THE STANFORD INTERNET OBSERVATORY

SIO 2022 Annual Report

Stanford
Internet Observatory
Cyber Policy Center

Photo: Aaron Kehoe

Table of Contents

Opening Letter

Our third year was a busy one for the Stanford Internet Observatory (SIO). We are an adaptive research lab, bringing the learnings from disparate academic fields together with the valuable experience of industry professionals and the critical advocacy of civil society organizations. In addition to our regular research outputs, including peer-reviewed academic papers, books, and blog posts, this year we launched new projects to build out the field of trust and safety. The **Journal of Online Trust and Safety** and the **Trust and Safety Research Conference** promote new research from academia, civil society, and industry in a rigorous but rapid peer review process. Investigators can apply our **Open-Sourced Tooling** to new and innovative projects. Our **Data Sets** similarly provide researchers with new source material for future studies. Finally, the framework developed in our **Election Integrity Partnership** forms a new model for collaborative cross-platform social media research on information interference.

On the research side, we continue to focus on four key areas—trust and safety, platform policy, information interference, and emerging technology. This year we published a major report summarizing the work of our **Virality Project**, which analyzed online discourse about COVID-19 vaccines over a seven-month period. We released white papers analyzing four growing **Alt-Tech Platforms**. We continued our collaborations with external partners at academic institutions and think tanks. Our researchers published 10 peer-reviewed journal articles, 10 reports analyzing platform takedowns of inauthentic networks, and 22 commentaries or op-eds. This tremendous output is detailed in the **Research** section of this report.

Our core team continues to grow as we diversify our research and projects. This year we saw off two fantastic postdoctoral scholars—Samantha Bradshaw, who joined the faculty at American University's School of International Service, and Josh Goldstein, who joined Georgetown University's Center for Security and Emerging Technology as a research fellow. We also proudly saw the departure of our research analyst Carly Miller, who joined the Oversight Board as a Data and Implementation Officer after completing her term at SIO.

We welcomed three new postdoctoral scholars who each bring in complementary and additive research projects. Karen Nershi joined us from the University of Pennsylvania, where her dissertation looked at cross-national anti-money laundering enforcement using data on cryptocurrency transactions. Ronald Robertson completed his PhD at Northeastern University, where he conducted research on algorithm audits and information seeking behavior in online search. Tongtong Zhang joined SIO in 2022 after completing her PhD at Stanford University. Her research studies online propaganda by authoritarian regimes. In addition to our researchers, we added two full-time staff members. Dan Bateyko joined us as a special projects manager, working on research directions for the study of online platforms, end-to-end encrypted messengers, and the Journal of Online Trust and Safety. John Perrino joined us as a policy analyst based in Washington D.C., where he leads SIO's policy engagement and communications efforts. Those efforts can be found in the **Policy** section of this report.

In 2023, we are delighted to welcome Professor **Jeffrey T. Hancock**, the Harry and Norman Chandler Professor of Communication and the founding director of the Stanford Social Media Lab, as our faculty director where he will continue to help guide and contribute to SIO research and scholarship.

As a program of Stanford University where many graduates go on to become technologists and executives, SIO is dedicated to educating the next generation to prevent the emergence of new harmful designs and products. We do this by teaching students about the mistakes of the past, how technology has amplified or accelerated societal trends, and how to spot similar abuses as they are reflected in new platforms. Our educational programs are explicitly cross-disciplinary: we teach public policy students how to hack systems and computer science students about how the social sciences must be considered in product design. We are working to make our courses freely available to students and professionals around the world and will be supporting other educational institutions with curricula and content as they build equivalent courses.

We welcome you to read this report featuring highlights of our outputs over the 2021–2022 academic year.

The Stanford Internet Observatory studies the tactics and tools by which online safety is undermined, and creates educational and technical responses to reduce harm.

By The Numbers

2022

COURSES

5

- CURRENT TOPICS IN TECHNOLOGY PLATFORM POLICY — 35 **students**
- HACK LAB — 100 **students**
- ONLINE OPEN SOURCE INVESTIGATIONS — 23 **students**
- THE POLITICS OF INTERNET ABUSE — 42 **students**
- TRUST AND SAFETY ENGINEERING — 156 **students**

356

STUDENTS ENROLLED IN SIO COURSES

38

STUDENT RAs WITH SIO

PUBLICATIONS

- 10 **Journal articles (3 with student co-authors)**
- 15 **White paper reports (13 with student co-authors)**
- 22 **Op-eds and news commentary pieces**
- 33 **Blog posts**
- 4 **Issues of the Journal of Online Trust and Safety**
  - **25 articles and 3 commentaries**
  - **53,290 total abstract views and 34,439 views**

5,398

MEDIA MENTIONS

Creating a Discipline in Online Trust and Safety

SIO set out to conduct research and teach best practices from industry, developing a new field of study and professional practice for trust and safety.

