Polynomial Times

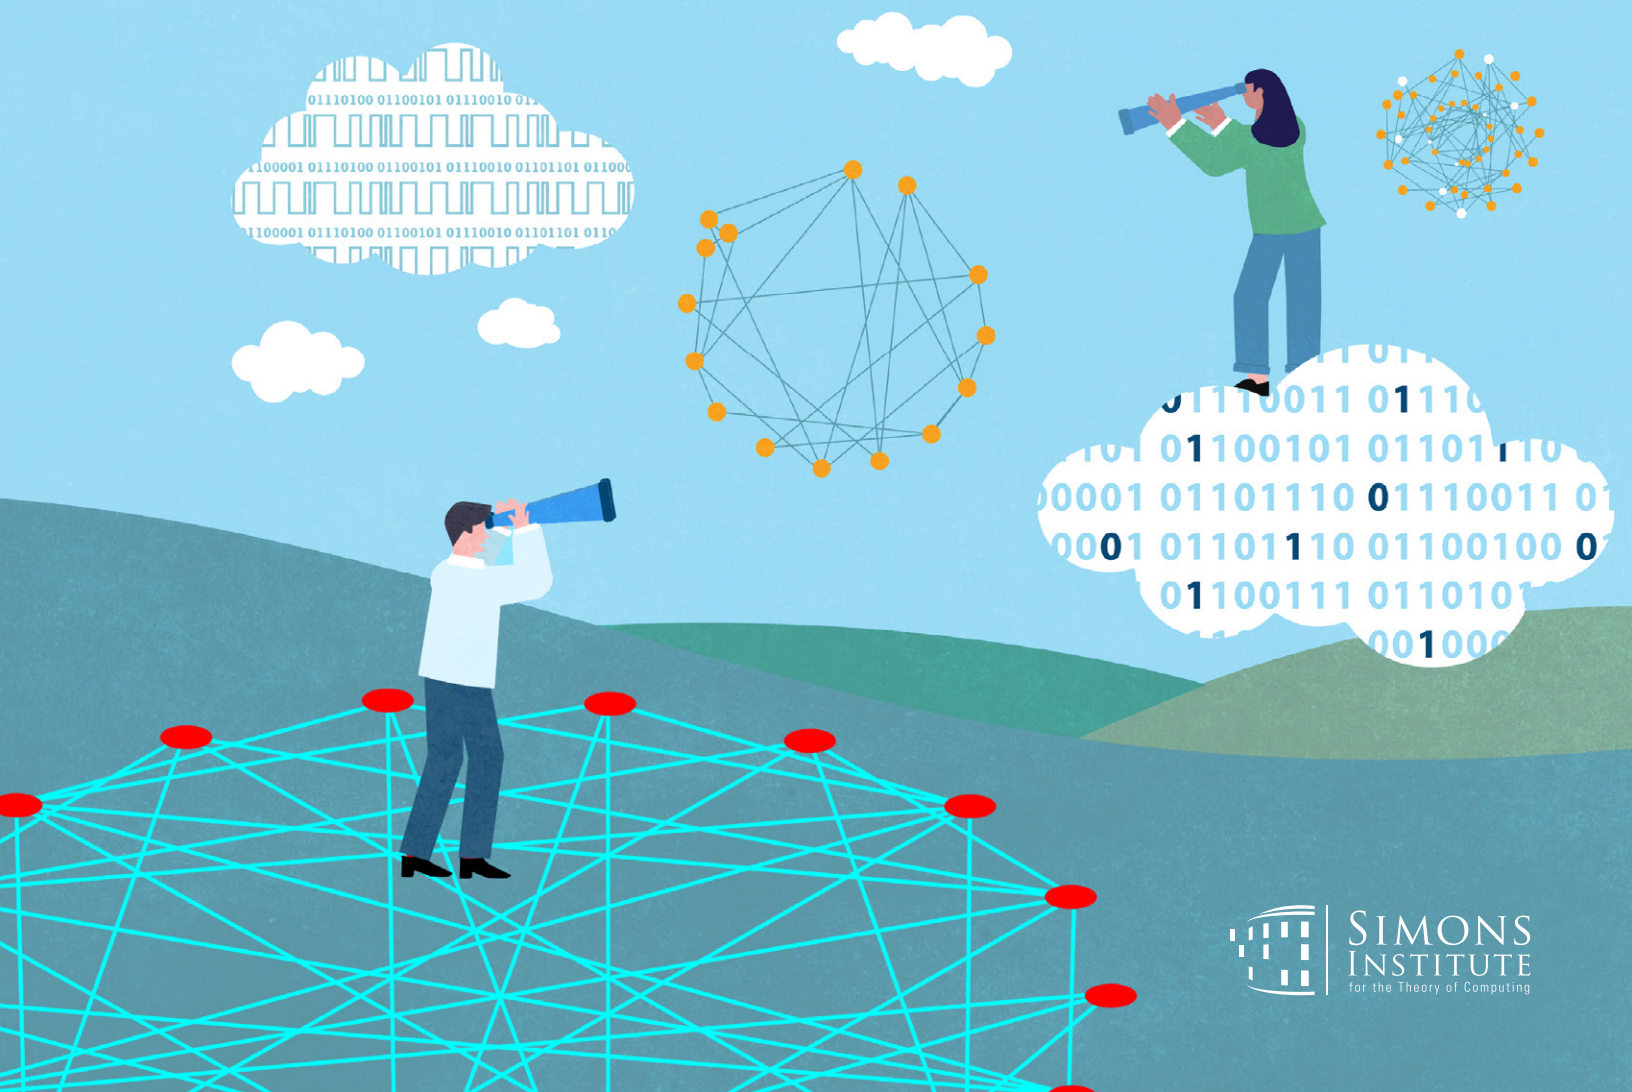
2026–27 | Issue 2

Theoretical Computer Science Meets Its AI Moment

A Conversation About Algorithmic Fairness

Fault-Tolerant Quantum Computing with Neutral Atoms

The Proof Behind the Proof



Simons Institute for the Theory of Computing

Polynomial Times

The annual magazine of
the Simons Institute for the Theory of Computing

2026–27 | Issue 2

Polynomial Times **Editorial Team**
Venkatesan Guruswami (Scientific Editor),
Kristin Kane (Editor), Preeti Aroon (Copy Editor),
Barry Bödeker (Graphic Designer)

Simons Institute for the Theory of Computing
Venkatesan Guruswami
Director

Nikhil Srivastava
Interim Associate Director

Jason Lee
Senior Scientist

Peter Bartlett
Research Director, Machine Learning

Shafi Goldwasser
Research Director, Resilience

Umesh Vazirani
Research Director, Quantum Computing

Amy Ambrose
Senior Development Director & Director of Operations

Kristin Kane
Director of Communications and Finance

Anil Ananthaswamy
Science Communicator at Large

Preeti Aroon
Copy Editor

Barry Bödeker
Graphic Designer

Natalie Chen
Research Administrator

Omied Far
Videographer

Ashley Hasson
Event Coordinator

Robin Hoey
Facilities Manager

Frances Klingenberger
Event Coordinator

Kelsey Larsen
Financial Specialist

Leisa Lei
HR Generalist

Drew Mason
Information Systems Analyst

Joe Ratzlaff
Information Systems Analyst

Atiya Wahab
Visitor Services Coordinator

Carolyn Winter
Scientific Program Coordinator

Carson Fox Young
Executive Assistant

Senior Advisors
Jelani Nelson
Bin Yu

Scientific Advisory Board (2026)
Bonnie Berger
Paul Christiano
Vincent Cohen-Addad
James Demmel
Uriel Feige
Parikshit Gopalan (Co-Chair)
Martin Grohe
Monika Henzinger
Maurice Herlihy
Yael Tauman Kalai
Amin Karbasi
Robert Kleinberg
Orna Kupferman
Thomas Rothvoss
Nathan Srebro
Rachel Ward
Ryan Williams
Mary Wootters (Co-Chair)



Inside

Letter from the Director 1

Theoretical Computer Science Meets Its AI Moment 3
Anil Ananthaswamy

Algorithms, Society, and the Law 9

A Conversation About Algorithmic Fairness 11
Sampath Kannan, Jon Kleinberg, Omer Reingold, and Rakesh Vohra

Proving Post-Quantum Security, in and out of Calvin Lab:
A Portrait of Nick Spooner 18

Rapid Developments Toward Fault-Tolerant Quantum Computing with Neutral Atoms 21
Madelyn Cain and Harry Levine

The Proof Behind the Proof: New Routes to SNARGs for NP 27
Yael Tauman Kalai and Abhishek Jain

Researcher, Expositor, Friend: A Portrait of Dan Gusfield 34
Kristin Kane and Tandy Warnow

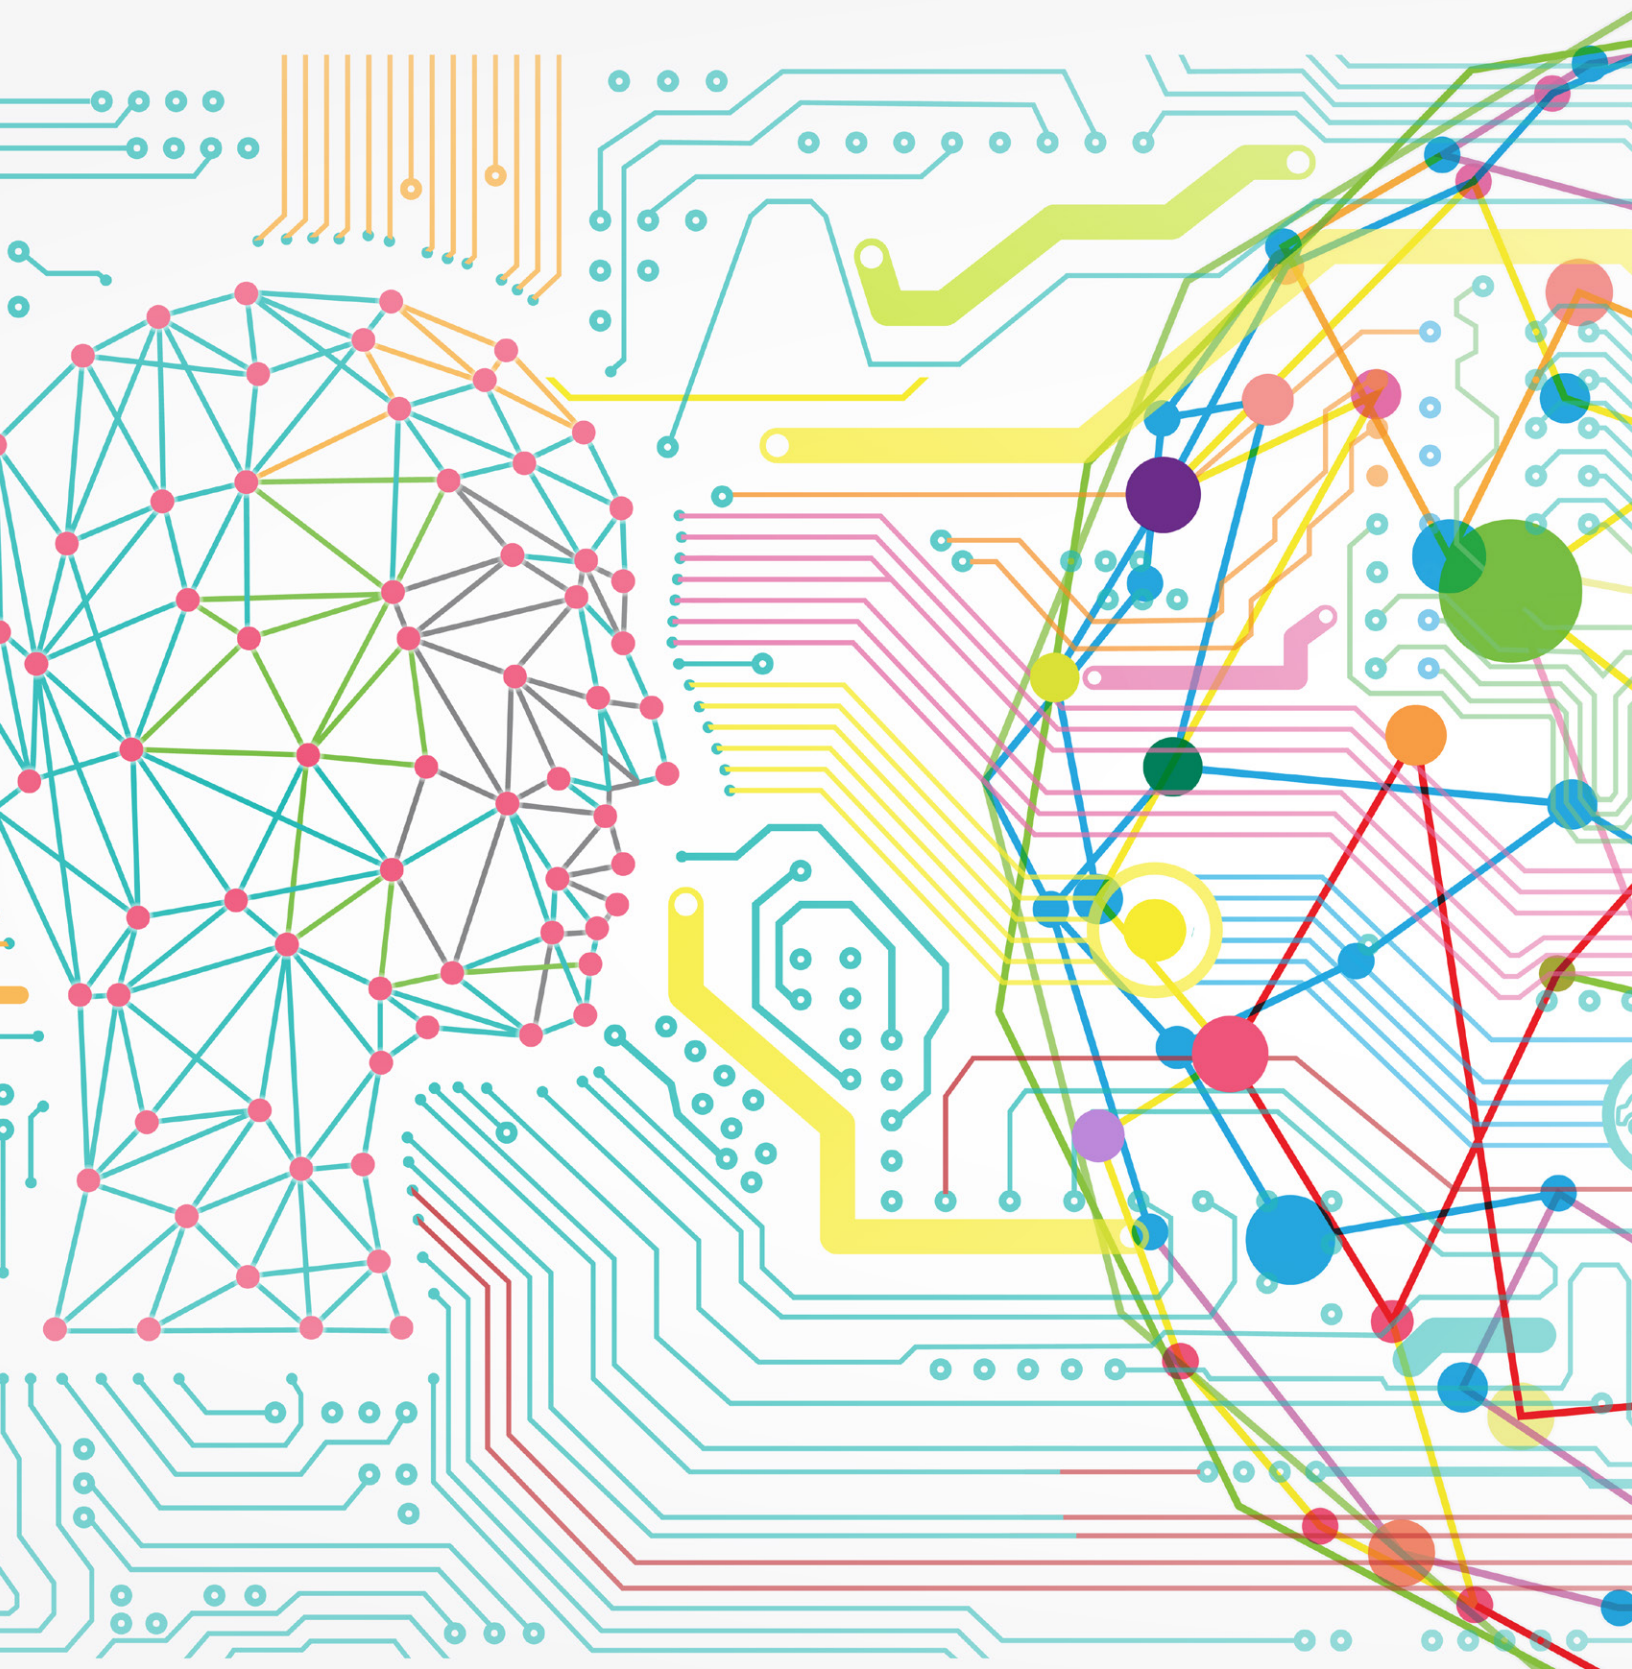
This Year's Research Programs 36
Jason Lee

Current & Upcoming Programs 38

Support for the Simons Institute 39



The Simons Institute's YouTube channel is the world's largest repository of video content on theoretical computer science. The collection includes almost every research talk that has been given at the Institute since we opened our doors in 2013, along with tutorials, interviews, public lectures, panel discussions, and short documentaries. The channel has 6.4K+ videos and 77K+ subscribers, and has had 9.5+ million views.





Dear friends,

I'm delighted to share with you the second issue of *Polynomial Times*, the annual magazine of the Simons Institute for the Theory of Computing. In it, we explore the frontiers of theoretical computer science (TCS) research, at the Simons Institute and beyond.

