Optimizing for What? Algorithmic Amplification and Society

April 28-29, 2023

KNIGHT
FIRST AMENDMENT
INSTITUTE at
COLUMBIA UNIVERSITY

WELCOME

Over the past six years, the Knight Institute has returned multiple times to the question of how new technology is shaping and distorting public discourse. Last year, as part of an initiative focused on the regulation of lies, we examined the problem of misinformation on social media, and considered what kinds of regulatory responses to that problem might be effective and consistent with democratic values. Previously, we've asked whether the power the technology giants exert over discourse online ought to be curtailed through antimonopoly measures, and explored alternatives to the internet's dominant ad-driven business model. For several years, we have represented researchers and journalists who study public discourse online, helping them pursue projects that illuminate the ways in which platforms determine who can speak, who gets heard, and which ideas get traction.

This symposium on algorithmic amplification advances and expands this work through an interdisciplinary collaboration with Princeton University's Arvind Narayanan, who is a visiting senior research scientist at the Knight Institute this year. Through the symposium and a series of commissioned papers, we are focusing on the technology that most directly shapes our online experience. We hope to demystify recommendation algorithms, asking what their implications might be for public discourse and for society more broadly, and how regulatory or other interventions might help us harness their benefits and mitigate their costs. These questions are of course at the heart of public debate right now—including the debate about banning TikTok, the controversy surrounding Elon Musk's takeover of Twitter, and the litigation pending before the U.S. Supreme Court concerning platforms' recommendation algorithms and Section 230 of the Communications Decency Act.

That so many leading scholars and researchers from so many fields and backgrounds are contributing to this symposium is a testament to the urgency of the topic—and also to Professor Narayanan's vision and leadership. We are very fortunate to have had the chance to work with him over the past months, and are grateful to him for the energy, deep knowledge, and seriousness of purpose he has brought to this important and pressing project.

Jameel Jaffer and Katy Glenn Bass

INTRODUCTION

Whatever happens to TikTok, its influence has already reshaped all of social media—recommendation algorithms play a central role in content propagation. Algorithms are also playing a bigger role in content moderation through methods such as downranking. To improve the quality of the information ecosystem, understanding the impacts of algorithms is essential.

Most platform recommendation algorithms today are designed to optimize for engagement. Algorithms are not neutral. Compared to non-algorithmic systems (such as a chronological feed), they amplify some speech and suppress others. Platforms are "complex systems," so amplification is an emergent and hard-to-predict effect of interactions between design and human behavior.

This symposium will tackle pressing empirical questions about amplification: How do algorithmic recommendations affect what real users see on social media, and how does platform design affect content creators? What research methods are available to study the feedback loop created by algorithms learning from users' behavior and users relying on algorithm-mediated social learning? How can we enable more audit research?

The normative questions are just as important. How do algorithmic platforms distribute attention and shape social relations? What kind of environment do we want the digital public square to be? What makes algorithmic amplification wrongful? What are the moral and political responsibilities of platforms?

The effects of platform algorithms are profound, because speech of every conceivable kind is carried by platforms, including entertainment, news, politics, social movements, educational content, public health information, scholarship, religious content, sports, cultural products such as music and art, and other commercial content about travel, restaurants, shopping, jobs, and more.

Some effects are positive, such as the decreased power of gatekeepers in identifying new talent. On the other hand, algorithmic logic affects everything ranging from the production and dissemination of science to the restaurant market. Each domain has its own notion of quality, refined over decades or centuries, but platforms reward unrelated factors. The algorithm-savvy are able to exploit the reach of platforms for their own purposes, while entities traditionally entrusted with information dissemination, such as public health agencies, are struggling to adapt to the new media environment.

The panels will discuss various ideas for reform, including nutrition labels, friction, algorithmic interventions, and decentralized alternatives. How can platforms go beyond engagement optimization? For example, how can they design recommender systems to bridge political divides? What can we learn from public service media on how to design recommendation engines that reflect cultural values and responsibly curate cultural content? How can platforms empower users to make better informed decisions about potential disinformation? Conversely, what design interfaces can allow users to actively teach platforms their preferences?

This symposium brings together technologists with legal scholars, sociologists, psychologists, philosophers, and others. The speakers include authors of 16 papers that were written specifically for the event, together with a few additional speakers who are also leading thinkers on the topic. Many of the speakers bring deep experience from having worked on these questions at platform companies. Broadly, the symposium aims to further the understanding of algorithmic amplification and to explore interventions, some situated within platforms and some outside them.

Arvind Narayanan

PROGRAM AT A GLANCE

Friday, April 28

Alfred Lerner Hall, Columbia University
Roone Arledge Cinema

9:00 am – 9:10 am — **Welcome**

Jameel Jaffer

9:10 am – 9:40 am — **Keynote and conversation**

Alondra Nelson
Jameel Jaffer

9:40 am – 10:50 am — **Panel 1: Level setting**

Panelists
Tarleton Gillespie
Daphne Keller
Kristian Lum

Moderator
Arvind Narayanan

10:50 am – 11:05 am — **Break**

11:05 am – 12:35 pm — **Panel 2: Audits**

Panelists
Fabian Baumann
William J. Brady
Smitha Milli
Inioluwa Deborah Raji

Moderator
Laura Edelson

12:35 pm – 2:00 pm — **Lunch**

See list of nearby lunch options on page 51.

2:00 pm – 3:10 pm — **Panel 3: Normative questions**

Panelists
Annie Dorsen
Benjamin Laufer
Seth Lazar

Moderator
Katy Glenn Bass

3:10 pm – 3:30 pm — **Break**

3:30 pm – 4:50 pm — **Panel 4: Reform part 1**

Panelists
Luca Belli
Brett Frischmann
Ravi Iyer
Yoel Roth

Moderator
Camille François

4:50 pm – 5:00 pm — **Visualizing virality**

Presenters
Samia Menon
Sahil Patel

Saturday, April 29

Faculty House, Columbia University
Presidential Room 2

9:30 am – 11:00 am	**Panel 5: Empirical look at user behavior**	
	Panelists Jason Burton Kevin Feng Benjamin Kaiser Angela Lai	**Moderator** Mor Naaman
11:00 am – 11:20 am	**Break**	
11:20 am – 12:30 pm	**Panel 6: Reform part 2**	
	Panelists Georgina Born Aviv Ovadya Alessandro Piscopo	**Moderator** Joe B. Bak-Coleman
12:30 pm – 1:30 pm	**Lunch**	
	Provided	

Keynote and conversation with Alondra Nelson

9:00 am – 9:40 am

Alondra Nelson will offer brief keynote remarks followed by a Q&A led by Jameel Jaffer.

Alondra Nelson is the Harold F. Linder Professor at the Institute for Advanced Study and a distinguished senior fellow at the Center for American Progress. As former deputy assistant to President Joe Biden, she served as principal deputy director for science and society and acting director of the White House Office of Science and Technology Policy (OSTP) from 2021-2023.

Including her in the list of Ten People Who Shaped Science in 2022, Nature said of Nelson, "this social scientist made strides for equity, integrity and open access." Nelson's work at OSTP also drove Biden-Harris administration strategy to create science and technology policy that expands economic opportunity, protects civil rights, enhances security, advances equity, and ensures that innovation works for, not against, our democratic values. During her tenure, Nelson led a team writing the landmark Blueprint for an AI Bill of Rights, which lays the groundwork for policymakers, technology developers, entrepreneurs, legislators, civil society, and others to better safeguard people's rights and access to opportunities as algorithms and AI reach further into our lives.

She was the 14th president and CEO of the Social Science Research Council, and in this role developed a series of programs that brought research to bear on the role of social media platforms on social relations and political culture.

An acclaimed social scientist, Nelson writes and lectures widely on the intersections of science, technology, medicine, and social inequality. She is the author of several books, including *The Social Life of DNA*. Her essays, reviews, and commentary have been featured in national and international media outlets, including The New York Times, The Washington Post, The Wall Street Journal, Wired, and Science.

Nelson is a member of the American Academy of Arts and Sciences, the American Association for the Advancement of Science, the American Philosophical Society, the National Academy of Medicine, and the Council of Foreign Relations.

Twitter: @alondra

PANEL 1

Level setting

9:40 am – 10:50 am

This panel will set the stage by discussing how platforms and platform algorithms work, laying out the issues at stake, reviewing recent developments, and looking at the legal questions relevant to possible reform options.

Panelists

Tarleton Gillespie, *Microsoft Research New England*
Daphne Keller, *Stanford University*
Kristian Lum, *University of Chicago*

Moderator

Arvind Narayanan, *Princeton University and the Knight First Amendment Institute at Columbia University*

PAPER ABSTRACTS

Do Not Recommend? Reduction as a form of content moderation

Tarleton Gillespie

Public debate about content moderation has overwhelmingly focused on removal: social media platforms deleting content and suspending users, or opting not to do so. However, removal is not the only available remedy. Reducing the visibility of problematic content is becoming a commonplace element of platform governance. Platforms use machine learning classifiers to identify content they judge misleading enough, risky enough, or offensive enough that, while it does not warrant removal according to the site guidelines, it warrants demoting them in algorithmic rankings and recommendations. In this essay, I document this shift and explain how reduction works. I then raise questions about what it means to use recommendation as a means of content moderation.

Amplification and Its Discontents

Daphne Keller

There is a popular line of reasoning in platform regulation discussions today that argues, "Platforms aren't responsible for what their users say, but they are responsible for what the platforms themselves choose to amplify." This provides a seemingly simple hook for regulating algorithmic amplification. However, for lawyers or policymakers trying to set rules for disinformation, hate speech, and other harmful or illegal content online, focusing on amplification won't make life any easier. It may increase, rather than decrease, the number of problems to be solved before arriving at well-crafted regulation. Models for regulating amplification have a great deal in common with the more familiar models from intermediary liability law, which defines platforms' responsibility for content posted by users. As with ordinary intermediary liability laws, the biggest questions may be practical: Who defines the rules for online speech, who enforces them, what incentives do they have, and what outcomes should we expect as a result? And as with those laws, some of the most important considerations—and, ultimately, limits on Congress's power—come from the First Amendment.

