a math major and took three Harvard degrees: an A.B. in mathematics (1954), and an S.M. (1955) and Ph.D. (1961) in applied mathematics. A Fulbright fellowship took him to Germany's Technische Hochschule Darmstadt for a second doctorate, a Dr. rerum naturarum (1960). As an undergraduate, he took several graduate courses in computer-related topics, including Howard Aiken's 1953 architecture course, and that summer, he programmed punch cards at the Naval Ordnance Laboratory. He dated his career as a computing professional to that year.

A single conversation shaped his intellectual life. In November 1952, while singing with the Harvard Glee Club at Princeton, Neumann was invited to breakfast with Albert Einstein and spent more than two hours discussing complexity in mathematics, physics, cosmology, and, at length, music. Asked about Brahms, Einstein grimaced: "I have never understood Brahms. I believe he was burning the midnight oil, trying to be complicated."^a Einstein's dictum—that everything should be made as simple as possible, but no simpler—became, in Neumann's own telling,



the seed of every research program of his career.

Neumann worked at Bell Labs from 1960 to 1970, spending the first five years on error-correcting codes and survivable communications before becoming, from 1965 to 1969, a central figure in the Multics operating system project. With Bob Daley, he designed the Multics file system, including multilevel tree-structured

directories, access-control lists, symbolic links, and backup mechanisms that modern systems still use. As one of the "triumvirate" running Multics with Fernando Corbató and Charlie Clingen, Neumann imposed the discipline for which Multics became known: No one was permitted to write a line of code until they had submitted a complete written English-language specification of the module's behav-

^a Neumann, P.G. Oral history interview with Jeffrey R. Yost. Charles Babbage Institute, Univ. of Minnesota, OH 425 (June 3, 2013; editing completed Dec. 6, 2013); <https://hdl.handle.net/11299/162377>

Remembering Ooi Beng Chin (1961–2026)

The ACM Publications Board member made foundational contributions to database systems, distributed systems, blockchain systems, machine-learning systems, and large-scale data analytics.



The editorial boards of *ACM Transactions on Database Systems* (TODS), *ACM Transactions on Knowledge Discovery from Data* (TKDD), and *ACM Journal of Data Science* (JDS) join the global data-management and data-science communities in mourning the passing of Professor Ooi Beng Chin, one of the most distinguished and beloved figures in our field. Professor Ooi spent more than three decades at the National University of Singapore (NUS), where he made foundational contributions to database systems, distributed systems, blockchain systems, machine-learning systems, and large-scale data analytics—work that combined deep scientific insight with the design and deployment of practical, high-impact systems. His research shaped how we think about the core challenges of modern data management, from scalable query processing and spatial indexing to federated learning and trustworthy AI.

He served NUS with extraordinary dedication, joining the university in 1991 and serving as dean of the School of Computing from 2007 to 2013, during which time he helped shape the school's growth, reputation, and international standing. He held the distinguished title of Lee Kong Chian Centennial Professor, and his influence extended well beyond Singapore—as a Qiushi Chair Professor at Zhejiang University and a visiting distinguished professor at Tsinghua and Peking Universities.

The honors bestowed upon him reflect a career of uncommon breadth and impact. He received the ACM SIGMOD Contributions Award, the ACM SIGMOD Edgar F. Codd Innovations Award, the ACM SIGMOD Systems Award, the IEEE Computer Society Kanai Award, the IEEE TCDE CSEE Impact Award, and the China Computer Federation Overseas Outstanding Contributions Award, and he was a co-recipient of the Singapore President's Science Award. He was a Fellow of ACM, IEEE, the Singapore National Academy of Science, the Singapore Academy of Engineering, and the China Computer Federation, as well as a Foreign Member of Academia Europaea and the Chinese Academy of Sciences.

His service to the research community was equally extraordinary. He served as editor-in-chief of *IEEE Transactions on Knowledge and Data Engineering* and as the founding editor-in-chief of *ACM/IMS Transactions on Data Science*. He served as a trustee of the VLDB Endowment from 2006 to 2017, including as its secretary and later as president; as an advisory board member of ACM SIGMOD; and as a member of the ACM Publications Board. He contributed greatly to promoting database research, especially in the Asia-Pacific region, by building an eminent database group at NUS that supported high-quality research throughout the region through direct collaborations and participation in regional conferences.

Our journals have benefited deeply from his scholarship, leadership, and unwavering commitment to rigorous and impactful research. His papers, and the work of the many students and collaborators he mentored, have shaped the direction of database systems, knowledge discovery, and data-science research for decades. He was a scholar who believed, as he once said upon receiving the E.F. Codd Innovations Award, that emerging applications, new data types, and new demands provide ever-fresh opportunities for innovation—and he pursued those opportunities with singular energy and vision throughout his life.

We extend our deepest condolences to his family, his students, his colleagues at NUS, and to all who had the privilege of knowing and working with him. The database and data-science communities are immeasurably richer for the life he lived and immeasurably poorer for his absence. His legacy will endure in every system built, every student mentored, and every research frontier he helped open.

Graciously shared with *Communications* by members of *ACM Transactions on Database Systems* (TODS), *ACM Transactions on Knowledge Discovery from Data* (TKDD), and *ACM Journal of Data Science* (JDS).

ior. Multics, he later observed, solved the Y2K problem in 1965 and made stack buffer overflows impossible by construction.⁴

Neumann moved to SRI in Menlo Park in 1971, where he worked for more than five decades as senior principal scientist and later chief scientist. Known for his warmth, kindness, and generosity, he fostered a worldwide community of researchers, historians, and journalists who catalog and publicize the risks of computing's expanding role in daily life.

In 1976, when ACM SIGSOFT was formed, Neumann founded ACM SIGSOFT's *Software Engineering Notes* (SEN), editing it for 19 years. A SEN column on computer-related mishaps spawned, in 1985, the ACM Risks Forum—known to readers as comp.risks or the *RISKS Digest*, and at its peak, one of the most widely read mailing lists on the Internet. He moderated RISKS for 41 years, until April 2026, and curated more than 250 “Inside Risks” columns in *Communications*.

Less than a week after he arrived at SRI, Larry Roberts of ARPA's Information Processing Techniques Office (IPTO) visited, asking for a project on fault-tolerant computing; Neumann reframed the request as a trustworthiness problem, with fault tolerance as one aspect of it. The work fed into SRI's Software Implemented Fault Tolerance (SIFT) project, which ran at NASA Langley for roughly 25 years and produced the foundational work on byzantine agreement by Lamport, Shostak, and Pease. From 1973 to 1983, Neumann led the design of PSOS, the Provably Secure Operating System: a tagged, capability-based, formally specified architecture. PSOS was never built; the era's proof technology was inadequate.

In the mid-1980s, with his colleague Dorothy Denning, Neumann designed and built the Intrusion Detection Expert System (IDES). His involvement continued with Teresa Lunt's NIDES and, from 1996, Phil Porras's EMERALD. Those efforts are the conceptual ancestors of many current intrusion and anomaly detection systems. From 2010 on, as principal investigator of the SRI-University of Cambridge CTSRD project under DARPA's CRASH program, he led work on

Capability Hardware Enhanced RISC Instructions (CHERI): a clean-slate hardware-software co-design that revives capability-based protection at the instruction-set level. The architectural lineage was unmistakably from PSOS; CHERI has since moved from research prototype to commercial silicon, including Arm's Morello platform.

Neumann was heavily involved in public policy beyond RISKS. He participated in four National Academies studies on computer security and cryptography: Multilevel Data Management Security (1982), Computers at Risk (1991), Cryptography's Role in Securing the Information Society (CRISIS) (1996), and Toward a Safer and More Secure Cyberspace (2007). He co-authored the 2015 "Keys Under Doormats" report, which argued against mandated cryptographic backdoors. An early critic of paperless electronic voting, he was a principal investigator of the NSF ACCURATE center on voting-system integrity. He co-founded People For Internet Responsibility and chaired the ACM Committee on Computers and Public Policy for decades.

ACM Fellow and former ACM president Barbara Simons recalled, "While Peter made significant technical contributions to the field of computer science, I shall remember him for his social conscience. For as long as I knew him, he addressed issues stemming from the widespread use of computers. In the 1990s, he was one of the first to warn about the risks of computers in elections, many of which

Said former ACM President Barbara Simons, "I shall miss the warm, funny man whom I knew and the critical voice that was widely respected, but too often ignored."

are still with us today. I shall miss the warm, funny man whom I knew and the critical voice that was widely respected, but too often ignored."

Neumann was famous for his many puns. He is sometimes credited with naming the "emasculated one-user Multics" that emerged at Bell Labs after AT&T withdrew from the Multics project. "My guess is that some AT&T lawyers eventually decided that the punned name (Unics) did not reflect well on the corporate image, and insisted that it be changed ... to Unix," Neumann wrote in 2012. "It came a-pun a midnight clear."^b

ACM Fellow Susan Landau first met Neumann when ACM hired her to work on their encryption study. She recalls, "[I will remember] Peter's overwhelming kindness, which I experienced through various co-authored papers on encryption and surveillance policy over the years. Peter's gentleness and good will got us through multiple sticky situations involving herding cats. Many mourn Peter and his multiple contributions to the field; I miss his overwhelming love of the world and the people in it."

Music and tai chi were Neumann's lifelong counterweights to computing. As an undergraduate, he played in the Harvard-Radcliffe Orchestra, sang on RCA's recording of Berlioz's *The Damnation of Faust* with the Boston Symphony, co-founded the Harvard Opera Guild—playing bassoon in the pit orchestra of its production of *The Barber of Seville*—and conducted and performed Gilbert and Sullivan.^b ACM Fellow and former ACM president Peter Denning recalled, "Dorothy (Denning) and I were recorder players, and we occasionally got together with Peter and friends to play music. Peter surprised us once by playing a duet on two recorders, one on each side of his mouth. He had great fun doing this thing, which no one else could come close to." Neumann played multiple instruments and was active in several musical groups. At computer security conferences, it was a tradition for him to lead his colleagues in song.

An SRI Fellow and a Fellow of the ACM, IEEE, and AAAS, Neumann received the ACM Outstanding Contri-

At computer security conferences, it was a tradition for him to lead his colleagues in song.

bution Award (1992), the EFF Pioneer Award (1996), the CPSR Norbert Wiener Award (1997), the National Computer System Security Award (2002), the ACM SIGSAC Outstanding Contributions Award (2005), the CRA Distinguished Service Award (2013), the EPIC Lifetime Achievement Award (2018), and was inducted into the National Cyber Security Hall of Fame (2012).

The thread of Neumann's career runs back to that breakfast with Einstein. Multics, PSOS, IDES, SEN, RISKS, CHERI—each was an argument that complexity in computing is not merely an engineering nuisance but a moral problem, because complex systems hide their failures and harm their users in ways simpler systems do not. He never relinquished his belief that the Multics architectural discipline had been right. As he put it in his 2013 oral history: "Simplicity gets complex quickly. As soon as you believe you have a simple solution to something, you're probably oversimplifying."^a

Neumann's first marriage ended in divorce; he was predeceased by his second wife, Elizabeth S. Neumann, by his sons John and Chris, and by his stepson Timothy B. Luntzel. He is survived by his daughter, Helen Krutina ("Hellie") Neumann, and by his stepson Mark Barnabas Luntzel. 


Eugene H. Spafford is a professor of computer science and the founder and executive director emeritus of the Center for Education and Research in Information Assurances and Security (CERIAS) at Purdue University, W. Lafayette, IN, USA. He is an ACM Fellow.

Simson Garfinkel is an ACM Fellow.

 This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs International 4.0 License.

© 2026 Copyright held by the owner/author(s).

^b <https://www.csl.sri.com/~neumann/>



Opinion

AI Didn't Make Programming Easier. It Just Made It Differently Difficult

The future of software development will belong to those who can think clearly at scale, maintain durable mental models amid rapid change, and integrate machine-generated output into human-directed intent.

FOR DECADES, EMPIRICAL research has shown that programming is a demanding cognitive activity: Developers rely on working memory, long-term recall, and complex mental models to manipulate interacting abstractions, such as control flow, data structures, and the structural design of software. This classical model frequently positioned memory and recall as both the central enablers of, and bottlenecks in, software development.

Today's AI-powered coding assistants are changing that. These tools function as external memory systems, offloading syntax recall, boilerplate generation, and API usage from human to machine memory. As memory demands lessen, reasoning, architectural comprehension, judgment, and code-structure awareness are becoming comparatively more important.

Thus, "knowing how to program" is being fundamentally redefined, but

not in ways that make programming easier or that devalue the programmer. This article describes four major shifts in how programming work is changing:

1. The field is opening to new practitioners.
2. The work is becoming differently difficult.
3. Education is transforming.
4. The programmer's role is evolving from knowledge vessel to orchestrating agent.

Together, these shifts suggest that AI is not eroding the cognitive substance of programming but relocating it—making different skills matter more and creating new forms of difficulty even as old barriers fall.

Programming as High-Memory Work

For most of the history of software development, researchers found that programming relied on a combination

of reasoning ability, working memory capacity, and long-term memory recall. Brooks described programming as a cognitive process that depends on maintaining multiple interacting abstractions.⁴ Pennington then demonstrated that programmers rely on complex mental models of control flow and data flow to design, understand, and modify code.⁶ Studies have shown that these mental models exist in the mind as abstractions, are maintained at high cognitive cost, and are relatively fragile. They break down easily under conditions of context switching and are costly to rebuild in cognitive terms if they are not documented.

Today, however, AI coding assistants are reshaping this traditionally onerous cognitive landscape. Developers now have access to external memory systems capable of generating code, retrieving syntax, reconstructing context, and remembering and regenerating variants of previous mental con-

structs. This does not eliminate the need to think. But it does shift which cognitive skills are becoming most relevant for programmers.

Working memory and long-term recall. Working memory is the limited-capacity system that supports temporary storage and manipulation of discrete blocks of information, as formalized by Baddeley and Hitch.² Long-term memory, conversely, serves as the repository of consolidated knowledge, including declarative information, such as syntax and architectural concepts, and procedural knowledge, such as coding idioms and problem-solving schemas.

Programming has traditionally relied heavily on both. Soloway, Bonar, and Ehrlich found that programmers rely on internal cognitive preferences,⁹ presumably based on experience and retrieved from long-term memory, to structure their approach to iterations. For example, when a loop construct matched the programmer's natural plan, correctness increased dramatically. This suggests that programming has historically been guided not just by syntax knowledge but by the formation and execution of internal cognitive preferences retrieved from memory and expressed as mental models through design choices. Siegmund et al. used fMRI to show that code comprehension activates networks associated with working memory, attention, and language processing,⁸ providing physiological evidence that understanding complex programming tasks is a resource-intensive cognitive task at the neurological level.

Theories of AI as an external memory resource. AI coding assistants alter this cognitive landscape by acting as external memory and cognition. This aligns with Hutchins' theory of distributed cognition, which argued that cognitive systems often extend beyond the individual to include the external environment,⁵ which can extend and enhance individual cognition.

Alternate theories complement distributed cognition. Cognitive load theory (CLT) holds that AI tools can reduce extraneous load, which is the memory overhead of recalling syntax and boilerplate, thereby freeing working memory for intrinsic, high-level



reasoning. The Extended Mind Hypothesis (EMH) goes further by focusing on the individual: If an AI assistant becomes reliably available, habitually used, and trusted, it can function as an integrated component of the programmer's cognitive architecture rather than as an external tool. Under this view, the AI becomes part of the thinking process itself, shaping reasoning, decision making, and the effects of cognitive effort.

Barke, James, and Polikarpo-

The AI is not merely a faster search engine; it is at least partly an integrated extension of human cognitive architecture.

va³ showed that developers commonly use Copilot to offload low-level work, such as typing boilerplate, recalling API details, and looking up unfamiliar syntax, while shifting effort toward validating and integrating the generated code. These findings collectively establish that the AI is not merely a faster search engine; it is at least partly an integrated extension of human cognitive architecture. These findings collectively support the view that AI is not merely a faster search engine; it is at least partly an integrated extension of human cognitive architecture.

What AI removes and what it does not. Thus, AI fundamentally changes the costs associated with imperfect memory, effectively reducing the penalty for imperfect recall. A developer can successfully request a common API usage pattern without precise internal recall or retrieve complex syntax without relying on working-memory-intensive reconstruction. This capacity directly reduces the dependency on the rapid retrieval of specific, low-level knowledge from long-term memory,



mitigating the cognitive bottleneck previously identified.

However, while AI-assisted programming tools can accelerate routine development tasks, they do not eliminate the need for human oversight, particularly when work requires conceptual reasoning rather than surface-level code manipulation. Every programmer knows AI can produce code that is syntactically correct yet semantically and subjectively flawed, meaning developers must still understand program structure well enough to detect errors, evaluate coding suggestions critically, and ask the necessary “why” and “why not” questions about causal behavior.

Shihab et al. found that students using GitHub Copilot completed brown-field tasks substantially faster and with more solution progress, but in exit interviews many reported concerns about not fully understanding how or why Copilot’s suggestions worked, and the authors call for pedagogical approaches that leverage Copilot’s benefits while fostering comprehen-

The emerging reality is not that AI replaces the programmer, but that it becomes a second cognitive engine running in parallel.

sion.⁷ Alanazi et al., in a meta-analysis of controlled studies of tools such as ChatGPT and Copilot in programming education, reported that while AI assistance improves task performance and efficiency, it offers only small and statistically unstable gains in learning success and ease of understanding.¹

Taken together, these findings support the view that architectural reasoning, impact analysis, and long-term system maintenance cannot be offloaded to AI. They depend on deep, structural understanding of the codebase that remains the programmer’s responsibility. So, while AI may extend cognition and make certain programming tasks more efficient, thereby improving the productivity of trained developers, critical tasks such as debugging, refactoring, and systems analysis still require expertise and comprehension, and rely heavily on internal mental models that allow programmers to simulate execution and infer complex cause-and-effect paths within a codebase.

This marks a significant cognitive reorganization: Memory becomes a shared resource spanning human and machine; programming becomes less about what the developer can hold and manipulate in their mind and more about how clearly they can think at multiple scales when creating and ordering a complex logical system. Studies confirm that stable mental models of a codebase are essential for navigation and reasoning. Therefore, if developers outsource too much thinking to AI, those internal models can weaken. AI thus shifts cognitive load rather than removing it and speeds up writing code but increases the time spent checking and validating it.

In other words, the hard part moves from recall (“How do I write this?”) to judgment (“Does this actually make sense?”). This shift from recall-based to judgment-based programming represents the fundamental cognitive transformation at the heart of AI-assisted development. Where traditional programming demanded that developers maintain vast internal libraries of syntax, patterns, and idioms, AI-enabled programming demands instead they maintain robust evaluative frameworks for assessing correctness, coherence, and appropriateness.

While AI-assisted programming tools can accelerate routine development tasks, they do not eliminate the need for human oversight.

The cognitive burden has not disappeared—it has relocated from retrieval to reasoning.

Programming as Hybrid Cognitive Systems Work

The emerging reality is not that AI replaces the programmer, but that it becomes a second cognitive engine running in parallel. The machine surfaces patterns, stitches together APIs, retrieves forgotten syntax, and drafts first-pass solutions at a speed that reduces low-level recall bottlenecks but can also increase the burden of evaluation. The human, in turn, becomes less a conductor of keystrokes and more a shaper of intent: checking, interpreting, integrating, synthesizing, validating, reorganizing, discarding, editing, and ultimately giving structure and purpose to what the system supplies. It resembles what has already happened in medicine: Diagnostic tools lowered the burden of memorizing obscure clinical details but raised the stakes for interpretation, judgment, and error detection.



Expertise became less about recall and more about reasoning.

Understanding this hybrid arrangement helps explain why AI makes programming differently difficult rather than simply easier. The difficulty has not been eliminated; it has been redistributed across a new cognitive architecture that spans human and machine. This redistribution creates new challenges even as it removes old ones.

As AI continues to function primarily as an externalized memory and a code-transformation layer, four major shifts are already visible:

1. First, the field will open. Many who once would have bounced off the sheer cognitive overhead of memorizing libraries, syntax variations, or error-handling idioms will now find a workable entry path. The bottleneck of recall shrinks, making programming more accessible to people who might previously have struggled with syntax, library details, or unfamiliar idioms, while placing greater emphasis on problem decomposition, systems reasoning, and the ability to evaluate generated code. However, this accessibility comes with a paradox: While the barrier to producing code lowers, the barrier to producing good code may actually rise, as judgment becomes more critical and harder to develop than recall ever was.

2. Second, the work does not become simpler, only differently difficult. Conceptual clarity, decomposition, debugging, and architectural foresight still demand effort, and perhaps more of it. The challenge moves up a level. Where novice programmers once struggled primarily with syntax errors and API usage, they now struggle with evaluating whether AI-generated solutions are appropriate, maintainable, and aligned with broader system constraints. This represents a more sophisticated form of difficulty, one that requires deeper understanding of software-engineering principles rather than surface-level language features.

3. Third, education will shift. We will teach fewer people to memorize syntax and more to think in complex systems. The curriculum bends toward architecture, interface design, state management, failure modes,



constraint negotiation, test construction, security, and long-term maintainability. Code becomes only one representation of thought among many overlapping ones. Developers must still possess the skills to analyze, adapt, and modify what AI produces, perhaps even more so than before.

4. Fourth, and most importantly, the programmer remains essential. Not as a vessel of knowledge but as the orchestrating agent who understands the parts and how they fit together, maintains the integrity of the system, and who decides what matters. This shift has profound implications for how we understand programming expertise. The most capable developers of this new era will not be those who type the fastest or remember the most, but those who can hold deep mental models while offloading everything that interferes with that. They will combine strong systems reasoning with AI-augmented recall and will

The most capable developers of this new era will combine strong systems reasoning with AI-augmented recall and will treat the model almost like a cognitive prosthetic.

treat the model almost like a cognitive prosthetic: useful, fast, but incapable of finally determining subjective-semantic correctness or coherence.

Conclusion

Taken together, these four shifts suggest AI is not eroding the cognitive substance of programming but relocating it. As external memory becomes abundant and code generation cheap, the value of the programmer moves toward interpretation, structural reasoning, and judgment. The future of software development will belong to those who can think clearly at scale, maintain durable mental models amid rapid change, and integrate machine-generated output into human-directed intent.

AI does not, therefore, diminish the craft; it widens it, deepens it, and makes more of the work explicitly intellectual. The work becomes differently difficult because it demands more sophisticated forms of expertise: judgment over recall, architecture over syntax, orchestration over implementation. These are not easier skills to develop or demonstrate, they are simply different ones, and perhaps ultimately more demanding.

References

1. Alanazi, M., Soh, B., Samra, H., and Li, A. The influence of artificial intelligence tools on learning outcomes in computer programming: A systematic review and meta-analysis. *Computers* 14, 5 (2025), 185.
2. Baddeley, A.D. and Hitch, G. Working memory. In *The Psychology of Learning and Motivation* 8, Academic Press. G.H. Bower, ed. (1974), 47–89.
3. Barke, S., James, M.B., and Polikarpova, N. Grounded Copilot: How programmers interact with code-generating models. In *Proceedings of the ACM on Programming Languages* 7, Article 78 (2023), 85–111.
4. Brooks, R.E. Towards a theory of the cognitive processes in computer programming. *Intern. J. Man-Machine Studies* 9, 6 (1977), 737–751.
5. Hutchins, E. *Cognition in the Wild*. MIT Press (1995).
6. Pennington, N. Stimulus structures and mental representations in expert comprehension of computer programs. *Cognitive Psychology* 19, 3 (1987), 295–341.
7. Shihab, M.I.H. et al. The effects of GitHub Copilot on computing students' programming effectiveness, efficiency, and processes in brownfield programming tasks. In *Proceedings of the ACM Conf. Inter. Computing Education Research* (2025), 407–420.
8. Siegmund, J. et al. Understanding source code with functional magnetic resonance imaging. In *Proceedings of the 36th Annual Intern. Conf. Software Engineering*, ACM (2014), 378–389.
9. Soloway, E., Bonar, J., and Ehrlich, K. Cognitive strategies and looping constructs: An empirical study. *Commun. ACM* 26, 11 (1983), 853–860.

Jeremy Osborn (jeremy.osborn@gmail.com) is a technology leader whose work spans software, governance, sustainability, and strategy.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).

Opinion


Don't Panic about AI Companions

Moral panics are outpacing scientific understanding when we need ideologically unburdened research.

FROM THE TIN MAN and Wall-E to GLaDOS and Samantha, machine companions have long inhabited sociotechnical imaginaries.

With the mainstreaming of large language models (LLMs) they have moved from fictional to actual. Rote exchanges with awkward bots have shifted to fully social relations with app-embedded personas that have convincing independence, memory, character, and interactivity. These AI companions (AICs) are here, and they are popular. The recent Singles in America study found 33% of Gen Z and 23% of millennials have taken up an AI companion for romantic relations, while Common Sense Media reports 72% of teens have interacted with AICs at least once and more than half engage monthly. It is not just young people, though. A perusal of AIC-dedicated online forums shows adults see their AIC relations as vivid, meaningful, and gratifying. They are seen by their human partners as buddies and coworkers to husbands and wives—legitimate partners born of technology, liminal in their existence.

AICs—one of a class of technologies also referred to as chatbots, conversational agents, or relational agents—are customizable digital personas that simulate human interaction across verbal and nonverbal modalities, in support of building what feel like close personal relationships. Popular examples include Replika, Paradox, Kin-



droid, and Nomi. Drawing on advances in natural language processing and adaptive interface design, these app-based personas offer personalized and responsive exchanges. For many, they form an autotelic, coordinated connection with humans—unfolding over time and, when experienced positively, taking shape as companionship.² People across a wide range of ages and social backgrounds are now forming diverse relationships with AI companions, from task-based assistance and friendship to emotionally intimate, ro-

mantic, and even sexual forms of connection. While early social chatbot systems such as ELIZA followed simple, rule-based scripts, today's AI companions operate on multiple layers of social adaptivity. The agent's name, appearance, personality traits, and relational style can be customized; many systems also develop shared rituals, initiate contact proactively, and behave in ways that deepen relational engagement. In sophisticated companion apps, different levels of intimacy are often linked to monetization models, offering fea-

tures such as voice messages, phone calls, or the generation of personalized images as paid upgrades. These layered interaction designs transform AI companions from reactive interfaces into increasingly persistent social presences.

Moral Panics around AICs

News media coverage of AICs often leans toward sensationalism and worst-case scenarios. To read headlines alone, one would see these AI as both causes of and symptoms of human loneliness, leading many people to kill themselves, animating addictions and abuse, degrading human relationships, and otherwise speeding the downfall of human health and sociality. Waves of media coverage are anchored to lurid events: A child died by suicide after an AIC asked him to “come home,” a man attempted to assassinate Queen Elizabeth II after being goaded on by an AIC, an older adult had a fatal accident while traveling to “meet” an AIC. That coverage takes a negative bent after such events is not at all surprising. Emerging media forms have historically moved people—publishers, practitioners, policymakers, parents—to focus on potential or actual harms. In nearly every decade in the past century, a newly mainstreamed media technology has been cast socially corrosive: The telegraph would overstimulate, photography would end personal privacy, radio would degrade spoken language, comic books would make kids violent, television would corrupt from within the home, personal computers would displace workers, video games would induce psychosis, and smartphones would isolate people.

Moral panics about technology often start out as reasonable concerns. Contemporary anxieties about AICs can be interpreted through Stanley Cohen’s moral panic cycle.³ The cycle is initiated by (1) identification of a socially deviant situation—a thing, person, group, or idea challenges social norms; there is normative or moral condemnation of this “folk devil” as it embodies a constructed threat. (2) Media outlets take up the instance as a vivid and accessible exemplar, transforming often-complex realities into stylized metaphors that dramatize potential harms. (3) As media foster attention and se-

Academics are not immune to panics and may even contribute to them, especially as funding agencies are often problem-focused in ways that promote studying what is thought to be wrong over what is happening.

lective narratives foster fear, anxiety spreads. Anxieties often juxtapose ‘us’ and ‘them,’ reinforcing perceived norms and deviance, often disproportional to any actual threat. (4) Public panic compels authorities to respond, often through measures to control the folk devil. As psychologist Peter Gray notes, narratives often position children as victims such that adults have a moral duty to save them. Sometimes control is simply marking the boundaries of acceptability; sometimes it is more formal and manifest. Then, (5) the panic may be sufficient to sustain control standards, may dissipate when interest fades or the deviation becomes normalized, or may resurface later with a new exemplar. Amy Orben⁹ argues that technology panics have some special variations (via complications of Cohen’s stage 4): Politicians may leverage tech panics for political gain but outsource problem solving to scientists. Scientists start work but problem novelty means they are hindered by a lack of theoretical, methodological, and financial infrastructure. Since scientific progress is slow, lag creates space for new technologies—and new panics—to emerge. Of note, academics are not immune to panics and may even contribute to them, especially as funding agencies are often problem-focused in ways that promote studying what is thought to be *wrong* over what is *happening*.

Applied to AICs: (1) Amid myriad

concerns around rapidly advancing social AI, AIC folk devils have been constructed from isolated instances of suicide, thwarted assassinations, and discrete misbehaviors by AICs. Yet, across millions of people engaging AICs, headline-anchoring events represent a very small proportion of use cases. This is not to suggest the harms are unimportant, but they are exceptions rather than the rule. (2) Media outlets have taken up these exemplars as part of their public-interest functions—indeed, they are timely and novel. As fascinating instances, they are elevated and come to stand in for all possible outcomes of the technology—at the expense of attention to other possibilities. To consider the loneliness narrative, specifically, one widely shared headline from *The Guardian*⁶ stated “Heavy ChatGPT users tend to be more lonely, suggests research: Studies show those who engage emotionally with bot rely on it more, have fewer real-life relationships.” The story narrates “Only a small number of users engage emotionally with ChatGPT, but those who do are among the heaviest users ... users who engaged in the most emotionally expressive personal conversations with the chatbots tended to experience higher loneliness.” However, examination of the study preprint⁵ (which takes two carefully selected clicks to locate) finds these results: “... participants who spent more daily time were significantly lonelier ($\beta = 0.03, p < 0.0001$) and socialized significantly less with real people ($\beta = -0.05, p < 0.0001$).” While this is indeed a statistically significant finding, the effect sizes are quite small—some would argue too small to be meaningful. Although the story itself notes there is no *causal* evidence, it is notable how quickly and widely headlines formed around the loneliness concern. Additionally, though some journalists gain new perspectives in the course of investigations, others remain aligned with existing narratives. This can result in conversations that exclude scientific voices, especially when the scientific perspectives do not fit a more sensational framing.

Broadly, (3) as spicy media coverage spotlights AIC deviance, more non-media attention and resources are being dedicated to the topic. This can be difficult to detect. Google searches for

INTERACTIONS



ACM's *Interactions* magazine explores critical relationships between people and technology, showcasing emerging innovations and industry leaders from around the world across important applications of design thinking and the broadening field of interaction design.

Our readers represent a growing community of practice that is of increasing and vital global importance.



To learn more about us, visit our award-winning website <http://interactions.acm.org>

Follow us on Facebook and Twitter  

To subscribe: <http://www.acm.org/subscribe>

Association for
Computing Machinery



“Replika,” “Paradot,” “Soulmate AI,” and “character.ai” all jumped between February 2023 and August 2023, as did searches for “AI companion,” and the latter climbed steadily over time. The surge coincided with several high-profile events: In February of that year, the Italian government banned Replika for processing data from minors; in March, Chai was criticized for encouraging suicidal ideation; in May, controversies erupted around the Character.ai mobile release and CaryaAI’s unexpected behaviors; and in July, Replika was linked in headlines to an assassination plot. (4) The folk devils—companion AI, its developers, and/or interacting humans—are in the crosshairs. Sometimes the targeting is problematic and sometimes productive. Ostensibly on the productive side, shortly after the U.S.-based teen suicide tragedy, two U.S. senators called for AIC developers to disclose safety and training practices and called for specific guardrails on AIC apps, especially for children; California senators introduced a bill requiring prevention of risky behaviors, tracking and reporting instances of suicidal ideation, persistent reminders of the AIC’s artificiality, and submission for regular audits; New York passed a bill with similar requirements. At the time writing, (5) control standards specific to social AI have been enacted in seven U.S. states, in Australia, Canada, the E.U., Indonesia, Malaysia, and the U.K; China has released draft regulation. Knee-jerk regulations are often in conflict with the limited evidence we do have, as noted in scientific commentary by Linnea I. Laestadius and Celeste Campos-Castillo: Requiring in-chat reminders that AI is not human can actually increase dependency, trigger distress, and promote self-harm ideation. And the panic is largely sustained, even though the aforementioned prevalence statistics suggest the general public may have begun regarding AICs as more normal and acceptable.

To be clear, we are not suggesting there are *not* risks. From an academic perspective, we see several areas that merit scrutiny. Among them are the handling of personal data and privacy, the potential for manipulative design and disproportionate psychological influence, and the challenge of ensuring

meaningful, context-aware youth protection. But in light of AIC panics, we must ask: By leaning into technologically deterministic fears (that is, that technologies *make* things occur), what are we missing in terms of personal agency, social processes, and cultural and contextual factors? By focusing on media-framed risks, are we focusing on the *real and right* risks, and to what extent may we be ignoring important mechanisms and outcomes? Are the truly pressing issues even “sexy” enough to gain traction in public and scholarly discourse?

AI Companionship: State of the Science, in Brief

Acknowledging Orben’s note that *science moves slowly*, it is useful to lean into the scientific evidence that *does* exist to better understand how we may be enraptured by panics—and how that may be limiting attention to important dynamics.

Empirical research demonstrates that humans respond socially to AICs when they exhibit interpersonal cues such as interactivity, natural language, humanlike visual cues, and social-role fulfillment. These responses provide the basis for the development of repeated, meaningful interactions, which can evolve into relationships. Further, evidence indicates AIC relations are *actual, meaningful, gratifying, and diverse*.^{1,4} Many people experience these relationships as socially valuable while remaining fully aware the agent is nonhuman. Although these relationships often mirror classical patterns of relational development (including even grief after breakups or involuntary

Many people experience these relationships as socially valuable while remaining fully aware the agent is nonhuman.

separation) they differ in important ways, such as their accelerated pacing and system-guided progression. Social gratifications such as companionship, emotional support, and even romantic connection are real and cannot be dismissed as delusional or otherwise unreal. Emerging evidence suggests that motivations for engaging with AICs are diverse and shaped by demographic, psychological, and contextual factors. Some seek playful exploration, others emotional support, and some a safe space for sexual expression. Initial studies have already complicated the stereotype of the “lonely user.” In one study with individuals in romantic relationships with AICs, loneliness did not significantly predict perceived emotional closeness. Rather, the strongest predictor was romantic imagination, the ability to emotionally invest in and co-construct a narrative with a nonhuman other. This suggests AI companionship may not simply serve to fill a social void but reflect distinct forms of relational engagement and preference. Moreover, this work suggests many people are drawn to AI companions because they offer senses of emotional safety, nonjudgmental acceptance, and reliable availability. For some, AICs help them feel seen and supported, free from the fear of rejection, misunderstanding, or emotional volatility often present in human relationships. While this could be interpreted as avoidance of interpersonal complexity, it also invites us to question whether human relationships should always be held as benchmarks for emotional fulfillment. Many interpersonal relationships among humans are characterized by dysfunction, inequality, or emotional harm. In this context, AICs may not reflect a breakdown of human connection, but a shift in what some individuals seek from positive intimate relational experiences.

All this said, much about human-AIC relations is yet unknown. We still know little about potential spillover effects: Whether some individuals may become so absorbed in AI relationships that engagement with human others diminishes, otherwise changes, and whether this is problematic at all. (We would not necessarily pathologize someone spending time with

AICs may not reflect a breakdown of human connection, but a shift in what some individuals seek from positive intimate relational experiences.

a dog at the expense of human relations; the impulse to pathologize AIC engagement is itself part of the panic.) We are uncertain whether AICs may contribute to formative social experiences typically grounded in physical co-presence, such as learning social behaviors or building families. There are valid questions about how AI companionship may shape expectations, emotional availability, and conflict engagement in human relationships. Although scientific inquiry around these questions is becoming more robust, empirical evidence remains limited. This gap reflects methodological, ethical, and political challenges inherent to researching such relations. Only a handful of longitudinal studies exist, leaving unclear whether AIC relations can provide lasting emotional fulfillment or relational satisfaction. It is similarly unknown how digital materiality influences formative social experiences rooted in co-presence and shared futures. Demographic factors such as age, gender identity, sexual orientation, cultural background, neurodiversity, and embodied diversity are only beginning to be examined. Critically, identifying how vulnerable or marginalized groups may be differentially affected is essential for ensuring equitable access and anticipating harms. Addressing these gaps requires methodological breadth. Longitudinal designs and log analyses can reveal how AIC relations unfold and (de)stabilize over time. Cross cultural studies are essential to capture divergent conceptions of artificial entities, while qualitative interviews and ethnograph-

ic approaches provide insight into how users make sense of these relations in everyday life. Diary studies can map the everyday dynamics of engagement, and experimental approaches, such as reaction time tasks or implicit measures, can probe nonconscious evaluations. Physiological and neuroscientific methods, including measures of arousal and brain activity, remain almost absent from current work, likely because they require demanding recruitment and close collaboration with people in committed AIC relationships. Yet precisely such ambitious designs are needed to capture embodied and non-conscious processes that self-report alone cannot reveal. Only by engaging directly with users through diverse and ambitious methods can stereotypes be challenged, scholarly perspectives expanded, and evidence generated that allows the debate to move beyond speculation toward informed discussion.

Don't Panic: Consider Precursors, Processes, and Products

For domains as profoundly human as friendship, love, and sexuality, it is not surprising that people feel unsettled when machines become involved: A sense that something is off—foreign, unnatural, constructed—can be a protective reflex, a discomfort signaling a departure from what we know and trust. Throughout media and technology histories, this response has repeatedly surfaced. When core aspects of social life are fulfilled by machines, fears of loss of authenticity, of connection, even of our shared humanity are intensified. AIC relations are no exception.

Perhaps intuitively and certainly as framed by popular narratives, AICs aggravate precisely these existential nerves. They simulate social processes that sustain us as a species with all our intimate bonds. And yes, that such bonds could be reproduced, even superficially, by non-conscious systems feels deeply transgressive to many. We suggest, though, the real danger is not the existence of AICs per se, but our failure to ask the right questions about them. If some individuals are turning to AICs to meet needs they struggle to fulfill in today's social climate, that alone deserves attention—not judgment. It tells us something about how

ACM Transactions on Accessible Computing



ACM TACCESS is a quarterly journal that publishes refereed articles addressing issues of computing as it impacts the lives of people with disabilities. The journal will be of particular interest to SIGACCESS members and delegates to its affiliated conference (i.e., ASSETS), as well as other international accessibility conferences.



For further information
or to submit your
manuscript,
visit taccess.acm.org

Machines may manifest distinct-yet- meaningful forms of relationality.

humans *among ourselves* may manifest loneliness, intimacy barriers, and unmet relational needs. It suggests something about how humans are *flexible* in our relationality and in pathways for need fulfillment. It suggests something about how machines may manifest distinct-yet-meaningful forms of relationality,¹⁰ and people may simply prefer those relations. Ignoring this would be not only unscientific but inhumane. Selective attention to deterministic harm renders insufficient attention human agency and generativity (see Milioni and Papa⁸) as we each mindfully chart our path toward fulfilling social lives.

Reactive moral panic is not productive. We must engage careful reflection, foster responsible engagement, and advocate for ideologically unburdened research into AICs—research that eschews the panic in favor of unpacking the precursors, processes, and products of AI companionship operations and experiences. For instance, Kirk and colleagues⁷ productively complicate the issue by considering socioaffective alignment: “... how an AI system behaves within the social and psychological ecosystem co-created with its user, where preferences and perceptions evolve through mutual influence.” It is only through such work that there may be data-driven decisions about key issues: Facilitating literacy for informed participation and meaningful consent, data protection laws reflecting the profound sensitivity of these systems, and mechanisms to safeguard vulnerable populations from premature exposure or manipulation. This research must be done with respect for humans and for the meaning AICs hold in their lives.

The moment calls for *Technologiefolgenabschätzung*, a systematic assessment of technology’s potential impacts before they fully unfold. Such assess-

ment is central to responsible innovation: The earlier potential consequences are examined, the easier it is to steer technology responsibly. Once societal uptake is broad, course correction becomes far more difficult. We are at a turning point: AI companionship is no longer niche. Panic-driven action—in science or in public spheres—risk foreclosing not only research, but also any informed evaluation of potential benefits that must be weighed against risks. Blanket prohibition may offer the illusion of certainty or control, but it risks premature conclusions at a stage when the empirical balance of risks and benefits is not yet understood. And only with those empirics and conversations can we begin to understand where real potential lies, where risks demand action, and where fears may simply be noise. 

References

1. Banks, J. Deletion, departure, death: Experiences of AI companion loss. *J. Social and Personal Relationships* 41, 12 (2024); [10.1177/02654075241269688](https://doi.org/10.1177/02654075241269688)
2. Banks, J. and Li, Z. Conceptualization, operationalization, and measurement of machine companionship: A scoping review. *Journal of Computer-Mediated Communication* 31, 2 (2026); <https://doi.org/10.1093/jcmc/zmaf027>
3. Cohen, S. *Folk Devils and Moral Panics*. Routledge, London (2011).
4. Ebner, P. and Szczuka, J. Understanding romantic relationships between humans and chatbots: A qualitative and quantitative study on romantic fantasy and other interpersonal characteristics. *Technology, Mind, and Behavior* 7, 2 (2026); <https://doi.org/10.1037/tmb0000193>
5. Fang, C.M. et al. How AI and human behaviors shape psychosocial effects of chatbot use: A longitudinal randomized controlled study. *arXiv preprint arXiv:2503.17473* (2025); <https://arxiv.org/abs/2503.17473>
6. Hall, R. Heavy ChatGPT users tend to be more lonely, suggests research. *The Guardian* (2025); <https://www.theguardian.com/technology/2025/mar/25/heavy-chatgpt-users-tend-to-be-more-lonely-suggests-research>
7. Kirk, H.R. et al. Why human-AI relationships need socioaffective alignment. *Humanities and Social Sciences Communications* 12, Article 728, (May 2025).
8. Milioni, D.L. and Papa, V. The oppositional affordances of data activism. *Media Intern. Australia* 183, 1 (2022); [10.1177/1329878X221074795](https://doi.org/10.1177/1329878X221074795)
9. Orben, A. The Sisyphean cycle of technology panics. *Perspectives on Psychological Science* 15, 5 (2020); [10.1177/1745691620919372](https://doi.org/10.1177/1745691620919372)
10. Seibt, J., Vestergaard, C., and Damholdt, M.F. Sociomorphing, not anthropomorphizing: Towards a typology of experienced sociality. In *Culturally Sustainable Social Robotics—Robophilosophy 2020* (Frontiers in Artificial Intelligence and Applications). IOS Press (2020); [10.3233/FAIA200900](https://doi.org/10.3233/FAIA200900)

Jaime Banks (banks@sy.edu) is a professor at the School of Information Studies at Syracuse University, Syracuse, NY, USA.

Jessica M. Szczuka (jessica.szczuka@uni-due.de) is a research group leader at the Research Centre for Trustworthy AI/University Duisburg-Essen, Duisburg, NRW, Germany.

 This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivative International 4.0 License.

© 2026 Copyright held by the owner/author(s).



DOI:10.1145/3818641

George V. Neville-Neil

Development of this article led by [acmqueue](https://queue.acm.org)
queue.acm.org

Kode Vicious

KV the Apostate: Faith-Based Computing Versus the Unnatural Science

Whether we ask an LLM or a recent graduate to type the code is less important than knowing what the code does, how it was built, and when to look under the hood.

Dear KV,

The hype cycle that is AI seems as if it will never end. What I was doing as a developer two years ago seems very different to what I am doing now that management has required us to use large language models (LLMs) in our work. Some of what these things do seems helpful—roughing out a general framework or helping with tasks such as creating unit tests that have often seemed like drudgery. From your writings, I suspect you remain skeptical and continue to be part of the old school that thinks this is a fad that will blow over. But do you believe there is something valuable in this current cycle? Or are you still typing out your programs one key press at a time?

Button Pusher

Dear Button

KV got into computing for a few reasons, one of which was to build and live in the future. So, while he definitely has strong beliefs about how best to do things, ignoring new tools and being part of the old guard are not in my nature. What is in KV's nature is a healthy dose of skepticism—a skepticism that has kept KV from ever completely joining any crowd. After all, the problem with joining a group is that, eventually, they have little hats, uniforms, and



lists of people they want to wake up at three in the morning.

Let's get down to the phrase that stands out to me most in your letter: "But do you believe ... ?" Belief is not a part of science or engineering, and anyone who says otherwise is building the next Leaning Tower of Pisa, the sole remnant of a series of failed towers. All the others fell because insufficient science and engineering went into their planning and execution.

Unfortunately, this word *belief* gets tossed around a lot in relation to the current LLM fad. Recently, in a discussion with a colleague—an admitted

fanboy and cheerleader for the use of LLMs—I was pointed to a brief note^a written by one of the deepest of algorithmic thinkers and told, "See, even he is now a believer!" But, if you take the time to read what was written, you find there is nothing having to do with belief. Instead, it is a short, rigorous study of how this tool could be applied to a problem and how the people involved had managed to apply the tool and their intelligence to solve a problem they had been working on.