The technology industry is facing a major challenge in building trustworthy technology that lives up to the expectations of citizens, governments, and the media. At the same time, academics, policymakers, and concerned citizens struggle to understand the options available to deal with problems inherent in billion-user social networks and the interventions that could make a positive impact. “Trust and Safety” is the term used in Silicon Valley for the internal teams that research, find, and stop abuse of their systems, and it is the term the Stanford Internet Observatory uses to describe a developing field of academic research.

At its founding, the Stanford Internet Observatory set out to conduct research and teach best practices from industry, developing a new field of study and professional practice for trust and safety. Over the past year, the Stanford Internet Observatory launched three key initiatives aimed at building the research capacity around trust and safety and training the next generation of researchers: launching the Journal of Online Trust and Safety, hosting the first annual Trust and Safety Research Conference, and expanding our teaching curriculum. These initiatives were designed to incentivize more robust, empirical research on understanding and reducing online harms. This requires bringing together researchers across disciplines and uniting academic researchers with practitioners in industry.

Journal of Online Trust and Safety

Academic publishing has not caught up to trust and safety—a rapidly evolving field that focuses on emergent behaviors and harms. Academic journals silo researchers within their disciplines—such as political science, computer science, or sociology—and journals can average two years between submission and publication. Papers sit behind steep paywalls, inaccessible to those outside of academic institutions. Scholars from industry, government, and civil society find little incentive to contribute to journals. In all, the process discourages multidisciplinary research and limits the accessibility and impact of research findings.

SIO launched the **Journal of Online Trust and Safety** (JOTS) this year to solve these issues for the newly established field of trust and safety. The journal is **cross-disciplinary**, with **rapid review** and **open access**. The independent editorial board—whose members also serve as reviewers—includes leading experts in trust and safety working both in academia and as practitioners. The editorial board ensures that published work is relevant and impactful.

The Journal of Online Trust and Safety's first year has been a tremendous success. It published three general issues and one special issue with a total of 25 articles and three commentaries. Articles have been referenced in news outlets including The Guardian, The Washington Post, NBC News, The Verge, Tech Policy Press, and more.

Trust and Safety Curriculum

Over the past three years, SIO director Alex Stamos has refined his very popular Trust and Safety Engineering course, teaching over 250 computer science students about both the technical and sociological roots of online harms and how online providers have responded. This year, SIO expanded the trust and safety curriculum with a new course, The Politics of Internet Abuse. The class explores political science research on how governments and other entities exploit platform design vulnerabilities and how online platforms currently respond to these threats. Students work together across the computer science and political science courses to gain an understanding of the most pressing challenges in global communication platforms and a strong interdisciplinary foundation for future research and work on mitigating these harms.

Together, these courses provide students with conceptual backgrounds to understand 12 common types of online harms, drive them to develop and apply that knowledge with practical exercises on how to design products defensively, and inspire them to explore further research in online trust and safety. This curriculum has directly led students to launch research projects, connected them with industry roles in trust and safety, and informed students developing new products.

Now, SIO is expanding the curriculum beyond the Stanford campus by launching a teaching consortium and training professors and instructors at other institutions on how to implement the trust and safety curriculum. In addition to modular video lectures and an open access textbook for use in classrooms, SIO has developed professional training with a two-hour tabletop exercise that simulates a real-life data and security breach at a social media company.

The Trust and Safety Research Conference

SIO organized the first-ever research conference focused on trust and safety, a sold-out event held on the Stanford campus from September 29–30, 2022. Trust and safety has an information-sharing problem. Practitioners benefit from academic insights, but have neither the time nor resources to engage with cutting-edge scholarship. Meanwhile, academics studying trust and safety issues lack a common language, missing opportunities to draw connections across disciplines and forge new, promising areas of study. Researchers from academia and civil society struggle to gain access to the data needed to answer their research questions, and have little visibility into the professional world of trust and safety teams. **The Trust and Safety Research Conference** bridges this gap by bringing together researchers and showcasing emerging studies in the newly established field of trust and safety. More than 90 presenters from across academia, civil society, government, and the tech industry discussed their work and shared cutting-edge research rooted in real-world problems that move the field of trust and safety forward.

Teaching

This year we worked with nearly 40 student research analysts and continue to have two-thirds of our publications co-authored by students.