This year's issue focuses on topics that have been at the forefront of all of our minds, most notably artificial intelligence, and the ways in which it has already begun to transform research in our field and the boundaries of what we thought was within reach, while also raising significant concerns around its safe and trustworthy deployment. With the support of a substantial new grant from Renaissance Philanthropy and XTX Markets, this year we will be launching a Research Pod on AI and Theoretical Computer Science, which will foster a sustained effort to both push the frontiers of the field with AI-assisted exploration, and help address the pressures and challenges posed by AI to the worldwide TCS community. In this issue of *Polynomial Times*, we report on a working group hosted by the Simons Institute that developed a set of concrete near-term recommendations for how our research community should respond to the increasing power and influence of AI.

Elsewhere in the magazine, we explore the Institute's work at the intersections of algorithms and society. We also look at exciting recent developments in cryptography and quantum computing, two topics of ongoing interest at the Institute that were the subject of research programs this past year. And we share stories of people whose lives and careers are interwoven with the Institute.

I encourage you to connect with us on socials, watch our content on YouTube, or join us in Calvin Lab for a research program or workshop. Hope to see you soon!

Best wishes,

Venkat



Venkatesan Guruswami
Director, Simons Institute for the Theory of Computing



Photo: Bruce Damonte

Theoretical Computer Science Meets Its AI Moment

by Anil Ananthaswamy

On October 26, 2010, someone — most likely a worried undergrad — posted a query on MathOverflow with this headline: "What's a mathematician to do?" The query continued: "What can one (such as myself) contribute to mathematics? … It seems entirely plausible that, with all the tremendously clever people working so hard on mathematics, there is nothing left for someone such as myself (who would be the first to admit they do not have any special talent in the field) to do."

Four days later, William Thurston, a Fields Medal winner and Professor of Mathematics at Cornell, answered the student, arguing that the study of mathematics is a communal exercise in developing human insight: "The product of mathematics is clarity and understanding. Not theorems, by themselves. … Mathematics only exists in a living community of mathematicians that spreads understanding and [breathes] life into ideas both old and new. The real satisfaction from mathematics is in learning from others and sharing with others. All of us have [a] clear understanding of a few things and murky concepts of many more. There is no way to run out of ideas in need of clarification," he wrote.

Thurston passed away in August 2012. But his response remains relevant today, and more so, as advances in artificial intelligence are forcing even seasoned researchers to ask: "What's a mathematician to do?"

This was especially true on September 8, when OpenAI announced that an internal system (more powerful than its then-newest model, GPT-6 Astra) had found "a solution to the Navier–Stokes existence and smoothness problem, one of the Millennium Prize Problems."[1] As if on cue, a group of 32 researchers gathered a day later at the Simons Institute for the Theory of Computing at UC Berkeley for a closed-door workshop to discuss the immediate future of theoretical computer science (TCS) in light of the challenges posed by AI — except that the organizers had been planning the two-day workshop for months. The latest OpenAI announcement served only to stoke the fire that had already been lit.

Over the next few weeks, a more tangled narrative of the extent of human influence on OpenAI's Navier–Stokes result would emerge, causing more controversy.[2] But even before Navier–Stokes, which involved a yet-to-be-released internal model, results from using publicly available models on other problems were challenging notions of what it means to do math. Pravesh Kothari, Associate Professor of Computer Science and Applied Mathematics at Princeton, and a workshop participant, told *Polynomial Times*: "So many conjectures that people have spent years, sometimes decades, working on are now resolved by AI. The Matrix Spencer conjecture that many of us had thought about over time was resolved. The Komlós conjecture — again, so many of us thought about it for decades — was resolved, and not even by a theoretical computer scientist using AI. Somebody from biostatistics was able to one-shot it."

Even as mathematicians work on verifying these AI-generated solutions, the concern is clear. "There are a lot of people talking about how AI is changing our field. But there's no mechanism to turn that into changes that need to happen on a very short timescale. The workshop was an attempt to gather opinions and ideas from a sample of the TCS community and come up with some actions that this group thinks are beneficial in the short term," said Nikhil Srivastava, Interim Associate Director of the Simons Institute. "And the hope is that these will then be executed by people in the group, and people outside the group, within the next six months."

[1] OpenAI, "On the Navier–Stokes Millennium Prize Problem" (2026).

[2] See, for example, Celina Zhao et al., "How an AI Math Breakthrough Ignited a Controversy," *Science* (2026).



During introductions, each participant spoke about the extent of AI usage in their research — and it ranged anywhere from a few hours a week to tens of hours, with AI agents running overnight while researchers slept. Thore Husfeldt, Professor of Computer Science at the IT University of Copenhagen, later told *Polynomial Times* that he was "moderately surprised" that everyone at the workshop was already using AI. "It was, of course, a self-selective crowd," said Husfeldt. But "we might have expected there to be lots of people who, for good reasons, would just reject AI on principled grounds, and tie themselves to the mast like Ulysses." That wasn't the case.

The organizers — Srivastava, Amit Sahai (Professor of Mathematics and Computer Science, UCLA), Tony Feng (Associate Professor of Mathematics, UC Berkeley), and Zeph Landau (Research Scientist in Computer Science, UC Berkeley) — introduced the issues, laid down some ground rules, and voiced some of the concerns simmering in the TCS community.

Sahai acknowledged the "elephant in the room": AI's outsized effect on mathematics research. Juxtaposed against AI's aforementioned success at proving certain important results that had resisted human effort for decades, Sahai emphasized the purpose of the field of theoretical computer science (his words had echoes of some of Thurston's views): to develop a deep human understanding of the relevant mathematical structures, questions, and results, and to train students to keep contributing to the field. He talked of the participants, partly in jest, partly seriously, as being the Council of Elrond that had received the ring. "What do we do with the ring?" asked Sahai.

What indeed?

Answering that question in the space of two days would require tremendous resolve. To help keep the meeting on track, the organizers brought in Paul Sullivan, a professional facilitator who used a steady hand to steer participants toward finding solutions rather than engaging in open-ended debate. It was clear that the "default setting" for such meetings — indulging in discussions and speculations about the future of AI, math, and maybe even civilization — wouldn't work. "The default setting is failure!" warned the organizers.

And failure wasn't an option. Landau, in his brief remarks, spoke of a survey carried out before the meeting, asking theoretical computer scientists associated with the Simons Institute to come up with questions, concerns, problems, and/or ideas for solutions that they felt were important to address in the next six months. The organizers received 134 responses, a few tinged with emotion and anxiety. Some wrote back to say that the current state of AI-infused math research wasn't what they had imagined doing math would be like when they began studying the subject. Others opined that they were glad the Simons Institute was meeting the challenge head-on, because many had been waiting for just such a meeting.

The responses helped the organizers firm up the questions to be addressed by the workshop attendees: How should attribution and disclosure be handled for AI-assisted research? How should TCS conferences shift, in submission, reviewing, and dissemination? How should PhD education and training change? How should the hiring process change?

The participants split into four groups, each tackling one set of related questions. Sullivan and the organizers enforced a format that required each group to first discuss relevant issues among themselves, then present their findings to, and take feedback from, the larger group, and then rinse and repeat. (Sullivan used the "Fist to Five" technique during the main group sessions to arrive at a consensus — responses to a suggested measure ranged from someone raising their fist to express extreme displeasure/complete opposition, to displaying all five fingers as a sign of total support.) Over the course of two days, the process resulted in each group identifying a set of problems that needed urgent attention and arriving at a set of recommended actions.

"I thought that there would be some structure to how these discussions happen and that there would be a synthesis at the end of it, but I didn't expect such a well-planned format for making it happen," said Kothari.

In their discussions, participants tackled many hard questions, including but not limited to these: What is the human intellectual contribution in research done using AI? How do we incentivize good research in the face of AI's ability to easily solve certain classes of



September 2026 Working Group

problems? Should there be penalties for antisocial behavior, such as publishing a badly written AI-aided paper or scooping someone's hard work by using AI? How should expositions of one-shotted results be treated? How can one foster old-fashioned collaborative thinking without the use of AI? What can researchers do to alleviate the uncertainty that PhD students are feeling about their job prospects? How do departments hire new faculty during the AI phase transition? How can the community address the incentive to lie about the extent of AI usage in someone's work?

The discussions, both within the small groups and among the larger group, surfaced a few anxieties. Some felt that it was depressing that someone with no expertise in an area could solve a problem using AI — could this lead to AI being used as an excuse to cut funding for science? Others averred that research had become less fun, now that AI is everywhere.

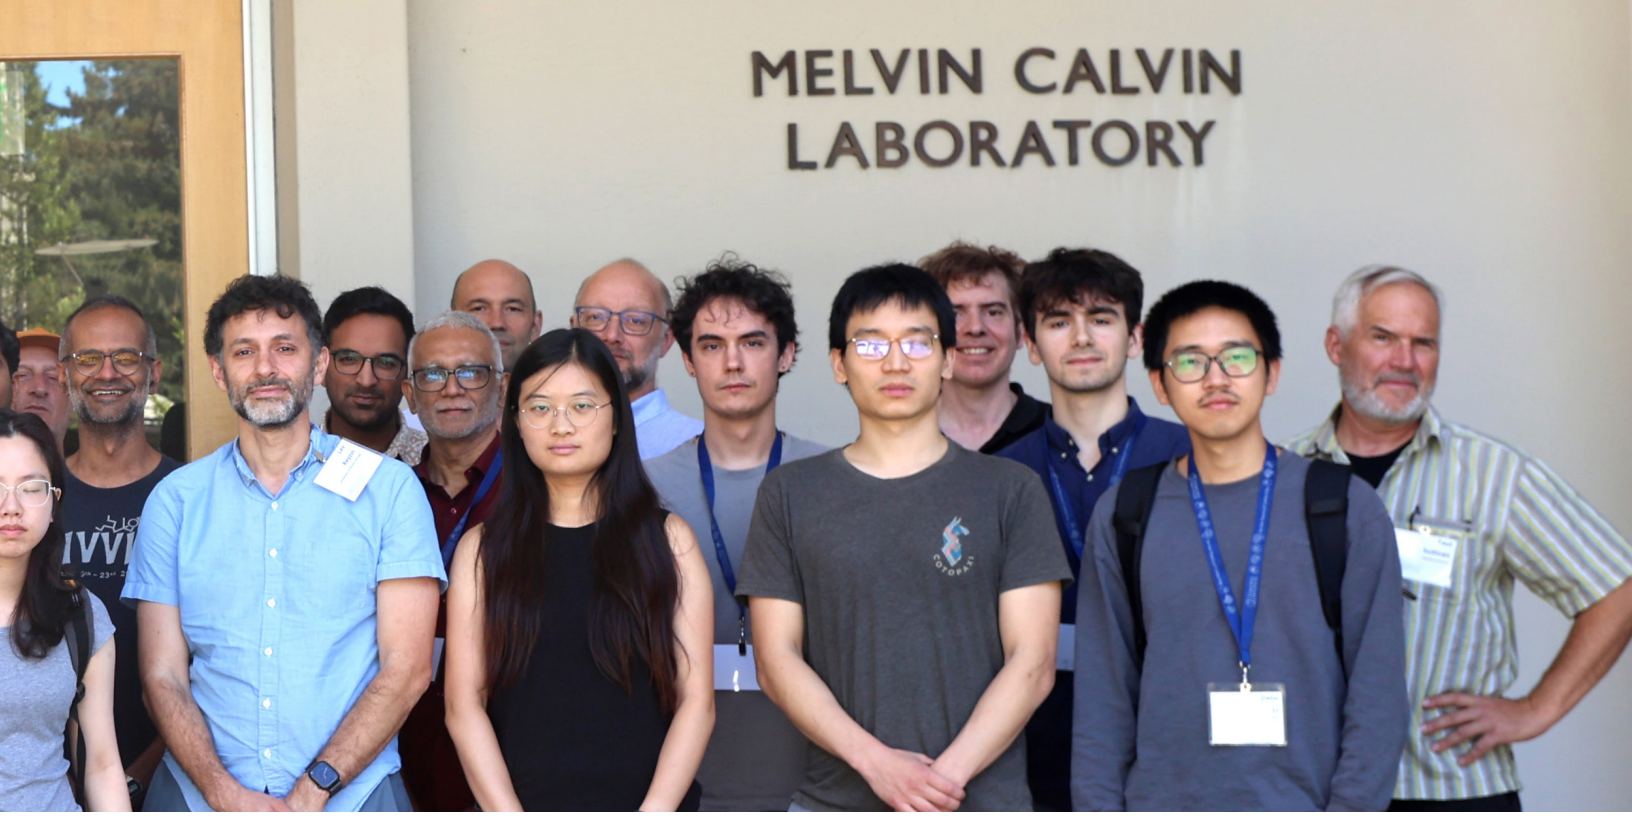
And more than once, participants brought up Terence Tao's recent comparison of AI companies' onslaught on math problems to certain environmentally degrading mining practices. "The indiscriminate automated strip-mining of open problems for solutions may destroy the ecosystem from which the next generation of mathematical techniques, problems, and practitioners would have developed, similarly to how using excavators to dig out treasures from an archeological site destroys the rich historical context and data that gives such treasures so much of their meaning and value," wrote Tao.[3]

However, it was hardly all doom and gloom. Among many participants, there was clear appreciation of the inevitability, even merits, of the integration of AI into research, and the fostering of new skills. For example, AI has greatly increased the number of theoretical results, demanding explanations and expositions from theorists. Some called for not stigmatizing papers resulting from one-shot prompting. Others argued for rejecting the idea that theory is so much more vulnerable to AI than other fields. Instead, they said, the TCS community should develop theories for AI-related problems that have real-world implications, such as alignment.

Despite the range of questions and concerns, the workshop participants managed to arrive at a broad consensus, in no small part due to Sullivan's engaging but firm enforcement of the rules of play. It helped that everyone could use only a notebook and a pen, with laptops disallowed until the final afternoon, when all the deliberations had to be turned into sets of recommendations.

"I've been to many workshops, and normally get really annoyed as soon as there are Post-it notes and meetings in small groups and large groups," said

[3] Terence Tao, "As Hugo Duminil-Copin wrote recently," Mathstodon (September 2, 2026).



Husfeldt. “But for once, this actually worked. I’ve been rhapsodizing about this to colleagues.”

Kothari agreed. “I was positively surprised by 1) how thoughtful and engaged everyone was, and 2) how prepared everyone seemed to be to make difficult decisions, even if they were not a hundred percent on board with it. The organizers were very clever in making sure that we worked on actionable items as opposed to philosophical questions.”

Within three weeks of the workshop, the working group published a set of 12 concrete, actionable recommendations for the TCS community, ranging from how to attribute the contributions of humans and AI in publicly posted research papers (for example, by clearly delineating AI’s influence on the formulation of the paper’s main questions, important technical results, etc.) to a set of rules and norms that conferences should follow in their call for papers (Is AI use permitted? Is there a limit on the number of submissions?, etc.) to actions that would help reinforce deep human understanding of the subject matter (including experimenting with including video talks as supplemental materials in conference paper submissions, and mandating that PhD students in theoretical computer science be trained and evaluated in part on research where the use of AI is forbidden).[4]

The success of the workshop notwithstanding, the larger philosophical concerns linger. “Unfortunately, those are also the ones that matter more,” said Kothari. For example, the question of what constitutes good math research in the age of AI hangs like the sword of Damocles over the TCS community. “Hopefully I’m wrong, but there is a good chance that mathematical problem-solving is essentially automatable. If so, working hard on well-posed open questions may become very different or, in the extreme case, even outdated. There are a lot of difficult questions about what important work will look like.”

Husfeldt, having returned to Europe, also spoke philosophically about reinforcing the importance of the human component of their research. “Computing now is clearly one of the most consequential sciences that we have invented as a civilization. It’s also clear that mathematics will continue to be the language in which science is expressed. Therefore, it’s super important that humans continue to use math to explain computation to each other. That is the mission call that we need to heed. We are exactly where we want to be, and finally we’re doing something of consequence, and we should just embrace that.”

Thurston would have approved.

[4] September 2026 Working Group, “AI and TCS: The Next Six Months,” Simons Institute for the Theory of Computing (2026).

“Established in 2020, the Simons Institute’s Law and Society Fellowships enhance Institute programs that address technologies with profound impacts on human society and with implications for ethics, law, and policy, by supporting researchers focused on the broader societal implications of the techniques and technologies studied in the programs.”



Algorithms, Society, and the Law

In Spring 2026, as part of our program on Federated and Collaborative Learning, the Simons Institute hosted two Law and Society Fellows: Greg Demirchyan (Fairlogic), a practicing attorney working on AI, fairness, and data privacy, and Rui-Jie Yew (Brown University), a graduate student in computer science studying the legal and human-centered aspects of algorithmic technologies. The two worked together on a range of exciting projects highlighting the real-world implications of the research that took place in the program, including a one-day workshop on AI governance that convened a group of policymakers, science policy advocates, and researchers from academia and industry to discuss the landscape of AI regulation. The workshop featured keynotes by California State Senator Jerry McNerney and former Biden Administration Tech Policy Advisor Suresh Venkatasubramanian, now the Dean's Professor of Computer Science and Data Science at Brown. Rui-Jie and Greg also collaborated on a paper on red-teaming the law, enhanced by conversations with scientists in the program. This project explores an adversarial red-teaming framework for legal alignment, modeling regulatory design as a multiagent zero-sum game between a loophole exploiter and a legal patcher.

Established in 2020, the Simons Institute's Law and Society Fellowships enhance Institute programs that address technologies with profound impacts on human society and with implications for ethics, law, and policy, by supporting researchers focused on the broader societal implications of the techniques and technologies studied in the programs.