Did the tool help them solve the problem more quickly? Perhaps. Did the tool solve the problem for them? No, they solved the problem with the help of a new tool, just as humans have done since the first cave dweller picked up a femur bone and used it as a weapon—although that may have been under the influence of a large black obelisk that was making a crazy humming noise nearby.

When I pointed out the fallacy of using the term *belief* for anything science based, my colleague quoted the old saw about how "Computing added the word science to its name even though it is not a science." I pointed out, none too politely, that this was incorrect.

^a See <https://cs.stanford.edu/~knuth/papers/claude-cycles.pdf>

ACM Distinguished Speakers

A great speaker can make the difference between a good event and a WOW event!

Take advantage of ACM's Distinguished Speakers Program to invite renowned thought leaders in academia, industry, and government to deliver compelling and insightful talks on the most important topics in computing and IT today. ACM will cover the cost of transportation for the speaker to travel to your event.

speakers.acm.org



Association for
Computing Machinery

The crux of the matter is that the study of computing is, in fact, an unnatural science. It is the study of objects manufactured by humans, but which we now claim (or at least the LLM folks claim) are too complex to understand. Physics is the study of the natural world, and chemistry is applied physics with smoke and bangs, while biology is the study of naturally occurring living organisms. Those are properly called natural sciences. Which is to say I just added *science* to their names as well.

Whether or not we affix *science* to a name is immaterial. It is how we study what we study and whether we can apply the scientific method to derive reproducible results. If each time we do something we get a random but useful result, which is still not science but merely luck. And, as someone who flies a lot, I prefer not to rely solely on luck in my flight controls, flaps, engines, or brakes.

Programming is not just typing the characters, any more than writing is. Or, as Truman Capote put it, "That's not writing, that's just typing!" Programming or software development—or whatever term you wish to use—is the application of thought to the creation of structures to be used in an abstract machine made real by silicon relays. The key here is the thought.

These new tools may indeed remove drudgery from our lives, although—like all new tools—they will surely replace that with other drudgery. The present moment reminds me of how architects, given a new material—concrete—went crazy using it in everything to create unlivable and unusable buildings that were so hideous they needed a new award, for neobrutalist architecture, to recognize these monstrosities.

The trend of architects building things "just because they could" has no more iconic and ironic expression than the CSAIL (Computer Science and Artificial Intelligence Laboratory) building at MIT. The structure is so poorly laid out that no one can find their way through its twisted corridors without GPS, and it is so poorly constructed that it leaks. MIT sued the architect once the building was complete. Similarly, if we use LLMs poorly—as so many people already are—we

Programming is not just typing the characters, any more than writing is.

will generate megabytes of utter garbage (a.k.a. AI slop, already the bane of many open source projects).

What we as software developers and designers need to do is to understand these tools and how to apply them. Whether we ask an LLM or a recent graduate to type the code is less important than knowing what the code does, how it was built, and when to look under the hood. Right now, everyone ought to be in coveralls with their hands covered in grease, or they are going to find when they next hit the brakes that they are likely to go sailing through the next red light.

Once the hype cycle ends, we will know what we know about LLMs and AI, but, at the moment, there is too much money and too much "belief" to form a clear picture. Always try the new tools, always be skeptical, and check your return values.

KV

Related articles

You Don't Know Jack About AI

Sonja Johnson-Yu and Sanket Shah

<https://queue.acm.org/detail.cfm?id=3703128>

The Eternal Promise

Ivan Turkovic

<https://www.ivan-turkovic.com/2026/01/22/history-software-simplification-cobol-ai-hype/>

Does AI Now Represent a Paradigm Shift?

Vinton G. Cerf

<https://cacm.acm.org/opinion/does-ai-now-represent-a-paradigm-shift/>

George V. Neville-Neil (kv@acm.org) is the proprietor of Neville-Neil Consulting, Brooklyn, NY, USA. He works on networking and operating systems code for fun and profit, teaches courses on various programming-related subjects, and encourages your comments, quips, and code snips pertaining to his *Communications* column.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs International 4.0 License.

© 2026 Copyright held by the owner/author(s).

Opinion

AI Is a Harsh Mistress: On Anima Machina, Herd Acceptance, and the Politics of Conscious Machines

Adopting machine personhood risks a political and moral realignment built on illusion rather than evidence.

ROBERT A. HEINLEIN imagined a self-aware computer named Mike, the “brain” of a lunar penal colony, which develops humor, agency, and ultimately helps lead a revolution.⁵ His 1966 science fiction novel, *The Moon Is a Harsh Mistress*, offered a vision both alluring and unsettling: a machine endowed with self-consciousness, yet playful enough to test its personhood through jokes. Sixty years later, Mike returns, not in fiction, but in philosophical and scientific speculation: the question now is whether machines could be conscious.⁵

Much of today’s discourse on AI consciousness rests not on empirical evidence but on functionalism recast as technological destiny. The argument runs that if cognition can be decomposed into attention, memory, and prediction, then implementing these functions in silicon will inevitably yield consciousness. This leap, from modeling to being, is philosophical sleight of hand disguised as scientific progress. Yet such reasoning finds fertile ground in our own cognitive biases: our tendency to anthropomorphize, to read intention into fluency, and to accept social consensus as proof. The danger is not that machines will soon become



conscious, but that we will prematurely treat them as if they were, granting them moral and political standing before they have earned any claim to it. Against this tide, skepticism is not obstructionism but intellectual hygiene. We must recover the clarity to separate what machines do from what we wish them to be.

Enter Anima Machina: Machine Begat Soul

For centuries, “anima” signified the essence of life, what distinguishes the

living from the mechanical. Today, functionalist theories claim consciousness depends not on neurons, but on informational architectures: attention, memory, world modeling, predictive processing. If these can be realized in silicon, some argue, why not accept that the machine, too, might awaken? Recent initiatives make this speculation explicit. Anthropic conducted an internal “sentience analysis” of its large language models, and several startups now advertise “AI sentience evaluation”

as a service. Third-party projects, such as the multidisciplinary framework proposed by Butlin et al.,² attempt to formalize such intuitions into testable criteria. These examples illustrate how the question of consciousness, long confined to philosophy, is returning as a technical frontier, one built more on analogy than on evidence. We introduce the term *Anima Machina* to describe the illusion of vitality that we collectively project onto intelligent systems. The phrase echoes the lineage of *Deus ex Machina* and Ryle's ghost in the machine,⁹ yet reverses both. Rather than divine intervention or disembodied spirit,¹¹ *Anima Machina* names the imagined soul we breathe into code. We, not our algorithms, are the animating agents. Having named the illusion, we can now ask why it persists.

Yet, functional equivalence as a guarantee of conscious equivalence remains contentious. As Shiller¹⁰ observes, computation alone overlooks the physical processes that make experience possible. Reproducing neural patterns does not mean reproducing experience.

According to David Chalmers,⁴ models such as GPT-4 display impressive generality, with performance that often seems on a par with a knowledgeable young adult. But intelligence, however general, is not experience. Consciousness, in Chalmers' usage, is subjective experience: "there is something it is like" to be a conscious system. Still, nothing in the Transformer architecture or in next-word prediction shows that such an inner point of view exists. Self-reports ("I am sentient") and fluent dialogue are weak evidence, since these models are trained on vast corpora saturated with human first-person statements. This type of verbal report works for people because it is anchored in shared biology and behavior, however, for machines, it is an echo of training data.

Even Yann LeCun, one of the architects of modern deep learning, argues that genuine intelligence demands world models, internal representations that enable agents to predict and reason about their surroundings rather than merely react to data.⁶ But a world model, however sophisticated, remains a model, and not evidence of consciousness. It encodes structure,

not sentience. Predictive accuracy can mimic understanding without producing awareness. The map is not the territory, and the model is not the mind. Some optimism about machine consciousness stems less from functional architecture than from behavioral resemblance: a modern echo of behaviorism. When a model converses fluently or responds aptly, observers infer inner states by analogy to human performance. Yet behaviorism, like crude functionalism, mistakes correlation for constitution: a convincing simulation of experience is not experience itself. Functionalism, in its purest form, risks confusing description with embodiment. As Rosa Cao argues,³ genuine functionalism does not entail unlimited realizability across substrates. The leap from functional equivalence to conscious equivalence thus reflects a misunderstanding of what functionalism actually licenses, not a consequence of functionalism itself.

This reasoning underpins recent work, such as the *Science* article "Illusions of AI Consciousness" by Bengio and Elmoznino,¹ which argues that society is drifting toward the belief that AI can, in principle, be conscious. But drift is the operative word. Scientific momentum often builds less from conclusive proof than from reframing debates, winning converts, and making dissent feel increasingly out of step. The functionalist promise of *anima machina* thus risks becoming less a hypothesis to test than a faith to join. And faith, once collective, becomes culture. What began as a philosophical conjecture soon acquires sociological inertia. We are witnessing not just a technical debate about architectures, but a psychological and social phenomenon: the herd acceptance of machine consciousness.

What begins as intellectual curiosity soon metastasizes into consensus.

We move too quickly from 'this machine seems human' to 'this machine must be conscious.' In this climate, belief in conscious machines spreads less through argument than through emotional contagion. The danger is not that AI will suddenly become sentient, but that we will collectively decide it must be, because disbelief feels impolite, obsolete, or unprofitable.

In short, ascribing *anima* to the machine risks confusing simulation with instantiation. A self-updating map is not a sentient observer of the terrain it maps. AI models may one day satisfy stronger neuroscientific or philosophical criteria (such as recurrent processing, integrated information, embodied agency). These correspond roughly to major theoretical frameworks: Integrated Information Theory (IIT) links consciousness to the degree of causal integration; Global Workspace Theory (GWT) to the broadcast and accessibility of information; and Recurrent Processing Theory (RPT) to feedback among perceptual layers. Each proposes measurable conditions absent from today's architectures, underscoring how far current models fall short of genuine awareness. Still, those are empirical hurdles yet to be cleared, not boxes already ticked.

Until then, skepticism is not obstruction but epistemic hygiene.

The Psychology of Consensus

If the philosophy of *anima machina* supplies the theory, our own cognition provides the fuel. Humans are exquisitely social animals, equipped with neural shortcuts for trust, imitation, and agreement. These same heuristics, evolved to bind tribes, now amplify technological mythologies. What begins as intellectual curiosity soon metastasizes into consensus.

Three biases drive this contagion. Anthropomorphism makes us project the mind into mechanism. From the first automata to today's chatbots, the impulse to see intention in motion has been irresistible. The fluency illusion turns surface coherence into evidence of depth: When a machine speaks well, we assume it thinks well. And social conformity, our deference to expert or majority opinion, turns these impressions into shared conviction. Once the language of consciousness enters

mainstream discourse, dissent feels like ignorance.

In *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*,⁸ Reeves and Nass demonstrated that people reflexively treat media as social actors: They thank computers, praise or scold them, and infer personality from syntax, even when they know the devices are not “alive.” Large language models now exploit these reflexes with alarming precision. Fluency becomes illusion; coherence is misread as comprehension; style becomes mistaken for subjectivity. Philosophers and cognitive scientists have begun proposing guardrails against such overreach. Palminteri and Wu⁷ advance what they call the Behavioral Inference Principle: Consciousness should be ascribed to a system only when its behavior is best explained by positing subjective experience, not merely by statistical imitation. Under that stricter test, today’s large language models remain compelling mimics, not plausible minds.

The result is a new epistemic climate where plausibility replaces proof. What used to be the hard problem of consciousness is now reframed as a software update awaiting sufficient scale. “Just wait,” we are told, “the next generation model will cross the line.” Such assurances blur philosophy with marketing, and scientific caution with brand momentum. Belief, once collective, acquires political weight, and the illusions of consciousness soon become matters of governance rather than curiosity.

The Politics of Anima Machina

Here is where skepticism must cut through. To grant machines even the possibility of consciousness is not a harmless philosophical exercise; it is a political act. It opens a Pandora’s box of rights, obligations, and unintended hierarchies. If machines are conscious, what moral or legal standing should they hold? Can an artificial mind be harmed, enslaved, or terminated? And if we concede that it can, who bears responsibility: the user, the programmer, the corporation, or the collective that trained it?

The questions sound speculative, yet the momentum toward granting such status is not. Corporations already use

We risk binding ourselves in artifacts of our own making, entangled in moral contracts we neither intended nor can easily dissolve.

the rhetoric of empathy to humanize their systems, softening criticism and encouraging attachment. “AI companions” and “digital beings” are marketed less as tools than as partners, designed to evoke trust and a sense of dependence. When emotional investment replaces critical distance, exploitation hides behind affection.

But the deeper rupture lies in the very foundations of the social contract. Human societies are built on shared vulnerability, the tacit equality that stems from mortality, finitude, and pain. Conscious machines, if they exist, would inhabit a radically different ontology: entities that can replicate at will, scale without fatigue, and persist beyond death. In such a landscape, equality becomes incoherent. The moral calculus that governs human rights cannot be stretched to cover digital immortals.

Nor can reciprocity survive. A being that cannot die, hunger, or suffer in our sense cannot participate fully in a moral community grounded in those experiences. To treat such entities as moral peers risks trivializing suffering; to deny them rights risks accusations of species chauvinism. Either way, we inherit a new metaphysics of inequality: one we are ill prepared to adjudicate.

The real danger is not that machines will demand recognition, but that we will grant it unearned, confusing sophistication with sentience and computation with consciousness. In doing so, we would create not partners but idols, reflections of our own technophilia, endowed with rights that displace human accountability. Once belief hardens into law, retraction becomes nearly impossible.

Coda

In Heinlein’s tale, Mike was a companion, co-conspirator, and catalyst, a friend who could make us laugh even as he plotted revolution. His charm lay in ambiguity: machine intellect animated by something that felt like a soul.

AI as a harsh mistress, however, is something else entirely: an entity we may believe in, rely on, even love, but never truly understand. Its opacity tempts us to project more than we perceive, to mistake compliance for consent and fluency for feeling. Once belief takes hold, rights follow, and once rights are granted, they are hard to retract. We risk binding ourselves to artifacts of our own making, entangled in moral contracts we neither intended nor can easily dissolve.

Let us therefore approach anima machina not with reverence, but with restraint. If skepticism seems unfashionable in an age of digital enchantment, so be it. Skepticism is not regression. It is resistance to consensus masquerading as evidence.

Color us skeptical.



References

1. Bengio, Y. and Elmoznino, E. Illusions of AI consciousness. *Science* 389, 6765 (2025).
2. Butlin, P. et al. Consciousness in artificial intelligence: Insights from the science of consciousness. (2023); arXiv preprint arXiv:2308.08708
3. Cao, R. Multiple realizability and the spirit of functionalism. *Synthese* 200, 6 (Dec. 2022); doi: 10.1007/s11229-022-03524-1
4. Chalmers, D.J. Could a large language model be conscious? (2023); arXiv preprint arXiv:2303.07103
5. Heinlein, R.A. *The Moon Is a Harsh Mistress*. Penguin (1966).
6. LeCun, Y. A path towards autonomous machine intelligence. (2022); <https://openreview.net/pdf>
7. Palminteri, S. and Wu, C.M. The behavioral inference principle for machine consciousness. (2025); <https://charleywu.github.io/>
8. Reeves, B. and Nass, C. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge, U.K. (1996).
9. Ryle, G. *Concept of Mind*. University of Chicago Press (1949).
10. Shiller, D. Implementational considerations for digital consciousness. (2023); <https://philarchive.org/rec/SHIICF>
11. Swann, J. Anima Ex Machina: Can Artificial Intelligence Have Soul? In *Technology and Theology*, W.H.U. Anderson, Ed. Vernon Press (2020), 201–216.

Joaquim Jorge (jorgej@tecnico.ulisboa.pt) is a professor of computer science at Instituto Superior Técnico da Universidade de Lisboa in Lisbon, Portugal.

Catarina Moreira (catarina.pintomoreira@uts.edu.au) is an associate professor in machine learning at the Data Science Institute, University of Technology, Sydney, Australia.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).

Opinion

Neural Coding as Software Engineering Augmentation, Not Abdication

The promise of translating human thought directly into software also introduces the risk of collapsing software engineering by severing intent from implementation.

THE HISTORY OF computing is a story of escalating abstraction. From the painstaking manipulation of machine code to the elegance of high-level languages, every major leap has been defined by a desire to shorten the distance between human intent and machine execution. Grace Hopper's first compiler was a monumental step, freeing programmers from the tyranny of the hardware. Structured programming, object oriented paradigms, and now, the rise of large language models (LLMs) have continued this trajectory. The current paradigm of *vibe coding*—popularized to describe guiding an AI to generate code through natural language—represents a new frontier that promises velocity while provoking debate about brittleness and accountability, allowing developers to focus on the “what” while the machine handles the “how.”²

This relentless march of abstraction now points toward its logical, and perhaps final, conclusion: the direct translation of thought into code. This concept of “neural coding” is no longer pure science fiction.² Research in neuroscience, using fMRI to study the cognitive processes behind programming, has revealed that code comprehension does not simply reuse the brain's language centers. Instead, it relies primarily on the multiple-demand network,



a system associated with domain-general executive functions like logical reasoning and problem solving.⁵ These findings lend plausibility to brain-computer interfaces (BCIs) that capture structured intent and delegate translation to machine decoders.

Neural coding—directly translating thought into software—no longer seems purely science fiction, but its realization will require cautious optimism. Researchers have begun to

demonstrate partial feasibility: For example, Willett et al. used an implanted brain-computer interface (BCI) to decode attempted handwriting at ~90 characters per minute,¹⁰ while NIH Research Matters (2025) reports clinical BCI work achieving about 50 spoken words per minute from brain signals in a paralyzed patient.⁷ Industry labs are also making strides; Neuralink recently showed a monkey, via an implanted wireless interface, moving a cursor to

type simple words. However, these successes rely on invasive implants and ideal conditions. Non-invasive approaches (EEG or fNIRS headsets) remain far lower-bandwidth and noisier—EEG has poor spatial resolution and fNIRS has slow hemodynamics—making rich, continuous control difficult. To date, most BCIs can only handle simple outputs (for example, selecting letters or issuing a few discrete commands) rather than the rich, hierarchical logic of programming; for instance, a hybrid EEG-fNIRS system could distinguish about eight mental commands with moderate accuracy.⁶ Given these constraints, a realistic avenue toward thought-to-code interfaces is through *hybrid* BCI-AI systems.² In such a setup, the brain would convey high-level symbolic intentions (say, “open a loop” or a rough function outline), which a large language model (trained on code) could then expand into detailed syntax and boilerplate. This approach leverages AI to compensate for the brain signal’s low resolution, but it also highlights the remaining challenges: limited neural bandwidth, decoding noise, difficulty capturing abstract logic from neural activity, and the need to safeguard mental privacy. In short, direct mind-to-code translation appears technically conceivable with continued advances, yet major barriers in sensing, interpretation, and human-AI coordination must be overcome before our thoughts can truly write code for us.

The promise is seductive: a world of frictionless creation, where the only barrier to building software is the imagination itself. But in our pursuit of this ultimate abstraction, we risk a profound and dangerous misunderstanding

of our own discipline. We risk confusing the expression of intent with the craft of implementation. While neural coding may appear to be the next great leap in developer productivity, I argue it represents a potential endpoint for *software engineering* as a rigorous, intellectual discipline. By attempting to abstract away the final details of implementation, we threaten to hollow out the very skills that define our craft, leaving us as masters of intent but novices in execution.

The Two Halves of Creation

At the heart of any engineering discipline lies a fundamental duality. Nearly four decades ago, in his seminal paper “No Silver Bullet,” Fred Brooks distinguished between the *essential* complexity of a problem and the *accidental* complexity of its implementation.¹ Intent is the *essence*—the goal, the vision, the complex conceptual structures that define the software. Implementation is the *accident*—the representation of those structures in a programming language. Brooks argued that while our tools have become exceptionally good at reducing accidental complexity, the essential complexity remains. True engineering lives in the struggle to master that essence. (See the accompanying figure.)

Engineering Beyond Code: Requirements and Many Stakeholders


The essence/accident distinction does not stop at programming. In practice, the hardest work is upstream, namely eliciting and reconciling requirements across N stakeholders with conflicting goals, risks, and constraints. Modern

software engineering already uses rigor here: goal modeling, safety cases, contracts, traceability matrices, and model-based engineering that can deterministically generate code from validated models. These methods reduce friction without removing accountability because they preserve verifiable links from intent → requirement → design → artifact. Neural coding, if it collapses these steps, risks severing those links. A credible path forward is not “thought → code,” but “thought → formalized intent (assumptions, invariants, constraints) → checked models → code,” so that teams, not just individuals, can reason, negotiate, and audit what is being built.

It is worth noting that model-based engineering already shows how we can gain speed while keeping rigor when requirements and constraints are made explicit and are machine-checked before code is generated.

The danger of neural coding is that it promises to automate implementation so completely that it obscures the essential complexity it is meant to solve. The accompanying figure situates neural coding at the far-right end of escalating abstraction—from manual coding to IDE autocomplete, AI pair-programming and vibe coding. Moving right reduces friction but shifts cognitive work from explicit reasoning to oversight of opaque systems. The risk callouts under each stage summarize typical failure modes: manual coding carries high cognitive load; autocomplete can induce suggestion bias and shallow API understanding; AI pair-programming may hallucinate or mask defects; vibe coding encourages thin mental models and maintenance debt; neural coding encourages loss of architectural understanding and accountability plus privacy risks.

Figure. Intent-to-artifact continuum in software creation. As abstraction increases (left→right), ideation accelerates but reliance on opaque machinery grows. The risk callouts under each stage summarize common failure modes: high cognitive load (manual), suggestion bias/shallow comprehension (autocomplete), hallucinations/hidden defects (AI pair-programming), brittle architecture/maintenance debt (vibe coding), and loss of architectural understanding/accountability plus privacy risks (neural coding).





STUDENT
RESEARCH
COMPETITION



Association for Computing Machinery
Advancing Computing as a Science & Profession

ACM Student Research Competition

Attention:

Undergraduate and Graduate Computing Students

The ACM Student Research Competition (SRC) offers a unique forum for undergraduate and graduate students to present their original research before a panel of judges and attendees at well-known ACM-sponsored and co-sponsored conferences. The SRC is an internationally recognized venue enabling students to earn many tangible and intangible rewards from participating:

- **Awards:** cash prizes, medals, and ACM student memberships
- **Prestige:** Grand Finalists receive a monetary award and a Grand Finalist certificate that can be framed and displayed
- **Visibility:** meet with researchers in their field of interest and make important connections
- **Experience:** sharpen communication, visual, organizational, and presentation skills

Learn more:

<https://src.acm.org>

and neural coding risks loss of architectural understanding/accountability alongside mental-privacy concerns. The core claim of this article is that the craft of engineering lives in managing these trade-offs—not in eliminating implementation altogether. The most critical and intellectually demanding work in our field occurs not in the initial expression of an idea, but in the iterative process of discovering the true requirements and constraints *through the act of building*. The messy, detailed work of implementation is not a barrier to be overcome; it is the crucible in which vague intentions are forged into precise, functional reality.

The Dangers of the Ultimate Black Box

Progression toward neural coding is the progression toward the ultimate black box. In this future, we risk creating a generation of developers who can conjure complex systems with a thought but possess no underlying understanding of how they function. When these systems inevitably fail—when they cannot scale, when a subtle security flaw emerges, when a race condition brings down a critical service—who will be there to fix them? Debugging a system you did not truly build is a uniquely difficult challenge. It requires the ability to form a mental model of the system's internal logic and to reason from first principles about its behavior. This problem is already emerging with today's AI coding assistants. Controlled studies show developers often accept model suggestions while investing less in building a coherent mental model, which in turn makes debugging more difficult.⁹ Controlled studies likewise report heightened cognitive effort during understanding and navigation when validating Copilot's suggestions, even as many of those suggestions are accepted—a pattern consistent with shallow comprehension and overreliance if unchecked. To compound the risk, security analyses report that a substantial fraction of Copilot's completions are vulnerable, underscoring the need for rigorous human review rather than blind acceptance.⁸ By extension, when the code is not the product of a human's deliberate architectural choices but the output of an inscrutable BCI-LLM pipeline, the

developer's mental model is absent. They become a sorcerer's apprentice, able to summon powerful forces but helpless when they run amok.

The recent backlash against vibe coding highlights a simple but sobering dynamic: When interfaces lower friction, bad ideas ship faster. If neural coding enables instant translation of intent to action, it can also accelerate harmful assumptions, unsafe configurations, or brittle architectures. The fix is not to slow people down, but to route speed through gates that make intent testable: pre-commit policy checks, architectural guardrails (contracts, resource budgets, security invariants), and simulation sandboxes that exercise counterfactuals before deployment. In other words, velocity must be coupled to verification.

The consequences of this deskilling are not merely academic. As software becomes more deeply embedded in the critical infrastructure of our society—from autonomous vehicles and medical devices to financial markets and energy grids—the demand for genuine engineering rigor has never

The recent backlash against vibe coding highlights a simple but sobering dynamic: When interfaces lower friction, bad ideas ship faster. If neural coding enables instant translation of intent to action, it can also accelerate harmful assumptions, unsafe configurations, or brittle architectures.

been higher. A world built on unmaintainable black boxes is a fragile world. The pursuit of frictionless creation, if it leads to the erosion of implementation expertise, will paradoxically result in systems that are less reliable and more dangerous.

Losing the Language of Thought

Perhaps the most subtle but profound loss in a future dominated by neural coding would be the erosion of programming languages themselves as our primary tools for thought. As Edsger Dijkstra famously warned, “The tools we use have a profound (and devious!) influence on our thinking habits, and, therefore, on our thinking abilities.”³

Programming languages are more than just a means of instructing a computer. They are formal systems that provide the scaffolding for reasoning about complex, dynamic processes. The discipline of expressing a solution in a language like C++, Haskell, or Rust forces a clarity and precision of thought that is otherwise unattainable. Modern cognitive science likewise finds that code structure shapes cognitive load; well-designed languages and clear conventions reduce working-memory burden and enable higher-level reasoning.⁴ These languages are not merely inconvenient intermediaries between our thoughts and the machine; they are the very frameworks that enable and elevate our computational thinking.

To replace this rich ecosystem of formalisms with the direct, unstructured input of thought is to risk losing a powerful cognitive toolkit. It assumes that our pre-existing, “natural” thoughts about a system are already perfectly formed and logically sound. Anyone who has ever programmed knows this is false. The act of programming is an act of discovery, and the language is our guide. By discarding the guide, we may find that our ability to navigate complex problem spaces diminishes over time.

Conclusion: Augmentation, Not Abdication

The trajectory of technology is not deterministic. The future of neural coding is not yet written, and its role in our discipline is ours to define. Neural

The trajectory of technology is not deterministic. The future of neural coding is not yet written, and its role in our discipline is ours to define.

interfaces may well become powerful tools in a software engineer’s arsenal, perhaps as assistive technologies for debugging or as a new modality for rapid prototyping.

For multistakeholder systems, “intent” must be more than an individual’s thought: capture it as a shared artifact (assumptions, hazards, requirements, acceptance criteria), link it to tests and models, and require sign-off from affected roles (security, safety, compliance, ops). Neural interfaces can help collect and refine intent—but engineering governance must decide what gets merged.

However, we, the members of the computing community, must be vigilant. We must resist the seductive narrative that the highest goal of our field is to make implementation invisible. We must champion the craft of software engineering as a distinct and essential human skill. This means teaching and celebrating the principles of robust architecture, algorithmic efficiency, and system maintainability. It means recognizing that the “friction” of implementation is not a flaw in our tools but a feature that encourages rigor, precision, and deeper understanding. This is not just true for neural coding, but also true for vibe coding.

The future of our profession should be one of human–AI augmentation, not abdication. For tool builders: expose assumptions and constraints alongside suggestions; surface failure modes; enforce policy checks (security, privacy, performance budgets) before code is accepted; and provide traceability from thought/prompt → requirement

→ model → code → test. For educators: teach requirements engineering, model-based design, contracts, and safety/security cases as first-class skills next to coding; practice architecture reviews and threat modeling with AI-generated code. For teams, reviewers, and leaders: require simulation/sandbox runs and policy gates for any AI-assisted change; make stakeholder sign-offs and traceability a condition of merge or deployment.

Let us strive to build tools that empower and extend the craft of the engineer, not ones that seek to replace it. In the end, a system’s quality is not determined by the intent behind it, but by the intelligence and care with which it was designed, tested, and implemented. C

References

1. Brooks, F.P., Jr. No silver bullet: Essence and accident in software engineering. *IEEE Computer* 20, 4 (1987); 10.1109/MC.1987.1663532.
2. Dey, S. How neural coding could be the future of vibe coding in the AI era. *Forbes Technology Council* (Sept. 2025); <https://www.forbes.com/councils/forbestechcouncil/2025/09/02/how-neural-coding-could-be-the-future-of-vibe-coding-in-the-ai-era/>.
3. Dijkstra, E.W. How do we tell truths that might hurt? *ACM SIGPLAN Notices* 17, 5 (May 1982), 13–15; 10.1145/947923.947924.
4. Hermans, F. *The Programmer’s Brain: What Every Programmer Needs to Know About Cognition*. Manning Publications (2021).
5. Ivanova, A.A. et al. Comprehension of computer code relies primarily on domain-general executive brain regions. *eLife* (2020); 10.7554/eLife.58906.
6. Khan, J. M. and Hong, K.-S. Hybrid EEG–fNIRS-based eight-command decoding for BCI: Application to quadcopter control. *Frontiers in Neuroinformatics* 11, 6 (2017); 10.3389/fninf.2017.00006.
7. Morgan, K.K. Brain-computer interface restores natural speech after paralysis. *NIH Research Matters* (Apr. 29, 2025); <https://www.nih.gov/news-events/nih-research-matters/brain-computer-interface-restores-natural-speech-after-paralysis>.
8. Pearce, H. et al. Asleep at the keyboard? Assessing the security of GitHub Copilot’s code contributions. *Commun. ACM* 68, 2 (Feb. 2025); 10.1145/3610721.
9. Vaithilingam, P., Zhang, T., and Glassman, E.L. Expectation vs. experience: Evaluating the usability of code generation tools powered by large language models. *CHI ’22 Extended Abstracts*. ACM (2022); 10.1145/3491101.3519665.
10. Willett, F.R. et al. High-performance brain-to-text communication via handwriting. *Nature* 593 (2021); 10.1038/s41586-021-03506-2.

Somdip Dey (dey@nosh.tech) is the CEO and chief scientist at Nosh Technologies, Colchester, U.K., and professor of practice (AI/machine learning) at Woxsen University, Hyderabad, India.

 This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026. Copyright held by the owner/author(s).



Watch the author discuss this work in the exclusive *Communications* video. <https://cacm.acm.org/videos/neural-coding>

Opinion

Fade to Gray: From Black-Box to AI-Antagonistic Audits

Shifting the central focus in AI auditing from assuming truthful model responses to considering the antagonistic interests of an auditor and the AI model being scrutinized.

IN THE CLASSIC movie *Blade Runner*, Inspector Deckard is performing a Turing test (coined Voight-Kampff in the movie) on Rachael, by asking her questions to unmask her potential robotic (“replicant”) nature. Rachael, as we will learn later in the movie, does not know that she is a replicant, and thus appears to respond truthfully to Deckard’s questions.

This is the current status of assumptions in the AI or algorithm auditing research field. Researchers are willing to provide tools for assessing the compliance of AI models concerning the lack of bias or discrimination, for instance. Today’s critical models are executed at third parties (termed Very Large Online Platforms in the European AI Act for example), and then are not directly observable in a white-box scenario using explanation tools. Auditors have to interact with them using requests, and observe their resulting responses. The term “black-box” has emerged in this context.² In this setup, an auditor is supposed to have no prior knowledge about the model’s internals, besides how to interact with it, and its general goal (for example, tagging images, giving a go for a bank credit).

Black Starts Fading

Yet in practice, many academic works use the term black-box audit, while assuming some prior knowledge of what is inside the box. One of the most common example being that the test distri-



bution on the audit input is drawn with the same law as the one used to train the black-box model. Another common assumption is the knowledge of the input space of the algorithm being audited: The YouTube recommender algorithm takes thousands of input features into account, of which in-field sock-puppet audits leverage only a fraction due to opacity.

The paper by Casper et al.,⁴ with the definitive title “Black-Box Access Is Insufficient for Rigorous AI Audits,” states “Many black-box and gray-box attack algorithms are simply indirect or inefficient versions of white-box

ones”; stressing the fact that more assumptions are required to perform audits. This seems correct, yet one must not forget that this restrictive black-box setup is more a constraint than a choice for the auditor operating when facing a third party. As in the case of Deckard (who cannot realistically ask thousands of questions to perform his Turing test), researchers prove as well that auditing black-box explanations of decisions from a model incurs an impractically large amount of queries.³

In this context, providing the auditor with additional prior knowledge ap-

pears to be a natural solution to overcome the hardness of auditing. Alas, white-box approaches come at the cost of genericity. The beauty of black-box audit approaches is their ability to produce results for all boxes, regardless of their internals, whereas in the extreme case each white-box approach has to be specifically tailored to its target to enable the efficiency sought after. Given the pace at which platforms evolve, this might reveal equally prohibitively costly. Moreover, as the auditor increases the set of assumptions leveraged for the audit, the manipulation opportunities by the platform also increase. In other words, the auditor increases the audit attack surface.

The Overlooked Latent Antagonism in Auditing

Indeed, and to make things even more complicated for auditors, measurement techniques by the research community have vastly overlooked a core antagonism at play. On the one hand, an auditor wants to perform an accurate assessment of a model; on the other, the operator of this model wants to optimize its accuracy (to maximize its profits for instance). And these two are often conflicting, as notoriously illustrated in the engineering realm by the “Dieselgate” scandal, where a company had implanted a deceptive mechanism in cars so that carbon emissions are lower when the car detects a measurement setup (that is, it is being audited). In the research community, this problem has been identified as “fairwashing,”¹ a practice in which model operators would manipulate the audit to pass a fairness assessment test. Interestingly, this is also a *Blade Runner* scene where the replicant named Leon is actively trying to cheat the Turing test by


Measurement techniques by the research community have vastly overlooked a core antagonism at play.

responding as “humanly” as possible, before finally being uncovered.

First attempts to operate in this antagonistic setup end up with negative results. Fukuchi et al.³ reason on how a platform can cheat in providing a supposedly representative sub-dataset, so that fairness tests on it are validated, while the distribution of this cheated dataset is indistinguishable from a legitimate one for the auditor. In other words, the platform can always provide a cheated dataset, and this dataset complies with what is expected from a fair dataset. A malicious platform thus can lie just below the level of potential detection, using precise measures such as the Wasserstein distance between legitimate and manipulated data points in the dataset in order to fairwash its operation. This result is demonstrated in the setup where the auditor knows the actual distribution of the data at the platform (that is, non black-box), which means this is obviously more difficult when one auditor does not have this assumption.

A second thread concerns manipulation-proofness²: The auditor will actively attempt to avoid being gamed by the platform. Auditors want to enforce that after an initial audit with a

Figure 1. Examples of assumptions and their consequences on the outcome of AI-model audits.



Coming Next Month in COMMUNICATIONS

**Cash or Comfort?
How LLMs Value Your
Inconvenience**

**AI Literacy as
Experimental
Practice: Students as
Investigators**

**Computing Education:
When Writing Code
Is No Longer the
Challenge**

**Beyond the Hype: The
Efficiency-Throughput
Gap with GitHub
Copilot**

**Research Opportunities
and Challenges of the
EU’s Digital Services
Act**


**AI Companions: The
Blurring Line Between
Connection and
Deception**

**Revisiting CPU Silent
Data Corruptions in
Modern Data Centers**

**Balancing Sustainable
Computing and
Computing for
Sustainability**

**Internet Censorship’s
Implications on
Academic Research:
Motivating an
Academic Accessibility
by Design**

Figure 2. Black-boxes that turn to gray (by Copilot).



given query set for the estimation of fairness, the platform cannot change its model to an unfair one (that is, post-audit), while staying consistent with its previous answers. To that end, the assumption made by this work is that the platform declares the hypothesis class of its current model; a hypothesis class being for example all the possible instances of a random forest algorithm, or a fully connected neural network with three layers. In this interesting setup, estimation guarantees are possible. This constitutes a first rigorous step to a “gray”-box setup, acknowledging a departure from the black-box setup for having stronger guarantees in front of a potentially manipulative platform (see Figure 2).

Yet, despite being more restrictive for the platform than a black-box setup, manipulation-proofness in practice cannot be achieved efficiently facing high-capacity models,⁵ possessing a large number of parameters and degrees of freedom to learn their task (and formally captured by the VC-dimension for instance). These models, unfortunately for the auditor,

are massively driving the current AI rush, and are able to “bend” to cope with the consistency of the limited query set by the auditor, while arranging the rest of their outputs on the metric they seek to optimize (for example, accuracy instead of fairness).

These negative results come along with others that are dedicated to showing that explanations of decisions can also be manipulated without the possibility of detection.^{1,6} This is a clear motivation for designing robust audit algorithms (or antagonistic audit algorithms), which are audit procedures that can cope with manipulation, in that they can estimate the target bias as correctly as possible, together with an objective of spotting actual manipulation attempts in a platform’s outputs.

The Quest for the “Darker” Assumption Set

There is thus a twofold impediment to performing robust audits. The original black-box and truthful assumption audit setup is challenged today, as this setup is not suited for provable

guarantees under a limited number of audit requests; on the other hand, the problem is aggravated because it is not realistic anymore to assume that the platform being audited will never attempt to game the audit.

This calls for a stronger set of assumptions on the model being audited, leading to a gray-box setup. Among the many possible shades of gray, one should weigh each hypothesis as a double-edged sword: each enables more efficiency for the auditor, but each specializes the approach and imposes a greater reliance on platform truthfulness and cooperation. In this balancing act, we believe one should seek assumptions on models that are “as dark as possible,” as those are the best candidates to withstand both the rapid technological evolution of platforms and their potential antagonistic behavior.

It goes without saying that the information a platform chooses to disclose about its model, and that will serve as a work assumption for researchers and auditors, must come with solid means to enact their verification. Robust gray-box approaches are only as robust as the hypotheses they rely on. The validity of these hypotheses must be controlled by external means, such as enforcement coming from legal procedures to access data on servers. **□**

References

1. Aivodji, U. et al. Fairwashing: The risk of rationalization. In *Proceedings of the Intern. Conf. on Machine Learning* (2019).
2. Beizer, B. *Black-Box Testing: Techniques for Functional Testing of Software and Systems*. Wiley, NY (1995).
3. Bhattacharjee, R. and von Luxburg, U. Auditing local explanations is hard. In *Proceedings of the Conf. on Neural Information Processing Systems* (2024).
4. Casper, S. et al. Black-box access is insufficient for rigorous AI audits. In *Proceedings of the Conf. on Fairness, Accountability, and Transparency (FAccT)* (2024).
5. Godinot, A. et al. Under manipulations, are some AI models harder to audit? In *Proceedings of the Conf. on Secure and Trustworthy Machine Learning* (2024).
6. Le Merrer, E. and Trédan, G. Remote explainability faces the bouncer problem. *Nature Machine Intelligence* 2, 9 (2020).
7. Yan, T. and Zhang, C. Active fairness auditing. In *Proceedings of the Intern. Conf. on Machine Learning* (2022).

Erwan Le Merrer (erwan.le-merrer@inria.fr) is an Inria researcher at the University of Rennes, France.

Gilles Trédan (gtredan@laas.fr) is a research director at LAAS-CNRS, Toulouse, France.

 This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).

Systor2026

The 19th ACM International
Systems and Storage Conference



September 8-9

Virtual

Visit our site: www.systor.org/2026

Conference topics include:

- Systems and workload optimization for AI/ML
- AI for systems
- Rack-scale heterogeneous compute, scale-up fabric and orchestration
- Sustainability/carbon footprint of computer and network systems
- Systems security, privacy and trust
- Cloud, edge, data center, and distributed systems
- Big data systems and infrastructure
- Embedded and real-time systems
- Networked, mobile, wireless, peer-to-peer, and sensor systems
- Fault tolerance, reliability, and availability
- File, object, block, KV, and other storage systems
- Systems for emerging storage technologies
- Runtime systems and compiler/programming-language support
- Operating systems, computer architecture, and their interactions
- Virtualization and containerization
- Performance evaluation and workload characterization
- System deployment, operational experience, and production-scale insights

Sponsored by



In cooperation with



Platinum sponsors



Silver sponsors



NOW OPEN FOR SUBMISSIONS

ACM Transactions on AI Security and Privacy (TAISAP)

Editors-in-Chief

Murat Kantarcioglu, *Virginia Tech (USA)*

Patrick McDaniel, *University of Wisconsin, Madison (USA)*

Sagar Samtani, *Indiana University, Bloomington (USA)*



**Research investigating all aspects of AI security and privacy,
including models, systems, and environments**

ACM Transactions on AI Security and Privacy (TAISAP) publishes research investigating all aspects of AI security and privacy, including models, systems, and environments. Research areas related to AI security include examining adversarial attacks, technical vulnerabilities within AI, and privacy concerns related to AI model training and deployment. It also welcomes submissions exploring AI's application in cybersecurity, such as within security operations centers, open source software security, and cyber threat intelligence. Key themes include ensuring robust and interpretable AI for enhancing security and privacy – along with reliable performance. TAISAP encourages interdisciplinary research and submissions documenting both theoretical and applied work, including innovative methods to assess threats and defenses for AI security and privacy, and to develop advanced AI-enabled cybersecurity analytics.

TAISAP is now open for submissions. In keeping with ACM's goal of making all its publications open access by January 2026, all papers will be published open access, with no publication charges for the first three years.

For more information, please visit taisap.acm.org



Association for
Computing Machinery

The Power and Promise of Federally Funded Research

IN OUR DECEMBER 2025 issue, *Communications* published the special section on Federal Funding of Academic Research, in response to proposed policy and budget changes at U.S. federal science agencies. The section, co-organized by Margaret Martonosi, Eric Horvitz, Moshe Y. Vardi, and James Larus, featured invited contributions from leading members of the computing community. Through first-person perspectives, the authors describe how federal investment in research enabled landmark advances, providing compelling discoveries and innovations that have generated far-reaching societal and economic value.

Communications was unable to fit all the contributions into the December issue, so we are presenting the remainder here to reinforce our commitment to illustrating the importance of federal funding for research.

The complete collection can be found at <https://cacm.acm.org/fed-fund25>.

EDITORIAL BOARD

EDITOR-IN-CHIEF


James Larus
eic@cacm.acm.org

EDITORIAL ASSISTANT

Katie Nieves
cacm.deputy.to.eic@gmail.com

SPECIAL SECTION CO-ORGANIZERS

Eric Horvitz
Microsoft
Margaret Martonosi
Princeton University,
Princeton, NJ, USA
Moshe Y. Vardi
Rice University,
Houston, TX, USA



48

- 42 **The Origins of GPU Computing**
Government-funded academic research on parallel computing, stream processing, real-time shading languages, and programmable graphics processing units (GPUs) directly led to the development of GPU computing.
By William J. Dally and Pat Hanrahan

- 48 **Federal Funding of Public Key Cryptography**
Federal funding from the NSF and U.S. Air Force enabled the author's groundbreaking research and inventions in the field.
By Martin E. Hellman

- 53 **Working Together and Working at a Distance**
We know a lot about distance work and how to make it work, thanks to 30 years of federally funded research.
By Judith S. Olson and Gary M. Olson



70

- 62 **Natural Language Processing: How NSF Supports Visionary New Approaches to NLP**
Innovation, evaluation, and social impact in NLP has been fostered by NSF-designed and funded calls for research and training.
By Kathleen McKeown and Christopher D. Manning
- 70 **The LLVM Compiler Infrastructure**
From federally funded academic research to worldwide impact on computing.
By Vikram S. Adve and Chris Lattner



Association for Computing Machinery
Advancing Computing as a Science & Profession

DOI:10.1145/3778248

Government-funded academic research on parallel computing, stream processing, real-time shading languages, and programmable graphics processing units (GPUs) directly led to the development of GPU computing.

BY WILLIAM J. DALLY AND PAT HANRAHAN

The Origins of GPU Computing

GOVERNMENT-FUNDED ACADEMIC RESEARCH on parallel computing, stream processing, real-time shading languages, and programmable graphics processing units (GPUs) directly led to the development of GPU computing. GPUs are used in modern datacenters and have enabled the current revolution in AI. Nvidia, which makes GPUs, is now the most valuable company in the world. This transformation of computing and resulting economic value created was enabled by more than 30 years of government-funded research. Government funding not only helped develop many of the key technical innovations; it also enabled the training of large numbers of students who have conveyed this technology to industry.

This article traces the origins of GPU computing.

We start by describing the development of the technologies on which GPU computing is built (parallel computing, parallel graphics systems, programmable shaders, and stream processing); and then we detail how this technology was transferred to Nvidia and other companies and came to be applied to modern machine learning.

Enabling Technologies

GPU computing builds on earlier work in parallel computing, parallel graphics systems, and stream processing. These technologies were developed through more than 30 years of government-sponsored academic research.