SIO thinks of its teaching mission both inside and outside of the classroom. Five academic courses at Stanford train students in both practical and theoretical topics around trust and safety, cybersecurity, technology policy, and online research techniques. SIO further mentors and trains students through its student analyst program, giving students the opportunity to work side by side with SIO's industry-leading researchers and allowing them to co-author and publish blogs, academic papers, and reports. Beyond the university, SIO's seminars, workshops, curriculum development, and simulation exercises make our pedagogy accessible to a broad range of academic, industry, and civil society partners.

SIO's alumni have gone on to careers in government at the U.S. Department of State and the Department of Homeland Security's Cybersecurity and Infrastructure Security Agency, and to serve Congress through the TechCongress fellowship program. In the private sector, SIO students and alumni have gone on to work on the trust and safety, policy, and investigations teams at TikTok, Meta, Twitter, Google, the Oversight Board, OpenAI, and many more.

This year we worked with nearly 40 student research analysts, continue to have two-thirds of our publications co-authored by students, and hosted two job fairs for students interested in careers in trust and safety.

Courses

Hack Lab, cotaught by SIO director Alex Stamos and research scholar Riana Pfefferkorn, aims to give students a solid understanding of the most common types of attacks used in cybercrime and cyberwarfare. Lectures cover the basics of an area of technology, how that technology has been misused in the past, and the legal and policy implications of those hacks. In lab sections, students apply this background to attack vulnerable, simulated targets using techniques and tools seen in the field. Students going into leadership positions will understand how events like ransomware attacks occur, how they as leaders can protect cyber assets, and what legal responsibilities they have. In two years, 210 students have completed the course. In 2022 the course enrolled 100 students.

Trust and Safety Engineering, also taught by Alex Stamos, exposes computer science undergraduates to the ways consumer internet services are abused to cause real human harm, along with potential operational, product, and engineering responses. Core to our trust and safety curriculum, the course covers topics such as spam, fraud, account takeovers, the use of social media by terrorists, misinformation, child exploitation, and harassment. This course is designed to give future entrepreneurs and product developers an understanding of how technology can be unintentionally misused and to expose them to the trade-offs and skills necessary to proactively mitigate future abuses. In 2022, 156 students completed Trust and Safety Engineering.

The Politics of Internet Abuse was newly added to SIO's teaching curriculum for 2022. The course, taught by SIO research scholar Shelby Grossman, serves as a social science complement to Trust and Safety Engineering. The course explores political science research on topics of online abuse and how online platforms respond to those threats. Students in the course complete a final project in partnership with computer science students in the Trust and Safety Engineering course where they design content policy around a specific type of online abuse. The class had 42 students in its first year.