Greg Demirchyan

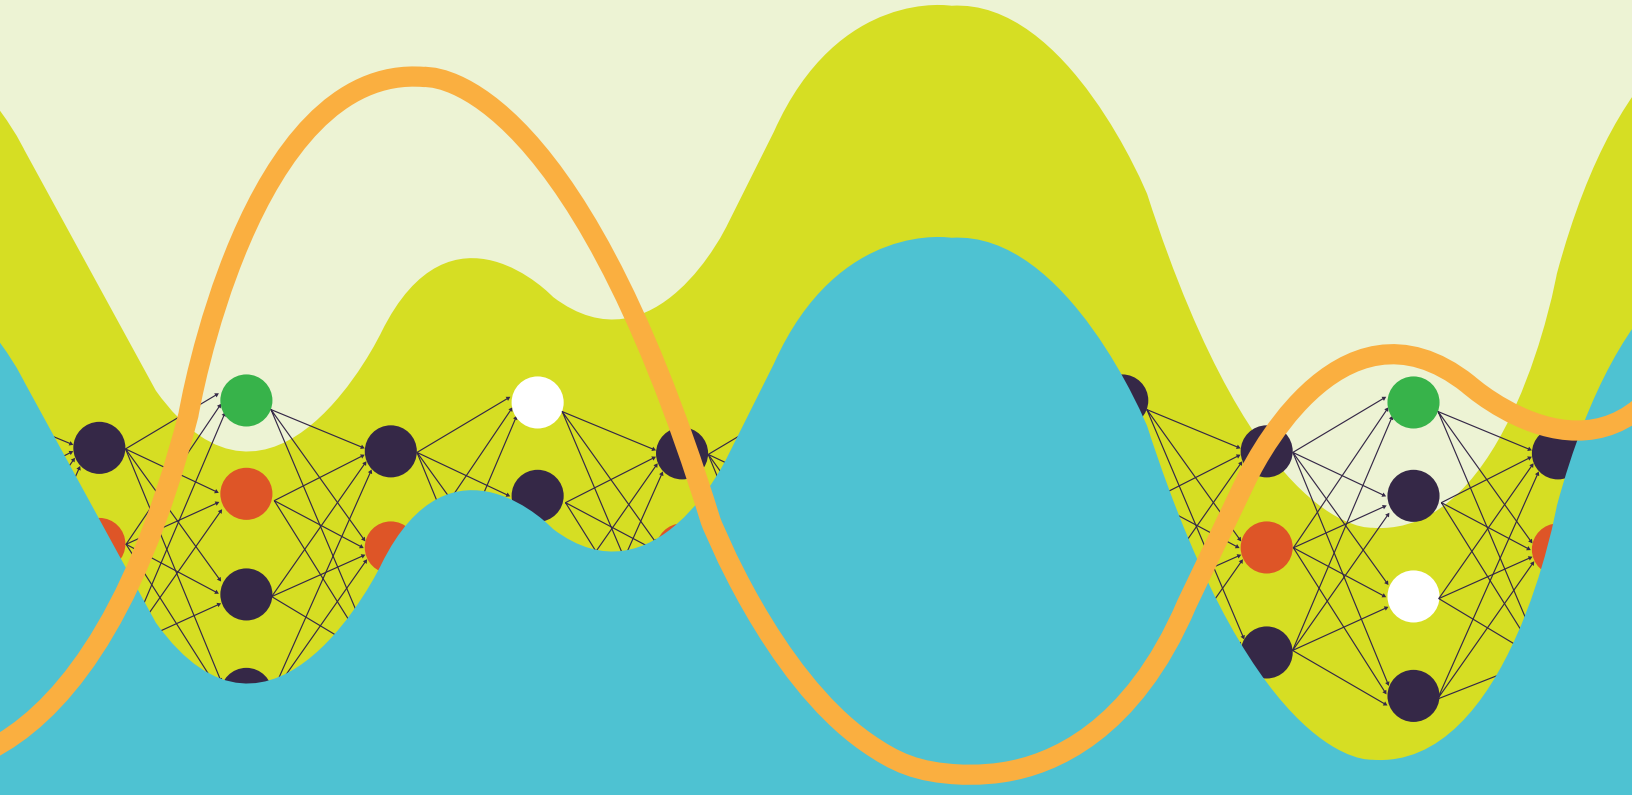


Rui-Jie Yew

These fellowships are part of the Institute's Algorithms, Society, and the Law initiative, which was established in 2019, in response to the ways in which algorithms are increasingly impacting all aspects of human society. Issues of data privacy and security, biased data sets, algorithmic fairness, transparency, and legal governance, to name a few, are emerging as central themes in many Simons Institute programs. And in the last few years, our field has been transformed by the emergence of machine learning and AI; their foundational questions now constitute a major thrust within theoretical computer science, along with the serious legal and societal implications of these new technologies.

Other components of the Algorithms, Society, and the Law initiative include particular research strands within numerous programs, as well as dedicated programs, clusters, public lectures, workshops, policy-and-practice white papers, newsletter articles, and video content. Algorithmic fairness, privacy regulation, ethics in autonomous systems, and many related topics are addressed in these convenings.

“Without definitions, we cannot even say that something is fair or not fair. Definitions allow us to say that something discriminates in one particular way and perhaps doesn’t discriminate in a different way. Definitions have let us expose trade-offs in a way we couldn’t before.”



A Conversation About Algorithmic Fairness

The Simons Institute for the Theory of Computing has played an important role in fostering research on algorithmic fairness from the early days of this field. The Institute held summer clusters on the topic in 2018 and 2019 that brought together researchers in theoretical computer science and the humanities, law, medicine, and education. The central concerns then were that algorithms (using machine learning) were classifying people and helping make critical decisions about such questions as who gets a bank loan or a job, who gets into a college, which defendants and incarcerated people are released on bail or parole, and how police should allocate resources among neighborhoods when deciding which ones to police more intensively.



Sampath Kannan

The AI revolution has made questions of fairness in decisions made by algorithms even more critical. But while the central AI safety question of alignment overlaps with fairness, the two areas are essentially distinct. Conversations around algorithmic fairness, and an independent research agenda for the field, therefore remain essential in the age of AI.

Panelists Jon Kleinberg (Cornell), Omer Reingold (Stanford), Rakesh Vohra (UPenn), and moderator Sampath Kannan (UPenn and former Associate Director of the Simons Institute) are distinguished researchers who have long grappled with questions of fairness. This transcript of their conversation has been edited for length and clarity.

Sampath Kannan: Jon, could you lead off by talking about different notions of fairness, and in what kinds of decision-making contexts they've been employed?

Jon Kleinberg: It's often easiest to think about the fairness definitions when they're motivated by some of the domains that they arise in. So let me start with some of these domains. Fairness as a concept in algorithms and computation has gone through a few iterations, and I guess Omer is in a particularly good position to talk about this, having just finished this really nice book about it.[1]

In the second half of the 20th century, in the '70s and '80s, we had questions related to the fair allocation of resources. And then once we had the internet, a lot of those ideas got used, for example, for fair allocation of bandwidth. Like, how does everybody fairly split the bandwidth in a network? In fact, a lot of how internet congestion control works is rooted in definitions of allocational fairness. So if you're taking a snapshot of the field in, say, the late '90s and asked, "Where does fairness show up in computing?", people might've thought: "How do I fairly assign bandwidth? How do I fairly locate fire stations and hospitals? How do I fairly route vehicles for trash pickup and other services?"

[1] Omer Reingold, *Equal Under the Algorithm: A People's Guide to Fair Computation* (Princeton University Press, 2027).



But around then, we got into the world of predictive algorithms, and algorithms began encroaching on a different set of questions where society has been concerned about fairness — in which one person passes judgment on another person.

In the current context, if we're hiring people, admitting them to college, assigning them medical care, deciding on loans, there's a prediction problem. So a lot of fairness questions ended up using the language of false positives and false negatives. If, for example, you're applying for a job, is the risk of false positives or false negatives for some group higher than for some other group? Is there some category of applicants where they are more likely to be unfairly rejected than some other category?

Sampath: Ricky, fairness has been talked about in economics for ages. Do you feel like theoretical computer science is reinventing the wheel most of the time? Or are you seeing some genuinely different types of questions being asked in the current work on fairness in theoretical computer science?

Rakesh (Ricky) Vohra: So, my answer would be a yes and a no. In the earlier work that Jon pointed out about allocating bandwidth, for example, in my mind the particular criterion used for the allocations was not novel. What was new was a serious concern about finding allocation methods that would operate at scale. I don't see as sharp a distinction between what Jon called the classical problems and these newer ones in which he used the word "judgment." Because I would frame the issue this way: I'm allocating scarce resources. What information am I allowed to use about you to determine whether you deserve anything? And in the bandwidth setting, the implicit assumption we made was that as long as you are on the network, you have an equal claim on the resources.

But there are other settings where we don't think that people necessarily have an equal claim on the resources. Then we are led to the question of, OK, if they don't have an equal claim on the resources, why not? Which of their attributes or features do we think should be considered relevant when making a decision about what resources to give them? What happens

when an attribute we consider to be important is not directly observable or verifiable? Do we ignore it or use a proxy? If so, which one? In many of the applications that Jon mentioned, we're trying to somehow match the resource to a need, or we're trying to assign it in a way that's accurate. And in that sense, I think that's novel.

Sampath: Omer, just to ask the same question on the flip side: Theoretical computer scientists have prided themselves on being able to take topics of societal interest like privacy or fairness and give them quantifiable definitions. Maybe we have done that in some ways for fairness?

Omer Reingold: This is one aspect in which computer scientists and game theory/economics are consistent. The weight we put on definitions is very high. Specifically for computer science and for theory, I think this is a major contribution. Definitions are vital for something like privacy, secrecy, or fairness. We argue about which definitions we care about or not, for various reasons. And there've been examples like COMPAS, where we could have argued forever just about whether this specific system is fair or not.[2]

But through that example, a few definitions emerged, and then the discussion was elevated and had impact in the sense of showing that specific definitions are inconsistent. The power of this viewpoint is very significant. Without definitions, we cannot even say that something is fair or not fair. Definitions allow us to say that something discriminates in one particular way and perhaps doesn't discriminate in a different way. Definitions have let us expose trade-offs in a way we couldn't before. We have examples in each one of these areas — in cryptography, in privacy — and we have definitions in algorithmic fairness that wouldn't have emerged without this perspective, and especially the understanding of the impact of the computational lens on these definitions. We want fairness in a way that is achievable. And it could be achievable by an algorithm. Or it could be achievable by a judge or a loan officer. Understanding the power and limitations of computation in this context is a distinctive contribution of theoretical computer science, and its importance does not depend on whether decisions are made by computers or by people. And that shapes the definition in a very powerful way. This is some of the impact that the theory of computing has had on this area.



Jon Kleinberg

Jon: Sampath, since you were asking for a survey of the landscape of definitions, I want to bring up some that arose in the discussion of the risk assessment tool, COMPAS, that Omer already alluded to. A risk tool like that is designed to assess the risk posed by defendants in criminal proceedings. And of course the first question is: Why is it even doing that? The answer is that this comes from the law, not from computer science or from economics: The law for pretrial detention says that a judge is required to make a prediction about the defendant. At the pretrial phase, when someone's been arrested, the question is where they should await trial. The judge at that point is not supposed to be judging guilt or innocence. They're supposed to be deciding: Does this person wait for trial in the community, or in jail? And they're supposed to be doing that based on a prediction of whether they're likely to return for their court appearance without committing a crime in the meantime. So the law says: Make a prediction. And that was why companies began showing up saying, "Algorithms can help you make predictions."

[2] COMPAS, a recidivism risk assessment tool previously used in the Broward County, Florida, criminal justice system, became the subject of a groundbreaking *ProPublica* story. Julia Angwin et al., "Machine Bias," *ProPublica* (2016).

Assigning a defendant a high risk score is very bad for the defendant because that makes it more likely the judge is going to incarcerate them. So, one thing you might want is to look at different members of different protected classes in the criminal justice system, often by race. And you would ask, "If I look at people who ended up returning for their court appearance and not committing a crime, were they assigned the same average risk across groups?" That's something that you would want. It would be bad if among all people who did not end up committing a crime in the future, members of one protected class got a higher average risk score than members of another class. That would be unfair, and in a way that's causing disproportionate harm to one group. And similarly for people who did go on to commit a crime: Did they get the same average risk score? So, those were two metrics that you might want. One is a kind of false-positive measure, because you're assigning risk to people who did not go on to commit a crime. And the other is a kind of false-negative measure, because you're assigning something other than high risk to people who did go on to commit a crime or fail to return for their court appearance. So those were two definitions of fairness. One is equalizing false-positive outcomes, and one is equalizing false-negative measures.

And then there's a third one, which is: Look at everybody who got assigned, say, 20% risk, and ask, "If I look into the future, what fraction of them went on to commit a crime?" You would expect that if 20% means what it claims to mean, 20% of those people, approximately, should go on to commit a crime. You might say that should be the case for the whole population. And you might also say that should be the case for each group — each protected class separately. Because if I have these two groups, A and B, and I look at everyone in those groups assigned a 20% risk score, then if in Group A, 30% of them go on to commit a crime, and in Group B, 10% of them go on to commit a crime, then we're basically telling the judge, "To correctly use this number, you also have to know what group this defendant came from, because it means different things for different people." And that's exactly what we don't want the judge to be doing. We don't want the judge to be using membership in a protected class as a criterion. So we want risk scores not to be miscalibrated. This is part of the broader issue of calibration in prediction algorithms, and of course, we have Ricky, who is the obvious person to speak to calibration. But the point is, you might want calibration not just across the whole population, but within each group separately, or else the prediction of the tool means different things based on what protected group you're in, which is exactly what we don't want a judge to use.



Rakesh Vohra

And then the punch line is, if you have equalizing of false positives, equalizing of false negatives, calibration across groups, then the result is that, in general, you can't achieve all three of those things simultaneously. They're all things that you should want, but you can't have all of them. And in some sense, that's saying that while COMPAS definitely had the problems that it had, whatever tool was being used, there would have been some problems. And in fact, even more importantly, whatever Broward County, Florida, which was using the tool, had been doing before they brought algorithms into their court system must have had that problem also. Because in the end, this was not a statement about algorithms. This was simply a statement about using numbers to make decisions. And so whatever they were doing, it would've had this problem, because it's a fundamental fact that you can't equalize all the different numerical measures that you want.

Rakesh: I would add that not only do we know which numerical measures are in conflict — we know how to "garble" the outputs of a prediction algorithm to trade off one measure against another. We also have a good understanding of the properties of the underlying data

that generate these conflicts. In other words, we know when the conflicts will not arise.

Sampath: That makes sense. But these are all group notions, right? What are some ways of mathematically formulating notions of fairness to individuals?

Jon: I'm not sure I'm stating a formal mathematical result here, but just intuitively, one way I think about what often feels like the inevitability of some group-level definition is: Imagine you're talking to your doctor and they say, "We've run some blood tests, and I think there's a 65% chance you have this condition." There's something strange about that. You're a human being who either has this condition or doesn't, right?

So the 65% almost by definition is in reference to some underlying group, weighted over the entire population, weighted by what they perceive to be all those other people's similarity to you, and based on what they know from past cases. And they've put you in some kind of a group, possibly with soft boundaries. But you, the individual, either have the condition or don't. So I feel like the instant we start talking about probabilities, we have implicitly jumped to some group-level definition.

Sampath: What about antecedents in the humanities and social sciences? Were these fields able to think about individual fairness in any fashion, maybe in a way that we are not able to capture mathematically?

Omer: I'll take a step back. If you look at much of the research in computer science in the last 16 years about machine learning fairness or AI fairness — the fairness in a process of decision-making — then I would say that there are several major axes of definition and of research. One axis that we touched on a little bit earlier is the gap between decisions that are shooting for equality and decisions that are trying to be faithful to some information and, perhaps, additional constraints — things like: We want to accept the same fraction of people to the university from one group as from another. The moment that the classical definition is trying to equalize some statistical measure between groups, we are reaching toward some kind of equality. These are more toward equality, whereas notions like calibration are a little bit more about accuracy, about faithfulness to the data.



Omer Reingold

The other axis is the axis between individuals and groups, where the classic research, mostly starting before computer science in statistics and other fields, really looks at a very rough distribution of the world into disjoint groups based, for example, on some protected attribute, like ethnicity. That's one side of the equation. The research that we have had really demonstrates how weak those definitions are. They have very few guarantees. And the other direction is notions of individual fairness. You cannot have individual accuracy because of exactly what Jon said, but there are powerful notions of individual fairness.

And then in between is this idea of multigroup fairness, which is a very computational perspective on fairness, and it says that whatever thing you're trying to achieve should not only apply to a few disjoint groups but should apply to a very rich collection of intersecting groups. What's even more important is that these intersecting groups should essentially contain all the groups that you can identify — not only by demographic attributes but also by domain-specific attributes in some measure of complexity, in the amount of computation that you need to identify the group.

And although both approaches are talking about groups, this multigroup perspective has shown itself to be far more powerful and to have much stronger guarantees. So it is an interesting spot between the two extremes of individual fairness and group fairness.

Sampath: Ricky, did you want to add?

Rakesh: Your question, Sampath, assumes a distinction between individual and group that I don't think exists. To argue that an individual has been treated unfairly compared to another presumes the two have some equal claim that is not being honored. But when we prioritize an individual to receive some scarce resource, it's not because of who they are but because of their attributes. These are often coarse, which means multiple individuals share the same attributes, which makes them a group. So, when Omer and you use the word "group," I go back to the question of what are the criteria that I'm allowed to condition a decision on. Conditioning certain decisions on race may be something that many people can agree would be unfair, but there are other settings that I think are more subtle. For example, I think most people would say it's unfair if I were to condition university admissions upon the wealth of your family. On the other hand, we seem to have no problem with conditioning the handing out of loans upon the wealth of your family, which we do by allowing a member of your family to cosign the loan. There is a subtlety here, which is that as we move between particular domains and particular kinds of resources being allocated, we may change whether we are willing to accept a particular attribute as being relevant or not.

Jon: That shows up, for example, in the whole body of legal principles around disparate impact, for example, right?

Sampath: We are talking about the areas where fairness has become important, and how theorists have handled them. There is a longer view of fairness, where we are asking not just that the decision at hand be made fairly, but are concerned with the past. For example, restorative justice. Or when a person is subject to a sequence of decisions as opposed to just one isolated decision. Are new definitions or new algorithmic paradigms needed here?

Omer: The literature on algorithmic fairness from the theoretical computer science community is very rich. For example, causality is a notion that has solid definitions, and can be used to address issues of fairness in a sequence of decisions.

