SIGecom Winter Highlights 2026

SARA FISH
Harvard University
and
VRINDA PAREEK
Columbia Law School
and
RUI-JIE YEW
Simons Institute for the Theory of Computing, UC Berkeley
Brown University

Sara Fish is a PhD Candidate in Applied Mathematics at Harvard University and Graduate Fellow at the Harvard Kempner Institute. Her research spans diverse topics at the intersection of EconCS and artificial intelligence, including algorithmic collusion, AI-enhanced democratic processes, and AI evaluations. Previously, she earned her B.S. in Mathematics from Caltech.

Vrinda Pareek holds an LL.M. degree from Columbia Law School. Her coursework focused on technology law, including copyright law and the regulation of artificial intelligence. She is an India-qualified lawyer, where she earned her B.A. and LL.B. degrees.

Rui-Jie Yew is a Law & Society Fellow at the Simons Institute for the Theory of Computing and a PhD candidate in computer science at Brown University. She conducts research on AI policy and incentives, alignment, and avoision in regulation. She holds an S.M. from the Massachusetts Institute of Technology and a B.A. from Scripps College.

The sixth annual ACM SIGecom Winter Meeting took place on 25 February, 2026. Organized by Rachel Cummings and Inbal Talgam-Cohen, the meeting explored topics at the confluence of Artificial Intelligence, Law, and Algorithmic Game Theory. We share some highlights from the 2026 Winter Meeting. Thanks to the organizers and all the participants!

The meeting kicked off with opening remarks from **Michal Feldman**. As part of these remarks, Feldman discussed topics of AI safety, AI regulation, and incentive alignment towards the prevention of harm and provided an overview of the meeting's structure. This year's meeting took the form of a virtual workshop and included talks, panels, and fireside chats with experts. These segments included presentations on motivating tne field of CS&LAW, discussion on published work, as well as discussion on future research ideas directions.

1. AT THE INTERSECTION OF COMPUTER SCIENCE AND LAW (DANNY WEITZNER)

As part of the Introductory Talks, we hear from Daniel Weitzner as the presenter and Kobbi Nissim as the discussant about what it means to do research at the intersection of Computer Science and Law.

Daniel Weitzner is a founding director of the MIT Internet Policy Research Initiative and is a Principal Research Scientist at MIT's CSAIL. He leads technically grounded research to inform critical internet public policy and data governance

research. He has worked at the National Telecommunications Information Agency, the Center for Democracy and Technology, the Electronic Frontier Foundation, and more. As part of this talk series, Weitzner discusses what it means to do research at the intersection of CS&LAW

Kobbi Nissim is a Professor of Computer Science at Georgetown University and an Affiliate Professor at Georgetown Law. Before that, he was a professor of computer science at Ben-Gurion University, and he spent several years at Harvard University in the Center for Research and Computation and Society. His work establishes rigorous practices for privacy in computation, including being one of the inventors of differential privacy. He also works at the intersection between privacy law and policy.

This talk is based on *At the Intersection of Computer Science and Law* [Feigenbaum and Weitzner 2026], a paper that he and **Joan Feingenbaum** submitted to ACM CS&LAW. In the paper, they lay down the motivations for launching the conference, which is now roughly in its fifth year. In this talk, there are two main problems that Weitzner presents as motivation for the studying of CS and law together:

(1) There are gaps in addressing the legal and policy challenges information systems have presented us.

(2) On the other hand, there is a lack of understanding in the building of governable systems or when systems are governable.

These are problems in both technology policy and system design that Weitzner has worked on and that a number of colleagues in CS&LAW have worked on that are still not particularly well-addressed. He discusses some research that answers to some of these gaps.

1.1 Starting from Law

This class of research starts from how can you tell if system architecture satisfies legal requirements. Some examples of this class of work include [Cohen and Nissim 2020], which defines the formal property of "predicate singling out" in understanding the notion of singling out in the GDPR. Another work includes [Meding and Sorge 2025] shows that a number of editing features will blur lines between editing and manipulation.

1.2 Starting from CS

This class of research starts from observations about properties of systems which may give rise to behaviors or functions that require some regulation/control. One example here includes the question of how courts should apply principles to modern encryption systems. Here, [Cohen et al. 2022] provides formal definitions that could help courts determine when a specific requirement to decrypt communication would run afoul of this constitutional protection. Another example of research that falls into this category is [Lee et al. 2024] which asks the "right questions" in considering where liability ought to be assigned in AI production under copyright law.

1.3 Broader Examples

Weitzner also discusses broader examples of research that may not fit neatly in either of the above two categories. For example, [Chan et al. 2025] tackles the question of balancing privacy and accountability in misconduct settlements and shows how use of cryptographic techniques like zero-knowledge proofs and private computation can help strike balance between privacy and accountability.

1.4 Intersection Rules

Finally, Weitzner discusses what it means to conduct "scholarship that makes a contribution in the intersection, as opposed to the union, of computer science and law": (1) **takes both disciplines seriously**, (2) **could be "publishable" contribution in both or either field**, (3) **but publishable value could be solely the intersection**.