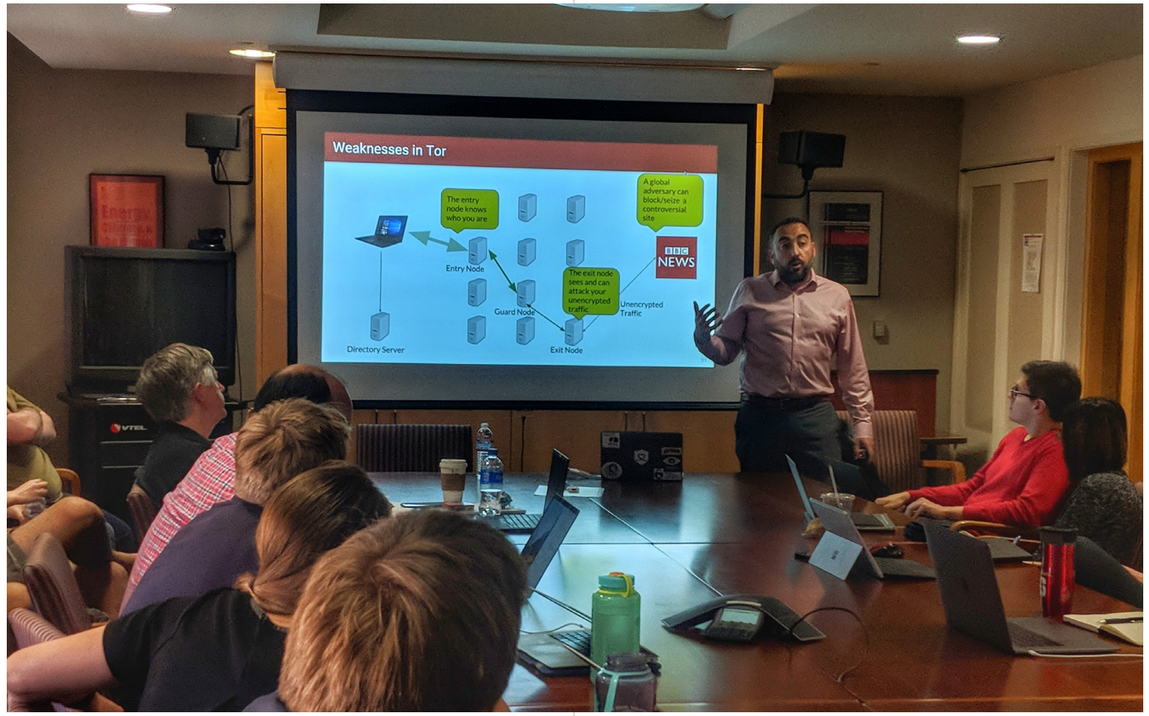

Online Open Source Investigations was developed by Shelby Grossman as a practical introduction to internet research using free and publicly available information. Dr. Grossman's syllabus blends best-in-practice open-source intelligence tools (OSINT) developed by organizations such as Bellingcat with SIO's own research practices. The course prepares students for online open-source research in jobs in the public sector, with technology companies and human rights organizations, and with other research and advocacy groups. SIO has also adapted the course into onboarding training for our research analysts. Intentionally small and hands-on, the course enrolled 23 students in 2022.

Current Topics in Technology Platform Policy was offered for a second time in the spring of 2022 to expose students to a broader set of outside experts and practitioners on the frontiers of emerging technology issues. Led by SIO director Alex Stamos and deputy director Elena Cryst, the class featured guest lectures by speakers with a wide range of distinguished professional and academic experiences, from U.S. government, to platform policy, and civil society oversight. The class enrolled 35 students in 2022.

Policy

A permanent presence in the nation’s capital has allowed our team to participate in meetings, workshops, and presentations on key policy and research issues.

Situated within the Cyber Policy Center at Stanford University, the Internet Observatory has since its inception engaged in dialogue around tech policy regulation and researcher data access in the U.S. and abroad. SIO invested in its policy engagement capacity this year, establishing a permanent Washington D.C. presence with two full-time staff to inform online trust and safety policy with SIO research and expertise at a pivotal moment for tech accountability policy and regulation around the world.

Deeply researched scholarship and reports are translated into timely and actionable policy briefings and recommendations for curbing online harms and abuse. A permanent presence in the nation’s capital has also allowed our team to participate in meetings, workshops, and presentations on key policy and research issues.

Over the past year, our team of scholars, researchers, and practitioners have testified before the U.S. Congress and addressed the United Kingdom and European Union parliaments on issues of online safety, election integrity, and social media transparency and accountability.

Our team has also provided written evidence for Congress, submitted expert comments to the Office of the Surgeon General, and authored comments for the Oversight Board. Policy commentary has appeared in Scientific American, Foreign Policy, The Hill, and Brookings TechStream, in addition to media appearances and quotes across dozens of news outlets.

In addition to sharing key research findings and recommendations on government initiatives and policy priorities, the SIO team regularly consults with policymakers and stakeholder communities working to address online trust and safety issues.