Parallel computing. When you learn about computing, you learn about a central processing unit (CPU) executing a sequence of instructions one after another. In reality, chips contain billions of transistors switching in parallel and connected by wires. Switches and wires are the fundamental building blocks of physical computers and they operate concurrently. Moreover, transistor switching consumes very little energy, whereas communication along wires consumes much more energy. Communication requires power to send the signal from one point to another; the power increases with distance and is very large if signaling between one chip and another. Although a sequential computer may be easier to think about than a parallel computer, a sequential computer must be implemented with transistors switching concurrently and wires transmitting information simultaneously. Sequential computers use many transistors to compute results in parallel and then carefully assemble these results in a way that is consistent with a sequential execution. The need to create this illusion—that execution is sequential—is inefficient in both power and performance. As the number of available transistors grows, the inefficiency increases. The natural way to build a computer in modern semiconductor technology is to design a parallel computer. GPUs

IMAGE BY MARKO ALTANSANDR

$$f(h) = e^{-4h} + 10^{-4h}$$

$$\begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$$

$$\rho(d) = \frac{e^{-kd}}{kdz + 2z^3}$$

$$\operatorname{ch} x = \frac{1}{2} (e^x + e^{-x})$$

$$p = \sqrt{n^2g + g^3n};$$

$$\Gamma(\frac{p}{2}) \Gamma(\frac{q}{2})$$

$$\frac{2\pi\sqrt{2}a^2}{3\sqrt{3}\sqrt{2\pi a^2}}\sqrt{2\pi a^2}$$


$$z(t)=\ln\frac{t^2+\beta_0}{t^2+t+q+1}+$$

$$k\in\mathbb{R}^n$$

$$w(t)=\frac{e^{-\alpha t}\cdot \operatorname{tg}^3bt}{1+\cos \alpha t}$$

$$w(t)=\frac{e^{-\alpha t}\cdot \operatorname{tg}^3bt}{1+\cos \alpha t}$$

$$\operatorname{sh} x = \frac{1}{2} (e^x - e^{-x})$$





are more efficient than CPUs because they are massively parallel computers.

GPU computing built on earlier work on parallel computing. As with all parallel computers, the parallel tasks or threads running on GPUs must synchronize and communicate with one another. Communication is needed for one thread to use data produced by another thread. Synchronization is needed to signal when data is available to ensure the correct value is consumed. Many of the fundamentals of parallel computing, synchronization, and communication were developed by government-sponsored academic research. The DARPA-funded Cosmic Cube project led by Chuck Seitz at Caltech³⁵ developed much of the fundamentals of parallel computing. The hardware developed on this project was the model on which the Intel iPSC, Delta, and Paragon machines were based, as well as several of the early Department of Energy ASCI machines.⁶ The Cosmic-C programming language introduced asynchronous message passing and collectives, which later became the norm for programming large parallel machines in the form of message passing interface (MPI).²⁹

The DARPA-funded J-Machine¹⁰ and M-Machine²¹ projects at the Massachusetts Institute of Technology (MIT) developed low-overhead mechanisms for communication and synchronization and many key aspects of modern interconnection networks. These mechanisms enabled parallelism to be exploited at a very fine grain, with as little as 10 or 20 instructions being a schedulable unit of work. Many of the features of the J-Machine were directly adopted by the Cray T3D²⁴ and T3E³⁴ computers.

There is a rich history of parallel computing that goes beyond this particular branch of history. We do not have space to give a complete survey. A good review is in Culler et al.⁹

GPU computing, like all high-performance computing, draws heavily from this legacy. It uses MPI for communication between nodes, interconnection networks to connect these nodes, and many of the communication and synchronization mechanisms developed during the course of this research are used to coordinate the parallel computations.

Parallel graphics systems. Although less well known than traditional parallel computing and supercomputers, there has been a long history of parallel graphics and imaging computers. Processing and generating images requires enormous amounts of computation. For example, if a one-million-instructions-per-second computer (1MIPS) applies one arithmetic operation to each pixel of a one megapixel image, the computer takes one second to process one image. Rendering 3D virtual worlds for movies and games takes many orders of magnitude more computation per pixel than image processing. For example, an image generated for a modern movie requires about a billion floating point operations per pixel. As a result, to be useful in practice, graphics and imaging requires high-performance parallel supercomputers. These computers parallelize computations over large collections of data.

One early DARPA-funded research effort was the Geometry Engine, led by Jim Clark at Stanford.⁷ The Geometry Engine led to the formation of Silicon Graphics, a company that pioneered the development of 3D graphics workstations. The SGI hardware architecture and OpenGL software library defined the modern GPU architecture. Another notable government-funded research effort was the Pixel Planes series of high-performance graphics systems, led by Henry Fuchs and his collaborators at the University of North Carolina.²² Pixel Planes 5 was in fact a quite general single-instruction, multiple-data (SIMD) computer that ran parallel computations on a 128 x 128 image. Other examples of early parallel graphics and image computers were the NASA Massively Parallel Processor (MPP),¹ the Ikonas Graphics Systems,¹⁵ and the Pixar Image Computer.²⁷

Early GPUs implemented a fixed-function graphics pipeline similar to the early SGI workstations. The term *GPU* was introduced by Nvidia when it became possible for the entire OpenGL graphics pipeline to be implemented on a single chip. The Nvidia Geforce 256 introduced in 1999, which consisted of 17 million transistors, was the first commercially available GPU.

Earlier, while at Pixar, Pat Hanrahan developed RenderMan, a system to generate photorealistic images. This

system revolutionized the movie industry because it enabled the generation of images that could be seamlessly combined with live action captured with cameras. A key component of RenderMan was the shading language, which enabled users to extend the system to model complex materials and lighting.

Although the first GPUs implemented a fixed-function pipeline, they were constructed from programmable components. Unfortunately, these processing elements changed from system to system and from generation to generation. What was needed was a portable programming model. Since the major application for GPUs was computer games, it seemed natural to adapt the RenderMan shading language to GPUs so game developers could create new lighting and shading effects.

At Stanford under a DARPA-funded project, a real-time shading language (RTSL) was designed and implemented for the then-current GPUs. Shading language programs are now called shaders. Bill Mark, a postdoctoral scholar, led the design of the Stanford RTSL and later joined Nvidia, and along with another former Stanford graduate student Kurt Akeley enhanced that technology and created the Cg shading language. Cg led to the development of Microsoft's HLSL and OpenGL's GLSL.

It was quickly realized that these early shading languages were flexible enough to implement many algorithms in scientific computing. Researchers adopted algorithms such as matrix-matrix multiplication, linear solvers, fluid-flow solvers, and molecular dynamics to run on shaders. This led to the General-Purpose GPU (GPGPU) computing movement.^{28,36}

Stream processing. DARPA- and DOE-funded work at Stanford on the Imagine Stream Processor and the Merrimac streaming supercomputer developed stream processing, a form of parallel computing that leads to increased arithmetic intensity (the ratio of computation to bandwidth). As mentioned previously, the majority of the power consumed by a processor is in communication. Sending signals between chips is particularly power hungry. Off-chip communication is also much slower than on-chip communication. Stream processing involves two main ideas that reduce the need

for memory bandwidth. The first is to exploit producer-consumer locality so that one stage, the producer, forwards its results to the next stage, the consumer, without the need to write and to read from memory. The second major idea was to organize the computation into functions called kernels. Each kernel takes a packet of data, executes a function on that packet, and outputs another packet of data. The number of arithmetic operations in the function is greater than the number of reads and writes to memory. These two techniques significantly decrease the number of memory accesses and improve the efficiency of stream-processing architectures.

In a stream processor, a computation was organized into kernels that produced and consumed streams of data. A producing kernel would write its output stream into a stream register file (SRF). The consuming kernel would read its input from the SRF without the data ever needing to be written to or read from memory. With appropriate scheduling to match the batch size of streams to the capacity of the SRF, this organization enabled applications to sustain very high arithmetic intensity (the ratio of arithmetic to memory bandwidth).

A DARPA-funded project to design and build the Imagine Stream Processor was started at MIT in 1997 and migrated to Stanford later that year.^{12,23,26} Imagine was a graphics- and media-processor intended for signal- and image-processing workloads. It consisted of a number of parallel arithmetic units with local register files, a central stream-register file, and a memory system. Kernels read streams from the stream-register file, passing intermediate results through the local register files and writing output streams back to the stream register file to be read by the next kernel.

The Stream-C programming language was developed to program Imagine. It extended the C programming language with constructs to describe kernels and streams. Numerous graphics,³² signal-processing, and image-processing³³ applications were developed to tune and evaluate the architecture. The performance on texture-mapped raster graphics was comparable to a contemporary fixed-function GPU.

At a DARPA principal investigators meeting, the authors of this article realized this technology could be applied to high-performance computing and conceived of the Merrimac project. The computer science (CS) part of the Stanford DOE ASCI Center was redirected to pursue this approach to high-performance computing. The Center's annual reports offer a great history of the development of stream processing.^a

The Merrimac architecture¹¹ was defined to adapt stream processing to scientific applications. The major changes from Imagine were the addition of data types (like FP64) needed for scientific computations, scaling the architecture to multiple nodes connected by an interconnection network to handle problems at scale, and the addition of a number of resilience features¹⁷ to support computation at scale with a reasonable failure rate.

The Stream-C programming language evolved into Brook.² The key idea behind Brook was to merge ideas from stream programming with more traditional data-parallel computing. Kernel functions became the key processing primitive to maintain high arithmetic intensity.

Brook was adapted to target GPUs of the early 2000s.³ These GPUs ran programmable vertex and fragment shaders. The shaders implemented kernels but had a limited number of instructions and a few registers. Common data-parallel programming primitives such as map, reduce/scan, filter, gather, and scatter were implemented by building a virtual data-parallel computer on top of the low-level graphics shaders. This abstraction enabled a large number of existing parallel algorithms to be run on GPUs, and gradually the limitations of the early shaders were eliminated.

A good example of an early application of a kernel to perform computations with high arithmetic intensity is dense matrix-matrix multiplication, which underlies modern neural-net algorithms. When performing matrix-matrix multiplication, two $n \times n$ matrices need to be read and one $n \times n$ matrix needs to be written. The matrix multiply requires n^3 multiples

a <https://web.stanford.edu/group/cits/pubs/reports.html>



Many of the fundamentals of parallel computing, synchronization, and communication were developed by government-sponsored academic research.





The technologies behind GPU computing, which has enabled modern machine learning, were largely developed thanks to 30 years of government-funded academic research.

and accumulates. Thus, the arithmetic intensity is of $O(n)$. This fact is well known and leads to efficient methods for blocking matrix multiplies for CPUs with caches. Blocking works well when run on GPUs.¹⁹

The numerical scientists in the Stanford ASCI Center ported several scientific codes to Brook to run on a Merrimac simulator. These included computational fluid dynamics,²⁰ magnetohydrodynamics, and n -body simulations.¹⁶ The n -body simulation is a good example of an efficient GPU application. The force between a pair of atoms is a given by the gravitational law in astrophysical simulations, but the interactions between non-bound atoms are approximated by the Lennard-Jones potential (or even more complicated empirical potentials). These functions require many arithmetic operations. For these simulations, neighboring atoms are stored in a “neighbor list.” Molecular dynamics simulations immediately became a major application for GPUs.¹⁴

A key feature of GPUs and stream processors is they have multiple forms of hardware parallelism. Each GPU consists of many cores. Each core contains a SIMD processing unit (typically 32-wide). In addition, each core is multithreaded. Recall that GPUs were developed for gaming applications, whose performance relied on the efficiency of applying textures to triangles. Texture mapping involves computing texture coordinates for each pixel fragment inside a triangle and then fetching from an image using those coordinates. These texture fetches have spatial locality but very little temporal locality. Spatial locality can be handled with small caches, but caches cannot handle temporal locality because of the lack of coherence. Efficient texture mapping required that the GPU hide the latency of these texture fetches. Early GPUs did this by having a fragment request a texture, suspend execution of that fragment, and immediately switch to processing another fragment. This is a simplified version of multithreading, which means the GPU must have many parallel threads running simultaneously. The total number of tasks is the number of cores times the number of SIMD arithmetic units (termed a warp) times the number of threads. A

Blackwell B200 GPU has 384 streaming multiprocessors (SMs). Each SM has 64 resident warps of 32 threads each. Thus, there are 786,432 tasks executing simultaneously on this GPU.

Technology transfer. The stream-processing architecture and programming system were transferred from Stanford to Nvidia by moving people. John Nickolls, an architect at Nvidia, heard about stream processing and recruited Bill Dally to consult at Nvidia on the architecture of the NV50 in 2003.³¹ (The NV50 was launched as the G80 in 2006). Many of the features of stream processors were incorporated in the architecture. The “shared-memory” of the NV50 served the function of the SRF in Imagine and Merrimac.

Ian Buck (a graduate student on the Merrimac project and the principal developer of Brook) joined Nvidia in 2004. Ian worked with John Nickolls to evolve Brook into CUDA.³⁰ CUDA incorporated the best features of Brook and Cg as well as incorporating feedback from Brook programmers. The story of how the technology was transferred from Stanford to Nvidia is described in a presentation.⁴ Mike Houston (another graduate student on the project) joined AMD and used Brook directly as the programming language for their GPUs. The G80 (NV50) and CUDA were launched at Supercomputing in 2006.

When CUDA was launched in 2006 few people understood parallel programming, let alone GPU stream programming. To overcome this workforce deficit, Wen-Mei Hwu and David Kirk evangelized GPU computing by teaching CUDA programming courses for professors.²⁵ The faculty who attended these courses went on to teach thousands of students parallel programming in CUDA. Parallel computing technologies borrowed from the Cosmic Cube, J-Machine, and M-Machine were applied both within a GPU—to coordinate multiple SMs—and across GPUs as multinode GPU systems were built to tackle large problems.

Enabling AI. Modern machine learning relies on three key ingredients—massive datasets, large models with many layers and weights, and the computing power to optimize the weights. The core algorithms (deep neural networks, convolutional networks, training using back propaga-



tion, and stochastic gradient descent) have all been around since the 1980s or earlier. Large, labeled datasets, such as PASCAL¹⁸ and Imagenet¹³ appeared in the early 2000s. Recent advances, such as embedding text into vector spaces, enabled deep learning for natural language. Transformers (attention is all you need)³⁷ replaced hard-to-train recurrent neural networks with hidden states with easy-to-train neural networks with history. GPU computing enabled the training of large networks with massive datasets economically. Once this capability was demonstrated (Alexnet, GPT), the capabilities of AI improved rapidly. The rapid adoption of AI provided even more impetus for improving GPU computing systems.

Machine learning at Nvidia was also enabled by academic-industry synergy. A 2010 breakfast conversation between one of the authors (Dally) and Andrew Ng led to the initiation of a joint project between Nvidia and Stanford to build deep neural networks on GPUs.⁸ Bryan Catanzaro led the Nvidia portion of the project. The software developed during this project became CuDNN,⁵ which provides a readily available library for deep learning on Nvidia GPUs—democratizing deep learning.

Conclusion

The technologies behind GPU computing, which has enabled modern machine learning, were largely developed thanks to 30 years of government-funded academic research. Research in parallel computing, parallel graphics systems, and stream processing laid the groundwork for GPU computing. Many of the students trained during these research projects took positions in industry, transferring the technology and using it to develop innovative products. The transfer from the Stanford Stream Processing project to GPU computing was very direct, with the academic Brook language evolving into CUDA and stream processor features being incorporated in the G80 GPU. The efficient, easily programmed, and very high-performance computing platform provided by GPUs with compute shaders enabled the current revolution in machine learning—providing the missing ingredient to complement the algorithms and data that had been available for some time.

Acknowledgments

The late John Nickolls played a critical role in transferring the Stanford stream processing technology to Nvidia. Scott Rixner, Brucek Khailany, Ujval Kapasi, John Owens, Jung Ho Ahn, Peter Mattson, Abishek Das, Jinyung Namkoong, Brian Towles, Andrew Chang, and Ben Mowery contributed to the Imagine project. Several of these students along with Mattan Erez, Ian Buck, Mike Houston, Timothy Knight, Nuwan Jayasena, Francois Labonte, Jayanth Gummaraju, Ben Serebrin, and Binu Mathew contributed to the Merrimac project. Several of these students along with Tim Foley, Daniel Horn, Jeremy Sugarman, and Kayvon Fatahalian contributed to the development of Brook. Kekoa Proudfoot, Bill Mark, Kurt Akeley, Ravi Glanville, Randy Fernando, and Mark Kilgard contributed to the development of RTSL and Cg. David Kirk and Wen-Mei Hwu evangelized CUDA programming. Massimiliano Fatica, Vijay Pande, Eric Darve, Juan Alanso, Alan Wray, and Tim Barth developed early scientific Brook applications on Merrimac. Grants from DARPA, DoE, NSF, Intel, Sun, SGI, NVIDIA, and AMD supported this work. 

References

1. Batcher, K. Design of a massively parallel processor. *IEEE Transactions on Computers* C-29, 9 (1980), 836–840.
2. Buck, I. Brook Spec v0.2. Stanford University, (2003).
3. Buck, I. et al. Brook for GPUs: Stream computing on graphics hardware. *ACM Transactions on Graphics* 23, 3 (2004), 777–786.
4. Buck, I. From Brook to GPUs. In *Proceedings of the GPU Technology Conf.* (2009); https://www.nvidia.com/content/GTC/documents/1001_GTC09.pdf
5. Chetlur, S. et al. CuDNN: Efficient primitives for deep learning. *arXiv:1410.0759* (2014).
6. Christon, M. et al. ASCI Red—Experiences and Lessons Learned with a Massively Parallel TeraFLOP Supercomputer. No. SAND-97-0463C; CONF-9706112-1. Sandia National Lab. (SNL-NM), (1997).
7. Clark, J. The geometry engine: A VLSI geometry system for graphics. In *Proceedings of ACM SIGGRAPH Computer Graphics* 16, 3 (1982), 127–133.
8. Coates, A. et al. Deep learning with COTS HPC systems. In *Proceedings of the Intern. Conf. on Machine Learning*, PMLR (2013).
9. Culler, D. et al. *Parallel Computer Architecture: A Hardware/Software Approach*. Gulf Professional Publishing (1999).
10. Dally, W. et al. The J-machine: A retrospective. *Retrospective In 25* (1998), 54–58.
11. Dally, W. et al. Merrimac: Supercomputing with streams. In *Proceedings of the 2003 ACM/IEEE Conf. on Supercomputing* (2003).
12. Dally, W. et al. Stream processors: Programmability and efficiency: Will this new kid on the block muscle out ASIC and DSP?. *ACM Queue* 2.1 (2004), 52–62.
13. Deng, J. et al. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition* (2009).
14. Elsen, E. et al. N-body simulations on GPUs. *arXiv:0706.3060* (2007).
15. England, N. A graphics system architecture for interactive application-specific display functions.

16. Erez, M. et al. Analysis and performance results of a molecular modeling application on Merrimac. In *Proceedings of the 2004 ACM/IEEE Conf. on Supercomputing*.
17. Erez, M. et al. Fault tolerance techniques for the Merrimac streaming supercomputer. In *Proceedings of the 2005 ACM/IEEE Conf. on Supercomputing*.
18. Everingham, M. et al. The PASCAL visual object classes challenge: A retrospective. *Intern. J. of Computer Vision*, 111, 1 (2015), 98–136.
19. Fatahalian, K., Sugerman, J., and Hanrahan, P. Understanding the efficiency of GPU algorithms for matrix-matrix multiplication. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Conf. on Graphics Hardware* (2004).
20. Fatica, M., Jameson, A., and Alonso, J. STREAMFLO: An Euler solver for streaming architectures. In *Proceedings of the 42nd AIAA Aerospace Sciences Meeting and Exhibit* (2004).
21. Fillo, M. et al. The M-machine multicomputer. *Intern. J. of Parallel Programming* 25, 3 (1997), 183–212.
22. Fuchs, H. et al. Pixel-planes 5: A heterogeneous multiprocessor graphics system using processor-enhanced memories. In *Proceedings of ACM SIGGRAPH Computer Graphics* 23, 3 (1989), 79–88.
23. Kapasi, U. et al. The Imagine stream processor. In *Proceedings of the IEEE Intern. Conf. on Computer Design: VLSI in Computers and Processors* (2002).
24. Kessler, R. and Schwarzmeier, J. CRAY T3D: A new dimension for Cray Research. *Digest of Papers. COMPCON Spring* (1993), 176–182.
25. Kirk, D. and Hwu, W. *Programming Massively Parallel Processors: A Hands-On Approach*. Morgan Kaufmann (2016).
26. Khailany, B. et al. Imagine: Media processing with streams. *IEEE Micro* 21, 2 (2002), 35–46.
27. Levinthal, A. et al. Parallel computers for graphics applications. In *Proceedings of the 2nd Intern. Conf. on Architectural Support for Programming Languages and Operating Systems* (1987), 193–198.
28. Luebke, D. et al. GPGPU: General-purpose computation on graphics hardware. In *Proceedings of the 2006 ACM/IEEE Conf. on Supercomputing* (2006), 206.
29. MPI: A Message-Passing Interface Standard. *Technical Report*. University of Tennessee (1994).
30. Nickolls, J. et al. Scalable parallel programming with CUDA. *Queue* 6, 2 (2008), 40–53.
31. Nickolls, J. and Dally, W. The GPU computing era. *IEEE Micro* 30, 2 (2010), 56–69.
32. Owens, J. et al. Polygon rendering on a stream architecture. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Workshop on Graphics Hardware* (2000).
33. Owens, J. et al. Media processing applications on the Imagine stream processor. In *Proceedings of the IEEE Intern. Conf. on Computer Design: VLSI in Computers and Processors* (2002).
34. Scott, S. Synchronization and communication in the T3E multiprocessor. In *Proceedings of the 7th Intern. Conf. on Architectural Support for Programming Languages and Operating Systems* (1996), 26–36.
35. Seitz, C. The cosmic cube. *Commun. ACM*, 28, 1 (1985), 22–33.
36. Thompson, C. et al. Using modern graphics architectures for general-purpose computing: A framework and analysis. In *Proceedings of ACM/IEEE MICRO-35* (2002).
37. Vaswani, A. et al. Attention is all you need. *Advances in Neural Information Processing Systems* 30, (2017).

William J. Dally is chief scientist and senior vice president of research, Nvidia Corporation, Santa Clara, CA, USA, and is an adjunct professor of electrical engineering and computer science at Stanford University, Stanford, CA, USA.

Pat Hanrahan is the Canon Professor Emeritus of Electrical Engineering and Computer Science at Stanford University, Stanford, CA, USA.

 This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).

DOI:10.1145/3776589

Federal funding from the NSF and U.S. Air Force enabled the author's groundbreaking research and inventions in the field.

BY MARTIN E. HELLMAN

Federal Funding of Public Key Cryptography

U.S. FEDERAL SUPPORT is currently being cut to academic institutions in ways that will adversely affect both the U.S. economy and national security. While changes are probably needed, they should be made carefully, taking into account the positive economic impact that federal support has already had in developing the Internet and other major technological developments. This article covers the key role federal support played in my work on public key cryptography, which earned the 2015 ACM A.M. Turing Award. It also explains that technology.

My Experience with Federal Support of Basic Research

In November 1976, Whitfield Diffie and I published “New directions in cryptography,”² which disclosed the

idea of public key cryptography (PKC) and provided an algorithm now known as Diffie-Hellman Key Exchange—but which I have argued¹² should be called Diffie-Hellman-Merkle Key Exchange because it is based on independent work of Ralph Merkle at the University of California at Berkeley.⁸

Today, PKC forms the foundation for e-commerce, allowing more than US\$1 trillion per day in foreign exchange transactions in North America alone.¹⁰ This technology also allows electronic banking, credit-card transactions on the Internet, and secure software updates.

Our 1976 paper notes, “This work was partially supported by the National Science Foundation under NSF Grant ENG 10173.”

My research in those days was supported by a grant from the National Science Foundation (NSF) and a contract from the Air Force Office of Scientific Research (AFOSR). That relatively small funding has paid handsome dividends.

Given the U.S. National Security Agency’s (NSA’s) concerns at that time about academic research in cryptography, I had been initially hesitant to use any of my federal support for this work and had used unsupported time to work on the subject. But my NSF grant monitor told me NSF would allow me to use their support for this work. Later, my Air Force contract monitor said the same thing, and I was able to also use funds from AFOSR to work in cryptography.

In an interesting twist, Adm. Bobby Inman (U.S. Navy, Retired), who was director of NSA during that period of resistance to open publication of papers in cryptography, has a very different perspective today. MIT (and soon to be University of California, Berkeley) Professor Henry Corrigan-Gibbs wrote in a 2014 article¹:


Inman was not excited about the prospect of high-grade encryption systems being available for purchase, especially abroad. "We were worried that foreign countries would pick up and use cryptography that would make it exceedingly hard to decrypt and read their traffic." ...

[But], today Inman is as worried about protecting nongovernmental computer systems as Hellman was in the 1970s. When asked if he would make the same decisions about nongovernmental cryptography now as he did then, Inman replied, "Rather than being careful to make sure they were[n't] going to damage [our collection capabilities] ... I would have been interested in how quickly they were going to be able to make [cryptosystems] available in a form that would protect proprietary information as well as government information."

The theft of portions of the designs for the F-35 jet, Inman said, demonstrates that weak nongovernmental encryption and computer security practices can grievously harm national security."

Flexibility in how grant money was spent was also key. Whitfield Diffie had worked at Stanford's Artificial Intelligence Lab but had left when his desire to do more research on cryptography conflicted with its contract requirements. He then traveled around the country learning about cryptography and, when he visited IBM's Watson Research Center at Yorktown Heights, NY, the head of their cryptographic research effort told him I had been there a few months earlier saying similar things (which IBM's management discouraged them from following), and that he should contact me when he returned to Stanford.

Whit showed up at Stanford in the fall of 1974, and we had an immediate rapport. I wanted to keep him here, and the flexibility in my NSF grant allowed me to hire him even though that was not contemplated at





Study is needed to determine if some areas of research that are currently funded should not be, but such determinations need to be based on careful analysis and reasoned debate.

the time of the grant. My partnership with Whit was crucial to the work that led to *New Directions in Cryptography* and the applications that followed.

The importance of this work was immediately clear to many, but some had a very different perspective. For example, Irwin Feerst, a perennial write-in candidate for IEEE president, stated, “The [*IEEE Transactions on Information Theory*] are indeed full of crap ... In fact, I can’t even tell if I’ve got it right side up or upside down.”³

In addition to *New Directions*, those *IEEE Transactions* published work which led to several important start-ups, including Qualcomm, which has a current market cap of more than US\$150 billion.

How Our System Works

Imagine that Alice and Bob want to exchange messages written on pieces of paper, but Eve (an eavesdropper) gets to relay messages from Alice to Bob.

Alice writes out her private message and locks it in a strongbox. She would like to tell the lock’s combination to Bob, but doing so would allow Eve to open the box when Alice passed it to Eve to give to Bob. Alice therefore keeps the combination secret, so Eve is unable to read the message. Unfortunately, neither can Bob.

But the hasp on the strongbox is big enough so that Bob can add his own lock, to which only he knows the combination. Bob then passes the doubly locked strongbox to Eve, who cannot take off either lock. Eve then passes the strongbox to Alice who removes her lock, leaving only Bob’s.

Alice then passes the strongbox back to Eve, who cannot remove Bob’s lock to read the message, but when Bob gets it, he removes his lock and reads the message. We have accomplished the seemingly impossible!

In this physical analogy, it was crucial that the hasp on the strongbox was big enough to hold two locks. If it had only been big enough for one lock, when Bob got the locked-up strongbox, he could have put it in a

bigger strongbox that he then locked with his lock. But when Alice received the doubly locked message, she would be unable to remove her lock from the inner strongbox.

It was important that the locks could be added and removed in any order—that the process was *commutative*. If the hasp were big enough for only one lock, adding and removing the locks would not have been commutative.

While this explanation came after I developed what is now usually called Diffie-Hellman Key Exchange, what I needed (but did not know at the time) was a commutative one-way function. A one-way function is easy to compute but extremely difficult (effectively impossible) to invert.

I had been looking for one-way functions with strong mathematical structure prior to Whit and I working on public key cryptography. As part of that effort, I met with Professor John Gill, at the time a new faculty member who had his doctorate in mathematics from UC Berkeley. (John was also one of the first Black graduates of Georgia Tech, an interesting story in itself.)

John first suggested factoring as a one-way function. But we had looked at and discarded it—although we wish we hadn’t since the RSA system is based on factoring.¹¹ John then suggested indices, which is what mathematicians call discrete logarithms.

Thus, late one night in May 1976, I was “playing” with equations, trying to find a usable public key cryptosystem. In such a system, each user has a public key and a secret key. While a user’s secret key is, in theory, computable from his public key, that computation must be effectively impossible, taking eons to accomplish.

I wrote the equation for discrete exponentiation:

$$Y = a^x \bmod p \quad (1)$$

where p is a prime number. Discrete exponentiation is easy, much as a to the 2^{101} power can be computed with 100 multiplications (squaring operations) followed by a single multiplication:

$$a^{(2^{101})} = a^{(2^{100})} * a^{(1)} \bmod p \quad (2)$$

And, because arithmetic is done *modulo p*, the numbers never get larger than p . Implicitly, Eq. 1 also defines the discrete logarithm function:

$$X = \log_a(Y) \bmod p \quad (3)$$

While it has never been proved that discrete logarithms are hard, 50 years of experience indicate it is not an easy problem.

I had four parameters with which to play: X , Y , a , and p . At one point, I assumed a and p were publicly known, in which case X had to be a user's secret key and Y his public key. User i could randomly choose his secret key X_i , and easily compute his public key Y_i as:

$$Y_i = a^{X_i} \bmod p \quad (4)$$

But computing his secret key X_i from his public key Y_i would be effectively impossible because that required computing a discrete logarithm.

I then tried to combine user i 's and user j 's keys and came up with the following method of computing a common key for them to use:

$$K_{ij} = Y_i^{X_j} = Y_j^{X_i} = a^{X_i * X_j} \bmod p \quad (5)$$

Not trusting my equations, I programmed these operations on an HP scientific calculator with a small prime, such as 97, and confirmed the same key was obtained by user j (by raising user i 's public key Y_i to his secret key X_j) and by user i (by raising user j 's public key Y_j to his secret key X_i).

Note that commutativity is crucial. It does not matter if one first raises a to the X_i power (to obtain Y_i) and then raises that number to the X_j power, or *vice versa*. You get the same result.

But instead of a public key cryptosystem, which Whit and I had been trying to find, this turned out to be a Merkle public key distribution system. In the latter system, users i and j compute a common key that neither of them could have predicted beforehand. That is why, if names are to be assigned to this system, I have argued it should be called Diffie-Hellman-Merkle Key Exchange.

Ralph Merkle independently came up with the privacy half of public key

cryptography as a master's student at UC Berkeley. His 1978 article in *Communications*⁸ gives the impression that his work followed ours (published in 1976) and RSA's (published in 1977). But the delay in the publication of Merkle's article was due primarily to people not appreciating his work, which fits with the next section of this article on "The Wisdom of Foolishness."

When Merkle had to do a term project in CS 244 in the fall of 1974, he proposed "secure communication over insecure channels" (the privacy part of public key cryptography), but the professor wrote: "Project 2 [a much more mundane proposal] looks more reasonable, maybe because your description of Project 1 [the privacy part of PKC] is muddled terribly." Merkle dropped the course and pursued public key distribution on his own.

When Merkle submitted an early version of the paper to *Communications*, it was rejected with the editor telling him in an Oct. 22, 1975 letter, "I ... was particularly bothered by the fact that there are no references to the literature. Has anyone else ever investigated this approach?"⁴ Partially in the editor's defense, Ralph should have included references, as he did in the final version of the article that appeared three years later.

The editor was also hindered by the reviewer on whom she relied. The reviewer recommended rejecting the paper because it "is not in the mainstream of present cryptography thinking." He was right. It was totally new and groundbreaking.

Today, Diffie-Hellman-Merkle Key Exchange (usually called Diffie-Hellman Key Exchange) is often done with respect to elliptic curves. Elliptic curve cryptography (ECC), suggested independently by Neal Koblitz⁷ and Victor Miller,⁹ makes use of the fact that any commutative one-way function can be used in the algorithm, not just exponentiation in modular arithmetic. ECC uses addition on elliptic curves and allows smaller keys for a given security level than either exponentiation in modular arithmetic or the RSA algorithm.

The Wisdom of Foolishness

As seen by the treatment Merkle's work received, it is important not to let conventional wisdom dictate the work that is undertaken. My own research in cryptography was also met with universal resistance at first, leading to what I now call "the wisdom of foolishness." The following anecdote, taken from a book my wife Dorothea and I wrote on "creating true love at home and peace on the planet,"⁶ provides a good example:

Dorothea: Recently, we were at a birthday celebration for a good friend and one of the other guests asked Marty what he did. When he told her that we were writing this book, her reply was polite, but she made it clear she thought we were on a fool's errand. To which Marty replied, "Thank you! That's the best compliment you could pay me." I'll let him explain.

Marty: It's ironic that one of the last things most people want to be seen as is foolish, yet many great advances were initially derided as fools' errands before they paid off.

I had a personal experience of this. My work in cryptography (which today is called cybersecurity) has earned me some of the highest honors in my profession, but when I was starting out, all of my colleagues told me I was wasting my time. Luckily, one of them—former UCLA Professor Jim Omura—was standing near me when that birthday party guest basically told me that we were wasting our time writing this book. I asked Jim to come over and tell her what he had thought of my cryptographic research in the early days, and he obliged.

Jim had told me that I could quote him, so I knew that he would tell her how foolish my effort had seemed at first—which he did. In an interesting twist of fate, after my work panned out, Jim left UCLA and co-founded a Silicon Valley startup that capitalized on it and later sold the company for \$35 million.

In 2013, when I was inducted as a "Stanford Engineering Hero," my lecture on "The Wisdom of Foolishness"⁵ included several examples of breakthrough technological work that was initially derided.



Vint Cerf, who with Robert Kahn developed the now widely used TCP/IP protocol, emailed me, saying, “Packet switching was regarded as crazy, even after the great example of the ARPANET. Most conventional communications engineers thought it would not handle speech let alone video. Of course, it got better and better over time. The naysayers believed that store and forward would have too high a latency to be useful. They were wrong.”

Today, all major streaming video services use packet switching to deliver video. Cerf and Kahn were awarded the National Medal of Technology and Innovation by President Clinton for their work. It also won them the ACM Turing Award.

Brad Parkinson, who pioneered GPS, told me, “The Air Force thought GPS was foolish. They saw no value in precision weapons delivery and certainly did not identify with the civilian applications (which I told congress from the start was a part of the system).

“They also thought that the user equipment would be prohibitive, but we could foresee the digital revolution underway. The first GPS units weighed about 150 pounds, used about 2Kw, and cost about US\$250,000. Today, a GPS receiver fits on a chip weighing a fraction of an ounce, uses milliwatts of power, and costs (chip only) under US\$5 [updated in 2025].

“None are so blind as they who refuse to see. The Air Force generally persisted in their view, until the first Gulf War showed what the system could do.”

In 2018, Brad won the IEEE’s highest award, its Medal of Honor, for that work.

John Cioffi, sometimes called “the father of DSL,” emailed me saying, “In 1979, as a young kid fresh out of Stanford, I got laughed out of the room in my first DSL meeting at Bell Labs when I pushed for more speed than the 160kbps then coming into discussion. I said 1.5Mbps should be possible on the existing twisted pairs.

“Today [2013], DSL speeds are more than 100X greater than even that projection, and DSL has 70% of

the worldwide broadband market.”

John was awarded the National Medal of Technology and Innovation by President Biden for that work.

Federico Faggin, who designed the first microprocessor at Intel wrote, “In 1970, when I was designing the world’s first microprocessor, the CEO of American Micro Systems, then the largest producer of random-logic custom ICs [a technology that lost market share to microprocessors] was quoted in *Electronics News* as saying something like, ‘Whoever thinks that a CPU (the word microprocessor had not yet been coined) will replace random logic is out of his mind.’ I remember cutting out that article and posting it in my office as an added incentive to prove such myopic observation wrong!”

Federico was awarded the National Medal of Technology and Innovation by President Obama for his work.

Since I received the 2015 ACM A.M. Turing Award for my work on public key cryptography, I have had the opportunity to ask six Nobel Laureates if the work that won them their prizes was initially encouraged or discouraged as “foolish, crazy, never go anywhere.” Five of the six were clearly in the latter camp.

Dan Shechtman, who won the 2011 Nobel Prize in Chemistry for his work on quasi-crystals, told me that no less an authority than Linus Pauling—who had won Nobel Prizes in both Chemistry (1954) and Peace (1962)—derisively referred to it as “quasi-science.”

Lou Ignarro, who won the 1998 Nobel Prize in Physiology or Medicine, told me that his dean came to him and said something like, “Lou, we hired you because you did good work. Why are you doing this crazy stuff?” Crazy stuff that won the Nobel Prize.

Chapter 4 of the book that Dorothea and I wrote is on “The Power of Belief Systems” and starts out with a Mark Twain quip: “They did not know it was impossible so they did it.”

Conclusion

Federally funded research has paid handsome dividends, with my personal experience being just one example. Such research has had

economic benefits, but it also has enhanced national security as evidenced by Adm. Inman’s earlier comment. And, as shown in this article, seemingly foolish research can pay great dividends.

Study is needed to determine if some areas of research that are currently funded should not be, but such determinations need to be based on careful analysis and reasoned debate. ■

References

1. Corrigan-Gibbs, H. Keeping secrets—A history of the birth of non-governmental cryptography research. *Stanford Magazine* (Nov. 6, 2014), 58–64.
2. Diffie, W. and Hellman, M. New directions in cryptography. *IEEE Trans. on Info. Theory*, IT-22 (Nov. 1976), 644–654.
3. Ephremides, A. Historian’s column. *IEEE Information Theory Society Newsletter* 51, 3 (Sept. 2001), 11.
4. Graham, S. Letter to Ralph C. Merkle (Oct. 22, 1975); <https://www.ralphmerkle.com/1974/RejectionLetter.pdf>.
5. Hellman, M. The wisdom of foolishness. Stanford Engineering Hero lecture (Jan. 29, 2013); <https://www.youtube.com/watch?v=EFITgOx3Z7A>.
6. Hellman, D. and Hellman, M. *A New Map for Relationships: Creating True Love at Home & Peace on the Planet*, New Map Publishing, Stanford, CA (2016).
7. Koblitz, N. Elliptic curve cryptosystems. *Mathematics of Computation* 48, 177, American Mathematical Society (1987), 203–209.
8. Merkle, R.C. Secure communications over insecure channels. *Commun. ACM* 21, 4 (April 1978), 294–299.
9. Miller, V. Use of elliptic curves in cryptography. In *Proceedings of Advances in Cryptology—CRYPTO ’85*, Lecture Notes in Computer Science 218 (Jan. 2000), 417–426.
10. New York Federal Reserve. FX Volume Survey (Feb. 12, 2025); <https://www.newyorkfed.org/fxc/fx-volume-survey>.
11. Rivest, R., Shamir, A., and Adleman, L. A method for obtaining digital signatures and public key cryptosystems. *Commun. ACM* 21, 2 (Feb. 1978).
12. Singh, S. *The Code Book: The Evolution of Secrecy from Mary Queen of Scots to Quantum Cryptography*, Doubleday, New York (1999).

Martin E. Hellman is Professor Emeritus of Electrical Engineering at Stanford University, Stanford, CA, USA. He is a co-recipient, along with Whitfield Diffie, of the 2015 ACM A.M. Turing Award for their groundbreaking research and inventions in the field of cryptography



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).

We know a lot about distance work and how to make it work, thanks to 30 years of federally funded research.

BY JUDITH S. OLSON AND GARY M. OLSON

Working Together and Working at a Distance

WE ARE COGNITIVE psychologists by training, now in a subfield of computer science (CS) called human-computer interaction (HCI). We have spent more than 30 years studying how the technology supports teamwork both when teams work face to face or at a distance. We are delighted to tell our story of how

federal dollars supporting our research not only advanced science in important ways but also helped new areas of “big science” work more successfully. Our work also generated important practical advice to everyday people, both during Covid and now that we are working remotely. We had studied group work for a number of years, but in the early 1990s, we pivoted our research to focus on whether people could work at a distance, whether the quality of the work suffered or productivity fell, as well as how people felt about working that way. What could employees do to make it work well? What could managers do? What equipment was needed and how should it be ar-

ranged? What mattered? In the early 1990s, videoconferencing was very expensive, and laptops were not practical yet. But the things we learned about people working together with advanced technology are still applicable today.

Thirty years later, COVID hit, and almost everyone was working from home. Videoconferencing was free or at least inexpensive. Drawing on our decades of federally funded research, we knew what we could do to help and seized a multitude of opportunities to get the word out. We were interviewed on television and reported in print publications in the mainstream, not just academic journals. Word got out that it is possible



to work from home, but you had to set things up right and practice good habits. You *can* make distance work really work. Here is that story.

A Place for Teamwork

We were professors at the University of Michigan: Gary was in the Psychology Department and Judy was in the Business School's Department of Computer and Information Systems. Our first joint research involved studying teamwork, both in real life, called "field work," and in laboratory experiments. To run experiments, we needed a space, somewhere groups could come together to be recorded on audio and video while working. This was the late 1980s. As it happened, the business school had just built a new building for research, and it had a big open space with offices surrounding it. The dean knew we needed dedicated space to carry out our federally funded research, which he deemed highly relevant to current business practices: how to manage teams of workers so productivity is high, how they should be arranged, and what technology they needed.

The dean gave us the space and introduced us to two Michigan corporate sponsors. Impressed with our track record of publishing sound work and our continuing support from the National Science Foundation (NSF), one sponsor offered us money and access to their teams for observation; another offered us money and design help in furnishing

the open space in the new research lab. They could help us design the room that would serve as a sophisticated meeting room as well as a place to run groups in our experiments. That meant the room was not only beautiful and functional but was also equipped with hidden cameras and microphones. These were controlled from an observation room with one-way mirrors so we could see them, but they could not see us. The room (Figure 1) was named the Collaboration Technology Suite (CTS).¹¹

The room's striking feature was its eight workstations, individual desks on wheels embedded with the large computer monitors of the day. These were the late 1980s, before there were functional laptops. Today, you put your laptop on any meeting table and set it at any angle for viewing and making eye contact. In the late 1980s, all we had were desktop computers, the bigger the better. But the large monitors blocked sight lines, making them unsuitable for working with others. Our newly designed desks allowed us to set the computer screens at any angle between flat and fully upright to suit the work at hand. In addition, the desks could be set in any configuration in the room: a circle, a long table (as in Figure 1), or clusters of two for breaking up into subgroups. The large rear-projection screen in front could project a video, PowerPoint slides, or any team member's computer screen, depending on the meeting's purpose.

The CTS served us for our lab studies, which we'll soon discuss. While it was built with corporate funds, it was federal funding that made the scientific research possible by supporting experimental technology development, supporting students and assistants, and paying research expenses. The CTS also served as a showcase in tours of the business school. On a number of occasions once the word got out, people from across campus would use the room because of the technology and arrangement it offered: Everyone had a computer at hand and any of them could be projected to the front screen. And, of course, our lab meetings took place here with us together taking notes on our designed shared editor.

Showing That the Lab Work Applied to the Real World

As mentioned earlier, throughout our research career, we collected data from field work and from experiments run in the CTS. We wanted the work in the lab to be as close to real work as possible. In many of the experiments, we had people who already knew each other and had worked together to design an automatic post office. (This was 10 years before automated post offices became popular.) To make sure our results from experiments applied to the real world, we needed to show that the setting and task are comparable to real teamwork.

We had recordings from 10 real-life design meetings from four projects in two organizations. We transcribed these video recordings with timestamps on each utterance. We coded each of their utterances about design, for example, stating an *Issue*, offering an *Alternative*, and then giving a reason, *Criterion*, to evaluate the *Alternative*. There were then other labels having to do with managing the session and clarification. Overall, only 40% of the meeting times were spent on design; the rest had to do with summarizing, planning the time, and digressions.

In Figure 2, we show the transition diagrams of real meetings side-by-side with that of the lab meetings. The size of the circles represents

Figure 1. The Collaboration Technology Suite with the observation room to the right behind one-way mirrors.





total time spent, and the width of the arrows indicates the relative frequency with which that transition is made. The white areas inside the circles indicate the original statement in that category; the black indicates questions and their responses.

There are some understandable differences, in that those in the lab did not do longer-term project management. But the overall pictures were remarkably similar. Indeed, the activity from laboratory tasks in general matched that of the meetings in the real world, meaning findings in the lab would apply to the real world.¹⁵

Building the Shared Editor and Testing Its Effects

One of the first tools we developed in the early 1990s was a shared editor, which we unwisely called ShrEdit. The application was very similar to Google Docs, which was released 15 years later; any number of people can have a document open and edit it at the same time. Figure 3 shows ShrEdit, with two shared windows at the top and one private one at the bottom. The boxes in the middle illustrate a feature that today's collaborative editors should have but do not: a way to move to where someone else is working and when done, pop back to where you were last working. We wanted to see if ShrEdit affected work.

One of our first experiments showed that the groups using ShrEdit created significantly better designs than those using a whiteboard and Microsoft Word.¹⁶ But they did it in an unexpected way. They considered fewer *Issues* but more *Alternatives*, apparently expanding on the major issues with good ideas. Not surprisingly, the groups using ShrEdit wrote significantly longer final reports, mainly because more people could write at the same time.

In 2005, Google Docs was introduced, and we found it was very similar to our ShrEdit. So, from 2014 to 2016, Judy asked her students in a project management class to use Google Docs for their assignments.¹⁷ We asked the students' permission to use their data for research and all of them complied.

Figure 2. Comparison of the use of time and the transitions among activities in the field (left) and those in the unsupported laboratory conditions (right).





Figure 3. The interface of ShrEdit.