Going forward, he hopes this area of research will help policymakers apply more rigorous system applications and conversely that it will help system designers build more governable systems. He hopes researchers continue to study the fit between law and CS with an eye toward effective laws and efficient systems.

1.5 Response: Scope and Mission of CS&Law: (Kobbi Nissim)

Nissim starts by motivating the urgency involved in answering questions at the intersection of the two fields. There are a growing number of decisions of legal consequence are made in sociotechnical systems and there are questions surrounding the scale of having to review the decisions made by sociotechnical systems. And, even if only a small fraction of the decisions made in sociotechnical systems required human review, they would quickly overwhelm our judiciary or administrative systems. When he first got started in this field, he noticed the two fields (CS and law) having: (1) **same words, different meaning**: PII, composition of datasets and (2) **different reasoning**: "and" meaning as a logical connector, **different/conflicting values and goals**: CS theory values precision while law values flexibility.

Nissim concludes with hope regarding the identification of paradigms for bridging between the two disciplines, while respecting their differing approaches and values. He looks forward to seeing the paper which does not satisfy intersection rule, but which does belong to CS&LAW.

2. CS&LAW MEETS ALGORITHMIC GAME THEORY (KATRINA LIGETT)

Katrina Ligett is a professor of computer science at the Hebrew University, a director of the Federal Center for the Study of Rationality, and a visiting associate at Caltech. She's one of the founders of the field of CS and law and led one of the first workshops at the Simon Institute back in 2019, which brought together experts from theoretical computer science, law, economics, and the social sciences more broadly. In this talk, Ligett catalogs questions that may be of interest to the triangle of CS, Law, and Econ. How do we walk into this simplex of disciplines? What are the questions and tools this can raise for us? Ligett starts from the less developed CS-law edge in the simplex and considers the value that a lens of algorithmic game theory (AGT) can bring.

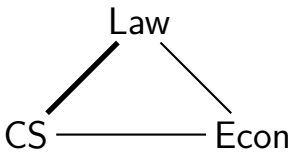

Fig. 1. Simplex of law, computer science, and economics with the CS-law edge highlighted.

Combing through the topics as part of the call for papers at the ACM CS&LAW conference [Feigenbaum and Weitzner 2026] with AGT "glasses", she identifies the following themes: **(1) Incentives**, **(2) Information**, **(3) Contracts**, and **(4) Attribution**. She presents the following research questions related to these themes:

2.1 How can AI and humans collaborate?

This body of questions concerns settings where you have different parties interacting. Here, *asymmetry* is a recurring theme and there are a variety of different asymmetries. Some examples include big tech companies and those regulating them, companies vs those using them, etc. These challenges are amenable to people with a toolkit (e.g., in economics and AGT) to model these asymmetries. An example of work concerning this topic is [Collina et al. 2026].

2.2 What *can* AI produce? How can we incentivize richness in the output of AI?

These types of questions relate to learning theory and results from learning theory about the limitations of the ability to recognize a language from samples is a computationally hard problem. [Kleinberg and Mullainathan 2024] show that it is possible to generate from a language given access to samples from that language. But there are also interesting challenges that are sort of implicitly raised here. Just because you can generate more samples from a language doesn't mean you're actually representing the full sort of breadth of that language. So, how do you potentially get better representation over the sort of the limits of getting richness in the outputs of AI?

2.3 How to incentivize (human) creative work?

This set of questions concerns the certification of creative work. A world where we can generate content cheaply and at scale raises a whole host of questions about how would you recognize creative work, how could you certify creative work, what is authorship, what rights should authors have in their work, how do you protect the rights of authors in their work. The existing infrastructure that we have for incentivizing human creative work, e.g., the copyright/patent family of tools on the regulatory legal side, seem to be insufficient to recognize and promote the things that we (potentially) want to promote in today's world. But there are a huge number of questions here, and some of them very quickly start to look like AGT questions: questions about what credit belongs to previous works, how do you identify intellectual contributions as people are building on each other's ideas. There are various initial works that chip off various pieces of these questions while echoing things that are familiar to us, like Shapley values.

2.4 How to balance freedom of expression with protections?

Now that we can produce lots of content very cheaply now with AI, we can also produce a lot of false, extremely convincing, personalized speech. Ligett asks the question here of: if we value meaningful, positive, constructive, useful, correct information and discourse, how can we promote that? What are the mechanisms that can be put in place? And again, there are questions about incentives, there are questions about recognizing this, there are questions about how do we build the systems that can promote this and protect this and ensure that there's room for expression that maybe we still value.

2.5 How can we ensure that AI companies' incentives align with societal objectives?

Stochasticity in modern machine learning and company incentives can make naively auditing AI safety difficult. There are also problems surrounding visibility. Most individual interactions with AI are not directly made public. Example works concerning this topic include [Saig et al. 2024; Singh et al. 2026; Hadfield-Menell and Hadfield 2019].