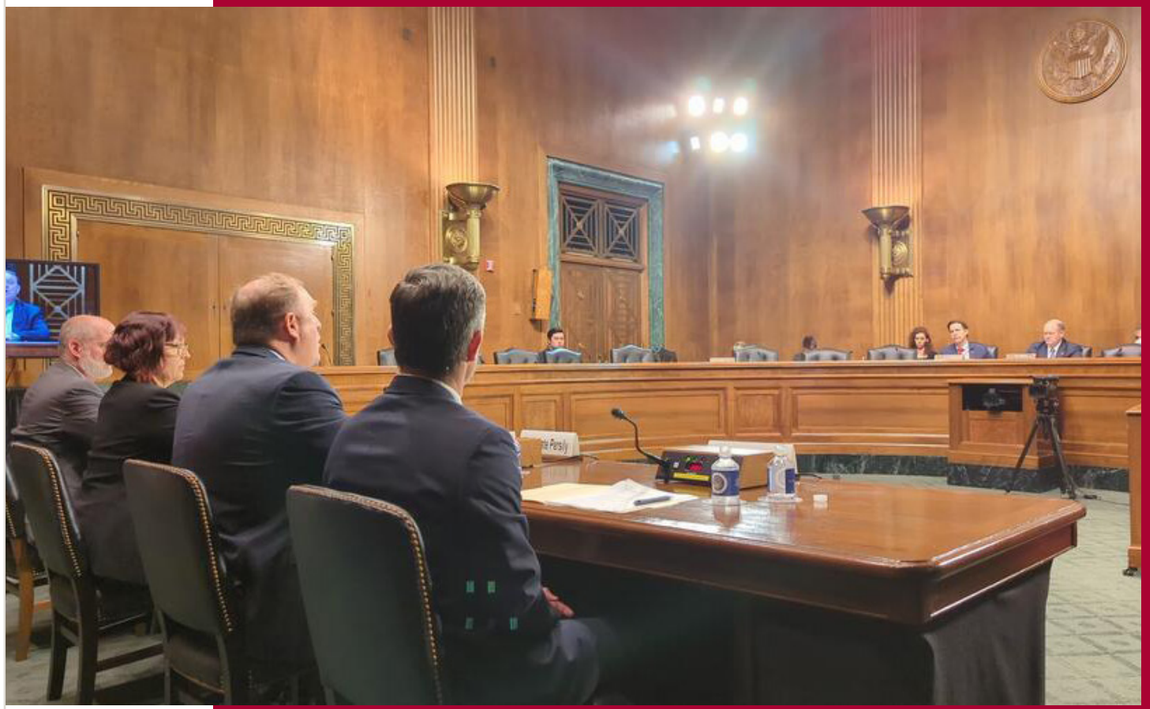

Emerging Threats to Election Administration
Senate Committee on Rules and Administration - October 26, 2021
Matt Masterson

Securing Democracy: Protecting against Threats to Election Infrastructure and Voter Confidence
House Committee on Homeland Security - January 20, 2022
Alex Stamos
Matt Masterson

Platform Transparency: Understanding the Impact of Social Media
Senate Committee on the Judiciary - May 4, 2022
Nate Persily
Daphne Keller

A Growing Threat: Foreign and Domestic Sources of Disinformation
Committee on House Administration - July 27, 2022
Renée DiResta

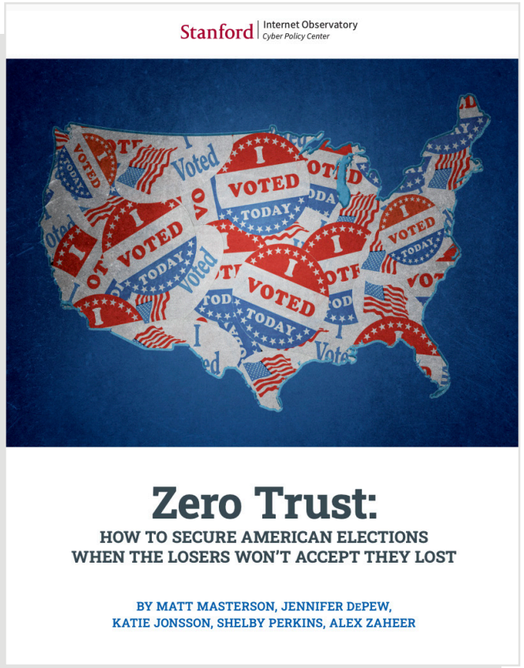

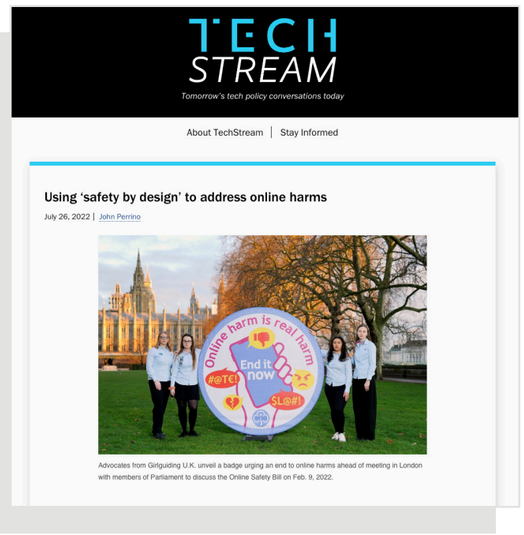

Looking Forward

The Stanford Internet Observatory has quickly become a leading voice on policy addressing social media abuse, online safety, and researcher access to social media data.

We will continue to expand on our mission to provide the policy community independent analysis on these issues with policy explainers and commentary that blends the technical and policy knowledge that decision makers need.

And we will serve as connectors and conveners with a reputation for clear, timely, and trusted feedback and policy recommendations based on rigorous policy analysis and academic research.

Research

The Stanford Internet Observatory conducts broad, interdisciplinary research that blends rigorous academic work with timely reports and commentary to inform decision makers on critical issues. Our research is loosely broken into four categories, although our outputs often integrate elements from multiple categories. This year, the team published 10 journal articles and 15 technical reports, as well as dozens of blog posts and op-eds. Highlights from this work are shared on the subsequent pages. Our researchers—from students to postdocs, young professionals to experienced experts—have contributed to, among other topics, significant findings on the tactics behind online information sharing, the dynamics of alt-platforms, and content moderation tools, and our team has continued to build technical tooling that facilitates better real-time online research.