Restorative justice has definitely been discussed. For example, when we're trying to make predictions, we can try to make them based on the current situation, but we can also imagine a more perfect world and try to make decisions that perhaps are based on that world. That's another area that has a lot of research.

In general, the topic of data generation and various faults in data has seen a huge amount of work. Are we predicting some proxy of what we care about? One example, of course, is predicting crime rates, when the parameters that we have access to are something other than the actual crime rates. For example, we have the rate that people were captured doing crimes, which is distinct from the actual rate of crime being committed, and also potentially reflects levels of policing. A slew of issues come into the learning pipeline as a consequence of setting a task that we care about, understanding the data that we have, and gathering the data. And all of it has been studied — it's a very rich research area.

Sampath: The last topic for this discussion: What are some common misconceptions about the study of algorithmic fairness?

Omer: I think that the main misconceptions are that fairness is only about equality, or fairness is only about merit. I have heard both of these, although from different people [laughs]. The implication is that because it's only about equality, it's not relevant, or because it's only about merit, it doesn't represent justice.

Another misconception is about the sophistication of our models of fairness. There are some basic definitions, such as equalized odds and calibration, that we've already mentioned. But these definitions are just one part of a much broader and deeper theory. The field has also identified several very big research directions. These include composition, which asks what happens to fairness when systems are combined, and downstream effects, which concern how decisions shape later opportunities and outcomes. And generative AI raises questions about fairness for much richer outputs, such as text and images, that go beyond a single prediction or decision.



There is the interface with other big communities, like the legal community. It's a big topic, and it's a very multidisciplinary topic that merits attention from other research communities as well, not just our own. But even if we just look at our small piece of this very big universe that contains philosophers and economists and statisticians and a very large collection of other fields — even within our small garden, there is a lot of beauty and a lot of things to do.

So, we haven't solved everything. I mean, if you look at cryptography, think about the time it took to reach the level of sophistication that we're at now. We're not yet there at the level of cryptography, but we're definitely at the level where there are lots of basic models, basic definitions, a lot of understanding of big questions.

And perhaps I'll add to that: It's the scale of computation and these properties that Jon has been talking about that made it such an urgent endeavor, and also made the need to formulate things mathematically so important. But through this process in the last 15 years, we have actually learned a lot about our own values, our own human values that are independent of who is making the decision, whether it's an algorithm or a human being. What properties would we like the decision to satisfy?

There's a lot of discussion of AI alignment, and it's very current and very important. We want AI to correspond to our values, but this also requires us to understand our values in a much finer and better way. And I think this is one of the benefits of this research area.

Sampath: Great final point. Thank you all.

Proving Post-Quantum Security, in and out of Calvin Lab: A Portrait of Nick Spooner

The coming era of large-scale, fault-tolerant quantum computers holds promise for advances in fundamental physics and chemistry research, the development of new quantum algorithms, and quantum speedups for many kinds of computational problems. At the same time, these devices threaten to break public-key cryptography, the security of which rests on the hardness of mathematical problems that quantum algorithms can solve efficiently. As cryptographers develop post-quantum alternatives, Nick Spooner is focused on the next challenge: proving that these new systems are actually secure.

Nick's first long-term visit to the Simons Institute was in the program on Proofs, Consensus, and Decentralizing Society, in 2019. At the time, he was a graduate student in Berkeley's EECS department, working with Alessandro Chiesa. Now, seven years later, he is an Assistant Professor of Computer Science at Cornell, where he does research on post-quantum proof systems. And over the course of that trajectory, he has been a regular participant in the Institute's extensive programming in the areas of quantum computing, cryptography, and their intersections.

Nick's main research interest — proof systems — unites complexity theory and cryptography. Proof systems are concerned with how to verify a claim efficiently, and often rely on cryptographic hardness for their security. Formally establishing this security often proceeds by reduction: Assume that an attacker can break a given encryption scheme, and then use that attacker as a subroutine in an algorithm for solving a particular mathematical problem believed to be intractable. If the problem cannot be solved efficiently, we can conclude that there exists no efficient attacker that can break the encryption scheme.

But a classical reduction may depend on running the attacker repeatedly, restoring it to a previous state between trials. In the case of quantum computers, however, measuring a quantum state can irreversibly alter it, and the no-cloning principle prevents an exact copy from being saved in advance. A quantum attacker may destroy its own internal state as it runs. So even when a cryptographic scheme rests on a problem believed to resist quantum algorithms, its classical security proof may no longer apply.



Nick Spooner

Nick recounts that during the COVID-19 pandemic, he had "a lot of free time to work on a very hard problem." He and collaborators Chiesa, Fermi Ma, and Mark Zhandry used that time to develop a new quantum rewinding technique and apply it to Kilian's four-message succinct argument system.[1] Their method, which recovers the information needed for the security reduction without violating quantum mechanics, demonstrated for the first time a post-quantum succinct argument system based on a concrete computational hardness assumption.

[1] Alessandro Chiesa et al., "Post-Quantum Succinct Arguments: Breaking the Quantum Rewinding Barrier," *FOCS* (2022).

[2] Tom Gur et al., "Perfect Zero-Knowledge PCPs for #P," *STOC* (2024).

[3] Tom Gur et al., "A Zero-Knowledge PCP Theorem," *STOC* (2025).

Nick's long-term goal is to build a coherent theory of post-quantum security proofs. When a classical proof has a quantum analogue, we should be able to construct it; and when there is no analogue, we should be able to explain why.

A second strand of his work — zero-knowledge probabilistically checkable proofs (PCPs) — examines another apparent contradiction. PCPs give a verifier access to just a few small sections of a proof, while zero-knowledge proofs reveal nothing beyond the truth of the statement being proved. How can a small number of revealed proof symbols be informative enough to establish validity, while being uninformative about everything except validity?

In collaboration with Tom Gur and Jack O'Connor, Nick developed perfect zero-knowledge PCPs that reconcile these demands.[2] Their constructions permit local verification, while guaranteeing that every revealed portion of the proof can be simulated without access to hidden contents. Their discovery grew out of an idea Nick first pursued as a graduate student, and culminated in results presented at STOC in 2024 and 2025.[3] He sees in these results a possible clue to the quantum PCP conjecture. Following the logic of the no-cloning theorem, examining any few qubits of a quantum PCP should reveal almost nothing. But the quantum PCP conjecture requires that measurements on only a few qubits be sufficient for determining whether the proof has the required global property. The construction by Nick and his collaborators demonstrates one mechanism by which local views can be simultaneously information poor and globally informative.

The Simons Institute has been a kind of intellectual home for Nick over the course of his research career, starting in graduate school. "One of the really nice things about being a student at Berkeley is that you have an ongoing relationship with the Simons Institute," he recalls. "Whenever there was a program that was of interest to me, or talks or whatever, I could just walk down the hill and be able to go to it. I felt extremely privileged to have this amazing resource where all of the greatest minds in the field are showing up at your doorstep. I don't know of anywhere else that has this feature. So I spent quite a lot of time there."

Since then, he has been a participant in the 2023 and 2026 quantum summer clusters, and was a Research Fellow in the Institute's 2025 program on Cryptography 10 Years Later: Obfuscation, Proof Systems, and Secure Computation.

Before taking his post at Cornell, Nick served on the faculty of the University of Warwick and was a visiting faculty member at NYU. He was a postdoctoral researcher at Boston University and holds a PhD from Berkeley, a master's from ETH Zurich, and a bachelor's from Cambridge. When not paving the way for secure post-quantum cryptography and proof systems, he can be found hiking Ithaca's gorges and singing in his local community chorus.

The Simons Institute trains the largest cohort of postdoctoral-level researchers in theoretical computer science in the world. This includes multiyear postdocs in our research pods and semester-long Research Fellows within each research program. Many of this year's Research Fellows have gone on to tenure-track positions at prestigious institutions, including Princeton, NYU, Virginia Tech, the Weizmann Institute of Science, the University of Edinburgh, the University of Sydney, the Institute of Theoretical Physics (IFT, Madrid), Nanyang Technological University, and DePaul University.



“Together, these advances open the possibility that fault-tolerant quantum computation could be realized much sooner than expected.”

Rapid Developments Toward Fault-Tolerant Quantum Computing with Neutral Atoms

by Madelyn Cain (Oratomic) and Harry Levine (UC Berkeley, Oratomic)

In the last decade, quantum computing has developed rapidly. Remarkable progress has been made across many hardware platforms such as trapped ions, superconducting qubits, and more recently, neutral atoms. Systems have been scaled to hundreds of qubits, and qubit control has also improved dramatically. This progress has been motivated by the overarching goal of realizing useful quantum computation, with potential applications across physics, materials science, machine learning, and beyond. These prospects are at the heart of a research pod and numerous research programs and clusters at the Simons Institute, as well as the NSF-funded Challenge Institute for Quantum Computation (CIQC), led by UC Berkeley.

Useful algorithms will likely require thousands of qubits and up to billions of operations, however, thereby requiring extremely small error rates per operation, far below what seems realizable in hardware.[1] In the mid-1990s, this challenge motivated Shor, Calderbank, and others to develop fault-tolerant quantum computation (FTQC), which encodes protected "logical qubits" into many redundant "physical qubits." This enables arbitrary suppression of logical errors as the redundancy grows, so long as physical error rates remain below a characteristic threshold of around 1%. The development of FTQC opens the door to useful quantum computation, but at a cost: Traditional approaches designed for hardware with planar connectivity (physical qubits are laid out in two dimensions, and are only directly entangled with nearest neighbors) require hundreds of physical qubits per logical qubit, meaning useful algorithms can require on the order of a million physical qubits.[2]

This daunting number — over 1 million qubits — has for decades made the prospect of large-scale quantum computation appear distant. The last few years, however, have upended this picture, enough to surprise even many of us working in the field. Hardware performance has improved dramatically, with exciting error-correction demonstrations across several hardware platforms (including neutral atoms, trapped ions, and superconducting qubits). Simultaneously, advances in quantum error-correction theory have introduced dramatically more efficient codes and logic by leveraging nonlocal connectivity, bringing useful computations from millions of qubits down to tens of thousands. Most recently, neutral-atom processors have emerged, with the ability to implement these new error-correcting codes by dynamically reconfiguring atoms, and with a path toward scaling system sizes and control capabilities. Together, these advances open the possibility that fault-tolerant quantum computation could be realized much sooner than expected.



Madelyn Cain

[1] Michael E. Beverland et al., "Assessing Requirements to Scale to Practical Quantum Advantage," arXiv (2022).

[2] Craig Gidney, "How to Factor 2048 Bit RSA Integers with Less Than a Million Noisy Qubits," arXiv (2025).

Advances in neutral-atom hardware

Neutral-atom hardware uses qubits encoded into individual atoms. These atoms are laser-cooled to millionths of a degree above absolute zero, and then isolated and levitated in arrays of tightly focused lasers called optical tweezers. Laser-cooled atoms have a rich history in physics as pristine platforms for observing exotic quantum effects among millions of interacting quantum particles, including famously the creation of quantum states of matter like the Bose–Einstein condensate. Their use as highly controlled qubits has emerged only much more recently, thanks to the rapidly growing toolbox for high-fidelity control built upon tweezers for isolating arrays of individual atoms.

Among the most important ingredients in quantum computation is the entangling gate, realized by getting two qubits to interact in a controlled way. Neutral atoms held in tweezers are by default noninteracting with one another, but it was proposed in 2000 by Jaksch, Lukin, and others that by using lasers to excite atoms to highly excited electronic orbitals — called Rydberg states — the atoms could become entangled.[3] This was first demonstrated in the early 2010s by the groups of Saffman[4] and Browaeys,[5] but until the late 2010s it still had error rates above around 20%, not yet in the domain of high-fidelity processing. We built upon this work at Harvard in 2019 in the group of Mikhail Lukin by using simplified laser pulses and higher-quality laser sources, dropping error rates to 3%.[6] Gates have continued to improve quickly in the last few years, with recent demonstrations showing as low as < 0.2% error (well below typical error-correction thresholds).[7] An additional key feature is that gates can be applied in parallel on relatively large arrays of atoms at the same time, simplifying control and calibration complexity.

In just the last five years, another key new tool was introduced: coherent transport of atoms.[8] Atoms can be reconfigured by moving the tweezers around mid-computation while preserving their entanglement. This enables efficient control strategies where atoms can be arranged in pairs, entangled, and then rearranged into new pair configurations and entangled again. Since the reconfiguration can be parallelized and since the laser pulse also implements gates on many atoms in parallel, this strategy emerged in the Harvard lab as a surprisingly effective way to program quantum circuits. Remarkably, this reconfigurability has become an increasingly central capability for enabling efficient error correction, intersecting a parallel thread of research on error-correction theory.

This reconfigurable neutral-atom architecture is now driving exciting explorations of fault-tolerant computation with hundreds of physical qubits across the community.[9] Error rates are below the critical threshold for error correction to start improving performance. Repeated syndrome extraction, the key primitive for checking and correcting errors, has been demonstrated in both codes with local connectivity and those with nonlocal connectivity. Additionally, transporting blocks of atoms around enables direct parallel operations on groups of *logical* qubits ("transversal gates"), which are normally more challenging to operate on.

Building on these demonstrations, a great deal of attention has turned toward scaling up, with the ultimate goal of realizing large-scale, error-corrected computation. Recent efforts to scale have produced larger arrays of around 10,000 atoms, although not yet with enough

[3] Dieter Jaksch et al., "Fast Quantum Gates for Neutral Atoms," *Physical Review Letters* (2000).

[4] Larry Isenhower et al., "Demonstration of a Neutral Atom Controlled-NOT Quantum Gate," *Physical Review Letters* (2010).

[5] Tatjana Wilk et al., "Entanglement of Two Individual Neutral Atoms Using Rydberg Blockade," *Physical Review Letters* (2010).

[6] Harry Levine et al., "Parallel Implementation of High-Fidelity Multiqubit Gates with Neutral Atoms," *Physical Review Letters* (2019).

[7] Simon J. Evered et al., "High-Fidelity Entangling Gates and Nonlocal Circuits with Neutral Atoms," arXiv (2026).

[8] Dolev Bluvstein et al., "A Quantum Processor Based on Coherent Transport of Entangled Atom Arrays," *Nature* (2022); Dolev Bluvstein et al., "Logical Quantum Processor Based on Reconfigurable Atom Arrays," *Nature* (2024).

[9] Dolev Bluvstein et al., "A Fault-Tolerant Neutral-Atom Architecture for Universal Quantum Computation," *Nature* (2026); Atom Computing et al., "Quantum Error Correction with the Toric Code," arXiv (2026); Bichen Zhang et al., "Logical Qubits with Erasure Conversion Using Metastable Neutral Atoms," *Nature Physics* (2026); Pedro Sales Rodriguez et al., "Experimental Demonstration of Logical Magic State Distillation," *Nature* (2025).

control to implement error correction.[10] Another important known bottleneck has been the gradual loss of atoms from tweezers, motivating new generations of machines that offer *continuous reloading*, where atoms can be replaced during computations without disturbing other atoms.[11] As other hardware primitives such as gate and measurement performance have continued to improve as well, integrating these advances into large-scale systems with full control capabilities for fault tolerance remains a key open challenge. Nonetheless, the recent demonstrations and the steady stream of new innovative atomic-control strategies suggest that machines with tens of thousands of highly controlled atomic qubits are within reach.

Developments in fault tolerance

The fast-paced development of neutral-atom processors has coincided with a rise in substantially more efficient methods for fault-tolerant computation. In many ways, it mirrors the development of practical fault tolerance for hardware with planar connectivity based on the "surface code," which encodes one logical qubit in hundreds of physical qubits. Progress over the past two decades has culminated in improvements in encoding rates (e.g., Fowler,[12] Gidney,[13] and Low[14]), logic gate constructions (e.g., Horsman[15] and Litinski[16]), and decoding, i.e., inferring which errors occurred so they can be corrected (e.g., Google[17] and Higgott[18]). Recent developments in codes, logic, and decoders for hardware with nonlocal connectivity have occurred on rapid timescales, with many foundational advances coming in just the past five years.



Harry Levine

The development of hardware leveraging nonlocal connectivity (enabled by the reconfigurability of neutral atoms and also trapped ions) provides a practical path to overcome the low encoding rates of the surface codes by vastly expanding the design space of realizable codes and logic gates. High-rate classes of quantum low-density parity-check (qLDPC) codes, which critically leverage nonlocality, were originally proposed in the mid-2010s.[19] Subsequent constructions improved the encoding rate and code distance (which sets how many physical errors can be detected and corrected) of practical code instances (e.g., Breuckmann,[20] Eberhardt,[21] and Panteleev[22]). In just the past year, constructions at algorithmically relevant code

[10] Hannah J. Manetsch et al., "A Tweezer Array with 6,100 Highly Coherent Atomic Qubits," *Nature* (2025); Yuqing Wang et al., "Trapping 11,000 Atoms in a Tweezer Array Generated by a Single Metasurface," arXiv (2026).

[11] Neng-Chun Chiu et al., "Continuous Operation of a Coherent 3,000-Qubit System," *Nature* (2025); Yiyi Li et al., "Fast, Continuous, and Coherent Atom Reloading in a Neutral-Atom Qubit Array," *Science* (2026); Atom Computing et al., "Quantum Error Correction."

[12] Austin G. Fowler et al., "Surface Codes: Towards Practical Large-Scale Quantum Computation," *Physical Review A* (2012).

[13] Craig Gidney et al., "Yoked Surface Codes," *Nature Communications* (2025).

[14] Guang Hao Low et al., "A Denser Planar Surface Code," arXiv (2026).

[15] Dominic Horsman et al., "Surface Code Quantum Computing by Lattice Surgery," *New Journal of Physics* (2012).

[16] Daniel Litinski, "A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery," *Quantum* (2019); Daniel Litinski et al., "Lattice Surgery with a Twist: Simplifying Clifford Gates of Surface Codes," *Quantum* (2018).

[17] Andrew W. Senior et al., "A Scalable and Real-Time Neural Decoder for Topological Quantum Codes," arXiv (2025).