2.6 Who/what is responsible for what?

This surrounds responsibility for speech/responsibility for actions. When something goes right/wrong who gets the credit/is held liable? The tools that we have to hold entities liable are somewhat limited. A related question here is: how can LLMs avoid violating copyright and attribution/credit/holding actors liable.

2.7 How to track/allocate trust in information?

Now, we are flooded with fake information, and so we are going to need some new infrastructure to help us understand how to get access to information we value. Some infrastructure we are missing is infrastructure that allows us to build networks of trust to make sense of claims/events/facts.

2.8 When algorithms are used for pricing: what's harmful, what's helpful, and what's legitimate?

We're going to see and we're already seeing personalization of prices. We have to decide what are the limits, what's acceptable, what's unacceptable. And then it raises all sorts of questions about how can we enforce, how can we detect, how can we ensure, how do we build incentives. There are a lot of concerns around the ability of AI, again, to be extremely personalized and extremely convincing, but also potentially to trigger our biases and take advantage of us in various ways. AI is also going to be in a position to potentially have access to much more information than individual humans, and potentially even individual humans who have their weaker AIs working on their behalf. So we have to be concerned about taking advantage of information deficits that consumers might have. We have to ask ourselves questions about when is personalization beneficial for social welfare, for fairness, for efficiency? What do we want to accept? What do we not want to accept? There are algorithmic and definitional questions here, as well as questions concerning the translation of sort of more abstract concepts that might come from ethics of the law or policy and trying to figure out what do they mean for us at a

technical level and how can we realize them. Example of works concerning this topic and algorithmic price discrimination in particular include [Bar-Gill et al. 2023].

2.9 How to define/detect/prevent algorithmic collusion?

There are really interesting questions around as these algorithms are often operating in wider and wider ranges of decision-making, we have a lot of potential for them to collude. They're going to be interacting with each other. They're going to be interacting with complex systems. And so there are a lot of questions about how do we define, how do we detect, how do we prevent algorithmic collusion? What does this look like? In particular, one tension between CS and law is that of intent in collusion. Example works concerning this topic include [Hartline et al. 2024a] [Fish et al. 2024].

2.10 Incentives: How to nudge the legal system (and the world) to adopt beneficial technologies?

So some examples are things like zero-knowledge proofs [Bitan et al. 2022], where we have a lot of technical capabilities that can do a lot of good, and that can improve a lot of pretty broken systems that we have. These questions surround filling that last gap between theory making things work in the real world.

2.11 In what sense does the law learn?

Common law has this concept of precedent where we sort of look to decisions that have been made in the past, and these help inform basically what the rules are today. But this can be framed as a really sort of interesting type of a learning problem where there's a decision to be made about whether to bring a case, and that decision encompasses a bunch of considerations, including possibly sort of the cost of bringing the case, the likelihood of succeeding. But if there is an interesting case that is never brought, then we never get precedent, and in some sense the law doesn't learn. Sometimes part of the incentive to bring a case is actually to set precedent. And so there's been some work on sort of this learning kind of algorithmic take, which has a kind of law and economics vibe in the sense that sort of trying to model the system that is the law [Hartline et al. 2022][Dutz et al. 2025]. Ligett also points out the computational and algorithmic flavor to this learning problem.

3. TALIA GILLIS

Talia Gillis is a professor of law at Columbia Law School, and studies the law and economics of consumer markets. She is interested in household financial behavior and how consumer welfare is shaped by technological changes. She is a recipient of the 2022 junior faculty grant and the Richard Paul Reichman Center for Business Law and Public Policy Grant. She completed her SJD degree at Harvard and a PhD in economics and clerked for Deputy Chief Justice Mer at the Supreme Court of Israel.

Gillis addressed law's (and lawyers') discomfort with ambiguity, and posited that regulatory ambiguity may, in fact, be an important strategic tool. Gillis examined the complex relationship between legal standards and algorithmic systems, focusing on the persistent problem of regulatory ambiguity. While she initially sought

to eliminate vagueness in laws governing consumer finance and anti-price discrimination, her views evolved through her engagements with policymakers. Now she thinks of this imprecision as a strategic tool used by regulators. Gillis argued that this flexibility in legal standards or regulatory approaches prevents firms from "discrimination hacking" or gaming fixed regulatory targets, and it allows regulatory oversight to adapt to rapidly evolving technological changes. By framing legal design as a mechanism design question, Gillis suggested that the optimal level of precision desirable in regulation depends on the specific strategic environment and the incentives of the actors involved. Ultimately, her work highlighted the tension between the desire for clear compliance rules and the necessity of maintaining enforcement discretion to prevent substantive harms in automated markets. Gillis encouraged participants to think of ambiguity in legal standards not as 'sloppy' drafting, but as a deliberate decision to preserve flexibility, deter gaming by regulated actors, or maintain enforcement discretion.