TRUST AND SAFETY

Internet platforms and services are built for users, but are often used in unexpected ways that result in real human harm. Our research continues to investigate operational, design, engineering, and policy processes that address the potential misuse and abuse of platform features.

PLATFORM POLICY

Technology companies and policymakers are increasingly addressing online harms in platform rules and public policy. SIO aims to ground government policy and regulation in technical reality while examining the impact and enforcement of platform policies and terms of service.

INFORMATION INTERFERENCE

SIO research works toward a better understanding of how rumors become misleading narratives spread across platforms and in the media. Our team studies the tactics and techniques used by networks of social media accounts in efforts to influence public opinion and the flow of information across online spaces and traditional media.

EMERGING TECHNOLOGY

With an eye towards the future, our research studies the potential use and abuse of machine learning and artificial intelligence technologies used to create media. This includes studying the underlying technology and novel application of deepfake video technologies and Generative Pre-trained Transformer (GPT) automated writing tools.

Alt-Platform Analyses

In the wake of the January 6, 2021, attack on the U.S. Capitol, the largest social media companies suspended former president Donald Trump's accounts and deplatformed many of his allies. To continue engaging with their audiences, many of these deplatformed accounts jumped to newer, less moderated social media platforms such as Parler, Gab, Gettr, Telegram, and Truth Social. These platforms saw a surge in active users and registered accounts. In this new environment, SIO conducted analyses of the dynamics of accounts and activities on three of these platforms: Parler, Gettr, and Gab.

Parler was launched in 2018 as a free speech alternative to Twitter. It had a modest user base of around 4 million users leading up to the 2020 election. That all changed after the 2020 election as the platform surged in popularity, growing to nearly 13 million users by January 8, 2021. The SIO analysis of Parler looked at the data structure of the platform to determine trends in behavior, and looked at elements such as users' bios and the content of their posts to understand the conversations on Parler. The point-in-time analysis reflected the state of the platform on January 9, 2021, just before it was deplatformed by its AWS hosting services.

- **Parler's moderation policies indicated that moderation occurs primarily based on user reports to the platform's mere 802 moderators, with minimal proactive monitoring.**
- **Many of the most active accounts on Parler used automated posting integrations—such as RSS feeds—that allow users to simultaneously post across multiple social media channels and personal blogs. This indicates that the users may not have been approaching Parler as their primary platform but rather hedging their bets and engaging with audiences on many fronts.**
- **A proliferation of fake and spam accounts on Parler promoted commercial content such as Trump coin scams and OnlyFans profiles.**

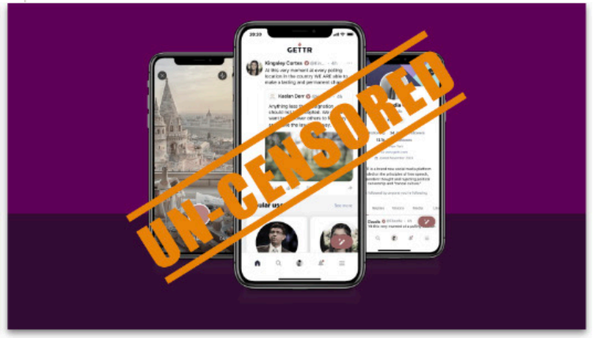

- **Several distinct account creation peaks on Parler attracted users from Brazil and Saudi Arabia, plus accounts in Chinese and Japanese, seemingly in response to increased content labeling and removal on Twitter.**

On July 1, 2021, Gettr, a new social network modeled after Twitter, was launched by former Trump spokesman Jason Miller, with assistance and promotion by exiled Chinese businessman Miles Guo, former Trump strategist Steve Bannon, and others. SIO charted Gettr's growth over its first month, examining the user community, content, structure, and dynamics. The analysis also highlighted some of the perils of launching such a network without trust and safety measures in place: the proliferation of gratuitous adult content, spam, and, most painfully, child exploitation imagery, all of which could be caught by cursory automated scanning systems.