In this essay, I will lay out why "regulating amplification" to restrict distribution of harmful or illegal content is hard. My goal in doing so is to keep smart people from wasting their time devising bad laws, and speed the day when we can figure out good ones. I will draw in part on novel regulatory models that are more developed in Europe. My analysis, though, will primarily use U.S. First Amendment law. I will conclude that many models for regulating amplification face serious constitutional hurdles, but that a few—grounded in content-neutral goals, including privacy or competition—may offer paths forward.

The Myth of "The Algorithm": A system-level view of algorithmic amplification

Kristian Lum and Tomo Lazovich

As people consume more content delivered by recommender systems, it has become increasingly important to understand how content is amplified by these recommendations. Much of the recent work to study algorithmic amplification implicitly assumes that *the algorithm* is a single machine learning model acting on an immutable corpus of content to be recommended. Additionally, there is an inherent assumption of a neutral "nonalgorithmic" baseline against which to compare. In actuality, there are several other components of the system that are not traditionally considered part of *the algorithm* that influence what ends up on a user's content feed and potentially corrupt the neutrality of any baseline measurement: upstream editorial policies or decisions that determine what content is eligible to be ranked by the algorithmic recommender system, including NSFW and toxicity filtering; peripheral models that shape the evolution of the social graph, such as account recommendation models; and explicit user preferences and behaviors. All of these components affect what ultimately gets amplified and can confound how we measure amplification.

Our proposed paper has three aims. First, we will enumerate some of these components that influence algorithmic amplification. Second, we will explore how the assumption of a "neutral" baseline that was not shaped by prior behavior of these components, particularly the "reverse chronological" content feed, can lead to poor measurement of amplification. Third, we will suggest some paths forward for measurement and mitigation that address the same concerns that underlie the recent discourse around algorithmic amplification but do not rely on the existence of a neutral baseline. We hope this work will be a call to action to the research community to consider otherwise overlooked areas that greatly influence how content is amplified on social platforms, and we see this workshop as an opportunity to gather input from the community on these areas.

PANEL 2

Audits

11:05 am - 12:35 pm

The panel will consider how algorithmic recommendations affect what real users see on social media, with deep dives into Twitter and YouTube. Panelists will discuss how platform design affects content creators and talk about research methods and ways to enable more audit research.

Panelists

Fabian Baumann, *Max Planck Institute for Human Development*
William J. Brady, *Northwestern University*
Smitha Milli, *Cornell Tech*
Inioluwa Deborah Raji, *University of California, Berkeley*

Moderator

Laura Edelson, *New York University*

PAPER ABSTRACTS

Field Experiments on the Impact of Algorithmic Curation on Content Consumption Behavior

Fabian Baumann and Philipp Lorenz-Spreen

Algorithms like search engines or recommender systems have the potential to decisively influence the drivers of cultural evolution, e.g., in information search and sharing. This influence is particularly evident on social media platforms, where direct peer-to-peer communication is typically mediated by algorithmically curated feeds that are optimized for engagement and that provide users with personalized content. Previously, algorithmic personalization on social media has been studied through a political lens focusing on the political content that users get recommended. Here we will take a more general perspective and focus on how recommended content (i.e., exposure) and user behavior (i.e., engagement) interplay across the broader cultural spectrum. The empirical study of the impact of algorithmic recommendations on culture poses fundamental challenges that connect to the classical trade-off between ecological validity and experimental control: When studied with observational data, the inherent coupling of algorithms and human behavior is impossible to disentangle, and when studied in the lab, the encountered content is often artificial and does not mimic realistic exposure.

Here, we will outline an approach to strike a good balance between the two, namely with field experiments on existing social media platforms. As one example, we present an experimental paradigm that makes use of Twitter's feature to switch between algorithmic ("For You") and chronological feed as an experimental manipulation. Screen recordings can be used for measuring the resulting exposure to content and Twitter's Application Programming Interface, or API, for monitoring the subsequent user behavior. While users in "reverse-chron" mode only experience their friends' activity, the "For You" feed adds personalized and presumably more engaging content from outside their social circle. After incentivizing participants to switch to one or the other feed and record their screen while browsing Twitter, we quantify the change of their behavior along various cultural dimensions. For instance, we can examine how users' content exposure changes, if users' engagement becomes more passive or active, and if the distribution of content becomes more heterogeneous or not. Our empirical results will help to get an ecologically valid measure of the causal effect of algorithmic curation on the statistical distribution of consumed and produced content.

Algorithm-Mediated Social Learning in Online Social Networks

William J. Brady, Joshua Conrad Jackson, Björn Lindström, and M.J. Crockett

Humans rely heavily on social learning to navigate the social and physical world. For the first time in history, we are interacting in online social networks where content algorithms filter social information, yet little is known about how these algorithms influence our social learning. In this review, we synthesize emerging insights into this "algorithm-mediated

social learning" and propose a framework that examines its consequences in terms of functional misalignment. We argue that the functions of human social learning and the goals of content algorithms are misaligned in practice. Algorithms exploit basic human social learning biases (i.e., a bias toward prestigious, in-group, moral, and emotional information, or PRIME information) as a side effect of their goals to sustain attention and maximize engagement on platforms. Social learning biases function to promote adaptive behaviors that foster cooperation and collective problem-solving. However, when social learning biases are exploited by algorithms, PRIME information becomes amplified in the digital social environment in ways that can stimulate conflict and spread misinformation. We show how this problem is ultimately driven by human-algorithm interactions where observational and reinforcement learning exacerbate algorithmic amplification, and how it may even escalate to impact cultural evolution. Finally, we discuss practical solutions for reducing functional misalignment in human-algorithm interactions via strategies that help algorithms promote more diverse and contextually sensitive information environments.

Emotional and Political Effects of Twitter's Ranking Algorithm

Smitha Milli, Micah Carroll, Sashrika Pandey, Yike Wang, and Anca Dragan

Social media ranking algorithms play a large role in curating the content seen by users online. Despite their potentially great impact, it is usually infeasible for third-party auditors to experimentally measure the effects of social media ranking algorithms. A randomized experiment would ordinarily require randomizing which users get content served by which algorithm, and thus, require insider access.

In this paper, we run a randomized experiment measuring effects of Twitter's ranking algorithm *without* internal access. To do so, we leverage a somewhat unique feature of Twitter: Users can opt-out of personalization and view their tweets in only a chronological order. For a large group of recruited participants (targeted at N=2,000), we will sample tweets from both their chronological and personalized timeline at the exact same moment in time. We then survey the participants about both sets of tweets in a randomized order. The questions concern the emotional and political content of the tweets as well as the user's change in emotions and political beliefs in response to the tweets. By comparing responses to tweets in personalized timeline to those for chronological timeline, we can measure the effect that the personalized ranking algorithm has (relative to a chronological feed).

Cycles of Symbol Production on Online Platforms

Inioluwa Deborah Raji, Fernando Diaz, and Irene Lo

While much of the emphasis in the current literature on algorithmic amplification focuses on how online platforms might distort content consumption patterns, less attention has been paid to how such platforms influence content creator actions and outputs.

Existing theoretical models of content producer dynamics on online platforms tend to be limited, anchored to similar narrow assumptions. Notably, past work assumes that online platforms operate mainly as content distributors and that content providers do not interact while competing for the finite attention of consumers. However, on many social media platforms in particular, this is not practically the case—producers effectively operate as just another class of users, interacting and influencing each other directly and indirectly, as well as adaptively updating their content at varying rates of production in response. In this work, we theorize that content creators on social media platforms do not just unilaterally compete but can in fact *collude* in various instances in order to amplify each other and minimize (rather than maximize) the diversity of the content available on an online platform, disrupting individualized amplification schemes.

We go further to explore the influence of the design space of various online platforms on the cycle of content homogenization and diversification. We find that platform characteristics such as user experience features, content discovery heuristics, and monetization schemes factor heavily into the degree of content creator interactions and collusion; the adaptability and rate of content production; and other creator norms, which in turn determine the length and nature of content homogenization cycles. We hope to support this finding with empirical evidence from TikTok, Facebook, and YouTube.

PANEL 3

Normative questions

2:00 pm - 3:10 pm

How do algorithmic platforms distribute attention and shape social relations? How have they influenced the arts? The public square? What makes algorithmic amplification wrongful? What are the moral and political responsibilities of platforms?

Panelists

Annie Dorsen, *Independent artist*
Benjamin Laufer, *Cornell Tech*
Seth Lazar, *Australian National University*

Moderator

Katy Glenn Bass, *Knight First Amendment Institute at Columbia University*

PAPER ABSTRACTS

The Work of Art in the Age of Digital Commodification: An analysis of the emerging digital political economy of the performing arts

Sam Gill and Annie Dorsen

This paper critically analyzes how digital technology may influence the performing arts. It does this by examining in some detail the conceptual foundations of the digital commodity form, and explores how the unique features and exigencies of digital commodification may influence a range of forms of creative expression—drawing in particular on the discourse surrounding the emergent "creator economy." The piece argues specifically that the digital commodity form is already having five impacts on the performing arts: (1) eroding the boundaries of art as a professional practice, (2) obliterating the line between creative producer and audience, (3) sublimating the aesthetic productive constraints and choices immanent in preset technologies, (4) replacement of legacy gatekeepers with digital operators, and (5) the deinstitutionalization of creative labor. As a result of these changes, the paper theorizes that the absorption of the performing arts into commercially driven digital systems will begin to reduce artistic and creative expression to homogenized, interchangeable content that has been shaped before inception by economic imperatives focused on human attention and engagement. The piece further worries that these shifts will corrode and overtake legacy cultural institutions, thereby eliminating any notion of authoritative aesthetic discernment and unleashing an intensification of creative labor similar to that seen in other sectors of the now digitized economy. It concludes with some reflections on generative artificial intelligence and a review of critical questions facing artists and cultural institutions as well as scholars and critics analyzing the rise of digital technology.