4. PANEL ON GENERATIVE AI, AGENTIC AI, LAW, AND CS

The 2026 Winter Meeting featured a dynamic panel discussion on GenAI, Agentic AI, Law, and CS, featuring four panelists with rich experience across economics, law, and computer science.

Vince Conitzer is a Professor of Computer Science at Carnegie Mellon University, where he directs the Foundations of Cooperative AI Lab (FOCAL). He is also Head of Technical AI Engagement at the Institute for Ethics in AI, and Professor of Computer Science and Philosophy, at the University of Oxford. He began his PhD in 2001 under Thomas Sandholm and has been an active member of both the EC and AI research communities throughout his career.

Peter Henderson is an assistant professor at Princeton University, with appointments in the Department of Computer Science and the School of Public and International Affairs, as well as the Center for Information Technology Policy. He received his J.D. from Stanford Law School and PhD in computer science from Stanford University.

Zoë Hitzig is a Junior Fellow at the Harvard Society of Fellows. She received her PhD in economics from Harvard in 2023. Between 2024 and 2026 she was Research Scientist at OpenAI, working on their safety research team and also working on their economic research team.

Georgios Piliouras is a senior staff research scientist at Google DeepMind, where he leads the Game Theory team, and Associate Professor in the Engineering Systems and Design Pillar at the Singapore University of Technology and Design (SUTD).

Below are selected questions and panelist responses, edited for clarity and brevity.

What do you see as the most pressing research questions in generative and agentic AI that can be addressed by the EconCS and CS&Law communities?

Georgios Piliouras: Sometimes we use syntactically similar languages to talk about rather different types of problems. For example, contracts. In CS, we talk about algorithmic contract theory, and we have some very well-designed problems in that space, on which we have made reasonable progress. But also, there are the

types of contracts that lawyers develop in practice. These two things are rather different. I would be interested to see if there are possible connections to make in this space.

Zoë Hitzig: I think that one of the most important questions for this community to be thinking about is data governance and property rights for data. Right now, the conversation about who owns what data and how different kinds of data get used in large AI systems has been co-opted by, first of all, the big labs who have an incentive to talk about it in a certain way, but also to some degree by lawyers who have one main apparatus for thinking about property rights for data, which is copyright. But as economists, we understand copyright from first principles: you want some mechanism like copyright so that people are incentivized to produce valuable stuff. So I think one of the most important and understudied questions right now is: what are the incentives that people have to produce the data that goes into large machine learning systems? And how do different mechanisms and institutions change the landscape and incentives that people had to produce that data in the first place? I think there are ways to talk about this both on a macro scale and on a micro scale.

Vince Conitzer: In my lab, we're thinking a lot about the interactions between AI systems, especially from a strategic perspective. We also think a lot about the ways in which AI agents are different from human agents and what that allows you to do. For example, an AI agent can be run in a simulation to see what it will do—you can't quite do that with a human. That gives different ways to achieve trust and cooperation. Finally, related to Zoë's points on incentives for content production, a very difficult—yet often legally significant—question is: can we understand why a language model produced any given output?

Peter Henderson: There are infinitely many things to work on right now because of how fast the field is moving. But I'll give three examples of things I consider high priority. First, the problem of specifying natural language rules and standards for models to follow. You can think about this from an economics perspective (e.g. contracts), from a legal perspective (e.g. statutes), or from a computer science perspective (e.g. reward specification). I think this is an overlooked problem, but one where all three fields can come together to make progress. Second, how do we reshape the law when it comes to copyright and fair use? This will require both a technical understanding of the models and how they're trained, as well as economic questions of impacts on the market and the value of the labor. Third, what are the right policies and legal tools by which to tackle potential labor market impacts? This again requires all three fields to come together to understand: How well are the models actually performing, so we can better project what might happen? What are the potential macroeconomic and microeconomic effects? And then, what are the legal and policy options that we have that we can actually intervene at a reasonable timescale?

As we approach AGI—that is, AI that can match or surpass human capabilities across virtually all cognitive tasks—there is a tension between utilizing this emerging technology, and protecting human jobs and expertise. What do you see as our way through this?

Vince Conitzer: When we should keep humans in the loop is a difficult question.

I think that it varies across different professions. You could make a good argument in law, that maybe there's a special reason to have humans still be part of it. More broadly, if AI doesn't advance too fast in the near future, maybe one of the main problems is disruption of hiring pipelines. We need people at the senior level who really understand their profession, but they won't have a path to get there if the entry-level work is done by AI. So maybe that's another reason to keep people in the process—to understand what the processes are.