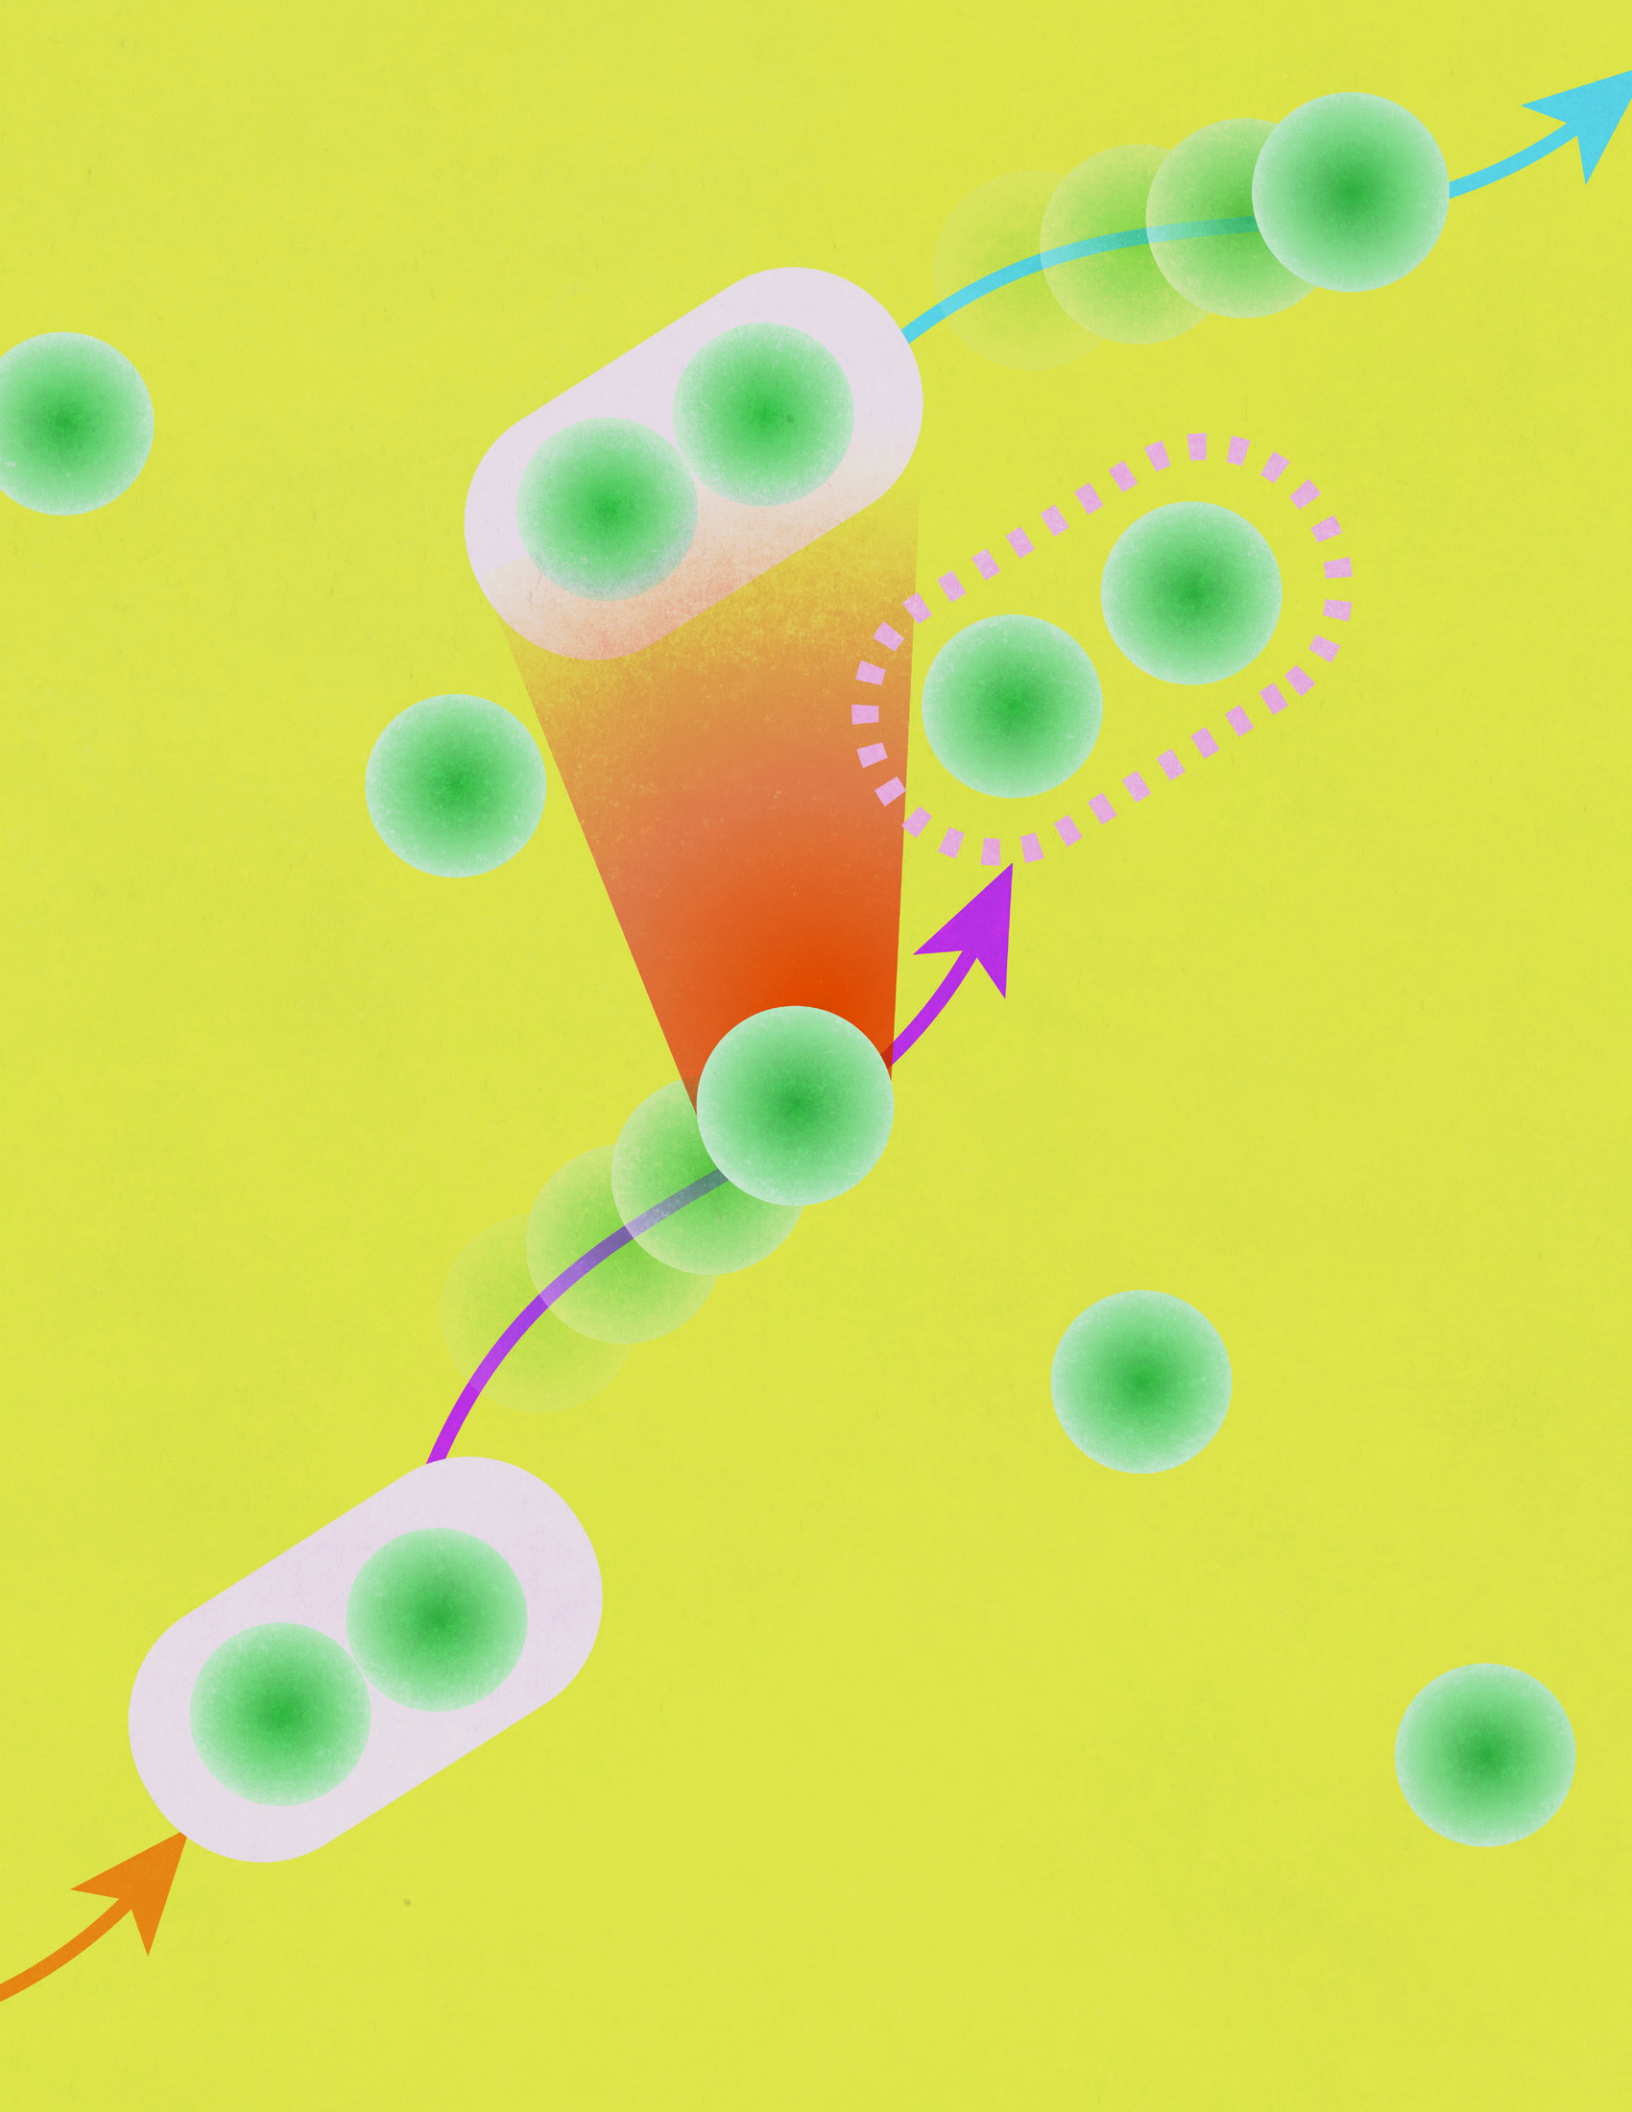
[18] Oscar Higgott et al., "Sparse Blossom: Correcting a Million Errors per Core Second with Minimum-Weight Matching," *Quantum* (2025).

[19] Jean-Pierre Tillich et al., "Quantum LDPC Codes with Positive Rate and Minimum Distance Proportional to the Square Root of the Blocklength," *IEEE Transactions on Information Theory* (2014).

[20] Nikolas P. Breuckmann et al., "Quantum Low-Density Parity-Check Codes," *PRX Quantum* (2021).

[21] Nikolas P. Breuckmann et al., "Balanced Product Quantum Codes," *IEEE Transactions on Information Theory* (2021).

[22] Pavel Panteleev et al., "Degenerate Quantum LDPC Codes with Good Finite Length Performance," *Quantum* (2021).



distances have been developed that encode one logical qubit in as few as two physical qubits,[23] together with hardware-efficient implementations.[24] This corresponds to a ~1,000x reduction in qubit overhead compared with the surface code, while achieving comparable or improved logical error rates, both for preserving information and for manipulating the encoded qubits (e.g., Cohen,[25] Williamson,[26] Cross,[27] and He[28]). Finally, continued algorithm development has reduced the logic gate depth of major quantum algorithms, putting them within the reach of daylong computations with the relatively slow computation speed of neutral-atom hardware.[29] Together, these advances enable the possibility of utility-scale computation with only tens of thousands of physical qubits.

The remarkable convergence of experimental and theoretical progress opens the door toward utility-scale quantum computation. Realizing such a computer in practice represents a substantial engineering challenge: The hardware will need to combine demonstrated state-of-the-art building blocks of fault-tolerant computation, such as high-fidelity operation, continuous reloading, and imaging, control, and trapping of large arrays. Furthermore, these components must be integrated at scales of tens of thousands of atomic qubits, beyond what has been realized in existing hardware. Nonetheless, such machines now appear within reach, and would be far more capable than was expected even several years ago.

The advent of large-scale quantum computers would open a new chapter of quantum technology. Continued work on quantum algorithms — in fields such as many-body quantum simulation, chemical modeling, artificial intelligence, and cryptography — is needed to broaden the applications and understand the capabilities of utility-scale computers. In key areas where algorithmic development remains challenging, such as quantum simulation and cryptanalysis of post-quantum schemes designed to be secure against quantum attacks, heuristics will be testable using quantum hardware. Looking ahead, next-generation advances in scaling and networking quantum hardware could enable applications in sensing, quantum communication, and distributed quantum computation.

Many of the recent contributions that have reshaped the outlook for quantum computing emerged from exploratory academic research. As the quantum industry grows and concentrates on large-scale systems, strong collaboration between industry and academia — including leading research hubs like the Simons Institute — will be essential to laying the groundwork for the next generation of advances in fault tolerance.

[23] Siyi Yang et al., "Spatially-Coupled QLDPC Codes," *Quantum* (2025); Kenta Kasai, "Breaking the Orthogonality Barrier in Quantum LDPC Codes," *Quantum* (2026); Madelyn Cain et al., "Shor's Algorithm Is Possible with as Few as 10,000 Reconfigurable Atomic Qubits," arXiv (2026).

[24] Chen Zhao et al., "Towards Ultra-High-Rate Quantum Error Correction with Reconfigurable Atom Arrays," arXiv (2026); Aditya Bhardwaj et al., "High-Rate qLDPC Processors," arXiv (2026).

[25] Lawrence Z. Cohen et al., "Low-Overhead Fault-Tolerant Quantum Computing Using Long-Range Connectivity," *Science Advances* (2022).

[26] Dominic J. Williamson et al., "Low-Overhead Fault-Tolerant Quantum Computation by Gauging Logical Operators," *Nature Physics* (2026).

[27] Andrew W. Cross et al., "Improved QLDPC Surgery: Logical Measurements and Bridging Codes," arXiv (2025).

[28] Zhiyang He et al., "Extractors: QLDPC Architectures for Efficient Pauli-Based Computation," arXiv (2025).

[29] Ryan Babbush et al., "Securing Elliptic Curve Cryptocurrencies Against Quantum Vulnerabilities: Resource Estimates and Mitigations," *PRX Quantum* (2026); Gidney, "How to Factor 2048 Bit RSA Integers"; Beverland et al., "Assessing Requirements"; André Schrottenloher, "Optimized Point Addition Circuits for Elliptic Curve Discrete Logarithms," arXiv (2026).



Photo: Bruce Damonte

The Proof Behind the Proof: New Routes to SNARGs for NP

by Yael Tauman Kalai (MIT) and Abhishek Jain (Johns Hopkins)

Societies are built on cooperation, which requires trust. But trust is a scarce commodity. Proofs offer an alternative: Rather than asking to be believed, the prover supplies evidence that a skeptical verifier can check for itself. Thus, even mutually distrustful parties can agree on the facts.

But there is a catch. Checking evidence takes work, and the work may grow with the complexity of the claim. Andrew Wiles's proof of Fermat's Last Theorem runs to 129 pages; its referees found a gap within months, and closing it took another year. Outside mathematics, matters are worse still: The most natural proof that a large computation was carried out correctly is the computation itself, and checking it means repeating it. Must the effort of checking a claim grow with the effort of establishing it? Or can every claim, however complex and however much evidence it rests on, come with a proof that is short and fast to verify? These are questions we explored during the Summer 2025 Simons Institute research program Cryptography 10 Years Later: Obfuscation, Proof Systems, and Secure Computation.

Complexity theory offers a discouraging first answer. For NP statements, proofs that are both short (say, of polylogarithmic size) and unconditionally sound are too much to hope for: If they existed, one could decide any NP problem by enumerating all short proofs, placing NP in quasi-polynomial time. Cryptography, however, is in large part the art of overcoming such impossibility results. The way around this is to ask for less: not that a false statement have no convincing proof, but that no *computationally bounded* prover be able to find one.

A short proof with this guarantee is called a *succinct non-interactive argument*, or SNARG. If a prover holds a (potentially large) witness w that can be checked in time T, a SNARG lets it convert w into a proof π whose length, and the verifier's running time, grow only polylogarithmically with T, beyond the unavoidable cost of reading the statement. Typically, the prover and the verifier share a public *common reference string* (CRS), generated before the statement is proved, that intuitively embeds a cryptographic challenge for the prover.