What Makes Algorithmic Amplification Wrongful?

Benjamin Laufer and Helen Nissenbaum

Increasingly concerned about the way in which content spreads on the internet, scholars reach for the concept of algorithmic amplification (AA) as both an explanation and a warning. Although these researchers frequently acknowledge the metaphorical and conceptual haziness around the term, they continue to rely on it to carry both descriptive and normative intent. In itself, haziness need not disqualify a concept, except when it hides substantive assumptions with decisive normative implications.

This paper offers foundational work to give AA conceptual precision and normative teeth. First, it resuscitates the historical context around the meaning of *amplify*, a transitive concept from signal processing and system dynamics. It then turns to the normative question: When is AA wrongful? A sound account of what makes amplification problematic is a necessary precondition for discussing what to do about it.

Research has found that AA may bring about negative social impacts including *disinformation*, *bias*, and *extremism*. These problems, tied to content rather than process, are harmful consequences of AA, but they are not constitutive. At the root of wrongful AA is the deterioration of existing trustworthy processes for justification and legitimation. Algorithmic decision-making can disrupt or distort these processes. By shattering long-standing norms crucial for maintaining a common stock of knowledge, AA can undermine democracy. Therefore, we contend that AA is problematic when information is distributed according to processes that were not arrived at through legitimate social deliberation.

Platform-mediated internet communications are particularly prone to wrongful forms of AA, for which platforms ought to be held responsible. Where we believe AA to be wrongful, we will demonstrate the mechanisms causing harm in the cases of climate science communication and vaccination campaigns.

Communicative Justice and the Distribution of Attention
Seth Lazar

I argue, first, that algorithmic intermediaries govern the digital public sphere through their architecture, amplification algorithms, and moderation practices, and that they have a responsibility to do so better. This means more than just enumerating and responding to pathologies such as misinformation, radicalization, and abuse. We also need a positive ideal to aim at. Political philosophy should offer such an ideal, but it tells us only when not to interfere in free speech, not how to shape public communication and distribute attention. In response, I introduce a new theory of communicative justice: an account of the communicative interests that those who govern the digital public sphere should promote, and the democratic egalitarian norms by which their doing so should be constrained. This can guide us in shaping public communication and distributing attention, in balancing the governing responsibilities of private and public actors, and in striving for procedural legitimacy in governance of the digital public sphere.

PANEL 4

Reform part 1
3:30 pm - 4:50 pm

Panelists will discuss various ideas for reforming, including nutrition labels, friction, algorithmic interventions, and decentralized alternatives, with a deep dive into one particular area: how to dampen conflict feedback loops.

Panelists
Luca Belli, *Sator Labs and University of California, Berkeley*
Brett Frischmann, *Villanova University*
Ravi Iyer, *Psychology of Technology Institute*
Yoel Roth, *University of California, Berkeley*

Moderator
Camille François, *Columbia University*

PAPER ABSTRACTS

What's in an Algorithm? Empowering users through nutrition labels for recommender systems

Luca Belli and Marlena Wisniak

Concerns around algorithmic amplification rightfully involve how they operate, particularly with regard to the harms they produce when optimized for engagement. Yet to effectively address adverse impacts of amplification on human rights and civic space, we need a nuanced and commonly agreed upon definition thereof.

Often amplification is treated as an innate property of an algorithmic system, maybe hard coded by the designers and developers to reflect their own values and (usually profit-driven) objectives. While human input certainly shapes what such models are optimized for, the reality is much more complex, as these sociotechnical systems respond to the users' behavior itself.

A corollary to defining amplification is to measure it. We argue that algorithmic amplification cannot be measured along one dimension only; rather it is a complex phenomenon that could be better understood via bringing multiple metrics together. Meaningful engagement with external stakeholders, especially marginalized groups and those living in the Global South, is urgently needed to map and understand diverse metrics and their limitations. Furthermore, amplification is relative to a baseline: Defining this common baseline is a pre-requirement that is often missed.

We propose to introduce "nutrition labels" for recommender models. Such a collection of agreed upon metrics could be useful to understand how these systems operate and ultimately ensure that their use protects and promotes human rights. Aiming to spark an inclusive conversation on how to measure amplification—with participation from civil society, academia, policymakers, international organizations, and the private sector—we offer a couple suggestions for how such measures could look.

Our research focuses on recommender systems for social media timelines, exploring metrics that would not only be relevant for today's dominant platforms, but also for alternative and emerging models such as the "Fediverse" and web3 technologies.

How Friction-in-Design Moderates, Amplifies, and Dictates Speech and Conduct

Brett Frischmann and Paul Ohm

Besides algorithmic determinations and human decisions to emphasize one speaker, message, product, or service over another, amplification is often the product of decisions to remove or inject friction in the design of digital interfaces and platforms. Thus, amplification (and prioritization and optimization) should be understood and evaluated in terms of countervailing design decisions regarding different types and degrees of friction. We are interested in optimization and amplification not only for speech and content but also as it increasingly shapes and dictates behavior and conduct.

Our project connects to an emerging literature that considers the roles friction plays in the design of platforms, software, and other technological systems, as a means to protect values such as security, privacy, competition, and consumer protection. We have written some foundational works on friction-in-design.

We consider case studies from social media platform design that highlight the roles friction plays in amplification and optimization. TikTok's infinite scroll removes friction to increase engagement. WhatsApp's limits on frequently forwarded messages use friction to reduce virality. Twitter warns users to reconsider retweeting links to unread articles.

Drawing on these case studies, we explore ways to inject friction into techno-social systems to address some of amplification's potential harms. Regulators might mandate limits on message forwarding or impose "rest stops" in infinite scrolls. Our earlier work analyzes how these approaches comport with the First Amendment. Engineers might be trained to better understand the risks of frictionless design and how, when, and where to inject purposeful friction to address these risks.

Besides algorithmic determinations and human decisions to emphasize one speaker, message, product, or service over another, amplification is often the product of decisions to remove or inject friction in the design of digital interfaces and platforms. Thus, amplification (and prioritization and optimization) should be understood and evaluated in terms of countervailing design decisions regarding different types and degrees of friction. We are interested in optimization and amplification not only for speech and content but also as it increasingly shapes and dictates behavior and conduct.

The Algorithmic Management of Polarization and Violence on Social Media

Ravi Iyer, Jonathan Stray, and Helena Puig Larrauri

Social media platforms are involved in all aspects of social life—including in conflict settings. These platforms are not equipped to make complex judgments about conflicts, but their incidental choices about how they are designed can have profound effects on people within conflict settings. At a minimum, they should not incentivize conflict actors toward more hateful and potentially violence-inducing speech, and they should not enable mass harassment and manipulation. They should provide reasonable affordances for empowering individuals within a conflict setting to keep themselves safe and informed. Evidence suggests these minimum conditions have not been met, though steps have been taken in the right direction. Platforms could be designed to dampen conflict feedback loops, and the resulting destructive escalation to polarization and violence. While content moderation has received considerable attention, it will never affect more than a small amount of objectively policy-violating content and expanding those efforts will only lead to more backtracking,

unfair over-enforcement, and controversy. In contrast, every experience of content that is consumed on social media platforms is influenced by the design of the user interface and algorithms of that platform. Platforms designed for business outcomes are not neutral with regard to conflict relevant behavior. In this paper, we will discuss evidence for how platforms and their algorithms are currently affecting polarization and violence. We will then make evidence-based suggestions for reforming platform design and suggest next steps for the many things that remain unknown.

PANEL 5

Empirical look at user behavior

9:30 am - 11:00 am

Algorithms learn from users' behavior, and users rely on algorithm-mediated social learning. What is the nature of the resulting feedback loop? How can platforms empower users to make better informed decisions about potential disinformation? Conversely, what design interfaces can allow users to actively teach platforms their preferences?

Panelists

Jason Burton, *Copenhagen Business School and Max Planck Institute for Human Development*
Kevin Feng, *University of Washington*
Benjamin Kaiser, *Princeton University*
Angela Lai, *New York University*

Moderator

Mor Naaman, *Cornell Tech*

PAPER ABSTRACTS

Algorithmic Amplification for Collective Intelligence

Jason Burton

The algorithmic amplification of online content is often framed as a danger to be mitigated, with the dominant "engagement-based ranking" approach frequently cited as cause for divisiveness and sensationalism in public discourse. In recent proof-of-concept studies, however, we show that algorithmic amplification can be designed to promote collective intelligence. Specifically, these studies show—through agent-based simulations and online multiplayer experiments—how systematic relationships between belief distributions and collective accuracy can be leveraged to algorithmically mediate online interactions and reduce error in collective estimations, even when the ground truth is unknown. In this work-in-progress, I expand on these studies by drawing from the literature on "wisdom of the crowd" effects and argumentation theory to design, deploy, and evaluate algorithms that curate content to support deliberation and improve the accuracy of people's beliefs. In doing so, this work targets both theoretical and practical implications: First, it provides further experimental evidence reaffirming the position that algorithmic amplification can influence the beliefs people form. Second, our findings aim to inform the design of new civic technologies in which algorithmic amplification plays a key role—for example, by proposing new features of deliberation and online behavior to be mined and amplified. Third, and most broadly, this work contributes to the ongoing conceptual discussion of how the design of recommendation systems and online platforms can be modified to better align with the democratic values that the internet once promised.