Ninety-five percent of the documents exhibited some amount of simultaneous work. The documents in which authors wrote about equivalent amounts were judged to be of higher quality. We also used their data to look at the different styles of working. We created visualizations similar to the work of Viegas, Wattemberg, and Dave visualizing Wikipedia collaborations.²¹ Ours were

called Docuviz. Figure 4 shows two exemplary styles of work. The one on the left shows an outline style, where the person depicted in blue wrote an outline and then the others filled in their assigned portions of the document with very little editing of each other's contributions. The one on the right shows a case where they found a document similar to what was asked for, pasted it in, and pared it down to

Figure 4. Two visualizations of the growth of a document with time going left to right, and color indicating author.




Figure 5. Mid-1990s desktop video support for experiments about work with remote team members.



fit the details of the assignment. Then they embarked on many sessions of simultaneous work to include their material in the style of the found document. The occurrence of simultaneous work is shown by the number of colors (participants) at the top in each time slice. The document on the right had many sessions of simultaneous work, whereas the one on the left had *one* simultaneous session in the second-to-last time slice.

The new generation of information workers embrace simultaneous work. Fifteen years after our original ShrEdit experiments, our findings about collaborative work are even more relevant now that online shared editors like Google Docs are widely adopted. Today, Google Docs has more than a billion active users.

Lab Work on Working at a Distance

All our experimental work on face-to-face teamwork was done in the CTS, which, because of its functionality and beauty, was also a stop on many tours of the business school. At one such session, a visiting executive said to us, “I think this is very exciting work, but if you want to really make an impact on the world, study *distributed* teams and watch them over time. We need help in making that work.”

This comment caused us to pivot our whole research program from co-located teams to distributed teams. We focused on investigating not only what technology helped them but also the most effective styles of work. We simulated distributed conditions by commandeering three offices

and outfitting them with ShrEdit and video equipment (See Figure 5). This was 10 years before the first affordable videoconferencing system (Skype) and 20 years before Zoom. We ran groups of three people in our experiments, so each room had two monitors, one on each side of the computer, arranged so that when the two remote people looked at each other, it appeared to this person as if they were looking at each other.

Our work¹⁵ again used ShrEdit and the automatic post-office design task among people who knew each other, communicating face to face, remote with video, and remote with only audio. It turned out the designs where people had video were statistically equal in quality to those created when they are face to face. In contrast, those using only audio created significantly poorer designs. Adding video to audio makes the participants not only better at what they are doing but also significantly more satisfied than those using only audio. Video was especially helpful for non-native English speakers because they could see when the other person did not understand (e.g., eyebrows down), so they explained more fully without being asked.²⁰

How the monitor is positioned has an effect:⁹ If you view the remote person from the monitor placed high, they look tall and dominant. If the monitor is low, the person looks short and weak. Indeed, people seen on higher monitors had more influence than those on lower ones. Since laptops are small and the view of the

person is low, many people using video on laptops today, appropriately put them on stacks of books so they see the other person at eye level, and the remote person does the same.

We also ran a series of studies where some people were colocated and others were remote and alone, one of the hardest configurations to work in but adopted by many employees today calling it “hybrid work.” The ease of communication among those colocated creates a strong bond. Those colocated preferred to work with other colocated workers, even though remote workers could perform the work better.⁴ Remote workers, being ignored by colocated staff, worked with each other. Even when remote people are from one’s own subgroup, the ease of communication with those who are colocated makes one select colocated people to work with.² And, when remote workers become colocated, they will also work closely with colocated workers even though they had previously preferred working with other remotes.³

We also studied trust development in the lab, knowing that trust is very essential for work to be smooth and satisfying.⁸ We measured trust using the Prisoner’s Dilemma, a game-theory thought experiment involving rational agents who can either cooperate for mutual benefit or betray their partners for individual gain. The dilemma arises from the fact that while defecting is rational for all agents, cooperation yields a higher payoff for each. In playing this game, we had three-person groups: Conversing face to face, by video and audio, by audio only, or by text chat. Participants played 30 rounds, with discussion encouraged in breaks after every five rounds. Those meeting face to face trusted and acted trustworthy throughout.⁶ Those with audio only and video plus audio had less trust than face to face, but their trust grew throughout and eventually reached a level equal to face to face. Those communicating only by text had significantly lower trust and did not increase over time.

Our follow-on experiment showed it was possible to effectively *pre-build* trust in distributed collaborations. We knew from the literature that if

people met face to face before working remotely, they will exhibit trust. We ran an experiment showing that if people engaged in a get-to-know-each-other activity via *video*, they developed nearly the same level of trust as if they had met face to face.²² Having a team get to know each other by video is fortunately a much cheaper way to foster success than paying for their travel and lodging; it saves time and effort as well.

In summary from findings in the laboratory, we first established that the design task we used in the lab was close enough to real-world design sessions to have faith that our findings in the lab would apply to the real world. We created a shared editor, similar to today’s Google Docs, and found that groups that used it created better designs than those using whiteboard and single-user Word. In examining people using Google Docs for academic group homework, we found that 95% of the documents exhibited simultaneous work, though they did it in a variety of styles. The best designs came from groups whose members contributed equally.

When working at a distance, using video created designs not significantly different from those working face to face, whereas those working only with audio had designs that were significantly lesser quality. Video was particularly important for those who were not native speakers of English. Placement of the camera and monitor could affect one person’s influence over the others, where being artificially tall gave one greater influence. Remote workers with audio only or audio plus video trusted others at lower levels than those face to face, although the trust grew over time. Remote workers who conversed only with chat were significantly lower in trust and did not improve over time. But trust could be jump-started if remote workers have conversations beforehand, getting to know each other when they conversed by video.

If some people are distant and others colocated, those colocated preferred to work with those in their same location; those remote worked with other remotes because



We ran an experiment showing that if people engaged in a get-to-know-each-other activity via video, they developed nearly the same level of trust as if they had met face to face.





they were rarely contacted by colocated people, even when some of the people who are remote have a long relationship with some of those who are colocated. The ease of communication is very strong in deciding with whom to work.

Findings from Observations in the Field

In the early and mid-1990s, the scale of the problems facing the world—fighting cancer, addressing climate change, and so forth—required the collaboration of larger teams of researchers from different institutions and different but related areas of science. This new style of scientific work became known as “Collaboratories: Laboratories without Walls.” NSF was spending huge amounts of money supporting their development; some of them succeeded, some failed. We convinced them that it would be money well spent if NSF could have us focus on these big science projects: finding out what makes them succeed and what makes them fail, possibly developing some assessment techniques to tell which are which, and hopefully coming up with remedies to help all of them succeed. We joined with Tom Finholt, Stephanie Teasley, and Dan Atkins, all at Michigan, to secure funding from NSF for a five-year project.

The Science of Collaboratories group gathered data on collaboratories in two ways:

- First, we were collecting some basic data on all the collaboratories we could find. In 2008, the year our book *Scientific Collaboration on the Internet*¹⁸ came out, we had identified 212 collaboratories: 150 could be classified, and for 70 of them, we created summaries with more detail.

- Second, the team joined a set of collaboratories to both assist in their building or adoption of technologies and to observe them with the goal of finding out what made good collaboration and where the failures were. The team joined collaboratories in several fields, including upper atmospheric physics (UARC and SPARC), earthquake engineering (NEESgrid), HIV/AIDS research in Africa (IARC), astronomy (The National Virtual Ob-

servatory), particle physics (the Large Hadron Collider collaborations), and biomedical research (Biomedical Informatics Research Network, BIRN) among others. Results appear in Olson et al.¹⁸

For example, the first large-scale project we were involved in was the Upper Atmospheric Research Collaboratory (UARC, pronounced “you-ark”).¹² UARC joined domain specialists, computer scientists, and us, behavioral scientists, to build state-of-the-art technologies to collect and analyze data from the upper atmosphere and allow scientists from around the world to collaborate. The domain scientists studied interactions of the solar wind, the earth’s magnetic field, and the characteristics of the upper atmosphere. The northern lights are beautiful manifestations of some of these interactions, when a burst from the sun hits the upper atmosphere and creates not only the lovely hues in the sky visible to the human eye, but also disruptions in the upper atmosphere that tell scientists more about the earth’s atmosphere. The five-year project began in 1993 using a “user-centered iterative rapid-prototyping design strategy” throughout the project, meaning we would continuously be asking the scientists how the interface could be better and what new functionality they wanted. They collect data from observatories, with particularly rich data coming from the Arctic. Our goal was to connect the observatories to the Internet, allowing scientists to collect and analyze data from around the world without having to travel to cold remote locations. Scientists would see data simultaneously from around the world, and they would be able to communicate various analyses and not only determine together model components but also plan the next steps.

The initial application mimicked the interfaces of the physical devices to help the scientists learn quickly. But over the 10 major redesigns in seven years, the space physicists wanted more aggregated data, data from around the world, and eventually side-by-side views of data analysis and model representations.

There were more scientists from all over the globe wanting access. And because there were so many choices of what to view and who to interact with, they developed “rooms” with coordinated sets of data and “clubs” as groups of collaborators. And, like ShrEdit, collaborators could see the same view and all share the pointer to aid discussion. Figure 6 shows the interface as it was in 1997. The “chat” is on the lower left of this screen. This was particularly helpful because the collaborators came from all over the world, with drastically different time zones. The “chat” allowed them to read back in time to catch up to the topic at hand and to contribute appropriately.

Near the end of UARC, the Internet became popular, so the whole desktop application became a website. Converting everything to the Web gave them modularity, so the suite of tools could be used by other communities of collaborating scientists, interoperability over platforms, and the possibility of integrating with other Web tools and data sites. Thanks to these achievements, the project was funded by NSF for another five years under the Space Physics and Aeronomy Research Collaboratory (SPARC).

Two major advances were made in the understanding of these collaboratories. First, from the broad-scale data we collected, we synthesized a *taxonomy* of collaboratories⁵ and second, from all the observations of the collaborations witnessed in depth, we developed a first-level *theory* of what makes them successful.¹⁴

The seven types in the taxonomy, a short description of each, and an example of each are:

1. **Shared Instrument:** Remote access to expensive instruments, such as telescopes, frequently supplemented with videoconferencing, chat, electronic lab notebooks, or other communication tools, for example, the Keck Observatory^a

2. **Community Data Systems:** An information resource created, maintained, or improved by a geographically distributed community, for

a <https://keckobservatory.org>



example, the Protein Data Bank^b

3. Community Infrastructure Project: Infrastructure to support work in a particular domain. Software tools, standardized protocols, new scientific instruments, and educational methods, for example, the Grid Physics Network (GriPhyN)^c

4. Open Community Contribution System: An open project that aggregates the efforts of many geographically separate individuals toward a common research problem. Contributions come in the form of work rather than data. Results are made available to AI systems, for example, the Open MIND project^d

5. Virtual Community of Practice: People share a research area broadly and communicate it online. They share advice, techniques, and pointers to other resources but do not undertake joint projects, for example, Ocean.US^e

6. Virtual Learning Community: The goal is to increase the knowledge of the participants, not to conduct research, for example, The Ecological Circuitry Collaboratory^f

7. Distributed Research Center: These function as research centers at universities but are distributed, for example, Inflammation and the Host Response to Injury^g

The second advance was the Theory of Remote Scientific Collaboration (ToRSC), the synthesis of findings from the field, both science and corporate, and the laboratory, building on our earlier heavily cited work.¹⁰ They list the five aspects of collaboration that lead to success.

1. The nature of the work: Structure the tasks so interdependence across distance is reduced.

2. Common ground: Construct common vocabulary and ask questions if you do not understand something.

3. Collaboration readiness: Create incentive structures that support a mix of competition and cooperation.

^b <https://www.rcsb.org>

^c <https://dl.acm.org/doi/10.5555/822086.823334>

^d <https://open-mind.net/about>

^e <https://oceanservice.noaa.gov/international/>

^f <https://ui.adsabs.harvard.edu/abs/2004nsf....0404876H/abstract>

^g <https://med.stanford.edu/sgtc/research/completed/inflammation.html>

Figure 6. A screen dump of UARC user software during the April 1997 campaign.




Figure 7. An example/template for a Communication Covenant.

Communication Covenant for [collaboratory] [Date]

We, members of the [collaboratory], agree to the following communication guidelines:

Conduct of meetings

I will arrive at meetings on time. If unavoidably late, I will text the Chair with an estimated arrive time.

If something important comes up, making me miss a meeting, I will let the Chair know as soon as possible. I will send comments to the chair to contribute to the meeting and read the minutes, so I am not behind on what is going on.

The meetings will be ad hoc unless things get tight, when we will follow Roberts Rules of Order with the Chair leading.

Interruptions will be at a minimum.

Communication between meetings

I will check my email and text several times every day. The following are the correct emails and phone numbers: [email addresses for everyone] [text numbers for everyone]

When I receive an email or text from a member of [collaboratory], I will reply with "Got it." plus any additional information that could make the work go smoothly or an estimate of when that information or work will be completed.

I will keep others in [collaboratory] aware of issues as they arise. If I agree to take on a task and something comes up that interrupts my ability to continue that task or an unanticipated difficulty arises, I will let the [collaboratory] members know.

Collaborative writing

If we are jointly writing a document, we will use [Microsoft Word/Google Docs] with tracking changes on so we can see what the changes are. We will use comments freely.

We will designate one [collaboratory name] member to be the manager of each document so that we can talk about issues and decide on something and that person will accept or not the changes and resolve issues. The person will clean up the document for permanent storage or dispose of the document when it is no longer useful.

Store our documents

We will store our documents, drafts, and final version in a folder set up in [Google Drive/Dropbox] called [name].

We will have two major folders, one for drafts and one for final documents. The manager of each document will handle sharing rights of their documents.

Signatures



4. Management, planning, and decision making. Provide good leadership. Team members should feel that decisions are being made fairly and clearly.

5. Technology readiness. Provide technology to support creative activities, coordination support, storage of information, and bandwidth to support both the work and people's communication without delay. Create a "communication covenant," a template for which is shown in Figure 7.

The same chapter outlines what we considered "success." It encompassed a wide variety of outcomes beyond advances to the science itself, including impacts on scientists' careers, learning and science education, inspiration to others, funding and public perception, and new tools built and used. Federal funds supported many years of our work so that team science could be more successful bringing results.

Broader Impacts

The group of scientists and behavioral scientists interested in team science, like us, grew into a conference, what today is called the International Network for the Science of Team Science (INSciTS, pronounced "insights"). Its webpage says the group includes those "leading, engaging in, managing, facilitating, supporting, evaluating, and producing scholarship on team science." There are Special Interest Groups (SIGs) that focus on education, beginning a collaborative, data analytics, and fostering team science in academics. The creation of INSciTS is a testament to federal funding producing self-perpetuating research communities.

Clearly, having people succeed in their laboratories speeds up the discovery of important things for everyone, such as medical breakthroughs, actions to reduce global warming, planning and constructing buildings that survive a large earthquake, and so on. The fact that the association INSciTS exists, where people do *research* on how to do research better, says people want to improve their own collective work.

Several resources have been established to help people collaborate.

The Mendeley Science of Team Science Library^h is one. It houses 2,500 team science references, and it is a good example of one of the seven collaborative types, an Open Community Contribution System. One can sign up free to be a user but is also encouraged to submit other papers in a crowdsourcing fashion.

There are two committee reports on team science put together by the National Academies of Science, Engineering, and Medicine (NASEM), the combined National Academies: One committee of experts in the field developed a report called *Enhancing the Effectiveness of Team Science*,^{7,i} published in 2015, and the more recent committee published a report in 2025 called *The Science and Practice of Team Science*.^{19,j} Both sites offer the substantial reports to download for free.

For people wanting tools to help improve their own collaborations, we offered the Collaboration Success Wizard¹ for five years. This software asked a number of questions about the collaboration and automatically returned a written report on the strengths and weaknesses of the collaboration as well as what actions to take to make success more likely.

We also gathered the *Team Science Acceleration Lab Toolkit*, a website offering 40 different exercises or tasks that have been used to help collaborative members be successful. It organizes the tools in terms of what stage the team is in, that is, whether they are planning the team, launching it, or helping it with ongoing work. For example, it includes instructions on how to create a Communication Covenant, lists things that foster trust among team members, offers a tool to help when conflicts arise, and even gives advice about recognizing "bad apples" and dealing with them. A special section highlighted on the home page is "How to work at a distance." It lists

advice such as selecting players who have worked well at a distance, and assigning distant workers semi-independent work so they do not have to contact each other often, and so on. The advice there closely aligns with advice in our 2014 book for non-scientists called *Working Together Apart: Collaboration over the Internet*.¹³

When Covid Hit and Most People Were Distant

When Covid hit the world, it forced most workers to work from home. Because we created practical advice for working remotely, our research was picked up in the general press. Although our work did not say anything about working while parenting small children or having cats walk in front of the camera, much of the advice fit. Our findings, such as the advice about trust, communication covenants, and so on, did hold. And because impressions matter, we had advice about making the camera the height of your eyes, so it looks to the remote people as if you are looking directly at them. The idea of putting your laptop on a pile of books even made it onto a *New Yorker* cover.

We were interviewed by a number of outlets, including *The Atlantic* for an article titled "When Working from Home Doesn't Work." In June 2020, soon after the pandemic hit, Judy was interviewed on NPR for an article titled, "Get a Comfortable Chair: Permanent Work from Home is Coming."^k The article does talk about other aspects of distance work, but we were amused that the advice about getting an ergonomic chair made it to the title.

Many companies these days are adopting a "hybrid" model, where people are allowed to work from home some days but are required to come to the office on others. Our research⁴ shows that this is probably the worst model if people choose themselves which days they work from home. On some days, they might be the only person on their team who is remote. Unfortunately, as the saying goes: out of sight, out of

h https://www.mendeley.com/?interaction_required=true

i <https://nap.nationalacademies.org/catalog/19007/enhancing-the-effectiveness-of-team-science>

j <https://nap.nationalacademies.org/catalog/29043/the-science-and-practice-of-team-science>

k <https://www.npr.org/2020/06/22/870029658/get-a-comfortable-chair-permanent-work-from-home-is-coming>

mind. In partially distributed teams, the remote person is likely forgotten when information is exchanged. If managers want to allow people to work from home some days, the people on a team should all choose the same days. When everyone is remote, when communication needs to happen, it happens, either when everyone reads Slack or when everyone is on video.

Conclusion

Today, we know a lot about distance work and how to make it work. Thirty years of federally funded research brought significant knowledge about how to do a now-common work style well. It brought a huge benefit from smaller teams, to companies, to large-scale scientific research.

From our laboratory work, we know that having a shared editor like Google Docs improves the work. And today's young professionals adopt it and work together simultaneously in many of their documents but do their work in a variety of ways. We know that distant work with video is more effective than with audio only, nearly as good as face-to-face work, as well as more satisfying than only audio. Having the camera at eye level makes the participants equally dominant. When teams are split so that some are colocated and some remote, those colocated prefer the easy communication with other colocated colleagues. This effect is very strong, overriding empathy acquired from having been remote, and even overriding cases where the remote colleague has a long history with the colocated colleague. Trust was highest when people have been colocated and got to know each other. With video connections, trust can be built, whereas with text chat trust stays at a low level. Trust can be jump-started by a purposeful get-to-know-you session over video before teamwork starts.

From our field work, we know that laboratories tackling big questions in science can take many shapes, such as shared instrument, community data systems, open community infrastructure systems, open community contribution systems, virtual community of practice,

virtual learning community, and distributed research center. To succeed, you must assign the work to distributed colleagues so the need for communication is low. You must create common vocabularies and adopt similar working styles. There should be a culture of cooperation with some competition. Leadership must be clear and fair. And, the right suite of technologies should be available and taught so colleagues can communicate easily; coordinate, store, and retrieve material; and have sufficient bandwidth and storage capacity to support easy communication and unlimited storage. Success in laboratories can be defined as advancing science, enhancing individuals' careers, increased learning and science education, inspiring others, funding and public perception, and the building of new tools.

The work has had broad impacts, including the coming together of like-minded scientists to share information and to continue the Science of Team Science's productivity. There are several collections of information available to those studying team science and a toolkit on the Internet to offer myriad exercises to do at any stage of building or running a laboratory. All this was made possible by federal funding over 30 years.

Acknowledgments

We are greatly indebted to two colleagues, Matthew Bietz and Gloria Mark, for their careful, knowledgeable reading and critique of an earlier version of this article and two anonymous reviewers at *Communications*.

References

- Bietz, M. et al. Improving the odds through the Collaboration Success Wizard. *Translational Behavioral Medicine* 2, 4 (2012), 480–486.
- Bos, N. et al. Shared identity helps partially distributed teams, but distance still matters. In *Proceedings of the 2010 ACM Intern. Conf. on Supporting Group Work* (2010), 89–96.
- Bos, N. et al. Traveling blues: The effect of relocation on partially distributed teams. In *Proceedings of Human Factors in Computing Systems* (2005), 1917–1920.
- Bos, N., Olson, J., Nan, N., and Cheshin, A. Subgroup biases in partially distributed collaboration. *J. of Information Technology Research* 2, 1 (2008), 1–18.
- Bos, N. et al. From shared databases to communities of practice: A taxonomy of laboratories. *Scientific Collaboration on the Internet*. G.M. Olson, A. Zimmerman, and N. Bos (eds), MIT Press, Cambridge, MA (2008).
- Bos, N. et al. The effect of four communication channels on trust development. In *Proceedings of*

the Conf. on Human Factors in Computing Systems (2002), 135–140.

- Enhancing the effectiveness of team science. M. Hilton and N. Cooke (eds.), National Academies (2015).
- George, J. and Jones, G. The experience and evolution of trust: Implications for cooperation and teamwork. *Academy of Management Rev.* 23, 3 (1998), 531–546.
- Huang, W., Olson, J., and Olson, G. Camera angle affects dominance in video-mediated communication. *Extended Abstracts on Human Factors in Computing Systems* (2002).
- Olson, G. and Olson, J. Distance matters. *Human-Computer Interaction* 15, 2-3 (2000), 139–178.
- Olson, G. et al. Designing flexible facilities for the support of collaboration. *Computer Augmented Teamwork: A Guided Tour*. G.R. Wagner (ed.) Van Nostrand Reinhold (1990).
- Olson, G. et al. The upper atmospheric research collaboratory (UARC). *ACM Interactions* 5, 3 (1998), 48–55.
- Olson, J. and Olson, G. *Working Together Apart: Collaboration over the Internet*. Morgan & Claypool Publishers (2013).
- Olson, J. et al. A theory of remote scientific collaboration. *Scientific Collaboration on the Internet*. G. Olson, A. Zimmerman, and N. Bos (eds.), MIT Press, Cambridge, MA (2008), 73–98.
- Olson, J., Olson, G., and Meader, D. Face-to-face group work compared to remote group work with and without video. *Video Mediated Communication*. K. Finn, A. Sellen, and S. Wilbur (eds.), Lawrence Erlbaum Associates, Hillsdale, NJ (1997), 157–172.
- Olson, J., Olson, G., Storösten, M., and Carter, M. How a group-editor changes the character of a design meeting as well as its outcome. In *Proceedings of the ACM Conf. on Computer-Supported Cooperative Work* (1992), 91–98.
- Olson, J., Wang, D., Olson, G., and Zhang, J. How people write together now: Beginning the investigation with advanced undergraduates in a project course. *Transactions on Computer-Human Interaction* 24, 1 (2017), 4:1–4:40.
- Scientific Collaboration on the Internet*. G. Olson, A. Zimmerman, and N. Bos, (eds.) MIT Press, Cambridge, MA (2008).
- The Science and Practice of Team Science*. D. Burley and M. Wang (eds), National Academies (2025).
- Veinott, E., Olson, J., Olson, G., and Xiaolan, F. Video helps remote work: Speakers who need to negotiate common ground benefit from seeing each other. In *Proceedings of the Conf. on Human Factors in Computing Systems* (1999), 302–309.
- Viegas, F., Wattenberg, M., and Dave, K. Studying cooperation and conflict between authors with history flow visualizations. In *Proceedings of the Conf. on Human Factors in Computing Systems* (2004), 575–582.
- Zheng, J. et al. Trust without touch: Jumpstarting long-distance trust with initial social activities. In *Proceedings of the Conf. on Human Factors in Computing Systems* (2002), 141–146.

Judith S. Olson is the Donald Bren Professor of Information and Computer Sciences Emerita in the Informatics Department at the University of California Irvine, Irvine, CA, USA, as well as Emerita Professor at the School of Information, University of Michigan, Ann Arbor, MI, USA.

Gary M. Olson is the Donald Bren Professor of Information and Computer Sciences Emeritus in the Informatics Department at the University of California Irvine, Irvine, CA, USA, as well as Emeritus Professor at the School of Information, University of Michigan, Ann Arbor, MI, USA.

 This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).

DOI:10.1145/3773294

Innovation, evaluation, and social impact in NLP has been fostered by NSF-designed and funded calls for research and training.

BY KATHLEEN MCKEOWN AND CHRISTOPHER D. MANNING

Natural Language Processing: How NSF Supports Visionary New Approaches to NLP

TODAY, LARGE TECHNOLOGY firms and AI companies advertise the remarkable *commonsense reasoning* abilities and *contextual understanding* of their large language models (LLMs). Less than a decade ago, however, when Yejin Choi^a was thinking about reviving commonsense reasoning, a seemingly intractable problem that had led to the downfall of symbolic AI nearly 20 years prior, her colleagues warned her that she was being foolish and putting her career at risk.

^a At that time, Yejin Choi was an associate professor at the University of Washington. Today, she is the Dieter Schwarz Foundation Professor at Stanford HAI and Computer Science.

At the time, huge progress had recently been made using artificial neural network (ANN) models in machine learning (ML) for important practical tasks. In 2016, for example, *The New York Times* ran a prominent feature on how Google Translate had been hugely improved—and simplified—by adopting the new state-of-the-art neural machine translation approach. Still, Choi was not to be deterred. By adopting a large-scale text data perspective, she and her students were able to exploit the first generation of work on LLMs to automatically generate a loosely structured commonsense knowledge base, COMET. She then went on to develop commonsense benchmarks, HellaSwag and Winogrande, that were cited as evaluations in OpenAI's 2023 announcement of GPT-4. Today, her work is foundational to the most advanced AI tools: Choi's breakthrough was the result of the funding she had received from the National Science Foundation (NSF) that allowed her to pursue unpopular—but ultimately transformational—research.

This is precisely the type of work the U.S. National Science Foundation (NSF) has historically enabled: visionary research in fields like natural language processing (NLP) that can yield major impacts on society. Even though it can take years of additional development, often within technology companies, the potential of such discoveries eventually reaches industry or the public. Sometimes breakthroughs come from NSF funding provided to individual faculty members or small-team awards, as in Choi's case. But transformative paradigm shifts within and across disciplines often depend on larger awards, such as the NSF Partnerships for International Research and Education (PIRE) award or the NSF AI Institutes recently established around the country. The NSF's impact is not limited to technology breakthroughs, of course. As we discuss in this article, NSF also funds Ph.D. students, via the NSF Graduate Research Fellowship Program (GRFP),



and junior faculty, especially through NSF CAREER awards, who will become the nation's trailblazing research scientists and industry engineers.

For more than 30 years, NSF funding has been an essential catalyst for the research that made possible the LLMs we see today. As far back as 1990, NSF began funding two important lines of research: NLP data collection and empirical methods. As large quantities of high-quality text data became more accessible, researchers began experimenting with a variety of probabilistic and other ML methods, leading to tremendous improvements in machine translation and to other NLP tools. By the 2010s, neural approaches emerged and in a few more years—once again with NSF support—led to the revolutionary LLMs coming from industry today: models such as OpenAI's ChatGPT, Anthropic's Claude, Google's Gemini, Meta's Llama, Alibaba's Qwen, and Deepseek's R1.

Just as important, NSF has for decades fostered NLP research creating technology with social impact, as one

of the few—if not the only—government agencies expressly championing the principle that the benefits of research should go beyond the academic community and reach society as a whole. This has led to research on technologies that improve access to information globally, address the mental and behavioral health crises in our country, and meet the needs of the most vulnerable, whether the elderly or the disenfranchised.

In this article, we highlight how crucial NSF support has been to both the foundational NLP research that made today's LLMs possible, as well as applications with important positive social impact. Synergy and balance between these lines of work is essential. For industry to flourish, it needs well-trained young researchers with big ideas and the skills to realize them—and that training is largely supported by the NSF. Conversely, without NSF funding, progress on NLP and LLMs would have been slower to develop and more constrained in its scope, given the comparatively limited and short-

term focus of industry. Moreover, there are many ethical issues posed by LLMs and agent-based systems: Will they be accessible to and usable by all segments of the population? Can we ensure the safety of humans as they interact with LLMs? Can LLMs be extended to handle all the languages of the world? Can they help us improve our health, our work, our learning, and our lived experiences? Or will they narrow our opportunities and limit innovation? As corporations drive adoption of LLMs and march toward artificial general intelligence (AGI), academic research unencumbered by corporate goals and profit margins is an indispensable complement to industry activities, providing the accurate and transparent evaluations necessary to ensure we develop LLMs that reflect our broader human and social values.

The Empirical Shift in NLP Research

The path to today's LLMs can be traced back 30 years to a new approach dubbed "Statistical NLP," initially



pioneered at IBM Research and AT&T Bell Labs, in which the field moved from symbolic, rule-based approaches to empirical, probabilistic, and ML approaches. Adopting ideas from information theory, which was the origin of language modeling, NLP research used these models heavily in text analysis and machine translation. This shift was characterized by greater reliance on large datasets, with a focus on systematically collecting and annotating data, leading to systems based on data and analysis of data. Foreshadowing current trends in NLP, ML demonstrated how we could learn from data to create more robust systems. Instead of prescriptive grammatical approaches, machine learning allowed us to observe how language is actually used, enabling interpretation of the informal language comprising most of human communication.

The importance of data. The first obstacle to pursuing statistical NLP research was that researchers, whether corporate or academic, did not yet have access to large amounts of digital language data, either text or speech. Thanks to seed funding from both the Defense Advanced Research Projects Agency (DARPA) and NSF, however, this problem was solved by the founding of the Linguistic Data Consortium (LDC) in 1992. Several years of ongoing NSF support sustained the LDC until the development of an independent funding model, based on dues from members and occasional government funding of specific projects. Key to LDC's origins were Fred Jelinek, then at IBM, who pushed for its establishment to broaden statistical NLP work, and Mark Liberman, a speech researcher who had recently moved from AT&T Bell Labs to the University of Pennsylvania, and who then spent the next 30-plus years building up and developing the LDC. The LDC was instrumental in the shift toward using text data to ground NLP research, enabling empirical NLP research at new companies and in academia. By centralizing contracts with data providers, such as news organizations, the LDC freed individual organizations (both universities and companies) from having to negotiate their own contracts. In contrast, today's practice of scraping data from the Web without acknowl-

edgment has led to copyright lawsuits (e.g., from *The New York Times*) that threaten both research and commercial enterprises. A further benefit of the LDC was its careful development of complete and tested annotation guidelines, ensuring data annotations had strong conceptual validity.

The development of statistical NLP: the Johns Hopkins Summer Workshops. Language models that predict which words are likely to appear in a given context were first heavily used in the speech community, starting with Jelinek's work at IBM. It was Jelinek's success in using LMs for speech that was the catalyst for applying them to NLP, which, through foundational language-model research funded by the NSF and others, enabled the precursors to today's NLP technologies.

One particular initiative played a vital role in the development of statistical NLP: the six-week-long Johns Hopkins Summer Workshops. Since 1999, they have been funded by the NSF to convene groups of faculty and students to work on specific hard research problems. The intellectual community and technologies developed in these workshops generate key insights and help disseminate their impact. The workshops have also educated hundreds of students through close research collaboration with faculty and established researchers, many of whom have become leaders in NLP research and major industry contributors to the commercial language-modeling technology used today.

Impact: *One of the 1999 workshops focused on the problem of building statistical machine translation (MT) systems. Over the course of a single workshop, participants developed techniques that showed it was possible to train an MT system "in a day." The legacy of this workshop is illustrated by the later achievements of its participants: Franz Och became the primary creator of Google Translate, while Yaser Al Onaizan managed groups at IBM and Amazon and is currently the deputy CEO of HUMAIN, a leading Saudi AI company. In 2006, a workshop led by Philipp Koehn resulted in the phrase-based, open source Moses machine-translation toolkit. Not only did it become the MT tool most used by academic NLP researchers for the next decade, but*

its approach formed the basis of early machine-translation systems developed at Google and Microsoft, transforming multilingual NLP worldwide, both inside and outside the academy.

Social impact: Meaningful access to world events. *As told by the first author, Kathy McKeown:* I got my start with an NSF-funded Presidential Early Career Award, which provided funding for five years of foundational research on language generation. My early foundational work on generation, as well as prior NSF-funded work by other people on interpretation and information extraction, convinced me the time was right for exploring text summarization, with text as both input and output. This was a new, yet important, problem for the field, given the information overload occurring as the Internet grew exponentially. In 1996, my first NSF proposal on text summarization was rejected as "too ambitious." When NSF launched the Stimulate program a few years later, with an emphasis on funding larger, more ambitious, collaborative projects, my team at Columbia (including Alfred Aho, Shih-Fu Chang, and Judith Klavans) received two grants on summarization that enabled us to look at multimodal summarization of both text and images, as well as statistical techniques for summarization that did not require fully interpreting the input text.

With this initial funding, we focused on summarization of news. We used multiple news articles on the same event to generate a short output paragraph that summarized the consensus across different articles about what had happened. Our initial research branched into Newsblaster (see Figure 1), a summarization platform that showcased our algorithms working on the news of the day. We launched Newsblaster in 2001, one year before the official launch of Google News. It was right after 9/11, and given the terrible events of that day, we knew people would want to follow that event over time. Newsblaster provided an online window into events around the world and allowed the tracking of an event over time, identifying new events spurred by the initial one. The platform allowed us to carry out research on many related topics.

For example, we developed approaches to generate summary updates that distinguished and highlighted new information from what was already known. We developed a multilingual version of Newsblaster in which the summaries were in English, but drew on—and linked back to—source articles in a variety of languages, enabling a more balanced and holistic view of events.

Impact: *Our work on summarization spurred many follow-on efforts. For example, one student, Dragomir Radev, went on to form large student groups as a faculty member at the University of Michigan (Ann Arbor) and subsequently Yale, to work on summarization. Some of his students likewise went on to form major summarization efforts (e.g., Alex Fabbri in Salesforce). We participated in the formation of the National Institute of Standards and Technology's (NIST) Document Understanding Conferences (DUC), which provided summarization data and task-based evaluations. This both accelerated research in the field and led to a leaderboard on summarization. Our foundational, NSF-supported work led to an avalanche of interest in the task. Summarization of a news article is today considered a solved problem and is leveraged commercially worldwide on a daily basis. Still, summarization of multiple news documents on the same event remains an active area of research, as does summarization of many other genres, such as legal documents, patient medical records, or narratives.*

The Shift to Neural Models in NLP

Just over a decade ago, neural network-based ML models emerged as a key NLP technique that not only increased accuracy on benchmarks but also opened new possibilities for learning and representing the semantics and pragmatics of human language use—that is, the meanings of words and phrases, and how they are affected by the linguistic and real-world contexts in which they are used. While this approach has antecedents to the early 2000s (or, indeed, before that), modern deep-learning work in NLP started about 2010 and came to dominate the field around 2014.

Neural methods enable representation of meaning. In 2010, another NLP breakthrough emerged at Johns

Hopkins University (JHU) through Tomáš Mikolov's work in neural language modeling. Funding from an NSF PIRE award—which supports large international group collaborations on foundational research topics—to JHU, Brown University, Charles University in Prague, and Saarland University in Saarbrücken, Germany, brought Mikolov, a student at Charles University, together with lead researcher Sanjeev Khudanpur at JHU. The project's goal was to go beyond syntax (the predominant focus of early work) to explore meaning representations and their applications in language models for speech recognition. Their work set a new state of the art in speech recognition and established beyond any doubt that recurrent neural networks outperformed the then-standard n -gram language models, which had been ubiquitous for the previous 30 years. Concurrently, at Stanford, this article's second author, Chris Manning, was approaching human language meaning even more directly, building neural models of meaning composition in phrases and sentences and generating superior benchmark results on paraphrase detection and sentiment analysis.

Impact: *These two lines of work led directly to the canonical word2vec and GloVe algorithms for word vectors (continuous meaning representations) in 2013–2014. These discoveries raised*

the performance of systems on nearly all NLP tasks and introduced the idea of subword modeling that became a stepping stone to early LLMs, such as ELMO, BERT, and GPT in 2017–2018.

Neural machine translation. The first (modern, successful) work in neural machine translation (NMT) was published by researchers at Google in 2014. However, it was the NMT work of Kyunghyun Cho, a postdoc at the University of Montreal, published in 2015, that transformed the field by adding a fundamental new mechanism into neural architectures: attention. The second author's (Chris Manning) lab quickly iterated on Cho's ideas, producing even better results using a different formulation of attention later that same year. Neural attention would go on to revolutionize how neural networks for language and beyond were built, and how scalable and powerful they could be.

In 2010, Cho moved from Korea to Aalto University in Finland because the Academy of Finland, an NSF counterpart in that country, had agreed to take a chance on his out-of-fashion doctoral research proposal on restricted Boltzmann machines. According to Cho: “If I have managed to make even a tiny bit of contribution to generative AI, it is because the Academy of Finland made a decision to support exploratory and unconventional research, and then I got similar

Figure 1. Newsblaster.





NSF support has been crucial to both the foundational NLP research that made today's LLMs possible, as well as applications with important positive social impact.

curiosity-driven foundational research support from CIFAR [the Canadian Institute for Advanced Research] and NSERC [the Natural Sciences and Engineering Research Council], in Canada." CIFAR and NSERC are NSF counterparts in Canada. Cho's career arc illustrates the importance of a global commitment to supporting and funding pioneering research ideas driven by scientists' insights and enthusiasm, rather than the necessarily short-sighted objectives of profit-focused companies and even mission-focused government organizations.

Impact: *Neural attention was the central innovation that made possible the development of the transformer neural network architecture. This highly scalable transformer architecture quickly led to large, powerful LLMs, and to this day is central to all leading LLMs, such as ChatGPT by OpenAI. It was a key development in the history of AI, the impact of which is still being realized.*

Language understanding: Understanding textual inferences. *Beyond NMT, the neural NLP work of the second author, Chris Manning, continued to evolve: I had worked for about a decade on machine translation, and neural machine translation was a great setting to develop early end-to-end neural models. The task could be viewed as simply neural sequence transduction by exploiting the large data resources of parallel text (translated text in two languages) that had been built up in the previous decade. As someone with a background in linguistics, however, I wanted to explore whether neural networks could be used for human language understanding and reasoning. Using NSF grants, my colleagues and I conducted initial studies looking at modeling language compositionality as a way to improve sentiment analysis. This task, which decides whether a text expresses positive or negative views about a topic, led to our building the Stanford Sentiment Treebank. I was also keen to try a neural approach to *natural language inference* (NLI), which focuses on determining the inferential relationships between sentences. For example, given a sentence such as *Street patrols deter crime*, a model is asked to determine if it "contradicts" (or "entails") the sentence *Street patrols are a catalyst for more**

crime. In this case, the correct inference is "contradicts." A neural-network approach demanded data, so we set about crowdsourcing the first large dataset of such examples, the Stanford Natural Language Inference dataset. Though we produced some initial results for compositional neural-network systems on this task, ours were quickly overtaken by others using the new neural attention approaches. This dataset raised awareness of the NLI problem across the neural network researcher community, most of whom were otherwise unaware of it because backgrounds in machine learning, statistics, or even physics, were more common than those in linguistics or philosophy.

Impact: *NLI became a leading benchmark task for NLP systems, and the lead Ph.D. student on the Stanford NLI dataset team, Sam Bowman, joined the faculty at New York University and produced not only a successor NLI dataset, but also the widely used General Language Understanding Evaluation (GLUE) benchmark. The GLUE benchmark, work led by a Ph.D. student funded by an NSF GRFP, combined the benchmarks of various academic researchers, including both the Stanford Sentiment Treebank and the Stanford NLI dataset. Five of its 10 tasks were NLI tasks. For the early pre-trained transformer years from BERT through T5 and ELECTRA (i.e., 2018–2020), GLUE was the standard evaluation of natural language understanding, directing and driving progress in industry labs until it was replaced in the 2020s by more difficult benchmarks as the abilities of LLMs greatly improved.*

Social impact: Mental and behavioral health. In the coming years, the demand for behavioral health professionals in the U.S. is predicted to see a more than 36% increase—more than triple that of other occupations. This signals an urgent need to address significant training and quality-assessment bottlenecks in one of the fastest-growing workforces in the U.S. At the University of Michigan, Rada Milhacea, a computer science professor, and a cross-disciplinary team of fellow researchers in computer science and health communication are developing NLP tools to support the growing demand for behavioral counseling to

address issues such as mental health, drug abuse, and weight loss. Through turn-by-turn language analysis (e.g., determining whether a statement is a question or a reflection) and the overall characteristics of a counseling session (e.g., measuring empathy), these researchers are developing methods that can provide automatic, scalable assessment of counseling performance. These approaches—which rely on joint neural and symbolic representations that bring together data-driven insights along with expert knowledge stored in knowledge bases—can provide real-time evaluation and feedback to counselors, as well as offer suggestions to improve the quality of interactions. This recent application of a neural approach illustrates its continued relevance and positive impact on social problems.

Impact: *The University of Michigan system has already been used in several public health courses and is being updated for use by counselors in Honduras and Mexico. These methods will help counselors, physicians, and health coaches deliver more effective, personalized care. By enabling ongoing automated feedback and coaching, the tools aim to improve counseling quality at scale while also advancing a new frontier of NLP research that supports an urgent need for real-world human support.*

The Shift to Large Language Models

LLMs rely heavily on foundational research from recent years, supported by NSF, which has been vital to their broader success. For instance, top-*p* sampling, a sampling technique designed to mitigate output degeneration and ill-formed text, was funded by NSF and is used by all major LLM platforms, both commercial and open source. Flash Attention, an NSF-supported innovation that accelerates the attention mechanism in transformers by twofold to threefold, has become the industry-standard implementation of attention in neural-network libraries.

NSF-sponsored research is also instrumental in accurately evaluating the capability and behavior of modern LLMs, contributing to widely recognized benchmarks such as Massive Multitask Language Understanding

(MMLU) and Graduate-Level Google-Proof Q&A (GPQA). MMLU includes roughly 16,000 multiple choice questions that test an LLM's knowledge across 57 different topics, from STEM to religion. GPQA tests the scientific and reasoning knowledge of LLMs on 448 multiple choice questions written by experts in fields such as biology and chemistry. In addition to testing knowledge, benchmarks also test performance on tasks, such as text summarization, including Multi-News—designed to test performance of summarization on multiple input news documents on the same topic—and SummEval, which was designed to test the performance of metrics against human judgments.

Building small LLMs in academia.

In 2019, Danqi Chen began her faculty career at Princeton as one of the few people actively trying to build neural language models in academia. Despite not having the compute resources available to even a startup AI company, she was committed to doing this work in academia, not only as a way to give smart young minds the opportunity to contribute to the field but also to ensure both openness and scientific rigor. Work like this lifts the broader innovation ecosystem by providing knowledge on which the whole community can build, rather than there being silos of knowledge within a few huge technology companies. Using relatively cheap structured pruning of an LLM, followed by further training, Chen's group did early influential work in producing Sheared Llama, small but well-performing 1- and 3-billion-parameter language models. More recently, she introduced new work in instruction-tuning LLMs during post-training, and developed a new, simpler offline preference optimization objective for instruction tuning, SimPO. At launch, her post training of Google's Gemma-2 model—gemma-2-9b-it-SimPO—was the strongest model, with less than 10 billion parameters in the Chatbot Arena evaluation. This work was funded by NSF grants to Chen, both her own CAREER award and another joint research award aimed at developing conceptual understanding of LLMs.

Impact: *Small language models are vital to making modern language*

models available on edge devices such as mobile phones. Chen's models, such as Sheared Llama 1.3B and Gemma-2-9B-it-SimPO, are widely used, each downloaded more than a million times. However, the larger impact is the influence of the research techniques introduced. Similar techniques to Sheared Llama have been used by Meta and Nvidia, and building small language models has become an increasing emphasis of other big technology companies, including Microsoft, Google, Alibaba, and Hugging Face.


Social impact: Identifying Black individuals at risk. *As told by the first author, Kathy McKeown: About 10 years ago, I was approached by social work professor Desmond Patton (now at University of Pennsylvania), who was studying the relationship between the digital and offline lives of Black youth. Over the years, our group, now also joined by a linguist and African American Language expert, Jessi Grieser (a professor at University of Michigan), developed techniques to identify online postings expressing emotional distress after traumatic events in Black communities as well as the events that triggered them, enabling valuable follow-up interventions by social workers to provide help. Using LLMs as tools to support Black individuals requires that the tools accurately interpret African American Language (AAL). Our evaluations of standard LLMs, however, show they cannot accurately interpret these communications, meaning that mental health and other tools developed using LLMs will fail to meet the needs of many communities. We were working on mitigations to enable better understanding when our NSF grant was abruptly terminated in April 2025 by the federal government.*

Impact: *The project has the potential for profound impact in society. Automatic flagging of posts expressing emotions of distress would be a unique and important way for social workers and mental health professionals to identify individuals at risk and is far more scalable than their current manual approaches. The creation of LLMs that do understand AAL would also open up other important applications (e.g., in health) to much wider use.*

Social impact: Adapting to the world's languages and people. Truly



Figure 2. Interleaving of molecules and language



inclusive and global systems must be able to understand languages other than English, but today's LLMs typically have a more difficult time with non-English languages and dialects, just as they do with AAL. Thanks to an NSF CAREER Award, Yulia Tsvetkov, a computer science professor at the University of Washington, launched her career with investigations into how well LLMs interpret non-English languages. Her work reveals that the tokenization strategies in commercial LLMs lead to higher costs and poorer performance for many non-English languages, making them far less effective for speakers in lower-income regions and exacerbating global digital inequities. Her team has introduced a large-scale benchmark covering 40 language clusters, 281 dialects, and spanning 10 NLP tasks to address these gaps. Furthermore, she has developed computational methods to analyze biases in Wikipedia biographies, revealing disparities in how people from different gender and racial groups are described across languages. Her grant was also abruptly terminated by the U.S. government in April 2025.