Zoë Hitzig: Hayek's key insight in *"The Use of Knowledge in Society"* was that we need people to have decision-making power on the fringes. We can think about this in terms of two theses. First, things need to be decentralized because people have tacit, immovable knowledge in their heads or because of where they are physically. Second, even if we could move all the information to some central source, that source wouldn't be able to process all of it. As Eric Brynjolfsson and I have argued [Brynjolfsson and Hitzig 2025], these two premises are eroding. AI is finding new ways to make knowledge that you couldn't codify before codifiable, and it's taking away the constraints on information processing. If we take that seriously, the economic reasoning behind decentralized markets is going away—centralization might be economically efficient. And that might not be great, because part of how people had bargaining power in the economy was that they had knowledge that couldn't be taken from them. So the question becomes: what can we do to give people bargaining power in the new economy? What are the countervailing institutions that we can build now for this completely new kind of economy?

If something goes wrong with AI, who is legally liable? The AI developers, the users or somewhere in between?

Peter Henderson: Even defining what systems are in scope is challenging—for example, autonomous vehicles and systems that act in the real world versus simple systems in video games or apps that might have very different types of downstream harms. To my mind, one of the most interesting short-term areas where we're going to have to figure this out is cybersecurity harms. It's pretty clear that these models can already be repurposed for third-party harms pretty easily, and on the technical safeguards question, it's pretty hard to identify these misuses in real time at scale. Now, if we look to past cybersecurity cases, we actually don't have that much litigation against providers of tools that already provide software to penetration test or to run attacks that can be used both offensively and defensively. I don't think there's going to be this beautiful framework in which we figure out how to allocate exact liability in every scenario. I think in practice it gets really messy really quickly, and we'll just have to wade through many of these cases. In the cybersecurity setting, for example, we do have duty of reasonable care standards and negligence standards that may well cover many of these scenarios and attribute liability in a perfectly reasonable way. I don't know that we need to go in and completely swap in a new framework when we have yet to work through the existing frameworks that may well apply already.

The EC community has had so much interaction with online ad auctions and allocating the resource of attention. Zoë, how might this relate to your remarks on bargaining power?

Zoë Hitzig: I think there is some way of creating the right institution that

gives people power over how they spend their attention. If we think about how we came to be auctioning off people's attention in the first place, it was a property rights issue—when Google said, hey, those search histories are mine. So there is this connection between the data rights issue and who has control over their own attention. As soon as you've given over tons of data, you have opened yourself to a certain kind of manipulation or advertising. There's an amazing book out called *Attensity* that is laying out a way to bring back our attention—it's not in our language, but I think it's a book that we might all be able to learn something from.

Vince Conitzer: Related to Zoë's response, there's an article by Marietje Schaake in the Financial Times about what she calls "botlash." A phrase associated with that is "resist and unsubscribe." The "unsubscribe" part really gets at using your attention for bargaining power.

Doesn't AI make cybersecurity more robust, since we can potentially find exploits before criminals do?

Peter Henderson: The key challenge is that defenders can use AI models to try to patch vulnerabilities, but each patch takes time to deploy. If both parties have equal access to the same model, there's a time delay to patching all the vulnerable systems, whereas the offensive side can start attacking right away. So if there was privileged early access, maybe accounting for this potential time delay, you could have it be a bit more balanced.

Any final thoughts?

Georgios Piliouras: I think change is coming, and that the rate of change is going to keep increasing. I have no predictions for the next five years, because one year already feels intense. We are in a very privileged minority to actually participate in these technical discussions and help steer this technology in a direction that is beneficial for all. I think we should take that opportunity seriously and collaborate.

5. BLAMELESS USERS IN A CLEAN ROOM: DEFINING COPYRIGHT PROTECTION FOR GENERATIVE MODELS (ALONI COHEN)

Aloni Cohen is an assistant professor of computer science and data science at the University of Chicago. Previously, he was a postdoctoral associate at Boston University with a joint appointment in the Institute for Computing and the School of Law. His research explores the interplay between theoretical cryptography, privacy, law, and policy. He completed his PhD at MIT, where he was advised by Shafi Goldwasser.

Cohen's core problem statement highlighted that current generative models risk "regurgitating" copyrighted training data, potentially making unwitting users liable for infringement. His research proposes a formal mathematical framework to define and provably prevent copyright-infringing outputs by focusing on the relationship between model training and "blameless" user behavior. Cohen identified gaps in the definition and conception of existing "Near Access Free" (NAF) models, which rely on a differential privacy-style comparison between models with and without a specific work. Cohen argued that NAF is insufficient because models can still learn a work through its derivatives (parodies, sequels) and because the guarantee fails

to compose across multiple user prompts.

To remedy these gaps, and drawing on the legal concept of "clean room design,"Cohen introduced a new standard for training algorithms. The framework defines a "blameless user" as one who, when provided only with the unprotected *ideas* of a work, has a negligible probability (β) of recreating the protected *expression*. A model is considered "kappa-beta clean" if it guarantees that such a user's probability of copying after interacting with the model remains below a specified threshold (κ). Cohen specified that provable protection under this framework would require two components: strong differential privacy, and golden data.