Teachable Agents for End-User Empowerment in Personalized Feed Curation

Kevin Feng, David McDonald, and Amy X. Zhang

As a small handful of platforms act as social architects—shaping user norms and behaviors through algorithmically curated feeds and interface affordances alike—the risks of curatorial centralization begin to emerge: Top-down, platformwide policies rob users of their sense of agency, fail to account for nuanced experiences, and overall marginalize users and communities with unique values and customs. Prior work has shown that users have attempted to reclaim their agency by deriving "algorithmic folk theories" to probe black-box feed curation algorithms and "teaching" algorithms to yield more satisfactory content through strategic interactions with their feed. Given this, we ask: How can users' inherent teaching abilities be more explicitly employed to empower personalized curation and transparent algorithmic customization in online social settings? We draw inspiration from the paradigm of interactive machine teaching and explore user-teachable agents for feed curation. To do this, we first conducted a formative study to understand how users would approach explicitly teaching an algorithmic agent about preferences in their social media feeds, as opposed to the agent implicitly learning them. Based on our findings, we propose in-feed affordances that allow users to execute a teaching loop by 1) explaining content preferences via examples of posts to a learnable agent, 2) evaluating the agent's effectiveness of learning, and 3) iteratively formulating a curriculum of teaching goals and examples. We conclude with a discussion of challenges and next steps, with an eye towards how our approach may be used to better align incentives of users and platforms in sociotechnical systems.

It's the Algorithm: A large-scale comparative field study of news quality interventions

Benjamin Kaiser and Jonathan Mayer

There is a widespread belief, and growing anecdotal evidence, that platforms' recommendation algorithms can contribute significantly to the spread or suppression of misinformation. But work by platforms and researchers to develop interventions to counter the spread of misinformation has overwhelmingly focused on user-facing, informative interventions like fact checks and content labels. There is little rigorous evidence to answer the question of whether algorithmic interventions may be more effective than informative interventions.

We conducted the first study analyzing both informative and algorithmic misinformation interventions deployed in the ordinary functionality of a major online platform. At large scale and across multiple countries, we compared the effects of informative and algorithmic interventions on user engagement with misinformation. We found that an algorithmic deamplification intervention reduced engagement with misinformation by over half, while informative interventions had statistically insignificant effects on engagement.

Based on our findings, we argue that research priorities should shift from informative interventions to algorithmic interventions, that platforms must be more transparent about what content their algorithms amplify and deamplify, and that research collaborations between platforms and academics—not just data sharing initiatives—are essential to learn how to effectively counter misinformation online.

Echo Chambers, Rabbit Holes, and Algorithmic Bias: How YouTube recommends content to real users

Megan Brown, James Bisbee, Angela Lai, Richard Bonneau, Jonathan Nagler, and Joshua A. Tucker

To what extent does the YouTube recommendation algorithm push users into echo chambers, ideologically biased content, or rabbit holes? Using a novel method to estimate the ideology of YouTube videos and an original experimental design to isolate the effect of the algorithm from user choice, we demonstrate that the YouTube recommendation algorithm does, in fact, push real users into mild ideological echo chambers where, by the

end of the data collection task, liberals and conservatives received different distributions of recommendations from each other, though this difference is small. While we find evidence that this difference increases the longer the user followed the recommendation algorithm, we do not find evidence that many go down rabbit holes that lead them to ideologically extreme content. Finally, we find that YouTube pushes all users, regardless of ideology, towards moderately conservative and an increasingly narrow range of ideological content the longer they follow YouTube's recommendations.

PANEL 6

Reform part 2

11:20 am - 12:30 pm

How can platforms go beyond engagement optimization? For example, how can they design recommender systems to bridge political divides? What can we learn from public service media on how to design recommendation engines that reflect cultural values and responsibly curate cultural content?

Panelists

Georgina Born, *University College London*
Aviv Ovadya, *Harvard University*
Alessandro Piscopo, *BBC Product Group*

Moderator

Joe B. Bak-Coleman, *Columbia University*

PAPER ABSTRACTS

A Public Service Media Perspective on the Algorithmic Amplification of Cultural Content

Fernando Diaz and Georgina Born

Streaming entertainment platforms curate cultural content such as music, film, and literature and significantly influence the nature of individual cultural experience. Recommender systems play an important role in this process, basing curatorial decisions on algorithms optimized for objectives such as engagement, retention, and advertising revenue. As a result, multiple studies have demonstrated that some genres or groups of content creators are amplified while others are overlooked. Although these studies describe distortions in the content people consume, they do not provide guidance on what appropriate curation of cultural content might look like. Considering this, we analyze algorithmic amplification specifically in the curation of cultural content, focusing on disparities between engagement and retention as goals of recommender systems and normative concerns about what kinds of algorithmic curation of cultural content can be developed to promote cultural experiences oriented to social justice and the public good. For guidance on such normative concerns, we turn to principles underlying public service media (PSM) systems in democratic societies. These principles, refined over decades in the programming of cultural content, expand the desiderata of recommender systems—both commercial and noncommercial—to include values furthering the democratic well-being and the cultural and social development of contemporary societies. Building on our recent work developing a metric to measure two PSM principles, commonality and diversity, in recommender systems, we propose a more comprehensive research program toward incorporating such principles into the design of recommender systems for cultural content, inviting the workshop to address how normative goals might transform processes of algorithmic amplification. Our proposed paper is a substantial expansion of our published work on public service media principles and the algorithmic curation of cultural goods (e.g., music, film, and literature). We are eager to share this collaboration with the symposium attendees and receive feedback.

Bridging Systems: Open problems for countering destructive divisiveness in ranking, recommenders, and governance

Aviv Ovadya and Luke Thorburn

Divisiveness appears to be increasing in much of the world, leading to concern about political violence and a decreasing capacity to collaboratively address large-scale societal challenges. In this working paper, we aim to articulate an interdisciplinary research and practice area focused on what we call bridging systems: systems which increase mutual understanding and trust across divides, creating space for productive conflict, deliberation, or cooperation. We give examples of bridging systems across three domains: recommender systems on social media, software for conducting civic forums, and human-facilitated group deliberation. We argue that these examples can be more meaningfully understood as processes for *attention-allocation* (as opposed to "content distribution" or "amplification") and develop a corresponding framework to explore similarities—and opportunities for bridging—across these seemingly disparate domains. We focus particularly on the potential of *bridging-based ranking* to bring the benefits of offline bridging into spaces which are already governed by algorithms. Throughout, we suggest research directions that could improve our capacity to incorporate bridging into a world increasingly mediated by algorithms and artificial intelligence.

Recommenders with Values: Developing recommendation engines in a public service organization

Alessandro Piscopo, Lianne Kerlin, North Kuras, James Fletcher, Calum Wiggins, Anna McGovern, and Megan Stamper

The BBC is the world's largest public service broadcaster. It reaches every week more than 80 percent of the U.K.'s adult population and 279 million people worldwide. In order to ensure that our audiences get the most engaging experience, our team develops recommender systems which aim to provide users with the most relevant pieces of content among the thousands the BBC publishes every day. All BBC output should serve the organization's mission to "act in the public interest, serving all audiences through the provision of impartial, high-quality and distinctive output and services which inform, educate, and entertain." Recommendations make no exception and, since they determine what our audiences see, they are in effect editorial choices at scale. How can we ensure that our recommendations are consistent with our mission and public service values, avoiding some of the harmful effects which might be associated with recommenders? In addressing this question, we identified two main challenges: (i) *methodological challenges*: Public service values are hard to pin down into a specific metric, therefore we have no clearly defined optimization function for our recommenders; (ii) *cultural/operational challenges*: Domain knowledge around public service values sits with our editorial staff, whereas data scientists are the recommendations specialists. We need to create a shared understanding of the problem and a common language to describe objectives and solutions across data science and editorial. Our paper describes the approach we devised to tackle these challenges, presenting a use case from our work on a BBC product, and reporting the lessons learned.

PARTICIPANTS

Joe B. Bak-Coleman is an associate research scholar at the Craig Newmark Center for Journalism Ethics and Security at Columbia University and a computational social scientist. He earned his Ph.D. in ecology and evolutionary biology at Princeton University in 2020, working with Iain Couzin and Dan Rubenstein, and recently completed a postdoctoral fellowship at the University of Washington Center for an Informed Public. His research focuses on understanding how collectives make decisions in the face of uncertainty. He's particularly interested in understanding what makes collective decision-making work and how it can go awry. Over the past decade, he has worked on collective decision-making in a range of contexts from animal groups and social media to metascience.

Twitter: @jbakcoleman

Fabian Baumann is a postdoc at the Max Planck Institute for Human Development, working in computational social science. His current focus is on how culture evolves as a collaborative process influenced by both social interactions and intelligent algorithms. In his research, he uses different methods, including dynamical systems modeling, natural language processing, as well as network science, which he studied during his Ph.D. in physics.

Twitter: @ftw_baumann

Luca Belli is the founder and CEO of Sator Labs; a University of California, Berkeley Tech Policy fellow; and a NIST AI visiting fellow. Previously, he was the co-founder and research lead for Twitter's Machine Learning, Ethics, Transparency, and Accountability team where he guided industry leading approaches for responsible machine learning practices and product changes. Before that, he operated as a data science and machine learning engineer at Conversant and WolframAlpha. His research interests lie at the intersection of feedback loops, algorithmic amplification (with a special eye on politics), and algorithmic audits. He holds a Ph.D. in mathematics from the University of Rome Tor Vergata.

Twitter: @__lucab; **Mastodon:** @__lucab@ macaw.social

James Bisbee is an assistant professor in Vanderbilt University's Department of Political Science and a faculty affiliate at the Data Science Institute. He is a political scientist who studies how public opinions and behaviors are influenced by information, ranging from local unemployment to elite cues, using all manner of empirical evidence. Previously, he was a postdoctoral fellow at New York University's Center for Social Media and Politics and at Princeton's Niehaus Center for Globalization and Governance. His research has been published in peer reviewed journals, including the American Political Science Review, the Journal of Politics, the Journal of Labor Economics, and International Organization, among others.