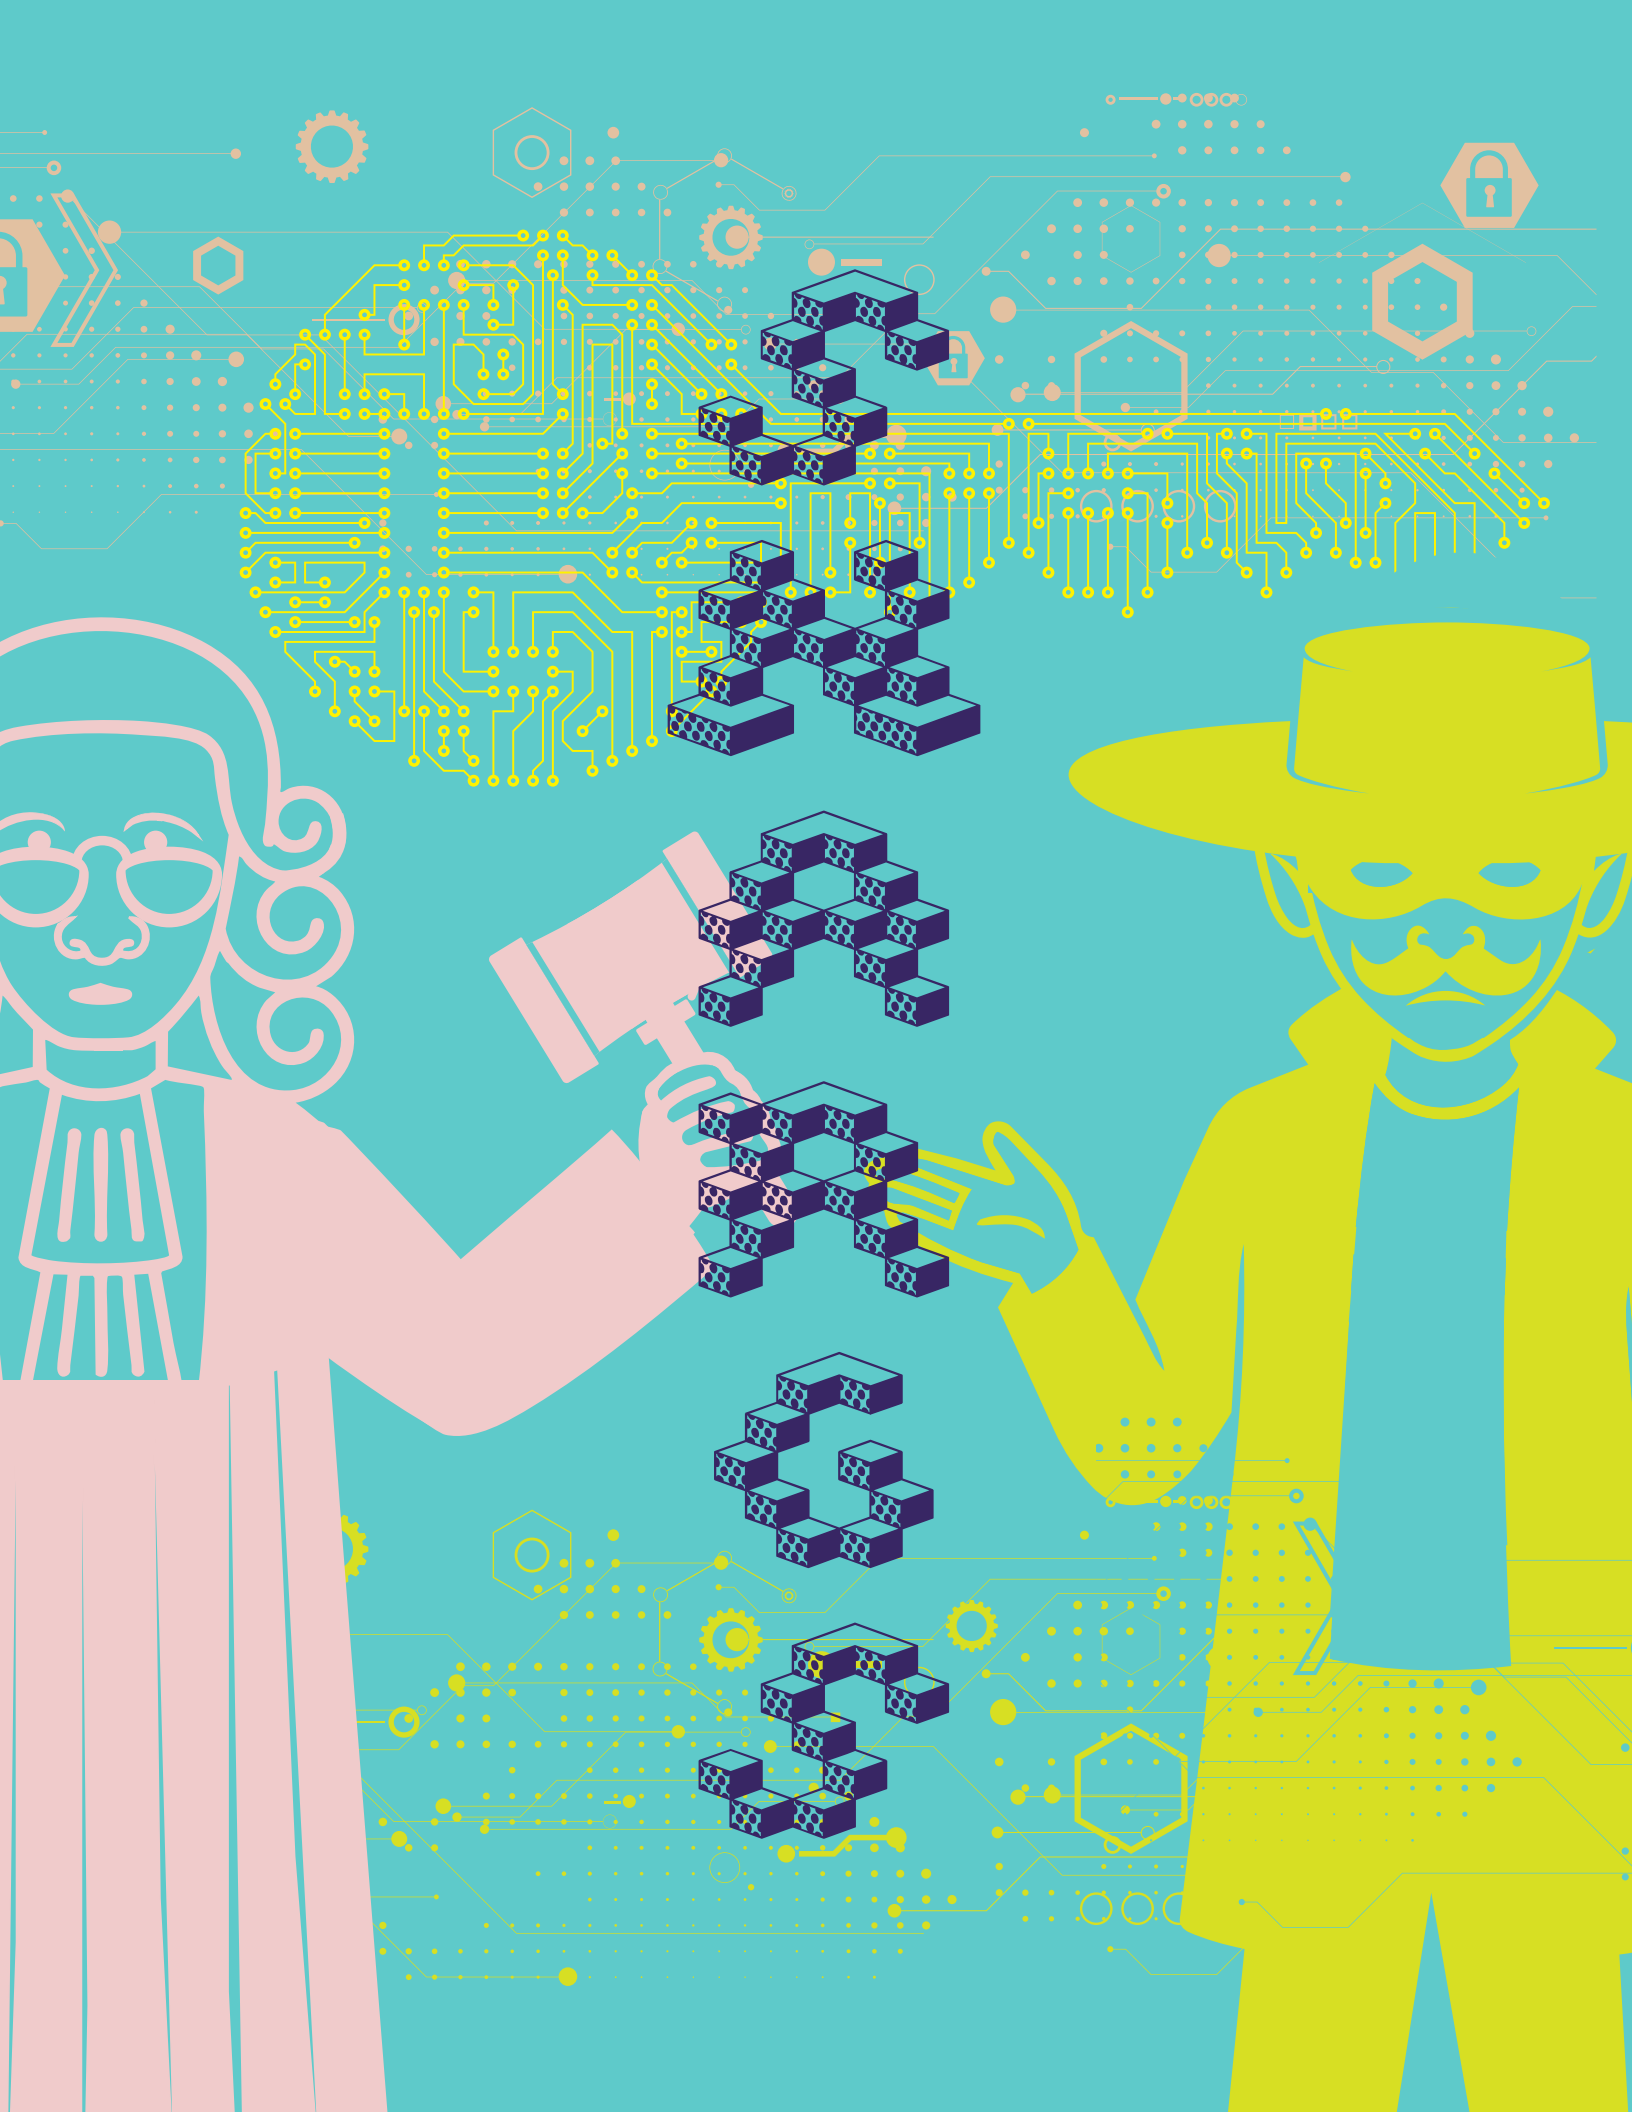
Yael Tauman Kalai

SNARGs have far-reaching applications. In a blockchain, thousands of computers must agree on a shared ledger; with SNARGs, an untrusted party can attest to the validity of a large batch of transactions by posting a short proof on the ledger that anyone can verify in an instant. And as computation and decision-making are increasingly delegated to machine learning systems, a short certificate of what was computed, by which model, and on which data, may soon be a necessity rather than a convenience.

Remarkably, SNARGs exist, at least heuristically, and their roots run deep in theoretical computer science. Building on the PCP theorem,[1] Kilian and Micali each showed how to build succinct arguments for NP.[2]

[1] Sanjeev Arora et al., "Probabilistic Checking of Proofs: A New Characterization of NP," *FOCS* (1992); Sanjeev Arora et al., "Proof Verification and Hardness of Approximation Problems," *FOCS* (1992).

[2] Joe Kilian, "A Note on Efficient Zero-Knowledge Proofs and Arguments," *STOC* (1992); Silvio Micali, "CS Proofs," *FOCS* (1994).



Their descendants, refined over three decades, are now deployed at scale, most visibly in the Ethereum ecosystem. Yet, the security of all these rests on heuristics or on nonstandard assumptions, which raises a foundational question: *Can we construct SNARGs for NP from standard, well-studied cryptographic assumptions?*

The obfuscation route

For more than a decade, the most general answer has come from indistinguishability obfuscation (iO). Very roughly, iO transforms a program into another program with the same input/output behavior, while guaranteeing that obfuscations of two functionally equivalent programs cannot be distinguished.

To see why obfuscation helps, imagine first that an all-powerful trusted party, call her Oprah, sits between the prover and the verifier. Oprah holds a secret signing key. Given a statement x and a candidate witness w, Oprah checks the witness and, if it is valid, publishes a short signature on x. The signature is the proof, and the verifier simply checks it. This is a perfect SNARG with an absurd setup. Obfuscation lets us keep the SNARG and drop the setup: Take the program that Oprah runs, with the signing key hardwired inside, obfuscate it, and publish the obfuscated program as the CRS. Anyone holding a valid witness can now run the program and obtain a signature, and as Sahai and Waters have shown, no efficient prover can obtain one for a false statement.[3]

The result was a landmark demonstration of the power of obfuscation. But iO is a heavy hammer: an exceptionally strong primitive whose constructions from well-studied assumptions are enormously complex and far from practical, and rely on bilinear maps, which quantum computers break.[4] The



Abhishek Jain

question of whether iO can be built from post-quantum assumptions such as learning with errors (LWE) remains open. This raises a natural question: Is obfuscation truly necessary, or has it merely been our most convenient tool?

A barrier

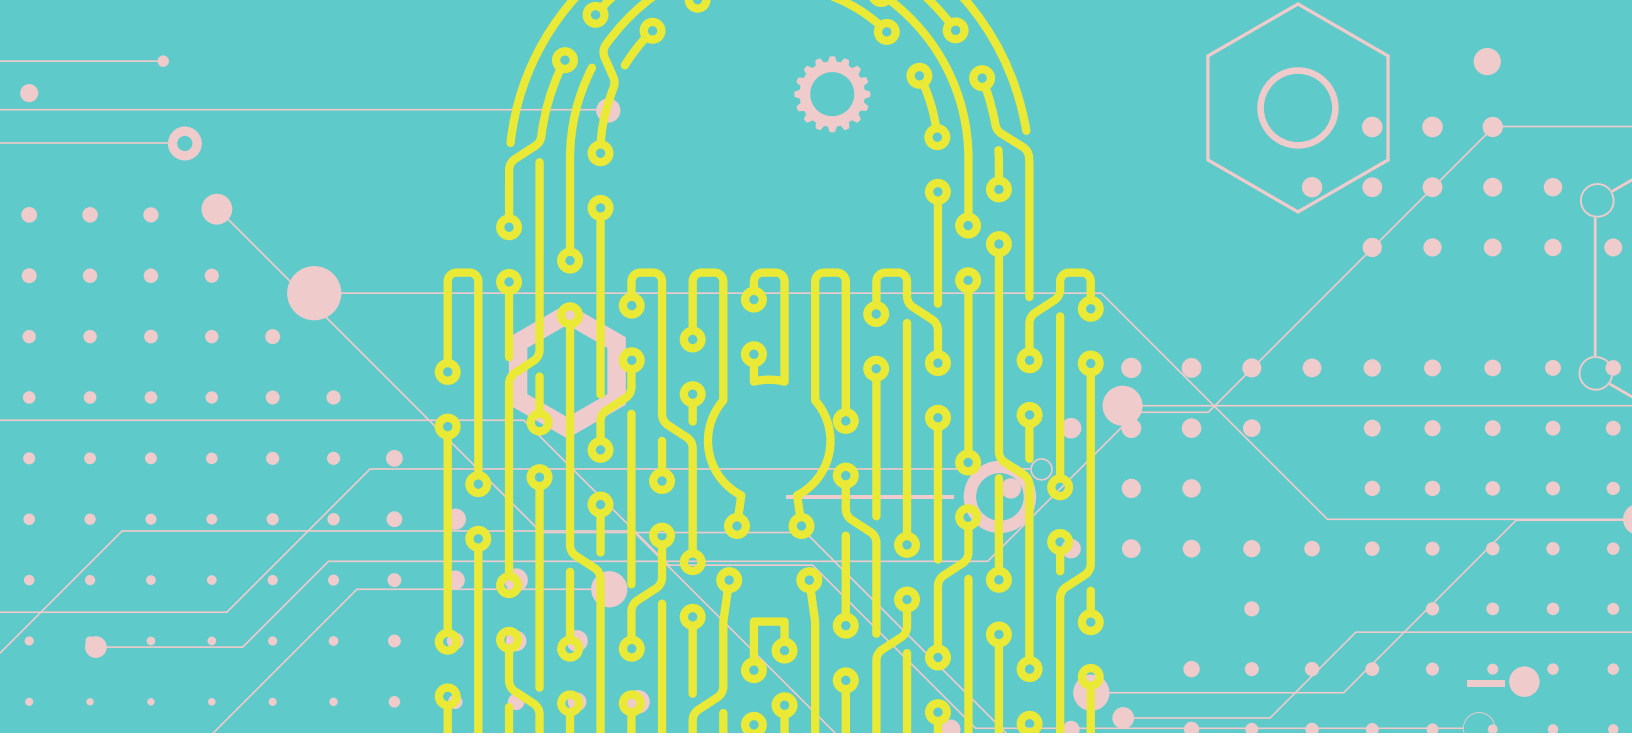
There was reason to think such powerful tools might be unavoidable. Gentry and Wichs identified a barrier to proving general SNARGs secure from standard assumptions.[5] It applies to security proofs that treat a cheating program as a *black box*: They run it and use its answers, but never inspect its code.[6] Assumptions such as the difficulty of factoring fall within this barrier: A claimed solution can be checked efficiently, a property that cryptographers call *falsifiability*. Obfuscation's security guarantee does not have this property, offering one escape. For years, similarly powerful assumptions seemed hard to avoid.

[3] They prove soundness when the statement is fixed before the CRS is published. Allowing the adversarial prover to choose the statement after seeing the CRS took another decade and considerably more machinery. Amit Sahai et al., "How to Use Indistinguishability Obfuscation: Deniable Encryption, and More," *STOC* (2014); Brent Waters et al., "Adaptively-Sound Succinct Arguments for NP from Indistinguishability Obfuscation," *STOC* (2024); Brent Waters et al., "Adaptive Security in SNARGs via iO and Lossy Functions," *CRYPTO* (2024); Brent Waters et al., "A Pure Indistinguishability Obfuscation Approach to Adaptively-Sound SNARGs for NP," *CRYPTO* (2025).

[4] Aayush Jain et al., "Indistinguishability Obfuscation from Well-Founded Assumptions," *STOC* (2021).

[5] Craig Gentry et al., "Separating Succinct Non-Interactive Arguments from All Falsifiable Assumptions," *STOC* (2011).