- **The parties responsible for the site and app were not transparent. While Jason Miller distanced Gettr from Miles Guo, the app appeared to still be developed by a Guo-linked development team.**
- **Gettr's early growth numbers as calculated from the platform data were lower than the numbers publicized by Jason Miller.**
- **Gettr shows similar cultural demographics as Parler: the site attracted far-right users in the United States and Brazil that were deplatformed by larger social media sites. The platform also had a sizable Arabic-speaking userbase.**
- **Gettr, like Parler, relied on a user reporting model to detect harmful content.**

Founded in 2016, Gab is the longest established of the three platforms SIO analyzed. Gab spent most of its existence on the toxic fringe of alternative social media platforms—something exacerbated by its self-identified white Christian nationalist user base, its links to the 2018 Pittsburgh synagogue shooting, and its fraught relationship with various app stores and service providers. The events of January 6, 2021, caused a massive spike in new users and income for Gab. This appears to have been a lifeline for the platform, which had stagnant user growth and increasing monetary losses; it resulted in thousands of new paid subscribers and donors, and a marked uptick in platform use. SIO's in-depth analysis of Gab raised questions about the knock-on effects of content moderation and deplatforming, and the balance between keeping mainstream communities safe and creating incubators for extremists.

- **The post-January 6 user growth on Gab very likely served as a critical financial and user lifeline for the platform.**
- **In 2021 and 2022 Gab was a significant tool for coordinating real-world events including anti-vaccine protests and "trucker convoy" groups.**
- **Extreme antisemitic, racist, and homophobic content is rife on the platform, with open praise of Nazism, encouragement of violence against minorities, and "Great Replacement" narratives. Many of the memes cited by the Buffalo shooter's manifesto are indistinguishable from content on Gab, and such content appears even in "mainstream" user groups.**
- **Gab contains much of the same toxic content as sites with more extreme reputations, such as Stormfront, but companies including Cloudflare and Epik—who refused service to Stormfront—continue to provide Gab with services.**

The Virality Project

Governments and the public health community faced an unprecedented challenge disseminating accurate health information during the COVID-19 global pandemic. Mis- and disinformation spread widely as experts and public officials struggled to deliver clear and consistent guidance.

To address these challenges, the Stanford Internet Observatory convened the Virality Project, a coalition of research organizations that observed and analyzed COVID-19-related vaccine content on social media from February to August 2021. Participants included the University of Washington Center for an Informed Public,

the Atlantic Council's Digital Forensic Research Lab, Graphika, the National Conference on Citizenship's Algorithmic Transparency Institute, and New York University's Center for Social Media and Politics and Tandon School of Engineering.

The coalition focused on identifying rumors about the safety, efficacy, or distribution of vaccines. Rapid analysis helped provide situational awareness and actionable insights for the public health community and other stakeholders tasked with informing the public and mitigating harm from false or misleading claims.

In a final report, *Memes, Magnets and Microchips: Narrative dynamics around COVID-19 vaccines*, the research team highlights findings that COVID-19 anti-vaccine narratives were not new, but rather adapted from long-used anti-vaccine storylines repeated by a small group of domestic activists and influencers. The report also offers analysis of social media platform policies and recommendations for a whole-of-society response to address harmful health misinformation going forward.

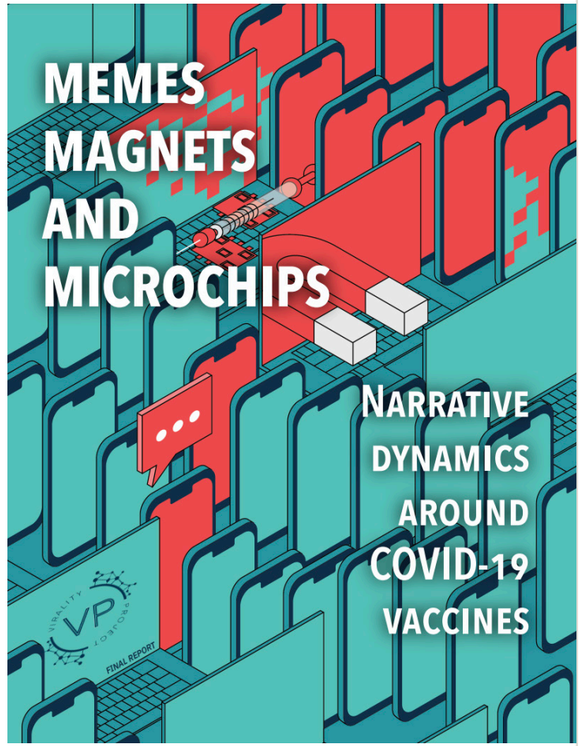

External Collaborations

Stanford Internet Observatory and its scholars proudly participate in collaborative efforts to build community and develop shared learnings on online information sharing and trust and safety. The following highlights a sample of our many collaborations.