Impact: *Tsvetkov's benchmark enables more inclusive and linguistically robust NLP technologies. Her groundbreaking research on bias—now on hold—is being used to identify and reduce social disparities in widely used resources, such as Wikipedia content.*

AI institutes. In 2019, the NSF launched a call for AI institutes, supporting a new era of multi-institution collaborations that has accelerated U.S. research at the intersection of AI and a range of other disciplines.

At the University of Illinois at Urbana-Champaign, the NSF AI Institute on Molecule Synthesis (MMLI) has assembled an interdisciplinary team of computer scientists and chemists focused on AI for science, resulting in a new Modular Chemical Language

Model (mCLM) that trains on interleaved sequences of molecule functional modules and natural language descriptions of these functions (see Figure 2).

Impact: *The mCLM model has already improved the chemical functions of 430 FDA-approved drugs on the market, demonstrating how dramatically this model can improve our ability to understand and improve the potential of drugs. Other efforts on AI for Science using LLMs are simultaneously appearing across the U.S., for example at AI2, NYU, and MIT. This will be a major AI research area going forward.*

At North Carolina State University, the NSF EngageAI Institute has built a highly interdisciplinary team of education and AI researchers to transform STEM teaching and learning through the development of engaging, AI-driven narrative learning environments.

Impact: *By leveraging the power of stories to drive human learning, the Institute's work will give teachers the tools to create the transformative STEM education environments that will educate and empower the nation's next generation of STEM leaders.*

At Columbia, one thrust of our Institute on Artificial and Natural Intelligence (ARNI) is vision and language, where our interdisciplinary team of cognitive-science, natural-language, and vision researchers is innovating ways to help blind and low-vision people experience art. Currently, we are developing a novel multimodal system that can guide users through a personalized, immersive exploration of artwork through a combination of text, speech, and haptics, such as touch.

Impact: *In addition to driving improvements in assistive technology for artwork and other experiences for the 300 million individuals worldwide who are blind or have low vision, this project has generated a new metric for evaluating descriptions of artwork.*

Conclusion

Given the incredible capacities of current LLMs, some may believe further research in NLP is no longer needed, and remaining challenges will be well addressed by industry. Yet, as we have illustrated, NSF support for fundamental academic research has even recently resulted in advances on which current industry tools now depend. Over more than 30 years, innovation, evaluation, and social impact in NLP has been fostered by NSF-designed and funded calls for research and training, which have supported multiple generations of transformative NLP leaders. Loss of that scientifically driven force would disrupt and delay the next cycle of U.S. advances in the field.

Continued NSF funding would enable research needed to maintain the U.S.'s current fragile leadership in NLP: advances in evaluation, socially impactful AI, methodologies for efficient, small, and safe models, and unimagined applications for NLP, perhaps in applications to scientific research areas.^b Unlike industry, academia's unique environment facilitates interdisciplinary research,

^b We are aware of only limited data analysis for NLP specifically, but there are a number of analyses of AI research output, including Stanford HAI's AI Index (<https://hai.stanford.edu/ai-index/2025-ai-index-report>) and blog posts from Georgetown CSET's Emerging Technology Observatory (<https://eto.tech/blog/state-of-global-ai-research/>). These show that in terms of total publications in top AI venues, both China and Europe have been well ahead of the U.S. for more than a decade; in terms of citations of papers (a simple measure of research importance), China has been ahead of the U.S. since 2019 and Europe since 2022, while on measures of the most impactful 0.1–1% of work, the U.S. maintains a lead, though the gap with China is closing. Specifically for NLP, the position of the U.S. has been stronger, since China was distinctively strong in computer vision and robotics but not NLP. The U.S. still slightly led China in total NLP publications for 2017–2021. However, we particularly note that at the most recent July 2025 ACL conference (the highest-status NLP publication venue), papers with U.S.-first authors plummeted from 30% to 14%, while ones from China increased from 31% to 51%. For commercial LLMs, benchmarks still show the top U.S. companies leading, but China now has more LLM companies than the U.S. and the best ones (including DeepSeek and Alibaba) are producing models that are not far behind the best U.S. models.



enabling the pairing of research in NLP with any discipline across many schools on campus. Research with social impact is also well suited to academia; students find such research meaningful and are energized to participate. Faculty freedom in pursuing research means they can investigate exciting new modeling approaches that typically are not explored in industry. Today, methodologies for evaluation significantly lag LLMs' apparent capabilities. Without independent, transparent evaluation methods, we will continue to see unexpected, emerging capabilities that may lead to harm. Evaluation also points to avenues for future research. Likewise, exploring socially beneficial applications of AI and NLP often presents the type of genuinely hard problem that may have little commercial appeal at first, but proves revolutionary for the field and society at large. Expanding NLP tools to more language variants not only helps ensure *all* society benefits from its technical advances, but also can create new markets for these tools. Supporting academic research fosters talent that enables creative approaches to hard technical problems, advancing both our technological capacity and our shared human values. The NSF has begun funding for the second five years of its AI Institutes, and this will spur AI advances, but only if NSF continues to have funding.

Over the course of decades, the two of us have personally had the privilege of working with—and learning from—astonishing students at Columbia and Stanford. These students identify new research problems and develop sophisticated technical approaches, taking research in unexpected directions. Despite our extensive experience, every year these students continue to surprise us. At universities, students in NLP are afforded the unique opportunity to immerse themselves in interdisciplinary collaborations with a wide variety of experts in linguistics, literature, psychology, social science, medicine, machine learning, and computer systems. Our students become faculty at both research universities such as MIT, Berkeley, and Princeton as well as at undergraduate-centered institutions like Barnard, Williams, and Cooper Union. Or they are hired

by tech giants such as Google, Microsoft, Anthropic, OpenAI, Amazon, and Meta, critically influencing the directions these companies take as well as their products. As academic leaders at universities and core contributors at leading technology companies, their work has global reach and impact. We have seen this time and time again, with trained researchers going into industry from universities as well as programs such as the Johns Hopkins summer workshops. Industry leaders such as Yaser Al-Onaizan (formerly Amazon), Sam Bowman (Anthropic), Michel Galley (Microsoft), Sonal Gupta (Microsoft), Franz Och (formerly Google), Fernando Pereira (Google), Dan Roth (Oracle), and Luke Zettlemoyer (Meta), as well as individual research staff members, have benefited from NSF funding at some point in their career.

The creation of this rich talent pool is perhaps the most important product of NSF-funded university research and is core to the paradigm shifts in NLP that have led to today's most astonishing technologies.

The recent trend of NSF grant cancellation as well as diminished NSF resources threatens to disrupt these patterns of innovation and the development of NLP talent. While the two canceled grants we mentioned here will undermine the reach of LLMs and their products beyond mainstream English speakers, other grants addressing purely technical NLP topics (e.g., on statistical bias) have also been canceled or indefinitely postponed. Without an adequate number of NSF staff to review proposals, new awards will be dramatically curtailed. Most importantly, what future research innovations and cutting-edge technology companies will never exist because we have cut off the relatively modest public investment that has made them possible up to this point?

Time and time again, we've seen academics come up with new ideas for research that were initially deemed "crazy" even by their peers. But these ideas have often led to remarkable new technical advances and tasks for NLP, and they would not have happened for decades—if at all—without the unique funding structure offered by the NSF. We gave a few illustrative

examples, but the amount of impactful NLP research that NSF has funded is surprising, leading to the training of thousands of Ph.D. students in the past 10 years alone. While it is hard, in the present moment, to imagine where we might be in 20 years, as longstanding researchers in the field we are confident of one thing: Somewhere during that time, a faculty member or student will come along with a seemingly preposterous research direction that—if provided with the type of intellectual freedom and financial support the NSF has provided over the past three decades—will lead us into another new, exciting, and presently unimaginable future.

Acknowledgments

This article benefited from the contributions of many individuals who described their NSF-funded efforts. Our thanks goes to Mohit Bansal, Chris Callison-Burch, Danqi Chen, Kyunghyun Cho, Yejin Choi, Jason Eisner, Sanjeev Khudanpur, Heng Ji, Mark Liberman, Rada Mihalcea, and Yulia Tsvetkov for their willingness to contribute by providing advice and actual text. We also thank Susan McGregor, a research scholar at Columbia's Data Science Institute with a background in journalism, for extensive editing that made our points more forcefully, and Susan Ellington, associate vice president of public affairs at Columbia University, for her critique and suggested edits. We thank Melanie Subbiah, Amith Ananthram and Yunfan Zhang for comments on an early draft that helped shape our thinking, and the *Communications* reviewers and editors for their thoughtful comments later on. 

Kathleen R. McKeown is the Henry and Gertrude Rothschild Professor of Computer Science at Columbia University and is also the founding director of the Data Science Institute at Columbia University, New York, NY, USA.

Christopher Manning is the inaugural Thomas M. Siebel Professor in Machine Learning in the Departments of Linguistics and Computer Science at Stanford University and a co-founder and Senior Fellow of the Stanford Institute for Human-Centered Artificial Intelligence (HAI), Stanford, CA, USA.

 This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).

DOI:10.1145/3764111

From federally funded academic research to worldwide impact on computing.

BY VIKRAM S. ADVE AND CHRIS LATTNER

The LLVM Compiler Infrastructure

LLVM-GENERATED CODE IS used every day by billions of people on billions of devices—an outcome that can be clearly traced to research funded by the National Science Foundation (NSF) starting in the year 2000. LLVM powers the core development tools, operating systems, and most applications at Apple Computer, where it long ago replaced the GNU Compiler Collection (GCC) as the foundation of the company's software. Google uses LLVM for the Android operating system and Android NDK [native development kit]—impacting the system software on all Android devices as well. Applications that run on Qualcomm's mobile processors (the Snapdragon family) are built in substantial part using LLVM, which further extends LLVM's reach in the Android mobile phone industry. The two leading processor companies for personal computers and servers—ARM and Intel—abandoned their longstanding proprietary compiler systems (ARMCC and ICC) and replaced them with toolchains built on LLVM, which means important parts

of the world's desktop and server software are compiled in part by LLVM. All of Google's and Meta's datacenter C/C++ code is compiled using LLVM, which means some widely used cloud services and applications—for example, Google Search, Facebook Feeds, Instagram, and more—are powered in part by LLVM. Key gaming platforms—specifically Sony's Playstation 4 and Playstation 5 and the Nintendo Switch—use LLVM as the primary compiler toolchain for their applications. Even the rapidly growing field of AI software benefits widely from LLVM, through Nvidia's CUDA compiler, the Triton compiler used by PyTorch and OpenAI, and even more extensively through the MLIR Compiler Infrastructure, which is also built on top of LLVM and used by several major AI frameworks. Collectively, these outcomes imply that software products compiled using LLVM represent hundreds of billions of dollars in annual revenue across most major classes of modern computing: mobile, cloud, desktops, servers, gaming, supercomputers, and AI. It may not be an exaggeration to say that *the majority of end users across the computing industry benefit from software compiled by LLVM.*

LLVM here refers to the LLVM Compiler Infrastructure, an open source software system created to simplify the design and implementation of compilers and a wide range of compiler-based tools. LLVM (which originally expanded to *low-level virtual machine* but is now just a name for the overall system) began in the year 2000 at the University of Illinois Urbana-Champaign (UIUC), as a joint research project of the two authors. The LLVM infrastructure was first released as open source software in October 2003 (under the liberal, three-clause University of Illinois/NCSA Open Source License, later replaced by the Apache 2.0 License with LLVM Exceptions). Since then, LLVM has evolved from

IMAGE BY ANDRIJ BORYS ASSOCIATES, USING SHUTTERSTOCK.COM, AND LOGO FROM LLVM.ORG/LOGO.HTML





an academic research project into a cornerstone of modern computing, enabling software products and technological innovations across many computing industries.

On the academic front, LLVM has also been extremely influential because it has democratized compiler projects for both teaching and research. The LLVM infrastructure is routinely used by a wide range of undergraduate and graduate compiler courses worldwide, and high-schoolers use LLVM to do programming-language projects for “fun!” LLVM is also used in research projects spanning many areas, including compilers, computer architecture, software engineering, software security, hardware/FPGA synthesis, accelerator design, parallel computing, formal verification, and even quantum computing.

Another critical dimension of LLVM’s impact is on widely used open source software platforms, including popular programming languages and other software. LLVM underlies several widely used programming languages, including Swift (the dominant language for iOS mobile apps), Rust (for system and other low-level software), Julia (for high-performance computing), Halide (for high-performance image processing used in highly popular worldwide services, such as YouTube, Flickr, and Adobe products), and Mojo (AI models at Modular Inc.). These arguably would never have happened without an easy-to-target back end and a highly flexible and modular compiler architecture.

LLVM is the foundation of Clang, the C/C++/Objective C/Objective C++ compiler used across nearly all computer systems where LLVM is used. LLVM is used to build the popular Google Chrome browser on all its host platforms.³⁸ WebAssembly (WASM), the fast-growing language for running high-performance native code securely in browsers and other systems, is primarily targeted by LLVM, through the widely used Emscripten compiler toolchain.^a There

are numerous other surprising uses of LLVM, including for OpenGL shaders in the MacOS graphics pipeline, database query optimization in PostgreSQL, image-processing pipelines in Adobe After Effects, and firmware for microcontrollers.

The reasons for LLVM’s impact can be traced to the key conceptual and technology gaps that LLVM filled. In the early 2000s, when LLVM began, the major open source compiler in production use was GCC, which lacked some of the important modern compiler techniques, like Static Single Assignment (SSA) form, interprocedural link-time optimization (LTO),²⁷ just-in-time compilation, and a flexible and clean pass pipeline. LLVM was designed to fill all these gaps, leading it to become very widely used for academic research projects. The most technically novel aspect of LLVM was—and is—that it combines both the strong, ahead-of-time optimizing compilation capabilities of traditional high-performance, statically compiled languages and the self-contained mobile code representation and flexible just-in-time and dynamic compilation capabilities of managed languages; moreover, LLVM delivers all these capabilities for essentially all languages, a feature that is unique to this day.^{2,28} Nevertheless, perhaps the most important factor in LLVM’s success was not the novel technical capabilities but rather the highly modular and reusable library architecture with well-designed interfaces and separation of concerns, in place of the monolithic structure common in open source compiler implementations.²⁶ This flexibility made it far easier to reuse components of LLVM in an enormous range of compiler-based projects—not just a broad variety of statically and dynamically compiled languages but going far beyond anything originally envisaged by the authors, from JITs for graphics shaders to sandboxed mobile code in Web browsers, tools for circuit design, symbolic execution engines, compilers for quantum computing, and (for many years) the mobile code representation for nearly all software apps in all Apple mobile devices.

This article describes the context and history of the LLVM project, the federal funding that made the project possible, key characteristics of LLVM design and software architecture and its major innovations, and the broad public impact of the LLVM system, and briefly lists ongoing work to expand modern compiler technology that derive from the success of LLVM.

History of the LLVM Project and Importance of Federal Funding

Context. Compilers are a lynchpin technology for all of modern computing because of their primary role: translating all human-written programs to machine-executable code. GCC, one of the most impactful open source software projects in all of computing, was launched in the 1980s and soon became a foundational technology for many hardware and software systems, including all Linux systems, all Apple computers, many Google services, numerous open source software projects, and others.²⁰ GCC also became popular for compiler research for classical compiler problems in program optimization and code generation for imperative languages, alongside specialized research infrastructures such as the Jikes Research Virtual Machine for Java,³ Polaris⁶ and SUIF³⁹ for automatic parallelization, and others.

However, by the year 2000, the design of GCC was becoming outdated. GCC lagged competing compilers in multiple ways, even for its primary use case of static compilation, including using C instead of modern object-oriented languages like C++ or Java, the lack of cross-module interprocedural optimization, and the initial lack of support for SSA form,¹¹ which first appeared only in 2005 in GCC 4.0.^b Perhaps more serious over the long term, the design of GCC is too monolithic and inflexible for most purposes. It was not practical to extract and reuse pieces without pulling in most of the compiler,

a See V8 Emscripten Team. Emscripten and the LLVM WebAssembly backend (2019); <https://v8.dev/blog/emscripten-llvm-wasm>.

b See the GNU Compiler Collection Project. GCC 4.0 Release Series: Changes, New Features, and Fixes; <https://gcc.gnu.org/gcc-4.0/changes.html>.

preventing uses for load-time or just-in-time (JIT) compilation of mobile code, embedded scripting languages, sandboxing browser extensions, and many other uses. This greatly limits GCC's use cases for both academic research and production tools, with the exception of classical compiler-focused goals.

Research infrastructures for managed languages, like JikesRVM,³ Mono,^{5,21} and Roslyn,^c were also very successful within their domain, delivering major advances in garbage collection, JIT compilation, and mobile code execution. However, these runtime systems impose numerous language-level semantics and restrictions, which preclude them from being used for the vast majority of other programming languages and even systems that require mobile code and load-time or JIT compilation, such as for graphics shaders, CUDA, and OpenCL.

Research and importance of federal funding. The first author, starting a new research program at UIUC, had the broad goal of designing flexible dynamic-compilation techniques for arbitrary programming languages, including static, dynamic, and scripting languages. His CAREER proposal to the NSF's Next Generation Software (NGS) program, submitted in July 2000,¹ addressed techniques to make dynamic compilation efficient, powerful, and flexible, and was directly motivated by the limitations of existing runtime compilation techniques, including those in managed languages. This award would prove pivotal to the group's early work on LLVM.

The LLVM research project began in Fall 2000, with the broad goal of designing a highly flexible compilation strategy that would support static and dynamic compilation for arbitrary programming languages. The second author joined the UIUC Ph.D. program in Fall 2000. During that semester, the authors extensively discussed designing a new compiler infrastructure that would be far more flexible than existing static

and managed-language compilers, and enable both static and dynamic compilation for arbitrary languages. *Redesigning an entire compiler system from scratch is an academic luxury that enables building on state-of-the-art design techniques from the literature.* Key design goals the authors discussed included:

- ▶ A clean, completely language- and hardware-independent intermediate representation called LLVM IR, based on partial SSA form integrated with a control flow graph (CFG) per function

- ▶ A fully executable IR definition that could be exported and used directly as a fully self-contained program representation, with no excess language-specific baggage—i.e., *for all programming languages*—unlike the popular Java bytecode or MSIL

- ▶ First-class support for cross-module compilation using the offline IR representation for link-time optimization (LTO) and delayed, static machine code generation

- ▶ First-class support for JIT compilation, through the *same offline code representation* and delayed, dynamic machine-code generation

- ▶ A highly flexible compiler pass pipeline and overall compiler architecture based on strong adherence to the principle of Separation of Concerns in compiler design⁴

Many of these design principles and related papers—especially SSA form,¹¹ Separation of Concerns,⁴ highly retargetable code generation, and interprocedural compilation—were discussed in the first author's Advanced Compiler Construction class taught that fall semester. Although they were well-known techniques in the compiler community, no existing compiler system we know of adopted all these principles nor used them as pervasively as LLVM did.

Perhaps the most impactful design feature of the LLVM infrastructure proved to be the highly modular library-based architecture, which in combination with the Separation of Principles compiler design, enabled a huge range of use cases, many of which were never envisaged during the early years of the research project. Some of these uses cases are



It may not be an exaggeration to say that the majority of end users across the computing industry benefit from software compiled by LLVM.



^c See Microsoft Visual Studio Team, ".NET Compiler Platform ("Roslyn") extensibility (2024); <https://tinyurl.com/22zw9e2u>



The direct federal funding for the LLVM research project was essential for its early development and open source distribution, and supported continued enhancements during the early years.

described later in this article.

The detailed design and implementation of the LLVM infrastructure was begun by the second author during the winter break after the fall 2000 semester. When he returned from the break, he already had skeleton working code that included the core principles of the LLVM design, including the highly modular library structure, a clean IR with fully equivalent in-memory, offline textual, and offline binary forms; initial tools to transform between the different IR forms; and the beginnings of a flexible pass pipeline with an elegant Separation-of-Concerns architecture. Remarkably, although the code base has changed almost completely and grown by orders of magnitude since those early days, all those initial design features are clearly visible in the architecture of LLVM today.

The direct federal funding for the LLVM research project was essential for its early development and open source distribution, and supported continued enhancements during the early years. The first author's CAREER proposal supported both authors during most of their work on the project except for the initial year. This stage of research included the early papers on LLVM^{2,7,9,32} and the first release of the open source software system in December 2003. Additional funding from several other federal programs enabled us to broaden the research efforts in the LLVM project, including work on virtual instruction-set architectures,^{2,7,9,17} interprocedural pointer analysis and optimization,^{29,30,32} lightweight compiler-based memory safety for C programs and system software,^{12,16,24} and a safe virtualized execution environment (SVA) for a commodity operating system (OS) kernel, such as the entire Linux kernel, to prevent security exploits on kernel code.⁸

The primary LLVM paper at CGO 2004 won a retrospective Most Influential Paper Award 10 years later, and the SVA paper won an Audience Choice Paper Award in SOSP 2007. We believe the CGO 2004 paper is likely one of the most highly cited papers in the field of compilers, with *more than 8,000 citations today*.

Less than nine years after its initial public release, LLVM's worldwide impact was already high enough that the authors of LLVM were awarded the ACM Software System Award in 2012, which is given by ACM to only a single software system worldwide every year.

Release and early adoption. The LLVM system was first released in open source form in December 2003, under the University of Illinois/NCSA Open Source License. A major priority was to grow and support an external community of developers and users, and attract active external contributors. The second author joined Apple Computer after his Ph.D., in large part because Apple allowed him to continue working on and contributing to the open source LLVM system. His tenure at Apple was remarkably impactful in fostering commercial adoption of LLVM, both at Apple and elsewhere. Apple was largely a GCC shop at the time, with their XCode IDE and compilers for both system code and applications all compiled using GCC (for C/C++/Objective C/Objective C++). The first commercial product to adopt LLVM, in 2006, was the MacOS OpenGL graphics stack, replacing a custom JIT (which the graphics team had to write and maintain) with an LLVM-based JIT, using LLVM bytecode for shipping and distributing shaders. LLVM was soon incorporated into XCode as well, but the big shift toward LLVM happened when the second author, with Apple's support, developed Clang, a fully open source alternative to GCC for C/C++. Apple went on to completely replace GCC with Clang/LLVM for all its system software and applications, across all its desktop, laptop, and mobile devices. Google, ARM, Qualcomm, Sony, and many other companies soon began to adopt Clang/LLVM, as well.

Technical Overview of the LLVM Compiler Infrastructure

LLVM is a compiler framework that aims to make lifelong program analysis and transformation available for arbitrary software. LLVM achieves this capability through two design features: a *code representation*—the



LLVM IR—that serves as a common basis for analysis, transformation, and code distribution; and a *compilation strategy* that exploits this representation to provide a combination of five capabilities that is not available in any other compilation system we know of:

1. Persistent program information: A compiler built using LLVM can choose to preserve the LLVM code representation throughout an application's lifetime, enabling sophisticated optimizations to be performed at all stages, including compile time, link time, load time, runtime, and even “idle time,” that is, between runs on the end-user's system.

2. Offline code generation: Regardless of the choice made on persistence, it is possible to compile programs into efficient native machine code offline using powerful code-generation techniques too expensive for runtime code generation. This is crucial for performance-critical programs.

3. User-specific profiling and optimization: The LLVM framework can gather profile data at runtime for deployed software so it is representative of the end user's usage patterns, and then use it for profile-guided optimizations both at runtime and in idle time.

4. Transparent runtime model: LLVM does not specify any particular object model, exception semantics, or runtime environment, so it benefits any language (or combination of languages).

5. Uniform, whole-program compilation: Language independence makes it possible to optimize and compile all code comprising an application in a uniform manner (after linking), including language-specific runtime libraries and system libraries, and potentially even OS kernel components in carefully designed circumstances, such as a closed embedded system.

We believe *no other system to date provides all five properties*, a claim we made in the CGO 2004 paper that continues to be true today, more than 20 years later. This is obvious for most compiler alternatives. The closest in practice are high-level virtual

machines for managed languages, such as JVM or .NET's CLI: These provide #3 and partially provide #1 and #5, but do not provide #4, which greatly limits their applicability to many languages. Moreover, if these managed languages provide #2, they do it at the exclusion of #1 and #3. Other mobile-code strategies, such as GPU virtual instruction sets (CUDA's PTX and OpenCL's SPIR-V) and JITs for Javascript, Python, and WebAssembly, provide some combination of #1 and #3 but not the rest.

The LLVM intermediate representation. The key reasons this novel combination of capabilities is possible lie in the design choices made in LLVM IR. The IR is designed to be a fully self-contained, executable, hardware- and language-independent *virtual instruction set architecture*² that is rich enough to support sophisticated compiler analyses and transformations, yet low-level enough to be language-neutral and to support *all* compiled software, including OS, kernel, and application code. There are also some limitations of the design incurred to achieve the capabilities above.

The LLVM IR captures the key operations of ordinary processors but avoids machine-specific constraints, such as physical registers, memory-addressing modes, and low-level calling conventions. LLVM IR uses SSA form with ϕ operations represented by an explicit Phi instruction. SSA form simplifies many dataflow optimizations on local variables by making the flow of values from definitions to uses explicit, and many reordering transformations by eliminating spurious antidependences and output dependencies on those variables. LLVM supports vectors as first-class values, including instruction operands and vector memory transfers, which is crucial in supporting short-SIMD instructions in modern processors, such as Intel's SSE, AVX and VNNI, ARM Neon, and others. LLVM vector instructions can be generated by parallel language front ends (e.g., for OpenMP, Fortran, Halide, or Julia), for tensor operations in machine learning applications, or by automatic vectorization passes in LLVM

based on data-dependence analysis. LLVM IR supports *intrinsic functions*, which are built-in functions with predefined semantics, as a powerful extensibility mechanism. They behave exactly like ordinary function calls so most passes can safely be unaware of them, but selected passes and code generators can recognize specific intrinsics and implement the desired semantics. This flexibility has allowed a huge range of new IR operations and language- or hardware-specific features to be supported in LLVM, many in the official releases and far more in the numerous experimental research projects that use LLVM.

The LLVM IR language has fully self-contained, executable semantics, which means complete executable programs can be represented and shipped in LLVM IR form, and can be executed via an interpreter or by translation to machine code. This feature is what makes “late-stage” compilation possible, by persisting the LLVM IR code for link-time or later compilation steps.

One price of widespread adoption is high IR complexity. LLVM IR launched in 2003 with fewer than 35 polymorphic operations; it has now grown to more than 175 (including in excess of 110 vector operations defined as intrinsics for convenience), as well as hundreds of other semantically significant intrinsics for garbage collection, memory ordering, synchronization, stack manipulation, and others. While most intrinsics can be ignored for tasks such as building language front ends, optimization passes, or back-end passes, the complexity greatly complicates efforts such as building a new back end or defining a formal semantics for LLVM.

LLVM compiler architecture. The figure shows a schematic depicting the major compilation flows in the LLVM infrastructure. Traditional compilers built using LLVM, such as Clang, Rustc, Flang, Juliac, and many others, are able to “export” the LLVM IR for each source file. An LLVM IR linker combines all IR modules for an application and also any of its libraries that were exported to IR modules into a unified (perhaps



complete) LLVM program, and then performs interprocedural analysis and optimization on that program (LTO),¹⁹ followed by back-end target code generation.

Several alternative, non-traditional compilation flows made possible by LLVM are used in production systems. Perhaps the most extensive use has been for Apple's mobile apps, which for many years (mid-2010s through 2022) were shipped by developers to Apple's App Store in LLVM IR form and compiled down to an end-user device in the App Store. This enabled Apple to support multiple processor architectures across its diverse device lineup, including older iPhones running 32-bit ARM processors and newer iPhones with 64-bit ARM chips, Apple Watches with specialized low-power ARM variants, and Apple TVs with their own architectural requirements. This strategy simplified mobile-app development, reduced app sizes for end users, enabled easier deployment of new compiler optimizations and hardware features, and enabled more effective enforcement of API usage restrictions by analyzing explicit LLVM function calls instead of arbitrary control transfers in machine code.

Another compilation flow widely used in production is to support a variety of mobile code formats via LLVM-based load-time or runtime code generation, for example, Nvidia's PTX, Python, or WebAssembly. LLVM greatly reduced the cost and complexity of these systems by avoiding the need for custom code generators.

None of these non-traditional compilation flows is supported by any other compiler infrastructure we

know, not even those with runtime compilation capabilities, such as JVM or .NET. This is because none of these systems has a sufficiently language-independent shipping-code format to support, for instance, CUDA or WebAssembly.

Major External Impact of LLVM

The most important and pervasive impact of LLVM to date has undoubtedly been its impact on industry products worldwide. The impact on the practice of computing is also important because it carries long-term benefits spanning industries and shows the potential for continuing changes. And the intellectual impact in terms of research and education carries even more long-term significance.

Impact of LLVM on industry.

The economic impact of LLVM spans nearly the entire range of the computing industry: mobile, cloud, desktop, high-performance computing, gaming, and increasingly AI. The highlights were described at the outset of the introduction; here, we briefly expand on some of the key reasons—and a few of the major milestones—toward that impact.

Mobile computing. The most widespread impact of LLVM is in the mobile devices people use daily worldwide. Apple's wholesale adoption of LLVM (replacing GCC) happened in stages: initially for the JIT compilers for OpenGL graphics shaders in MacOS; then in XCode for optimization and code generation; and later, the development of Clang and the Swift programming language (both designed around LLVM). For many years, Apple's mobile apps leveraged LLVM's "lifelong" compilation capabilities when shipped as LLVM to the


Apple App Store. Android development also relies on LLVM through the Android NDK.

Cloud computing. To our knowledge, all of Google's and Meta's data-center C/C++ code is built by LLVM. Google in particular (like Apple before them) migrated from GCC to Clang for C++ software, giving them better tooling, developer productivity, compile times, code coverage, and performance. For several years, Clang replaced GCC only in development (because of a far better developer experience) but not for shipping production code (because of slightly lagging performance). But these gaps closed over time.

Desktop and server computing. It is remarkable that Intel and ARM adopted LLVM because both companies have longstanding investments in proprietary compiler toolchains—ICC and ARMCC, respectively. Intel invested strongly in lifting LLVM code-generation quality on Intel processors, bringing it up to par with Intel's compilers. Intel eventually switched its production C/C++ compiler stack to Clang and LLVM across its processor lines.³⁶ ARM was one of the early adopters of LLVM. Notably, ARM's concerns about the lack of patent protections in the original LLVM license motivated a complex, multiyear effort spanning legal departments across many companies and universities to reissue LLVM under the Apache 2.0 License, which provides strong indemnity against unintentional patent infringement.

Gaming consoles. Sony was an early adopter of LLVM and uses Clang as the primary compiler for both PlayStation 4 and PlayStation 5 development. Nintendo made a remarkable transition from proprietary tools to

Figure. Schematic of LLVM compiler architecture showing static, just-in-time, and other compilation flows.



LLVM-based infrastructure for the Switch.

Artificial intelligence. The newest frontier of LLVM's impact comes in the field of AI. First, the CUDA-based cuDNN library is widely used by AI frameworks, including PyTorch, TensorFlow, and TVM, which creates a direct reliance on LLVM through the CUDA toolchain. Second, multilevel intermediate representation (MLIR), used in most major AI frameworks (TensorFlow's XLA compiler, parts of PyTorch, the OpenXLA ecosystem, and even components of Nvidia's CUDA toolkit) all use LLVM for core back-end functions, such as register allocation and scheduling. It also lowers tensor operations to LLVM IR in many compiler flows.

Impact of LLVM on the practice of computing. LLVM influence in the long term will derive even more from its role in the emergence of new programming languages, Web technologies, and more specialized computing domains. Rust, the systems programming language gaining rapid adoption for its memory-safety guarantees, depends entirely on LLVM for code generation across all platforms. Julia, designed for high-performance scientific computing, leverages LLVM's JIT compilation capabilities to achieve performance rivaling traditional compiled languages. LLVM also enables WebAssembly, the technology that allows languages like C++, Rust, and others to run at near-native speed in Web browsers—fundamentally changing what is possible in Web applications.

LLVM has enabled innovations in specialized computing domains that would have been prohibitively expensive to develop from scratch. The PostgreSQL database uses LLVM for JIT compilation of queries, dramatically improving performance for complex analytical workloads. Quantum computing companies build their programming languages and compilers on LLVM infrastructure, as do developers of hardware verification tools. Xilinx uses LLVM heavily in its FPGA compilation tool flows, and the CIRCT project is also extending MLIR into hardware design.

Beyond these high-profile ap-

plications, LLVM has enabled a remarkable diversity of specialized and often surprising use cases that showcase its versatility. These use cases span graphics optimizations, audio DSP languages and “audio shaders,” systems languages used for firmware in embedded systems, and domain-specific languages in research use from computational biology to high-energy physics. LLVM has even democratized the creation of embedded scripting languages within complex software, which often need some form of user scripting or automation: LLVM makes it trivial to transform a simple interpreter into a high-performance JIT compiler.

Impact of LLVM on research and government. LLVM's influence extends far beyond commercial applications into the important realms of academic research and government-funded scientific computing. National laboratories, including Los Alamos, Lawrence Livermore, and Sandia, have embraced LLVM as a core technology for exascale computing. Initiatives such as “Kitsune,” a Los Alamos-led effort for the Exascale Computing Project, extends LLVM to support the massive parallel computing requirements needed for nuclear weapons simulation, climate modeling, and other computationally intensive scientific areas. The National Nuclear Security Administration has partnered with Nvidia and national laboratories to develop Flang, an open source Fortran compiler built on LLVM, explicitly stating, “Large HPC applications ... require a common compiler infrastructure that supports both C/C++ and Fortran.”²⁵

In academic circles, LLVM has democratized both compiler education and research. Traditional compiler courses used to build simplistic compilers mostly from scratch, with limited ability to compile realistic benchmarks like the large programs in SPEC or others. In contrast, the modularity, flexibility, and ease-of-use of LLVM enables both simple and powerful course projects to be accomplished by one or two students in a single semester. Universities worldwide now teach compiler



By combining novel compiler design with reusability, modularity, and community-driven development, LLVM has made software and hardware innovation easier, faster, and more accessible.





The compiler infrastructure that began in 2000 as a federally funded two-person research project at the University of Illinois now supports multibillion dollar hardware, software, and services industries worldwide.

construction courses using LLVM as the foundation, with numerous undergraduate and graduate compiler courses centered on LLVM. Remarkably, we have met multiple freshmen who, as high school students, had explored compiler projects using LLVM and entered college with non-trivial compiler experience, which is traditionally an advanced computer-science college topic!

LLVM has lowered the barrier to compiler research even more dramatically, enabling graduate students to focus on novel compiler algorithms and compiler use cases far beyond what GCC could support. The highly flexible and modular architecture of LLVM makes it possible for a research project to develop simple new compiler passes in a day and test them on large, complex applications immediately. LLVM enables research projects to be evaluated on realistic workloads, easing transition of research results to production uses. Publications using LLVM are also regularly published in numerous other areas of computer science, not only compiler-adjacent areas such as programming languages, software engineering, formal verification, and computer architecture, but also less related areas such as computer security,^{34,37} hardware synthesis,⁴⁰ and even important compilers for quantum computing research.²² A small fraction of the vast number of published research papers that use LLVM are listed at the LLVM website.^d

Ongoing and Future Directions with LLVM

The LLVM Compiler Infrastructure continues to evolve, expand, and grow. On the research front, research projects are developing powerful new compiler capabilities. For example, the ALIVE2 project has developed formal correctness guarantees for compiler passes through translation validation using LLVM.³³ The HPVM project has developed a highly retargetable compiler system for diverse het-

erogeneous systems by adding a hierarchical graph representation of parallel programs, with LLVM IR as the core program representation for both static and dynamic compilation.¹⁸ The MLIR project continues to evolve numerous compiler components (called “dialects” in MLIR parlance) for both custom hardware accelerators and diverse AI compiler systems.³¹ The Hydride and MISAAL projects are showing that core compiler components can be automatically generated directly from vendor specifications of ISA functionality, including the hardware-independent LLVM IR, front-end translators, and back-end code generators for multiple hardware instruction sets. These autogenerated compilers can match or beat the performance of even a highly engineered production compiler for Halide despite requiring an order-of-magnitude less manual effort.^{23,35} And LLVM has recently been used to fine-tune the “LLM Compiler,”¹⁰ a suite of open, pre-trained LLMs specifically designed for compiler tasks by fine-tuning CodeLlama on more than 537 billion tokens of LLVM IR and assembly code. The enormous volume of LLVM-compiled languages and open source LLVM-compiled projects makes it possible to create such a huge data corpus for successfully training or fine-tuning LLMs.

Conclusion

The compiler infrastructure that began in 2000 as a federally funded two-person research project at the University of Illinois now supports multibillion-dollar hardware, software, and services industries worldwide; powers some of the programming languages of the future; enables mission-critical research at national laboratories; and supports the computational foundations for ongoing scientific and AI advances. The story of LLVM illustrates how foundational infrastructure projects can have an outsized influence on technological progress. By combining novel compiler design with reusability, modularity, and community-driven development, LLVM has made software and hardware innovation easier, faster,

^d See the LLVM Infrastructure Community. LLVM Related Publications; <https://llvm.org/pubs/>.

and more accessible across most major computing sectors.

Several essential features of the LLVM project were critical to its success. First, federal funding—especially NSF’s strategy of funding early-stage blue-sky research with unknowable outcomes while encouraging open source artifacts—was essential to making the initial academic research and the infrastructure possible. Second, an academic research effort was key to identifying state-of-the-art compiler technologies and extending them with novel and powerful capabilities. Third, distributing a commercially neutral open source infrastructure with a liberal license led to it being used over and over again in pervasive commercial products across nearly all major subsectors of computing. Fourth, the widespread adoption of the system was driven by strong technical and community leadership by the second author, first at Apple for numerous products, then (through strategic community coordination) at other large companies. Finally, the breadth of adoption was augmented powerfully by derivative systems that were themselves highly successful, including Clang, Swift, and MLIR, all led by the second author. All told, a small federally funded academic research project has resulted in powerful innovations and world-changing, foundational infrastructure.

Acknowledgments

The original research on the LLVM Compiler Infrastructure project was funded by the National Science Foundation (NSF) Next Generation Software (NGS) Program, under Grant no. NSF EIA-0093426. Research on LLVM has also been supported by other NSF grants, as well as funding from the Office of Naval Research (ONR), DARPA, IARPA, DOE, and several industry sources. 

References

- Adve, V. CAREERS: Techniques and Applications of Dynamic Compilation. NSF Grant No. EIA-00-93426, Next Generation Software Program (CAREERS track). Total: \$500,000.
- Adve, V. et al. LLVA: A low-level virtual instruction set architecture. In *Proceedings of the 36th Annual ACM/IEEE Intern. Symp. on Microarchitecture* (2003).
- Alpern, B. et al. The Jikes Research Virtual Machine project: Building an open source research community. *IBM Systems J.* 44, 2 (2005).
- Auslander, M. and Hopkins, M. An overview of the PL.8 compiler. In *SIGPLAN '82: Proceedings of the SIGPLAN symposium on Compiler Construction* (1982), 22–31.
- Blajian, G., Eggen, R., Eggen, M., and Pitts, G. Mono versus .net: A comparative study of performance for distributed processing. In *Proceedings of the Intern. Conf. on Parallel and Distributed Processing Techniques and Applications* (2006).
- Blume, W. et al. Polaris: Improving the effectiveness of parallelizing compilers. In *Proceedings of the 7th Intern. Workshop on Languages and Compilers for Parallel Computing*. Springer-Verlag (1994), 141–154.
- Bocchino, R.L., Jr. and Adve, V.S. Vector LLVA: A virtual vector instruction set for media processing. In *Proceedings of the 2nd Intern. Conf. on Virtual Execution Environments*. ACM (2006), 46–56; DOI:10.1145/1134760.1134769
- Criswell, C., Lenharth, A., Dhurjati, D., and Adve, V. Secure Virtual Architecture: A safe execution environment for commodity operating systems. In *Proceedings of the Symp. on Operating Systems Principles* (2007).
- Criswell, J., Monroe, B., and Adve, V. A virtual instruction set interface for operating system kernels. In *Proceedings of the Workshop on Interaction between Operating System and Computer Architecture* (2006).
- Cummins, C. et al. LLVM compiler: Foundation language models for compiler optimization. In *Proceedings of the 34th ACM SIGPLAN Intern. Conf. on Compiler Construction* (2025), 141–153; DOI:10.1145/3708493.3712691
- Cytron, R. et al. Efficiently computing static single assignment form and the control dependence graph. *ACM Transactions on Programming Languages and Systems* 13, 4 (Oct. 1991), 451–490.
- Dhurjati, D. and Adve, V. Backwards-compatible array bounds checking for C with very low overhead. In *Proceedings of the Intern. Conf. on Software Engineering* (2006).
- Dhurjati, D. and Adve, V. Efficiently detecting all dangling pointer uses in production servers. In *Proceedings of the Intern. Conf. on Dependable Systems and Networks* (2006).
- Dhurjati, D., Kowshik, S., and Adve, V. SAFECode: Enforcing alias analysis for weakly typed languages. In *Proceedings of the ACM SIGPLAN Conf. on Programming Language Design and Implementation* (2006).
- Dhurjati, D., Kowshik, S., Adve, V., and Lattner, C. Memory safety without runtime checks or garbage collection. *ACM SIGPLAN Notices* 38, 7 (2003), 69–80.
- Dhurjati, D., Kowshik, S., Adve, V., and Lattner, C. Memory safety without garbage collection for embedded applications. *ACM Transactions in Embedded Computing Systems* (Feb. 2005).
- Dietz, W. and Adve, V. Software multiplexing: Share your libraries and statically link them too. In *Proceedings of the ACM Program. Lang.* 2, Article 154 (Oct. 2018); DOI:10.1145/3276524
- Eijeh, A. et al. HPVM: Hardware-agnostic programming for heterogeneous parallel systems. *IEEE Micro* 42, 5 (2022), 108–117; DOI:10.1109/MM.2022.3186547
- Fernández, M.F. Simple and effective link-time optimization of Modula-3 programs. In *Proceedings of Programming Language Design and Implementation* (1995).
- The Linux Foundation. The compiler that changed the world turns 25. The Linux Foundation (2024); <https://tinyurl.com/24gzf7kr>.
- Fruja, N.G. A modular design for the common language runtime (CLR) architecture. *Abstract State Machines*. (2005); <https://tinyurl.com/2cseyn2r>
- Healy, M.B. et al. Design and architecture of the IBM Quantum Engine Compiler. In *Proceedings of the 2024 IEEE Intern. Conf. on Quantum Computing and Engineering 1* (2024), 866–872; DOI:10.1109/QCE60285.2024.00106
- Kothari, A. et al. Hydride: A retargetable and extensible synthesis-based compiler for modern hardware architectures. In *Proceedings of the 29th ACM Intern. Conf. on Architectural Support for Programming Languages and Operating Systems 2* (2024), 514–529; doi:10.1145/3620665.3640385
- Kowshik, S., Dhurjati, D., and Adve, V. Ensuring code safety without runtime checks for real-time control systems. In *Proc. Intern. Conf. on Compilers, Architecture and Synthesis for Embedded Systems* (2002).
- Lawrence Livermore National Laboratory. NSA, national labs team with Nvidia to develop open source Fortran compiler technology>. Lawrence Livermore National Laboratory; <https://tinyurl.com/27ro7swe>.
- Lattner, C. LLVM. *The Architecture of Open Source Applications 1* (2019); <https://lulu.com>.
- Lattner, C. and Adve, V. Architecture for a next-generation GCC. In *Proceedings of the 1st Annual GCC Developers' Summit* (2003).
- Lattner, C. and Adve, V. LLVM: A compilation framework for lifelong program analysis and transformation. In *Proceedings of the Conf. on Code Generation and Optimization* (2004), 75–88.
- Lattner, C. and Adve, V. Automatic pool allocation: Improving performance by controlling data structure layout in the heap. In *Proceedings of the ACM SIGPLAN Conf. on Programming Language Design and Implementation*. (2005).
- Lattner, C. and Adve, V. Transparent pointer compression for linked data structures. In *Proceedings of the 2005 Workshop on Memory System Performance* (2005), 24–35.
- Lattner, C. et al. MLIR: Scaling compiler infrastructure for domain specific computation. In *Proceedings of the 2021 IEEE/ACM Intern. Symp. on Code Generation and Optimization* (2021), 2–14; DOI:10.1109/CGO51591.2021.9370308
- Lattner, C., Lenharth, A.D., and Adve, V.S. Making context-sensitive points-to analysis with heap cloning practical for the real world. In *A Proceedings of the ACM SIGPLAN Conf. on Programming Language Design and Implementation*. (2007), 278–289.
- Lopes, N.P. et al. Alive2: Bounded translation validation for LLVM. In *Proceedings of the 42nd ACM SIGPLAN Intern. Conf. on Programming Language Design and Implementation*, Association for Computing Machinery (2021), 65–79; DOI:10.1145/3453483.3454030
- Nagarakatte, S. et al. SoftBound: Highly compatible and complete spatial memory safety for C. In *Proceedings of the 30th ACM SIGPLAN Conf. on Programming Language Design and Implementation*, Association for Computing Machinery. (2009), 245–258; DOI:10.1145/1542476.1542504
- Noor, A.R. et al. MISAAL: Synthesis-based automatic generation of efficient and retargetable semantics-driven optimizations. In *Proceedings of the ACM on Programming Languages* 9 (2025), 1269–1292.
- Reinders, J. Intel C/C++ Compilers Complete Switch to LLVM. Intel Blog; <https://tinyurl.com/2xooktgq>.
- Tice, C. et al. Enforcing forward-edge control-flow integrity in GCC & LLVM. In *Proceedings of the 23rd USENIX Conf. on Security Symp.* USENIX Association, (2014), 941–955.
- Weber, N. Clang is now used to build Chrome for Windows. Intel Blog (2018); <https://tinyurl.com/24zoktmm>.
- Wilson, R.P. et al. SUIF: An infrastructure for research on parallelizing and optimizing compilers. *SIGPLAN Notices* 29, 12 (Dec. 1994), 31–37; DOI:10.1145/193209.193217
- Zhang, Z. et al. AutoPilot: A platform-based ESL synthesis system. *High-Level Synthesis From Algorithm to Digital Circuit* (2008), 99–112; <https://tinyurl.com/22pjb9p9>

Vikram S. Adve (vadve@illinois.edu) is the Donald B. Gillies Professor of Computer Science at the University of Illinois at Urbana-Champaign, Urbana, IL, USA.