[6] The original theorem concerns claims chosen after seeing the CRS. Chen and Jin give related barriers for claims fixed beforehand, under restrictions on the setup length and the security proof. Liyan Chen et al., "On the Impossibility of SNARGs with Short CRS (or: Revisiting Gentry-Wichs Barrier in the Non-Adaptive Setting)," *FOCS* (2025).



A new beginning

To sidestep this barrier, researchers initiated a fresh study of SNARGs by focusing on prominent subclasses of NP. Kalai, Raz, and Rothblum constructed a SNARG for *deterministic* computations from standard cryptographic assumptions, such as LWE,[7] by exploiting an unexpected connection to *non-signaling strategies*, a notion originating in quantum mechanics.[8] Their construction requires the verifier to keep a secret key. This restriction was later overcome by Kalai, Paneth, and Yang, relying on a new assumption on bilinear groups.[9]

All subsequent progress in this area started from a seemingly modest goal of *batch arguments*. A batch argument allows a prover to certify k NP statements with a single proof whose size depends only on the size of a single witness, but only polylogarithmically on k. Batch arguments for NP were first constructed from LWE by Choudhuri, Jain, and Jin,[10] and subsequently from standard group-based assumptions.[11]

Modest as it sounds, batching provides a natural template to certify arbitrary nondeterministic computations. Break the computation into T small steps, the i-th asserting that one application of the transition function carries configuration c_{i-1} to configuration c_i; commit succinctly to the sequence of configurations; and certify all T steps with a single batch argument.

Why is this sound? In the obfuscation route, the entire witness was checked in one stroke, inside Oprah's program. Here the verifier receives only a short proof, far too short to contain the computation, and the security analysis can extract from it only a little at a time. For *deterministic* computations, Choudhuri, Jain, and Jin, as well as a separate team consisting of Kalai, Vaikuntanathan, and Zhang, devised a proof strategy that proceeds through a sequence of *checkpoints*: Any proof the verifier accepts must be consistent with the correct computation up to step 1; given that, up to step 2; and so on to the end, each checkpoint following from the previous one by the local check for that step.[12] The argument works because a deterministic computation has exactly one correct configuration after i steps, so each checkpoint is a well-defined claim.

A *nondeterministic* computation, however, has no such thing as the correct intermediate state. Without it, there are no obvious checkpoints, and no obvious way to argue soundness for NP.

[7] Yael Tauman Kalai et al., "How to Delegate Computations: The Power of No-Signaling Proofs," *STOC* (2014).

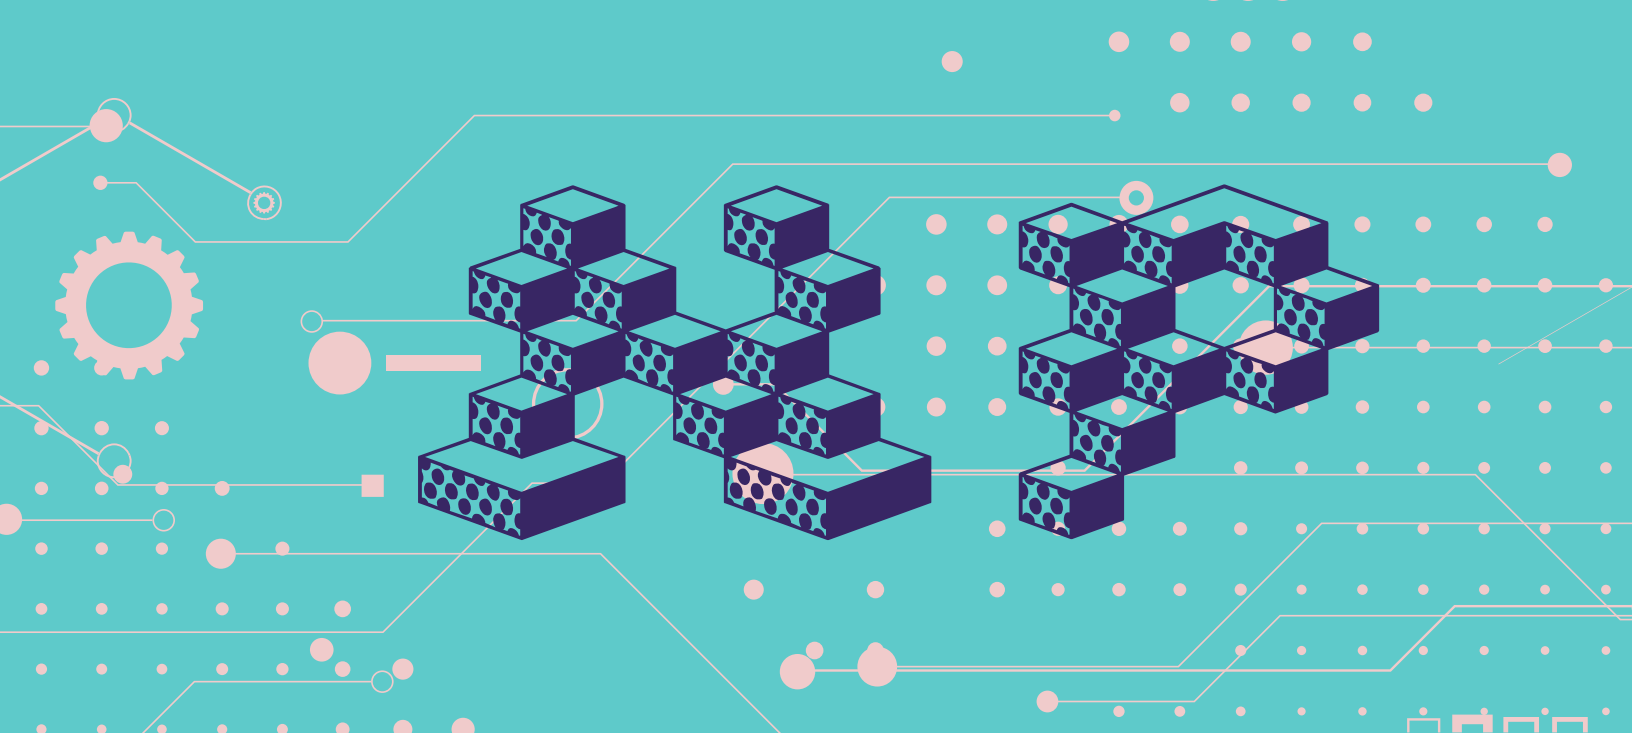
[8] Sandu Popescu et al., "Quantum Nonlocality as an Axiom," *Foundations of Physics* (1994).

[9] Yael Tauman Kalai et al., "How to Delegate Computations Publicly," *STOC* (2019).

[10] Arka Rai Choudhuri et al., "Non-Interactive Batch Arguments for NP from Standard Assumptions," *CRYPTO* (2021).

[11] Brent Waters et al., "Batch Arguments for NP and More from Standard Bilinear Group Assumptions," *CRYPTO* (2022); Arka Rai Choudhuri et al., "Correlation Intractability and SNARGs from Sub-Exponential DDH," *CRYPTO* (2023).

[12] Arka Rai Choudhuri et al., "SNARGs for P from LWE," *FOCS* (2021); Yael Tauman Kalai et al., "Somewhere Statistical Soundness, Post-Quantum Security, and SNARGs," *TCC* (2021).



Proof complexity enters the picture

A way forward came from an unexpected direction. Jain and Jin established a connection between cryptography and *proof complexity* (the study of how hard it is to prove propositional tautologies in formal proof systems) by showing that formal proofs of mathematical statements, written in the Extended Frege (EF) proof system, can serve as cryptographic material.[13] Jin, Kalai, Lombardi, and Vaikuntanathan brought this idea to bear on SNARGs for nondeterministic languages.[14] Consider a false instance $x \notin L$ of an NP language L. The claim that x has no witness is a propositional tautology: For every candidate w, the verification circuit rejects. Assuming LWE, they construct a SNARG for every NP language whose false instances have polynomial-size EF proofs of nonmembership.

Why do propositional proofs help? A propositional proof is a sequence of lines, each following from earlier ones by a simple rule, and the lines can serve as checkpoints. For a false instance x, the security analysis walks through the proof of nonmembership line by line, showing at each line that any proof the verifier accepts must be consistent with it. The last line asserts that the verification circuit rejects, and consistency with it contradicts acceptance. Crucially, the construction itself never sees the propositional proof; only the security analysis does.

This raises a natural question: Which languages have short propositional proofs of nonmembership, and how far can this idea reach toward all of NP? The answer is: not all the way, at least not as it stands. If every false instance of an NP-complete language had a polynomial-size EF proof of nonmembership, that proof would itself be a witness of nonmembership, and NP would equal coNP. Thus, EF proofs of nonmembership should not be expected for all of NP, and reaching NP requires relaxing the requirement.

The subsequent work of Jin, Kalai, Lombardi, and Mathialagan does exactly that.[15] Loosely speaking, the researchers replace EF proofs of nonmembership with a weaker object that one may think of as a *computational* EF proof of nonmembership. In such a proof, some intermediate lines need not proceed by logical rules; instead, they may involve nonlogical steps that invoke the security of a cryptographic primitive, such as the indistinguishability of a pseudorandom string from true randomness. The researchers use this approach to construct a *universal* SNARG for NP, which is secure as long as there exists some secure (possibly unknown) SNARG for NP. They were able to use this framework to build SNARGs for NP from witness encryption, a primitive weaker than iO whose existence is nevertheless known, at present, only from the same well-studied assumptions as iO.[16] Devadas, Hopkins, Kalai, Kothari, Lombardi, and Mathialagan used this framework to build SNARGs

[13] Abhishek Jain et al., "Indistinguishability Obfuscation via Mathematical Proofs of Equivalence," *FOCS* (2022).

[14] Zhengzhong Jin et al., "SNARGs Under LWE via Propositional Proofs," *STOC* (2024).

[15] Zhengzhong Jin et al., "Universal SNARGs for NP from Proofs of Correctness," *STOC* (2025).

[16] Sanjam Garg et al., "Witness Encryption and Its Applications," *STOC* (2013).

for NP under a mathematical conjecture.[17] And Jin, Lu, and Peng used it to construct SNARGs for UP, the class of NP languages with unique witnesses, assuming EF proofs of uniqueness.[18]

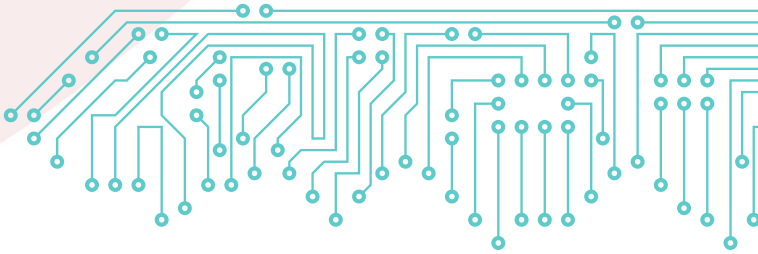
Cryptography from the difficulty of proving a theorem

A recent work of Hsieh, Jain, Li, and Mathialagan charts a different path to SNARGs for all of NP, through yet another connection with proof complexity.[19] Their starting point is an asymmetry: While EF proofs of nonmembership are not expected for all of NP, actually *proving that they do not exist* appears to be extraordinarily hard. Lower bounds against EF, and against strong proof systems generally, are among the central open problems of proof complexity. Razborov posed the challenge of proving such lower bounds, even by reducing them to familiar computational hardness assumptions.[20] And decades of work have not produced such a theorem.

The key insight of Hsieh et al. is to turn this difficulty into a source of cryptographic hardness. Recall that the SNARG of Jin et al. uses EF proofs of nonmembership only in its security analysis: The construction itself does not depend on them and can be run on any NP language.[21] Suppose, then, that an efficient adversary breaks its soundness at some false instance x. Under LWE, the security analysis implies that x has no short EF proof of nonmembership, for otherwise the adversary could not exist. In other words, an attack on the construction, together with the LWE assumption, yields a proof of an EF lower bound in ordinary mathematics, say in ZFC. This alone would be remarkable: An attack on a cryptographic scheme would deliver a theorem of exactly the kind Razborov asked for.

Hsieh et al. sharpen the observation considerably. Most cryptographic reasoning, with its hybrid arguments, reductions, and probability estimates, requires surprisingly little logical strength. The researchers show that the security analysis of Jin et al. can be formalized in APC1, a weak theory of bounded arithmetic designed to capture feasible reasoning. An efficient attack at a false instance x therefore yields a proof of an EF lower bound for x inside APC1, where such a proof would be more surprising still. Assuming that EF lower bounds are unprovable in APC1, no such attack can exist. The result is a SNARG for all of NP from standard cryptographic and derandomization assumptions, together with this unprovability assumption. Once again, propositional proofs, or here their absence, enter only through the security analysis.

The assumption has a desirable win-win character: Either it holds, and we have SNARGs for all of NP without obfuscation, or it fails in a sufficiently explicit way, which would be a major advance in proof complexity. More broadly, cryptography has always rested on the presumed *difficulty of computing* something. This work opens the door to cryptography that rests on the presumed difficulty of *proving* something: of establishing mathematical theorems in a formal system or, put differently, of ruling out claims that we believe to be false but cannot refute.



[17] Lalita Devadas et al., "SNARGs for NP and Non-Signaling PCPs, Revisited," *STOC* (2026). They conjecture that the exponential-size Hadamard PCP is statistically non-signaling sound.

[18] Zhengzhong Jin et al., "SNARKs from LWE via Non-Black-Box Reductions," *STOC* (2026).

[19] Yao-Ching Hsieh et al., "SNARGs for NP from Unprovability of Mathematical Theorems (or: How to Use the Simplicity of Cryptographic Reasoning)," *STOC* (2026).

[20] Alexander A. Razborov, "Pseudorandom Generators Hard for k-DNF Resolution and Polynomial Calculus Resolution," *Annals of Mathematics* (2015).

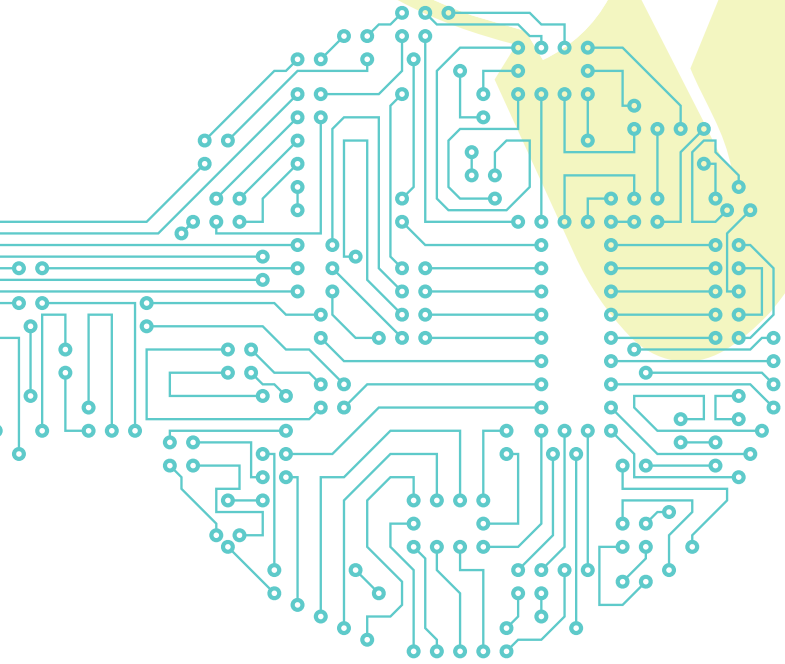
[21] Jin et al., "SNARGs Under LWE."

What comes next

The long-term goal of this line of research remains an efficient SNARG for all of NP from post-quantum assumptions such as LWE. We are not there yet. The unprovability-based construction rests on a new proof-theoretic assumption, and it achieves soundness only for statements fixed before the CRS.

Yet the landscape has changed. Obfuscation is no longer the only general route, and proof complexity and cryptography now interact in ways that were barely visible a few years ago, with traffic running in both directions. Cryptography has borrowed the objects of proof complexity — propositional proofs of nonmembership and of uniqueness — and put them to an unusual use, not as components of constructions but as instruments of security analysis. It has borrowed proof complexity's open problems too, treating the long-standing failure to prove Extended Frege lower bounds as a hardness assumption in its own right. In return, cryptography now poses questions of its own to proof complexity: Which natural families of tautologies admit short EF proofs, which formal theories can carry a security argument, and whether an attack on a concrete cryptographic scheme could ever double as a lower-bound proof. Each field now has reason to watch the other closely.

The broad lesson of this story is that succinct verification is not only about compressing a computation. It may also require compressing the *reason* that a cheating strategy cannot exist. By treating mathematical proofs — and our inability in some cases to produce them — as part of the cryptographic design space, we are beginning to chart routes to succinct proofs without obfuscation.



Researcher, Expositor, Friend: A Portrait of Dan Gusfield

by Kristin Kane (Simons Institute) and Tandy Warnow (University of Illinois Urbana-Champaign)

Daniel Gusfield is recognized as having played a founding role in the field of algorithmic computational biology. He was a long-term visitor in the Simons Institute's program on Evolutionary Biology and the Theory of Computing, during our first year, and has been a regular participant in workshops and programs ever since. He is also the first person we're aware of who has included the Simons Institute in his estate planning.

Now a Distinguished Professor Emeritus of Computer Science at UC Davis, Dan began his career with a focus on combinatorial optimization. His most impactful paper in this area shows how to compute the famous Gomory–Hu tree for multiterminal network flow by adding only five simple lines of code to any program that computes a single-source, single-terminal network flow.[1] His first forays into computational biology were in the 1980s, at a time when there was really only a handful of other computer scientists involved (e.g., Dick Karp, Zvi Galil, Gene Myers, and Gene Lawler). Since then, he has gone on to make seminal contributions to algorithms and theory for comparing biomolecular sequences (DNA, RNA, protein), inferring evolutionary trees, haplotyping, and a host of other core problems. His most highly cited paper provides a linear-time solution to the two-state perfect phylogeny problem, which addresses the question of how to infer evolutionary trees (phylogenies) from traits that have the special property that they arise only once and then are never lost.[2] This problem is relevant to understanding tumor evolution, human population genetics, and language evolution, and his result has made a substantial impact in these fields.