6. AI SUPPRESSION: E-DISCOVERY SOFTWARE AND BRADY (REBECCA WEXLER)

Rebecca Wexler is a professor of law at Columbia Law School whose teaching and research sit at the intersection of law and technology. She has a specific focus on privacy and secrecy in the context of the criminal legal system. She has served as a senior policy advisor in the White House Office of Science and Technology Policy and has testified before both the House and the Senate Judiciary Committees.

Wexler presented a paper co-authored with Jason Hartline, Liren Shan, and Alex Sun. It explores how machine learning-based e-discovery software, which helps litigators review large amounts of evidence, might either facilitate or hinder the discovery of exculpatory evidence of innocence in criminal cases, depending on how it is configured.

Wexler highlighted the complete lack of regulatory guidance on the use of machine learning (ML) e-discovery tools in criminal proceedings, and specifically regarding compliance with the *Brady* rule that obligates prosecutors to disclose exculpatory evidence to the defense in criminal proceedings. Wexler (and the authors of the paper) share the concern that current ML workflows in the field, specifically Continuous Active Learning (CAL), prioritizes documents it is most confident are relevant to the human reviewer, unlike classic machine learning that might prioritize data near the classifier line to learn. Through simulations on synthetic data, the research identified a technical bias where inculpatory evidence tends to cluster, while exculpatory evidence (such as alibis or impeachment material) is often more dispersed. Because CAL surfaces high-confidence documents first, there is a high probability that a prosecutor will find sufficient evidence of guilt and cease review prematurely before the algorithm surfaces the more "distant" evidence of innocence.

To mitigate the risk of suppressed exculpatory evidence, Wexler proposed three primary interventions for "technology-assisted review" (TAR) in criminal cases:

—**Two-Pass Search Strategy:** Prosecutors should be required to run separate, independent searches for inculpatory and exculpatory materials rather than training a single classifier on both.

—**Multi-calibration:** Implementing multi-calibration ensures the algorithm performs reliably across all data subsets, preventing the neglect of less-obvious exculpatory clusters.

—**Minimization Protocols:** In cases where Fourth Amendment constraints limit searches, protocols must explicitly allow for the "tagging" of Brady material during the document-sorting phase.

In conclusion, the research calls for courts to clarify the strict liability standard for evidence that the government possesses digitally but has not reviewed. The speaker concludes that a failure to resolve these doctrinal ambiguities leaves a gray area that may incentivize poor data hygiene and result in the continued failure to disclose evidence of innocence.

7. ALGORITHMIC COLLUSION WITHOUT THREATS (JUBA ZIANI)

Juba Ziani is an Assistant Professor in the H. Milton Stewart School of Industrial and Systems Engineering and an Adjunct Professor in the School of Computer Science. His research lies at the intersection of Computer Science, Operations Research, and Economics. His paper *Algorithmic Collusion without Threats* is joint work with Eshwar Ram, Natalie Collina, Sampath Kannan, and Aaron Roth.

Juba's talk began with a discussion of the differences between *classic price collusion*, in which collusion arises via human interaction (e.g., explicit agreement), and *algorithmic collusion*, in which collusion is an emergent behavior of algorithms in equilibrium. One mechanism by which algorithmic collusion can arise is if the algorithms involved use *threats* (e.g., grim trigger strategies). The main question their paper asks is: *Are threats the only mechanism by which (individually rational) sellers can engage in autonomous algorithmic collusion?*

They show, contrary to standard economic intuition, that supra-competitive prices can arise in equilibrium, even when both players are using algorithms which do not encode threats. An example such equilbrium involves one seller using a no-swap-regret algorithm, and a second seller setting a fixed price, equal to the Stackelberg leader strategy in the one-shot game. This example is particularly striking because the no-swap-regret property has been proposed as a "definitionally non-threatening" property by which one might audit pricing algorithms for antitrust compliance [Chassang and Ortner 2023; Hartline et al. 2024b]. Juba's talk closed with a discussion highlighting how this work calls for a more nuanced understanding of algorithmic collusion.

8. FIRESIDE CHAT: JASON HARTLINE AND NICOLE IMMORLICA

The 2026 Winter Meeting closed with a fireside chat, in which Nicole Immorlica interviewed Jason Hartline about his perspective on CS&Law.

Nicole Immorlica: What do you see as the main opportunity with CS&Law?

Jason Hartline: I tend to think of computer science and law as having three main touch points: CS technology might advance the law; as CS tools impact society, law needs to keep up; and viewing the legal system itself as a computational system. As computer scientists trying to understand the outcomes of computational systems, we should apply our tools to the legal system to try to understand its outcomes. I see a lot of parallels between this and the early days of EconCS.

Nicole Immorlica: What do you see as the challenges to collaborating with legal scholars, and how can we approach that from the CS or EconCS side?