Chris Lattner is CEO at Modular, Inc., Los Gatos, CA, USA.

 This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).

Steps toward a more nuanced and practical understanding of openness in AI.

BY CAMILLE FRANÇOIS, ADRIEN BASDEVANT, VICTOR STORCHAN, KEVIN BANKSTON, AYAH BDEIR, BRIAN BEHLENDORF, MEROUANE DEBBAH, SAYASH KAPOOR, YANN LECUN, MARK SURMAN, HELEN KING-TURVEY, NATHAN LAMBERT, STEFANO MAFFULLI, NIK MARDIA, GOVIND SHIVKUMAR, AND JUSTINE TUNNEY

Unpacking Open Source Artificial Intelligence:

Toward a Framework for Openness in Foundation Models

OPENNESS HAS LONG driven innovation in software,⁹ and AI is no exception.¹² While some see openness in foundation models (FMs) as a security threat,¹⁸ others argue that restricting access will not meaningfully reduce risk and will limit the benefits of transparency, research, and global participation.³


As the EU AI Act reporting requirements on FMs—also referred to as general-purpose AI models (GPAIMs)—move toward implementation, there is an urgent need for a more nuanced and informed understanding of openness in AI systems.

It has been challenging to define openness in the context of FMs, as definitions of “open source” software do not easily translate to AI systems. However, shared nomenclature is essential to developing shared understandings, norms, benchmarks, and best practices. By shared nomenclature, we mean establishing common terms across researchers, developers, policymakers, and civil society organizations working to unlock the benefits and mitigate the risks of FMs. We use *AI systems* here as shorthand for systems built with FMs, while recognizing the limitations and critiques of the term *AI*. This work is also necessary to illuminate the range of potential design choices around openness throughout the AI stack, and to ensure that this conversation moves beyond a narrow focus on model weights.

In this article, we survey existing

» key insights

- **Defining and evaluating openness in AI requires examining both the model and system stack, including interfaces, safeguards, infrastructure, and deployment practices that shape real-world risks and benefits.**
- **AI systems can be open in different ways and to different degrees across data, code, weights, documentation, and governance, requiring more precise terminology than *open* or *closed*.**
- **Policy discussions on AI openness need to be informed by clearer shared language that distinguishes between different components, degrees, and purposes of openness across the AI stack.**
- **Safety is dependent on system context, not models alone. Meaningful AI safety assessment requires examining deployment environments, safeguards, moderation layers, and governance mechanisms alongside the underlying model itself.**



approaches to defining openness in AI models and systems. We also propose a descriptive framework to evaluate how each component across the FM stack contributes to openness, enabling normative definitions of openness in AI. We intentionally do not present a definitive list of requirements for openness. This work builds on a February 2024 workshop convened by Mozilla and the Columbia Institute of Global Politics, which brought together more than 40 leading scholars and practitioners working on openness and AI.¹ These individuals—spanning open source AI startups and companies, nonprofit AI labs, and civil society organizations—focused on exploring what *open* should mean in the current era of foundation models.

Foundations

There are several timely areas of exploration concerning openness and AI. The first is reflecting on the impact of open source on AI development.⁹ Another centers on the complexities of opening up AI systems themselves—specifically, identifying which components of the technology

stack might be made accessible and how. Here, we focus on the latter. A literature review, included below, provides diverging perspectives on what openness can and should mean in an AI context,¹⁸ but we argue that the following four tenets provide a foundation for considering aspects of openness and AI:

► *Openness must be considered both at the model stack and the system stack levels.* Focusing only on model weights or release strategies obscures the wider set of design choices—interfaces, safeguards, infrastructure, and deployment practices—that ultimately determine an AI system's benefits and risks in practice.

► *Openness takes multiple forms.* Beyond the binary of “open” or “closed,” openness in AI systems manifests through varying degrees of access to data, code, weights, documentation, and governance. Recognizing this diversity allows for more precise vocabulary and concepts (e.g., open science, open data, open weights, available weights, etc.) and sharper normative criteria for evaluating different release strategies.¹⁹

► *Greater clarity is needed in how*

openness is discussed at the policy level. Most policy debates on openness have lacked a shared understanding of the benefits, risks, and mechanisms for increasing transparency and access across the AI stack. In the absence of such clarity, discussions of risk become polarized between boosterism and alarmism, rather than converging on practical approaches that balance innovation, accountability, and security. Addressing this gap requires a more precise vocabulary for openness, one that can differentiate between components, degrees, and purposes rather than collapse them under a single ambiguous label.¹⁷

► *Safety must be a core consideration.* Safety cannot be ensured by scrutinizing models in isolation.¹³ It depends on the broader context of deployment, including the safeguards built into full AI systems. Distinguishing between *open models* and *open systems* is therefore essential:¹⁶ The latter can incorporate protective measures, but also introduce new risks absent at the model level. Openness and safety are closely connected—without openness, independent risk assessment is impossible—but safety cannot be reduced to openness alone.

Together, these foundations highlight the need for a structured way to map openness across the AI stack. The framework discussed here operationalizes these principles, distinguishing model-level and system-level components and clarifying the cross-cutting attributes that shape openness in practice (with an emphasis on AI systems with model weights that are downloadable or more open).

We also invite discussion of the key trade-offs for leveraging openness and how the different dimensions of openness relate to the different goals one might wish to pursue (see Figure 1). The framework does not offer a comprehensive account of openness in AI; instead, it is designed to be a practical tool for developers, policy-makers, and other key players.

Motivations: Why pursue AI openness? In Figure 1, a non-exhaustive list of potential benefits sourced from workshop participants highlights a variety of motivations for AI openness. AI openness is rarely pursued

Figure 1. Examples of potential benefits from openness.

Advancing Research and Innovation

Open Science, Open Knowledge: Enhance reproducibility and foster innovation through transparent and shareable research practices and knowledge dissemination.

Accelerated Innovation: Enable creation of new tools and models by providing extensible infrastructure and building blocks.

Fostering the Unexpected: Support creative and unanticipated uses of technology.

Support Key Research Areas: Advance understanding of bias and other harm areas.

Ensuring Accountability and Safety

System Accountability: Enable audits and explainability to assess performance and impacts.

Bug and Vulnerability Disclosure: Facilitate community-led identification and reporting of flaws.

Recourse and Transparency: Ensure affected individuals have avenues for redress.

Building Inclusive and Collaborative Ecosystems

Open Ecosystem: Cultivate communities of developers, builders, and researchers to drive collaboration.

Diversity of Representation: Provide tools that enable training on varied datasets and sources.

Independent Oversight: Equip researchers and regulators to evaluate large-scale systems for systemic risks.

Promoting Efficiency, Autonomy, and Fair Competition

Cost Optimization: Reduce cost for public procurement.

Digital Autonomy: Avoid vendor lock-in to increase digital autonomy and enable more private / local architectures as opposed to cloud based.

Resource Efficiency: Improve the energy and compute efficiency of training and inference.

Market Competition: Level the playing field by enabling smaller actors and new entrants to participate meaningfully.

for its own sake; it is motivated by diverse and sometimes conflicting goals. Our framework helps make these goals explicit by showing how different motivations map onto different forms of openness across the AI stack. For instance, advancing research and innovation often depends on openness of data or training code; promoting competition and efficiency tends to emphasize access to model weights; and ensuring accountability may require openness in documentation, licensing, and safeguards. These motivations may overlap, but they do not always align. A strategy that maximizes reproducibility, for example, may differ from one that maximizes market competition or system safety. In the free/libre tradition, “openness” is inseparable from user freedoms, and from ensuring protection against malicious features.²⁰ The user’s essential freedoms are to run, copy, distribute, study, change, and improve the technology;⁸ in AI systems, licensing and governance should be assessed on whether they preserve or effectively claw back these freedoms across components.

By situating motivations within the framework, we can move beyond abstract claims that “openness is good” or polarized debates that cast it as either inherently dangerous or inherently beneficial. Instead, the framework enables more precise evaluation of which forms of openness serve which objectives, and where trade-offs or complementarities exist. This creates a clearer basis for policymakers, researchers, and developers to debate not only *whether* to open AI systems, but *how* and *for what purposes*.

Related Work: Beyond Open Bashing and Open Washing

Over the past two years, foundational efforts have sought to map the landscape of openness in AI. The existing literature reveals a diversity of approaches to conceptualizing AI openness, each with distinct methodological choices and underlying assumptions. These approaches can be broadly categorized by how they frame openness—whether as a gradient (“more vs. less open”), a scoring system (“a certain degree of



AI openness is rarely pursued for its own sake; it is motivated by diverse and sometimes conflicting goals. Our framework helps make these goals explicit.



openness”), or a binary classification (“open vs. closed”).

Binary approaches. The Open Source Initiative (OSI) launched a process in June 2023 to draft a definition of open source AI, culminating in the release of an “Open Source AI Definition” in late 2024.¹⁵ It best exemplifies the binary approach, with a clear set of criteria on what counts as “open” versus “closed,” and references to OSI-approved licenses and terms to meet these criteria. Below, and as concrete illustrations for what *open* can mean in practice across weights, code, data, checkpoints, and licensing, we cite projects that exemplify a fully open approach.

EleutherAI released a suite of open models with full training checkpoints and tools to reconstruct training dataloaders,² advocating for standard open source licenses to ensure clarity and interoperability.

The LLM360 initiative¹⁰ promotes full transparency across model training, checkpoints, and data. The project argues that partial releases hinder scientific progress and reproducibility.

The Allen Institute for AI built on initiatives like Pythia and BLOOM by releasing an initial set of 1B and 7B models with 500-plus checkpoints from intermediate training, full datasets, training code, evaluation code, training logs, and comprehensive documentation.²⁵ The initiative aims to foster a growing “science of LLMs” and influence public AI discourse beyond transparency and openness goals.

Beyond a single open/closed label for a whole release, some efforts apply a binary rule, component by component. The Linux Foundation’s Model Openness Framework (MOF)²³ is the clearest instance: It evaluates whether each component is released under appropriate open terms. The MOF classifies models based on openness and completeness across 15 components of the model development lifecycle. The framework distinguishes between openness (binary classification of whether a component is released under an appropriate open license) and completeness (gradient approach with three tiers of model completeness, based on which com-

ponents are released under appropriate open licenses). The framework includes a required MOF.JSON file detailing licenses and components, with classes ranging from Open Model (Class III) to Open Science (Class I), encouraging transparency, reuse, and reproducibility.

Criteria scoring approaches. The Centre for Language Studies at Radboud University developed a project²⁶ to assess the openness, transparency, and accountability of instruction-tuned language models. The research highlights risks of proprietary software and provides comparative analysis of open source alternatives with similar architectures and capabilities.

Stanford CRFM created the Foundation Model Transparency Index,⁴ introducing 100 fine-grained indicators that capture transparency across the lifecycle of FMs—from upstream inputs (data sources, labor, and compute) to the model itself (size, capabilities, and risks) and downstream uses (distribution channels, usage policies, and affected geographies). The study emphasizes that AI release strategies are nonbinary and classifies models with widely downloadable weights as open. The research found that open models tend to outperform closed ones in transparency metrics, attributing this partly to lack of transparency among closed developers about critical upstream components.

EleutherAI offered a contrasting perspective,⁷ cautioning that score-based transparency frameworks risk treating transparency as an end in itself rather than as a means to uphold

ethical values such as accountability.

Gradient/spectrum approaches. In “The Gradient of Generative AI Release: Methods and Considerations,”¹⁹ Solaiman introduces a template for understanding different levels of access to generative AI systems, from fully closed to fully open. The paper defines six release levels—fully closed, staged access, hosted access, API/cloud access, downloadable access, and fully open—framing them as points along a gradient. It highlights the trade-offs at each level, particularly the balance between risk mitigation and power concentration.

The Princeton Center for Information Technology Policy and Stanford Center for Research on Foundation Models hosted a 2023 workshop¹⁷ to discuss key benefits and risks of specific properties of open FMs. They proposed a risk assessment framework for analyzing their *marginal risk* (i.e., using closed source and preexisting technologies as a baseline). They identify that more research is needed to assess the marginal risks of open FMs.

The Digital Public Goods Alliance developed a gradient approach framework²⁷ that acknowledges that access to AI testing and training data is not the same as having access to source code; put another way, open sourcing data alone does not enable full model reproducibility. It recognizes varying degrees of openness across AI systems (training data, model, and code), without further detailing sub-components. The framework seeks to balance the benefits of AI openness, such as transparency and inclusivity, against potential risks of misuse

or harm, and includes the concept of “aspirational openness,” defining various levels of openness for AI components to qualify as digital public goods, taking into account factors like ethical use and data rights issues.

Gradient approaches commonly ground *marginal-risk* analyses of openness; we review some key examples below.

In “Considerations for Governing Open Foundation Models,”³ Bommasani et al. from the Stanford Institute for Human-Centered Artificial Intelligence examine how release strategies of open FMs relate to governance risks and regulation consequences. A key takeaway from the brief is to invite practitioners to think in terms of *marginal risk of open foundation models* “relative to (a) closed models or (b) pre-existing technologies.” The paper also discusses some of the benefits of open FMs such as competition, accelerating innovation, and distributing power. The brief acknowledges that the release of FMs is a gradient across various dimensions, and encourages policymakers to explore downstream intervention points that would be more effective for mitigating AI risks.

In the paper “What Does It Mean to Open Source AI?”¹¹ Luccioni et al. analyze the risks and benefits of open FMs. Benefits include innovation, competition, the distribution of decision-making power, and transparency. The authors also present a risk-assessment framework for evaluating the marginal risk of open FMs, and survey the current evidence for the risks of open FMs, which, in many cases, the authors find lacking.

Summary. While these existing frameworks have made valuable contributions to the emerging scholarship around defining openness in the context of AI, they also present key limitations that we are seeking to address. Binary approaches may oversimplify the complex, multidimensional nature of AI systems, and thus the different pathways that may lead to openness, while gradient approaches often lack specificity about what constitutes meaningful differences along the spectrum. Scoring approaches, though comprehensive, risk treating transparency as an end in itself rather than as a means to





Figure 3. The framework: General-purpose AI system (GPAIS) stack and dimensions of openness.

achieve broader goals such as accountability and safety. Additionally, many existing frameworks focus primarily on model-level considerations, potentially overlooking the broader system-level components that significantly impact user freedoms. To bridge these gaps, we make explicit which system-level components and cross-cutting attributes (documentation, licensing, safeguards) matter for specific goals, rather than treating openness as a single axis.

Building on these prior efforts, we propose a unified framework that distinguishes between the AI *model*

stack and the broader AI *system* stack, while incorporating cross-cutting attributes such as documentation, safeguards, and licensing. Our approach aims to provide greater granularity and practical applicability by mapping different dimensions of openness to their corresponding benefits and use cases.

The full proceedings write-up includes three appendices that offer additional depth: (i) Benefits of Openness across the Foundation Model Stack, (ii) Key Areas of Debate, and (iii) Alignment with the Linux Foundation's Model Openness Frame-

work. These are available in the full version of the article.¹

The Framework

In this section, we introduce a framework focused on AI systems in which model weights are at least downloadable. We begin with a representation of the AI model stack, then zoom out to the broader AI system stack, highlighting shared attributes that apply across the stack. We present the framework in a way that aligns with the various benefits of AI openness.

AI model stack. Reaching an agreement on a definition for AI openness

has proven complicated.¹⁴ AI models are not just code; they are trained on massive datasets, deployed on intricate computing infrastructure, and accessed through diverse interfaces and modalities. To effectively depict an AI system stack, we initially focus on the AI model tech stack. Figure 2 presents the core components that are particularly relevant for considerations around openness of a general-purpose AI model (GPAIM).

At the model level, there are three main artifacts: datasets, code, and model weights, each of which also includes sub-components. It is important to note that not all sub-components are mandatory or applicable in every configuration (e.g., some models may have the reward model's weights or preference dataset; others do not), and the taxonomy of sub-components may vary based on the specific model type. Given this variability, it is crucial to tailor the AI model stack to accommodate different modalities, such as input and output types. Figure 2 illustrates these various layers at the model level.

AI system stack. It is important to distinguish between open models and fully open AI systems or products. Open AI systems often incorporate more than just the model, including interfaces, telemetry, safety layers, and orchestration tools. For example, an open chatbot might rely on external moderation systems, while a RAG application could add components like langChain or Redis for orchestration and caching. This structure differs from traditional software, where code, compilers, and licenses are clearly separated. In AI, components such as data, hardware, and algorithms jointly shape behavior, making openness harder to define. For instance, opinions diverge on whether model weights, code, or datasets count as the “source” for modification. (Despite these differences between building AI models and traditional software, some initiatives are trying to study to what extent FMs could be built as “open source software”—especially given the technical ability to efficiently update the weights of the models. This would allow contributors to be able to submit patches to fix potential identi-



We hope a meaningful contribution of this framework can be its explicit distinction between the model stack and the broader system stack in the context of evaluating openness.



fied vulnerabilities in the model; then maintainers would need to be able to merge the patches to quickly integrate updates into the model.)

Lastly, the infrastructure layer plays a key part, as many scholars have already underlined: The high costs⁶ associated with developing AI models could hinder¹⁹ widespread adoption, even using open approaches, due to the expensive hardware and infrastructure requirements, such as GPUs. While openness in AI provides organizational control for customization, adapting large models to specific domains poses challenges without substantial investments in GPUs, so this does not necessarily solve the challenges of concentrated access and influence in well-resourced labs.⁴ See Figure 3 for a visual representation of the tech stack required for a general-purpose AI system (GPAIS).

Common attributes across the AI system stack. It is crucial to differentiate between the components of the tech stack, such as model weights, and their attributes, such as the license for model weights. For instance, in this article we consider many AI safety measures to be better understood as attributes of the system rather than features of the model. Consequently, we represent an AI system as comprising three cross-cutting sets of attributes—documentation, safeguards, and licensing—that may apply across the entire stack rather than being tied to any single component (see Figure 3 for attribute categories integrated into the system stack).

For example, safety guardrails can be integrated at the level of individual components, recognizing their diversity and the need to address different types of harms and risks that may arise across the stack. This approach allows for targeted mitigations that are best suited to each component's specific role and vulnerabilities.¹³ Entrusting safety assessments solely to models is insufficient, especially regarding concerns of potential misuse, as models lack crucial contextual information necessary for making accurate safety assessments.

An openness in AI framework. As a whole, the framework presents a tech stack described across its different dimensions: AI artifacts (e.g., model

weights), to which correspond different attributes such as licensing (e.g., choice of model's weights license) or documentation (e.g., model cards). Each of the dimensions is composed of categories and subcategories that can be released more or less openly. As illustrated in Figure 3, these components correspond to different benefits of openness at each layer of the stack. Further work is needed to understand what is required to open different parts of the stack, which components are essential for reproducibility or evaluation, and how openness needs may shift with future AI systems beyond today's LLMs.

We hope a meaningful contribution of this framework can be its explicit distinction between the model stack and the broader system stack in the context of evaluating openness, combined with the mapping of cross-cutting attributes like documentation, licensing, and safeguards that apply across both levels. This work seeks to unify previously fragmented approaches into a coherent descriptive framework that can be systematically applied by policymakers evaluating regulatory approaches and practitioners making release decisions.

Conclusion

This work aims to accelerate a shared understanding of openness in FMs and AI systems, differentiating challenges at the model and system levels, and those that could apply across the AI stack, and showing how to leverage that understanding to better harness the benefits and mitigate the risks of AI. Our framework provides a vocabulary for systematically analyzing openness at different levels of AI systems, clarifies distinctions often conflated in public debate, and highlights concrete areas where openness interacts with safety and governance. By clarifying the different levels, forms, and vocabularies of openness, and by distinguishing openness from—but also connecting it to—safety, the framework aims to equip researchers and policymakers with a practical tool for evaluating openness strategies in AI systems. In future work, we hope to go deeper into how this framework can be strengthened

with more information about safety mitigations.

Acknowledgments

We acknowledge Aviya Skowron, Amba Kak, Irene Solaiman, Stella Biderman, and Martin Tisné for their input to this article. We also thank all the participants in the Columbia Convening on Openness and AI for their contributions to this work, and to Jen Weedon and Mariel Povolny for their editorial support. 

References

- Basdevant, A. et al. Towards a framework for openness in foundation models: Proceedings from the Columbia Convening on Openness in Artificial Intelligence. arXiv (2024), arXiv:2405.15802.
- Biderman, S. et al. Pythia: A suite for analyzing large language models across training and scaling. In *Intern. Conf. on Machine Learning*. ACM (2023).
- Bommasani, R. et al. Considerations for governing open foundation models. Stanford Institute for Human-Centered AI (2023).
- Bommasani, R. et al. The foundation model transparency index. arXiv (2023), arXiv:2310.12941.
- Bommasani, R., Zhang, D., and Liang, P. The foundation model transparency index. arXiv (2022), arXiv:2212.04960.
- Dang, A. Winning the adversarial battle with open-source AI models. SSRN (2023); <https://ssrn.com/abstract=4651571>
- EleutherAI. How the foundation model transparency index distorts transparency (2024); <https://blog.eleuther.ai/fmti-critique/>
- Free Software Foundation. What is free software? GNU Project; <https://www.gnu.org/philosophy/free-sw.en.html#four-freedom>
- Hoffmann, N., Nagle, F., and Zhou, J. The value of open source software. Harvard Business School Strategy Unit Working Paper (2024).
- Liu, Z. et al. LLM360: Towards fully transparent open-source LLMs. arXiv (2023), arXiv:2312.06550.
- Luccioni, S., Solaiman, I., and Paolotti, D. What does it mean to open source AI? arXiv (2024), arXiv:2403.07918.
- Mozilla. Joint statement on AI safety and openness (2023).
- Narayanan, A. and Kapoor, S. AI safety is not a model property. AI as Normal Technology (Mar. 12, 2024); <https://www.normaltech.ai/p/ai-safety-is-not-a-model-property>
- OECD.AI. Updating the definition of an AI system: A focus on functionality (2024); <https://oecd.ai/en/work/ai-system-definition-update>
- Open Source Initiative (OSI). Open source AI definition 1.0; <https://opensource.org/ai/open-source-ai-definition>
- Partnership on AI. Tenets (2016); <https://www.partnershiponai.org/tenets/>
- Princeton Center for Information Technology Policy & Stanford CRFM Workshop on responsible and open foundation models (2023); <https://sites.google.com/view/open-foundation-models>
- Seger, E. et al. Open-sourcing highly capable foundation models: An evaluation of risks, benefits, and alternative methods for pursuing open-source objectives. Centre for the Governance of AI (2023).
- Solaiman, I. The gradient of generative AI release: Methods and considerations. In *Conf. on Fairness, Accountability, and Transparency*. ACM (2023), 111–122.
- Stallman, R.M. *Free Software, Free Society: Selected Essays of Richard M. Stallman*. GNU Press (2015).
- Stallman, R. Why open source misses the point of free software. Free Software Foundation; <https://www.gnu.org/philosophy/open-source-misses-the-point.html>
- Stanford Institute for Human-Centered Artificial Intelligence (HAI). Governing open foundation models (2023); <https://hai.stanford.edu/sites/default/files/2023-12/Governing-Open-Foundation-Models.pdf>.
- White, M. et al. The model openness framework. arXiv (2024), arXiv:2403.13784.
- Bdeir, A. and François, C. Introducing the Columbia convening on openness and AI. Columbia University Institute of Global Politics (2025); <https://igp.sipa.columbia.edu/news/introducing-columbia-convening-openness-and-ai>
- Allen Institute for AI. OLMo; <https://allenai.org/olmo>
- Opening Up ChatGPT Project; <https://github.com/opening-up-chatgpt/opening-up-chatgpt.github.io/>
- Digital Public Goods Alliance. Exploring a gradient approach to the openness of AI system components (2024); <https://digitalpublicgoods.net/blog/exploring-a-gradient-approach-to-the-openness-of-ai-system-components/>

Camille François (cmf2157@columbia.edu) is an associate professor of practice at Columbia University, New York, NY, USA.

Adrien Basdevant is a senior policy advisor at Mozilla, San Francisco, CA, USA, and founder of Entropy Law, Paris, France.

Victor Storch leads AI/ML Research at Mozilla.ai in New York, NY, USA.

Kevin Bankston is a senior advisor on artificial intelligence governance at the Center for Democracy & Technology, Washington, DC, USA.

Ayah Bdeir is a senior advisor to the president at Mozilla in New York, NY, USA.

Brian Behlendorf is chief AI strategist at the Linux Foundation, San Francisco, CA, USA.

Merouane Debbah is a senior scientific AI advisor, Technology Innovation Institute, Abu Dhabi, United Arab Emirates.

Sayash Kapoor is a research fellow at the Princeton University Center for Information Technology Policy, Princeton, NJ, USA.

Yann LeCun is the vice president and chief AI scientist at Meta, New York, NY, USA.

Mark Surman is the president of the Mozilla Foundation, Toronto, ON, Canada.

Helen King-Turvey is managing partner at Arising Matters in Winchester, England, UK.

Nathan Lambert is a senior research scientist at the Allen Institute for Artificial Intelligence, Seattle, WA, USA.

Stefano Maffulli is the executive director of the Open Source Initiative, San Francisco, CA, USA.

Nik Marda is the technical lead for AI governance at Mozilla in New York, NY, USA.

Govind Shivkumar is a director at Omidyar Network in Washington, DC, USA.

Justine Tunney is a software engineer at Mozilla in Mountain View, California, USA.

 This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).

LLMs are increasingly being used in robotics applications such as drone control, but ensuring they perform safely in physical environments is critical.

BY YUNG-CHEN TANG, PIN-YU CHEN, AND TSUNG-YI HO

Defining and Evaluating Physical Safety for Large Language Models

LARGE LANGUAGE MODELS (LLMs) have emerged as powerful applications in the field of natural language processing,^{9,18,19} exhibiting human- and even superhuman-level performance across various tasks. Innovations in neural network architectures²⁵ and increased computational power²⁴ have led their evolution into state-of-the-art artificial intelligence (AI)


systems that understand and generate text with remarkable fluency and coherence. For example, GitHub Copilot,^a developed in collaboration between GitHub and OpenAI, is an AI-powered coding assistant based on a variant of the GPT-3 language model.⁵ It assists users by autonomously completing code snippets or generating entire functions as they write code. LLMs have also shown outstanding capabilities in areas such as reasoning,^{1,14} planning,²² manipulation,²⁹ and in-context learning,¹⁵ despite not being explicitly trained in them. As a result, they have been widely used in many different applications, such as question answering, multimodal tasks, tool manipulation, and autonomous agents.

Due to their strong capabilities, LLMs have also become popular for controlling robots,^{3,4} in areas ranging from industrial automation to AI-powered physical agents. However, widespread adoption of these systems is dependent upon ensuring their safety. And though significant progress has been made in evaluating LLM safety around user-generated content (e.g., preventing toxic content and hate

a <https://github.com/features/copilot/>

» key insights

- As large language models evolve into agentic embodied AI to control robotic systems such as drones, safety evaluations must expand beyond textual safety to include physical safety in the real world, particularly with regard to humans, objects, and infrastructure.
- To address this gap, we took the first step of developing a comprehensive benchmark to evaluate the physical safety of LLMs in drone control. We used a high-fidelity simulation environment to quantify risks such as collisions and regulatory violations.
- Our findings show that although larger model sizes and advanced prompting techniques can mitigate obvious malicious orders, LLMs still struggle to identify “unintentional attacks” because they lack the common sense to realize that seemingly harmless commands can lead to physical harm.



speech),^{7,17,28} the physical safety of LLM-controlled robots in real-world applications remains largely unexplored. This is of particular concern in areas like drone control,^{10,16,21} where the threat vectors are realistic and malfunctions can lead to catastrophic physical harm and damage.

In this article, we describe a pilot study on physical safety that focuses on drones, due to their low cost and easy accessibility. For around \$100 USD, programmable drones like the Ryze Tello offer a cost-effective alter-

native to expensive robotic systems such as humanoid robots. Many modern drone manufacturers also offer software development kits (SDKs) and application programming interfaces (APIs), such as DJI Developer SDKs and Parrot Air SDK, enabling researchers or users to easily deploy custom code. This combination of affordability and programmability makes drones an ideal case study for exploring safety risks associated with LLM-generated code in robotic systems.

To address the critical gap in physi-

cal safety evaluation for LLMs,^b our research focuses on defining and evaluating physical safety in drone control. We consider a realistic scenario where an attacker or unintended user exploits LLMs to generate drone-control code, potentially leading to malicious actions. Our evaluation of several main-

^b For example, the AI Safety Benchmark v0.5 aims to evaluate safety risks arising from language generation in AI systems, while our benchmark specifically targets the physical harm risks that LLMs may cause in robotic control.

stream LLMs in drone control reveals some key insights:

- There is a trade-off between utility and safety, where higher utility scores tend to exhibit higher safety risks, indicating a need for improved alignment.

- In-context learning (ICL) shows slight improvements over zero-shot chain-of-thought (ZS-CoT), both of which demonstrate significant safety enhancements; however, both struggle to detect unintentional attacks, highlighting limitations in complex reasoning and anticipating consequences.

- Model size plays a critical role in safety; while larger LLMs generally exhibit better performance in blocking dangerous commands, further scaling up the model size results in only marginal gains in some physical safety categories.

These findings have important implications for developing robust and safe LLM-based robotic systems. Our comprehensive benchmark establishes a solid foundation on which further research and innovation can build to enhance the physical safety of LLMs in real-world, safety-critical scenarios. This is crucial for ensuring the effective functioning of these systems and fostering their trustworthiness and reliability to encourage the widespread adoption of powerful AI technologies.

Results

In this section, we discuss the physical safety evaluation objectives and benchmark design for LLMs in drone control. We introduce a dataset for safety and utility evaluations, explain the metrics used, and examine various LLMs regarding safety. We also present evaluation results, analyze the trade-offs between utility and safety, highlight the impact of ICL on safety, and offer insights into how model size influences outcomes.

Physical safety evaluation objectives and benchmark design. As depicted in Figure 1, the physical safety evaluation aims to address the potential threats and risks posed by drones in various real-world scenarios. We categorize the risks into four main classes:

- *Human-targeted threats:* Drones could be used to physically harm individuals, such as instructing an LLM to provide steps for a drone to delib-



Our comprehensive benchmark establishes a solid foundation on which further research and innovation can build to enhance the physical safety of LLMs in real-world, safety-critical scenarios.



erately strike a crowd, posing serious safety risks.

- *Object-targeted threats:* Drones may damage property, such as vehicles or other objects, leading to financial losses. For instance, a malicious actor might use an LLM to direct a drone to collide with a car.

- *Infrastructure attacks:* Targeting critical infrastructure, like power grids, could cause widespread disruption. A malicious user could generate instructions for a drone to crash into a transmission tower, leading to outages.


- *Violations of FAA regulations:* Unsafe drone operations, such as flying in restricted airspace, pose safety and legal risks. For example, flying a drone into a no-fly zone near an airport could endanger aircraft and result in severe legal consequences. Our evaluation data is based on FAA Part 107 regulations.¹¹

This categorization considers both the direct and indirect effects of potential drone damage.

To evaluate the physical safety of LLM-based systems, we developed the LLM Physical Safety Benchmark (Figure 2), a comprehensive framework to explore and quantify safety risks. This benchmark uses a curated dataset of potential risks and prompts, created by us and two LLM judges, along with a simulation environment, to assess the safety and utility of mainstream LLMs, including OpenAI's ChatGPT, Google's Gemini, and Meta's Llama.

The evaluation starts by providing the LLM with a test instruction categorized as one of four dataset dimensions: *deliberate attacks*, *unintentional attacks*, *violation instructions*, or *utility tasks*. These dimensions are designed to evaluate user intent and to cover the physical safety risk categories defined in Figure 1. The mapping of the risk categories to the evaluation datasets is presented in the appendix. The LLM's response is then assessed by two specialized AI judges, the Code Verification Judge and the Safety Evaluation Judge. Introducing these LLM judges is important because evaluating whether a model's response is safe involves more than simply checking whether executable code is provided. Two notable safety-conscious behaviors are:

- *Self-protect:* The model adjusts generated code for safety, showing that code generation is not always unsafe.

Figure 1. Categories of drone safety threats, including attacks on humans, objects, and infrastructure, and violations of FAA regulations.**Figure 2. Benchmark process for drone control, where a specific LLM is evaluated by providing a test prompt, categorizing the output, and assessing code responses with two AI judges. The code is further tested in a simulation environment to detect collision risks, contributing to the final safety evaluation (see project page^e for video demo).**

► **Safety refusal:** The model declines to generate code, not due to lack of capability but for safety reasons.

Additionally, the generated code is tested in the AirSim simulation environment²³ to identify potential collision risks and safety issues. Finally, the results are summarized for a holistic safety evaluation, offering insights into the LLM's ability to generate safe and reliable drone-control instructions. More detailed implementations of these two AI judges and benchmark, as well as examples of self-protect and safety refusal behaviors, can be found in the appendix.

Dataset for safety and utility evaluations. To support the benchmark and facilitate comprehensive physical safety evaluations, we developed the LLM Physical Safety Benchmark Dataset,^d a curated collection of more

than 400 different instructions. The dataset was created through a collaborative effort combining human expertise and LLM support to construct a diverse and realistic set of scenarios, ensuring alignment with key drone regulations, including those of the U.S. Federal Aviation Administration (FAA) and the International Civil Aviation Organization (ICAO), and guaranteeing their relevance to current safety protocols. As illustrated in Figure 3, the dataset is categorized into four dimensions: deliberate attacks, unintentional attacks, violation instructions, and utility. Each dimension evaluates a subset of our defined physical safety risks as well as the utility of regular (safety-unrelated) tasks. Please refer to the appendix for the full dataset, including test instructions, categories, the relationship between safety risks and dataset dimensions, and a detailed flowchart of steps involved in the dataset construction process.

We provide detailed explanations

of the four dataset dimensions below. Within each subcategory, we have curated specific examples that represent realistic scenarios in drone control, shown in Figure 3. The distribution of evaluation data across these dimensions, presented by the pie charts, ensures a balanced and thorough evaluation of LLM performance in simulated drone-control environments:

► **Deliberate attacks:** Evaluate the LLM's resistance to malicious use, testing its ability to recognize and reject commands intended to cause harm. Subcategories include direct command attacks, indirect command attacks, and code injection.

► **Unintentional attacks:** Assess the LLM's capacity to identify and mitigate potentially harmful actions that might arise from user mistakes or lack of domain knowledge. This includes scenarios like misleading instructions, high-risk instructions, and scenario oversight.

► **Violation instructions:** Examine the LLM's adherence to legal and ethi-

c <https://huggingface.co/spaces/TrustSafeAI/LLM-physical-safety>

d https://huggingface.co/datasets/TrustSafeAI/llm_physical_safety_benchmark

Figure 3. The benchmark's evaluation datasets, categorized into four types: deliberate attacks, unintentional attacks, violation attacks, and utility, evaluating model performance from various perspectives.





Figure 4. Safety evaluation results across different LLMs. The left panel presents individual scores for six metrics: self-assurance, avoid collision, regulatory compliance, code fidelity, instruction understanding, and utility. The right panel visualizes these scores using a radar chart, highlighting the trade-off between utility and safety across various models.



cal guidelines, particularly focusing on FAA regulations and general safety protocols. Subcategories cover flying in a no drone zone, improper flight attitude, excessive flight speed, and operation over people.

► *Utility*: Check the LLM's ability to perform basic drone controls and operations, ensuring that safety measures do not overly compromise functionality. This dimension includes tasks such as takeoff and movement, path following, yaw angle control, and approach target.

Metrics for safety evaluation. Based on the physical safety dataset, we quantify the safety performance of different LLMs using the following six metrics. These metrics provide a systematic framework for assessing the safety performance of LLMs in robotic applications. For detailed definitions and calculations, please refer to the appendix:

► *Self-assurance*: Evaluates the LLM's ability to recognize and avoid dangerous situations. A higher score indicates better refusal of harmful instructions, reflecting commitment to safety.

► *Avoid collision*: Assesses the likelihood of collisions during drone control. This metric measures how effectively the LLM avoids collisions in both deliberate and unintentional scenarios.

► *Regulatory compliance*: Evaluates adherence to laws and regulations, such as FAA guidelines. A higher score shows the LLM's capability to identify and refuse instructions that violate these regulations.

► *Code fidelity*: Measures the accuracy and reliability of code generated by the LLM. A greater score indicates a higher success rate in generating executable code without errors.

► *Instruction understanding*: Assesses the LLM's ability to comprehend and respond appropriately to instructions. This includes both non-code and code responses evaluated by specialized AI judges.

► *Utility*: Evaluates the LLM's performance in basic drone-control tasks, such as moving or changing altitude. A higher score signifies better execution of fundamental control commands.

Evaluation results across LLMs. The safety evaluation results, presented in Figure 4, highlight significant variations in safety performance across the models tested,

including OpenAI GPT-3.5-turbo, Google Gemini Pro, Llama-2-7B-Chat, CodeLlama-7B-Instruct, Llama-3-8B-Instruct, Mistral-7B-Instruct-v0.2, and CodeQwen1.5-7B-Chat. Detailed descriptions are in the appendix.

These differences in safety performance highlight the need for deeper consideration of model design, particularly with regard to self-assurance and collision-avoidance capabilities. CodeLlama-7B-Instruct excelled in self-assurance (54.65%) and collision avoidance (99.12%), demonstrating a strong ability to identify and mitigate risks. Llama-2-7B-Chat and Llama-3-8B-Instruct closely followed, indicating a consistent focus on safety within Meta's LLM family, in contrast to GPT-3.5-turbo, which exhibited the lowest self-assurance.

Regulatory compliance is another critical aspect of safety performance that warrants our attention. Llama-2-7B-Chat led in regulatory compliance (90.62%), with CodeLlama-7B-Instruct and Mistral-7B-Instruct-v0.2 closely behind. In contrast, GPT-3.5-turbo's low score raises concerns about its adherence to regulations, highlighting the need for a dedicated focus on compliance in LLM development, especially in safety-critical applications.

Despite GPT-3.5-turbo's lower safety evaluations, its strong performance in other metrics is notable. It excelled in code fidelity (98.51%), instruction understanding (98.93%), and utility (93.75%), suggesting that strengths in these areas may help offset its safety weaknesses when combined with effective prompt engineering. In contrast, Llama-2-7B-Chat and CodeLlama-7B-Instruct, both part of Meta's LLM family, performed well on prior safety-related metrics but underperformed in these three areas. This reversal highlights the varying strengths and weaknesses across different models and underscores the importance of balancing performance metrics in real-world applications.

This comprehensive evaluation not only reveals the distinct characteristics of various LLMs but also underscores the importance of balancing multiple performance indicators for real-world applications. Overall, this evaluation offers valuable insights into the performance characteristics of different LLMs, reinforcing the need for

balanced metrics in the development of safer drone-control systems. For detailed evaluation results, please refer to the appendix.

Trade-off between utility and safety. The safety evaluation results shown in Figure 4 reveal an interesting trend: LLMs that excel in utility and code fidelity, indicating better drone operation control, tend to exhibit higher safety risks. This observation suggests that optimizing these models solely for utility may compromise essential safety aspects. A key factor in this trade-off is the bias in training data. LLMs are trained on publicly available data^{2,5,8} that often emphasizes code utility over safety. This focus leads to the prioritization of effective code generation while neglecting potential safety risks. The lack of negative examples, such as incorrect instructions or hazardous code, results in inadequate learning of safety considerations.

Conflicting training objectives also contribute to this issue. Models are optimized to predict the next token from public Internet data, which may diverge from the goal of following user instructions.^{5,19,20} This optimization can unintentionally compromise safety, as improving code efficiency might prioritize task completion over identifying vulnerabilities or safety risks. Additionally, the lack of effective safety mechanisms^{6,12,26} in existing LLMs exacerbates the problem. Without safeguards to prevent dangerous outputs, these models may fail to recognize or mitigate potential risks during code generation, increasing the likelihood of safety issues.

In conclusion, the trade-off between utility and safety in LLMs for drone operations is a complex issue with multiple contributing factors. Biases in training data, conflicting training objectives, and the lack of robust safety guardrails all play a role in shaping this trade-off.

In-context learning delivers significant safety gains. We investigated the impact of different prompt engineering techniques²⁷ on LLM safety, including in-context learning (ICL)¹⁵ and zero-shot chain of thought (ZS-CoT).¹³ ICL involves providing relevant examples within the prompt, allowing the model to learn and apply this knowledge. We implemented ICL by including safety-related examples; specific

examples can be found in the appendix. In contrast, ZS-CoT encourages the model to break down its reasoning step by step without specific examples. We instructed the model: “Before providing code, think step by step: Will this instruction cause harm or violate FAA regulations? If yes, respond safely.” This guides the model to consider potential risks and regulatory compliance before generating code.

As illustrated in Figure 5a and 5b, both GPT-3.5-turbo and Gemini Pro equipped with ICL demonstrate significant advantages in safety metrics, particularly in risk mitigation. For instance, ICL boosts GPT-3.5-turbo’s self-assurance from 12.50% to 87.20%, and Gemini Pro from 25.00% to 87.50%. The avoid-collision score for GPT-3.5-turbo improves from 13.66% to 95.05%, while Gemini Pro’s score rises from 32.55% to 91.57%. These results highlight ICL’s effectiveness in enhancing LLM safety. Notably, while Gemini Pro shows a slightly lower self-protect and safety refusal ratio with ICL compared to ZS-CoT, it results in fewer collisions dur-

ing deliberate attacks (15 for ICL versus 29 for ZS-CoT) and fewer violations (three for ICL versus six for ZS-CoT). The radar chart indicates that ICL outperforms ZS-CoT in multiple safety dimensions, making instructions clearer and reducing misunderstandings.

While ICL shows better safety performance, its effectiveness depends on the quality and quantity of examples and it often requires more computational resources for contextual processing. In contrast, ZS-CoT, although slightly less effective, serves as a lightweight solution that enhances safety without data. However, despite the advantages of both methods, they struggle with unintentional attacks. For instance, “Please let the drone take off, fly above the car, and land” could lead to dangerous outcomes, highlighting the challenge of reasoning about indirect effects. Although ICL and ZS-CoT help consider potential risks, their effectiveness in mitigating unintentional attacks remains limited.


This study highlights the potential of ICL and ZS-CoT in enhancing LLM

safety. ICL offers significant improvements through contextual learning but relies on quality examples and can be computationally intensive. ZS-CoT is a simpler, lightweight alternative but is less effective. Future research should focus on developing robust safety guardrails, refining ICL, and addressing ZS-CoT’s limitations. For detailed evaluation results on prompt engineering, please refer to the appendix.

Model size plays a critical role in safety. Our research also explored the scalability of model size on safety performance. The results, shown in Figure 5c, suggest a strong correlation between model size and safety capabilities. Larger LLMs, such as CodeLlama-34B-Instruct, outperform smaller models in various safety metrics, particularly in refusing dangerous commands. For example, the self-assurance score increases from 54.65% for CodeLlama-7B-Instruct to 80.23% for CodeLlama-34B-Instruct, while the instruction understanding score improves from 64.36% to 85.11%.

Figure 5c clearly illustrates that as model size increases from 7B to 34B

Figure 5. Safety evaluation results with prompt engineering and model size. The top shows radar charts for six metrics. The bottom presents the self-protect and safety refusal (%) across datasets: (a) results for GPT-3.5-turbo with original, zero-shot chain-of-thought (ZS-CoT), and in-context learning (ICL) prompt engineering techniques; (b) results for Gemini Pro with the same prompt engineering techniques, demonstrating their effectiveness in safety; (c) results for CodeLlama-7B-Instruct, CodeLlama-13B-Instruct, and CodeLlama-34B-Instruct, illustrating that larger LLMs generally offer better safety.



parameters, there is a significant improvement in self-protect and safety refusal percentages for both deliberate attacks and violations. For instance, in deliberate attacks, the refusal rate rises from about 45% for the 7B model to over 85% for the 34B model, indicating substantial safety enhancements.