Twitter: @JamesBisbee

Richard Bonneau is a faculty research affiliate at New York University's Center for Social Media and Politics (CSMaP). His expertise in data science, leading large-scale systems biology consortia, motivates many contributions to the CSMaP lab. His experience with lab-based science, industry collaboration, and network science are key to the lab's innovative construction. Bonneau was selected by Discover magazine as one of the top 20 scientific minds under 40, and a review in the top biology journal, Cell, lists Bonneau's 2007 paper on the prediction of global dynamic regulatory networks as a landmark paper in the field of systems biology. Bonneau is a founding member of the Flatiron Institute, a new large-scale effort to create an intramural data science center at the Simons Foundation, a principal investigator on the initial Moore-Sloan data science environments grant, and part of the group of faculty at NYU that created the new Center for Data Science at NYU.

Twitter: @RichBonneauNYC

Georgina Born is a professor of anthropology and music in the Department of Anthropology, University College London. From 2010-2021, she was professor of music and anthropology in the Faculty of Music, University of Oxford, and from 2006-2010, she was professor of sociology, anthropology, and music at the University of Cambridge. Earlier, she had a professional life as a musician in experimental rock, jazz, and improvised music. She has held visiting professorships as follows: Bloch professor, University of California, Berkeley Department of Music (2014); Schulich distinguished visiting professor in music, McGill University (2015); visiting professor jointly in the Schools of Arts, Humanities, and Social Sciences at University of California, Irvine (2019-2022); professor II at the University of Oslo, Department of Musicology (2014-2019); visiting professor, Aarhus University, School of Communication and Culture (2017); Marie Curie senior research fellow at the Aarhus Institute of Advanced Studies (2018-2019); and distinguished global scholar, Department of Music, Princeton University (2020-2022).

Twitter: @georgieborn

PARTICIPANTS

William J. Brady is an assistant professor of management and organizations at Northwestern University. His research examines the dynamics of emotion at the social network level and their consequences for group behavior. His recent work studies how human psychology and technology-mediated social contexts interact to shape our emotions and intergroup attitudes. Combining tools of behavioral science and computational social science, his research aims to develop person-centered and design-centered interventions to improve our digital social interactions. He has been selected for the SAGE Emerging Scholar Award. Brady earned his B.A. in psychology and philosophy, with distinction, from the University of North Carolina at Chapel Hill, his Ph.D. in social psychology at New York University, and was awarded a postdoctoral fellowship from the National Science Foundation where he worked at Yale University.

Twitter: @william__brady; **Mastodon:** @williambrady@hci.social

Megan Brown is a senior research engineer and research scientist at New York University's Center for Social Media and Politics. As a research engineer, Brown collects and maintains large-scale collections of social media and digital trace data for the purposes of social science research. In her research endeavors, she studies cross-platform media manipulation, political bias in algorithmic systems, and the effect of platform governance and moderation policies on the spread of political content.

Twitter: @m_dot_brown; **Mastodon:** @m_dot_brown@mastodon.social

Jason Burton is an assistant professor at Copenhagen Business School and an Alexander von Humboldt Research fellow at the Max Planck Institute for Human Development. His research uses computational methods to study human behavior in the context of a digital society. This involves topics ranging from how people fundamentally reason with information and make decisions, to how the structure of online information environments influences the beliefs people form, to how online environments can be (re)designed to promote collective intelligence. Burton holds a Ph.D. in psychology from Birkbeck, University of London, and an M.S. in organizational psychiatry and psychology from King's College London.

Twitter: @jason_w_burton

Micah Carroll is an artificial intelligence Ph.D. student at University of California, Berkeley. He is broadly interested in ensuring that AI systems interfacing with humans will lead to beneficial individual and societal outcomes. Carroll is currently working on recommender systems, investigating how the choice of algorithm might affect users. Another interest is that of human-AI collaboration, with the goal of improving the quality and robustness of human models and of agents trained to collaborate with humans.

Twitter: @MicahCarroll

M. J. Crockett is an associate professor at Princeton University's Department of Psychology and University Center for Human Values. Prior to joining Princeton, Crockett was an associate professor of psychology at Yale University, associate professor of experimental psychology at the University of Oxford, and a fellow of Jesus College. They hold a B.S. in psychobiology from University of California, Los Angeles, and a Ph.D. in experimental psychology from the University of Cambridge and completed a Henry Wellcome Trust postdoctoral fellowship with economists and neuroscientists at the University of Zurich and University College London.

Twitter: @mollycrockett

Fernando Diaz is a research scientist at Google. Diaz's research focuses on the design of information access systems, including search engines, music recommendation services, and crisis response platforms. He is particularly interested in understanding and addressing the societal implications of artificial intelligence (AI) more generally. Previously, Diaz was the assistant managing director of Microsoft Research Montréal, where he also led FATE Montréal, and a director of research at Spotify, where he helped establish its research organization on recommendation, search, and personalization. He received his B.S. in computer science from the University of Michigan, Ann Arbor, and his M.S. and Ph.D. from the University of Massachusetts at Amherst.

Twitter: @841io

Annie Dorsen is a theater director working at the intersection of algorithmic art and live performance. Her most recent piece, Prometheus Firebringer, will have its New York premiere in May. Prior algorithmic performance projects include Infinite Sun (2019), The Great Outdoors (2017), Yesterday Tomorrow (2015), A Piece of Work (2013), and Hello Hi There (2010). These pieces have been presented at numerous theaters and festivals world-wide, including at the Brooklyn Academy of Music (New York), Hebbel am Ufer (Berlin), the Holland Festival (Amsterdam), and Festival d'Automne (Paris). She has taught at University of Chicago and Bard College, and been a frequent guest lecturer at numerous universities and art schools. Dorsen is the recipient of a MacArthur Fellowship, the Spalding Gray Award, a Guggenheim Fellowship, a Foundation for Contemporary Arts Grant to

Artists Award, and the Herb Alpert Award in the Arts. She is currently completing her second year at NYU School of Law.

Twitter: @AnnieDorsen; **Mastodon:** @anniedorsen@mastodon.social

Anca Dragan is an associate professor at University of California, Berkeley. Dragan runs the InterACT lab, focusing on enabling robots to work with, around, and in support of people. Her goal is for robots to autonomously generate their behavior in a way that moves beyond functionality and formally accounts for interaction with humans. This combines optimal control, machine learning, and cognitive science, with applications in collaborative manipulation and autonomous driving. Dragan got her Ph.D. in the Robotics Institute at Carnegie Mellon University on planning intent-expressive robot motion.

Twitter: @ancadianadragan

Laura Edelson co-leads the Cybersecurity for Democracy project at New York University's Tandon School of Engineering where she is also a postdoctoral researcher. Edelson is a computer scientist with expertise in large online networks and extensive real-world experience building big data machine learning augmented systems. Her current research involves large-scale analysis of online paid political content on major platforms, such as Facebook, Google, and Twitter, and the development of methods to detect inauthentic content and fraudulent actors. She also currently serves as the chief technologist of the Antitrust Division of the Department of Justice. Previously, Edelson was a software engineer for Palantir and Factset. During her time in industry, her work focused on applied machine learning and big data.

Twitter: @LauraEdelson2; **Mastodon:** @whiskeyocelot@mastodon.social

Kevin Feng is a Ph.D. student in the Human Centered Design and Engineering Department at the University of Washington. His interests are at the intersection of social computing and human-centered machine learning. He is motivated by the observation that key stakeholders are increasingly at risk of agency loss from black-boxed but powerful technologies in sociotechnical systems. As such, his research strives to embed the reclamation of stakeholder agency as a vital and meaningful pillar of these systems. His work is supported by a 2022 Herbold Fellowship in data science and computation. He graduated from Princeton University with a B.S.E. in computer science and a minor in visual arts in 2021, where he was affiliated with the Center for Information Technology Policy and worked on tools to democratize web advertisement research at scale.

Twitter: @kjfeng_; **Mastodon:** @kjfeng@hci.social

James Fletcher leads on Responsible Data and AI for the BBC, ensuring that the use of data and AI/ML aligns with the BBC's values and legal and regulatory obligations. He looks after responsible AI/ML governance through the BBC's Machine Learning Engine Principles and checklist, building tools, capacity, and culture to turn principles into practice with teams throughout the organization.

Twitter: @radiojay

Camille François is on the adjunct faculty of Columbia University's School of International and Public Affairs, where she teaches a course on cyber conflict and information operations. She is the chair of the advisory board and former chief innovation officer at Graphika, where she led the company's work to detect and mitigate disinformation, media manipulation, and harassment. François was also previously the principal researcher at Jigsaw, an innovation unit at Google that builds technology to address global security challenges and protect vulnerable users. She served as a special advisor to the chief technology officer of France in the prime minister's office. François is a Mozilla fellow, a Berkman Klein Center affiliate, and a Fulbright scholar. She holds a master's degree in human rights from the French Institute of Political Sciences and a master's degree in international security from the School of International and Public Affairs at Columbia University.

Twitter: @camillefrancois

Brett Frischmann is the Charles Widger Endowed University professor in law, business, and economics at Villanova. His interdisciplinary research on the relationships between infrastructural resources, governance, commons, and spillovers includes a series of foundational books and articles. His 2018 book, *Re-Engineering Humanity*, co-authored with Evan Selinger, thoroughly examines various mechanisms for techno-social engineering of humans, including algorithmic amplification, as well as the normative conflict between engineered optimality and commitments to human freedom and pluralism. Frischmann is an affiliated scholar of the Center for Internet and Society at Stanford, an affiliated faculty member of the Vincent and Elinor Ostrom Workshop in political theory and policy analysis at Indiana University, and a trustee for the Nexa Center for Internet & Society, Politecnico di Torino. Frischmann also served as the Microsoft visiting professor of information and technology policy at Princeton University's Center for Information Technology Policy.

Twitter: @brettfrischmann

PARTICIPANTS

Sam Gill is CEO of the Doris Duke Foundation, a New York-headquartered, national philanthropic organization that operates five national grantmaking programs—in the performing arts, the environment, medical research, child and family well-being, and mutual understanding between communities—as well as Duke Farms and Shangri La, two centers that serve the public directly. Previously, Gill was senior vice president and chief program officer at the John S. and James L. Knight Foundation, where he oversaw more than $100 million in annual grantmaking, in addition to managing the Knight Foundation's research and assessment portfolio and its grants administration function. He was also vice president of Freedman Consulting, LLC. Gill served on the board of the Philip and Patricia Frost Museum of Science in Miami and on the Commission on the Practice of Democratic Citizenship. He attended the University of Chicago and the University of Oxford, where he was a Rhodes Scholar.