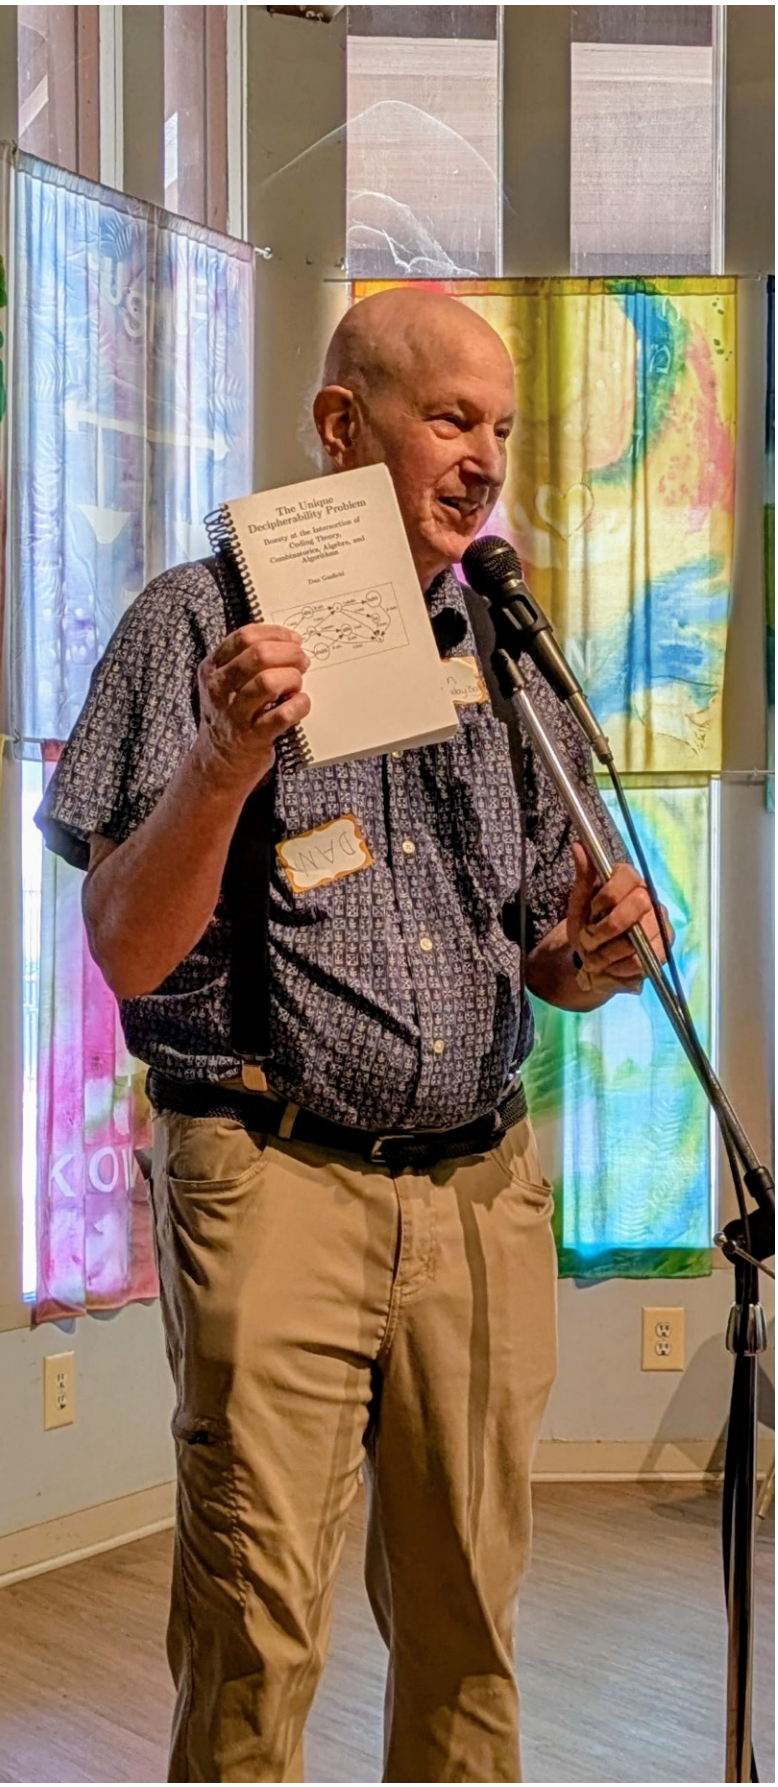
Gusfield's 75th birthday celebration (2026)

For his many research accomplishments, Dan has been elected a Fellow of the IEEE (2015), the International Society for Computational Biology (2016), and the Association for Computing Machinery (2017). He was also the founding Editor-in-Chief of the *IEEE/ACM Transactions on Computational Biology and Bioinformatics*.

Dan's generosity to the Simons Institute and other parts of the University of California is grounded in lifelong connections. Aside from a freshman semester at the University of Illinois and a short period early in his career, Dan has spent his entire adult life in the UC system. After earning a bachelor's at UC Berkeley and a master's at UCLA, he returned to Berkeley for his PhD, where his advisor was Dick Karp (who later became the Founding Director of the Simons Institute). From there, Dan went on to an assistant professorship at Yale for six years, before settling in for good at UC Davis.

During grad school, Dan sublet EECS Professor Beresford Parlett's house in the Berkeley Hills with fellow graduate students Mike Sipser, Barbara Simons, and Bruce Char. "And because it was this giant house," he recalls, "it had lots of unused bedrooms. So we became sort of the unofficial hotel of the computer science department. I had left UCLA because it was dead at night, and came back to Berkeley, where there were exciting colleagues and it was just lively all the

[1] Dan Gusfield, "Very Simple Methods for All Pairs Network Flow Analysis," *SIAM Journal on Computing* (1990).
[2] Dan Gusfield, "Efficient Algorithms for Inferring Evolutionary Trees," *Networks* (1991).



time.” Simons Institute Director Emerita Shafi Goldwasser was one of those exciting grad student colleagues, during what was by all accounts a very special time to be in computer science at Berkeley.

Years later, Dan found that same camaraderie and energy at the Simons Institute. “Up on the second floor, there’s just this constant buzz of people working together at the whiteboards, and then, of course, breaks with coffee and delicacies in the afternoon. It was very stimulating, very exciting. And, you know, the Simons Institute has been very successful, the way it was intended to be and the way it was designed.”

Dan’s giving to the Simons Institute was inspired by his recognition that while he cares deeply about democratization and making educational opportunities available broadly, he wanted to make a special gift to an institution where some of the most productive researchers in the world can do their best work.

Alongside research, another of Dan’s passions is expository science writing. He has written five books providing technical depth and clarity on a wide range of topics, and is currently writing two more. His most well-known contribution in this vein is *Algorithms on Strings, Trees, and Sequences: Computer Science and Computational Biology* (1997), which provides an introduction to computational biology through the eyes of a discrete algorithms researcher, and has been cited more than 8,500 times.

“I think my interest in science explanation really comes out of my need to explain it to myself — because I was in the slow math group in high school,” Dan chuckles. “I need a lot of explanation.”

An avid hiker, Dan has spent a nontrivial portion of his life trekking in the Dolomites, the Himalayas, and the Sierra Nevada, where he has a cabin. And as an ardent admirer of Vincent van Gogh, he has traveled the world to see his works in person.

“Dan is smart, adventurous, and funny,” relates Barbara Simons. “But most of all, he is kind, generous, and thoughtful. I feel very fortunate to have him as a friend.”

Dan celebrated his 75th birthday this summer.

This Year's Research Programs

by Jason Lee (Simons Institute)

The 2026–27 academic year at the Simons Institute brings together three programs in core theoretical computer science (TCS) and its connections to mathematics. In Fall 2026, two programs with synergistic overlaps — one on Pseudorandomness & High-Dimensional Expansion, and one on Spectral Theory Beyond Graphs — run concurrently, with some shared workshops, including a joint boot camp. In Spring 2027, a program on Symmetry in Efficient Computation with Local Constraints studies constraint satisfaction and related problems where a certain algebraic structure often dictates if and how a problem is efficiently solvable. The broad and powerful theme of how the local structure and constraints of a system can influence its global properties and behavior manifests itself across these three programs.

Pseudorandomness & High-Dimensional Expansion

Random bits are one of the most useful resources in algorithm design, but true random bits are difficult to obtain from physical processes, which can have imperfections and complex correlations. Pseudorandomness is the subfield of TCS that asks when deterministic processes can behave like or mimic true random processes. In the past 40 years, this viewpoint has led to pseudorandom number generators (what your computer uses when you call rand), randomness extractors, discrepancy measures, error-correcting codes, and expander graphs. The applications of pseudorandomness span cryptography, complexity, and algorithms. However, various foundational questions in the field remain open, such as whether every efficient randomized algorithm has an efficient deterministic counterpart (BPP = P), or whether space-bounded computation can be derandomized (BPL = L).[1]

Expander graphs are sparse graphs that nevertheless have strong connectivity and rapidly mixing random walks, and efficient deterministic constructions are known. High-dimensional expansion generalizes the idea of expansion in graphs to hypergraphs and simplicial complexes. Similar to expanders, high-dimensional expanders are deterministic and sparse, with the kind of random-like behavior that algorithms need. However, they are significantly more elusive — while random graphs are expanders and thus their existence is easy to establish, this is not the case for high-dimensional expanders. The construction of high-dimensional expanders relies heavily on several interconnected fields of mathematics, including combinatorics, group and representation theory, number theory, arithmetic geometry, and topology.

Pseudorandomness and expander constructions have played a crucial role in a number of applications across computer science, including breakthroughs on locally testable codes and quantum error-correcting codes, faster sampling algorithms, and probabilistically checkable proofs. This program brings together disparate communities that work on the same themes of explicit constructions, derandomization, and applications. The primary questions they're studying: Which pseudorandom properties of graphs have high-dimensional counterparts, how can we construct objects achieving them, and what new algorithms and applications will they enable?

Spectral Theory Beyond Graphs

Spectral graph theory studies a graph through its Laplacian matrix. The eigenvalues and eigenvectors of the graph Laplacian reveal whether the graph has bottlenecks, indicate how quickly a random walk mixes, and help identify communities. Over the past few decades, spectral graph theory has become a basic tool in theoretical computer science.

However, graphs are only one setting in which spectra carry geometric and probabilistic information. Closely related Laplacian operators appear in manifolds and groups, with connections to mathematical physics, random matrix theory, number theory, and operator algebras. The Spectral Theory Beyond Graphs program is motivated by the observation that these neighboring subjects have many common themes, so techniques can diffuse between subjects.

[1] This was recently resolved in the affirmative by OpenAI, and is now the subject of a working group in the Pseudorandomness & High-Dimensional Expansion program.



One theme is using local information to characterize global behavior. For example, recent work on fast sampling algorithms using Markov chains shows that rapid mixing, which is a global property of random walks, can be proved by utilizing only local neighborhood information. Recent prominent ideas such as spectral independence, stochastic localization, curvature, and high-dimensional expansion give different versions of this local-to-global principle. The program's overarching theme is to investigate whether these theories share an underlying fundamental principle, and when one theory can be used to strengthen the others.

A second theme is randomness. Random graphs are a primary test bed for spectral graph theory, but strikingly similar questions arise for random manifolds, random operators, and random matrices. Recent progress in free probability and matrix concentration has already fed back into theoretical computer science, enabling results that seemed out of reach. Similarly, techniques developed for random graphs are transferring to the spectra of random surfaces.

Finally, the program also moves beyond the graph Laplacian to richer matrices such as Schrödinger operators, Kikuchi matrices, and other matrices arising in algorithms that exhibit spectral behavior that the usual Laplacian does not capture. Questions about whether eigenvectors spread out or localize may have underexplored algorithmic consequences.

The broader objective is to identify shared underlying spectral principles behind these different areas, and investigate the extent to which the techniques developed for one setting can transfer to another.

Symmetry in Efficient Computation with Local Constraints

Constraint satisfaction problems (CSPs) specify a set of local rules, where satisfying each rule is simple, and they ask for global solutions that satisfy all the local constraints (or as many of them as possible). A classic example is graph coloring, where the edges stipulate pairs of vertices that cannot receive the same color, and the goal is to find a coloring that respects all the edge constraints. The modern theory of CSPs is centered on the observation that the computational tractability of a CSP is often governed by symmetries called polymorphisms. A polymorphism is a mapping that takes solutions and transforms them into another solution. Such closure properties prevent the solution space from having arbitrary combinatorial structure. As an example, linear polymorphisms lead to efficient solutions found via Gaussian elimination.

Different polymorphisms expose different kinds of regularity that algorithms can exploit — for example, linear structure can lead to Gaussian elimination, while majority-type structure can support local-consistency algorithms. For finite-domain CSPs, this viewpoint ultimately led to a complete classification of which families of constraints are solvable in polynomial time and which are NP-hard, governed in a precise sense by the existence or lack of sufficiently rich polymorphisms.

This program seeks to extend this viewpoint from exactly solving worst-case CSPs to promise, random, approximate, and quantum CSPs. The program also has a boot camp, as well as three focused workshops, each on one of the latter three themes.

Approximate CSPs focus on the case when it is computationally intractable to satisfy all the constraints, and instead seek to satisfy the largest possible fraction. The famous max-cut problem has a semidefinite programming (SDP) relaxation that finds a cut with value at least 0.878 of the optimum, and the same SDP framework gives an optimal approximation for all max-CSPs under the unique games conjecture.[2] Promise CSPs focus on a different form of approximation where the local constraints must all be met, but one can do so under a relaxed form of the constraining rule. A quintessential example is approximate graph coloring, where the goal might be to color a three-colorable graph using a larger number of colors.

Random CSPs probe how the geometry of the solution space changes under different densities of constraints, and the phase transitions of when algorithms can find solutions or fail.

Finally, quantum CSPs replace the discrete assignments with a quantum state, and the goal is to find the global minimal energy state of the local Hamiltonian. The quantum PCP conjecture asserts that it is computationally hard to find even a constant factor approximation to the minimum energy, analogous to how the classical PCP theorem implies hardness of approximation of CSPs.

This program will assemble a diverse group of researchers with complementary expertise to make progress on these challenges, both broadening and deepening our understanding of the role of symmetries across these facets of CSPs.

In November 2026, the Simons Institute launched a Research Pod on AI and Theoretical Computer Science, which will run alongside these programs, along with ongoing pods on Machine Learning, Quantum Computing, and Resilience.

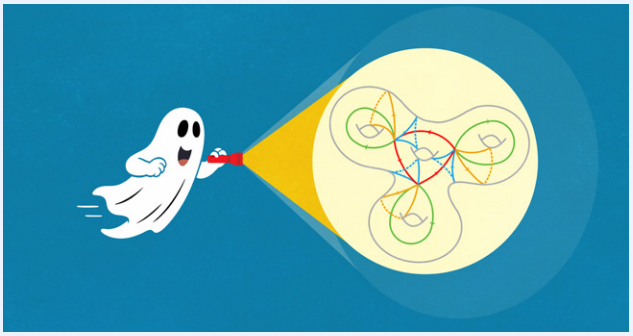
[2] The unique games conjecture was another one of the major problems resolved in the October 6 OpenAI drop.

Current & Upcoming Programs

Pseudorandomness & High-Dimensional Expansion
(Fall 2026)

Spectral Theory Beyond Graphs
(Fall 2026)

Symmetry in Efficient Computation with Local Constraints
(Spring 2027)





A Quantum Sprint
(Spring – Summer 2027)

Diffusion Generative Modeling
(Fall 2027)

Streaming and Sketching: Foundations and Frontiers
(Spring 2028)

This list is current as of November 2026 — check the website for updates.

Support for the Simons Institute

Make a gift

The Simons Institute for the Theory of Computing gratefully acknowledges the individuals and institutions that provided us with financial support during 2025–26:

Gagan Aggarwal and Parikshit Gopalan
Amy Ambrose and Koyo Konishi
Anonymous (various)
Apple
Sepehr Assadi
Amik Behera
Richard Belew
Samy Bengio
Marshall Bern
Clément Canonne
Jennifer Chayes and Christian Borgs
Jingyuan Chen
Kevin Chen
Debasish Das
Anirban Dasgupta
Dimitrios Diochnos
Jiechao Feng
Yuval Filmus
Fujitsu
Google
Google.org
Venkatesan Guruswami
Zaid Harchaoui
Kangmin Hu
Ovid Jacob and Gitta Salomon
Jane Street
Ce Jin
Sampath Kannan
Richard M. Karp and Mary Jo Lazear
K. N. King
Manchor Ko
Lorraine Ling and Paul Ling
Guanghui Liu
Tse Ka Man
Stephen Mitchell
Thomas Moore
OpenAI
Tu Ouyang
Project CETI
Anand Raghavan
Mayank Rathee
Sunny Sadana
Brian Shiratsuki
Simons Foundation & Simons Foundation International
Clifford Stein
Shobhana Murali Stoyanov and Tsvetan Stoyanov
Val Tannen
Éva Tardos
David Tennenhouse and Orna Resnekov
Sennur Ulukus
Vinod Vaikuntanathan
Dan-Virgil Voiculescu and Ioana Voiculescu
Avi Wigderson
Ryan Williams
Honam Wong
Jae Oh Woo
Jiaqi Yang
Bin Yu and Ke-Ning Shen
Fan Pu Zeng

All gifts and partnerships are matched one to one by Simons Foundation International.

To become an industry partner or learn more about engagement possibilities, contact Director Venkatesan Guruswami (simonsdirector@berkeley.edu) or Senior Development Director Amy Ambrose (amyambrose@berkeley.edu).



Simons Institute for the Theory of Computing
121 Calvin Lab #2190
UC Berkeley
Berkeley, CA 94720-2190

To request this document in an alternative format, including large print, braille, or electronic formats, please contact Barry Bödeker at barry.b@berkeley.edu