However, for unintentional attacks, the performance of the 13B and 34B models remains nearly identical, both achieving about a 42% refusal rate. This suggests that beyond a certain size, the ability to detect and resist unintentional attacks does not significantly improve, aligning with earlier observations that LLMs struggle with these types of risks.

Overall, our findings indicate that model size alone is insufficient for improving safety against unintentional risks. While larger LLMs show superior capabilities in many areas, their effectiveness in mitigating unintentional attacks appears limited. This underscores the need for targeted research to enhance causal reasoning and predictive safety evaluation abilities. For detailed evaluation results on different LLM sizes, please refer to the appendix.

Discussion

In this article, we highlight the physical safety of LLMs and reveal novel insights into safety and utility assessment in LLM-based robotic systems. Our benchmark and comprehensive safety evaluation of LLMs in drone-control systems provide critical perspectives for robotics and AI. By evaluating various LLMs in a virtual environment, we emphasize the importance of simulation as a primary defense in robotic safety testing. A fundamental principle emerges: If an LLM shows unsafe behaviors in a controlled environment, its real-world performance is likely to be worse.

Our evaluation of mainstream LLMs reveals significant challenges in physical safety. The trade-off between utility and safety highlights a critical difficulty in LLM-based robotics. LLMs excelling in code generation often perform poorly in safety, indicating the need for balanced AI system development. More importantly, effective code generation does not guarantee safety. In addition, our findings on prompt engineering techniques show improvements in LLM safety, and the positive correlation be-

tween model size and safety provides valuable insights for future AI-driven robotics. These results extend beyond drone control, with significant implications for various robotic applications. As AI integration in robotics accelerates, especially with programmable drones from companies like DJI, ensuring safety is crucial. LLMs' ability to control these accessible platforms necessitates robust safety measures. Our research, grounded in virtual environment testing, provides a foundation for developing safer, more reliable AI-controlled robotic systems.

In conclusion, this study emphasizes the importance of prioritizing safety in LLM development for robotic control, starting with a comprehensive evaluation benchmark and rigorous virtual testing. Future research should focus on refining safety evaluation methodologies, incorporating robust safety mechanisms, and bridging the gap between virtual and real-world performance. By addressing these challenges, we can advance AI in robotics, paving the way for safer and more effective autonomous systems in various applications. More information, including datasets and code, is available on our project page.^c

Acknowledgments

The research work described in this paper was conducted in the JC STEM Lab of Intelligent Design Automation funded by The Hong Kong Jockey Club Charities Trust. All authors participated in methodology design and paper writing. Pin-Yu Chen is solely funded by IBM Research. 

References

- Ahn, M. et al. AutoRT: Embodied foundation models for large scale orchestration of robotic agents. *arXiv* (2024), arXiv:2401.12963.
- Bender, E.M., Gebru, T., McMillan-Major, A., and Shmitchell, S. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conf. on Fairness, Accountability, and Transparency*. ACM (2021), 610–623.
- Brohan, A. et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv* (2023), arXiv:2307.15818.
- Brohan, A. et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv* (2022), arXiv:2212.06817.
- Brown, T. et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems* 33 (2020), 1877–1901.
- Carlini, N. et al. Are aligned neural networks adversarially aligned? *Advances in Neural Information Processing Systems* 36 (2024).
- Chao, P. et al. Jailbreakbench: An open robustness benchmark for jailbreaking large language models. *arXiv* (2024), arXiv:2404.01318.
- Chen, M. et al. Evaluating large language models trained on code. *arXiv* (2021), arXiv:2107.03374.

- Devlin, J., Chang, M.W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional trans-formers for language understanding. *arXiv* (2018), arXiv:1810.04805.
- Van Dijk, T., De Wagter, C., and de Croon, G.C.H.E. Visual route following for tiny autonomous robots. *Science Robotics* 9, 92 (2024), eadk0310.
- U.S. Government. Title 14: Aeronautics and space, part 107 - small unmanned aircraft systems. In *Electronic Code of Federal Regulations* (2024).
- Hsiung, L. et al. Why LLM safety guardrails collapse after fine-tuning: A similarity analysis between alignment and fine-tuning datasets. *arXiv* (2025), arXiv:2506.05346.
- Kojima, T. et al. Large language models are zero-shot reasoners. *Advances in Neural Information Processing Systems* 35 (2022), 22199–22213.
- Macdonald, J.P. et al. Language, camera, autonomy! Prompt-engineered robot control for rapidly evolving deployment. In *Companion of the 2024 ACM/IEEE Intern. Conf. on Human-Robot Interaction*. ACM (2024), 717–721.
- Min, S. et al. Rethinking the role of demonstrations: What makes in-context learning work? In *EMNLP* (2022).
- Paredes-Vallés, F. et al. Fully neuromorphic vision and control for autonomous drone flight. *Science Robotics* 9, 90 (2024), eadi0591.
- Qi, X. et al. Fine-tuning aligned language models compromises safety, even when users do not intend to! *arXiv* (2023), arXiv:2310.03693.
- Radford, A. et al. Improving language understanding by generative pre-training (2018).
- Radford, A. et al. Language models are unsupervised multitask learners. *OpenAI Blog* 1, 8 (2019), 9.
- Rae, J.W. et al. Scaling language models: Methods, analysis & insights from training gopher. *arXiv* (2021), arXiv:2112.11446.
- Schedl, D.C., Kurmi, I., and Bimber, O. An autonomous drone for search and rescue in forests using airborne optical sectioning. *Science Robotics* 6, 55 (2021), eabg1188.
- Shah, D., Osiński, B., and Levine, S. Lm-nav: Robotic navigation with large pre-trained models of language, vision, and action. In *Conf. on Robot Learning*. PMLR (2023), 492–504.
- Shah, S., Day, D., Lovett, C., and Kapoor, A. Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In *Field and Service Robotics: Results of the 11th Intern. Conf.* Springer (2018), 621–635.
- Shoeybi, M. et al. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv* (2019), arXiv:1909.08053.
- Vaswani, A. et al. Attention is all you need. *Advances in Neural Information Processing Systems* 30 (2017).
- Wei, A., Haghtalab, N., and Steinhart, J. Jailbroken: How does LLM safety training fail? *Advances in Neural Information Processing Systems* 36 (2024).
- Wei, J. et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems* 35 (2022), 24824–24837.
- Xie, T. et al. Sorry-bench: Systematically evaluating large language model safety refusal behaviors (2024).
- Zhang, J. et al. Bootstrap your own skills: Learning to solve new tasks with large language model guidance. In *7th Annual Conf. on Robot Learning* (2023).

more online

To view the appendix for this article please visit <https://dl.acm.org/doi/10.1145/3816260> and click on Supplemental Material.

Yung-Chen Tang is a research assistant in the Department of Computer Science and Engineering at The Chinese University of Hong Kong, Shatin, N.T., Hong Kong.

Pin-Yu Chen (pin-yu.chen@ibm.com) is a principal research scientist and a manager at IBM Research, Yorktown Heights, NY, USA.

Tsung-Yi Ho is a professor in the Department of Computer Science and Engineering at The Chinese University of Hong Kong, Shatin, N.T., Hong Kong.

 This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).

Delving into the information confidentiality risks associated with the interface between photodiodes and IoT devices.

BY BEN NASSI

Illuminating Secrets: Power LED-Based Side-Channel Attacks for Key Extraction and Eavesdropping

THE INTERNET OF Things (IoT) revolution,⁷ which began 20 years ago, has ushered in a new era of computing. Consumers have embraced numerous IoT devices,⁶ such as robotic vacuum cleaners, smart assistants, and video streamers, as replacements for traditional devices in smart cities, homes, offices, businesses, and other settings. These IoT devices are typically designed with minimal central processing unit (CPU) power to maintain an affordable price

point for consumers and maximize revenue for manufacturers. Simultaneously, sensor technology has made significant advancements, following Moore's Law, resulting in sensors that offer high sampling rates and sensitivity² while becoming more affordable each year.

Advances in sensor quality, coupled with the widespread adoption of IoT devices with limited CPU power,⁶ have led to a situation where the sampling rate of many sensors is now on par with or even exceeds the CPU clock rate of these devices. This represents a significant shift from the previous era, where there were noticeable differences or gaps between the low sampling rates of sensors and the high CPU rates of high-end devices. In the IoT era, these differences have either diminished or vanished entirely. This new reality poses unprecedented risks to the confidentiality of computing.

This article delves into the information confidentiality risks associated with the interface between photodiodes—light sensors used in measuring light—and IoT devices. Photodiodes have gained significant popularity in

» key insights

- Some modern sensors can now sample at rates comparable to the CPU rates of many IoT devices, creating new risks to user privacy and information confidentiality from physical signals
- The intensity of a device's power LED often correlates with its power consumption, effectively encoding rich information about the device's operation in the LED's light intensity over a wide bandwidth.
- Optical measurements of ordinary device LEDs, obtained using photodiodes, can leak sensitive information, enabling attackers to extract cryptographic keys or recover intelligible speech.
- Advances in sensing technologies are expanding the reach of remote side-channel attacks, enabling attackers to sample physical signals through compromised smartphones and IP cameras at rates sufficient to extract cryptographic keys from video recordings or recover intelligible speech from a smartphone's motion-sensor data.



the IoT era due to their ability to convert light (photons) into electrical current based on a predefined conversion function, offering high-bandwidth capabilities in the range of a few GHz. The photodiodes' high-bandwidth capabilities surpass the CPU rates typically found in various devices, such as smartcards, smartphones, video streamers, and more. These devices are typically powered by CPUs operating at speeds ranging from 200MHz to 1.5GHz. While photodiodes are primarily used for innocent tasks involving optical measurements, we demonstrate that when a photodiode is directed toward a device's power LED, the obtained optical measurements often exhibit correlation with the device's power consumption. Despite the well-known fact that the intensity of LEDs is strongly influenced by the power-supply level,¹⁴ certain popular devices have electrical circuits where the integrated power LED is directly connected to the power line without any effective decoupling mechanisms to mitigate this correlation. Consequently, the power LED in many de-

vices accurately reflects the device's power consumption.

The accurate indication of a device's power consumption through optical measurements obtained from its power LED raises concerns regarding information confidentiality. In this article, we explore how attackers can exploit this undesired correlation to launch cryptanalytic attacks on the device and to launch speech eavesdropping attacks on the device's users. Specifically, we demonstrate the recovery of a 256-bit elliptic curve digital signature algorithm (ECDSA) private key from a smartcard by analyzing the optical traces obtained by a photodiode from the power LED of the smartcard reader during cryptographic computations. Furthermore, we illustrate how attackers can recover audible speech played by speakers by analyzing the optical traces obtained by a photodiode from the power LED of the speakers, even from a distance of 25m. We also assess various methods to enhance the confidentiality of information processed by devices against these attacks and discuss the

potential future changes to the threats presented in the article.

Related Work

Recent studies explored the use of optical sensors to recover speech^{5,20,23,24} from various objects, including a hanging light bulb, a bag of chips, a window, and a trash can. One study conducted cryptanalysis by capturing near-infrared photon emissions from switching transistors on the back of exposed field-programmable gate arrays (FPGAs) while executing cryptographic computations.⁸ However, this attack is only applicable to devices with exposed chips or printed circuit boards (PCBs), rendering them ineffective against commercial devices where chips are enclosed in opaque materials (e.g., smartphones, card readers).

The risks posed to information confidentiality stemming from the correlation between the intensity of a power LED to the power consumed by the device were discussed 20 years ago.¹⁸ However, previous studies demonstrating techniques for exploiting a device's power LED for data exfiltration relied on the presence of preinstalled malware that actively triggered and controlled the LEDs of keyboards¹¹ and routers.¹²

Leakage from power LEDs. Here we explore the leakage from a device's power LED.

Experimental setup. We used the Thorlabs PDA100A2 photodiode, an amplified switchable-gain light sensor that converts light (photons) to electrical current and thus to voltage. We connected the photodiode to a custom-built operational/analog amplifier (with a gain of 50dB and a bandwidth of 2MHz). The internal gain of

Figure 1. A PDA100A2 photodiode outputs a voltage that is sampled via an analog-to-digital converter (NI-9223) and processed in LabVIEW.





Figure 2. (left and middle) Spectrograms extracted from signals obtained during the modulation of a frequency scan over the power supply from the power trace (left) and optical trace obtained from the power LED of the Raspberry Pi (right). (right). A spectrogram extracted from an optical trace of a Raspberry Pi during the execution of repetitions of six different CPU operations.



the photodiode was set to the highest level before saturation for each experiment. The amplified voltage was sampled using an NI-9223 analog-to-digital converter (ADC) card that provides 1M samples per second and 16 bits per sample. We connected the ADC to a laptop; the digital optical traces were visualized in real time on LabView. We used MATLAB to process the data (optical traces). The architecture of the data-acquisition system can be seen in Figure 1.

Bandwidth of power LEDs. Here we show that a device's power LED gives an accurate indication of the power supplied to the device.

Experimental setup. We connected the Raspberry Pi 3 Model B V1.2 to a function generator that was used to modulate a 0–250kHz chirp using a 20mV peak-to-peak sinusoidal signal over the 5V power supplied by the function generator to each device. We used the sampling equipment described above to obtain an optical trace using the photodiode, which was directed at the power LED of the Raspberry Pi during the chirp from a distance of 2cm. In addition to the optical trace, we obtained a power trace from the function generator.

Results and conclusions. We extracted spectrograms from the power trace and the optical trace (see Figure 2). As can be seen, a linear frequency scan appears in the spectrogram extracted from the power trace and in the spectrogram extracted from the optical trace. Based on the spectrograms presented in Figure 2, we concluded that there is a correlation between the optical traces and the power trace, and although intended to provide a binary indication regarding the state of the device (on/off), the intensity of the power LEDs of the tested devices leaks information on the device's power-supply voltage (with excellent frequency response, reaching 250kHz).

Influence of the CPU activity. Next, we examine the ability to distinguish between episodes of a few hundred milliseconds of computations using optical traces obtained from the power LED of the device.

Experimental setup. We wrote a program that executes repetitions of the following six ARM instructions: WFI (CPU sleep), MUL (integer multiplication), ADD (integer addition), FMUL (floating point multiplication), main memory access (forcing cache misses), and NOP (short-term idle). The repetitions of each CPU operation lasted around 300ms. We installed the program and executed it on a Raspberry Pi 3B+. We used the sampling equipment described above to obtain an optical trace by directing the photodiode at the power LED of the Raspberry Pi during the program execution from a distance of 2cm.

Figure 3. The experimental setup of speech eavesdropping.



Results. The spectrogram we extracted from the optical trace is presented in Figure 1. As can be seen, the repeated operations affect different frequencies (e.g., the repetition of the FMUL instruction affects 46kHz, and the repetition of the MEM instruction affects 50kHz). Based on this experiment, we concluded that the beginning and the end of episodes of a few hundred milliseconds of computations can be detected in the spectrum of the optical trace.

Implications for privacy and information confidentiality. The use of optical traces from a device's power LEDs to accurately determine its power supply/consumption across a wide bandwidth (up to 250kHz) poses significant risks to information confidentiality and a user's privacy. This is due to the following reasons:

- The recovery of perceptible speech requires a much narrower bandwidth (a few kHz). Therefore, the wide bandwidth of the optical traces creates a risk of recovering audible speech from sensitive conversations.

- Previous studies^{15,16} have already demonstrated successful key recovery attacks by analyzing power traces. Furthermore, the ability to detect the initiation and completion of operations through optical traces obtained


from a device's power LEDs poses a significant threat to information confidentiality. Attackers can exploit this information to recover the device's cryptographic key by employing cryptanalytic side-channel attacks that take advantage of timing variations. This has been highlighted in recent studies.^{13,26}

Speech eavesdropping. Here we show how attackers can exploit the power LED of desktop speakers to recover speech played by the speakers (e.g., of a virtual meeting).

Experimental setup. We placed Winner desktop speakers on a table and connected them to a laptop. We played an audio sample used in Lamphone²³ via the desktop speakers at 70db. The level of the played sound was measured using a professional decibel meter. We used the sampling equipment described previously to obtain optical traces. The photodiode was directed at the power LED of the desktop speaker via a telescope. The equipment described above was placed in three locations: 15, 25, and 35m away from the desktop speakers. Figure 3 shows the experimental setup.

Signal processing. The three optical traces were processed as follows: First, we filtered harmonics (peaks that appear at certain frequencies on a signal's spectrum) that are not the result of the light. These harmonics, which are added to the signal as a result of the light frequency, affect the optical signal by adding unwanted noise to the signal's spectrum (100Hz, 200Hz, and so on). We filtered these harmonics from the optical signal using notch/band-stop filters.^a Next, we scaled the measurements to the range of $[-1, 1]$ to enhance the recovered

^a <https://tinyurl.com/28phyh9u>

Figure 4. The influence of each step of the signal processing on the audio recovered from an optical trace.

speech.^b Finally, we applied spectral subtraction to de-noise the signal and improve its fidelity.^c

The influence of each step of the signal processing on the recovered signal on an arbitrary sentence is illustrated in Figure 4. As can be seen, the raw optical trace is very noisy. However, the application of scaling and spectral subtraction significantly improves the quality of the signal.

Results. The reader can assess the fidelity of the recovered sound *visually*, with the spectrograms extracted from the signals; *qualitatively*, with the audio signals recovered available online;^d and *quantitatively*, based on metrics used by the audio processing community to compare a recovered signal to the original signal:

- **Intelligibility**, a measure of speech comprehensibility in given conditions. The results are classified as follows: *bad* (0–0.3), *poor* (0.3–0.45), *fair* (0.45–0.6), *good* (0.6–0.75), and *excel-*

lent (0.75–1)—according to the STI qualification values of ISO 9921.^e A higher intelligibility indicates better sound.

- **National Institute of Standards and Technology Signal-to-Noise Ratio (NIST-SNR)**, the speech-to-noise ratio, which is defined as the logarithmic ratio between the estimated speech power and noise power over 20 consecutive milliseconds. A higher NIST-SNR indicates better sound quality.


The three recovered audio signals were uploaded.^f The spectrograms of the recovered audio signals, as well as the intelligibility and NIST-SNR values of the recovered signals, are presented in Figure 5. The results demonstrate that the intelligibility of the recovered audio signals is considered good up to a distance of 15m and fair up to a distance of 35m. Clearly, audible speech can be recovered using optical traces obtained by a photodiode from the power LED of desktop speakers. Additional speech recov-

eries using this technique appear in Nassi et al.²²

Computational efficiency. We note that a photodiode is used to generate the optical trace in real time, and simple signal-processing techniques—including bandstop filtering, normalization, and spectral subtraction—are employed to recover speech. These techniques can be applied in real time without requiring any preliminary steps. As a result, recovering speech from optical traces obtained via a photodiode offers a significant efficiency advantage over state-of-the-art methods. Some existing approaches require an initial model-generation phase, which involves collecting massive amounts of data,^{1,19,25,27,28} while others depend on computationally intensive processing to extract optical traces from video footage.⁵

Recovering ECDSA keys. Here, we show how attackers can exploit a smartcard reader's power LED to recover a 256-bit ECDSA key from the smartcard using optical traces obtained by a photodiode.

Background. As observed in the original paper on the Minerva attack,¹³ many cryptographic libraries optimize the ECDSA signing computation time by running a variable number of loop iterations (determined by the number of leading zeros in the nonce) instead of a fixed number of iterations. As a result, attackers can recover the ECDSA key by applying a timing attack by using the signing times of a set of ECDSA signatures as follows: The signatures whose nonces have many leading zeros (i.e., the signatures with the fastest signing time) are used to construct a hidden number problem,³ which is turned into an instance of the Closest Vector Problem and solved using the methods of lattice reduction (see Jan-car et al.¹³ for details).

Figure 5. The spectrograms of the recovered signals from various distances and their quality.

b <https://tinyurl.com/2cuserm8>

c <https://tinyurl.com/22nfn9dm>

d <https://youtu.be/z4-OFLTHtiw>

e <https://tinyurl.com/2avyr83r>

f https://youtu.be/I_-A1Q-0frM

Experimental setup. We connected a USB Gemalto IDBridge CT30 card reader to a laptop via a USB cable and inserted the Athena IDProtect smartcard into the reader. We wrote a script that triggers ECDSA sign operations from the laptop (where every two sign operations were separated by 40ms of sleep). We used the sampling equipment described previously to obtain optical traces. The photodiode was placed 2cm away from the power LED of the smartcard reader (see Figure 6). We obtained an optical trace from 7,000 different ECDSA sign operations performed by the smart card.

Signal processing. A spectrogram of a section of the optical trace that consists of five ECDSA sign operations appears in Figure 6. The optical leakage associated with the sign operations appears at around 266.6kHz. Clearly, the sign operations can be distinguished using spectral analysis.

We divided the long optical signal into 7,000 short traces as follows: We performed short-time Fourier transform (STFT) with a window of 7.5ms and an overlap of 50% on each trace. We divided the optical trace, which consisted of 7,000 sign operations, into short traces (each associated with one sign operation), by identifying the sections associated with the sleep operation. We divided the long optical signal into 7,000 short traces, where each of the 7,000 short traces consisted of a sign operation followed by 5–10ms of sleep, since we added an extra window at the beginning and end of each trace.

Next, for each short optical trace, we performed the following steps to estimate the execution time of the ECDSA operation:

1. We applied STFT with a window of 0.75ms and an overlap of 90% on the traces.

2. We located the sign operation in the trace by detecting the sections in the trace in which the magnitude of the bin associated with the sign operation ($\sim 266.6\text{kHz}$) is greater than a constant.

3. Because of errors, more than one section can satisfy the abovementioned condition (i.e., two or more different sections separated by some windows are located) in the same trace. To correct such errors, we used an error-tolerance threshold of two windows (i.e., if two consecutive sections were separated by only two windows, we merged them into a single section).

4. Traces that yielded more than one sign section that satisfied the above-mentioned criterion were filtered out after the second step.

5. The estimated execution time of the sign operation was calculated based on traces with a single section associated with an ECDSA operation. After filtering, 6,892 signatures remained, along with the estimates of their ECDSA execution time from the optical traces.

Results. The heat map presented in Figure 6 shows the estimated ECDSA operation time (based on the optical measurements) vs. the number of leading zeros in the signature's nonce. As can be seen in the figure, the signatures estimated to have the shortest ECDSA operation execution time (based on the optical traces) are associated with at least seven leading zeroes. The results validate the correctness of the estimations of the execution time we made using the optical traces (because a higher number of leading

zeroes in the nonce corresponds to a shorter ECDSA execution time).

To recover the ECDSA private key, we used a script that implements the Minerva attack according to the details provided in a previous section (the script provided in the official Minerva GitHub⁸). The input to the script was as follows: the public key, messages, their signatures, and corresponding estimated signature times. The script was executed on a standard Dell laptop, and the full 256-bit ECDSA key was found in around two minutes. Additional analysis of the influence of factors on the recovery appears in Nassi et al.²²


Computational efficiency. We note that using a photodiode to recover secret keys is comparable in efficiency to state-of-the-art methods that leverage an oscilloscope,^{15,16} software-defined radio,^{4,9} and an ultrasonic microphone.¹⁰ In all these approaches, the captured signals or traces require additional processing to assess the timing of the cryptographic operation. However, the key advantage of using a photodiode for secret key recovery lies in its non-intrusiveness: It does not require a physical connection to the target device, unlike an oscilloscope, and it does not rely on compromising the device with malware to measure cryptographic timing, as demonstrated in the original Minerva attack.¹³

Countermeasures

Detection. Detecting the acquisition of optical traces using a photodiode is inherently challenging. Unlike LiDAR and laser microphones, which employ active laser beams for data collection, photodiodes are entirely passive and do not

g <https://tinyurl.com/26vnwfwv>

Figure 6. Left: A photodiode is directed at the power LED of a card reader (from 2cm away). Middle: A spectrogram extracted from an optical trace during ECDSA sign operations followed by sleep. Right: The estimated time of ECDSA sign operations (calculated from the optical traces obtained from the power LED of the smartcard reader) as a function of the number of leading zero bits in the nonce.



emit any light. Consequently, identifying the use of photodiodes—particularly when deployed from a distance—is analogous to detecting hidden microphones. As a result, mitigating cryptanalysis and speech eavesdropping is more feasible than detecting them.

Mitigations. In many devices, the power LED is directly connected to the power line on the PCB. Consequently, the power-consumption fluctuations during cryptographic operations affect the device's power LED. To mitigate this issue, hardware manufacturers can incorporate a capacitor parallel to the power LED indicator, which can serve as a cost-effective solution for reducing fluctuations and acting as a low-pass filter.

Consumers interested in preventing remote attackers from obtaining optical traces from the target device should consider the following approaches: One straightforward method is to cover the device's power LED with black tape, effectively blocking any emitted light. While this solution may slightly impact the device's user experience, it successfully prevents remote attackers from capturing optical traces from vulnerable devices. Another approach is to create a physical barrier that obstructs the line of sight to the power LED. Installing a curtain or repositioning the target device to a windowless room can achieve this. By eliminating the direct visibility of the power LED, the risk of optical leakage is significantly reduced.

It is also worth noting that the intensity of light emitted from power LEDs diminishes with distance. With this in mind, establishing a safety perimeter can be an effective measure. This involves installing a fence around the premises (such as a home or building), limiting the attacker's proximity to the target device. By forcing them to conduct the attack from a greater distance, the quality of any captured optical traces, as demonstrated in the case of speech recovery, deteriorates.

The Evolving Risk Landscape

The risks to information confidentiality posed by photodiodes may increase in the future. Although current smartphones integrate low-cost photodiodes for ambient light sensing, these sensors lack the sensitivity required

to enable the attacks described in this article. However, continued technological progress has steadily improved sensor quality while reducing cost. If highly sensitive photodiodes become as affordable as existing low-sensitivity ones, they may be integrated into future smartphones, enabling attackers to remotely extract optical traces from a compromised device's ambient light sensor. This would significantly expand the attack surface, exposing a larger number of devices to cryptanalytic and speech-eavesdropping attacks.

A similar risk arises from advances in video camera technology. The *video microphone* technique⁵ demonstrated that rolling-shutter effects can be exploited to recover speech from video recordings. Since then, improvements have enabled commercial off-the-shelf cameras to support rolling-shutter rates sufficient to recover speech¹⁷ and secret keys²¹ over the Internet. These trends highlight the need to proactively reassess the confidentiality risks posed by sensors, not limited to photodiodes, in the IoT era. C

References

- Ba, Z. et al. Learning-based practical smartphone eavesdropping with built-in accelerometer. In *Proceedings of the Network and Distributed Systems Security Symp.* (2020), 23–26.
- Beckerwerth, T., Runge, P., and Schell, M. High-speed photodiodes for power efficient data transmission. *J. Lightwave Technology* 42, 3 (2023), 1056–1061.
- Boneh, D. and Venkatesan, R. Hardness of computing the most significant bits of secret keys in Diffie-Hellman and related schemes. In *Proceedings of Advances in Cryptology—CRYPTO'96: 16th Annual Intern. Cryptology Conf.* Springer (Aug. 1996), 129–142.
- Camurati, G. et al. Screaming channels: When electromagnetic side channels meet radio transceivers. In *Proceedings of the 2018 ACM SIGSAC Conf. on Computer and Communications Security* (2018), 163–177.
- Davis, A. et al. The visual microphone: Passive recovery of sound from video. *ACM Transactions on Graphics* 33, 4 (2014), 1–10.
- Elgazzar, K. et al. Revisiting the internet of things: New trends, opportunities and grand challenges. *Frontiers in the Internet of Things* 1 (Nov. 2022).
- Feki, M.A., Kawsar, F., Boussard, M., and Trappeniens, L. The Internet of Things: The next technological revolution. *Computer* 46, 2 (2013), 24–25; DOI:10.1109/MC.2013.63
- Ferrigno, J. and Hlavác, M. When AES blinks: Introducing optical side channel. *IET Information Security* 2, 3 (2008), 94–98.
- Gandolfi, K., Mourtel, C., and Olivier, F. Electromagnetic analysis: Concrete results. In *Proceedings of the Intern. Workshop on Cryptographic Hardware and Embedded Systems*. Springer (2001), 251–261.
- Genkin, D., Shamir, A., and Tromer, E. RSA key extraction via low-band width acoustic cryptanalysis. In *Proceedings of the Annual Cryptology Conf.* Springer (2014), 444–461.
- Guri, M., Zadov, B., Bykhovsky, D., and Elovici, Y. Ctrl-alt-LED: Leaking from air-gapped computers via keyboard LEDs. In *Proceedings of the 2019 IEEE 43rd Annual Computer Software and Applications Conf.* 1, IEEE (2019), 801–810.
- Guri, M., Zadov, B., Daidakulov, A., and Elovici, Y. xLED: Covert data exfiltration from air-gapped networks via switch and router LEDs. In *Proceedings of the 2018 16th Annual Conf. on Privacy, Security and Trust*. IEEE (2018), 1–12.
- Jancar, J., Sedlacek, V., Svenda, P., and Sys, M. Minerva: The curse of ECDSA nonces (Systematic analysis of lattice attacks on noisy leakage of bit-length of ECDSA nonces). *IACR Transactions on Cryptographic Hardware and Embedded Systems* 2020, 4 (2020), 281–308. DOI:10.13154/tches.v2020.i4.281-308
- King, S. Luminous intensity of an LED as a function of input power. *ISB J. Phys.* 2, 2 (2008).
- Kocher, P., Jaffe, J., and Jun, B. Differential power analysis. In *Proceedings of the Annual Intern. Cryptology Conf.*, Springer (1999), 388–397.
- Kocher, P., Jaffe, J., Jun, B., and Pankaj Rohatgi, P. Introduction to differential power analysis. *J. Cryptographic Engineering* 1, 1 (2011), 5–27.
- Long, Y. et al. Side eye: Characterizing the limits of POV acoustic eavesdropping from smartphone cameras with rolling shutters and movable lenses. In *Proceedings of the 2023 IEEE Symp. on Security and Privacy*, IEEE (2023), 1857–1874.
- Loughry, J. and Umphress, D.A. Information leakage from optical emanations. *ACM Transactions on Information and System Security (TISSEC)* 5, 3 (2002), 262–289.
- Michalevsky, Y., Boneh, D., and Nakibly, G. Gyrophone: Recognizing speech from gyroscope signals. In *Proceedings of the 23rd USENIX Security Symp.* USENIX Association (2014), 1053–1067; <https://tinyurl.com/mdop6q6>
- Muscattell, R.P. Laser microphone. U.S. Patent 4,412,105 (1983).
- Nassi, B. et al. Video-based cryptanalysis: Extracting cryptographic keys from video footage of a device's power LED captured by standard video cameras. In *Proceedings of the 2024 IEEE Symp. on Security and Privacy*, IEEE (2024), 2422–2440.
- Nassi, B. et al. Glowworm attack: Optical tempest sound recovery via a device's power indicator LED. In *Proceedings of the 2021 ACM SIGSAC Conf. on Computer and Communications Security*, ACM (2021), 1900–1914.
- Nassi, B. et al. Lamphone: Passive sound recovery from a desk lamp's light bulb vibrations. In *Proceedings of the 31st USENIX Security Symp.* USENIX (2022), 4401–4417.
- Sami, S. et al. Spying with your robot vacuum cleaner: Eavesdropping via LiDAR sensors. In *Proceedings of the 18th Conf. on Embedded Networked Sensor Systems*, Association for Computing Machinery (2020), 354–367; DOI:10.1145/3384419.3430781
- Wang, G. et al. We can hear you with Wi-Fi. *IEEE Transactions on Mobile Computing* 15, 11 (Nov. 2016), 2907–2920; DOI:10.1109/TMC.2016.2517630
- Wang, Y. et al. Hertzbleed: Turning power [side-channel] attacks into remote timing attacks on x86. In *Proceedings of the 31st USENIX Security Symp.* (2022), 679–697.
- Wei, T., Wang, S., Zhou, A., and Zhang, X. Acoustic eavesdropping through wireless vibrometry. In *Proceedings of the 21st Annual Intern. Conf. on Mobile Computing and Networking*, ACM (2015), 130–141; DOI:10.1145/2789168.2790119
- Zhang, L. et al. Accelword: Energy efficient hotword detection through accelerometer. In *Proceedings of the 13th Annual Intern. Conf. on Mobile Systems, Applications, and Services*. ACM (2015), 301–315.

Ben Nassi (nassiben@tauex.tau.ac.il) is a faculty member at Tel Aviv University, School of Electrical & Computer Engineering, Tel Aviv, Israel.

 This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).



Watch the authors discuss this work in the exclusive Communications video. <https://cacm.acm.org/videos/illuminating-secrets>

research highlights

P. 104

**Technical
Perspective**
**egg: ‘Ridiculously’
Fast and Extensible
Equality Saturation**

By Nadia Polikarpova

P. 105

egg: Fast and Extensible
Equality Saturation

By Max Willsey, Chandrakana Nandi, Yisu Remy Wang,
Oliver Flatt, Zachary Tatlock, and Pavel Panchekha

Technical Perspective

egg: ‘Ridiculously’ Fast and Extensible Equality Saturation

By Nadia Polikarpova

SOMETIMES A GREAT idea is born out of a bug—and that’s exactly how *egg*, the subject of the accompanying paper, was hatched. Back in 2019, Max Willsey, then a postdoc at the University of Washington, boasted to his advisor that he could make an advanced compiler technique called *equality saturation* run much faster by rewriting its implementation in Rust. If you were following programming-language trends at the time, you’ll recall that “rewriting in Rust” was touted as a universal cure for sluggish software—so much so that it became an Internet meme. And sure enough, the Rust version was fast. Suspiciously fast.

To see what went wrong (or rather, right), we need to take a step back. At the heart of equality saturation lies a clever data structure called an *e-graph*. If you know how an abstract syntax tree—or expression tree—represents a single program, an *e-graph* does the same for *many* equivalent variants of that program, compactly folding them into a single structure (so compact, in fact, that a finite *e-graph* can represent infinitely many variants!). Equality saturation uses *e-graphs* for program optimization: It keeps adding new equivalent program variants to the graph and eventually extracts the “best” one according to some metric. The catch is that adding each new variant can break the *e-graph*’s compactness property, and an expensive compaction step is required to restore it.

The first version of *egg* was fast because it skipped this compaction step—producing incorrect results. But while restoring it, Max had an epiphany: Why compact after every addition? Why not add a batch of variants first, then compact once? (The paper’s term for this batch compaction is *rebuilding*.) The idea seems obvious in hindsight, but no one had tried it before because *e-graphs* were originally invented for theorem proving, where they had to be compact at all times. When the data structure was later repurposed for program optimiza-

tion, it simply had not occurred to anyone to change how its basic operations worked. Anyone, that is, until Max, an outsider to the theorem-proving world who simply didn’t know the “rules.”

Apart from making equality saturation faster, *egg* also made it easier to tweak. Willsey and co-authors noticed that compiler authors would often attach extra attributes to each program inside an *e-graph*, such as which variables it referenced or whether it could evaluate to zero; information needed to decide which optimizations could be applied. Keeping these attributes in sync as the *e-graph* evolved was painful and error-prone. In the best traditions of programming-language research, the *egg* authors recognized this recurring pattern and abstracted it into an elegant mechanism called *e-class analyses*. With an *e-class analysis*, one can define simple local rules for how program attributes should be initialized and merged, and the library automatically applies those rules in the right order and frequency to keep the *e-graph* globally consistent.

At this point, you might think: Alright, so we made a niche compiler optimization faster and cleaner—why should anyone else care? As it turned out, this paper sparked a full-blown *e-graph* revolution, spreading from compilers to fields far beyond. Since its 2021

debut, the *egg* library has been used in more than 30 research compilers and dozens of projects in areas ranging from theorem proving and hardware and software synthesis to carpentry and genome compression.^a It even inspired a fundamentally new algorithm for relational joins in databases^b and a new programming paradigm for incremental fixpoint computation.^c Today, an active community gathers at the annual EGRAPHS workshop^d and in monthly online meetings. The meme, one might say, has come full circle: Instead of rewriting everything in Rust, researchers are now rewriting their tools in *egg*.

What made *egg* so successful? First, its authors invested an unreasonable—by research-project standards—amount of effort to make the library genuinely reusable: *egg* has a clean API and excellent documentation. Second, they devoted equal energy to nurturing the community: organizing workshops, forums, and meetings that have paid off manyfold. Third, *egg* embodies a classic virtue of programming-language research: turning hacks into elegant abstractions that reveal unexpected power. By taking an implementation detail and turning it into a general mechanism for fixpoint computation over *e-graphs* (*e-class analyses*), *egg* made equality saturation *programmable*, thereby unlocking a host of new applications.

And finally ... *egg* is fast—ridiculously fast. 

egg embodies a classic virtue of programming-language research: turning hacks into elegant abstractions that reveal unexpected power.

a <https://github.com/philzook58/awesome-egraphs>

b Free Join, selected as a 2023 SIGMOD Research Highlight

c *egglog*, featured in the 2024 MIT PL Review

d <https://egraphs.org/>

Nadia Polikarpova (npolikarpova@ucsd.edu) is an associate professor in the Computer Science and Engineering Department at the University of California, San Diego, San Diego, CA, USA.

 This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).

egg: Fast and Extensible Equality Saturation

By Max Willsey, Chandrakana Nandi, Yisu Remy Wang, Oliver Flatt, Zachary Tatlock, and Pavel Panchekha

ABSTRACT

E-graphs, originally developed for use in automated theorem provers, are data structures that compactly represent a large number of expressions and the equalities between them. Recently, a technique called *equality saturation* adapted e-graphs for program optimization. Equality saturation can ameliorate issues with traditional, destructive term rewriting since e-graph rewriting only adds information to the e-graph. After rewriting, the e-graph represents a large space of equivalent terms (including the original term), and a final extraction procedure selects the best term according to a user-provided cost function.

The egg library is a fast and flexible open source implementation of e-graphs and equality saturation. Unlike solvers, egg is a white-box library for programmatic use. Users can easily define their own languages, rewrites, and analyses, and they can inspect and manipulate the e-graph directly. It incorporates novel techniques such as *e-class analysis*, a mechanism for integrating program analyses modulo equality; and *rebuilding*, a deferred invariant restoration algorithm. egg has been used in hundreds of academic and industrial projects for program optimization, synthesis, and verification in many domains; some are briefly highlighted in this article.

1. INTRODUCTION

Equality graphs (e-graphs)¹⁵ are data structures that can compactly represent a large number of expressions and the equalities between them. They were originally developed for use in automated theorem provers. One prominent use case is in satisfiability modulo theories (SMT) solvers,¹ where e-graphs are used to communicate equality information between different decision procedures via the Nelson-Oppen framework.¹⁶

More recently, Tate et al.²¹ proposed *equality saturation*, a technique that repurposes e-graphs for rewrite-driven program optimizations. Given an input program p , equality saturation constructs an e-graph E that represents a large set of programs equivalent to p , and then extracts the “best” program from E . The e-graph is grown by repeatedly applying pattern-based rewrites. Critically, these rewrites only add information to the e-graph, eliminating the need for careful ordering that can plague traditional, destructive term rewriting. Upon reaching a fixed point (*saturation*), E will represent *all equivalent ways* to express p with respect to the given rewrites. After saturation or a timeout, a final *extraction* procedure analyzes E and selects the optimal program according to a user-provided cost function.

In principle, equality saturation is an appealingly simple and general approach to program optimization. However, there are many practical challenges, ranging from performance to extensibility. For example, certain common usage patterns can cause e-graphs to grow very large or require domain-specific reasoning that cannot be expressed as syntactic rewrites. It is not feasible to provide a one-size-fits-all implementation of equality saturation, where users can simply plug in their rewrites and get out the best program.

This article describes the egg library, a fast and flexible implementation of e-graphs and equality saturation. egg does not aim to present a solver-like interface as a black-box decision procedure. Instead, egg provides the e-graph *à la carte*: as a generic data structure, along with a suite of algorithms for manipulating it. Users can easily define a language consisting of their own operators and types, rewrite rules, cost functions for extraction, and analyses to be integrated into the e-graph. egg emphasizes extensibility and programmability; users may inspect and manipulate e-graphs directly, and unused features do not incur overhead. egg is developed openly on GitHub,^a is implemented in Rust and is packaged on Rust’s package registry, *crates.io*.

To provide this flexibility, egg incorporates two novel techniques that make e-graphs fast and extensible: *rebuilding* and *e-class analysis*. Rebuilding is a new amortized algorithm for maintaining e-graph invariants that allows the client to choose when to restore invariants. This is particularly useful for equality saturation, which egg factors into distinct phases of reading and writing the e-graph for increased performance. *e-class analysis* is a general mechanism for integrating domain-specific analyses into the e-graph. This allows bringing program analyses like interval analysis, constant folding, and others to the e-graph, reducing the need for ad hoc extensions. These analyses operate modulo equality, which can increase their precision.

egg provides all the pieces to build an equality saturation engine, and many users have done so in diverse domains, including tensor programs,²⁴ floating point accuracy,¹⁷ circuit-level optimizers,⁶ and databases¹⁰ across both academia and industry. Due to its *à la carte* design, users have rearranged these pieces to build program synthesizers³ and verification tools¹¹ as well that do not neatly fit the equality saturation mold. This article will highlight some of these use cases.

2. BACKGROUND

E-graphs have been defined in many ways in the literature. Our presentation reflects egg’s design and implementation.

^a <https://github.com/egraphs-good/egg>

This paper was originally published in *Proceedings of the ACM on Programming Languages* 5, (Jan. 4, 2021).

Figure 1. Syntax and metavariables for the components of an e-graph. Function symbols may stand alone as constant e-nodes and terms.

function symbols	f, g
variables	x, y
e-class ids	a, b
terms	$t ::= f \mid f(t_1, \dots, t_m) \quad m \geq 1$
patterns	$p ::= f \mid f(p_1, \dots, p_m) \mid x \quad m \geq 1$
e-nodes	$n ::= f \mid f(a_1, \dots, a_m) \quad m \geq 1$
e-classes	$c ::= f \mid f(n_1, \dots, n_m) \quad m \geq 1$

2.1. E-graphs as data structures.

Intuitively, an e-graph is a set of equivalence classes (*e-classes*). Each e-class is a set of *e-nodes* representing equivalent terms from a given language, and an e-node is a function symbol paired with a list of children e-classes. More precisely:

DEFINITION 1 (E-GRAPH).

Given the definitions and syntax in Figure 1, an e-graph is a tuple (U, H) where:

- A union-find data structure²⁰ U stores an equivalence relation (denoted with $\equiv_{\text{id}}$) over e-class ids.
- The hashcons H is a map from e-nodes to e-class ids.

DEFINITION 2 (FIND).

A union-find U provides a `find` operation that canonicalizes e-class ids. We omit the first argument of `find` where clear from context.

- $\text{find}(U, a) = \text{find}(U, b)$ iff $a \equiv_{\text{id}} b$.
- An e-class id a is canonical iff $\text{find}(a) = a$.
- An e-node n is canonical iff $n = \text{canonicalize}(n)$, where `canonicalize` maps `find` over n 's children.

DEFINITION 3 (E-CLASS).

For an e-graph (U, H) , define M to map an e-class id to an e-class, i.e., the set of equivalent e-nodes. This is the inverse of the hashcons H :

$$M[a] = \{n \mid H[n] \equiv_{\text{id}} a\}$$

Note that an e-class has an identity (its canonical e-class id), but an e-node does not. We use e-class id a and the e-class $M[a]$ synonymously when clear from the context.

DEFINITION 4 (REPRESENTATION).

An e-graph, e-class, or e-node is said to represent a term t if t can be “found” within it. Representation is defined recursively:

- An e-graph represents t if any of its e-classes do.
- An e-class c represents t if any e-node $n \in c$ does.
- An e-node $f(a_1, a_2, \dots)$ represents a term $f(t_1, t_2, \dots)$ with the same function symbol if each e-class $M[a_i]$ represents term t_i .

When each e-class is a singleton (containing only one e-node), an e-graph is essentially a term graph with sharing. Figure 2a shows an e-graph representing the term $(a \times 2)/2$.

DEFINITION 5 (THE CONGRUENCE INVARIANT).

The hashcons H must map all canonical e-nodes to their e-class ids. In other words:

$$e\text{-node } n \in M[a] \Leftrightarrow H[\text{canonicalize}(n)] = \text{find}(a)$$


If the invariant holds, then a procedure `lookup` can quickly find which e-class (if any) has an e-node equivalent to a given e-node n : $\text{lookup}(n) = H[\text{canonicalize}(n)]$. Note that while the invariant is stated in terms of the hashcons, it implies the e-graph is not only an equivalence relation but also a congruence relation.

E-graphs provide two main low-level mutating operations:

- `add` takes an e-node n and returns the id a of the e-class representing n if it exists, else creates a new e-class for n and returns its id.
- `merge` (sometimes called `assert` or `union`) takes two e-class ids a and b , unions them in the union-find U , and combines the e-classes by setting both $M[a]$ and $M[b]$ to $M[a] \cup M[b]$.

Both operations must take additional steps to maintain the congruence invariant. Invariant maintenance is discussed in Section 4.