Twitter: @thesamgill

Tarleton Gillespie is a senior principal researcher at Microsoft Research New England, part of the Social Media Collective. He also retains an affiliated associate professor position with Cornell University, in the Department of Communication and the Department of Information Science. Gillespie's current work investigates how social media platforms and other algorithmic information systems shape public discourse. His most recent book is *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media (*Yale, 2018)

Twitter: @TarletonG; **Mastodon:** @tarleton@mastodon.social

Katy Glenn Bass is the research director of the Knight First Amendment Institute at Columbia University. She is responsible for conceptualizing and executing all of the Institute's research initiatives, including the production of scholarship and research materials, the organization of conferences and symposia, and the Institute's Visiting Research Scholars program. Prior to joining the Institute, Glenn Bass worked at PEN America, where she supervised the production of reports analyzing free expression issues. She has also taught at NYU Law's Center for Constitutional Transitions and at the Walter Leitner International Human Rights Clinic at Fordham Law School. She holds a B.A. from Princeton University and a J.D. from Harvard Law School, where she received the Kaufman Pro Bono Service Award.

Twitter: @KGlennBass; **Mastodon:** @kgb@mas.to

Ravi Iyer is the managing director of the Psychology of Technology Institute, which is a project of University of Southern California's Neely Center. Before that, he led data science, research, and product teams across Facebook toward improving the societal impact of social media. He helped co-found and build the initial algorithms for Ranker.com, which is a profitable publisher of crowdsourced lists that serves tens of millions of unique visitors monthly and employs 125-plus people. He has a Ph.D. in social psychology from the University of Southern California and has co-authored dozens of highly cited empirical articles about individual values, political opinions, polarization, and technology. He has worked with numerous nonprofits fighting polarization. His work across tech and academia has been featured in The Wall Street Journal, The New York Times, The Atlantic, South by Southwest, and numerous other venues.

Twitter: @RaviDataIyer; **Mastodon:** @riyer@techhub.social

Joshua Conrad Jackson is a postdoctoral researcher at the Kellogg School of Management at Northwestern University. He will join the Booth School of Business at the University of Chicago as an assistant professor of behavioral science in 2023. He is interested in how cultural evolution shapes psychology, particularly the way that we perceive our minds and other people's minds. There is now well-documented variation across history and culture in how people perceive and communicate about emotion, morality, prejudice, religious belief, and many other forms of social cognition. He believes that models of cultural evolution can help social scientists understand how this variation occurs, and how it may influence the development of political movements, intergroup conflicts, and economic systems. In his research, he has tested basic questions about the cultural evolution of the mind, written about underused methods of studying cultural evolution, and shown how basic findings about cultural variation could have implications for policy and intergroup relations.

Twitter: @josh_c_jackson

Jameel Jaffer is the executive director of the Knight First Amendment Institute at Columbia University. Under his leadership, the Institute has filed precedent-setting litigation, undertaken major interdisciplinary research initiatives, and become an influential voice in debates about the freedoms of speech and the press in the digital age. Until August 2016, Jaffer served as deputy legal director at the ACLU, where he oversaw the organization's work on free speech, privacy, technology, national security, and international human rights. Jaffer's recent writing has appeared in The New York Times, The New Yorker, The Washington Post, and the Yale Law Journal Forum. He is an executive editor of Just Security,

a national security blog, and his most recent book, *The Drone Memos*, was one of The Guardian's "Best Books of 2016." Jaffer is a graduate of Williams College, Cambridge University, and Harvard Law School, where he was an editor of the Harvard Law Review. He served as a law clerk to the Honorable Amalya L. Kearse of the U.S. Court of Appeals for the Second Circuit, and then to the Right Honorable Beverley McLachlin, Chief Justice of Canada.

Twitter: @JameelJaffer; **Mastodon:** @jameeljaffer@mastodon.social

Benjamin Kaiser is a Ph.D. candidate in computer science at Princeton University in the Center for Information Technology Policy. His research spans topics in usable security, privacy, and information integrity. Before attending Princeton, Kaiser worked at the Massachusetts Institute of Technology Lincoln Laboratory, where he conducted applied security and privacy research. Kaiser holds a B.S. and M.S. in computer science from Rensselaer Polytechnic Institute.

Twitter: @benhkaiser

Daphne Keller is the director of the Program on Platform Regulation at Stanford's Cyber Policy Center. Her work focuses on platform regulation and internet users' rights. She has published both academically and in the popular press, testified and participated in legislative processes, and taught and lectured extensively. Her recent work focuses on legal protections for users' free expression rights when state and private power intersect, particularly through platforms' enforcement of terms of service or use of algorithmic ranking and recommendations. Until 2020, Keller was the director of intermediary liability at Stanford's Center for Internet and Society. She also served until 2015, as associate general counsel for Google, where she had primary responsibility for the company's search products. Keller has taught internet law at Stanford, Berkeley, and Duke law schools. She is a graduate of Yale Law School, Brown University, and Head Start.

Twitter: @daphnek; **Mastodon:** @daphnehk@mastodon.social

Lianne Kerlin is a lead researcher in the BBC's Research and Development team, where she is passionate about enhancing experiences in a digital world. Her research interests lie largely in understanding the impact of digital technologies on people and societies, spanning areas around human values, ethics, and metrics. Most recently, she has been working in the BBC's Product Group on a project to define public service value metrics for the BBC's digital product portfolio.

Twitter: @lianne_k123

North Kuras is a senior user experience and design architect in the BBC's User Experience and Design team. Before joining the BBC, he ran a small agency that specialized in digital and print design for charities and not-for-profits. He is interested in the technical and human systems that underpin great user experiences.

Angela Lai is a Ph.D. candidate in New York University's Center for Data Science and a graduate research associate at New York University's Center for Social Media and Politics. Her research interests include natural language processing, network analysis, and political and social behavior and its interaction with social media.

Twitter: @angelaight

Benjamin Laufer is a Ph.D. candidate at Cornell Tech. He studies the values and politics embedded in technological systems, particularly those deployed in high-impact, high-complexity domains in the public realm. He is a doctoral fellow at the Digital Life Initiative and an affiliate of the Artificial Intelligence, Policy, and Practice group at Cornell University. He is advised by Helen Nissenbaum and Jon Kleinberg. He holds a B.S.E. (cum laude) in operations research and financial engineering from Princeton University.

Twitter: @BenDLaufer; **Mastodon:** @benlaufer@mastodon.social

Seth Lazar is a professor of philosophy at the Australian National University, an Australian Research Council future fellow, and a distinguished research fellow of the University of Oxford Institute for Ethics in AI. He has worked on the ethics of war, self-defense, and risk and now leads the Machine Intelligence and Normative Theory Lab, where he directs research projects on normative philosophy of computing. He was general co-chair for the Association for Computing Machinery (ACM) Fairness, Accountability, and Transparency conference 2022; program co-chair for the ACM/Association for the Advancement of Artificial Intelligence's AI, Ethics, and Society conference in 2021; and is one of the authors of a study by the U.S. National Academies of Science, Engineering, and Medicine on the ethics and governance of responsible computing research. He gave the 2022 Mala and Solomon Kamm lecture in ethics at Harvard University and the 2023 Tanner Lectures on AI and human values at Stanford University.

Twitter: @sethlazar; **Mastodon:** @sethlazar@mastodon.social

PARTICIPANTS

Tomo Lazovich is a senior research scientist at the Institute for Experiential AI at Northeastern University. They were a senior machine learning researcher at Twitter, developing a suite of metrics to measure inequality in outcomes for the Machine Learning, Ethics, Transparency, and Accountability team. Lazovich worked as Machine Learning team lead at Lightmatter, adapting algorithms to the computer hardware startup's novel photonics-based hybrid digital-analog architecture, and at nonprofit engineering company Draper, building deep learning architectures for the identification and repair of bugs in source code as part of the DARPA-funded Mining and Understanding Software Enclaves program. Lazovich holds a Ph.D. in physics from Harvard University, where their thesis was based on the discovery and subsequent study of the Higgs Boson particle at the Large Hadron Collider in Switzerland. In addition, they are also currently a part-time J.D. candidate at Northeastern University, hoping to fuse their technical knowledge with legal expertise to build practical regulatory solutions for AI.

Twitter: @laughsovich; **Mastodon:** @laughsovich@mastodon.social

Björn Lindström is a senior researcher and principal investigator at the Department of Clinical Neuroscience, Karolinska Institute, Stockholm, Sweden. Prior to joining the Karolinska Institute, he was a tenured assistant professor at Vrije Universiteit Amsterdam. He was also a postdoctoral scholar at the University of Amsterdam and at the Center for Neuroeconomics at the University of Zurich. His research investigates the psychological, computational, and neural mechanisms of social learning and cultural evolution. His research is funded by a starting grant from the European Research Council and a Wallenberg Academy Fellowship, awarded by the Knut and Alice Wallenberg Foundation.

Twitter: @B_Lindstroem

Irene Lo is an assistant professor in management science and engineering at Stanford University. Her research is on designing matching markets and assignment processes to improve market outcomes, with a focus on public sector applications and socially responsible operations research. She is also interested in mechanism design for social good and graph theory. She was previously a postdoctoral fellow at the Economics Department at Stanford University. She obtained her Ph.D. in operations research from Columbia University and graduated from Princeton University with an A.B. in mathematics.