- In collaboration with the **National Democratic Institute** and the **International Republican Institute**, SIO published **Combating Information Manipulation: A Playbook for Elections and Beyond** on September 20, 2021. The playbook was co-authored by SIO postdoctoral scholar **Samantha Bradshaw**.
- The **Aspen Institute's Commission on Information Interference** was created to address the conditions under which bad information becomes as prevalent and persuasive as good information. The commission published its **final report** on November 15, 2021. SIO director **Alex Stamos** served as a commissioner, and research manager **Renée DiResta** served as a technical advisor.
- SIO convened the **Virality Project** with partners at the **University of Washington's Center for an Informed Public, New York University's Tandon School of Engineering and Center for Social Media and Politics, Graphika**, and the **Atlantic Council's Digital Forensic Research Lab**.
- SIO again partnered with the **University of Washington's Center for an Informed Public** to relaunch the **Election Integrity Partnership** during the 2022 midterm elections.

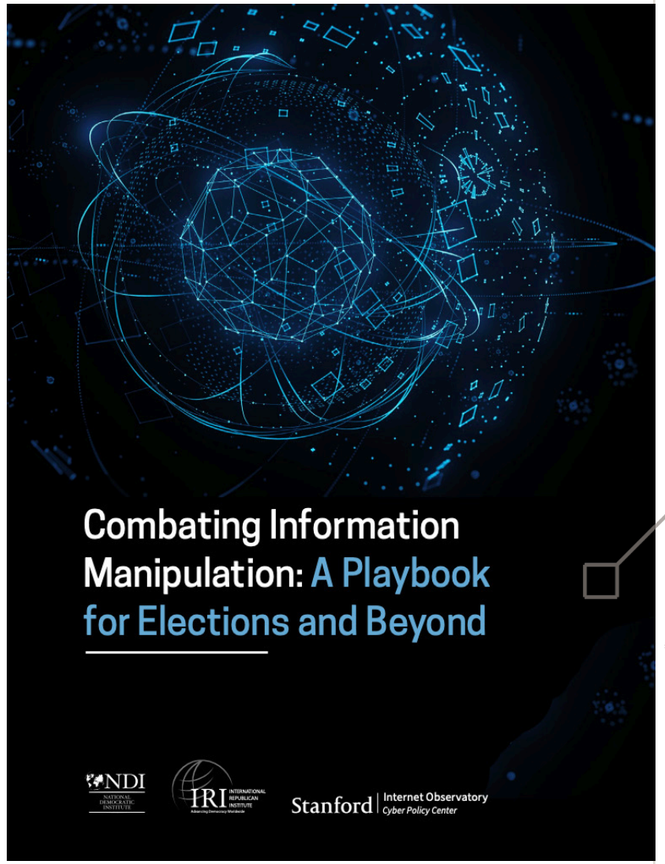

Technical Capabilities

SIO's innovative and ever-expanding technical capabilities are what make our research possible. Under the leadership of industry veteran and SIO Chief Technologist David Thiel, our team of technical research assistants develop tools for rapid ingest and analysis of social media data, and publish open-source tools for the research community. SIO now has a robust and rapid analysis toolkit for ingesting and understanding narrative spread on Twitter using the platform's API. Researchers can ingest both historical and real-time Telegram channels and apply optical character recognition on images shared to the channels. These tools have facilitated real-time research on narratives about the war in Ukraine and expanded capability for observing election and vaccine narratives across platforms.

The SIO tech team has released several collaborative open-source tools for analysis of alt-platforms including Gab and Truth Social. They have also contributed to other open-source projects to better enable online research. These tools are released on the SIO GitHub repository and have been used by other researchers and investigators. This contributory work builds out SIO's contribution to the research community, both at Stanford and around the world.

Thank You

Our work could not have been possible without the many contributions of our staff, students, and colleagues.

Thanks to our generous donors:
Craig Newmark Philanthropies
Charles Koch Foundation
Omidyar Network Funds
Theodore Schlein
Cognizant Corporation
Hewlett Foundation
National Science Foundation
Election Trust Initiative

Thanks to our staff:
Dan Bateyko
Elena Cryst
Renée DiResta
Shelby Grossman
Karen Nershi
John Perrino
Riana Pfefferkorn
Ronald Robertson
Alex Stamos
David Thiel
Tongtong Zhang

Thanks to our faculty partners:
Michael McFaul, director, Freeman Spogli Institute for International Studies
Nathaniel Persily, co-director, Cyber Policy Center
Dan Boneh, co-director, Cyber Policy Center
Jeff Hancock, founding director, Stanford Social Media Lab

… and to over a hundred students, colleagues, and friends who have been part of our team over the past three years.

Stanford | Internet Observatory
Cyber Policy Center