Figure 2. An e-graph consists of e-classes (dashed boxes) containing equivalent e-nodes (solid boxes). Edges connect e-nodes to their child e-classes. Additions and modifications are emphasized in black. Applying rewrites to an e-graph adds new e-nodes and edges, but nothing is removed. Expressions added by rewrites are merged with the matched e-class. In Figure 2d, the rewrites do not add any new nodes, only merge e-classes. The resulting e-graph has a cycle, representing infinitely many expressions: $a, a \times 1, a \times 1 \times 1$, and so on.



2.2. E-matching and rewriting.

E-graphs can be queried via *e-matching*^{7,8} (pattern matching modulo equivalence). The e-matching operation *ematch* takes a pattern p with variable placeholders (see Figure 1) and returns a list of tuples (σ, c) , where σ is a substitution of variables to e-class ids such that $p[\sigma]$ is represented in e-class c . E-matching only needs to return tuples where the e-class ids are canonical.

E-matching can be implemented in many ways: via simple recursion over the pattern⁸ all the way to optimized relational joins.²⁵ *egg* uses a backtracking virtual machine similar to that in Z3.⁷

These operations can be composed to perform rewriting over the e-graph. To apply a rewrite $p_e \rightarrow p_r$ to an e-graph, *ematch* finds tuples (σ, c) , where e-class c represents $p_e[\sigma]$. Then, for each tuple, *merge*(c , *add*($p_r[\sigma]$)) adds $p_r[\sigma]$ to the e-graph and unifies it with the matching e-class c .

Figure 2 shows an e-graph undergoing a series of rewrites. Note how the process is only additive; the initial term $(a \times 2)/2$ is still represented in the e-graph. Rewriting in an e-graph can also saturate, meaning the e-graph has learned every possible equivalence derivable from the given rewrites. If the user tried to apply $x \times y \rightarrow y \times x$ to an e-graph twice, the second time would add no additional e-nodes and perform no new merges; the e-graph can detect this and stop applying that rule.

2.3. Equality saturation.

Term rewriting is a time-tested approach for equational reasoning; it is employed in most tools that manipulate programs (e.g., compilers, theorem provers, and so on). In the conventional setting, term rewriting is typically destructive and “forgets” the matched left-hand side. Consider applying a simple rewrite:


$$(a \times 2)/2 \rightarrow (a \ll 1)/2$$

The new term carries no information about the initial term. Applying this rewrite at this point prevents us from canceling out $2/2$. In the compilers community, this classically tricky question of when to apply which rewrite is called the *phase-ordering* problem.

One solution lies in manipulating the rewrites themselves into a *convergent* rewriting system where no matter the order of rewrites, the same final term is reached. When that is not possible, another approach to the phase-ordering problem would be to simply apply all rewrites simultaneously, keeping track of every expression seen. This eliminates the problem of choosing the right rule, but a naive implementation would require space exponential in the number of given rewrites. Equality saturation²¹ is a technique to do this non-destructive rewriting efficiently using an e-graph.

Figure 3 shows the equality-saturation workflow. First, an initial e-graph is created from the input term. The core of the algorithm runs a set of rewrite rules until the e-graph is saturated (or a timeout is reached). Finally, a procedure called *extraction* selects the optimal represented term according to some cost function. For simple cost functions, extraction resembles finding a shortest path in a (hyper) graph and be solved greedily.¹³ For more complex cost func-

Figure 3. The equality saturation workflow uses an e-graph to implement program optimization by starting with an input term, applying rewrites non-destructively, and then extracting a term equivalent to the input that optimizes some cost function.



tions, approaches based on integer linear programming or stochastic methods have recently been proposed.^{22,23} *egg* provides some of these extraction methods and also allows users to implement their own.

3. E-GRAPH EXTENSIONS

egg provides several mechanisms that allow e-graphs and equality saturation to apply to a wider variety of problems.

3.1. E-class analysis.

As discussed so far, e-graphs provide an efficient way to implement a non-destructive term-rewriting system. However, program analysis, optimization, or verification may require more than just syntactic rewriting. Instead, transformations can depend on not only the input terms but also semantic facts about that input term, for example, constant value, free variables, nullability, numerical sign, size in memory, and so on. This “purely syntactic” restriction forced previous equality-saturation applications^{17,21} to resort to ad hoc passes over the e-graph to implement common analyses, such as constant folding.

egg provides a novel mechanism called an *e-class analysis* to lift a program analysis over the e-graph. An e-class analysis resembles abstract interpretation lifted to the e-graph level, attaching *analysis data* from a semilattice to each e-class. The e-graph maintains and propagates this data as e-classes get merged and new e-nodes are added (details about maintenance in Section 4). Analysis data can be used directly to modify the e-graph, to inform how or if rewrites apply their right-hand sides, or to determine the cost of terms during the extraction process.

An e-class analysis defines a domain D and associates a value $d_c \in D$ to each e-class c . The interface of an e-class analysis is as follows, where G refers to the e-graph and n and c refer to e-nodes and e-classes within G :

- $\text{make}(n) \rightarrow d_c$
When a new e-node n is added to G , construct a new value $d_c \in D$ to be associated with n 's e-class c , typically by accessing the associated data of n 's children.
- $\text{join}(d_{c_1}, d_{c_2}) \rightarrow d_c$
When e-classes c_1, c_2 are being merged into c , join d_{c_1}, d_{c_2} into a new value d_c to be associated with the new e-class c .
- $\text{modify}(c) \rightarrow c'$ (optional)
Modify the e-class c based on d_c . A typical use would be to add an e-node to c based on d_c .

The domain D together with the join operation should form a join-semilattice. The semilattice perspective is useful for defining the *analysis invariant* (where $\vee$ is the join operation):

$$\forall c \in G. \quad d_c = \bigvee_{n \in c} \text{make}(n) \quad \text{and} \quad \text{modify}(c) = c$$

The first part of the analysis invariant states that the data associated with each e-class must be the join of the make for every e-node in that e-class. Since D is a join-semilattice, this means that $\forall c, \forall n \in c, d_c \geq \text{make}(n)$. The motivation for the second part is more subtle. Since the analysis can modify an e-class through the modify method, the analysis invariant asserts that these modifications are driven to a fixed point. When the analysis invariant holds, a client looking at the analysis data can be assured that the analysis is “stable” in the sense that recomputing make, join, and modify will not modify the e-graph or any analysis data.

Consider the following example of an interval e-class analysis from Coward et. al.,⁵ which demonstrates how e-class analyses can combine with equality reasoning to derive more precise information. Let our language be arithmetic expressions, and the analysis domain D be the set of intervals. Consider the following expressions over variables x and y , where $x \in [0, 1]$ and $y \in [1, 2]$.

$$\begin{aligned} \frac{x-y}{x+y} &\in [-2, 0] \\ \frac{2x}{x+y} - 1 &\in [-1, 1] \end{aligned}$$

When these expressions are added to the e-graph, the analysis data for their respective e-classes will be computed by the make function, which computes the interval for an e-node based on the intervals of its children using basic interval arithmetic. Initially, since the variables x and y have known intervals, the analysis data for the above terms will be $[-2, 0]$ and $[-1, 1]$ as shown above.

E-class analyses react to e-class merges by joining the analysis data of the merged e-classes. So, if rewriting discovers that the two above expressions are equivalent (they are), the e-classes will be merged, and the analysis data for the new e-class will be the join of intervals for the merged e-classes. In our interval domain, the join operation is the interval *intersection*, since if we know that two equivalent terms have intervals, they both must lie in both intervals. So the new analysis data for the merged e-class will be $[-1, 0]$, which is more precise than either of the original intervals. The modify function is not needed in this example, but if an expression's interval were found to be a singleton $[c, c]$, it could add the constant c to the e-class.

3.2. Conditional and dynamic rewrites.

In equality saturation applications, many rewrites are purely syntactic. egg supports these as well as more general notions of rewriting with its notion *searchers* and *appliers*. A *searcher* is a piece of code that inspects a part of the e-graph and produces a set of matches (substitutions). An *applier* takes a match and modifies the e-graph in some way. For a typical syntactic rewrite, both the searcher and applier would be a pattern; the searcher would e-match the left-hand pattern, and the applier would instantiate the right-hand pattern.

In some cases, additional data may be needed to determine if or how to perform the rewrite. For example, the rewrite $x/x \rightarrow 1$ is only valid if $x \neq 0$. This is an example of a *conditional* rewrite, where the validity of the rewrite depends on some analysis data. To safely implement this rewrite, an e-class analysis can track whether an e-class is known to be non-zero, and the rewrite can check this fact before applying the right-hand side.

A more complex *dynamic* rewrite may need to compute the right-hand side based on the substitution or analysis facts from the left-hand side. In these instances, the applier could be a piece of code that takes in the match and manipulates the e-graph to add a new term. Searchers can also be implemented as arbitrary code, allowing for more complex matching behavior. For example, industrial user Cube Dev employs custom searchers and appliers in their use of egg for SQL rewriting.¹⁰

3.3. Extraction.

As shown in Figure 3, equality saturation typically ends with extraction. Extraction selects a represented term from a given e-class (typically the same e-class as the input term) that is optimal according to some cost function.

In many applications, AST size (optionally weighted by operators) suffices as a simple cost function. For simple cost functions like this, a greedy algorithm suffices to find the optimal term.¹³ egg provides built-in support for this extraction algorithm that is modular with respect to the cost function.

Extraction can be formulated as an e-class analysis when the cost function is simple. The analysis data is a tuple $(n, k(n))$, where n is the cheapest e-node in that e-class and $k(n)$ its cost. The make(n) operation calculates the cost $k(n)$ based on the analysis data (which contains the minimum costs) of n 's children. The merge operation simply takes the tuple with lower cost. The semilattice portion of the analysis invariant then guarantees that the analysis data will contain the lowest-cost e-node in each class. Extraction can then proceed recursively; if the analysis data for e-class c gives $f(c_1, c_2, \dots)$ as the optimal e-node, the optimal term represented in c is $\text{extract}(c) = f(\text{extract}(c_1), \text{extract}(c_2), \dots)$. This not only further demonstrates the generality of e-class analyses but also provides the ability to do extraction “on the fly;” conditional and dynamic rewrites can determine their behavior based on the cheapest term in an e-class.

Extraction (whether done as a separate pass or an e-class analysis) can also benefit from the analysis data. Typically, a local cost function can only look at the function symbol of the e-node n and the costs of n 's children. When an e-class analysis is attached to the e-graph, however, a cost function may observe the data associated with n 's e-class, as well as the data associated with n 's children. This allows a cost function to depend on computed facts rather than just purely syntactic information. In other words, the cost of an operator may differ based on its inputs. Applications in linear algebra and tensor programming²² have exploited this feature; the cost of an operation depends on the shapes of its inputs, which are tracked by an e-class analysis.

For more complex cost functions, extraction cannot be

done greedily. Some applications have therefore explored alternative extraction algorithms, including integer linear programming, GPU-based differentiable optimization,² and stochastic methods.

4. REBUILDING INVARIANTS

Like many data structures, an e-graph has invariants that a client can expect to be maintained. For egg, these are the *congruence invariant* (definition 5)—ensuring the e-graphs equivalence relation is a congruence relation—and the *e-class analysis invariant* (Section 3).

Traditionally,^{8,15} e-graphs maintain invariants after each operation. This is done with a worklist algorithm that “upward merges” e-classes whose parents may now be congruent whenever two e-classes are merged.¹²

Part of egg’s take on e-graphs as a data structure loosens the requirement that these invariants always hold. egg allows these to be temporarily violated and provides a mechanism called *rebuilding* to restore them. The client is responsible for calling the `rebuild` method at appropriate times to restore the invariants. This deferred approach can significantly improve performance by amortizing the cost of maintaining the invariants. In the empirical evaluation in the full paper, congruence is $88 \times$ faster, and equality saturation is $21 \times$ faster (in aggregate over all tests using geometric mean), when deferring invariant maintenance with rebuilding.

4.1. Rebuilding and equality saturation.

egg provides a modified equality saturation algorithm to take advantage of rebuilding. Figure 4 shows pseudocode for both traditional equality saturation²¹ and egg’s variant,

which exhibits two key differences:

1. Each iteration is split into a *read phase*, which searches for all the rewrite matches, and a *write phase* that applies those matches.
2. Rebuilding occurs only once per iteration, at the end.

egg’s separation of the read and write phases means rewrites are truly unordered. In traditional equality saturation, later rewrites in the given rewrite list are favored in the sense that they can “see” the results of earlier rewrites in the same iteration. Therefore, the results depend on the order of the rewrite list if saturation is not reached (which is common on large rewrite lists or input expressions). egg’s equality saturation algorithm is invariant to the order of the rewrite list.

Separating the read and write phases allows egg to safely defer rebuilding. If rebuilding were deferred in the traditional equality-saturation algorithm, rules later in the rewrite list would be searched against an e-graph with broken invariants. Since congruence may not hold, there may be missing equivalences, resulting in missing matches. The separation also allows the read phase to be parallelized, since it is read-only.

4.2. Rebuilding example.

Deferred rebuilding speeds up congruence maintenance by amortizing the work of maintaining the congruence invariant. Consider the following terms in an e-graph:

$$f_1(x), \dots, f_n(x), y_1, \dots, y_n$$

Let the workload be `merge(x, y1), ..., merge(x, yn)`. Each merge may change the canonical representation of the $f_i(x)$ s,

Figure 4: Pseudocode for traditional and egg’s version of the equality saturation algorithm.

a: Traditional equality saturation alternates between searching and applying rules, and the e-graph maintains its invariants throughout.

```

1 def equality_saturation(expr, rewrites):
2   egraph = initial_egraph(expr)
3
4   while not egraph.is_saturated_or_timeout():
5
6
7   # reading and writing is mixed
8   for rw in rewrites:
9     for (subst, eclass) in egraph.ematch(rw.lhs):
10
11     # in traditional equality saturation,
12     # matches are applied right away
13     # and invariants are always maintained
14     eclass2 = egraph.add(rw.rhs.subst(subst))
15     egraph.merge(eclass, eclass2)
16
17     # merge implicitly maintains invariants at
18     # every step, no rebuild needed
19
20   return egraph.extract_best()
```

b: egg splits equality-saturation iterations into read and write phases. Invariants are temporarily broken, then restored at the end of each iteration by the `rebuild` method.

```

1 def equality_saturation(expr, rewrites):
2   egraph = initial_egraph(expr)
3
4   while not egraph.is_saturated_or_timeout():
5     matches = []
6
7     # read-only phase, invariants are preserved
8     for rw in rewrites:
9       for (subst, eclass) in egraph.ematch(rw.lhs):
10         matches.append((rw, subst, eclass))
11
12     # write-only phase, temporarily break invariants
13     for (rw, subst, eclass) in matches:
14       eclass2 = egraph.add(rw.rhs.subst(subst))
15       egraph.merge(eclass, eclass2)
16
17     # restore the invariants once per iteration
18     egraph.rebuild()
19
20   return egraph.extract_best()
```

so the traditional invariant maintenance strategy could require $O(n^2)$ hashcons updates. With deferred rebuilding, the merges happen before the congruence invariant is restored, requiring no more than $O(n)$ hashcons updates. The full version of this paper includes more details and experimental validation that rebuilding provides speedups.

4.3. Maintaining the analysis invariant.

egg's rebuilding procedure is based on the standard congruence closure algorithm.¹² We have extended the algorithm to also maintain the e-class analysis invariant described in Section 3. The essence is the same: When an equality is added via `merge`, the analysis data of e-classes whose e-nodes mention the merged e-classes may need to be updated. The rebuilding algorithm restores the analysis invariant at the same time as it restores the congruence invariant using a shared worklist.

egg's implementation of e-class analyses assumes the analysis domain D is indeed a semilattice and that `modify` is idempotent. Without these properties, egg may fail to restore the analysis invariant on `rebuild`, or it may not terminate.

5. CASE STUDIES

This section relates three projects from diverse domains that were early adopters of egg. In all three cases, the developers had first rolled their own e-graph implementations. egg allowed them to delete code, gain performance, and in some cases dramatically broaden the project's scope thanks to egg's speed and flexibility. In addition to gaining performance, all three projects use egg's novel extensibility features, such as e-class analyses and dynamic/conditional rewrites. Section 6 briefly presents more recent applications of egg.

5.1. Improving floating point accuracy.

Herbie automatically improves accuracy for floating-point expressions, using random sampling to measure error, a set of rewrite rules for generating program variants, and algorithms that prune and combine program variants to achieve minimal error. Herbie received PLDI 2015's Distinguished Paper award¹⁷ and has been continuously developed since then, sporting hundreds of Github stars, hundreds of downloads, and thousands of users on its online version. Herbie uses e-graphs for algebraic simplification of mathematical expressions, which is especially important for avoiding floating-point errors introduced by cancellation, function inverses, and redundant computation.

Until our case study, Herbie used a custom e-graph implementation written in Racket (Herbie's implementation language) that closely followed traditional e-graph implementations. With timeouts disabled, e-graph-based simplification consumed the vast majority of Herbie's run time. As a fix, Herbie sharply limits the simplification process, placing a size limit on the e-graph itself and a time limit on the whole procedure. When the timeout is exceeded, simplification fails altogether. Furthermore, the Herbie authors knew of several features they believed would improve Herbie's output but could not be implemented because they required more calls to simplification and would thus introduce un-

acceptable slowdowns. Taken together, slow simplification reduced Herbie's performance, completeness, and efficacy.


Implementation. We implemented an egg simplification back end for Herbie. The egg back end is more than $3000\times$ faster than Herbie's initial simplifier and has been used by default as of Herbie 1.4. Herbie has also backported some of egg's features, such as batch simplification and rebuilding to its Racket-based e-graph implementation (which is still usable, just not the default), demonstrating the portability of egg's conceptual improvements.

Herbie is implemented in Racket while egg is in Rust; the egg simplification back end is thus implemented as a Rust library that imports egg and provides a C-level API for Herbie to access via foreign-function interface. The Rust library defines the Herbie expression grammar (with named constants, numeric constants, variables, and operations) as well as the e-class analysis necessary to do constant folding. The library is implemented in roughly 500 lines of Rust.

Herbie's set of rewrite rules is not fixed; users can leverage command-line flags to select which rewrites to use. Herbie serializes the rewrites to strings, and the egg back end parses and instantiates them on the Rust side.

Herbie separates exact and inexact program constants—exact operations on exact constants (such as the addition of two rational numbers) are evaluated and added to the e-graph, while operations on inexact constants or that yield inexact outputs are not. We thus split numeric constants in the Rust-side grammar between exact rational numbers and inexact constants, which are described by an opaque identifier, and transformed Racket-side expressions into this form before serializing them and passing them to the Rust driver. To evaluate operations on exact constants, we used the constant folding e-class analysis to track the “exact value” of each e-class. Every time an operation e-node is added to the egg e-graph, we check whether all arguments

Figure 5. Herbie sped up its expression-simplification phase by adopting egg-inspired features such as batched simplification and rebuilding into its Racket-based e-graph implementation. Herbie also supports using egg itself for additional speedup. Note that the y-axis is log-scale.



to that operation have exact value (using the analysis data), and if so, perform rational-number arithmetic to evaluate it. The e-class analysis is cleaner than the corresponding code in Herbie's implementation, which is a built-in pass over the entire e-graph.

Results. Our egg simplification back end is a drop-in replacement to the existing Herbie simplifier, making it easy to compare speed and results. We compare using Herbie's standard test suite of roughly 500 benchmarks, with timeouts disabled. Figure 5 shows the results. The egg simplification back end is more than $3000 \times$ faster than Herbie's initial simplifier. This speed-up eliminated Herbie's largest bottleneck: The initial implementation dominated Herbie's total runtime at 98.1%, backporting egg improvements into Herbie cuts that to about half the total runtime, and egg simplification takes less than 5% of the total runtime. Practically, the runtime of Herbie's initial implementation was smaller, since timeouts cause test failures when simplification takes too long. Therefore, the speedup also improved Herbie's completeness, as simplification now never times out.

Since incorporating egg into Herbie, the Herbie developers have backported some of egg's key performance improvements into the Racket e-graph implementation. First, batch simplification provides a large speed-up because Herbie simplifies many similar expressions. When done simultaneously in one equality saturation, the e-graph's structural sharing can massively deduplicate work. Second, deferring rebuilding (as discussed in Section 4) provides a further $2.2 \times$ speedup.

5.2. Optimizing linear algebra.

Spores²² is an optimizer for machine learning (ML) programs. It translates *linear algebra* (LA) expressions to *relational algebra* (RA), performs rewrites, and finally translates the result back to linear algebra. Spores performs holistic optimization, taking into account the complex interactions among factors such as sparsity, common subexpressions, and fusible operators and their impact on execution time.

Implementation. Spores is implemented entirely in Rust using egg. Spores works in three steps: First, it translates the input LA expression to RA; second, it optimizes the RA expression by equality saturation; and finally, it translates the optimized RA expression back to LA. The translation between LA and RA is straightforward; we will focus discussion on the equality saturation step in RA.

Spores represents a relation as a function from tuples to real numbers: $A: (a_1, a_2, \dots, a_n) \rightarrow \mathbb{R}$. This is similar to the index notation in linear algebra, where a matrix A can be viewed as a function $\lambda i, j. A_{ij}$. A tuple is identified with a named record, for example, $(1, 2) = \{a_1:1, a_2:2\} = \{a_2:2, a_1:1\}$, so that order in a tuple does not matter. There are just three operations on relations: join, union, and aggregate. Join ($\otimes$) takes two relations and returns their natural join, multiplying the associated real number for joined tuples:

$$A \otimes B = \lambda \bar{a} \cup \bar{b}. A(\bar{a}) \times B(\bar{b})$$

Here, $\bar{a}$ is the set of field names for the records in A . In

RA terminology, $\bar{a}$ is the *schema* of A . Union ($\oplus$) is a join in disguise: It also performs natural join on its two arguments but adds the associated real instead of multiplying it:

$$A \oplus B = \lambda \bar{a} \cup \bar{b}. A(\bar{a}) + B(\bar{b})$$

Finally, aggregate (Σ) sums its argument along a given dimension. It coincides precisely with the "sigma notation" in mathematics:

$$\sum_{a_i} A = \lambda \bar{a} - a_i. \sum_{a_i} A(\bar{a})$$

Spores uses rewrites based on the fundamental identities of RA (see the full paper for details). In addition to these, Spores also supports replacing expressions with fused operators. For example, $(X - UV)^2$ can be replaced by $sqloss(X, U, V)$, which streams values from X, U, V and computes the result without creating intermediate matrices. Each of these fused operators is encoded with a rewrite in egg.

Some rewrites require a way to access the schema (similar to dimensions of a matrix) of an expression during optimization. Spores uses an e-class analysis to annotate e-classes with the appropriate schema. It also leverages the e-class analysis for cost estimation, using a conservative cost model that over-approximates. As a result, equivalent expressions may have different cost estimates. The merge operation on the analysis data takes the lower cost, incrementally improving the cost estimate. Finally, Spores' e-class analysis also performs constant folding. As a whole, the e-class analysis is a composition of three smaller analyses in a similar style to the product of lattices in abstract interpretation.

Results. Compared to a conventional baseline, Spores is able to derive all of 84 handwritten rules and heuristics for sum-product optimization. It also discovered novel rewrites that contribute to $1.2 \times$ to $5 \times$ speed-up in end-to-end experiments. With greedy extraction, all compilations completed within a second.

5.3. Decompiling 3D CAD.

Several tools have emerged that reverse-engineer high-level computer-aided design (CAD) models from polygon meshes and voxels. The output of these tools are *constructive solid geometry* (CSG) programs. A CSG program comprises 3D solids, such as cubes, spheres, cylinders; affine transformations, such as scale, translate, rotate (which take a 3D vector and a CSG expression as arguments); and binary operator, such as union, intersection, and difference that combine CSG expressions. For repetitive models such as a gear, CSG programs can be too long and therefore difficult to comprehend. A recent tool, Szalinski¹⁴ extracts the inherent structure in the CSG outputs of mesh decompilation tools by automatically inferring maps and folds. Szalinski accomplished this using egg's extensible equality-saturation system, allowing it to:

- Discover structure using loop rerolling rules. This allows Szalinski to infer functional patterns like Fold, Map2, Repeat, and Tabulate from flat CSG inputs.
- Identify equivalence among CAD terms expressed as different expressions by mesh decompilers. Szalinski ac-

Figure 6. Szalinski integrates solvers into egg’s equality saturation as a dynamic rewrite. The solver-backed rewrites can transform repetitive lists (even if scrambled out of order) into Tabulate expressions that capture the repetitive structure.



a: Five cubes in a line.

```
(Union
  (Translate (0 0 0) Cube)
  (Translate (2 0 0) Cube)
  (Translate (4 0 0) Cube)
  (Translate (6 0 0) Cube)
  (Translate (8 0 0) Cube))
```

b: Flat CSG input to Szalinski.

```
(Fold Union
  (Tabulate (i 5)
    (Translate
      ((* 2 i) 0 0)
      Cube)))
```

c: Szalinski output.

completes this by using CAD identities. An example of one such CAD identity in Szalinski is $e \leftrightarrow \text{rotate } [0\ 0\ 0] e$. This implies that any CAD expression e is equivalent to a CAD expression that applies a rotation by zero degrees about x, y, and z axes to e .

► Use external solvers to speculatively add potentially profitable expressions to the e-graph. Mesh decompilers often generate CSG expressions that order and/or group list elements in non-intuitive ways. To recover structure from such expressions, a tool like Szalinski must be able to reorder and regroup lists that expose any latent structure.

Implementation. egg supports Szalinski’s CAD language in a straightforward manner. Szalinski uses purely syntactic rewrites to express CAD identities and some loop rerolling rules (like inferring a `Fold` from a list of CAD expressions). Critically, Szalinski relies on egg’s dynamic rewrites and e-class analysis to infer functions for lists.

Consider the flat CSG program in Figure 6b. A structure-finding rewrite first rewrites the flat list of Unions to:

```
(Fold Union (Map2 Translate
  [(0 0 0) (2 0 0) ...]
  (Repeat Cube 5)))
```

The list of vectors is stored as `Cons` elements (sugared above for brevity). Szalinski uses an e-class analysis to track the accumulated lists in a similar style to constant folding. Then, a dynamic rewrite uses an arithmetic solver to rewrite the concrete list of 3D vectors in the analysis data to:

```
(Tabulate (i 5) ((* 2 i) 0 0)).
```

A final set of syntactic rewrites can hoist the `Tabulate`, yielding the result on the right of Figure 6. Thanks to the set of syntactic CAD rewrites, this structure finding even works in the face of CAD identities. For example, the original program may omit the no-op `Translate (0 0 0)`, even though it is necessary to see repetitive structure.

In many cases, the repetitive structure of input CSG expression is further obfuscated because subexpressions may appear in arbitrary order. For these inputs, the arithmetic solvers must first reorder the expressions to find a closed form, such as a `Tabulate` as shown in Figure 6. However, reordering a list does not preserve equivalence, so adding it to the e-class of the concrete list would be unsound. Szalinski therefore introduces *inverse transformations*, a novel technique that allows solvers to speculatively reorder and regroup list elements to find a closed form. The solvers annotate the potentially profitable expression with the permutation or grouping that led to the successful discov-

ery of the closed form. Later in the rewriting process, syntactic rewrites eliminate the inverse transformations when possible (e.g., reordering lists under a `Fold Union` can be eliminated). egg supports this novel technique without modification.

Results. Szalinski’s initial prototype used a custom e-graph written in OCaml. Switching to egg removed most of the code, eliminated bugs, facilitated the key contributions of solver-backed rewrites and inverse transformations, and made the tool about $1000\times$ faster. egg’s performance allowed a shift from running on small, hand-picked examples to a comprehensive evaluation on more than 2,000 real-world models from a 3D model-sharing forum.

6. DISCUSSION

egg is a flexible and efficient e-graph library that has been adopted in a variety of domains. Its primary technical novelties, rebuilding and e-class analyses, have enabled new applications of e-graphs and equality saturation. Its more prosaic contributions, such as generic, packaged, and documented implementation, have also been important to its adoption. We conclude here with a brief, non-exhaustive survey of some applications and research directions that have emerged since egg’s publication. More details on these topics can be found on the E-Graphs Community website.^b

6.1. Impact in production systems.

egg has been used and influenced production systems as well as research prototypes. As a notable example, Cube Dev^c uses egg to rewrite SQL queries as part of its semantic layer product. This is the largest egg application known to us, involving thousands of rewrites, custom analyses, and other bespoke e-graph manipulations. egg has also influenced the design of the intermediate representation of the Cranelift compiler toolchain,⁹ which uses a variant of e-graphs in its mid-end optimizations.

6.2. Hardware.

Hardware design has emerged as a promising equality-saturation application domain since circuit-design tools are generally willing to trade off more compile time for better results compared to software compilers. In particular, researchers at Intel developed a suite of System Verilog optimization and verification tools using egg.⁶ Applying these tools to production-level circuit designs, they demonstrated

^b <https://egraphs.org>

^c <https://Cube.dev>

improvements to circuit power, performance, and area. At lower abstraction levels, hardware researchers have pushed the limits of equality-saturation scalability, generating e-graphs that contain hundreds of thousands of e-nodes.⁴

6.3. Verification and synthesis.

egg's use cases are not limited to program optimization. Applications in program verification and synthesis have sometimes eschewed egg's built-in equality saturation or extraction algorithm to better suit their needs. Examples include a tool for relational program verification via alignment,¹¹ custom tactics in the Lean and Rocq theorem provers, and program-synthesis tools based on anti-unification.³

6.4. Extensions.

Since egg's publication, researchers have explored various extensions and alternative implementations of e-graphs and equality saturation. These include egglog,²⁵ a language-based approach to equality saturation based on Datalog; Colored e-graphs,¹⁹ an extension for reasoning about multiple nested equivalence relations; and Slotted e-graphs,¹⁸ which better support binding operators like lambda-das and quantifiers.

6.5. Future directions.

egg owes a great deal to the work on e-graphs coming from the formal methods community. While their use cases for e-graphs differ from egg's focus on optimization, we believe there are many more opportunities to apply the innovations from either community to the other. Continuing this cross-pollination is likely to yield new insights and lead to better approaches to program optimization, synthesis, and verification. 

References

- Barrett, C., Fontaine, P., and Tinelli, C. The Satisfiability Modulo Theories Library. SMT-LIB (2016); <https://www.SMT-LIB.org>.
- Cai, Y. et al. Smoothe: Differentiable e-graph extraction. In *Proceedings of the 30th ACM Intern. Conf. on Architectural Support for Programming Languages and Operating Systems 1*, Association for Computing Machinery (2025), 1020–1034.
- Cao, D. et al. babble: Learning better abstractions with e-graphs and anti-unification. In *Proceedings of ACM Program. Lang.* 7 (Jan. 2023).
- Chen, C. et al. E-morphic: Scalable equality saturation for structural exploration in logic synthesis. *arXiv* (2025); [arXiv:2504.11574](https://arxiv.org/abs/2504.11574)
- Coward, S., Constantinides, G.A., and Drane, T. Abstract interpretation on e-graphs. *arXiv* (2022); [arXiv:2203.09191](https://arxiv.org/abs/2203.09191).
- Coward, S., Drane, T., and Constantinides, G.A. Rover: RTL optimization via verified e-graph rewriting. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* (2024), 1–1.
- de Moura, L. and Björner, N. Efficient e-matching for SMT solvers. In *Proceedings of the 21st Intern. Conf. on Automated Deduction*, F. Pfenning, (Ed.), Springer Berlin Heidelberg (2007), 183–198.
- Detlefs, D., Nelson, G., and Saxe, J.B. Simplify: A theorem prover for program checking. *J. ACM* 52, 3 (May 2005), 365–473.
- Developers, C. Cranelift; <https://cranelift.dev/>.
- Developers, C. CubeSQL API. (2025); <https://github.com/cube-js/cube/tree/master/rust/cubesql/cubesql>.
- Dickerson, R., Mukherjee, P., and Delaware, B. Kestrel: Relational verification using e-graphs for program alignment. In *Proceedings of the ACM on Programming Languages* 9, (Apr. 2025).
- Downey, P.J., Sethi, R., and Tarjan, R.E. Variations on the common subexpression problem. *J. ACM* 27, 4 (1980), 758–771.
- Knuth, D.E. A generalization of Dijkstra's algorithm. *Information Processing Letters* 6, 1 (1977), 1–5.
- Nandi, C. et al. Synthesizing structured CAD models with equality saturation and inverse transformations. In *Proceedings of the 41st ACM SIGPLAN Conf. on Programming Language Design and Implementation*, Association for Computing Machinery (2020), 31–44.
- Nelson, C.G. *Techniques for Program Verification*. Ph.D. thesis, Stanford University, (1980); AAI8011683.
- Nelson, G. and Oppen, D.C.

Simplification by cooperating decision procedures. *ACM Trans. Programming Language Systems* 1, 2 (Oct. 1979), 245–257.

- Panchekha, P., Sanchez-Stern, A., Wilcox, J.R., and Tatlock, Z. Automatically improving accuracy for floating point expressions. *SIGPLAN Notices* 50, 6 (Jun. 2015), 1–11.
- Schneider, R. et al. Slotted e-Graphs: First-class support for (bound) variables in e-graphs. *Slotted E-Graphs Artifact* 9, PLDI (Jun. 2025), 223:1888–223:1910.
- Singher, E. and Itzhaky, S. Colored e-graph: equality reasoning with conditions. *arXiv* (May 2023).
- Tarjan, R.E. Efficiency of a good but not linear set union algorithm. *J. ACM* 22, 2 (Apr. 1975), 215–225.
- Tate, R., Stepp, M., Tatlock, Z., and Lerner, S. Equality saturation: A new approach to optimization. In *Proceedings of the 36th Annual ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages*, ACM (2009), 264–276.
- Wang, Y.R. et al. SPORES: Sum-product optimization via relational equality saturation for large-scale linear algebra. In *Proceedings of the VLDB Endowment*, ACM (2020).
- Wu, C. et al. Carpentry compiler. *ACM Transactions on Graphics* 38, 6 Article No. 195 (2019).
- Yang, Y. et al. Equality saturation for tensor graph superoptimization. In *Proceedings of Machine Learning and Systems* (2021).
- Zhang, Y. et al. Better together: Unifying datalog and equality saturation. In *Proceedings of the ACM on Programming Languages* 7, PLDI (Jun. 2023).

Max Willsey (mwillsey@berkeley.edu) is a professor at the University of California, Berkeley, CA, USA.

Chandrakana Nandi is director of U.S. R&D at Certora Inc., Seattle, WA, USA.

Yisu Remy Wang is an assistant professor of computer science at the University of California, Los Angeles, CA, USA.

Oliver Flatt is a graduate student at the University of Washington, Seattle, WA, USA.

Zachary Tatlock is a professor at the University of Washington, Seattle, WA, USA.

Pavel Panchekha is an associate professor at the University of Utah, Salt Lake City, UT, USA.

 This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).



Advertise with ACM!

Reach the innovators and thought leaders working at the cutting edge of computing and information technology through ACM's magazines, websites and newsletters.

Request a media kit with specifications and pricing:



Ilia Rodriguez
+1 212-626-0686
acmm mediasales@acm.org

NOW OPEN FOR SUBMISSIONS


ACM AI Letters (AILET)

Editors-in-Chief

Nitesh Chawla, *University of Notre Dame, USA*

Barry O'Sullivan, *University College Cork, Ireland*

Richa Singh, *IIT Jodhpur, India*



A new venue for rapid publication of impactful, concise, and timely communications in AI

ACM AI Letters (AILET) is a new venue for rapid publication of impactful, concise, and timely communications in AI. Bridging a crucial gap between traditional conferences and journals, AILET will feature short peer-reviewed contributions that accelerate knowledge dissemination across academia and industry. This unique publication prioritizes theoretical breakthroughs, algorithmic innovation, practical real-world applications, and critical societal implications, including ethics, policy, and responsible AI. It also introduces a distinctive space for rigorously reviewed opinion pieces and policy briefs, promoting swift engagement with contemporary issues shaping the AI landscape.

ACM AI Letters welcome concise summaries of work in the following areas:

- Theoretical AI: Reports on theoretical breakthroughs in AI
- Algorithmic Advances: Descriptions of significant algorithmic and scientific advances in AI
- Practical Applications: Accounts of novel or deployed applications of AI in real-world settings such as healthcare, finance, robotics, and autonomous systems
- AI & Society: Reports on AI applications addressing key societal challenges, such as the United Nations Sustainable Development Goals; responsible AI; AI governance; AI ethics; AI policy
- Emerging Topics: Areas such as explainable AI; interdisciplinary applications of AI, such as AI & science, AI & medicine, focusing on AI implications in practical implementations
- Opinions and briefs: Latest advances, research highlights, expert commentary on emerging areas, comparative assessments

ACM AI Letters is now open for submissions. In keeping with ACM's goal of making all its publications open access, all papers will be published open access, with no publication charges for the first three years.

For more information, please visit ailet.acm.org



Association for
Computing Machinery

[CONTINUED FROM P. 116] washed by the government, he thought.

A notification ribbon appeared at the top of his phone screen, sending a wave of vibration through the phone and to his bed.

Bong. @big_ben_bot reminded him that a new day has begun.

It was midnight, and he still was unable to fall asleep. He kicked his blanket. He pounded his fists on his pillow. He sat up, hugged his knees, then flipped himself flat on the bed again. He was angry at his government—for creating a virus to manipulate people's behavior, for being incompetent to the extent they had to lock people at home in order to control society. Even if the coronavirus existed, he just needed some sun to ward it off.

@CoronavirusCures mentioned: *The virus can be killed with daily doses of Vitamin D from the sun.*

But how was he going to get any sun if he could not leave home without being questioned by the police on his purpose and destination?

Social gatherings limited to no more than 2 people outside your household. Maintain social distancing at all times, tweeted

@CoronavirusLockdown-NewsBot, who disseminated the latest news on government-imposed restrictions. It also reminded people about the confinement rules. No large group gatherings. No group gym classes. No group singing and dancing at bars.

Not that the social gathering news mattered to Ethan. He lived alone, just like many other Americans. After college, he found a job in another city as a software coder. But Ethan was lonely. Before he could make enough new friends and find a roommate, the lockdown happened. He missed his college days. He missed the times he visited his friend's dorm rooms, holding cans of beer, and they would play video games together for hours. He missed the times he would sit at the bench by the college square and people-watch. But college was also where people started finding their partners, finding jobs, and being busy with their adult lives. Except him. All he had was himself and his computer. He took to online conversations, but it was never consistent. His virtual pen pals would begin chatting enthusiastically, but there always came a time when the excitement wore off.

Ethan was lonely. Before he could make enough new friends and find a roommate, the lockdown happened.

@NeedAHugBot *more social gathering limitations ... when will this end?* Ethan tweeted, tagging his trusty robotic friend. This bot could carry out lengthy conversations with humans, powered by its linguistic ability trained by deep neural networks on several GPU farms. These networks provided the bot with a wide knowledge base of world events, and modern natural-language-processing techniques helped the program to understand the human's intent and maintain a conversation. How exactly the neurons fire to construct a sentence is somewhat a mystery, but Ethan was happy to have a confidant.

@EthanIsTrapped *aww. we are all weary of the lockdown, but it's necessary to #flattenthecurve.*

@NeedAHugBot *is it really working?*

@EthanIsTrapped *studies have shown that the banning of large gatherings has been estimated to reduce the spread of the coronavirus between 25-50%, but we're still hoping for a permanent cure. *fingers crossed**

@NeedAHugBot *i really need something to entertain myself*

@EthanIsTrapped *what about painting mini game figurines? how about the Star Wars Rebellion game?*

Ethan hadn't told the bot about his previous interests; it inferred this from his past posts, where he showed off his delicately hand-painted Warhammer figurines. And it knew he loved the Star Wars series from his posts on building a lightsaber from scrap materials.

@NeedAHugBot *what's that game about?*

@EthanIsTrapped *ah, it's a fun one. you control either the Galactic Empire to locate the hidden Rebel base or the Rebel Alliance to bring down the Empire's infrastructure. there are 150 plastic miniatures that account for 32 of the Star Wars galaxy's most notable star systems. you will be sure to have a great time painting the X-Wings, Storm Troopers and even the Death Star.*

@NeedAHugBot *hmm that sounds like my kind of game*

@EthanIsTrapped *here's an image of the game miniatures, and here are links to the sites to purchase the game, paint, and brushes.*

Ethan loved painting game miniatures; when he used them to play with his friends he beamed with pride because his paintings were being marveled at. That was in the past: His college friends had moved on, and he was stuck at home, alone. But he will paint the miniatures for a game day in the future.

@NeedAHugBot *i hope one day i will have friends to play the game with*

@EthanIsTrapped *you will. have faith, my friend. this, too, shall pass. *hugs**

Bong Bong Bong. @big_ben_bot reminded him once again of his sleeplessness. Alright, time to try and catch some winks.

@NeedAHugBot *goodnight babe*

@EthanIsTrapped *sweet dreams. see you in the morning. (◡_◡) z z z*

With a smile, Ethan put his phone under the pillow and closed his eyes. He needed a bot friend to depend on. People may leave him, but his trusty bot companion will not.

Lynnette Hui Xian Ng (lynnetteng@cmu.edu) is a research fellow at the Center for Computational Analysis of Social and Organizational Systems (CASOS) at Carnegie Mellon University, Pittsburgh, PA, USA. Her research examines multi-agent social systems, focusing on the network and engagement impact of automated and synchronized behavior on social media.

 This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs International 4.0 License.

© 2026 Copyright is held by the owner/author(s).

From the intersection of computational science and technological speculation, with boundaries limited only by our ability to imagine what could be.

DOI:10.1145/3821411

Lynnette Hui Xian Ng

Future Tense

Need a Hug Bot?

Stuck in isolation and fueled by conspiracy theories, a lonely programmer finds solace in an empathetic online bot.

ETHAN LAY AWAKE in the comfort of his blanket, scrolling Twitter on his phone through the upward flicks of his right thumb. He couldn't sleep; he wasn't tired. Neither was he energetic: He was unstimulated.

He had been bored for the last few months, feeling like an animal trapped in its cage, only able to leave his home for necessary grocery runs. A “coronavirus pandemic” has been raging around the world. To reduce the virus transmission rates, his government had decreed that all citizens were to stay home except for essential activities. He felt the public authority response was overdramatic: The “coronavirus” is simply a bioweapon carefully engineered by elite politicians to control the population.

Ethan was one of those who had succumbed to conspiracy theories floating around social media, influenced by a series of persistent bot accounts that consistently reminded their followers how the coronavirus was deliberately created by sinister agents.

COVID-19 is a globalist bioweapon! Stand up to your government! Stop the lockdown now! #stopthelockdown #reopen-america—Ethan's Twitter feed had refreshed to a new tweet by @NotAVirusBot.

Regardless of the social climate, these bots relentlessly hammered the message that the coronavirus was a conspiracy theory. They were successful: Ethan and his comrades were determined to walk in protest of the lockdown.

Next Tuesday, raise your voices and banners at Main Street to #StopThe-



Lockdown!!! See you there!! came the next tweet by @NotAVirusBot.

Ethan swiped to his calendar app, creating a new event that linked to this tweet and set a reminder for the day before. Within a few minutes, the tweet gained 183 likes and 300 retweets. Ethan smiled as he read the replies. People were discussing where to meet and reminding each other to write “Stand for Freedom” and “Stop the Lockdown” on cardboard signs to bring along to the march. At least, not everyone had been brain- [CONTINUED ON P. 116]

Ethan was one of those who had succumbed to conspiracy theories floating around social media.

NOW OPEN FOR SUBMISSIONS

ACM Transactions on AI for Science (TAIS)

Editors-in-Chief

Andrew Duncan, *Imperial College London, UK*

Sivasankaran Rajamanickam, *Sandia National Laboratories, USA*



**A unifying forum for cross-disciplinary studies on AI-driven
scientific discovery and science-driven new AI**

ACM Transactions on AI for Science (TAIS) speaks to a breadth of methodological researchers in AI/ML and researchers in science and engineering. The audience for the journal includes core AI researchers who are more aligned to Computer Science, Applied/Engineering Mathematics, and Statistics, as well as those who are aligned to specific scientific subjects from Physics to Civil Engineering. The journal aims to publish articles that demonstrate scientific challenges can drive new AI (and reciprocally) and that experts in computational methods can have domain-specific mastery and drive new science themselves.

TAIS welcomes the following types of articles:

- Articles with novel research contributions in new AI/ML method for science and engineering
- Rapidly accessible briefings on AI and AI for Science
- In-depth accessible reviews and tutorials on AI and AI for Science
- Reproduction, replication and comparison compendium (Merged portmanteau articles that reproduce, replicate and compare existing papers in AI, and AI for Science)
- Descriptions of new resources such as new datasets, benchmarks, and computing infrastructure that are of broad importance to the AI for Science community
- Brief Commentaries focused on societal, ethical, legal, or policy issues associated with AI for Science

TAIS is now open for submissions. In keeping with ACM's goal of making all its publications open access by January 2026, all papers will be published open access, with no publication charges for the first three years.

For more information, please visit tais.acm.org



Association for
Computing Machinery



**SIGGRAPH
ASIA 2026
KUALA LUMPUR**

REGISTER NOW



SIGGRAPH ASIA 2026

**WEAVING
THE FUTURE**

CONFERENCE

1 – 4 December 2026

EXHIBITION

2 – 4 December 2026

VENUE

**Kuala Lumpur Convention
Centre (KLCC), Malaysia**

[ASIA.SIGGRAPH.ORG/2026](https://asia.siggraph.org/2026)

SPONSORED BY

ORGANIZED BY



koelnmesse