Philipp Lorenz-Spreen is a research scientist at the Max Planck Institute for Human Development. He is a network scientist researching the self-organized online discourse and how to empower democratic and autonomous decision-making through platform design and boosting. His aim is to better understand the interplay between human behavior and the connectivity and functioning of online platforms, in particular, not only how this affects our public discourse and thus our democracy but also to question the current status quo and explore how this technology offers untapped opportunities for an improved information landscape and participatory democracy lived online. Previously, he did his Ph.D. at the Technical University Berlin on empirical methods and theoretical models to describe the dynamics of collective attention from online data sets. At the Ludwig Maximilian University of Munich, he studied physics with a focus on systems biophysics.

Twitter: @lorenz_spreen; **Mastodon:** @lorenz_spreen@social.dev-wiki.de

Kristian Lum is an associate research professor at the University of Chicago's Data Science Institute. She has worked at Twitter as part of its Machine Learning Ethics, Transparency, and Accountability Team; the Human Rights Data Analysis Group; the University of Pennsylvania; and the Virginia Bioinformatics Institute at Virginia Institute of Technology. She is a founding member of the executive committee of the ACM conference on Fairness, Accountability, and Transparency and was named an emerging leader in statistics by the Committee of Presidents of Statistical Societies and a Kavli fellow by the National Academy of Sciences.

Twitter: @KLdivergence; **Mastodon:** @KLdivergence@mastodon.social

Jonathan Mayer is an assistant professor at Princeton University, where he holds appointments in the Department of Computer Science and the Princeton School of Public and International Affairs. Before joining the Princeton faculty, he served as the technology law and policy advisor to U.S. Sen. Kamala Harris and as the chief technologist of the Federal Communications Commission Enforcement Bureau. Mayer's research centers on the intersection of technology and law, with emphasis on national security, criminal procedure, and consumer privacy. He is both a computer scientist and a lawyer, and he holds a Ph.D. in computer science from Stanford University and a J.D. from Stanford Law School.

Twitter: @jonathanmayer; **Mastodon:** @jonathanmayer@mastodon.social

David McDonald is a professor in the Department of Human Centered Design and Engineering at University of Washington. McDonald has ongoing projects studying collaboration in Wikipedia, social matching, and fostering healthy behavior all through systems that interleave computation with human activity. He has published research on ubiquitous sensing for behavior change, collaboration in distributed contributor systems, collaborative authoring, recommendation systems, and public use of large screen displays. His general research interests

span computer-supported cooperative work and human-computer interaction. McDonald earned his Ph.D. in information and computer science at the University of California, Irvine (UC, Irvine). At UC, Irvine, he was part of the Computing, Organizations, Policy and Society group. He worked at FX Palo Alto Laboratory in the Personal and Mobile technology group and at AT&T Labs, Human Computer Interaction group. McDonald recently finished serving as a program director for the Human Centered Computing, Network Science and Engineering, and Social Computational Systems programs at the National Science Foundation.

Twitter: @dwmcphd

Anna McGovern is a lead automated curation specialist at BBC Product Group. She works with editorial teams in BBC News, BBC iPlayer, BBC Sport, BBC Sounds, and BBC World Service on machine learning and automation projects. She also works on codifying editorial wisdom into the BBC's algorithms so that content is surfaced in rich, interesting, and appropriate ways that is valued by the audience and meets the BBC's commitment to responsible machine learning. Additionally, McGovern also focuses on evaluating and testing automatically curated output (such as recommendations and personalization) within various contexts, which requires assessment of complex editorial matters, compliance of legal issues, and regulatory responsibilities as outlined by Ofcom.

Samia Menon is an undergraduate at Columbia University's Fu Foundation School of Engineering and Applied Sciences.

Smitha Milli is a postdoctoral associate at Cornell Tech. They received their B.S. and Ph.D. in electrical engineering and computer science from University of California, Berkeley.

Twitter: @SmithaMilli

Mor Naaman is a professor of information science at Cornell Tech where he currently serves as an associate dean for technical programs. His research area is technology, media, and democracy, with a particular focus on the trustworthiness of our information ecosystem. His lab uses a wide range of tools—from machine learning, to computational social science, to online experiments, to qualitative methods—to understand and try to address these challenges, with an increasing focus on the impact and role of AI-mediated communication. Previously, Naaman was on the faculty at Rutgers SC&I, led a research team at Yahoo! Research Berkeley, received a Ph.D. in computer science from Stanford University, and played professional basketball for Hapoel Tel Aviv. He is a recipient of a NSF Early Faculty CAREER Award, research awards and grants from numerous corporations including Microsoft and Google, and multiple best paper awards.

Twitter: @informor; **Mastodon:** @mor@hci.social

Jonathan Nagler is a professor of politics and a co-director of New York University's Center for Social Media and Politics. He received his B.A. in government from Harvard University in 1982, and his Ph.D. from California Institute of Technology (Caltech) in 1989. He has been a visiting associate professor at Caltech and Harvard and has taught at the Summer Program, European Consortium for Political Research, Essex University, England, and the Summer Program, Inter-university Consortium for Political and Social Research, University of Michigan, as well as the Economic and Social Research Council Oxford Spring School in quantitative methods for social research. In 2012, Nagler was a Fernand Braudel fellow at the European University Institute. Nagler's research focuses on voting and elections.

Twitter: @Jonathan_Nagler

Arvind Narayanan is a professor of computer science at Princeton University. He co-authored a textbook on fairness and machine learning and is currently co-authoring a book on AI snake oil. He led the Princeton Web Transparency and Accountability Project to uncover how companies collect and use our personal information. His work was among the first to show how machine learning reflects cultural stereotypes, and his doctoral research showed the fundamental limits of de-identification. Narayanan is a recipient of the Presidential Early Career Award for Scientists and Engineers (PECASE), twice a recipient of the Privacy Enhancing Technologies Award, and thrice a recipient of the Privacy Papers for Policy Makers Award.

Twitter: @random_walker; **Mastodon:** @randomwalker@mastodon.social

Helen Nissenbaum is a professor of information science and founding director of the Digital Life Initiative at Cornell Tech. Her work focuses on ethical and political implications of digital technologies on issues such as privacy, bias in digital systems, trust online, ethics in design, and accountability in computational and algorithmic systems. Grants from the National Science Foundation, Air Force Office of Scientific Research, the U.S. Department of Health and Human Services Office of the National Coordinator, McArthur Foundation, Defense Advanced Research Projects Agency (DARPA), and the National Security Agency have supported her research. Recipient of the 2014 Barwise Prize of the American Philosophical Association and the International Association of

PARTICIPANTS

Computing and Philosophy Covey Award for computing, ethics, and philosophy, Nissenbaum has contributed to privacy-enhancing free software, TrackMeNot (protecting against profiling based on web searches) and AdNauseam (protecting against profiling based on ad clicks). She holds a Ph.D. in philosophy from Stanford University and a B.A. (Hons) in philosophy and mathematics from the University of the Witwatersrand, South Africa.

Twitter: @HNissenbaum

Paul Ohm is a professor of law at the Georgetown University Law Center. Ohm builds bridges between computer science and law, utilizing his training and experience as a lawyer, policymaker, computer programmer, and network systems administrator. His research focuses on information privacy, computer crime law, surveillance, technology and the law, and artificial intelligence and the law. Ohm serves as a faculty director for the Institute for Technology Law and Policy, the Center on Privacy and Technology at Georgetown Law, and the Georgetown University Tech & Society Initiative. He has testified before committees of both houses of Congress and advised numerous government agencies, including the Federal Trade Commission, Federal Communications Commission, and several state attorneys general. He serves on the boards of directors for two of the most impactful and innovative nonprofits in technology policy: Upturn and The Markup.

Twitter: @paulohm

Aviv Ovadya is affiliated with Harvard's Berkman Klein Center, a visiting researcher at Cambridge University's Center for the Future of Intelligence, and consults for technology companies, civil society organizations, and funders. He researches and supports approaches for technology governance, building on insights from offline deliberative democracy and AI-augmented deliberation technology. After receiving his B.S. and M.Eng. degrees in computer science from MIT and a stint in industry, Ovadya transitioned to focusing exclusively on the societal implications of technology, through academic fellowships, working in roles such as the chief technologist for the Center for Social Media Responsibility (University of Michigan), and consulting for and advising organizations, including the Partnership on AI, the Civic Health Project, and Cohere. His work has been covered by NPR, The Economist, and The New York Times, and his writing has been published by Bloomberg, Harvard Business Review, the MIT Technology Review, and The Washington Post.

Twitter: @metaviv; **Mastodon:** @aviv@mastodon.online

Sashrika Pandey is an undergraduate researcher at University of California, Berkeley, interested in studying human-AI systems and recommender systems.

Sahil Patel is a graduate student at Columbia University's Fu Foundation School of Engineering and Applied Sciences.

Alessandro Piscopo is a lead data scientist in the BBC's Product Group. His team, Datalab, develops recommendation engines across the organization and has so far deployed live recommenders on products such as BBC Sounds, BBC World Service, and the BBC News app. His research interests lie in the area at the intersection of peer-production communities, collaborative knowledge engineering, data quality, and responsible AI. He received his Ph.D. from the University of Southampton in 2019.

Twitter: @aliossandro; **Mastodon:** @aliossandro@mastodon.social

Helena Puig Larrauri is strategy lead and co-founder of Build Up. She has over a decade of experience advising and working with civil society actors and multilateral organizations in conflict contexts and polarized environments. She specializes in the integration of digital technology and innovation processes to peace processes and civic dialogues, and in understanding how conflict drivers show up in digital spaces. She has published and spoken on these subject matters extensively. She is also an advisor on digital technologies and mediation to the United Nations Mediation Support Unit and an Ashoka fellow. She previously served on the boards of Elva Community Engagement, International Alert, Impact Hub Barcelona, and the Standby Task Force, and she currently serves on the boards of Public Sentiment and Digital Peace Now. She holds a B.A. in politics, philosophy, and economics from Oxford University and a M.A. in public policy (economics) from Princeton University.