Jason Hartline: I've been involved in the EC community since its very beginning. At the beginning, economists didn't come. Then there was a switch point around 2010, where we started to get more and more of them trickling over. Now a large fraction of papers in EC are authored by economists. It's been wonderful to

have them on board. One of the things I love about the CS&Law community right now is it feels like EC did back in 2000, where every time you go, there's someone doing a brand new connection between computer science and law that you've never heard of before. There's a lot of creativity happening. Legal scholars are hungry to bring computer science in. They've been doing law and tech stuff for decades, and we haven't been paying attention. So compared to CS and economics, you almost switch roles for CS and law.

Nicole Immorlica: I think one of the main frictions between CS and law is that legal scholarship and CS scholarship look very, very different. Computer scientists tend to have actually done something technical with every new paper we write. That's not the case with legal scholarship—they typically have an interpretation or they're combining ideas in a way that hasn't previously been discussed. Thinking about CS and economics, we might disagree about what approximation means or what efficiency means, but when we write down the math, we then agree. The math says what the math says. In a lot of the legal work, they haven't written down the math yet, so we can't even know if we agree. This insistence on understanding what the definitions are that allow computer scientists to make progress on legal problems is super important. But my impression of law is it's a lot of interpretation, and that's almost by design, because the flexibility in interpretation enables law to modernize. I wonder if the CS framework is going to perhaps box it in too much. Do you have thoughts on that?

Jason Hartline: It's interesting that you would say box it in. We have a law of copyright, and then LLMs come along, and when they invented copyright, they weren't thinking about LLMs. You can ask two kinds of questions: what does the existing law of copyright say about LLMs? And also, if we wanted to have a more LLM-savvy law of copyright, what should it say about LLMs? I think these are both important questions. But I think being future-proof is a strength of computer science. One of my favorite examples is when the designers of the internet came up with TCP/IP, they had no way of anticipating the things that we'd be doing with it today. And somehow those protocols that were invented in the '70s are still working and quite effective. This insistence of designing things that work out of the context that we were designing them for is a strength of computer science and algorithms.

Nicole Immorlica: What are some potential big wins that we could be proud of as an interdisciplinary field?

Jason Hartline: One thing that I've been quite impressed with about talking with legal scholars is that they seem to be much closer to practical impact than we tend to be in computer science. They frequently get put in front of panels of legislatures or judges, and they have to give opinions and write opinions. So the connections between legal scholarship and practice are actually often quite impactful.

I think one possible big win is having algorithmic sanity prevail in legal struggles relating to algorithms. For example, thinking about algorithmic collusion: the standard definition of competitive is Nash equilibrium, and we know by the PPAD completeness result that Nash equilibrium is not something algorithms can find quickly. So if we had a law saying algorithms playing against each other had

to reach Nash equilibrium, otherwise they were colluding, that would be an algorithmically insane law. The big wins are making sure that the regulations are algorithmically sane. A compelling example is a paper by Aloni Cohen and Kobbi Nissim on understanding singling out in the GDPR. The GDPR says that if you can't be singled out, then it's not private data. Their study showed that if you just literally try to interpret what the law is saying, that definition made no sense because obvious things might result in singling out in ways that are statistically uninteresting. They changed the definition to something they call *predicate singling out*, which means: can you single out someone more so than you could by random chance? Then they were able to argue that things like differential privacy protect from being predicate singled out, whereas other commonly recommended practices, like k-anonymity, do not. That kind of work—clarifying what definitions should be—is very impactful.

Nicole Immorlica: Can you pick out another greatest hit?

Jason Hartline: Let's talk about zero knowledge for criminal evidence. This is a paper with Rebecca Wexler and her co-authors. Rebecca Wexler is an expert on privilege—the idea that the law can't make you do certain kinds of things. Some privileges are at odds with getting the truth out in court proceedings. And so the law has to figure out how to get the truth out while respecting privileges. In computer science, we have novel ideas about what a proof is—interactive proofs, zero knowledge proofs. Zero knowledge proofs are proofs where the verifier is able to be certain that the prover can establish something, but doesn't learn anything extra about other stuff that they shouldn't be allowed to learn. So there's a possibility that we could have the best of both worlds: defendants in cases can interrogate the technology that's implicating them, whereas the technology owners and operators don't have to divulge their intellectual property, which they have privilege over. The law is often about getting a proof of something, and theoretical computer science has very sophisticated ideas about what a proof is. I think this is a huge opportunity for work.

Nicole Immorlica: I really appreciate that this paper shifts the conversation from trying to decide what objects must be disclosed, to what must we be able to prove about these objects.

Jason Hartline: And Rebecca Wexler made the point that this paper proves to judges that it should not be enough for someone to say "we have intellectual property, so we don't want to disclose this." Because it is possible to have zero-knowledge proofs, we should demand more from the technologies that people want to use in courts.

This is very reminiscent of the kind of meta-reasoning we like in computer science: this problem is NP-complete, so I don't know that no one can solve it, but all these people who tried to solve other NP-complete problems also can't solve my problem, so we think it's intractable. This way of doing reasoning between computer scientists and legal scholars I think is really quite compelling.