Twitter: @helenapuigl

Inioluwa Deborah Raji is a Ph.D. student at the University of California, Berkeley, and Mozilla fellow. She is interested in questions on algorithmic auditing and evaluation. In the past, she worked closely with the Algorithmic Justice League initiative to highlight bias in deployed AI products. She has also worked with Google's Ethical AI team and been a research fellow at the Partnership on AI and AI Now Institute at New York University, working on various projects to operationalize ethical considerations in ML engineering practice. Recently, she was named to Forbes 30 Under 30 and MIT Tech Review 35 Under 35 Innovators.

PARTICIPANTS

Twitter: @rajiinio

Yoel Roth is a technology policy fellow at the University of California, Berkeley, Goldman School of Public Policy and the former head of trust & safety at Twitter. For more than seven years, he led the teams responsible for Twitter's content moderation, integrity, and platform security efforts, including policy development, threat investigation, product, design, research, and operations. His current research and writing focus on the trust and safety industry and how technology companies manage the conflicting values and incentives built into content moderation at scale. Before joining Twitter, Yoel received his Ph.D. from the Annenberg School for Communication at the University of Pennsylvania. His research and teaching examined the technical, policy, business, and cultural dynamics of social networking and dating apps at the dawn of the "App Store" age.

Mastodon: @yoyoel@macaw.social

Megan Stamper is head of Data Science for the BBC's Product Group, the part of the BBC that builds and delivers audience-facing products such as iPlayer. Stamper has been a data scientist working in the media sector throughout her career, from the Financial Times to the BBC, focused on how to use data and machine learning to build the best digital experiences for millions of users worldwide.

Twitter: @megzaz; **Mastodon:** @megzaz@mastodon.scot

Jonathan Stray is a senior scientist at the Center for Human Compatible AI at University of California, Berkeley, where he works on the design of recommender systems for better personalized news and information. He previously taught the dual master's degree in computer science and journalism at Columbia University, built several pieces of software for investigative journalism, worked as an editor at the Associated Press, and developed graphics algorithms for Adobe. He holds a M.S. in computer science from the University of Toronto and a M.A. in journalism from the University of Hong Kong.

Twitter: @jonathanstray; **Mastodon:** @jonathanstray@mastodon.social

Luke Thorburn is a Ph.D. candidate in Safe & Trusted AI at King's College London. His core research interests lie at the intersection of AI, epistemology, and conflict. He tries to understand how technology can be used to reduce conflict risks and improve collective decision-making, as well as to better understand the extent to which these goals are important. He co-authors the Understanding Recommenders project at the Center for Human-Compatible AI at University of California, Berkeley, has worked with Ofcom on the evaluation of recommender systems, and contributed to the Prefiguration Working Group for the International Observatory on Information and Democracy. His background is in probability and statistics.

Twitter: @LukeThorburn_; **Mastodon:** @lukethorburn@social.coop

Joshua A. Tucker is professor of politics, director of New York University's Jordan Center for Advanced Study of Russia, and co-director of the New York University Center for Social Media and Politics. His research focuses on the intersection of social media and politics, including partisan echo chambers, online hate speech, the effects of exposure to social media on political knowledge, online networks and protest, disinformation and fake news, how authoritarian regimes respond to online opposition, and Russian bots and trolls. He is the co-chair of the external academic team for the U.S. 2020 Facebook & Instagram Election Study and serves on the advisory board of the American National Election Study and the Comparative Study of Electoral Systems. His most recent books are the co-authored *Communism's Shadow: Historical Legacies and Contemporary Political Attitudes* (Princeton University Press, 2017) and the co-edited *Social Media and Democracy: The State of the Field* (Cambridge University Press, 2020).

Twitter: @j_a_tucker; **Mastodon:** @j_a_tucker@mastodon.social

Yike Wang is an undergraduate researcher at University of California, Berkeley.

Calum Wiggins is executive product manager for Recommendations, part of the team that builds content recommenders for the BBC's audience-facing products, such as BBC iPlayer and BBC Sounds. In his career to date, Wiggins has worked on a broad range of products, from technical platforms to user interfaces, in both commercial and not-for-profit organizations.

Marlena Wisniak is senior legal consultant at the European Center for Not-for-Profit Law, where she leads advocacy and policy in technology and human rights with a focus on algorithmic systems. Until recently, she oversaw content governance on Twitter's legal team and had previously led the civil society and academic portfolios at the Partnership on AI. As a lawyer, Wisniak has over a decade of experience working globally across sectors at the intersection of technology and human rights, with expertise in emerging technologies and business and human rights. Admitted to the Swiss Bar, Wisniak graduated from

the University of Geneva and Stanford Law School.

Twitter: @marle_wi

Amy X. Zhang is an assistant professor at University of Washington's Allen School of Computer Science and Engineering, where she leads the Social Futures Lab, a group dedicated to reimagining social and collaborative systems to empower people and improve society. She has been a Google Research scholar, a Belfer fellow at the Anti-Defamation League, a Berkman Klein fellow, a Google Ph.D. fellow, and a National Science Foundation graduate research fellow. She is a visiting researcher on the Semantic Scholar team at AI2, and prior to UW, she was a Stanford postdoctoral researcher after completing a Ph.D. at the Massachusetts Institute of Technology (MIT) Computer Science and Artificial Intelligence Laboratory, where she received the George Sprowls Best Thesis Award at MIT in computer science. She received a Master of Philosophy in computer science at the University of Cambridge on a Gates Fellowship and a B.S. in computer science at Rutgers University, where she was captain of the Division I Women's tennis team.

Twitter: @amyxzh; **Mastodon:** @axz@hci.social

NEARBY LUNCH OPTIONS

During Friday's lunch break (12:35 pm - 2:00 pm), check out our favorite spots for quick meals near campus. We've also included sit-down options for dinner in the neighborhood.

Quick Meals

Ban Ban Shop
(Banh mi, bao buns, and more) – 2-minute walk
2911 Broadway, New York, NY 10025 | Broadway between 113th & 114th

Junzi Kitchen
(Chinese savory pancakes, noodles, and more) – 2-minute walk
2896 Broadway, New York, NY 10025 | Broadway & 113th

Milano Market Westside
(Sandwiches) – 2-minute walk
2892 Broadway, New York, NY 10025 | Broadway between 112th & 113th

Sweetgreen
(Salads) – 2-minute walk
2937 Broadway, New York, NY 10025 | Broadway & 115th

Wu and Nussbaum
(Bagels, dumplings, and more) – 2-minute walk
2897 Broadway, New York, NY 10025 | Broadway & 113th

Pret A Manger
(Cafe) – 3-minute walk
2955 Broadway, New York, NY 10025 | Broadway & 116th

Shake Shack
(Burgers) – 3-minute walk
2957 Broadway, New York, NY 10025 | Broadway & 116th

Hamilton Deli
(Sandwiches) – 5-minute walk
1129 Amsterdam Ave, New York, NY 10025 | Amsterdam between 115th & 116th

Hula Poke
(Poke) – 7-minute walk
1028 Amsterdam Ave, New York, NY 10025 | Amsterdam between 110th & 111th

Roti Roll
(Indian street food) – 9-minute walk
994 Amsterdam Ave, New York, NY 10025 | Amsterdam & 109th

Falafel on Broadway
(Falafel) – 11-minute walk
3151 Broadway, New York, NY 10027 | Broadway between La Salle & Tiemann

In addition, there are food trucks offering a variety of options on Broadway near 116th Street.

Sit-Down Meals

Community Food & Juice
(Organic) – 3-minute walk
2893 Broadway, New York, NY 10025 | Broadway between 112th & 113th

Le Monde
(French) – 3-minute walk
2885 Broadway, New York, NY 10025 | Broadway between 112th & 113th

Atlas Kitchen
(Chinese) – 6-minute walk
258 W 109th St, New York, NY 10025 | Broadway & 109th

Tartina
(Italian) – 7-minute walk
1034 Amsterdam Ave, New York, NY 10025 | Amsterdam & 111th

V & T
(Italian) – 7-minute walk
1024 Amsterdam Ave, New York, NY 10025 | Amsterdam between 110th & 111th

Massawa
(Eritrean and Ethiopian) – 10-minute walk
1239 Amsterdam Ave, New York, NY 10027 | Amsterdam & 121st

Taqueria Y Fonda La Mexicana
(Mexican) – 10-minute walk
968 Amsterdam Ave, New York, NY 10025 | Amsterdam between 107th & 108th

Bánh
(Vietnamese) – 11-minute walk
942 Amsterdam Ave, New York, NY 10025 | Amsterdam between 106th & 107th

Pisticci
(Italian) – 11-minute walk
125 La Salle St, New York, NY 10027 | La Salle between Claremont & Broadway

Oliva Tapas Españolas
(Tapas) – 14-minute walk
3229 Broadway, New York, NY 10027 | Broadway between 129th & 130th

The Knight First Amendment Institute at Columbia University defends the freedoms of speech and the press in the digital age through strategic litigation, research, and public education. It promotes a system of free expression that is open and inclusive, that broadens and elevates public discourse, and that fosters creativity, accountability, and effective self-government.

The Knight Institute is grateful for the support of Arnold Ventures; the Carnegie Corporation of New York; Columbia University; Abby and Rick Elbaum; the Ford Foundation; the William and Flora Hewlett Foundation; the John S. and James L. Knight Foundation; the Charles Koch Foundation; the John D. and Catherine T. MacArthur Foundation; Netgain Partnership, a project of Media Democracy Fund; The Omidyar Group; Open Society Foundations; the Pritzker Pucker Family Foundation; the Bernard and Anne Spitzer Charitable Trust; and the Zegar Family Foundation.

If you plan on posting about the symposium on social media, please include the hashtag #OptimizingForWhat.

knightcolumbia.org

The Knight Institute would also like to thank the Brown Institute for Media Innovation and Columbia Law School for their support and partnership.

Design: Astrid Da Silva / Illustration © Emilie Flamme

KNIGHT
FIRST AMENDMENT
INSTITUTE at
COLUMBIA UNIVERSITY