Nicole Immorlica: One thing that might be missing from the CS&Law intersection is the EconCS&Law intersection. Do you think is this CS and law interdisciplinary space needs more injection of other different fields?

Jason Hartline: The SIGecom community has brought computer science and

economics much closer together, and that makes the field of law and economics very accessible to us—we can just readily get involved. We have the same interest in putting formality and arguing about incentives in the law, so I don't see any reason to be separate. Already, economic papers on antitrust and algorithmic collusion are showing up in the law and economics literature.

Nicole Immorlica: What would you recommend we all read if we're wanting to dive into this a little bit more?

Jason Hartline: I would say come next week to the ACM CS&Law Conference in Berkeley.

9. CONCLUSION

This conference concluded with remarks from Inbal Talgam-Cohen, who highlighted the contributions of the numerous participants. Perhaps the most important takeaway from this year's research meeting is that the burgeoning field of CS-law is rich ground for the computer science and economics thinking that the field of AGT offers. And, not only is this field amenable to the AGT toolkit, but, with the rising influence of AI in our society, perhaps, also has a lot to gain from it.

REFERENCES

Bar-Gill, O., Sunstein, C. R., and Talgam-Cohen, I. 2023. Algorithmic harm in consumer markets. *Journal of Legal Analysis 15,* 1, 1–47.

Bitan, D., Canetti, R., Goldwasser, S., and Wexler, R. 2022. Using zero-knowledge to reconcile law enforcement secrecy and fair trial rights in criminal cases. In *Proceedings of the 2022 Symposium on Computer Science and Law*. 9–22.

Brynjolfsson, E. and Hitzig, Z. 2025. *AI's Use of Knowledge in Society*. University of Chicago Press.

Chan, P. K., Carrel, A., Varia, M., and Wang, X. 2025. Murmurs of the silenced: Secure reporting of misconduct settlements. In *Proceedings of the 2025 Symposium on Computer Science and Law*. 121–135.

Chassang, S. and Ortner, J. 2023. Regulating Collusion. *Annual Review of Economics 15,* Volume 15, 2023 (Sept.), 177–204. Publisher: Annual Reviews.

Cohen, A. and Nissim, K. 2020. Towards formalizing the gdpr's notion of singling out. *Proceedings of the National Academy of Sciences 117,* 15, 8344–8352.

Cohen, A., Scheffler, S., and Varia, M. 2022. Can the government compel decryption? don't trust-verify. In *Proceedings of the 2022 Symposium on Computer Science and Law*. 23–36.

Collina, N., Globus-Harris, I., Goel, S., Gupta, V., Roth, A., and Shi, M. 2026. Collaborative prediction: Tractable information aggregation via agreement. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 4712–4798.

Dutz, M., Shao, H., Blum, A., and Cohen, A. 2025. A machine learning theory perspective on strategic litigation. *arXiv preprint arXiv:2506.03411*.

Feigenbaum, J. and Weitzner, D. J. 2026. At the intersection of computer science and law.

Fish, S., Gonczarowski, Y. A., and Shorrer, R. I. 2024. Algorithmic collusion by large language models. *arXiv preprint arXiv:2404.00806 7,* 2, 5.

Hadfield-Menell, D. and Hadfield, G. K. 2019. Incomplete contracting and ai alignment. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 417–422.

Hartline, J. D., Linna Jr, D. W., Shan, L., and Tang, A. 2022. Algorithmic learning foundations for common law. In *Proceedings of the 2022 Symposium on Computer Science and Law*. 109–117.

Hartline, J. D., Long, S., and Zhang, C. 2024a. Regulation of algorithmic collusion. In *Proceedings of the 2024 Symposium on Computer Science and Law*. 98–108.

Hartline, J. D., Long, S., and Zhang, C. 2024b. Regulation of Algorithmic Collusion. In *Proceedings of the 2024 Symposium on Computer Science and Law*. CSLAW '24. Association for Computing Machinery, New York, NY, USA, 98–108.

Kleinberg, J. and Mullainathan, S. 2024. Language generation in the limit. *Advances in Neural Information Processing Systems 37*, 66058–66079.

Lee, K., Cooper, A. F., and Grimmelmann, J. 2024. Talkin"bout ai generation: copyright and the generative-ai supply chain (the short version). In *Proceedings of the 2024 Symposium on Computer Science and Law*. 48–63.

Meding, K. and Sorge, C. 2025. What constitutes a deep fake? the blurry line between legitimate processing and manipulation under the eu ai act. In *Proceedings of the 2025 Symposium on Computer Science and Law*. 152–159.

Saig, E., Einav, O., and Talgam-Cohen, I. 2024. Incentivizing quality text generation via statistical contracts. *Advances in Neural Information Processing Systems 37*, 51196–51222.

Singh, A., Rodemann, J., Verma, R., Chau, S. L., and Muandet, K. 2026. Incentive aware ai regulations: A credal characterisation. *arXiv preprint arXiv:2603.05175*.